

APPENDIX A



MATRIX ALGEBRA

A.1 TERMINOLOGY

A **matrix** is a rectangular array of numbers, denoted

$$\mathbf{A} = [a_{ik}] = [\mathbf{A}]_{ik} = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1K} \\ a_{21} & a_{22} & \cdots & a_{2K} \\ \cdots & \cdots & \cdots & \cdots \\ a_{n1} & a_{n2} & \cdots & a_{nK} \end{bmatrix} \quad (\text{A-1})$$

The typical element is used to denote the matrix. A subscripted element of a matrix is always read as $a_{\text{row}, \text{column}}$. An example is given in Table A.1. In these data, the rows are identified with years and the columns with particular variables.

A **vector** is an ordered set of numbers arranged either in a row or a column. In view of the preceding, a **row vector** is also a matrix with one row, whereas a **column vector** is a matrix with one column. Thus, in Table A.1, the five variables observed for 1972 (including the date) constitute a row vector, whereas the time series of nine values for consumption is a column vector.

A matrix can also be viewed as a set of column vectors or as a set of row vectors.¹ The **dimensions** of a matrix are the numbers of rows and columns it contains. “**A** is an $n \times K$ matrix” (read “ n by K ”) will always mean that **A** has n rows and K columns. If n equals K , then **A** is a **square matrix**. Several particular types of square matrices occur frequently in econometrics.

- A **symmetric matrix** is one in which $a_{ik} = a_{ki}$ for all i and k .
- A **diagonal matrix** is a square matrix whose only nonzero elements appear on the **main diagonal**, that is, moving from upper left to lower right.
- A **scalar matrix** is a diagonal matrix with the same value in all diagonal elements.
- An **identity matrix** is a scalar matrix with ones on the diagonal. This matrix is always denoted **I**. A subscript is sometimes included to indicate its size, or **order**. For example,
- A **triangular matrix** is one that has only zeros either above or below the main diagonal. If the zeros are above the diagonal, the matrix is **lower triangular**.

A.2 ALGEBRAIC MANIPULATION OF MATRICES

A.2.1 EQUALITY OF MATRICES

Matrices (or vectors) **A** and **B** are equal if and only if they have the same dimensions and each element of **A** equals the corresponding element of **B**. That is,

$$\mathbf{A} = \mathbf{B} \quad \text{if and only if } a_{ik} = b_{ik} \quad \text{for all } i \text{ and } k. \quad (\text{A-2})$$

¹Henceforth, we shall denote a matrix by a boldfaced capital letter, as is **A** in (A-1), and a vector as a boldfaced lowercase letter, as in **a**. Unless otherwise noted, a vector will always be assumed to be a *column vector*.

TABLE A.1 Matrix of Macroeconomic Data

Row	Column				
	1 Year	2 Consumption (billions of dollars)	3 GNP (billions of dollars)	4 GNP Deflator	5 Discount Rate (N.Y Fed., avg.)
1	1972	737.1	1185.9	1.0000	4.50
2	1973	812.0	1326.4	1.0575	6.44
3	1974	808.1	1434.2	1.1508	7.83
4	1975	976.4	1549.2	1.2579	6.25
5	1976	1084.3	1718.0	1.3234	5.50
6	1977	1204.4	1918.3	1.4005	5.46
7	1978	1346.5	2163.9	1.5042	7.46
8	1979	1507.2	2417.8	1.6342	10.28
9	1980	1667.2	2633.1	1.7864	11.77

Source: Data from the *Economic Report of the President* (Washington, D.C.: U.S. Government Printing Office, 1983).

A.2.2 TRANSPOSITION

The **transpose** of a matrix \mathbf{A} , denoted \mathbf{A}' , is obtained by creating the matrix whose k th row is the k th column of the original matrix. Thus, if $\mathbf{B} = \mathbf{A}'$, then each column of \mathbf{A} will appear as the corresponding row of \mathbf{B} . If \mathbf{A} is $n \times K$, then \mathbf{A}' is $K \times n$.

An equivalent definition of the transpose of a matrix is

$$\mathbf{B} = \mathbf{A}' \Leftrightarrow b_{ik} = a_{ki} \quad \text{for all } i \text{ and } k. \quad (\text{A-3})$$

The definition of a symmetric matrix implies that

$$\text{if (and only if) } \mathbf{A} \text{ is symmetric, then } \mathbf{A} = \mathbf{A}'. \quad (\text{A-4})$$

It also follows from the definition that for any \mathbf{A} ,

$$(\mathbf{A}')' = \mathbf{A}. \quad (\text{A-5})$$

Finally, the transpose of a column vector, \mathbf{a} , is a row vector:

$$\mathbf{a}' = [a_1 \quad a_2 \quad \cdots \quad a_n].$$

A.2.3 MATRIX ADDITION

The operations of addition and subtraction are extended to matrices by defining

$$\mathbf{C} = \mathbf{A} + \mathbf{B} = [a_{ik} + b_{ik}]. \quad (\text{A-6})$$

$$\mathbf{A} - \mathbf{B} = [a_{ik} - b_{ik}]. \quad (\text{A-7})$$

Matrices cannot be added unless they have the same dimensions, in which case they are said to be **conformable for addition**. A **zero matrix** or **null matrix** is one whose elements are all zero. In the addition of matrices, the zero matrix plays the same role as the scalar 0 in scalar addition; that is,

$$\mathbf{A} + \mathbf{0} = \mathbf{A}. \quad (\text{A-8})$$

It follows from (A-6) that matrix addition is commutative,

$$\mathbf{A} + \mathbf{B} = \mathbf{B} + \mathbf{A}. \quad (\text{A-9})$$

and associative,

$$(\mathbf{A} + \mathbf{B}) + \mathbf{C} = \mathbf{A} + (\mathbf{B} + \mathbf{C}), \quad (\text{A-10})$$

and that

$$(\mathbf{A} + \mathbf{B})' = \mathbf{A}' + \mathbf{B}'. \quad (\text{A-11})$$

A.2.4 VECTOR MULTIPLICATION

Matrices are multiplied by using the **inner product**. The inner product, or **dot product**, of two vectors, \mathbf{a} and \mathbf{b} , is a scalar and is written

$$\mathbf{a}'\mathbf{b} = a_1b_1 + a_2b_2 + \cdots + a_nb_n. \quad (\text{A-12})$$

Note that the inner product is written as the transpose of vector \mathbf{a} times vector \mathbf{b} , a row vector times a column vector. In (A-12), each term a_jb_j equals b_ja_j ; hence

$$\mathbf{a}'\mathbf{b} = \mathbf{b}'\mathbf{a}. \quad (\text{A-13})$$

A.2.5 A NOTATION FOR ROWS AND COLUMNS OF A MATRIX

We need a notation for the i th row of a matrix. Throughout this book, an untransposed vector will always be a column vector. However, we will often require a notation for the column vector that is the transpose of a row of a matrix. This has the potential to create some ambiguity, but the following convention based on the subscripts will suffice for our work throughout this text:

- \mathbf{a}_k , or \mathbf{a}_l or \mathbf{a}_m , will denote column k , l , or m of the matrix \mathbf{A} ,
- \mathbf{a}_i , or \mathbf{a}_j or \mathbf{a}_t or \mathbf{a}_s , will denote the column vector formed by the transpose of row i , j , t , or s of matrix \mathbf{A} . Thus, \mathbf{a}_i' is row i of \mathbf{A} . (A-14)

For example, from the data in Table A.1 it might be convenient to speak of $\mathbf{x}_i = 1972$ as the 5×1 vector containing the five variables measured for the year 1972, that is, the transpose of the 1972 row of the matrix. In our applications, the common association of subscripts “ i ” and “ j ” with individual i or j , and “ t ” and “ s ” with time periods t and s will be natural.

A.2.6 MATRIX MULTIPLICATION AND SCALAR MULTIPLICATION

For an $n \times K$ matrix \mathbf{A} and a $K \times M$ matrix \mathbf{B} , the product matrix, $\mathbf{C} = \mathbf{AB}$, is an $n \times M$ matrix whose ik th element is the inner product of row i of \mathbf{A} and column k of \mathbf{B} . Thus, the product matrix \mathbf{C} is

$$\mathbf{C} = \mathbf{AB} \Rightarrow c_{ik} = \mathbf{a}_i'\mathbf{b}_k. \quad (\text{A-15})$$

[Note our use of (A-14) in (A-15).] To multiply two matrices, the number of columns in the first must be the same as the number of rows in the second, in which case they are **conformable for multiplication**.² Multiplication of matrices is generally not commutative. In some cases, \mathbf{AB} may exist, but \mathbf{BA} may be undefined or, if it does exist, may have different dimensions. In general, however, even if \mathbf{AB} and \mathbf{BA} do have the same dimensions, they will not be equal. In view of this, we define **premultiplication** and **postmultiplication** of matrices. In the product \mathbf{AB} , \mathbf{B} is *premultiplied* by \mathbf{A} , whereas \mathbf{A} is *postmultiplied* by \mathbf{B} .

²A simple way to check the conformability of two matrices for multiplication is to write down the dimensions of the operation, for example, $(n \times K)$ times $(K \times M)$. The inner dimensions must be equal; the result has dimensions equal to the outer values.

806 APPENDIX A ♦ Matrix Algebra

Scalar multiplication of a matrix is the operation of multiplying every element of the matrix by a given scalar. For scalar c and matrix \mathbf{A} ,

$$c\mathbf{A} = [ca_{ik}]. \tag{A-16}$$

The product of a matrix and a vector is written

$$\mathbf{c} = \mathbf{A}\mathbf{b}.$$

The number of elements in \mathbf{b} must equal the number of columns in \mathbf{A} ; the result is a vector with number of elements equal to the number of rows in \mathbf{A} . For example,

$$\begin{bmatrix} 5 \\ 4 \\ 1 \end{bmatrix} = \begin{bmatrix} 4 & 2 & 1 \\ 2 & 6 & 1 \\ 1 & 1 & 0 \end{bmatrix} \begin{bmatrix} a \\ b \\ c \end{bmatrix}.$$

We can interpret this in two ways. First, it is a compact way of writing the three equations

$$\begin{aligned} 5 &= 4a + 2b + 1c, \\ 4 &= 2a + 6b + 1c, \\ 1 &= 1a + 1b + 0c. \end{aligned}$$

Second, by writing the set of equations as

$$\begin{bmatrix} 5 \\ 4 \\ 1 \end{bmatrix} = a \begin{bmatrix} 4 \\ 2 \\ 1 \end{bmatrix} + b \begin{bmatrix} 2 \\ 6 \\ 1 \end{bmatrix} + c \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix},$$

we see that the right-hand side is a **linear combination** of the columns of the matrix where the coefficients are the elements of the vector. For the general case,

$$\mathbf{c} = \mathbf{A}\mathbf{b} = b_1\mathbf{a}_1 + b_2\mathbf{a}_2 + \cdots + b_k\mathbf{a}_k. \tag{A-17}$$

In the calculation of a matrix product $\mathbf{C} = \mathbf{A}\mathbf{B}$, each column of \mathbf{C} is a linear combination of the columns of \mathbf{A} , where the coefficients are the elements in the corresponding column of \mathbf{B} . That is,

$$\mathbf{C} = \mathbf{A}\mathbf{B} \Leftrightarrow \mathbf{c}_k = \mathbf{A}\mathbf{b}_k. \tag{A-18}$$

Let \mathbf{e}_k be a column vector that has zeros everywhere except for a one in the k th position. Then $\mathbf{A}\mathbf{e}_k$ is a linear combination of the columns of \mathbf{A} in which the coefficient on every column but the k th is zero, whereas that on the k th is one. The result is

$$\mathbf{a}_k = \mathbf{A}\mathbf{e}_k. \tag{A-19}$$

Combining this result with (A-17) produces

$$(\mathbf{a}_1 \ \mathbf{a}_2 \ \cdots \ \mathbf{a}_n) = \mathbf{A}(\mathbf{e}_1 \ \mathbf{e}_2 \ \cdots \ \mathbf{e}_n) = \mathbf{A}\mathbf{I} = \mathbf{A}. \tag{A-20}$$

In matrix multiplication, the identity matrix is analogous to the scalar 1. For any matrix or vector \mathbf{A} , $\mathbf{A}\mathbf{I} = \mathbf{A}$. In addition, $\mathbf{I}\mathbf{A} = \mathbf{A}$, although if \mathbf{A} is not a square matrix, the two identity matrices are of different orders.

A conformable matrix of zeros produces the expected result: $\mathbf{A}\mathbf{0} = \mathbf{0}$.

Some general rules for matrix multiplication are as follows:

- **Associative law:** $(\mathbf{A}\mathbf{B})\mathbf{C} = \mathbf{A}(\mathbf{B}\mathbf{C}).$ (A-21)

- **Distributive law:** $\mathbf{A}(\mathbf{B} + \mathbf{C}) = \mathbf{A}\mathbf{B} + \mathbf{A}\mathbf{C}.$ (A-22)

• **Transpose of a product:** $(AB)' = B'A'$. (A-23)

• **Transpose of an extended product:** $(ABC)' = C'B'A'$. (A-24)

A.2.7 SUMS OF VALUES

Denote by \mathbf{i} a vector that contains a column of ones. Then,

$$\sum_{i=1}^n x_i = x_1 + x_2 + \dots + x_n = \mathbf{i}'\mathbf{x}. \tag{A-25}$$

If all elements in \mathbf{x} are equal to the same constant a , then $\mathbf{x} = a\mathbf{i}$ and

$$\sum_{i=1}^n x_i = \mathbf{i}'(a\mathbf{i}) = a(\mathbf{i}'\mathbf{i}) = na. \tag{A-26}$$

For any constant a and vector \mathbf{x} ,

$$\sum_{i=1}^n ax_i = a \sum_{i=1}^n x_i = a\mathbf{i}'\mathbf{x}. \tag{A-27}$$

If $a = 1/n$, then we obtain the arithmetic mean,

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i = \frac{1}{n} \mathbf{i}'\mathbf{x}, \tag{A-28}$$

from which it follows that

$$\sum_{i=1}^n x_i = \mathbf{i}'\mathbf{x} = n\bar{x}.$$

The sum of squares of the elements in a vector \mathbf{x} is

$$\sum_{i=1}^n x_i^2 = \mathbf{x}'\mathbf{x}; \tag{A-29}$$

while the sum of the products of the n elements in vectors \mathbf{x} and \mathbf{y} is

$$\sum_{i=1}^n x_i y_i = \mathbf{x}'\mathbf{y}. \tag{A-30}$$

By the definition of matrix multiplication,

$$[\mathbf{X}'\mathbf{X}]_{kl} = [\mathbf{x}'_k \mathbf{x}_l] \tag{A-31}$$

is the inner product of the k th and l th columns of \mathbf{X} . For example, for the data set given in Table A.1, if we define \mathbf{X} as the 9×3 matrix containing (year, consumption, GNP), then

$$\begin{aligned} [\mathbf{X}'\mathbf{X}]_{23} &= \sum_{t=1972}^{1980} \text{consumption}_t \text{GNP}_t = 737.1(1185.9) + \dots + 1667.2(2633.1) \\ &= 19,743,711.34. \end{aligned}$$

If \mathbf{X} is $n \times K$, then [again using (A-14)]

$$\mathbf{X}'\mathbf{X} = \sum_{i=1}^n \mathbf{x}_i \mathbf{x}'_i.$$

808 APPENDIX A ♦ Matrix Algebra

This form shows that the $K \times K$ matrix $\mathbf{X}'\mathbf{X}$ is the sum of n $K \times K$ matrices, each formed from a single row (year) of \mathbf{X} . For the example given earlier, this sum is of nine 3×3 matrices, each formed from one row (year) of the original data matrix.

A.2.8 A USEFUL IDEMPOTENT MATRIX

A fundamental matrix in statistics is the one that is used to transform data to deviations from their mean. First,

$$\mathbf{i}\bar{x} = \frac{1}{n}\mathbf{i}'\mathbf{x} = \begin{bmatrix} \bar{x} \\ \bar{x} \\ \vdots \\ \bar{x} \end{bmatrix} = \frac{1}{n}\mathbf{i}\mathbf{i}'\mathbf{x}. \quad (\text{A-32})$$

The matrix $(1/n)\mathbf{i}\mathbf{i}'$ is an $n \times n$ matrix with every element equal to $1/n$. The set of values in deviations form is

$$\begin{bmatrix} x_1 - \bar{x} \\ x_2 - \bar{x} \\ \dots \\ x_n - \bar{x} \end{bmatrix} = [\mathbf{x} - \mathbf{i}\bar{x}] = \left[\mathbf{x} - \frac{1}{n}\mathbf{i}\mathbf{i}'\mathbf{x} \right]. \quad (\text{A-33})$$

Since $\mathbf{x} = \mathbf{I}\mathbf{x}$,

$$\left[\mathbf{x} - \frac{1}{n}\mathbf{i}\mathbf{i}'\mathbf{x} \right] = \left[\mathbf{I}\mathbf{x} - \frac{1}{n}\mathbf{i}\mathbf{i}'\mathbf{x} \right] = \left[\mathbf{I} - \frac{1}{n}\mathbf{i}\mathbf{i}' \right] \mathbf{x} = \mathbf{M}^0\mathbf{x}. \quad (\text{A-34})$$

Henceforth, the symbol \mathbf{M}^0 will be used only for this matrix. Its diagonal elements are all $(1 - 1/n)$, and its off-diagonal elements are $-1/n$. The matrix \mathbf{M}^0 is primarily useful in computing sums of squared deviations. Some computations are simplified by the result

$$\mathbf{M}^0\mathbf{i} = \left[\mathbf{I} - \frac{1}{n}\mathbf{i}\mathbf{i}' \right] \mathbf{i} = \mathbf{i} - \frac{1}{n}\mathbf{i}(\mathbf{i}'\mathbf{i}) = \mathbf{0},$$

which implies that $\mathbf{i}'\mathbf{M}^0 = \mathbf{0}'$. The sum of deviations about the mean is then

$$\sum_{i=1}^n (x_i - \bar{x}) = \mathbf{i}'[\mathbf{M}^0\mathbf{x}] = \mathbf{0}'\mathbf{x} = 0. \quad (\text{A-35})$$

For a single variable \mathbf{x} , the sum of squared deviations about the mean is

$$\sum_{i=1}^n (x_i - \bar{x})^2 = \left(\sum_{i=1}^n x_i^2 \right) - n\bar{x}^2. \quad (\text{A-36})$$

In matrix terms,

$$\sum_{i=1}^n (x_i - \bar{x})^2 = (\mathbf{x} - \bar{x}\mathbf{i})'(\mathbf{x} - \bar{x}\mathbf{i}) = (\mathbf{M}^0\mathbf{x})'(\mathbf{M}^0\mathbf{x}) = \mathbf{x}'\mathbf{M}^0\mathbf{M}^0\mathbf{x}.$$

Two properties of \mathbf{M}^0 are useful at this point. First, since all off-diagonal elements of \mathbf{M}^0 equal $-1/n$, \mathbf{M}^0 is symmetric. Second, as can easily be verified by multiplication, \mathbf{M}^0 is equal to its square; $\mathbf{M}^0\mathbf{M}^0 = \mathbf{M}^0$.

DEFINITION A.1 Idempotent Matrix

An *idempotent matrix*, \mathbf{M} , is one that is equal to its square, that is, $\mathbf{M}^2 = \mathbf{M}\mathbf{M} = \mathbf{M}$. If \mathbf{M} is a symmetric idempotent matrix (all of the idempotent matrices we shall encounter are asymmetric), then $\mathbf{M}'\mathbf{M} = \mathbf{M}$.

Thus, \mathbf{M}^0 is a symmetric idempotent matrix. Combining results, we obtain

$$\sum_{i=1}^n (x_i - \bar{x})^2 = \mathbf{x}'\mathbf{M}^0\mathbf{x}. \tag{A-37}$$

Consider constructing a matrix of sums of squares and cross products in deviations from the column means. For two vectors \mathbf{x} and \mathbf{y} ,

$$\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y}) = (\mathbf{M}^0\mathbf{x})'(\mathbf{M}^0\mathbf{y}), \tag{A-38}$$

so

$$\begin{bmatrix} \sum_{i=1}^n (x_i - \bar{x})^2 & \sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y}) \\ \sum_{i=1}^n (y_i - \bar{y})(x_i - \bar{x}) & \sum_{i=1}^n (y_i - \bar{y})^2 \end{bmatrix} = \begin{bmatrix} \mathbf{x}'\mathbf{M}^0\mathbf{x} & \mathbf{x}'\mathbf{M}^0\mathbf{y} \\ \mathbf{y}'\mathbf{M}^0\mathbf{x} & \mathbf{y}'\mathbf{M}^0\mathbf{y} \end{bmatrix}. \tag{A-39}$$

If we put the two column vectors \mathbf{x} and \mathbf{y} in an $n \times 2$ matrix $\mathbf{Z} = [\mathbf{x}, \mathbf{y}]$, then $\mathbf{M}^0\mathbf{Z}$ is the $n \times 2$ matrix in which the two columns of data are in mean deviation form. Then

$$(\mathbf{M}^0\mathbf{Z})'(\mathbf{M}^0\mathbf{Z}) = \mathbf{Z}'\mathbf{M}^0\mathbf{M}^0\mathbf{Z} = \mathbf{Z}'\mathbf{M}^0\mathbf{Z}.$$

A.3 GEOMETRY OF MATRICES

A.3.1 VECTOR SPACES

The K elements of a column vector

$$\mathbf{a} = \begin{bmatrix} a_1 \\ a_2 \\ \dots \\ a_K \end{bmatrix}$$

can be viewed as the coordinates of a point in a K -dimensional space, as shown in Figure A.1 for two dimensions, or as the definition of the line segment connecting the origin and the point defined by \mathbf{a} .

Two basic arithmetic operations are defined for vectors, **scalar multiplication** and **addition**. A scalar multiple of a vector, \mathbf{a} , is another vector, say \mathbf{a}^* , whose coordinates are the scalar multiple of \mathbf{a} 's coordinates. Thus, in Figure A.1,

$$\mathbf{a} = \begin{bmatrix} 1 \\ 2 \end{bmatrix}, \quad \mathbf{a}^* = 2\mathbf{a} = \begin{bmatrix} 2 \\ 4 \end{bmatrix}, \quad \mathbf{a}^{**} = -\frac{1}{2}\mathbf{a} = \begin{bmatrix} -\frac{1}{2} \\ -1 \end{bmatrix}.$$




FIGURE A.1 Vector Space.

The set of all possible scalar multiples of \mathbf{a} is the line through the origin, $\mathbf{0}$ and \mathbf{a} . Any scalar multiple of \mathbf{a} is a segment of this line. The sum of two vectors \mathbf{a} and \mathbf{b} is a third vector whose coordinates are the sums of the corresponding coordinates of \mathbf{a} and \mathbf{b} . For example,

$$\mathbf{c} = \mathbf{a} + \mathbf{b} = \begin{bmatrix} 1 \\ 2 \end{bmatrix} + \begin{bmatrix} 2 \\ 1 \end{bmatrix} = \begin{bmatrix} 3 \\ 3 \end{bmatrix}.$$

Geometrically, \mathbf{c} is obtained by moving in the distance and direction defined by \mathbf{b} from the tip of \mathbf{a} or, because addition is commutative, from the tip of \mathbf{b} in the distance and direction of \mathbf{a} .

The two-dimensional plane is the set of all vectors with two real-valued coordinates. We label this set \mathbb{R}^2 (“R two,” not “R squared”). It has two important properties.

- \mathbb{R}^2 is closed under scalar multiplication; every scalar multiple of a vector in \mathbb{R}^2 is also in \mathbb{R}^2 .
- \mathbb{R}^2 is closed under addition; the sum of any two vectors in the plane is always a vector in \mathbb{R}^2 .

DEFINITION A.2 Vector Space

A *vector space* is any set of vectors that is closed under scalar multiplication and addition.

Another example is the set of all real numbers, that is, \mathbb{R}^1 , that is, the set of vectors with one real element. In general, that set of K -element vectors all of whose elements are real numbers is a K -dimensional vector space, denoted \mathbb{R}^K . The preceding examples are drawn in \mathbb{R}^2 .

A.3.2 LINEAR COMBINATIONS OF VECTORS AND BASIS VECTORS

In Figure A.2, $\mathbf{c} = \mathbf{a} + \mathbf{b}$ and $\mathbf{d} = \mathbf{a}^* + \mathbf{b}$. But since $\mathbf{a}^* = 2\mathbf{a}$, $\mathbf{d} = 2\mathbf{a} + \mathbf{b}$. Also, $\mathbf{e} = \mathbf{a} + 2\mathbf{b}$ and $\mathbf{f} = \mathbf{b} + (-\mathbf{a}) = \mathbf{b} - \mathbf{a}$. As this exercise suggests, any vector in \mathbb{R}^2 could be obtained as a **linear combination** of \mathbf{a} and \mathbf{b} .

DEFINITION A.3 Basis Vectors

A set of vectors in a vector space is a **basis** for that vector space if any vector in the vector space can be written as a linear combination of that set of vectors.

As is suggested by Figure A.2, any pair of two-element vectors, including \mathbf{a} and \mathbf{b} , that point in different directions will form a basis for \mathbb{R}^2 . Consider an arbitrary set of vectors in \mathbb{R}^2 , \mathbf{a} , \mathbf{b} , and \mathbf{c} . If \mathbf{a} and \mathbf{b} are a basis, then we can find numbers α_1 and α_2 such that $\mathbf{c} = \alpha_1\mathbf{a} + \alpha_2\mathbf{b}$. Let

$$\mathbf{a} = \begin{bmatrix} a_1 \\ a_2 \end{bmatrix}, \quad \mathbf{b} = \begin{bmatrix} b_1 \\ b_2 \end{bmatrix}, \quad \mathbf{c} = \begin{bmatrix} c_1 \\ c_2 \end{bmatrix}.$$

Then

$$\begin{aligned} c_1 &= \alpha_1 a_1 + \alpha_2 b_1, \\ c_2 &= \alpha_1 a_2 + \alpha_2 b_2. \end{aligned} \tag{A-40}$$

The solutions to this pair of equations are

$$\alpha_1 = \frac{b_2 c_1 - b_1 c_2}{a_1 b_2 - b_1 a_2}, \quad \alpha_2 = \frac{a_1 c_2 - a_2 c_1}{a_1 b_2 - b_1 a_2}. \tag{A-41}$$

This result gives a unique solution unless $(a_1 b_2 - b_1 a_2) = 0$. If $(a_1 b_2 - b_1 a_2) = 0$, then $a_1/a_2 = b_1/b_2$, which means that \mathbf{b} is just a multiple of \mathbf{a} . This returns us to our original condition, that \mathbf{a} and \mathbf{b} must point in different directions. The implication is that if \mathbf{a} and \mathbf{b} are any pair of vectors for which the denominator in (A-41) is not zero, then any other vector \mathbf{c} can be formed as a *unique* linear combination of \mathbf{a} and \mathbf{b} . The basis of a vector space is not unique, since any set of vectors that satisfies the definition will do. But for any particular basis, only one linear combination of them will produce another particular vector in the vector space.

A.3.3 LINEAR DEPENDENCE

As the preceding should suggest, K vectors are required to form a basis for \mathbb{R}^K . Although the basis for a vector space is not unique, not every set of K vectors will suffice. In Figure A.2, \mathbf{a} and \mathbf{b} form a basis for \mathbb{R}^2 , but \mathbf{a} and \mathbf{a}^* do not. The difference between these two pairs is that \mathbf{a} and \mathbf{b} are linearly *independent*, whereas \mathbf{a} and \mathbf{a}^* are linearly *dependent*.

DEFINITION A.4 Linear Dependence

A set of vectors is **linearly dependent** if any one of the vectors in the set can be written as a linear combination of the others.




FIGURE A.2 Linear Combinations of Vectors.

Since \mathbf{a}^* is a multiple of \mathbf{a} , \mathbf{a} and \mathbf{a}^* are linearly dependent. For another example, if

$$\mathbf{a} = \begin{bmatrix} 1 \\ 2 \end{bmatrix}, \quad \mathbf{b} = \begin{bmatrix} 3 \\ 3 \end{bmatrix}, \quad \text{and} \quad \mathbf{c} = \begin{bmatrix} 10 \\ 14 \end{bmatrix},$$

then

$$2\mathbf{a} + \mathbf{b} - \frac{1}{2}\mathbf{c} = \mathbf{0},$$

so \mathbf{a} , \mathbf{b} , and \mathbf{c} are linearly dependent. Any of the three possible pairs of them, however, are linearly independent.

DEFINITION A.5 Linear Independence

A set of vectors is *linearly independent* if and only if the only solution to

$$\alpha_1 \mathbf{a}_1 + \alpha_2 \mathbf{a}_2 + \cdots + \alpha_K \mathbf{a}_K = \mathbf{0}$$

is

$$\alpha_1 = \alpha_2 = \cdots = \alpha_K = 0.$$

The preceding implies the following equivalent definition of a basis.

DEFINITION A.6 Basis for a Vector Space

A basis for a vector space of K dimensions is any set of K linearly independent vectors in that vector space.

Since any $(K + 1)$ st vector can be written as a linear combination of the K basis vectors, it follows that any set of more than K vectors in \mathbb{R}^K must be linearly dependent.

A.3.4 SUBSPACES**DEFINITION A.7 Spanning Vectors**

*The set of all linear combinations of a set of vectors is the vector space that is **spanned** by those vectors.*

For example, by definition, the space spanned by a basis for \mathbb{R}^K is \mathbb{R}^K . An implication of this is that if **a** and **b** are a basis for \mathbb{R}^2 and **c** is another vector in \mathbb{R}^2 , the space spanned by [**a**, **b**, **c**] is, again, \mathbb{R}^2 . Of course, **c** is superfluous. Nonetheless, any vector in \mathbb{R}^2 can be expressed as a linear combination of **a**, **b**, and **c**. (The linear combination will not be unique. Suppose, for example, that **a** and **c** are also a basis for \mathbb{R}^2 .)

Consider the set of three coordinate vectors whose third element is zero. In particular,

$$\mathbf{a}' = [a_1 \quad a_2 \quad 0] \quad \text{and} \quad \mathbf{b}' = [b_1 \quad b_2 \quad 0].$$

Vectors **a** and **b** do not span the three-dimensional space \mathbb{R}^3 . Every linear combination of **a** and **b** has a third coordinate equal to zero; thus, for instance, $\mathbf{c}' = [1 \quad 2 \quad 3]$ could not be written as a linear combination of **a** and **b**. If $(a_1b_2 - a_2b_1)$ is not equal to zero [see (A-41)], however, then any vector whose third element is zero can be expressed as a linear combination of **a** and **b**. So, although **a** and **b** do not span \mathbb{R}^3 , they do span something, the set of vectors in \mathbb{R}^3 whose third element is zero. This area is a plane (the “floor” of the box in a three-dimensional figure). This plane in \mathbb{R}^3 is a **subspace**, in this instance, a two-dimensional subspace. Note that it is *not* \mathbb{R}^2 ; it is the set of vectors in \mathbb{R}^3 whose third coordinate is 0. Any plane in \mathbb{R}^3 , regardless of how it is oriented, forms a two-dimensional subspace. Any two independent vectors that lie in that subspace will span it. But without a third vector that points in some other direction, we cannot span any more of \mathbb{R}^3 than this two-dimensional part of it. By the same logic, any line in \mathbb{R}^3 is a one-dimensional subspace, in this case, the set of all vectors in \mathbb{R}^3 whose coordinates are multiples of those of the vector that define the line. A subspace is a vector space in all the respects in which we have defined it. We emphasize that it is *not* a vector space of lower dimension. For example, \mathbb{R}^2 is not a subspace of \mathbb{R}^3 . The essential difference is the number of dimensions in the vectors. The vectors in \mathbb{R}^3 that form a two-dimensional subspace are still three-element vectors; they all just happen to lie in the same plane.

The space spanned by a set of vectors in \mathbb{R}^K has at most K dimensions. If this space has fewer than K dimensions, it is a subspace, or **hyperplane**. But the important point in the preceding discussion is that every set of vectors spans some space; it may be the entire space in which the vectors reside, or it may be some subspace of it.

A.3.5 RANK OF A MATRIX

We view a matrix as a set of column vectors. The number of columns in the matrix equals the number of vectors in the set, and the number of rows equals the number of coordinates in each column vector.

DEFINITION A.8 Column Space

The *column space* of a matrix is the vector space that is spanned by its column vectors.

If the matrix contains K rows, its column space might have K dimensions. But, as we have seen, it might have fewer dimensions; the column vectors might be linearly dependent, or there might be fewer than K of them. Consider the matrix

$$\mathbf{A} = \begin{bmatrix} 1 & 5 & 6 \\ 2 & 6 & 8 \\ 7 & 1 & 8 \end{bmatrix}.$$

It contains three vectors from \mathbb{R}^3 , but the third is the sum of the first two, so the column space of this matrix cannot have three dimensions. Nor does it have only one, since the three columns are not all scalar multiples of one another. Hence, it has two, and the column space of this matrix is a two-dimensional subspace of \mathbb{R}^3 .

DEFINITION A.9 Column Rank

The *column rank* of a matrix is the dimension of the vector space that is spanned by its column vectors.

It follows that the column rank of a matrix is equal to the largest number of linearly independent column vectors it contains. The column rank of \mathbf{A} is 2. For another specific example, consider

$$\mathbf{B} = \begin{bmatrix} 1 & 2 & 3 \\ 5 & 1 & 5 \\ 6 & 4 & 5 \\ 3 & 1 & 4 \end{bmatrix}.$$

It can be shown (we shall see how later) that this matrix has a column rank equal to 3. Since each column of \mathbf{B} is a vector in \mathbb{R}^4 , the column space of \mathbf{B} is a three-dimensional subspace of \mathbb{R}^4 .

Consider, instead, the set of vectors obtained by using the *rows* of \mathbf{B} instead of the columns. The new matrix would be

$$\mathbf{C} = \begin{bmatrix} 1 & 5 & 6 & 3 \\ 2 & 1 & 4 & 1 \\ 3 & 5 & 5 & 4 \end{bmatrix}.$$

This matrix is composed of four column vectors from \mathbb{R}^3 . (Note that \mathbf{C} is \mathbf{B}' .) The column space of \mathbf{C} is at most \mathbb{R}^3 , since four vectors in \mathbb{R}^3 must be linearly dependent. In fact, the column space of

\mathbf{C} is \mathbb{R}^3 . Although this is not the same as the column space of \mathbf{B} , it does have the same dimension. Thus, the column rank of \mathbf{C} and the column rank of \mathbf{B} are the same. But the columns of \mathbf{C} are the rows of \mathbf{B} . Thus, the column rank of \mathbf{C} equals the **row rank** of \mathbf{B} . That the column and row ranks of \mathbf{B} are the same is not a coincidence. The general results (which are equivalent) are as follows.

THEOREM A.1 Equality of Row and Column Rank

The column rank and row rank of a matrix are equal. By the definition of row rank and its counterpart for column rank, we obtain the corollary,

the row space and column space of a matrix have the same dimension. (A-42)

Theorem A.1 holds regardless of the actual row and column rank. If the column rank of a matrix happens to equal the number of columns it contains, then the matrix is said to have **full column rank**. **Full row rank** is defined likewise. Since the row and column ranks of a matrix are always equal, we can speak unambiguously of the **rank of a matrix**. For either the row rank or the column rank (and, at this point, we shall drop the distinction),

$$\text{rank}(\mathbf{A}) = \text{rank}(\mathbf{A}') \leq \min(\text{number of rows, number of columns}). \quad (\text{A-43})$$

In most contexts, we shall be interested in the columns of the matrices we manipulate. We shall use the term **full rank** to describe a matrix whose rank is equal to the number of columns it contains.

Of particular interest will be the distinction between **full rank** and **short rank matrices**. The distinction turns on the solutions to $\mathbf{Ax} = \mathbf{0}$. If a nonzero \mathbf{x} for which $\mathbf{Ax} = \mathbf{0}$ exists, then \mathbf{A} does not have full rank. Equivalently, if the nonzero \mathbf{x} exists, then the columns of \mathbf{A} are linearly dependent and at least one of them can be expressed as a linear combination of the others. For example, a nonzero set of solutions to

$$\begin{bmatrix} 1 & 3 & 10 \\ 2 & 3 & 14 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

is any multiple of $\mathbf{x}' = (2, 1, -\frac{1}{2})$.

In a product matrix $\mathbf{C} = \mathbf{AB}$, every column of \mathbf{C} is a linear combination of the columns of \mathbf{A} , so each column of \mathbf{C} is in the column space of \mathbf{A} . It is possible that the set of columns in \mathbf{C} could span this space, but it is not possible for them to span a higher-dimensional space. At best, they could be a full set of linearly independent vectors in \mathbf{A} 's column space. We conclude that the column rank of \mathbf{C} could not be greater than that of \mathbf{A} . Now, apply the same logic to the rows of \mathbf{C} , which are all linear combinations of the rows of \mathbf{B} . For the same reason that the column rank of \mathbf{C} cannot exceed the column rank of \mathbf{A} , the row rank of \mathbf{C} cannot exceed the row rank of \mathbf{B} . Since row and column ranks are always equal, we conclude that

$$\text{rank}(\mathbf{AB}) \leq \min(\text{rank}(\mathbf{A}), \text{rank}(\mathbf{B})). \quad (\text{A-44})$$

A useful corollary to (A-44) is:

$$\text{If } \mathbf{A} \text{ is } M \times n \text{ and } \mathbf{B} \text{ is a square matrix of rank } n, \text{ then } \text{rank}(\mathbf{AB}) = \text{rank}(\mathbf{A}). \quad (\text{A-45})$$

Another application that plays a central role in the development of regression analysis is, for any matrix \mathbf{A} ,

$$\text{rank}(\mathbf{A}) = \text{rank}(\mathbf{A}'\mathbf{A}) = \text{rank}(\mathbf{A}\mathbf{A}'). \quad (\text{A-46})$$

A.3.6 DETERMINANT OF A MATRIX

The determinant of a square matrix—determinants are not defined for nonsquare matrices—is a function of the elements of the matrix. There are various definitions, most of which are not useful for our work. Determinants figure into our results in several ways, however, that we can enumerate before we need formally to define the computations.

PROPOSITION

The determinant of a matrix is nonzero if and only if it has full rank.

Full rank and short rank matrices can be distinguished by whether or not their determinants are nonzero. There are some settings in which the value of the determinant is also of interest, so we now consider some algebraic results.

It is most convenient to begin with a diagonal matrix

$$\mathbf{D} = \begin{bmatrix} d_1 & 0 & 0 & \cdots & 0 \\ 0 & d_2 & 0 & \cdots & 0 \\ & & & \cdots & \\ 0 & 0 & 0 & \cdots & d_K \end{bmatrix}.$$

The column vectors of \mathbf{D} define a “box” in \mathbb{R}^K whose sides are all at right angles to one another.³ Its “volume,” or determinant, is simply the product of the lengths of the sides, which we denote

$$|\mathbf{D}| = d_1 d_2 \cdots d_K = \prod_{k=1}^K d_k. \quad (\text{A-47})$$

A special case is the identity matrix, which has, regardless of K , $|\mathbf{I}_K| = 1$. Multiplying \mathbf{D} by a scalar c is equivalent to multiplying the length of each side of the box by c , which would multiply its volume by c^K . Thus,

$$|c\mathbf{D}| = c^K |\mathbf{D}|. \quad (\text{A-48})$$

Continuing with this admittedly special case, we suppose that only one column of \mathbf{D} is multiplied by c . In two dimensions, this would make the box wider but not higher, or vice versa. Hence, the “volume” (area) would also be multiplied by c . Now, suppose that each side of the box were multiplied by a different c , the first by c_1 , the second by c_2 , and so on. The volume would, by an obvious extension, now be $c_1 c_2 \cdots c_K |\mathbf{D}|$. The matrix with columns defined by $[c_1 \mathbf{d}_1 \ c_2 \mathbf{d}_2 \ \cdots]$ is just \mathbf{DC} , where \mathbf{C} is a diagonal matrix with c_i as its i th diagonal element. The computation just described is, therefore,

$$|\mathbf{DC}| = |\mathbf{D}| \cdot |\mathbf{C}|. \quad (\text{A-49})$$

(The determinant of \mathbf{C} is the product of the c_i s since \mathbf{C} , like \mathbf{D} , is a diagonal matrix.) In particular, note what happens to the whole thing if one of the c_i s is zero.

³Each column vector defines a segment on one of the axes.

For 2×2 matrices, the computation of the determinant is

$$\begin{vmatrix} a & c \\ b & d \end{vmatrix} = ad - bc. \quad (\text{A-50})$$

Notice that it is a function of all the elements of the matrix. This statement will be true, in general. For more than two dimensions, the determinant can be obtained by using an **expansion by cofactors**. Using *any* row, say i , we obtain

$$|\mathbf{A}| = \sum_{k=1}^K a_{ik}(-1)^{i+k}|\mathbf{A}_{ik}|, \quad k = 1, \dots, K, \quad (\text{A-51})$$

where \mathbf{A}_{ik} is the matrix obtained from \mathbf{A} by deleting row i and column k . The determinant of \mathbf{A}_{ik} is called a **minor** of \mathbf{A} .⁴ When the correct sign, $(-1)^{i+k}$, is added, it becomes a **cofactor**. This operation can be done using any column as well. For example, a 4×4 determinant becomes a sum of four 3×3 s, whereas a 5×5 is a sum of five 4×4 s, each of which is a sum of four 3×3 s, and so on. Obviously, it is a good idea to base (A-51) on a row or column with many zeros in it, if possible. In practice, this rapidly becomes a heavy burden. It is unlikely, though, that you will ever calculate any determinants over 3×3 without a computer. A 3×3 , however, might be computed on occasion; if so, the following shortcut will prove useful:

$$\begin{vmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{vmatrix} = a_{11}a_{22}a_{33} + a_{12}a_{23}a_{31} + a_{13}a_{32}a_{21} - a_{31}a_{22}a_{13} - a_{21}a_{12}a_{33} - a_{11}a_{23}a_{32}.$$

Although (A-48) and (A-49) were given for diagonal matrices, they hold for general matrices \mathbf{C} and \mathbf{D} . One special case of (A-48) to note is that of $c = -1$. Multiplying a matrix by -1 does not necessarily change the sign of its determinant. It does so only if the order of the matrix is odd. By using the expansion by cofactors formula, an additional result can be shown:

$$|\mathbf{A}| = |\mathbf{A}'| \quad (\text{A-52})$$

A.3.7 A LEAST SQUARES PROBLEM

Given a vector \mathbf{y} and a matrix \mathbf{X} , we are interested in expressing \mathbf{y} as a linear combination of the columns of \mathbf{X} . There are two possibilities. If \mathbf{y} lies in the column space of \mathbf{X} , then we shall be able to find a vector \mathbf{b} such that

$$\mathbf{y} = \mathbf{Xb}. \quad (\text{A-53})$$

Figure A.3 illustrates such a case for three dimensions in which the two columns of \mathbf{X} both have a third coordinate equal to zero. Only \mathbf{y} s whose third coordinate is zero, such as \mathbf{y}^0 in the figure, can be expressed as \mathbf{Xb} for some \mathbf{b} . For the general case, assuming that \mathbf{y} is, indeed, in the column space of \mathbf{X} , we can find the coefficients \mathbf{b} by solving the set of equations in (A-53). The solution is discussed in the next section.

Suppose, however, that \mathbf{y} is not in the column space of \mathbf{X} . In the context of this example, suppose that \mathbf{y} 's third component is not zero. Then there is no \mathbf{b} such that (A-53) holds. We can, however, write

$$\mathbf{y} = \mathbf{Xb} + \mathbf{e}, \quad (\text{A-54})$$

where \mathbf{e} is the difference between \mathbf{y} and \mathbf{Xb} . By this construction, we find an \mathbf{Xb} that is in the column space of \mathbf{X} , and \mathbf{e} is the difference, or "residual." Figure A.3 shows two examples, \mathbf{y} and

⁴If i equals j , then the determinant is a **principal minor**.




FIGURE A.3 Least Squares Projections.

y^* . For the present, we consider only y . We are interested in finding the \mathbf{b} such that \mathbf{y} is as close as possible to \mathbf{Xb} in the sense that \mathbf{e} is as short as possible.

DEFINITION A.10 Length of a Vector

The length, or *norm*, of a vector \mathbf{e} is

$$\|\mathbf{e}\| = \sqrt{\mathbf{e}'\mathbf{e}}. \tag{A-55}$$

The problem is to find the \mathbf{b} for which

$$\|\mathbf{e}\| = \|\mathbf{y} - \mathbf{Xb}\|$$

is as small as possible. The solution is that \mathbf{b} that makes \mathbf{e} perpendicular, or *orthogonal*, to \mathbf{Xb} .

DEFINITION A.11 Orthogonal Vectors

Two nonzero vectors \mathbf{a} and \mathbf{b} are *orthogonal*, written $\mathbf{a} \perp \mathbf{b}$, if and only if

$$\mathbf{a}'\mathbf{b} = \mathbf{b}'\mathbf{a} = 0.$$

Returning once again to our fitting problem, we find that the \mathbf{b} we seek is that for which

$$\mathbf{e} \perp \mathbf{Xb}.$$

Expanding this set of equations gives the requirement

$$\begin{aligned} (\mathbf{Xb})'\mathbf{e} &= 0 \\ &= \mathbf{b}'\mathbf{X}'\mathbf{y} - \mathbf{b}'\mathbf{X}'\mathbf{Xb} \\ &= \mathbf{b}'[\mathbf{X}'\mathbf{y} - \mathbf{X}'\mathbf{Xb}] \end{aligned}$$

or, assuming \mathbf{b} is not $\mathbf{0}$, the set of equations

$$\mathbf{X}'\mathbf{y} = \mathbf{X}'\mathbf{X}\mathbf{b}.$$

The means of solving such a set of equations is the subject of Section A.5.

In Figure A.3, the linear combination $\mathbf{X}\mathbf{b}$ is called the **projection** of \mathbf{y} into the column space of \mathbf{X} . The figure is drawn so that, although \mathbf{y} and \mathbf{y}^* are different, they are similar in that the projection of \mathbf{y} lies on top of that of \mathbf{y}^* . The question we wish to pursue here is, Which vector, \mathbf{y} or \mathbf{y}^* , is closer to its projection in the column space of \mathbf{X} ? Superficially, it would appear that \mathbf{y} is closer, since \mathbf{e} is shorter than \mathbf{e}^* . Yet \mathbf{y}^* is much more nearly parallel to its projection than \mathbf{y} , so the only reason that its residual vector is longer is that \mathbf{y}^* is longer compared with \mathbf{y} . A measure of comparison that would be unaffected by the length of the vectors is the angle between the vector and its projection (assuming that angle is not zero). By this measure, θ^* is smaller than θ , which would reverse the earlier conclusion.

THEOREM A.2 The Cosine Law

The angle θ between two vectors \mathbf{a} and \mathbf{b} satisfies

$$\cos \theta = \frac{\mathbf{a}'\mathbf{b}}{\|\mathbf{a}\| \cdot \|\mathbf{b}\|}.$$

The two vectors in the calculation would be \mathbf{y} or \mathbf{y}^* and $\mathbf{X}\mathbf{b}$ or $(\mathbf{X}\mathbf{b})^*$. A zero cosine implies that the vectors are orthogonal. If the cosine is one, then the angle is zero, which means that the vectors are the same. (They would be if \mathbf{y} were in the column space of \mathbf{X} .) By dividing by the lengths, we automatically compensate for the length of \mathbf{y} . By this measure, we find in Figure A.3 that \mathbf{y}^* is closer to its projection, $(\mathbf{X}\mathbf{b})^*$ than \mathbf{y} is to its projection, $\mathbf{X}\mathbf{b}$.

A.4 SOLUTION OF A SYSTEM OF LINEAR EQUATIONS

Consider the set of n linear equations

$$\mathbf{A}\mathbf{x} = \mathbf{b}, \tag{A-56}$$

in which the K elements of \mathbf{x} constitute the unknowns. \mathbf{A} is a known matrix of coefficients, and \mathbf{b} is a specified vector of values. We are interested in knowing whether a solution exists; if so, then how to obtain it; and finally, if it does exist, then whether it is unique.

A.4.1 SYSTEMS OF LINEAR EQUATIONS

For most of our applications, we shall consider only square systems of equations, that is, those in which \mathbf{A} is a square matrix. In what follows, therefore, we take n to equal K . Since the number of rows in \mathbf{A} is the number of equations, whereas the number of columns in \mathbf{A} is the number of variables, this case is the familiar one of “ n equations in n unknowns.”

There are two types of systems of equations.

DEFINITION A.12 Homogeneous Equation System

A homogeneous system is of the form $\mathbf{Ax} = \mathbf{0}$.

By definition, a nonzero solution to such a system will exist if and only if \mathbf{A} does not have full rank. If so, then for at least one column of \mathbf{A} , we can write the preceding as

$$\mathbf{a}_k = - \sum_{m \neq k} \frac{x_m}{x_k} \mathbf{a}_m.$$

This means, as we know, that the columns of \mathbf{A} are linearly dependent and that $|\mathbf{A}| = 0$.

DEFINITION A.13 Nonhomogeneous Equation System

A nonhomogeneous system of equations is of the form $\mathbf{Ax} = \mathbf{b}$, where \mathbf{b} is a nonzero vector.

The vector \mathbf{b} is chosen arbitrarily and is to be expressed as a linear combination of the columns of \mathbf{A} . Since \mathbf{b} has K elements, this solution will exist only if the columns of \mathbf{A} span the entire K -dimensional space, \mathbb{R}^K .⁵ Equivalently, we shall require that the columns of \mathbf{A} be linearly independent or that $|\mathbf{A}|$ not be equal to zero.

A.4.2 INVERSE MATRICES

To solve the system $\mathbf{Ax} = \mathbf{b}$ for \mathbf{x} , something akin to division by a matrix is needed. Suppose that we could find a square matrix \mathbf{B} such that $\mathbf{BA} = \mathbf{I}$. If the equation system is premultiplied by this \mathbf{B} , then the following would be obtained:

$$\mathbf{BAx} = \mathbf{Ix} = \mathbf{x} = \mathbf{Bb}. \quad (\text{A-57})$$

If the matrix \mathbf{B} exists, then it is the **inverse** of \mathbf{A} , denoted

$$\mathbf{B} = \mathbf{A}^{-1}.$$

From the definition,

$$\mathbf{A}^{-1}\mathbf{A} = \mathbf{I}.$$

In addition, by premultiplying by \mathbf{A} , postmultiplying by \mathbf{A}^{-1} , and then canceling terms, we find

$$\mathbf{AA}^{-1} = \mathbf{I}$$

as well.

If the inverse exists, then it must be unique. Suppose that it is not and that \mathbf{C} is a different inverse of \mathbf{A} . Then $\mathbf{CAB} = \mathbf{CAB}$, but $(\mathbf{CA})\mathbf{B} = \mathbf{IB} = \mathbf{B}$ and $\mathbf{C}(\mathbf{AB}) = \mathbf{C}$, which would be a

⁵If \mathbf{A} does not have full rank, then the nonhomogeneous system will have solutions for *some* vectors \mathbf{b} , namely, any \mathbf{b} in the column space of \mathbf{A} . But we are interested in the case in which there are solutions for *all* nonzero vectors \mathbf{b} , which requires \mathbf{A} to have full rank.

contradiction if \mathbf{C} did not equal \mathbf{B} . Since, by (A-57), the solution is $\mathbf{x} = \mathbf{A}^{-1}\mathbf{b}$, the solution to the equation system is unique as well.

We now consider the calculation of the inverse matrix. For a 2×2 matrix, $\mathbf{AB} = \mathbf{I}$ implies that

$$\begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad \text{or} \quad \begin{cases} a_{11}b_{11} + a_{12}b_{21} = 1 \\ a_{11}b_{12} + a_{12}b_{22} = 0 \\ a_{21}b_{11} + a_{22}b_{21} = 0 \\ a_{21}b_{12} + a_{22}b_{22} = 1 \end{cases}$$

The solutions are

$$\begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix} = \frac{1}{a_{11}a_{22} - a_{12}a_{21}} \begin{bmatrix} a_{22} & -a_{12} \\ -a_{21} & a_{11} \end{bmatrix} = \frac{1}{|\mathbf{A}|} \begin{bmatrix} a_{22} & -a_{12} \\ -a_{21} & a_{11} \end{bmatrix}. \tag{A-58}$$

Notice the presence of the reciprocal of $|\mathbf{A}|$ in \mathbf{A}^{-1} . This result is not specific to the 2×2 case. We infer from it that if the determinant is zero, then the inverse does not exist.

DEFINITION A.14 Nonsingular Matrix
A matrix is nonsingular if and only if its inverse exists.

The simplest inverse matrix to compute is that of a diagonal matrix. If

$$\mathbf{D} = \begin{bmatrix} d_1 & 0 & 0 & \dots & 0 \\ 0 & d_2 & 0 & \dots & 0 \\ & & \dots & & \\ 0 & 0 & 0 & \dots & d_K \end{bmatrix}, \quad \text{then} \quad \mathbf{D}^{-1} = \begin{bmatrix} 1/d_1 & 0 & 0 & \dots & 0 \\ 0 & 1/d_2 & 0 & \dots & 0 \\ & & \dots & & \\ 0 & 0 & 0 & \dots & 1/d_K \end{bmatrix},$$

which shows, incidentally, that $\mathbf{I}^{-1} = \mathbf{I}$.

We shall use a^{ik} to indicate the ik th element of \mathbf{A}^{-1} . The general formula for computing an inverse matrix is

$$a^{ik} = \frac{|C_{ki}|}{|\mathbf{A}|}, \tag{A-59}$$

where $|C_{ki}|$ is the ki th cofactor of \mathbf{A} . It follows, therefore, that for \mathbf{A} to be nonsingular, $|\mathbf{A}|$ must be nonzero. Notice the reversal of the subscripts

Some computational results involving inverses are

$$|\mathbf{A}^{-1}| = \frac{1}{|\mathbf{A}|}, \tag{A-60}$$

$$(\mathbf{A}^{-1})^{-1} = \mathbf{A}, \tag{A-61}$$

$$(\mathbf{A}^{-1})' = (\mathbf{A}')^{-1}. \tag{A-62}$$

$$\text{If } \mathbf{A} \text{ is symmetric, then } \mathbf{A}^{-1} \text{ is symmetric.} \tag{A-63}$$

When both inverse matrices exist,

$$(\mathbf{AB})^{-1} = \mathbf{B}^{-1}\mathbf{A}^{-1}. \tag{A-64}$$

822 APPENDIX A ♦ Matrix Algebra

Note the condition preceding (A-64). It may be that \mathbf{AB} is a square, nonsingular matrix when neither \mathbf{A} nor \mathbf{B} are even square. (Consider, for example, $\mathbf{A}'\mathbf{A}$.) Extending (A-64), we have

$$(\mathbf{ABC})^{-1} = \mathbf{C}^{-1}(\mathbf{AB})^{-1} = \mathbf{C}^{-1}\mathbf{B}^{-1}\mathbf{A}^{-1}. \quad (\text{A-65})$$

Recall that for a data matrix \mathbf{X} , $\mathbf{X}'\mathbf{X}$ is the sum of the *outer products* of the rows \mathbf{X} . Suppose that we have already computed $\mathbf{S} = (\mathbf{X}'\mathbf{X})^{-1}$ for a number of years of data, such as those given at the beginning of this appendix. The following result, which is called an **updating formula**, shows how to compute the new \mathbf{S} that would result when a new row is added to \mathbf{X} :

$$[\mathbf{A} \pm \mathbf{bb}']^{-1} = \mathbf{A}^{-1} \mp \left[\frac{1}{1 \pm \mathbf{b}'\mathbf{A}^{-1}\mathbf{b}} \right] \mathbf{A}^{-1}\mathbf{bb}'\mathbf{A}^{-1}. \quad (\text{A-66})$$

Note the reversal of the sign in the inverse. Two more general forms of (A-66) that are occasionally useful are

$$[\mathbf{A} \pm \mathbf{bc}']^{-1} = \mathbf{A}^{-1} \mp \left[\frac{1}{1 \pm \mathbf{c}'\mathbf{A}^{-1}\mathbf{b}} \right] \mathbf{A}^{-1}\mathbf{bc}'\mathbf{A}^{-1}. \quad (\text{A-66a})$$

$$[\mathbf{A} \pm \mathbf{BCB}']^{-1} = \mathbf{A}^{-1} \mp \mathbf{A}^{-1}\mathbf{B}[\mathbf{C}^{-1} \pm \mathbf{B}'\mathbf{A}^{-1}\mathbf{B}]^{-1}\mathbf{B}'\mathbf{A}^{-1}. \quad (\text{A-66b})$$

A.4.3 NONHOMOGENEOUS SYSTEMS OF EQUATIONS

For the nonhomogeneous system

$$\mathbf{Ax} = \mathbf{b},$$

if \mathbf{A} is nonsingular, then the unique solution is

$$\mathbf{x} = \mathbf{A}^{-1}\mathbf{b}.$$

A.4.4 SOLVING THE LEAST SQUARES PROBLEM

We now have the tool needed to solve the least squares problem posed in Section A3.7. We found the solution vector, \mathbf{b} to be the solution to the nonhomogenous system $\mathbf{X}'\mathbf{y} = \mathbf{X}'\mathbf{X}\mathbf{b}$. Let \mathbf{z} equal the vector $\mathbf{X}'\mathbf{y}$ and let \mathbf{A} equal the square matrix $\mathbf{X}'\mathbf{X}$. The equation system is then

$$\mathbf{Ab} = \mathbf{a}.$$

By the results above, if \mathbf{A} is nonsingular, then

$$\mathbf{b} = \mathbf{A}^{-1}\mathbf{a} = (\mathbf{X}'\mathbf{X})^{-1}(\mathbf{X}'\mathbf{y})$$

assuming that the matrix to be inverted is nonsingular. We have reached the irreducible minimum. If the columns of \mathbf{X} are linearly independent, that is, if \mathbf{X} has full rank, then this is the solution to the least squares problem. If the columns of \mathbf{X} are linearly dependent, then this system has no unique solution.

A.5 PARTITIONED MATRICES

In formulating the elements of a matrix—it is sometimes useful to group some of the elements in **submatrices**. Let

$$\mathbf{A} = \left[\begin{array}{cc|c} 1 & 4 & 5 \\ 2 & 9 & 3 \\ 8 & 9 & 6 \end{array} \right] = \begin{bmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix}.$$

\mathbf{A} is a **partitioned matrix**. The subscripts of the submatrices are defined in the same fashion as those for the elements of a matrix. A common special case is the **block diagonal matrix**:

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_{11} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{22} \end{bmatrix},$$

where \mathbf{A}_{11} and \mathbf{A}_{22} are square matrices.

A.5.1 ADDITION AND MULTIPLICATION OF PARTITIONED MATRICES

For conformably partitioned matrices \mathbf{A} and \mathbf{B} ,

$$\mathbf{A} + \mathbf{B} = \begin{bmatrix} \mathbf{A}_{11} + \mathbf{B}_{11} & \mathbf{A}_{12} + \mathbf{B}_{12} \\ \mathbf{A}_{21} + \mathbf{B}_{21} & \mathbf{A}_{22} + \mathbf{B}_{22} \end{bmatrix} \quad (\text{A-67})$$

and

$$\mathbf{AB} = \begin{bmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix} \begin{bmatrix} \mathbf{B}_{11} & \mathbf{B}_{12} \\ \mathbf{B}_{21} & \mathbf{B}_{22} \end{bmatrix} = \begin{bmatrix} \mathbf{A}_{11}\mathbf{B}_{11} + \mathbf{A}_{12}\mathbf{B}_{21} & \mathbf{A}_{11}\mathbf{B}_{12} + \mathbf{A}_{12}\mathbf{B}_{22} \\ \mathbf{A}_{21}\mathbf{B}_{11} + \mathbf{A}_{22}\mathbf{B}_{21} & \mathbf{A}_{21}\mathbf{B}_{12} + \mathbf{A}_{22}\mathbf{B}_{22} \end{bmatrix}. \quad (\text{A-68})$$

In all these, the matrices must be conformable for the operations involved. For addition, the dimensions of \mathbf{A}_{ik} and \mathbf{B}_{ik} must be the same. For multiplication, the number of columns in \mathbf{A}_{ij} must equal the number of rows in \mathbf{B}_{ji} for all pairs i and j . That is, all the necessary matrix products of the submatrices must be defined. Two cases frequently encountered are of the form

$$\begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{bmatrix}' \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{bmatrix} = [\mathbf{A}'_1 \quad \mathbf{A}'_2] \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{bmatrix} = [\mathbf{A}'_1\mathbf{A}_1 + \mathbf{A}'_2\mathbf{A}_2] \quad (\text{A-69})$$

and

$$\begin{bmatrix} \mathbf{A}_{11} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{22} \end{bmatrix}' \begin{bmatrix} \mathbf{A}_{11} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{22} \end{bmatrix} = \begin{bmatrix} \mathbf{A}'_{11}\mathbf{A}_{11} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}'_{22}\mathbf{A}_{22} \end{bmatrix}. \quad (\text{A-70})$$

A.5.2 DETERMINANTS OF PARTITIONED MATRICES

The determinant of a block diagonal matrix is obtained analogously to that of a diagonal matrix:

$$\begin{vmatrix} \mathbf{A}_{11} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{22} \end{vmatrix} = |\mathbf{A}_{11}| \cdot |\mathbf{A}_{22}|. \quad (\text{A-71})$$

The determinant of a general 2×2 partitioned matrix is

$$\begin{vmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{vmatrix} = |\mathbf{A}_{22}| \cdot |\mathbf{A}_{11} - \mathbf{A}_{12}\mathbf{A}_{22}^{-1}\mathbf{A}_{21}| = |\mathbf{A}_{11}| \cdot |\mathbf{A}_{22} - \mathbf{A}_{21}\mathbf{A}_{11}^{-1}\mathbf{A}_{12}|. \quad (\text{A-72})$$

A.5.3 INVERSES OF PARTITIONED MATRICES

The inverse of a block diagonal matrix is

$$\begin{bmatrix} \mathbf{A}_{11} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{22} \end{bmatrix}^{-1} = \begin{bmatrix} \mathbf{A}_{11}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{22}^{-1} \end{bmatrix}, \quad (\text{A-73})$$

which can be verified by direct multiplication.

824 APPENDIX A ♦ Matrix Algebra

For the general 2×2 partitioned matrix, one form of the **partitioned inverse** is

$$\begin{bmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix}^{-1} = \begin{bmatrix} \mathbf{A}_{11}^{-1}(\mathbf{I} + \mathbf{A}_{12}\mathbf{F}_2\mathbf{A}_{21}\mathbf{A}_{11}^{-1}) & -\mathbf{A}_{11}^{-1}\mathbf{A}_{12}\mathbf{F}_2 \\ -\mathbf{F}_2\mathbf{A}_{21}\mathbf{A}_{11}^{-1} & \mathbf{F}_2 \end{bmatrix}, \quad (\text{A-74})$$

where

$$\mathbf{F}_2 = (\mathbf{A}_{22} - \mathbf{A}_{21}\mathbf{A}_{11}^{-1}\mathbf{A}_{12})^{-1}.$$

The upper left block could also be written as

$$\mathbf{F}_1 = (\mathbf{A}_{11} - \mathbf{A}_{12}\mathbf{A}_{22}^{-1}\mathbf{A}_{21})^{-1}.$$

A.5.4 DEVIATIONS FROM MEANS

Suppose that we begin with a column vector of n values \mathbf{x} and let

$$\mathbf{A} = \begin{bmatrix} n & \sum_{i=1}^n x_i \\ \sum_{i=1}^n x_i & \sum_{i=1}^n x_i^2 \end{bmatrix} = \begin{bmatrix} \mathbf{i}'\mathbf{i} & \mathbf{i}'\mathbf{x} \\ \mathbf{x}'\mathbf{i} & \mathbf{x}'\mathbf{x} \end{bmatrix}.$$

We are interested in the lower right-hand element of \mathbf{A}^{-1} . Upon using the definition of \mathbf{F}_2 in (A-74), this is

$$\begin{aligned} \mathbf{F}_2 &= [\mathbf{x}'\mathbf{x} - (\mathbf{x}'\mathbf{i})(\mathbf{i}'\mathbf{i})^{-1}(\mathbf{i}'\mathbf{x})]^{-1} = \left\{ \mathbf{x}' \left[\mathbf{I}\mathbf{x} - \mathbf{i} \left(\frac{1}{n} \right) \mathbf{i}'\mathbf{x} \right] \right\}^{-1} \\ &= \left\{ \mathbf{x}' \left[\mathbf{I} - \left(\frac{1}{n} \right) \mathbf{i}\mathbf{i}' \right] \mathbf{x} \right\}^{-1} = (\mathbf{x}'\mathbf{M}^0\mathbf{x})^{-1}. \end{aligned}$$

Therefore, the lower right-hand value in the inverse matrix is

$$(\mathbf{x}'\mathbf{M}^0\mathbf{x})^{-1} = \frac{1}{\sum_{i=1}^n (x_i - \bar{x})^2} = a^{22}.$$

Now, suppose that we replace \mathbf{x} with \mathbf{X} , a matrix with several columns. We seek the lower right block of $(\mathbf{Z}'\mathbf{Z})^{-1}$, where $\mathbf{Z} = [\mathbf{i}, \mathbf{X}]$. The analogous result is

$$(\mathbf{Z}'\mathbf{Z})^{22} = [\mathbf{X}'\mathbf{X} - \mathbf{X}'\mathbf{i}(\mathbf{i}'\mathbf{i})^{-1}\mathbf{i}'\mathbf{X}]^{-1} = (\mathbf{X}'\mathbf{M}^0\mathbf{X})^{-1},$$

which implies that the $K \times K$ matrix in the lower right corner of $(\mathbf{Z}'\mathbf{Z})^{-1}$ is the inverse of the $K \times K$ matrix whose jk th element is $\sum_{i=1}^n (x_{ij} - \bar{x}_j)(x_{ik} - \bar{x}_k)$. Thus, when a data matrix contains a column of ones, the elements of the inverse of the matrix of sums of squares and cross products will be computed from the original data in the form of deviations from the respective column means.

A.5.5 KRONECKER PRODUCTS

A calculation that helps to condense the notation when dealing with sets of regression models (see Chapters 14 and 15) is the **Kronecker product**. For general matrices \mathbf{A} and \mathbf{B} ,

$$\mathbf{A} \otimes \mathbf{B} = \begin{bmatrix} a_{11}\mathbf{B} & a_{12}\mathbf{B} & \cdots & a_{1K}\mathbf{B} \\ a_{21}\mathbf{B} & a_{22}\mathbf{B} & \cdots & a_{2K}\mathbf{B} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1}\mathbf{B} & a_{n2}\mathbf{B} & \cdots & a_{nK}\mathbf{B} \end{bmatrix}. \quad (\text{A-75})$$

Notice that there is no requirement for conformability in this operation. The Kronecker product can be computed for any pair of matrices. If \mathbf{A} is $K \times L$ and \mathbf{B} is $m \times n$, then $\mathbf{A} \otimes \mathbf{B}$ is $(Km) \times (Ln)$.

For the Kronecker product,

$$(\mathbf{A} \otimes \mathbf{B})^{-1} = (\mathbf{A}^{-1} \otimes \mathbf{B}^{-1}), \tag{A-76}$$

If \mathbf{A} is $M \times M$ and \mathbf{B} is $n \times n$, then

$$|\mathbf{A} \otimes \mathbf{B}| = |\mathbf{A}|^n |\mathbf{B}|^M,$$

$$(\mathbf{A} \otimes \mathbf{B})' = \mathbf{A}' \otimes \mathbf{B}'$$

$$\text{trace}(\mathbf{A} \otimes \mathbf{B}) = \text{tr}(\mathbf{A})\text{tr}(\mathbf{B}).$$

For \mathbf{A} , \mathbf{B} , \mathbf{C} , and \mathbf{D} such that the products are defined is

$$(\mathbf{A} \otimes \mathbf{B})(\mathbf{C} \otimes \mathbf{D}) = \mathbf{AC} \otimes \mathbf{BD}.$$

A.6 CHARACTERISTIC ROOTS AND VECTORS

A useful set of results for analyzing a square matrix \mathbf{A} arises from the solutions to the set of equations

$$\mathbf{Ac} = \lambda \mathbf{c}. \tag{A-77}$$

The pairs of solutions are the **characteristic vectors** \mathbf{c} and **characteristic roots** λ . If \mathbf{c} is any solution vector, then $k\mathbf{c}$ is also for any value of k . To remove the indeterminacy, \mathbf{c} is **normalized** so that $\mathbf{c}'\mathbf{c} = 1$.

The solution then consists of λ and the $n - 1$ unknown elements in \mathbf{c} .

A.6.1 THE CHARACTERISTIC EQUATION

Solving (A-77) can, in principle, proceed as follows. First, (A-77) implies that

$$\mathbf{Ac} = \lambda \mathbf{Ic}$$

or that

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

This equation is a homogeneous system that has a nonzero solution only if the matrix $(\mathbf{A} - \lambda \mathbf{I})$ is singular or has a zero determinant. Therefore, if λ is a solution, then

$$|\mathbf{A} - \lambda \mathbf{I}| = 0. \tag{A-78}$$

This polynomial in λ is the **characteristic equation** of \mathbf{A} . For example, if

$$\mathbf{A} = \begin{bmatrix} 5 & 1 \\ 2 & 4 \end{bmatrix},$$

then

$$|\mathbf{A} - \lambda \mathbf{I}| = \begin{vmatrix} 5 - \lambda & 1 \\ 2 & 4 - \lambda \end{vmatrix} = (5 - \lambda)(4 - \lambda) - 2(1) = \lambda^2 - 9\lambda + 18.$$

The two solutions are $\lambda = 6$ and $\lambda = 3$.

In solving the characteristic equation, there is no guarantee that the characteristic roots will be real. In the preceding example, if the 2 in the lower left-hand corner of the matrix were -2 instead, then the solution would be a pair of complex values. The same result can emerge in the general $n \times n$ case. The characteristic roots of a symmetric matrix are real, however.⁶ This result will be convenient because most of our applications will involve the characteristic roots and vectors of symmetric matrices.

For an $n \times n$ matrix, the characteristic equation is an n th-order polynomial in λ . Its solutions may be n distinct values, as in the preceding example, or may contain repeated values of λ , and may contain some zeros as well.

A.6.2 CHARACTERISTIC VECTORS

With λ in hand, the characteristic vectors are derived from the original problem,

$$\mathbf{A}\mathbf{c} = \lambda\mathbf{c}$$

or

$$(\mathbf{A} - \lambda\mathbf{I})\mathbf{c} = \mathbf{0}. \quad (\text{A-79})$$

Neither pair determines the values of c_1 and c_2 . But this result was to be expected; it was the reason $\mathbf{c}'\mathbf{c} = 1$ was specified at the outset. The additional equation $\mathbf{c}'\mathbf{c} = 1$, however, produces complete solutions for the vectors.

A.6.3 GENERAL RESULTS FOR CHARACTERISTIC ROOTS AND VECTORS

A $K \times K$ symmetric matrix has K distinct characteristic vectors, $\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_K$. The corresponding characteristic roots, $\lambda_1, \lambda_2, \dots, \lambda_K$, although real, need not be distinct. The characteristic vectors of a symmetric matrix are orthogonal,⁷ which implies that for every $i \neq j$, $\mathbf{c}_i'\mathbf{c}_j = 0$.⁸ It is convenient to collect the K -characteristic vectors in a $K \times K$ matrix whose i th column is the \mathbf{c}_i corresponding to λ_i ,

$$\mathbf{C} = [\mathbf{c}_1 \quad \mathbf{c}_2 \quad \dots \quad \mathbf{c}_K],$$

and the K -characteristic roots in the same order, in a diagonal matrix,

$$\mathbf{\Lambda} = \begin{bmatrix} \lambda_1 & 0 & 0 & \dots & 0 \\ 0 & \lambda_2 & 0 & \dots & 0 \\ & & & \dots & \\ 0 & 0 & 0 & \dots & \lambda_K \end{bmatrix}.$$

Then, the full set of equations

$$\mathbf{A}\mathbf{c}_k = \lambda_k\mathbf{c}_k$$

is contained in

$$\mathbf{A}\mathbf{C} = \mathbf{C}\mathbf{\Lambda}. \quad (\text{A-80})$$

⁶A proof may be found in Theil (1971).

⁷For proofs of these propositions, see Strang (1988).

⁸This statement is not true if the matrix is not symmetric. For instance, it does not hold for the characteristic vectors computed in the first example. For nonsymmetric matrices, there is also a distinction between "right" characteristic vectors, $\mathbf{A}\mathbf{c} = \lambda\mathbf{c}$, and "left" characteristic vectors, $\mathbf{d}'\mathbf{A} = \lambda\mathbf{d}'$, which may not be equal.

Since the vectors are orthogonal and $\mathbf{c}'_i \mathbf{c}_i = 1$, we have

$$\mathbf{C}'\mathbf{C} = \begin{bmatrix} \mathbf{c}'_1 \mathbf{c}_1 & \mathbf{c}'_1 \mathbf{c}_2 & \cdots & \mathbf{c}'_1 \mathbf{c}_K \\ \mathbf{c}'_2 \mathbf{c}_1 & \mathbf{c}'_2 \mathbf{c}_2 & \cdots & \mathbf{c}'_2 \mathbf{c}_K \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{c}'_K \mathbf{c}_1 & \mathbf{c}'_K \mathbf{c}_2 & \cdots & \mathbf{c}'_K \mathbf{c}_K \end{bmatrix} = \mathbf{I}. \quad (\text{A-81})$$

Result (A-81) implies that

$$\mathbf{C}' = \mathbf{C}^{-1}. \quad (\text{A-82})$$

Consequently,

$$\mathbf{C}\mathbf{C}' = \mathbf{C}\mathbf{C}^{-1} = \mathbf{I} \quad (\text{A-83})$$

as well, so the rows as well as the columns of \mathbf{C} are orthogonal.

A.6.4 DIAGONALIZATION AND SPECTRAL DECOMPOSITION OF A MATRIX

By premultiplying (A-80) by \mathbf{C}' and using (A-81), we can extract the characteristic roots of \mathbf{A} .

DEFINITION A.15 Diagonalization of a Matrix

The *diagonalization* of a matrix \mathbf{A} is

$$\mathbf{C}'\mathbf{A}\mathbf{C} = \mathbf{C}'\mathbf{C}\mathbf{A} = \mathbf{I}\mathbf{A} = \mathbf{A}. \quad (\text{A-84})$$

Alternatively, by *post*multiplying (A-80) by \mathbf{C}' and using (A-83), we obtain a useful representation of \mathbf{A} .

DEFINITION A.16 Spectral Decomposition of a Matrix

The *spectral decomposition* of \mathbf{A} is

$$\mathbf{A} = \mathbf{C}\mathbf{A}\mathbf{C}' = \sum_{k=1}^K \lambda_k \mathbf{c}_k \mathbf{c}'_k. \quad (\text{A-85})$$

In this representation, the $K \times K$ matrix \mathbf{A} is written as a sum of K rank one matrices. This sum is also called the **eigenvalue** (or, “own” value) decomposition of \mathbf{A} . In this connection, the term *signature* of the matrix is sometimes used to describe the characteristic roots and vectors. Yet another pair of terms for the parts of this decomposition are the **latent roots** and **latent vectors** of \mathbf{A} .

A.6.5 RANK OF A MATRIX

The diagonalization result enables us to obtain the rank of a matrix very easily. To do so, we can use the following result.

THEOREM A.3 Rank of a Product

For any matrix \mathbf{A} and nonsingular matrices \mathbf{B} and \mathbf{C} , the rank of \mathbf{BAC} is equal to the rank of \mathbf{A} . The proof is simple. By (A-45), $\text{rank}(\mathbf{BAC}) = \text{rank}[(\mathbf{BA})\mathbf{C}] = \text{rank}(\mathbf{BA})$. By (A-43), $\text{rank}(\mathbf{BA}) = \text{rank}(\mathbf{A}'\mathbf{B}')$, and applying (A-45) again, $\text{rank}(\mathbf{A}'\mathbf{B}') = \text{rank}(\mathbf{A}')$ since \mathbf{B}' is nonsingular if \mathbf{B} is nonsingular (once again, by A-43). Finally, applying (A-43) again to obtain $\text{rank}(\mathbf{A}') = \text{rank}(\mathbf{A})$ gives the result.

Since \mathbf{C} and \mathbf{C}' are nonsingular, we can use them to apply this result to (A-84). By an obvious substitution,

$$\text{rank}(\mathbf{A}) = \text{rank}(\mathbf{A}). \quad (\text{A-86})$$

Finding the rank of \mathbf{A} is trivial. Since \mathbf{A} is a diagonal matrix, its rank is just the number of nonzero values on its diagonal. By extending this result, we can prove the following theorems. (Proofs are brief and are left for the reader.)

THEOREM A.4 Rank of a Symmetric Matrix

The rank of a symmetric matrix is the number of nonzero characteristic roots it contains.

Note how this result enters the spectral decomposition given above. If any of the characteristic roots are zero, then the number of rank one matrices in the sum is reduced correspondingly. It would appear that this simple rule will not be useful if \mathbf{A} is not square. But recall that

$$\text{rank}(\mathbf{A}) = \text{rank}(\mathbf{A}'\mathbf{A}). \quad (\text{A-87})$$

Since $\mathbf{A}'\mathbf{A}$ is always square, we can use it instead of \mathbf{A} . Indeed, we can use it even if \mathbf{A} is square, which leads to a fully general result.

THEOREM A.5 Rank of a Matrix

The rank of any matrix \mathbf{A} equals the number of nonzero characteristic roots in $\mathbf{A}'\mathbf{A}$.

Since the row rank and column rank of a matrix are equal, we should be able to apply Theorem A.5 to \mathbf{AA}' as well. This process, however, requires an additional result.

THEOREM A.6 Roots of an Outer Product Matrix

The nonzero characteristic roots of \mathbf{AA}' are the same as those of $\mathbf{A}'\mathbf{A}$.

The proof is left as an exercise. A useful special case the reader can examine is the characteristic roots of \mathbf{aa}' and $\mathbf{a}'\mathbf{a}$, where \mathbf{a} is an $n \times 1$ vector.

If a characteristic root of a matrix is zero, then we have $\mathbf{Ac} = \mathbf{0}$. Thus, if the matrix has a zero root, it must be singular. Otherwise, no nonzero \mathbf{c} would exist. In general, therefore, a matrix is singular; that is, it does not have full rank if and only if it has at least one zero root.

A.6.6 CONDITION NUMBER OF A MATRIX

As the preceding might suggest, there is a discrete difference between full rank and short rank matrices. In analyzing data matrices such as the one in Section A.2, however, we shall often encounter cases in which a matrix is not quite short ranked, because it has all nonzero roots, but it is close. That is, by some measure, we can come very close to being able to write one column as a linear combination of the others. This case is important; we shall examine it at length in our discussion of multicollinearity. Our definitions of rank and determinant will fail to indicate this possibility, but an alternative measure, the **condition number**, is designed for that purpose. Formally, the condition number for a square matrix \mathbf{A} is

$$\gamma = \left[\frac{\text{maximum root}}{\text{minimum root}} \right]^{1/2} \quad (\text{A-88})$$

For nonsquare matrices \mathbf{X} , such as the data matrix in the example, we use $\mathbf{A} = \mathbf{X}'\mathbf{X}$. As a further refinement, because the characteristic roots are affected by the scaling of the columns of \mathbf{X} , we scale the columns to have length 1 by dividing each column by its norm [see (A-55)]. For the \mathbf{X} in Section A.2, the largest characteristic root of \mathbf{A} is 4.9255 and the smallest is 0.0001543. Therefore, the condition number is 178.67, which is extremely large. (Values greater than 20 are large.) That the smallest root is close to zero compared with the largest means that this matrix is nearly singular. Matrices with large condition numbers are difficult to invert accurately.

A.6.7 TRACE OF A MATRIX

The **trace** of a square $K \times K$ matrix is the sum of its diagonal elements:

$$\text{tr}(\mathbf{A}) = \sum_{k=1}^K a_{kk}.$$

Some easily proven results are

$$\text{tr}(c\mathbf{A}) = c(\text{tr}(\mathbf{A})), \quad (\text{A-89})$$

$$\text{tr}(\mathbf{A}') = \text{tr}(\mathbf{A}), \quad (\text{A-90})$$

$$\text{tr}(\mathbf{A} + \mathbf{B}) = \text{tr}(\mathbf{A}) + \text{tr}(\mathbf{B}), \quad (\text{A-91})$$

$$\text{tr}(\mathbf{I}_K) = K. \quad (\text{A-92})$$

$$\text{tr}(\mathbf{AB}) = \text{tr}(\mathbf{BA}). \quad (\text{A-93})$$

$$\mathbf{a}'\mathbf{a} = \text{tr}(\mathbf{a}'\mathbf{a}) = \text{tr}(\mathbf{aa}')$$

$$\text{tr}(\mathbf{A}'\mathbf{A}) = \sum_{k=1}^K \mathbf{a}'_k \mathbf{a}_k = \sum_{i=1}^K \sum_{k=1}^K a_{ik}^2.$$

The permutation rule can be extended to any *cyclic* permutation in a product:

$$\text{tr}(\mathbf{ABCD}) = \text{tr}(\mathbf{BCDA}) = \text{tr}(\mathbf{CDAB}) = \text{tr}(\mathbf{DABC}). \quad (\text{A-94})$$

By using (A-84), we obtain

$$\text{tr}(\mathbf{C}'\mathbf{A}\mathbf{C}) = \text{tr}(\mathbf{A}\mathbf{C}\mathbf{C}') = \text{tr}(\mathbf{A}\mathbf{I}) = \text{tr}(\mathbf{A}) = \text{tr}(\mathbf{A}). \quad (\text{A-95})$$

Since \mathbf{A} is diagonal with the roots of \mathbf{A} on its diagonal, the general result is the following.

THEOREM A.7 Trace of a Matrix

The trace of a matrix equals the sum of its characteristic roots.

(A-96)

A.6.8 DETERMINANT OF A MATRIX

Recalling how tedious the calculation of a determinant promised to be, we find that the following is particularly useful. Since

$$\begin{aligned} \mathbf{C}'\mathbf{A}\mathbf{C} &= \mathbf{A}, \\ |\mathbf{C}'\mathbf{A}\mathbf{C}| &= |\mathbf{A}|. \end{aligned} \quad (\text{A-97})$$

Using a number of earlier results, we have, for orthogonal matrix \mathbf{C} ,

$$\begin{aligned} |\mathbf{C}'\mathbf{A}\mathbf{C}| &= |\mathbf{C}'| \cdot |\mathbf{A}| \cdot |\mathbf{C}| = |\mathbf{C}'| \cdot |\mathbf{C}| \cdot |\mathbf{A}| = |\mathbf{C}'\mathbf{C}| \cdot |\mathbf{A}| = |\mathbf{I}| \cdot |\mathbf{A}| = 1 \cdot |\mathbf{A}| \\ &= |\mathbf{A}| \\ &= |\mathbf{A}|. \end{aligned} \quad (\text{A-98})$$

Since $|\mathbf{A}|$ is just the product of its diagonal elements, the following is implied.

THEOREM A.8 Determinant of a Matrix

The determinant of a matrix equals the product of its characteristic roots.

(A-99)

Notice that we get the expected result if any of these roots is zero. Since the determinant is the product of the roots, it follows that a matrix is singular if and only if its determinant is zero and, in turn, if and only if it has at least one zero characteristic root.

A.6.9 POWERS OF A MATRIX

We often use expressions involving powers of matrices, such as $\mathbf{A}\mathbf{A} = \mathbf{A}^2$. For positive integer powers, these expressions can be computed by repeated multiplication. But this does not show how to handle a problem such as finding a \mathbf{B} such that $\mathbf{B}^2 = \mathbf{A}$, that is, the square root of a matrix. The characteristic roots and vectors provide a solution. Consider first

$$\begin{aligned} \mathbf{A}\mathbf{A} &= \mathbf{A}^2 = (\mathbf{C}\mathbf{A}\mathbf{C}')(\mathbf{C}\mathbf{A}\mathbf{C}') = \mathbf{C}\mathbf{A}\mathbf{C}'\mathbf{C}\mathbf{A}\mathbf{C}' = \mathbf{C}\mathbf{A}\mathbf{I}\mathbf{A}\mathbf{C}' = \mathbf{C}\mathbf{A}^2\mathbf{C}' \\ &= \mathbf{C}\mathbf{A}^2\mathbf{C}'. \end{aligned} \quad (\text{A-100})$$

Two results follow. Since \mathbf{A}^2 is a diagonal matrix whose nonzero elements are the squares of those in \mathbf{A} , the following is implied.

For any symmetric matrix, the characteristic roots of \mathbf{A}^2 are the squares of those of \mathbf{A} , and the characteristic vectors are the same.

(A-101)

The proof is obtained by observing that the second line in (A-100) is the spectral decomposition of the matrix $\mathbf{B} = \mathbf{A}\mathbf{A}$. Since $\mathbf{A}^3 = \mathbf{A}\mathbf{A}^2$ and so on, (A-101) extends to any positive integer. By convention, for any \mathbf{A} , $\mathbf{A}^0 = \mathbf{I}$. Thus, for any symmetric matrix \mathbf{A} , $\mathbf{A}^K = \mathbf{C}\mathbf{\Lambda}^K\mathbf{C}'$, $K = 0, 1, \dots$. Hence, the characteristic roots of \mathbf{A}^K are λ^K , whereas the characteristic vectors are the same as those of \mathbf{A} . If \mathbf{A} is nonsingular, so that all its roots λ_i are nonzero, then this proof can be extended to negative powers as well.

If \mathbf{A}^{-1} exists, then

$$\mathbf{A}^{-1} = (\mathbf{C}\mathbf{\Lambda}\mathbf{C}')^{-1} = (\mathbf{C}')^{-1}\mathbf{\Lambda}^{-1}\mathbf{C}^{-1} = \mathbf{C}\mathbf{\Lambda}^{-1}\mathbf{C}', \tag{A-102}$$

where we have used the earlier result, $\mathbf{C}' = \mathbf{C}^{-1}$. This gives an important result that is useful for analyzing inverse matrices.

THEOREM A.9 Characteristic Roots of an Inverse Matrix

If \mathbf{A}^{-1} exists, then the characteristic roots of \mathbf{A}^{-1} are the reciprocals of those of \mathbf{A} , and the characteristic vectors are the same.

By extending the notion of repeated multiplication, we now have a more general result.

THEOREM A.10 Characteristic Roots of a Matrix Power

For any nonsingular symmetric matrix $\mathbf{A} = \mathbf{C}\mathbf{\Lambda}\mathbf{C}'$, $\mathbf{A}^K = \mathbf{C}\mathbf{\Lambda}^K\mathbf{C}'$, $K = \dots, -2, -1, 0, 1, 2, \dots$

We now turn to the general problem of how to compute the square root of a matrix. In the scalar case, the value would have to be nonnegative. The matrix analog to this requirement is that all the characteristic roots are nonnegative. Consider, then, the candidate

$$\mathbf{A}^{1/2} = \mathbf{C}\mathbf{\Lambda}^{1/2}\mathbf{C}' = \mathbf{C} \begin{bmatrix} \sqrt{\lambda_1} & 0 & \dots & 0 \\ 0 & \sqrt{\lambda_2} & \dots & 0 \\ & & \dots & \\ 0 & 0 & \dots & \sqrt{\lambda_n} \end{bmatrix} \mathbf{C}'. \tag{A-103}$$

This equation satisfies the requirement for a square root, since

$$\mathbf{A}^{1/2}\mathbf{A}^{1/2} = \mathbf{C}\mathbf{\Lambda}^{1/2}\mathbf{C}'\mathbf{C}\mathbf{\Lambda}^{1/2}\mathbf{C}' = \mathbf{C}\mathbf{\Lambda}\mathbf{C}' = \mathbf{A}. \tag{A-104}$$

If we continue in this fashion, we can define the powers of a matrix more generally, still assuming that all the characteristic roots are nonnegative. For example, $\mathbf{A}^{1/3} = \mathbf{C}\mathbf{\Lambda}^{1/3}\mathbf{C}'$. If all the roots are strictly positive, we can go one step further and extend the result to any real power. For reasons that will be made clear in the next section, we say that a matrix with positive characteristic roots is **positive definite**. It is the matrix analog to a positive number.

DEFINITION A.17 Real Powers of a Positive Definite Matrix

For a positive definite matrix \mathbf{A} , $\mathbf{A}^r = \mathbf{C}\mathbf{\Lambda}^r\mathbf{C}'$, for any real number, r . (A-105)

The characteristic roots of \mathbf{A}^r are the r th power of those of \mathbf{A} , and the characteristic vectors are the same.

If \mathbf{A} is only **nonnegative definite**—that is, has roots that are either zero or positive—then (A-105) holds only for nonnegative r .

A.6.10 IDEMPOTENT MATRICES

Idempotent matrices are equal to their squares [see (A-37) to (A-39)]. In view of their importance in econometrics, we collect a few results related to idempotent matrices at this point. First, (A-101) implies that if λ is a characteristic root of an idempotent matrix, then $\lambda = \lambda^K$ for all nonnegative integers K . As such, if \mathbf{A} is a symmetric idempotent matrix, then all its roots are one or zero. Assume that all the roots of \mathbf{A} are one. Then $\mathbf{A} = \mathbf{I}$, and $\mathbf{A} = \mathbf{C}\mathbf{A}\mathbf{C}' = \mathbf{C}\mathbf{I}\mathbf{C}' = \mathbf{C}\mathbf{C}' = \mathbf{I}$. If the roots are not all one, then one or more are zero. Consequently, we have the following results for symmetric idempotent matrices:⁹

- *The only full rank, symmetric idempotent matrix is the identity matrix \mathbf{I} .* (A-106)
- *All symmetric idempotent matrices except the identity matrix are singular.* (A-107)

The final result on idempotent matrices is obtained by observing that the count of the nonzero roots of \mathbf{A} is also equal to their sum. By combining Theorems A.5 and A.7 with the result that for an idempotent matrix, the roots are all zero or one, we obtain this result:

- *The rank of a symmetric idempotent matrix is equal to its trace.* (A-108)

A.6.11 FACTORING A MATRIX

In some applications, we shall require a matrix \mathbf{P} such that

$$\mathbf{P}'\mathbf{P} = \mathbf{A}^{-1}.$$

One choice is

$$\mathbf{P} = \mathbf{A}^{-1/2}\mathbf{C}'$$

so that,

$$\mathbf{P}'\mathbf{P} = (\mathbf{C}')'(\mathbf{A}^{-1/2})'\mathbf{A}^{-1/2}\mathbf{C}' = \mathbf{C}\mathbf{A}^{-1}\mathbf{C}',$$

as desired.¹⁰ Thus, the **spectral decomposition** of \mathbf{A} , $\mathbf{A} = \mathbf{C}\mathbf{A}\mathbf{C}'$ is a useful result for this kind of computation.

The **Cholesky factorization** of a symmetric positive definite matrix is an alternative representation that is useful in regression analysis. Any symmetric positive definite matrix \mathbf{A} may be written as the product of a **lower triangular matrix** \mathbf{L} and its transpose (which is an **upper triangular matrix**) $\mathbf{L}' = \mathbf{U}$. Thus, $\mathbf{A} = \mathbf{L}\mathbf{U}$. This result is the Cholesky decomposition of \mathbf{A} . The square roots of the diagonal elements of \mathbf{L} , d_i , are the **Cholesky values** of \mathbf{A} . By arraying these in a diagonal matrix \mathbf{D} , we may also write $\mathbf{A} = \mathbf{L}\mathbf{D}^{-1}\mathbf{D}^2\mathbf{D}^{-1}\mathbf{U} = \mathbf{L}^*\mathbf{D}^2\mathbf{U}^*$, which is similar to the spectral decomposition in (A-85). The usefulness of this formulation arises when the inverse of \mathbf{A} is required. Once \mathbf{L} is

⁹Not all idempotent matrices are symmetric. We shall not encounter any asymmetric ones in our work, however.

¹⁰We say that this is “one” choice because if \mathbf{A} is symmetric, as it will be in all our applications, there are other candidates. The reader can easily verify that $\mathbf{C}\mathbf{A}^{-1/2}\mathbf{C}' = \mathbf{A}^{-1/2}$ works as well.

computed, finding $\mathbf{A}^{-1} = \mathbf{U}^{-1}\mathbf{L}^{-1}$ is also straightforward as well as extremely fast and accurate. Most recently developed econometric software packages use this technique for inverting positive definite matrices.

A third type of decomposition of a matrix is useful for numerical analysis when the inverse is difficult to obtain because the columns of \mathbf{A} are “nearly” collinear. Any $n \times K$ matrix \mathbf{A} for which $n \geq K$ can be written in the form $\mathbf{A} = \mathbf{U}\mathbf{W}\mathbf{V}'$, where \mathbf{U} is an orthogonal $n \times K$ matrix—that is, $\mathbf{U}'\mathbf{U} = \mathbf{I}_K$ — \mathbf{W} is a $K \times K$ diagonal matrix such that $w_i \geq 0$, and \mathbf{V} is a $K \times K$ matrix such that $\mathbf{V}'\mathbf{V} = \mathbf{I}_K$. This result is called the **singular value decomposition (SVD)** of \mathbf{A} , and w_i are the singular values of \mathbf{A} .¹¹ (Note that if \mathbf{A} is square, then the spectral decomposition is a singular value decomposition.) As with the Cholesky decomposition, the usefulness of the SVD arises in inversion, in this case, of $\mathbf{A}'\mathbf{A}$. By multiplying it out, we obtain that $(\mathbf{A}'\mathbf{A})^{-1}$ is simply $\mathbf{V}\mathbf{W}^{-2}\mathbf{V}'$. Once the SVD of \mathbf{A} is computed, the inversion is trivial. The other advantage of this format is its numerical stability, which is discussed at length in Press et al. (1986).

Press et al. (1986) recommend the SVD approach as the method of choice for solving least squares problems because of its accuracy and numerical stability. A commonly used alternative method similar to the SVD approach is the QR decomposition. Any $n \times K$ matrix, \mathbf{X} , with $n \geq K$ can be written in the form $\mathbf{X} = \mathbf{Q}\mathbf{R}$ in which the columns of \mathbf{Q} are orthonormal ($\mathbf{Q}'\mathbf{Q} = \mathbf{I}$) and \mathbf{R} is an upper triangular matrix. Decomposing \mathbf{X} in this fashion allows an extremely accurate solution to the least squares problem that does not involve inversion or direct solution of the normal equations. Press et al. suggest that this method may have problems with rounding errors in problems when \mathbf{X} is nearly of short rank, but based on other published results, this concern seems relatively minor.¹²

A.6.12 THE GENERALIZED INVERSE OF A MATRIX

Inverse matrices are fundamental in econometrics. Although we shall not require them much in our treatment in this book, there are more general forms of inverse matrices than we have considered thus far. A **generalized inverse** of a matrix \mathbf{A} is another matrix \mathbf{A}^+ that satisfies the following requirements:

1. $\mathbf{A}\mathbf{A}^+\mathbf{A} = \mathbf{A}$.
2. $\mathbf{A}^+\mathbf{A}\mathbf{A}^+ = \mathbf{A}^+$.
3. $\mathbf{A}^+\mathbf{A}$ is symmetric.
4. $\mathbf{A}\mathbf{A}^+$ is symmetric.

A unique \mathbf{A}^+ can be found for any matrix, whether \mathbf{A} is singular or not, or even if \mathbf{A} is not square.¹³ The unique matrix that satisfies all four requirements is called the **Moore—Penrose inverse** or **pseudoinverse** of \mathbf{A} . If \mathbf{A} happens to be square and nonsingular, then the generalized inverse will be the familiar ordinary inverse. But if \mathbf{A}^{-1} does not exist, then \mathbf{A}^+ can still be computed.

An important special case is the overdetermined system of equations

$$\mathbf{A}\mathbf{b} = \mathbf{y},$$

¹¹Discussion of the singular value decomposition (and listings of computer programs for the computations) may be found in Press et al. (1986).

¹²The National Institute of Standards and Technology (NIST) has published a suite of benchmark problems that test the accuracy of least squares computations (<http://www.nist.gov/itl/div898/strd>). Using these problems, which include some extremely difficult, ill-conditioned data sets, we found that the QR method would reproduce all the NIST certified solutions to 15 digits of accuracy, which suggests that the QR method should be satisfactory for all but the worst problems.

¹³A proof of uniqueness, with several other results, may be found in Theil (1983).

834 APPENDIX A ♦ Matrix Algebra

where \mathbf{A} has n rows, $K < n$ columns, and column rank equal to $R \leq K$. Suppose that R equals K , so that $(\mathbf{A}'\mathbf{A})^{-1}$ exists. Then the Moore—Penrose inverse of \mathbf{A} is

$$\mathbf{A}^+ = (\mathbf{A}'\mathbf{A})^{-1}\mathbf{A}',$$

which can be verified by multiplication. A “solution” to the system of equations can be written

$$\mathbf{b} = \mathbf{A}^+\mathbf{y}.$$

This is the vector that minimizes the length of $\mathbf{A}\mathbf{b} - \mathbf{y}$. Recall this was the solution to the least squares problem obtained in Section A.4.4. If \mathbf{y} lies in the column space of \mathbf{A} , this vector will be zero, but otherwise, it will not.

Now suppose that \mathbf{A} does not have full rank. The previous solution cannot be computed. An alternative solution can be obtained, however. We continue to use the matrix $\mathbf{A}'\mathbf{A}$. In the spectral decomposition of Section A.6.4, if \mathbf{A} has rank R , then there are R terms in the summation in (A-85). In (A-102), the spectral decomposition using the reciprocals of the characteristic roots is used to compute the inverse. To compute the Moore—Penrose inverse, we apply this calculation to $\mathbf{A}'\mathbf{A}$, using only the nonzero roots, then postmultiply the result by \mathbf{A}' . Let \mathbf{C}_1 be the R characteristic vectors corresponding to the nonzero roots, which we array in the diagonal matrix, Λ_1 . Then the Moore—Penrose inverse is

$$\mathbf{A}^+ = \mathbf{C}_1\Lambda_1^{-1}\mathbf{C}_1'\mathbf{A}',$$

which is very similar to the previous result.

If \mathbf{A} is a symmetric matrix with rank $R \leq K$, the Moore—Penrose inverse is computed precisely as in the preceding equation without postmultiplying by \mathbf{A}' . Thus, for a symmetric matrix \mathbf{A} ,

$$\mathbf{A}^+ = \mathbf{C}_1\Lambda_1^{-1}\mathbf{C}_1',$$

where Λ_1^{-1} is a diagonal matrix containing the reciprocals of the *nonzero* roots of \mathbf{A} .

A.7 QUADRATIC FORMS AND DEFINITE MATRICES

Many optimization problems involve double sums of the form

$$q = \sum_{i=1}^n \sum_{j=1}^n x_i x_j a_{ij}. \quad (\text{A-109})$$

This **quadratic form** can be written

$$q = \mathbf{x}'\mathbf{A}\mathbf{x},$$

where \mathbf{A} is a symmetric matrix. In general, q may be positive, negative, or zero; it depends on \mathbf{A} and \mathbf{x} . There are some matrices, however, for which q will be positive regardless of \mathbf{x} , and others for which q will always be negative (or nonnegative or nonpositive). For a given matrix \mathbf{A} ,

1. If $\mathbf{x}'\mathbf{A}\mathbf{x} > (<) 0$ for all nonzero \mathbf{x} , then \mathbf{A} is **positive (negative) definite**.
2. If $\mathbf{x}'\mathbf{A}\mathbf{x} \geq (\leq) 0$ for all nonzero \mathbf{x} , then \mathbf{A} is **nonnegative definite** or **positive semidefinite** (nonpositive definite).

It might seem that it would be impossible to check a matrix for definiteness, since \mathbf{x} can be chosen arbitrarily. But we have already used the set of results necessary to do so. Recall that a

symmetric matrix can be decomposed into

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

Therefore, the quadratic form can be written as

$$\mathbf{x}'\mathbf{A}\mathbf{x} = \mathbf{x}'\mathbf{C}\mathbf{A}\mathbf{C}'\mathbf{x}.$$

Let $\mathbf{y} = \mathbf{C}'\mathbf{x}$. Then

$$\mathbf{x}'\mathbf{A}\mathbf{x} = \mathbf{y}'\mathbf{\Lambda}\mathbf{y} = \sum_{i=1}^n \lambda_i y_i^2. \tag{A-110}$$

If λ_i is positive for all i , then regardless of \mathbf{y} —that is, regardless of \mathbf{x} — q will be positive. This case was identified earlier as a positive definite matrix. Continuing this line of reasoning, we obtain the following theorem.

THEOREM A.11 Definite Matrices

*Let \mathbf{A} be a symmetric matrix. If all the characteristic roots of \mathbf{A} are positive (negative), then \mathbf{A} is **positive definite (negative definite)**. If some of the roots are zero, then \mathbf{A} is **nonnegative (nonpositive) definite** if the remainder are positive (negative). If \mathbf{A} has both negative and positive roots, then \mathbf{A} is **indefinite**.*

The preceding statements give, in each case, the “if” parts of the theorem. To establish the “only if” parts, assume that the condition on the roots does not hold. This must lead to a contradiction. For example, if some λ can be negative, then $\mathbf{y}'\mathbf{\Lambda}\mathbf{y}$ could be negative for some \mathbf{y} , so \mathbf{A} cannot be positive definite.

A.7.1 NONNEGATIVE DEFINITE MATRICES

A case of particular interest is that of nonnegative definite matrices. Theorem A.11 implies a number of related results.

- If \mathbf{A} is nonnegative definite, then $|\mathbf{A}| \geq 0$. (A-111)

Proof: The determinant is the product of the roots, which are nonnegative.

The converse, however, is not true. For example, a 2×2 matrix with two negative roots is clearly not positive definite, but it does have a positive determinant.

- If \mathbf{A} is positive definite, so is \mathbf{A}^{-1} . (A-112)

Proof: The roots are the reciprocals of those of \mathbf{A} , which are, therefore positive.

- The identity matrix \mathbf{I} is positive definite. (A-113)

Proof: $\mathbf{x}'\mathbf{I}\mathbf{x} = \mathbf{x}'\mathbf{x} > 0$ if $\mathbf{x} \neq \mathbf{0}$.

A very important result for regression analysis is

- If \mathbf{A} is $n \times K$ with full column rank and $n > K$, then $\mathbf{A}'\mathbf{A}$ is positive definite and $\mathbf{A}\mathbf{A}'$ is nonnegative definite. (A-114)

Proof: By assumption, $\mathbf{A}\mathbf{x} \neq \mathbf{0}$. So $\mathbf{x}'\mathbf{A}'\mathbf{A}\mathbf{x} = (\mathbf{A}\mathbf{x})'(\mathbf{A}\mathbf{x}) = \mathbf{y}'\mathbf{y} = \sum_j y_j^2 > 0$.

A similar proof establishes the nonnegative definiteness of $\mathbf{A}\mathbf{A}'$. The difference in the latter case is that because \mathbf{A} has more rows than columns there is an \mathbf{x} such that $\mathbf{A}'\mathbf{x} = \mathbf{0}$. Thus, in the proof, we only have $\mathbf{y}'\mathbf{y} \geq 0$. The case in which \mathbf{A} does not have full column rank is the same as that of $\mathbf{A}\mathbf{A}'$.

- If \mathbf{A} is positive definite and \mathbf{B} is a nonsingular matrix, then $\mathbf{B}'\mathbf{A}\mathbf{B}$ is positive definite. (A-115)

Proof: $\mathbf{x}'\mathbf{B}'\mathbf{A}\mathbf{B}\mathbf{x} = \mathbf{y}'\mathbf{A}\mathbf{y} > 0$, where $\mathbf{y} = \mathbf{B}\mathbf{x}$. But \mathbf{y} cannot be $\mathbf{0}$ because \mathbf{B} is nonsingular.

Finally, note that for \mathbf{A} to be negative definite, all \mathbf{A} 's characteristic roots must be negative. But, in this case, $|\mathbf{A}|$ is positive if \mathbf{A} is of even order and negative if \mathbf{A} is of odd order.

A.7.2 IDEMPOTENT QUADRATIC FORMS

Quadratic forms in idempotent matrices play an important role in the distributions of many test statistics. As such, we shall encounter them fairly often. Two central results are of interest.

- Every symmetric idempotent matrix is nonnegative definite. (A-116)

Proof: All roots are one or zero; hence, the matrix is nonnegative definite by definition.

Combining this with some earlier results yields a result used in determining the sampling distribution of most of the standard test statistics.

- If \mathbf{A} is symmetric and idempotent, $n \times n$ with rank J , then every quadratic form in \mathbf{A} can be written $\mathbf{x}'\mathbf{A}\mathbf{x} = \sum_{j=1}^J y_j^2$ (A-117)

Proof: This result is (A-110) with $\lambda =$ one or zero.

A.7.3 COMPARING MATRICES

Derivations in econometrics often focus on whether one matrix is "larger" than another. We now consider how to make such a comparison. As a starting point, the two matrices must have the same dimensions. A useful comparison is based on

$$d = \mathbf{x}'\mathbf{A}\mathbf{x} - \mathbf{x}'\mathbf{B}\mathbf{x} = \mathbf{x}'(\mathbf{A} - \mathbf{B})\mathbf{x}.$$

If d is always positive for any nonzero vector, \mathbf{x} , then by this criterion, we can say that \mathbf{A} is larger than \mathbf{B} . The reverse would apply if d is always negative. It follows from the definition that

$$\text{if } d > 0 \text{ for all nonzero } \mathbf{x}, \text{ then } \mathbf{A} - \mathbf{B} \text{ is positive definite.} \quad (\text{A-118})$$

If d is only greater than or equal to zero, then $\mathbf{A} - \mathbf{B}$ is nonnegative definite. The ordering is not complete. For some pairs of matrices, d could have either sign, depending on \mathbf{x} . In this case, there is no simple comparison.

A particular case of the general result which we will encounter frequently is:

$$\begin{aligned} &\text{If } \mathbf{A} \text{ is positive definite and } \mathbf{B} \text{ is nonnegative definite,} \\ &\text{then } \mathbf{A} + \mathbf{B} \geq \mathbf{A}. \end{aligned} \quad (\text{A-119})$$

Consider, for example, the "updating formula" introduced in (A-66). This uses a matrix

$$\mathbf{A} = \mathbf{B}'\mathbf{B} + \mathbf{b}\mathbf{b}' \geq \mathbf{B}'\mathbf{B}.$$

Finally, in comparing matrices, it may be more convenient to compare their inverses. The result analogous to a familiar result for scalars is:

$$\text{If } \mathbf{A} > \mathbf{B}, \text{ then } \mathbf{B}^{-1} > \mathbf{A}^{-1}. \quad (\text{A-120})$$

In order to establish this intuitive result, we would make use of the following, which is proved in Goldberger (1964, Chapter 2):

THEOREM A.12 Ordering for Positive Definite Matrices

If \mathbf{A} and \mathbf{B} are two positive definite matrices with the same dimensions and if every characteristic root of \mathbf{A} is larger than (at least as large as) the corresponding characteristic root of \mathbf{B} when both sets of roots are ordered from largest to smallest, then $\mathbf{A} - \mathbf{B}$ is positive (nonnegative) definite.

The roots of the inverse are the reciprocals of the roots of the original matrix, so the theorem can be applied to the inverse matrices.

A.8 CALCULUS AND MATRIX ALGEBRA¹⁴

A.8.1 DIFFERENTIATION AND THE TAYLOR SERIES

A variable y is a function of another variable x written

$$y = f(x), \quad y = g(x), \quad y = y(x),$$

and so on, if each value of x is associated with a single value of y . In this relationship, y and x are sometimes labeled the **dependent variable** and the **independent variable**, respectively. Assuming that the function $f(x)$ is continuous and differentiable, we obtain the following derivatives:

$$f'(x) = \frac{dy}{dx}, \quad f''(x) = \frac{d^2y}{dx^2},$$

and so on.

A frequent use of the derivatives of $f(x)$ is in the **Taylor series approximation**. A Taylor series is a polynomial approximation to $f(x)$. Letting x^0 be an arbitrarily chosen expansion point

$$f(x) \approx f(x^0) + \sum_{i=1}^p \frac{1}{i!} \frac{d^i f(x^0)}{d(x^0)^i} (x - x^0)^i. \quad (\text{A-121})$$

The choice of the number of terms is arbitrary; the more that are used, the more accurate the approximation will be. The approximation used most frequently in econometrics is the **linear approximation**,

$$f(x) \approx \alpha + \beta x \quad (\text{A-122})$$

where, by collecting terms in (A-121), $\alpha = [f(x^0) - f'(x^0)x^0]$ and $\beta = f'(x^0)$. The superscript "0" indicates that the function is evaluated at x^0 . The **quadratic approximation** is

$$f(x) \approx \alpha + \beta x + \gamma x^2, \quad (\text{A-123})$$

where $\alpha = [f^0 - f'^0 x^0 + \frac{1}{2} f''^0 (x^0)^2]$, $\beta = [f'^0 - f''^0 x^0]$ and $\gamma = \frac{1}{2} f''^0$.

¹⁴For a complete exposition, see Magnus and Neudecker (1988).

We can regard a function $y = f(x_1, x_2, \dots, x_n)$ as a **scalar-valued function** of a vector; that is, $y = f(\mathbf{x})$. The vector of partial derivatives, or **gradient vector**, or simply **gradient**, is

$$\frac{\partial f(\mathbf{x})}{\partial \mathbf{x}} = \begin{bmatrix} \partial y / \partial x_1 \\ \partial y / \partial x_2 \\ \dots \\ \partial y / \partial x_n \end{bmatrix} = \begin{bmatrix} f_1 \\ f_2 \\ \dots \\ f_n \end{bmatrix}. \tag{A-124}$$

The vector $\mathbf{g}(\mathbf{x})$ or \mathbf{g} is used to represent the gradient. Notice that it is a column vector. The shape of the derivative is determined by the denominator of the derivative.

A **second derivatives matrix** or **Hessian** is computed as

$$\mathbf{H} = \begin{bmatrix} \partial^2 y / \partial x_1 \partial x_1 & \partial^2 y / \partial x_1 \partial x_2 & \dots & \partial^2 y / \partial x_1 \partial x_n \\ \partial^2 y / \partial x_2 \partial x_1 & \partial^2 y / \partial x_2 \partial x_2 & \dots & \partial^2 y / \partial x_2 \partial x_n \\ \dots & \dots & \dots & \dots \\ \partial^2 y / \partial x_n \partial x_1 & \partial^2 y / \partial x_n \partial x_2 & \dots & \partial^2 y / \partial x_n \partial x_n \end{bmatrix} = [f_{ij}]. \tag{A-125}$$

In general, \mathbf{H} is a square, symmetric matrix. (The symmetry is obtained for continuous and continuously differentiable functions from Young's theorem.) Each column of \mathbf{H} is the derivative of \mathbf{g} with respect to the corresponding variable in \mathbf{x}' . Therefore,

$$\mathbf{H} = \begin{bmatrix} \partial(\partial y / \partial \mathbf{x}) / \partial x_1 & \partial(\partial y / \partial \mathbf{x}) / \partial x_2 & \dots & \partial(\partial y / \partial \mathbf{x}) / \partial x_n \end{bmatrix} = \frac{\partial(\partial y / \partial \mathbf{x})}{\partial(x_1 \ x_2 \ \dots \ x_n)} = \frac{\partial(\partial y / \partial \mathbf{x})}{\partial \mathbf{x}'} = \frac{\partial^2 y}{\partial \mathbf{x} \partial \mathbf{x}'}$$

The first-order, or linear Taylor series approximation is

$$y \approx f(\mathbf{x}^0) + \sum_{i=1}^n f_i(\mathbf{x}^0)(x_i - x_i^0). \tag{A-126}$$

The right-hand side is

$$f(\mathbf{x}^0) + \left[\frac{\partial f(\mathbf{x}^0)}{\partial \mathbf{x}^0} \right]' (\mathbf{x} - \mathbf{x}^0) = [f(\mathbf{x}^0) - \mathbf{g}(\mathbf{x}^0)' \mathbf{x}^0] + \mathbf{g}(\mathbf{x}^0)' \mathbf{x} = [f^0 - \mathbf{g}^0' \mathbf{x}^0] + \mathbf{g}^0' \mathbf{x}.$$

This produces the linear approximation,

$$y \approx \alpha + \beta' \mathbf{x}.$$

The second-order, or quadratic, approximation adds the second-order terms in the expansion,

$$\frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n f_{ij}^0 (x_i - x_i^0)(x_j - x_j^0) = \frac{1}{2} (\mathbf{x} - \mathbf{x}^0)' \mathbf{H}^0 (\mathbf{x} - \mathbf{x}^0),$$

to the preceding one. Collecting terms in the same manner as in (A-126), we have

$$y \approx \alpha + \beta' \mathbf{x} + \frac{1}{2} \mathbf{x}' \Gamma \mathbf{x}, \tag{A-127}$$

where

$$\alpha = f^0 - \mathbf{g}^0' \mathbf{x}^0 + \frac{1}{2} \mathbf{x}^0' \mathbf{H}^0 \mathbf{x}^0, \quad \beta = \mathbf{g}^0 - \mathbf{H}^0 \mathbf{x}^0 \quad \text{and} \quad \Gamma = \mathbf{H}^0.$$

A linear function can be written

$$y = \mathbf{a}' \mathbf{x} = \mathbf{x}' \mathbf{a} = \sum_{i=1}^n a_i x_i,$$

so

$$\frac{\partial(\mathbf{a}'\mathbf{x})}{\partial\mathbf{x}} = \mathbf{a}. \tag{A-128}$$

Note, in particular, that $\partial(\mathbf{a}'\mathbf{x})/\partial\mathbf{x} = \mathbf{a}$, not \mathbf{a}' . In a set of linear functions

$$\mathbf{y} = \mathbf{Ax},$$

each element y_i of \mathbf{y} is

$$y_i = \mathbf{a}'_i\mathbf{x},$$

where \mathbf{a}'_i is the i th row of \mathbf{A} [see (A-14)]. Therefore,

$$\frac{\partial y_i}{\partial\mathbf{x}} = \mathbf{a}'_i = \text{transpose of } i\text{th row of } \mathbf{A},$$

and

$$\begin{bmatrix} \partial y_1/\partial\mathbf{x}' \\ \partial y_2/\partial\mathbf{x}' \\ \dots \\ \partial y_n/\partial\mathbf{x}' \end{bmatrix} = \begin{bmatrix} \mathbf{a}'_1 \\ \mathbf{a}'_2 \\ \dots \\ \mathbf{a}'_n \end{bmatrix}.$$

Collecting all terms, we find that $\partial\mathbf{Ax}/\partial\mathbf{x}' = \mathbf{A}$, whereas the more familiar form will be

$$\frac{\partial\mathbf{Ax}}{\partial\mathbf{x}} = \mathbf{A}'. \tag{A-129}$$

A quadratic form is written

$$\mathbf{x}'\mathbf{Ax} = \sum_{i=1}^n \sum_{j=1}^n x_i x_j a_{ij}. \tag{A-130}$$

For example,

$$\mathbf{A} = \begin{bmatrix} 1 & 3 \\ 3 & 4 \end{bmatrix},$$

so that

$$\mathbf{x}'\mathbf{Ax} = 1x_1^2 + 4x_2^2 + 6x_1x_2.$$

Then

$$\frac{\partial\mathbf{x}'\mathbf{Ax}}{\partial\mathbf{x}} = \begin{bmatrix} 2x_1 + 6x_2 \\ 6x_1 + 8x_2 \end{bmatrix} = \begin{bmatrix} 2 & 6 \\ 6 & 8 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = 2\mathbf{Ax}, \tag{A-131}$$

which is the general result when \mathbf{A} is a symmetric matrix. If \mathbf{A} is not symmetric, then

$$\frac{\partial(\mathbf{x}'\mathbf{Ax})}{\partial\mathbf{x}} = (\mathbf{A} + \mathbf{A}')\mathbf{x}. \tag{A-132}$$

Referring to the preceding double summation, we find that for each term, the coefficient on a_{ij} is $x_i x_j$. Therefore,

$$\frac{\partial(\mathbf{x}'\mathbf{Ax})}{\partial a_{ij}} = x_i x_j.$$

840 APPENDIX A ♦ Matrix Algebra

The square matrix whose ij th element is $x_i x_j$ is \mathbf{xx}' , so

$$\frac{\partial(\mathbf{x}'\mathbf{A}\mathbf{x})}{\partial\mathbf{A}} = \mathbf{xx}' \quad (\text{A-133})$$

Derivatives involving determinants appear in maximum likelihood estimation. From the cofactor expansion in (A-51),

$$\frac{\partial|\mathbf{A}|}{\partial a_{ij}} = (-1)^{i+j} |\mathbf{A}_{ji}| = c_{ij}$$

where $|\mathbf{C}_{ji}|$ is the ji th cofactor in \mathbf{A} . The inverse of \mathbf{A} can be computed using

$$\mathbf{A}_{ij}^{-1} = \frac{(-1)^{i+j} |\mathbf{C}_{ji}|}{|\mathbf{A}|}$$

(note the reversal of the subscripts), which implies that

$$\frac{\partial \ln|\mathbf{A}|}{\partial a_{ij}} = \frac{(-1)^{i+j} |\mathbf{C}_{ji}|}{|\mathbf{A}|}$$

or, collecting terms,

$$\frac{\partial \ln|\mathbf{A}|}{\partial \mathbf{A}} = \mathbf{A}^{-1}$$

Since the matrices for which we shall make use of this calculation will be symmetric in our applications, the transposition will be unnecessary.

A.8.2 OPTIMIZATION

Consider finding the x where $f(x)$ is maximized or minimized. Since $f'(x)$ is the slope of $f(x)$, either optimum must occur where $f'(x) = 0$. Otherwise, the function will be increasing or decreasing at x . This result implies the **first-order** or **necessary condition for an optimum** (maximum or minimum):

$$\frac{dy}{dx} = 0 \quad (\text{A-134})$$

For a maximum, the function must be concave; for a minimum, it must be convex. The **sufficient condition for an optimum** is:

$$\begin{aligned} \text{For a maximum, } \frac{d^2y}{dx^2} < 0; \\ \text{for a minimum, } \frac{d^2y}{dx^2} > 0. \end{aligned} \quad (\text{A-135})$$

Some functions, such as the sine and cosine functions, have many **local optima**, that is, many minima and maxima. A function such as $(\cos x)/(1+x^2)$, which is a damped cosine wave, does as well but differs in that although it has many local maxima, it has one, at $x = 0$, at which $f(x)$ is greater than it is at any other point. Thus, $x = 0$ is the **global maximum**, whereas the other maxima are only **local maxima**. Certain functions, such as a quadratic, have only a single optimum. These functions are **globally concave** if the optimum is a maximum and **globally convex** if it is a minimum.

For maximizing or minimizing a function of several variables, the first-order conditions are

$$\frac{\partial f(\mathbf{x})}{\partial \mathbf{x}} = \mathbf{0} \quad (\text{A-136})$$

This result is interpreted in the same manner as the necessary condition in the univariate case. At the optimum, it must be true that no small change in any variable leads to an improvement in the function value. In the single-variable case, d^2y/dx^2 must be positive for a minimum and negative for a maximum. The second-order condition for an optimum in the multivariate case is that, at the optimizing value,

$$\mathbf{H} = \frac{\partial^2 f(\mathbf{x})}{\partial \mathbf{x} \partial \mathbf{x}'} \tag{A-137}$$

must be positive definite for a minimum and negative definite for a maximum.

In a single-variable problem, the second-order condition can usually be verified by inspection. This situation will not generally be true in the multivariate case. As discussed earlier, checking the definiteness of a matrix is, in general, a difficult problem. For most of the problems encountered in econometrics, however, the second-order condition will be implied by the structure of the problem. That is, the matrix \mathbf{H} will usually be of such a form that it is always definite.

For an example of the preceding, consider the problem

$$\text{maximize } R = \mathbf{a}'\mathbf{x} - \mathbf{x}'\mathbf{A}\mathbf{x},$$

where

$$\mathbf{a}' = (5 \quad 4 \quad 2)$$

and

$$\mathbf{A} = \begin{bmatrix} 2 & 1 & 3 \\ 1 & 3 & 2 \\ 3 & 2 & 5 \end{bmatrix}$$

Using some now familiar results, we obtain

$$\frac{\partial R}{\partial \mathbf{x}} = \mathbf{a} - 2\mathbf{A}\mathbf{x} = \begin{bmatrix} 5 \\ 4 \\ 2 \end{bmatrix} - \begin{bmatrix} 4 & 2 & 6 \\ 2 & 6 & 4 \\ 6 & 4 & 10 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \mathbf{0}. \tag{A-138}$$

The solutions are

$$\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} 4 & 2 & 6 \\ 2 & 6 & 4 \\ 6 & 4 & 10 \end{bmatrix}^{-1} \begin{bmatrix} 5 \\ 4 \\ 2 \end{bmatrix} = \begin{bmatrix} 11.25 \\ 1.75 \\ -7.25 \end{bmatrix}$$

The sufficient condition is that

$$\frac{\partial^2 R(\mathbf{x})}{\partial \mathbf{x} \partial \mathbf{x}'} = -2\mathbf{A} = \begin{bmatrix} -4 & -2 & -6 \\ -2 & -6 & -4 \\ -6 & -4 & -10 \end{bmatrix} \tag{A-139}$$

must be negative definite. The three characteristic roots of this matrix are -15.746 , -4 , and -0.25403 . Since all three roots are negative, the matrix is negative definite, as required.

In the preceding, it was necessary to compute the characteristic roots of the Hessian to verify the sufficient condition. For a general matrix of order larger than 2, this will normally require a computer. Suppose, however, that \mathbf{A} is of the form

$$\mathbf{A} = \mathbf{B}'\mathbf{B},$$

where \mathbf{B} is some known matrix. Then, as shown earlier, we know that \mathbf{A} will always be positive definite (assuming that \mathbf{B} has full rank). In this case, it is not necessary to calculate the characteristic roots of \mathbf{A} to verify the sufficient conditions.

A.8.3 CONSTRAINED OPTIMIZATION

It is often necessary to solve an optimization problem subject to some constraints on the solution. One method is merely to “solve out” the constraints. For example, in the maximization problem considered earlier, suppose that the constraint $x_1 = x_2 - x_3$ is imposed on the solution. For a single constraint such as this one, it is possible merely to substitute the right-hand side of this equation for x_1 in the objective function and solve the resulting problem as a function of the remaining two variables. For more general constraints, however, or when there is more than one constraint, the method of Lagrange multipliers provides a more straightforward method of solving the problem. We

$$\begin{aligned} \text{maximize}_x f(\mathbf{x}) \text{ subject to } c_1(\mathbf{x}) &= 0, \\ c_2(\mathbf{x}) &= 0, \\ &\dots \\ c_J(\mathbf{x}) &= 0. \end{aligned} \tag{A-140}$$

The Lagrangean approach to this problem is to find the stationary points—that is, the points at which the derivatives are zero—of

$$L^*(\mathbf{x}, \lambda) = f(\mathbf{x}) + \sum_{j=1}^J \lambda_j c_j(\mathbf{x}) = f(\mathbf{x}) + \lambda' \mathbf{c}(\mathbf{x}). \tag{A-141}$$

The solutions satisfy the equations

$$\begin{aligned} \frac{\partial L^*}{\partial \mathbf{x}} &= \frac{\partial f(\mathbf{x})}{\partial \mathbf{x}} + \frac{\partial \lambda' \mathbf{c}(\mathbf{x})}{\partial \mathbf{x}} = \mathbf{0} \ (n \times 1), \\ \frac{\partial L^*}{\partial \lambda} &= \mathbf{c}(\mathbf{x}) = \mathbf{0} \ (J \times 1). \end{aligned} \tag{A-142}$$

The second term in $\partial L^*/\partial \mathbf{x}$ is

$$\frac{\partial \lambda' \mathbf{c}(\mathbf{x})}{\partial \mathbf{x}} = \frac{\partial \mathbf{c}(\mathbf{x})' \lambda}{\partial \mathbf{x}} = \left[\frac{\partial \mathbf{c}(\mathbf{x})'}{\partial \mathbf{x}} \right] \lambda = \mathbf{C}' \lambda, \tag{A-143}$$

where \mathbf{C} is the matrix of derivatives of the constraints with respect to \mathbf{x} . The j th row of the $J \times n$ matrix \mathbf{C} is the vector of derivatives of the j th constraint, $c_j(\mathbf{x})$, with respect to \mathbf{x}' . Upon collecting terms, the first-order conditions are

$$\begin{aligned} \frac{\partial L^*}{\partial \mathbf{x}} &= \frac{\partial f(\mathbf{x})}{\partial \mathbf{x}} + \mathbf{C}' \lambda = \mathbf{0}, \\ \frac{\partial L^*}{\partial \lambda} &= \mathbf{c}(\mathbf{x}) = \mathbf{0}. \end{aligned} \tag{A-144}$$

There is one very important aspect of the constrained solution to consider. In the unconstrained solution, we have $\partial f(\mathbf{x})/\partial \mathbf{x} = \mathbf{0}$. From (A-144), we obtain, for a constrained solution,

$$\frac{\partial f(\mathbf{x})}{\partial \mathbf{x}} = -\mathbf{C}' \lambda, \tag{A-145}$$

which will not equal $\mathbf{0}$ unless $\lambda = \mathbf{0}$. This result has two important implications:

- The constrained solution cannot be superior to the unconstrained solution. This is implied by the nonzero gradient at the constrained solution. (That is, unless $\mathbf{C} = \mathbf{0}$ which could happen if the constraints were nonlinear. But, even if so, the solution is still no better than the unconstrained optimum.)
- If the Lagrange multipliers are zero, then the constrained solution will equal the unconstrained solution.

To continue the example begun earlier, suppose that we add the following conditions:

$$x_1 - x_2 + x_3 = 0,$$

$$x_1 + x_2 + x_3 = 0.$$

To put this in the format of the general problem, write the constraints as $\mathbf{c}(\mathbf{x}) = \mathbf{C}\mathbf{x} = \mathbf{0}$, where

$$\mathbf{C} = \begin{bmatrix} 1 & -1 & 1 \\ 1 & 1 & 1 \end{bmatrix}.$$

The Lagrangean function is

$$R^*(\mathbf{x}, \lambda) = \mathbf{a}'\mathbf{x} - \mathbf{x}'\mathbf{A}\mathbf{x} + \lambda'\mathbf{C}\mathbf{x}.$$

Note the dimensions and arrangement of the various parts. In particular, \mathbf{C} is a 2×3 matrix, with one row for each constraint and one column for each variable in the objective function. The vector of Lagrange multipliers thus has two elements, one for each constraint. The necessary conditions are

$$\mathbf{a} - 2\mathbf{A}\mathbf{x} + \mathbf{C}'\lambda = \mathbf{0} \quad (\text{three equations}) \tag{A-146}$$

and

$$\mathbf{C}\mathbf{x} = \mathbf{0} \quad (\text{two equations}).$$

These may be combined in the single equation

$$\begin{bmatrix} -2\mathbf{A} & \mathbf{C}' \\ \mathbf{C} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \lambda \end{bmatrix} = \begin{bmatrix} -\mathbf{a} \\ \mathbf{0} \end{bmatrix}.$$

Using the partitioned inverse of (A-74) produces the solutions

$$\lambda = -[\mathbf{C}\mathbf{A}^{-1}\mathbf{C}']^{-1}\mathbf{C}\mathbf{A}^{-1}\mathbf{a} \tag{A-147}$$

and

$$\mathbf{x} = \frac{1}{2}\mathbf{A}^{-1}[\mathbf{I} - \mathbf{C}'(\mathbf{C}\mathbf{A}^{-1}\mathbf{C}')^{-1}\mathbf{C}\mathbf{A}^{-1}]\mathbf{a}. \tag{A-148}$$

The two results, (A-147) and (A-148), yield analytic solutions for λ and \mathbf{x} . For the specific matrices and vectors of the example, these are $\lambda = [-0.5 \ -7.5]'$, and the constrained solution vector, $\mathbf{x}^* = [1.5 \ 0 \ -1.5]'$. Note that in computing the solution to this sort of problem, it is not necessary to use the rather cumbersome form of (A-148). Once λ is obtained from (A-147), the solution can be inserted in (A-146) for a much simpler computation. The solution

$$\mathbf{x} = \frac{1}{2}\mathbf{A}^{-1}\mathbf{a} + \frac{1}{2}\mathbf{A}^{-1}\mathbf{C}'\lambda$$

suggests a useful result for the constrained optimum:

$$\text{constrained solution} = \text{unconstrained solution} + [2\mathbf{A}]^{-1}\mathbf{C}'\lambda. \tag{A-149}$$

Finally, by inserting the two solutions in the original function, we find that $R = 24.375$ and $R^* = 2.25$, which illustrates again that the constrained solution (in this *maximization* problem) is inferior to the unconstrained solution.

A.8.4 TRANSFORMATIONS

If a function is strictly monotonic, then it is a **one-to-one function**. Each y is associated with exactly one value of x , and vice versa. In this case, an **inverse function** exists, which expresses x as a function of y , written

$$y = f(x)$$

and

$$x = f^{-1}(y).$$

An example is the inverse relationship between the log and the exponential functions.

The slope of the inverse function,

$$J = \frac{dx}{dy} = \frac{df^{-1}(y)}{dy} = f^{-1\prime}(y),$$

is the **Jacobian** of the transformation from y to x . For example, if

$$y = a + bx,$$

then

$$x = -\frac{a}{b} + \left[\frac{1}{b}\right] y$$

is the inverse transformation and

$$J = \frac{dx}{dy} = \frac{1}{b}.$$

Looking ahead to the statistical application of this concept, we observe that if $y = f(x)$ were *vertical*, then this would no longer be a functional relationship. The same x would be associated with more than one value of y . In this case, at this value of x , we would find that $J = 0$, indicating a singularity in the function.

If \mathbf{y} is a column vector of functions, $\mathbf{y} = \mathbf{f}(\mathbf{x})$, then

$$\mathbf{J} = \frac{\partial \mathbf{x}}{\partial \mathbf{y}'} = \begin{bmatrix} \partial x_1 / \partial y_1 & \partial x_1 / \partial y_2 & \cdots & \partial x_1 / \partial y_n \\ \partial x_2 / \partial y_1 & \partial x_2 / \partial y_2 & \cdots & \partial x_2 / \partial y_n \\ \vdots & \vdots & \ddots & \vdots \\ \partial x_n / \partial y_1 & \partial x_n / \partial y_2 & \cdots & \partial x_n / \partial y_n \end{bmatrix}$$

Consider the set of linear functions $\mathbf{y} = \mathbf{A}\mathbf{x} = \mathbf{f}(\mathbf{x})$. The inverse transformation is $\mathbf{x} = \mathbf{f}^{-1}(\mathbf{y})$, which will be

$$\mathbf{x} = \mathbf{A}^{-1}\mathbf{y},$$

if \mathbf{A} is nonsingular. If \mathbf{A} is singular, then there is no inverse transformation. Let \mathbf{J} be the matrix of partial derivatives of the inverse functions:

$$\mathbf{J} = \begin{bmatrix} \partial x_i \\ \partial y_j \end{bmatrix}.$$

The absolute value of the determinant of \mathbf{J} ,

$$\text{abs}(|\mathbf{J}|) = \left| \frac{\partial \mathbf{x}}{\partial \mathbf{y}'} \right|,$$

is the **Jacobian** determinant of the transformation from \mathbf{y} to \mathbf{x} . In the nonsingular case,

$$\text{abs}(|\mathbf{J}|) = \text{abs}(|\mathbf{A}^{-1}|) = \frac{1}{\text{abs}(|\mathbf{A}|)}.$$

In the singular case, the matrix of partial derivatives will be singular and the determinant of the Jacobian will be zero. In this instance, the singular Jacobian implies that \mathbf{A} is singular or, equivalently, that the transformations from \mathbf{x} to \mathbf{y} are functionally dependent. The singular case is analogous to the single-variable case.

Clearly, if the vector \mathbf{x} is given, then $\mathbf{y} = \mathbf{A}\mathbf{x}$ can be computed from \mathbf{x} . Whether \mathbf{x} can be deduced from \mathbf{y} is another question. Evidently, it depends on the Jacobian. If the Jacobian is not zero, then the inverse transformations exist, and we can obtain \mathbf{x} . If not, then we cannot obtain \mathbf{x} .

APPENDIX B



PROBABILITY AND DISTRIBUTION THEORY

B.1 INTRODUCTION

This appendix reviews the distribution theory used later in the book. Since a previous course in statistics is assumed, most of the results will be stated without proof. The more advanced results in the later sections will be developed in greater detail.

B.2 RANDOM VARIABLES

We view our observation on some aspect of the economy as the **outcome** of a random process which is almost never under our (the analyst's) control. In the current literature, the descriptive (and perspective laden) term **data generating process**, or DGP is often used for this underlying mechanism. The observed (measured) outcomes of the process are assigned unique numeric values. The assignment is one to one; each outcome gets one value, and no two distinct outcomes receive the same value. This outcome variable, X , is a **random variable** because, until the data are actually observed, it is uncertain what value X will take. Probabilities are associated with outcomes to quantify this uncertainty. We usually use capital letters for the "name" of a random variable and lowercase letters for the values it takes. Thus, the probability that X takes a particular value x might be denoted $\text{Prob}(X = x)$.

A random variable is **discrete** if the set of outcomes is either finite in number or countably infinite. The random variable is **continuous** if the set of outcomes is infinitely divisible and, hence, not countable. These definitions will correspond to the types of data we observe in practice. Counts of occurrences will provide observations on discrete random variables, whereas measurements such as time or income will give observations on continuous random variables.

B.2.1 PROBABILITY DISTRIBUTIONS

A listing of the values x taken by a random variable X and their associated probabilities is a **probability distribution**, $f(x)$. For a discrete random variable,

$$f(x) = \text{Prob}(X = x). \quad (\mathbf{B-1})$$

846 APPENDIX B ♦ Probability and Distribution Theory

The axioms of probability require that

1. $0 \leq \text{Prob}(X = x) \leq 1.$ (B-2)

2. $\sum_x f(x) = 1.$ (B-3)

For the continuous case, the probability associated with any particular point is zero, and we can only assign positive probabilities to intervals in the range of x . The **probability density function (pdf)** is defined so that $f(x) \geq 0$ and

1. $\text{Prob}(a \leq x \leq b) = \int_a^b f(x) dx \geq 0.$ (B-4)

This result is the area under $f(x)$ in the range from a to b . For a continuous variable,

2. $\int_{-\infty}^{+\infty} f(x) dx = 1.$ (B-5)

If the range of x is not infinite, then it is understood that $f(x) = 0$ anywhere outside the appropriate range. Since the probability associated with any individual point is 0,

$$\begin{aligned} \text{Prob}(a \leq x \leq b) &= \text{Prob}(a \leq x < b) \\ &= \text{Prob}(a < x \leq b) \\ &= \text{Prob}(a < x < b). \end{aligned}$$

B.2.2 CUMULATIVE DISTRIBUTION FUNCTION

For any random variable X , the probability that X is less than or equal to a is denoted $F(a)$. $F(x)$ is the **cumulative distribution function (cdf)**. For a discrete random variable,

$$F(x) = \sum_{X \leq x} f(X) = \text{Prob}(X \leq x). \quad (\text{B-6})$$

In view of the definition of $f(x)$,

$$f(x_i) = F(x_i) - F(x_{i-1}). \quad (\text{B-7})$$

For a continuous random variable,

$$F(x) = \int_{-\infty}^x f(t) dt \quad (\text{B-8})$$

and

$$f(x) = \frac{dF(x)}{dx}. \quad (\text{B-9})$$

In both the continuous and discrete cases, $F(x)$ must satisfy the following properties:

1. $0 \leq F(x) \leq 1.$
2. If $x > y$, then $F(x) \geq F(y).$
3. $F(+\infty) = 1.$
4. $F(-\infty) = 0.$

From the definition of the cdf,

$$\text{Prob}(a < x \leq b) = F(b) - F(a). \quad (\text{B-10})$$

Any valid pdf will imply a valid cdf, so there is no need to verify these conditions separately.

B.3 EXPECTATIONS OF A RANDOM VARIABLE

DEFINITION B.1 Mean of a Random Variable

The *mean, or expected value, of a random variable* is

$$E[x] = \begin{cases} \sum_x x f(x) & \text{if } x \text{ is discrete,} \\ \int_x x f(x) dx & \text{if } x \text{ is continuous.} \end{cases} \quad (\text{B-11})$$

The notation \sum_x or \int_x , used henceforth, means the sum or integral over the entire range of values of x . The mean is usually denoted μ . It is a weighted average of the values taken by x , where the weights are the respective probabilities. It is not necessarily a value actually taken by the random variable. For example, the expected number of heads in one toss of a fair coin is $\frac{1}{2}$.

Other **measures of central tendency** are the **median**, which is the value m such that $\text{Prob}(X \leq m) \geq \frac{1}{2}$ and $\text{Prob}(X \geq m) \geq \frac{1}{2}$, and the **mode**, which is the value of x at which $f(x)$ takes its maximum. The first of these measures is more frequently used than the second. Loosely speaking, the median corresponds more closely than the mean to the middle of a distribution. It is unaffected by extreme values. In the discrete case, the modal value of x has the highest probability of occurring.

Let $g(x)$ be a function of x . The function that gives the expected value of $g(x)$ is denoted

$$E[g(x)] = \begin{cases} \sum_x g(x) \text{Prob}(X = x) & \text{if } X \text{ is discrete,} \\ \int_x g(x) f(x) dx & \text{if } X \text{ is continuous.} \end{cases} \quad (\text{B-12})$$

If $g(x) = a + bx$ for constants a and b , then

$$E[a + bx] = a + bE[x].$$

An important case is the expected value of a constant a , which is just a .

DEFINITION B.2 Variance of a Random Variable

The *variance of a random variable* is

$$\begin{aligned} \text{Var}[x] &= E[(x - \mu)^2] \\ &= \begin{cases} \sum_x (x - \mu)^2 f(x) & \text{if } x \text{ is discrete,} \\ \int_x (x - \mu)^2 f(x) dx & \text{if } x \text{ is continuous.} \end{cases} \end{aligned} \quad (\text{B-13})$$

$\text{Var}[x]$, which must be positive, is usually denoted σ^2 . This function is a measure of the dispersion of a distribution. Computation of the variance is simplified by using the following

848 APPENDIX B ♦ Probability and Distribution Theory

important result:

$$\text{Var}[x] = E[x^2] - \mu^2. \tag{B-14}$$

A convenient corollary to (B-14) is

$$E[x^2] = \sigma^2 + \mu^2. \tag{B-15}$$

By inserting $y = a + bx$ in (B-13) and expanding, we find that

$$\text{Var}[a + bx] = b^2 \text{Var}[x], \tag{B-16}$$

which implies, for any constant a , that

$$\text{Var}[a] = 0. \tag{B-17}$$

To describe a distribution, we usually use σ , the positive square root, which is the **standard deviation** of x . The standard deviation can be interpreted as having the same units of measurement as x and μ . For any random variable x and any positive constant k , the **Chebychev inequality** states that

$$\text{Prob}(\mu - k\sigma \leq x \leq \mu + k\sigma) \geq 1 - \frac{1}{k^2}. \tag{B-18}$$

Two other measures often used to describe a probability distribution are

$$\text{skewness} = E[(x - \mu)^3]$$

and

$$\text{kurtosis} = E[(x - \mu)^4].$$

Skewness is a measure of the asymmetry of a distribution. For symmetric distributions,

$$f(\mu - x) = f(\mu + x)$$

and

$$\text{skewness} = 0.$$

For asymmetric distributions, the skewness will be positive if the “long tail” is in the positive direction. Kurtosis is a measure of the thickness of the tails of the distribution. A shorthand expression for other **central moments** is

$$\mu_r = E[(x - \mu)^r].$$

Since μ_r tends to explode as r grows, the normalized measure, μ_r/σ^r , is often used for description. Two common measures are

$$\text{skewness coefficient} = \frac{\mu_3}{\sigma^3}$$

and

$$\text{degree of excess} = \frac{\mu_4}{\sigma^4} - 3.$$

The second is based on the normal distribution, which has excess of zero.

For any two functions $g_1(x)$ and $g_2(x)$,

$$E[g_1(x) + g_2(x)] = E[g_1(x)] + E[g_2(x)]. \tag{B-19}$$

For the general case of a possibly nonlinear $g(x)$,

$$E[g(x)] = \int_x g(x) f(x) dx \tag{B-20}$$

and

$$\text{Var}[g(x)] = \int_x (g(x) - E[g(x)])^2 f(x) dx. \quad (\text{B-21})$$

(For convenience, we shall omit the equivalent definitions for discrete variables in the following discussion and use the integral to mean either integration or summation, whichever is appropriate.)

A device used to approximate $E[g(x)]$ and $\text{Var}[g(x)]$ is the linear Taylor series approximation:

$$g(x) \approx [g(x^0) - g'(x^0)x^0] + g'(x^0)x = \beta_1 + \beta_2x = g^*(x). \quad (\text{B-22})$$

If the approximation is reasonably accurate, then the mean and variance of $g^*(x)$ will be approximately equal to the mean and variance of $g(x)$. A natural choice for the expansion point is $x^0 = \mu = E(x)$. Inserting this value in (B-22) gives

$$g(x) \approx [g(\mu) - g'(\mu)\mu] + g'(\mu)x, \quad (\text{B-23})$$

so that

$$E[g(x)] \approx g(\mu) \quad (\text{B-24})$$

and

$$\text{Var}[g(x)] \approx [g'(\mu)]^2 \text{Var}[x]. \quad (\text{B-25})$$

A point to note in view of (B-22) to (B-24) is that $E[g(x)]$ will generally not equal $g(E[x])$. For the special case in which $g(x)$ is concave—that is, where $g''(x) < 0$ —we know from **Jensen's inequality** that $E[g(x)] \leq g(E[x])$. For example, $E[\log(x)] \leq \log(E[x])$.

B.4 SOME SPECIFIC PROBABILITY DISTRIBUTIONS

Certain experimental situations naturally give rise to specific probability distributions. In the majority of cases in economics, however, the distributions used are merely models of the observed phenomena. Although the normal distribution, which we shall discuss at length, is the mainstay of econometric research, economists have used a wide variety of other distributions. A few are discussed here.¹

B.4.1 THE NORMAL DISTRIBUTION

The general form of the normal distribution with mean μ and standard deviation σ is

$$f(x | \mu, \sigma^2) = \frac{1}{\sigma\sqrt{2\pi}} e^{-1/2[(x-\mu)^2/\sigma^2]}. \quad (\text{B-26})$$

This result is usually denoted $x \sim N[\mu, \sigma^2]$. The standard notation $x \sim f(x)$ is used to state that "x has probability distribution $f(x)$." Among the most useful properties of the normal distribution

¹A much more complete listing appears in Maddala (1977, Chaps. 3 and 18) and in most mathematical statistics textbooks. See also Poirier (1995) and Stuart and Ord (1989). Another useful reference is Evans, Hastings and Peacock (1993). Johnson et al. (1974, 1993, 1994, 1995, 1997) is an encyclopedic reference on the subject of statistical distributions.

is its preservation under linear transformation.

$$\text{If } x \sim N[\mu, \sigma^2], \text{ then } (a + bx) \sim N[a + b\mu, b^2\sigma^2]. \tag{B-27}$$

One particularly convenient transformation is $a = -\mu/\sigma$ and $b = 1/\sigma$. The resulting variable $z = (x - \mu)/\sigma$ has the **standard normal distribution**, denoted $N[0, 1]$, with density

$$\phi(z) = \frac{1}{\sqrt{2\pi}} e^{-z^2/2}. \tag{B-28}$$

The specific notation $\phi(z)$ is often used for this distribution and $\Phi(z)$ for its cdf. It follows from the definitions above that if $x \sim N[\mu, \sigma^2]$, then

$$f(x) = \frac{1}{\sigma} \phi\left[\frac{x - \mu}{\sigma}\right].$$


Figure B.1 shows the densities of the standard normal distribution and the normal distribution with mean 0.5, which shifts the distribution to the right, and standard deviation 1.3 which, it can be seen, scales the density so that it is shorter but wider. (The graph is a bit deceiving unless you look closely; both densities are symmetric.)

Tables of the standard normal cdf appear in most statistics and econometrics textbooks. Because the form of the distribution does not change under a linear transformation, it is not necessary to tabulate the distribution for other values of μ and σ . For any normally distributed variable,

$$\text{Prob}(a < x < b) = \text{Prob}\left(\frac{a - \mu}{\sigma} < \frac{x - \mu}{\sigma} < \frac{b - \mu}{\sigma}\right), \tag{B-29}$$

which can always be read from a table of the standard normal distribution. In addition, because the distribution is symmetric, $\Phi(-z) = 1 - \Phi(z)$. Hence, it is not necessary to tabulate both the negative and positive halves of the distribution.

FIGURE B.1 The Normal Distribution.



B.4.2 THE CHI-SQUARED, *t*, AND *F* DISTRIBUTIONS

The chi-squared, *t*, and *F* distributions are derived from the normal distribution. They arise in econometrics as sums of *n* or *n*₁ and *n*₂ other variables. These three distributions have associated with them one or two “degrees of freedom” parameters, which for our purposes will be the number of variables in the relevant sum.

The first of the essential results is

- If $z \sim N[0, 1]$, then $x = z^2 \sim$ chi-squared[1]—that is, **chi-squared** with one degree of freedom—denoted

$$z^2 \sim \chi^2[1]. \tag{B-30}$$

This distribution is a skewed distribution with mean 1 and variance 2. The second result is

- If x_1, \dots, x_n are *n* independent chi-squared[1] variables, then

$$\sum_{i=1}^n x_i \sim \text{chi-squared}[n]. \tag{B-31}$$

The mean and variance of a chi-squared variable with *n* degrees of freedom are *n* and 2*n*, respectively. A number of useful corollaries can be derived using (B-30) and (B-31).

- If $z_i, i = 1, \dots, n$, are independent $N[0, 1]$ variables, then

$$\sum_{i=1}^n z_i^2 \sim \chi^2[n]. \tag{B-32}$$

- If $z_i, i = 1, \dots, n$, are independent $N[0, \sigma^2]$ variables, then

$$\sum_{i=1}^n (z_i/\sigma)^2 \sim \chi^2[n]. \tag{B-33}$$

- If x_1 and x_2 are independent chi-squared variables with *n*₁ and *n*₂ degrees of freedom, respectively, then

$$x_1 + x_2 \sim \chi^2[n_1 + n_2]. \tag{B-34}$$

This result can be generalized to the sum of an arbitrary number of independent chi-squared variables.

Figure B.2 shows the chi-squared density for three degrees of freedom. The amount of skewness declines as the number of degrees of freedom rises. Unlike the normal distribution, a separate table is required for the chi-squared distribution for each value of *n*. Typically, only a few percentage points of the distribution are tabulated for each *n*. Table G.3 in Appendix G of this book gives upper (right) tail areas for a number of values.

- If x_1 and x_2 are two independent chi-squared variables with degrees of freedom parameters *n*₁ and *n*₂, respectively, then the ratio

$$F[n_1, n_2] = \frac{x_1/n_1}{x_2/n_2} \tag{B-35}$$

has the ***F* distribution** with *n*₁ and *n*₂ degrees of freedom.

The two degrees of freedom parameters *n*₁ and *n*₂ are the numerator and denominator degrees of freedom, respectively. Tables of the *F* distribution must be computed for each pair of values of (*n*₁, *n*₂). As such, only one or two specific values, such as the 95 percent and 99 percent upper tail values, are tabulated in most cases.




FIGURE B.2 The Chi-squared [3] Distribution.

- If z is an $N[0, 1]$ variable and x is $\chi^2[n]$ and is independent of z , then the ratio

$$t[n] = \frac{z}{\sqrt{x/n}} \tag{B-36}$$

has the t distribution with n degrees of freedom.

The t distribution has the same shape as the normal distribution but has thicker tails. Figure B.3 illustrates the t distributions with three and 10 degrees of freedom with the standard normal distribution. Two effects that can be seen in the figure are how the distribution changes as the degrees of freedom increases, and, overall, the similarity of the t distribution to the standard normal. This distribution is tabulated in the same manner as the chi-squared distribution, with several specific cutoff points corresponding to specified tail areas for various values of the degrees of freedom parameter.

Comparing (B-35) with $n_1 = 1$ and (B-36), we see the useful relationship between the t and F distributions:

- If $t \sim t[n]$, then $t^2 \sim F[1, n]$.

If the numerator in (B-36) has a nonzero mean, then the random variable in (B-36) has a noncentral t distribution and its square has a noncentral F distribution. These distributions arise in the F tests of linear restrictions [see (6-6)] when the restrictions do not hold as follows:

1. **Noncentral chi-squared distribution.** If z has a normal distribution with mean μ and standard deviation 1, then the distribution of z^2 is noncentral chi-squared with parameters 1 and $\mu^2/2$.
 - a. If $\mathbf{z} \sim N[\mu, \Sigma]$ with J elements, then $\mathbf{z}'\Sigma^{-1}\mathbf{z}$ has a noncentral chi-squared distribution with J degrees of freedom and noncentrality parameter $\mu'\Sigma^{-1}\mu/2$, which we denote $\chi_*^2[J, \mu'\Sigma^{-1}\mu/2]$.
 - b. If $\mathbf{z} \sim N[\mu, \mathbf{I}]$ and \mathbf{M} is an idempotent matrix with rank J , then $\mathbf{z}'\mathbf{M}\mathbf{z} \sim \chi_*^2[J, \mu'\mathbf{M}\mu/2]$.




FIGURE B.3 The Standard Normal, $t[3]$ and $t[10]$ Distributions.

2. *Noncentral F distribution.* If X_1 has a noncentral chi-squared distribution with noncentrality parameter λ and degrees of freedom n_1 and X_2 has a central chi-squared distribution with degrees of freedom n_2 and is independent of X_1 , then

$$F_* = \frac{X_1/n_1}{X_2/n_2}$$

has a noncentral F distribution with parameters n_1 , n_2 , and λ .² Note that in each of these cases, the statistic and the distribution are the familiar ones, except that the effect of the nonzero mean, which induces the noncentrality, is to push the distribution to the right.

B.4.3 DISTRIBUTIONS WITH LARGE DEGREES OF FREEDOM

The chi-squared, t , and F distributions usually arise in connection with sums of sample observations. The degrees of freedom parameter in each case grows with the number of observations. We often deal with larger degrees of freedom than are shown in the tables. Thus, the standard tables are often inadequate. In all cases, however, there are **limiting distributions** that we can use when the degrees of freedom parameter grows large. The simplest case is the t distribution. The t distribution with infinite degrees of freedom is equivalent to the standard normal distribution. Beyond about 100 degrees of freedom, they are almost indistinguishable.

For degrees of freedom greater than 30, a reasonably good approximation for the distribution of the chi-squared variable x is

$$z = (2x)^{1/2} - (2n - 1)^{1/2}, \quad (\text{B-37})$$

²The denominator chi-squared could also be noncentral, but we shall not use any statistics with doubly noncentral distributions.

854 APPENDIX B ♦ Probability and Distribution Theory

which is approximately standard normally distributed. Thus,

$$\text{Prob}(\chi^2[n] \leq a) \approx \Phi[(2a)^{1/2} - (2n - 1)^{1/2}].$$

As used in econometrics, the F distribution with a large-denominator degrees of freedom is common. As n_2 becomes infinite, the denominator of F converges identically to one, so we can treat the variable

$$x = n_1 F \tag{B-38}$$

as a chi-squared variable with n_1 degrees of freedom. Since the numerator degree of freedom will typically be small, this approximation will suffice for the types of applications we are likely to encounter.³ If not, then the approximation given earlier for the chi-squared distribution can be applied to $n_1 F$.

B.4.4 SIZE DISTRIBUTIONS: THE LOGNORMAL DISTRIBUTION

In modeling size distributions, such as the distribution of firm sizes in an industry or the distribution of income in a country, the **lognormal distribution**, denoted $LN[\mu, \sigma^2]$, has been particularly useful.⁴

$$f(x) = \frac{1}{\sqrt{2\pi} \sigma x} e^{-1/2(\ln x - \mu)/\sigma)^2}, \quad x > 0.$$

A lognormal variable x has

$$E[x] = e^{\mu + \sigma^2/2}$$

and

$$\text{Var}[x] = e^{2\mu + \sigma^2} (e^{\sigma^2} - 1).$$

The relation between the normal and lognormal distributions is

$$\text{If } y \sim LN[\mu, \sigma^2], \quad \ln y \sim N[\mu, \sigma^2].$$

A useful result for transformations is given as follows:

If x has a lognormal distribution with mean θ and variance λ^2 , then

$$\ln x \sim N(\mu, \sigma^2), \quad \text{where } \mu = \ln \theta^2 - \frac{1}{2} \ln(\theta^2 + \lambda^2) \quad \text{and} \quad \sigma^2 = \ln(1 + \lambda^2/\theta^2).$$

Since the normal distribution is preserved under linear transformation,

$$\text{if } y \sim LN[\mu, \sigma^2], \quad \text{then } \ln y^r \sim N[r\mu, r^2\sigma^2].$$

If y_1 and y_2 are independent lognormal variables with $y_1 \sim LN[\mu_1, \sigma_1^2]$ and $y_2 \sim LN[\mu_2, \sigma_2^2]$, then

$$y_1 y_2 \sim LN[\mu_1 + \mu_2, \sigma_1^2 + \sigma_2^2].$$

³See Johnson and Kotz (1994) for other approximations.

⁴A study of applications of the lognormal distribution appears in Aitchison and Brown (1969).

B.4.5 THE GAMMA AND EXPONENTIAL DISTRIBUTIONS

The **gamma distribution** has been used in a variety of settings, including the study of income distribution⁵ and production functions.⁶ The general form of the distribution is

$$f(x) = \frac{\lambda^P}{\Gamma(P)} e^{-\lambda x} x^{P-1}, \quad x \geq 0, \lambda > 0, P > 0. \quad (\text{B-39})$$

Many familiar distributions are special cases, including the **exponential distribution** ($P=1$) and chi-squared ($\lambda = \frac{1}{2}, P = \frac{n}{2}$). The **Erlang distribution** results if P is a positive integer. The mean is P/λ , and the variance is P/λ^2 .

B.4.6 THE BETA DISTRIBUTION

Distributions for models are often chosen on the basis of the range within which the random variable is constrained to vary. The lognormal distribution, for example, is sometimes used to model a variable that is always nonnegative. For a variable constrained between 0 and $c > 0$, the **beta distribution** has proved useful. Its density is

$$f(x) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} \left(\frac{x}{c}\right)^{\alpha-1} \left(1 - \frac{x}{c}\right)^{\beta-1} \frac{1}{c}. \quad (\text{B-40})$$

This functional form is extremely flexible in the shapes it will accommodate. It is symmetric if $\alpha = \beta$, asymmetric otherwise, and can be hump-shaped or U-shaped. The mean is $c\alpha/(\alpha + \beta)$, and the variance is $c^2\alpha\beta/[(\alpha + \beta + 1)(\alpha + \beta)^2]$. The beta distribution has been applied in the study of labor force participation rates.⁷

B.4.7 THE LOGISTIC DISTRIBUTION

The normal distribution is ubiquitous in econometrics. But researchers have found that for some microeconomic applications, there does not appear to be enough mass in the tails of the normal distribution; observations that a model based on normality would classify as “unusual” seem not to be very unusual at all. One approach has been to use thicker-tailed symmetric distributions. The **logistic distribution** is one candidate; the cdf for a logistic random variable is denoted

$$F(x) = \Lambda(x) = \frac{1}{1 + e^{-x}}.$$

The density is $f(x) = \Lambda(x)[1 - \Lambda(x)]$. The mean and variance of this random variable are zero and $\pi^2/3$.

B.4.8 DISCRETE RANDOM VARIABLES

Modeling in economics frequently involves random variables that take integer values. In these cases, the distributions listed thus far only provide approximations that are sometimes quite inappropriate. We can build up a class of models for **discrete random variables** from the **Bernoulli distribution** for a single binomial outcome (trial)

$$\text{Prob}(x = 1) = \alpha,$$

$$\text{Prob}(x = 0) = 1 - \alpha,$$

⁵Salem and Mount (1974).

⁶Greene (1980a).

⁷Heckman and Willis (1976).




FIGURE B.4 The Poisson [3] Distribution.

where $0 \leq \alpha \leq 1$. The modeling aspect of this specification would be the assumptions that the success probability α is constant from one trial to the next and that successive trials are independent. If so, then the distribution for x successes in n trials is the **binomial distribution**,

$$\text{Prob}(X = x) = \binom{n}{x} \alpha^x (1 - \alpha)^{n-x}, \quad x = 0, 1, \dots, n.$$

The mean and variance of x are $n\alpha$ and $n\alpha(1 - \alpha)$, respectively. If the number of trials becomes large at the same time that the success probability becomes small so that the mean $n\alpha$ is stable, then, the limiting form of the binomial distribution is the **Poisson distribution**,

$$\text{Prob}(X = x) = \frac{e^{-\lambda} \lambda^x}{x!}.$$

The Poisson distribution has seen wide use in econometrics in, for example, modeling patents, crime, recreation demand, and demand for health services. An example is shown in Figure B.4.

B.5 THE DISTRIBUTION OF A FUNCTION OF A RANDOM VARIABLE

We considered finding the expected value of a function of a random variable. It is fairly common to analyze the random variable itself, which results when we compute a function of some random variable. There are three types of transformation to consider. One discrete random variable may be transformed into another, a continuous variable may be transformed into a discrete one, and one continuous variable may be transformed into another.




FIGURE B.5 Censored Distribution.

The simplest case is the first one. The probabilities associated with the new variable are computed according to the laws of probability. If y is derived from x and the function is one to one, then the probability that $Y = y(x)$ equals the probability that $X = x$. If several values of x yield the same value of y , then $\text{Prob}(Y = y)$ is the sum of the corresponding probabilities for x .

The second type of transformation is illustrated by the way individual data on income are typically obtained in a survey. Income in the population can be expected to be distributed according to some skewed, continuous distribution such as the one shown in Figure B.5.

Data are often reported categorically, as shown in the lower part of the figure. Thus, the random variable corresponding to observed income is a discrete transformation of the actual underlying continuous random variable. Suppose, for example, that the transformed variable y is the mean income in the respective interval. Then

$$\text{Prob}(Y = \mu_1) = P(-\infty < X \leq a),$$

$$\text{Prob}(Y = \mu_2) = P(a < X \leq b),$$

$$\text{Prob}(Y = \mu_3) = P(b < X \leq c),$$

and so on, which illustrates the general procedure.

If x is a continuous random variable with pdf $f_x(x)$ and if $y = g(x)$ is a continuous monotonic function of x , then the density of y is obtained by using the change of variable technique to find

the cdf of y :

$$\text{Prob}(y \leq b) = \int_{-\infty}^b f_x(g^{-1}(y))|g^{-1\prime}(y)| dy.$$

This equation can now be written as

$$\text{Prob}(y \leq b) = \int_{-\infty}^b f_y(y) dy.$$

Hence,

$$f_y(y) = f_x(g^{-1}(y))|g^{-1\prime}(y)|. \quad \text{(B-41)}$$

To avoid the possibility of a negative pdf if $g(x)$ is decreasing, we use the absolute value of the derivative in the previous expression. The term $|g^{-1\prime}(y)|$ must be nonzero for the density of y to be nonzero. In words, the probabilities associated with intervals in the range of y must be associated with intervals in the range of x . If the derivative is zero, the correspondence $y = g(x)$ is vertical, and hence all values of y in the given range are associated with the same value of x . This single point must have probability zero.

One of the most useful applications of the preceding result is the linear transformation of a normally distributed variable. If $x \sim N[\mu, \sigma^2]$, then the distribution of

$$y = \frac{x - \mu}{\sigma}$$

is found using the result above. First, the derivative is obtained from the inverse transformation

$$y = \frac{x}{\sigma} - \frac{\mu}{\sigma} \Rightarrow x = \sigma y + \mu \Rightarrow f^{-1\prime}(y) = \frac{dx}{dy} = \sigma.$$

Therefore,

$$f_y(y) = \frac{1}{\sqrt{2\pi}\sigma} e^{-[(\sigma y + \mu) - \mu]^2 / (2\sigma^2)} |\sigma| = \frac{1}{\sqrt{2\pi}} e^{-y^2/2}.$$

This is the density of a normally distributed variable with mean zero and unit standard deviation one. This is this result which makes it unnecessary to have separate tables for the different normal distributions which result from different means and variances.

B.6 REPRESENTATIONS OF A PROBABILITY DISTRIBUTION

The probability density function (pdf) is a natural and familiar way to formulate the distribution of a random variable. But, there are many other functions that are used to identify or characterize a random variable, depending on the setting. In each of these cases, we can identify some other function of the random variable that has a one to one relationship with the density. We have already used one of these quite heavily in the preceding discussion. For a random variable which has density function $f(x)$, the distribution function or pdf, $F(x)$, is an equally informative function that identifies the distribution; the relationship between $f(x)$ and $F(x)$ is defined in (B-6) for a discrete random variable and (B-8) for a continuous one. We now consider several other related functions.

For a continuous random variable, the **survival function** is $S(x) = 1 - F(x) = \text{Prob}[X > x]$. This function is widely used in epidemiology where x is time until some transition, such as recovery

from a disease. The **hazard function** for a random variable is

$$h(x) = \frac{f(x)}{S(x)} = \frac{f(x)}{1 - F(x)}$$

The hazard function is a conditional probability;

$$h(x) = \lim_{t \downarrow 0} \text{Prob}(X \leq x + t \mid X \geq x)$$

hazards have been used in econometrics in studying the duration of spells, or conditions, such as unemployment, strikes, time until business failures, and so on. The connection between the hazard and the other functions is $h(x) = -d \ln S(x)/dx$. As an exercise, you might want to verify the interesting special case of $h(x) = 1/\lambda$, a constant—the only distribution which has this characteristic is the exponential distribution noted in Section B.4.5.

For the random variable X , with probability density function $f(x)$, if the function

$$M(t) = E[e^{tx}]$$

exists, then it is the **moment-generating function**. Assuming the function exists, it can be shown that

$$d^r M(t)/dt^r \big|_{t=0} = E[x^r].$$

The moment generating function, like the survival and the hazard functions, is a unique characterization of a probability distribution. When it exists, the moment-generating function has a one-to-one correspondence with the distribution. Thus, for example, if we begin with some random variable and find that a transformation of it has a particular MGF, then we may infer that the function of the random variable has the distribution associated with that MGF. A convenient application of this result is the MGF for the normal distribution. The MGF for the standard normal distribution is $M_z(t) = e^{t^2/2}$.

A useful feature of MGFs is the following:

if x and y are independent, then the MGF of $x + y$ is $M_x(t)M_y(t)$.

This result has been used to establish the **contagion** property of some distributions, that is, the property that sums of random variables with a given distribution have that same distribution. The normal distribution is a familiar example. This is usually not the case. It is for Poisson and chi-squared random variables.

One qualification of all of the preceding is that in order for these results to hold, the MGF must exist. It will for the distributions that we will encounter in our work, but in at least one important case, we cannot be sure of this. When computing sums of random variables which may have different distributions and whose specific distributions need not be so well behaved, it is likely that the MGF of the sum does not exist. However, the characteristic function,

$$\phi(t) = E[e^{itx}]$$

will always exist, at least for relatively small t . The characteristic function is the device used to prove that certain sums of random variables converge to a normally distributed variable—that is, the characteristic function is a fundamental tool in proofs of the central limit theorem.

B.7 JOINT DISTRIBUTIONS

The **joint density function** for two random variables X and Y denoted $f(x, y)$ is defined so that

$$\text{Prob}(a \leq x \leq b, c \leq y \leq d) = \begin{cases} \sum_{a \leq x \leq b} \sum_{c \leq y \leq d} f(x, y) & \text{if } x \text{ and } y \text{ are discrete,} \\ \int_a^b \int_c^d f(x, y) dy dx & \text{if } x \text{ and } y \text{ are continuous.} \end{cases} \quad (\text{B-42})$$

The counterparts of the requirements for a univariate probability density are

$$\begin{aligned} f(x, y) &\geq 0, \\ \sum_x \sum_y f(x, y) &= 1 \quad \text{if } x \text{ and } y \text{ are discrete,} \\ \int_x \int_y f(x, y) dy dx &= 1 \quad \text{if } x \text{ and } y \text{ are continuous.} \end{aligned} \quad (\text{B-43})$$

The cumulative probability is likewise the probability of a joint event:

$$\begin{aligned} F(x, y) &= \text{Prob}(X \leq x, Y \leq y) \\ &= \begin{cases} \sum_{X \leq x} \sum_{Y \leq y} f(x, y) & \text{in the discrete case} \\ \int_{-\infty}^x \int_{-\infty}^y f(t, s) ds dt & \text{in the continuous case.} \end{cases} \end{aligned} \quad (\text{B-44})$$

B.7.1 MARGINAL DISTRIBUTIONS

A **marginal probability density** or marginal probability distribution is defined with respect to an individual variable. To obtain the marginal distributions from the joint density, it is necessary to sum or integrate out the other variable:

$$f_x(x) = \begin{cases} \sum_y f(x, y) & \text{in the discrete case} \\ \int_y f(x, s) ds & \text{in the continuous case.} \end{cases} \quad (\text{B-45})$$

and similarly for $f_y(y)$.

Two random variables are statistically independent if and only if their joint density is the product of the marginal densities:

$$f(x, y) = f_x(x) f_y(y) \Leftrightarrow x \text{ and } y \text{ are independent.} \quad (\text{B-46})$$

If (and only if) x and y are independent, then the cdf factors as well as the pdf:

$$F(x, y) = F_x(x) F_y(y) \quad (\text{B-47})$$

or

$$\text{Prob}(X \leq x, Y \leq y) = \text{Prob}(X \leq x) \text{Prob}(Y \leq y).$$

B.7.2 EXPECTATIONS IN A JOINT DISTRIBUTION

The means, variances, and higher moments of the variables in a joint distribution are defined with respect to the marginal distributions. For the mean of x in a discrete distribution,

$$\begin{aligned}
 E[x] &= \sum_x x f_x(x) \\
 &= \sum_x x \left[\sum_y f(x, y) \right] \\
 &= \sum_x \sum_y x f(x, y).
 \end{aligned}
 \tag{B-48}$$

The means of the variables in a continuous distribution are defined likewise, using integration instead of summation:

$$\begin{aligned}
 E[x] &= \int_x x f_x(x) dx \\
 &= \int_x \int_y x f(x, y) dy dx.
 \end{aligned}
 \tag{B-49}$$

Variances are computed in the same manner:

$$\begin{aligned}
 \text{Var}[x] &= \sum_x (x - E[x])^2 f_x(x) \\
 &= \sum_x \sum_y (x - E[x])^2 f(x, y).
 \end{aligned}
 \tag{B-50}$$

B.7.3 COVARIANCE AND CORRELATION

For any function $g(x, y)$,

$$E[g(x, y)] = \begin{cases} \sum_x \sum_y g(x, y) f(x, y) & \text{in the discrete case,} \\ \int_x \int_y g(x, y) f(x, y) dy dx & \text{in the continuous case.} \end{cases}
 \tag{B-51}$$

The covariance of x and y is a special case:

$$\begin{aligned}
 \text{Cov}[x, y] &= E[(x - \mu_x)(y - \mu_y)] \\
 &= E[xy] - \mu_x \mu_y \\
 &= \sigma_{xy}.
 \end{aligned}
 \tag{B-52}$$

If x and y are independent, then $f(x, y) = f_x(x) f_y(y)$ and

$$\begin{aligned}
 \sigma_{xy} &= \sum_x \sum_y f_x(x) f_y(y) (x - \mu_x)(y - \mu_y) \\
 &= \sum_x (x - \mu_x) f_x(x) \sum_y (y - \mu_y) f_y(y) \\
 &= E[x - \mu_x] E[y - \mu_y] \\
 &= 0.
 \end{aligned}$$

862 APPENDIX B ♦ Probability and Distribution Theory

The sign of the covariance will indicate the direction of covariation of X and Y . Its magnitude depends on the scales of measurement, however. In view of this fact, a preferable measure is the correlation coefficient:

$$r[x, y] = \rho_{xy} = \frac{\sigma_{xy}}{\sigma_x \sigma_y}, \quad (\text{B-53})$$

where σ_x and σ_y are the standard deviations of x and y , respectively. The correlation coefficient has the same sign as the covariance but is always between -1 and 1 and is thus unaffected by any scaling of the variables.

Variables that are uncorrelated are not necessarily independent. For example, in the discrete distribution $f(-1, 1) = f(0, 0) = f(1, 1) = \frac{1}{3}$, the correlation is zero, but $f(1, 1)$ does not equal $f_x(1)f_y(1) = (\frac{1}{3})(\frac{2}{3})$. An important exception is the joint normal distribution discussed subsequently, in which lack of correlation does imply independence.

Some general results regarding expectations in a joint distribution, which can be verified by applying the appropriate definitions, are

$$E[ax + by + c] = aE[x] + bE[y] + c, \quad (\text{B-54})$$

$$\begin{aligned} \text{Var}[ax + by + c] &= a^2 \text{Var}[x] + b^2 \text{Var}[y] + 2ab \text{Cov}[x, y] \\ &= \text{Var}[ax + by], \end{aligned} \quad (\text{B-55})$$

and

$$\text{Cov}[ax + by, cx + dy] = ac \text{Var}[x] + bd \text{Var}[y] + (ad + bc) \text{Cov}[x, y]. \quad (\text{B-56})$$

If X and Y are uncorrelated, then

$$\begin{aligned} \text{Var}[x + y] &= \text{Var}[x - y] \\ &= \text{Var}[x] + \text{Var}[y]. \end{aligned} \quad (\text{B-57})$$

For any two functions $g_1(x)$ and $g_2(y)$, if x and y are independent, then

$$E[g_1(x)g_2(y)] = E[g_1(x)]E[g_2(y)]. \quad (\text{B-58})$$

B.7.4 DISTRIBUTION OF A FUNCTION OF BIVARIATE RANDOM VARIABLES

The result for a function of a random variable in (B-41) must be modified for a joint distribution. Suppose that x_1 and x_2 have a joint distribution $f_x(x_1, x_2)$ and that y_1 and y_2 are two monotonic functions of x_1 and x_2 :

$$\begin{aligned} y_1 &= y_1(x_1, x_2) \\ y_2 &= y_2(x_1, x_2). \end{aligned}$$

Since the functions are monotonic, the inverse transformations,

$$\begin{aligned} x_1 &= x_1(y_1, y_2) \\ x_2 &= x_2(y_1, y_2), \end{aligned}$$

exist. The Jacobian of the transformations is the matrix of partial derivatives,

$$J = \begin{bmatrix} \partial x_1 / \partial y_1 & \partial x_1 / \partial y_2 \\ \partial x_2 / \partial y_1 & \partial x_2 / \partial y_2 \end{bmatrix} = \begin{bmatrix} \partial \mathbf{x} \\ \partial \mathbf{y}' \end{bmatrix}.$$

The joint distribution of y_1 and y_2 is

$$f_y(y_1, y_2) = f_x[x_1(y_1, y_2), x_2(y_1, y_2)] \text{abs}(|J|).$$

The determinant of the Jacobian must be nonzero for the transformation to exist. A zero determinant implies that the two transformations are functionally dependent.

Certainly the most common application of the preceding in econometrics is the linear transformation of a set of random variables. Suppose that x_1 and x_2 are independently distributed $N[0, 1]$, and the transformations are

$$y_1 = \alpha_1 + \beta_{11}x_1 + \beta_{12}x_2,$$

$$y_2 = \alpha_2 + \beta_{21}x_1 + \beta_{22}x_2.$$

To obtain the joint distribution of y_1 and y_2 , we first write the transformations as

$$\mathbf{y} = \mathbf{a} + \mathbf{B}\mathbf{x}.$$

The inverse transformation is

$$\mathbf{x} = \mathbf{B}^{-1}(\mathbf{y} - \mathbf{a}),$$

so the absolute value of the determinant of the Jacobian is

$$\text{abs}|J| = \text{abs}|\mathbf{B}^{-1}| = \frac{1}{\text{abs}|\mathbf{B}|}.$$

The joint distribution of \mathbf{x} is the product of the marginal distributions since they are independent. Thus,

$$f_x(\mathbf{x}) = (2\pi)^{-1} e^{-(x_1^2 + x_2^2)/2} = (2\pi)^{-1} e^{-\mathbf{x}'\mathbf{x}/2}.$$

Inserting the results for $\mathbf{x}(\mathbf{y})$ and J into $f_y(y_1, y_2)$ gives

$$f_y(\mathbf{y}) = (2\pi)^{-1} \frac{1}{\text{abs}|\mathbf{B}|} e^{-(\mathbf{y}-\mathbf{a})'(\mathbf{B}\mathbf{B}')^{-1}(\mathbf{y}-\mathbf{a})/2}.$$

This **bivariate normal distribution** is the subject of Section B.9. Note that by formulating it as we did above, we can generalize easily to the multivariate case, that is, with an arbitrary number of variables.

Perhaps the more common situation is that in which it is necessary to find the distribution of one function of two (or more) random variables. A strategy that often works in this case is to form the joint distribution of the transformed variable and one of the original variables, then integrate (or sum) the latter out of the joint distribution to obtain the marginal distribution. Thus, to find the distribution of $y_1(x_1, x_2)$, we might formulate

$$y_1 = y_1(x_1, x_2)$$

$$y_2 = x_2.$$

The absolute value of the determinant of the Jacobian would then be

$$J = \text{abs} \begin{vmatrix} \frac{\partial x_1}{\partial y_1} & \frac{\partial x_1}{\partial y_2} \\ 0 & 1 \end{vmatrix} = \text{abs} \left(\frac{\partial x_1}{\partial y_1} \right)$$

The density of y_1 would then be

$$f_{y_1}(y_1) = \int_{y_2} f_x[x_1(y_1, y_2), y_2] \text{abs}|J| dy_2.$$

B.8 CONDITIONING IN A BIVARIATE DISTRIBUTION

Conditioning and the use of conditional distributions play a pivotal role in econometric modeling. We consider some general results for a bivariate distribution. (All these results can be extended directly to the multivariate case.)

In a bivariate distribution, there is a **conditional distribution** over y for each value of x . The conditional densities are

$$f(y|x) = \frac{f(x, y)}{f_x(x)} \tag{B-59}$$

and

$$f(x|y) = \frac{f(x, y)}{f_y(y)}.$$

It follows from (B-46) that:

$$\text{If } x \text{ and } y \text{ are independent, then } f(y|x) = f_y(y) \text{ and } f(x|y) = f_x(x). \tag{B-60}$$

The interpretation is that if the variables are independent, the probabilities of events relating to one variable are unrelated to the other. The definition of conditional densities implies the important result

$$\begin{aligned} f(x, y) &= f(y|x) f_x(x) \\ &= f(x|y) f_y(y). \end{aligned} \tag{B-61}$$

B.8.1 REGRESSION: THE CONDITIONAL MEAN

A **conditional mean** is the mean of the conditional distribution and is defined by

$$E[y|x] = \begin{cases} \int_y y f(y|x) dy & \text{if } y \text{ is continuous,} \\ \sum_y y f(y|x) & \text{if } y \text{ is discrete.} \end{cases} \tag{B-62}$$

The conditional mean function $E[y|x]$ is called the **regression** of y on x .

A random variable may always be written as

$$\begin{aligned} y &= E[y|x] + (y - E[y|x]) \\ &= E[y|x] + \varepsilon. \end{aligned}$$

B.8.2 CONDITIONAL VARIANCE

A conditional variance is the variance of the conditional distribution:

$$\begin{aligned} \text{Var}[y|x] &= E[(y - E[y|x])^2 | x] \\ &= \int_y (y - E[y|x])^2 f(y|x) dy, \quad \text{if } y \text{ is continuous} \end{aligned} \tag{B-63}$$

or

$$\text{Var}[y|x] = \sum_y (y - E[y|x])^2 f(y|x), \quad \text{if } y \text{ is discrete.} \tag{B-64}$$

The computation can be simplified by using

$$\text{Var}[y|x] = E[y^2|x] - (E[y|x])^2. \tag{B-65}$$

The conditional variance is called the **scedastic function** and, like the regression, is generally a function of x . Unlike the conditional mean function, however, it is common for the conditional variance not to vary with x . We shall examine a particular case. This case does not imply, however, that $\text{Var}[y|x]$ equals $\text{Var}[y]$, which will usually not be true. It implies only that the conditional variance is a constant. The case in which the conditional variance does not vary with x is called **homoscedasticity** (same variance).

B.8.3 RELATIONSHIPS AMONG MARGINAL AND CONDITIONAL MOMENTS

Some useful results for the moments of a conditional distribution are given in the following theorems.

THEOREM B.1 Law of Iterated Expectations

$$E[y] = E_x[E[y|x]]. \tag{B-66}$$

The notation $E_x[\cdot]$ indicates the expectation over the values of x . Note that $E[y|x]$ is a function of x .

THEOREM B.2 Covariance

In any bivariate distribution,

$$\text{Cov}[x, y] = \text{Cov}_x[x, E[y|x]] = \int_x (x - E[x]) E[y|x] f_x(x) dx. \tag{B-67}$$

(Note that this is the covariance of x and a function of x .)

The preceding results provide an additional, extremely useful result for the special case in which the conditional mean function is linear in x .

THEOREM B.3 Moments in a Linear Regression

If $E[y|x] = \alpha + \beta x$, then

$$\alpha = E[y] - \beta E[x]$$

and

$$\beta = \frac{\text{Cov}[x, y]}{\text{Var}[x]}. \quad (\text{B-68})$$

The proof follows from (B-66).

The preceding theorems relate to the conditional mean in a bivariate distribution. The following theorems which also appears in various forms in regression analysis describe the conditional variance.

THEOREM B.4 Decomposition of Variance

In a joint distribution,

$$\text{Var}[y] = \text{Var}_x[E[y|x]] + E_x[\text{Var}[y|x]]. \quad (\text{B-69})$$

The notation $\text{Var}_x[\cdot]$ indicates the variance over the distribution of x . This equation states that in a bivariate distribution, the variance of y decomposes into the variance of the conditional mean function plus the expected variance around the conditional mean.

THEOREM B.5 Residual Variance in a Regression

In any bivariate distribution,

$$E_x[\text{Var}[y|x]] = \text{Var}[y] - \text{Var}_x[E[y|x]]. \quad (\text{B-70})$$

On average, conditioning reduces the variance of the variable subject to the conditioning. For example, if y is homoscedastic, then we have the unambiguous result that the variance of the conditional distribution(s) is less than or equal to the unconditional variance of y . Going a step further, we have the result that appears prominently in the bivariate normal distribution (Section B.9).

THEOREM B.6 Linear Regression and Homoscedasticity

In a bivariate distribution, if $E[y|x] = \alpha + \beta x$ and if $\text{Var}[y|x]$ is a constant, then

$$\text{Var}[y|x] = \text{Var}[y](1 - \text{Corr}^2[y, x]) = \sigma_y^2(1 - \rho_{xy}^2). \quad (\text{B-71})$$

The proof is straightforward using Theorems B.2 to B.4.

B.8.4 THE ANALYSIS OF VARIANCE

The variance decomposition result implies that in a bivariate distribution, variation in y arises from two sources:

1. Variation because $E[y|x]$ varies with x :

$$\text{regression variance} = \text{Var}_x[E[y|x]]. \quad (\text{B-72})$$

2. Variation because, in each conditional distribution, y varies around the conditional mean:

$$\text{residual variance} = E_x[\text{Var}[y|x]]. \quad (\text{B-73})$$

Thus,

$$\text{Var}[y] = \text{regression variance} + \text{residual variance}. \quad (\text{B-74})$$

In analyzing a regression, we shall usually be interested in which of the two parts of the total variance, $\text{Var}[y]$, is the larger one. A natural measure is the ratio

$$\text{coefficient of determination} = \frac{\text{regression variance}}{\text{total variance}}. \quad (\text{B-75})$$

In the setting of a linear regression, (B-75) arises from another relationship that emphasizes the interpretation of the correlation coefficient.

$$\text{If } E[y|x] = \alpha + \beta x, \text{ then the coefficient of determination} = \text{COD} = \rho^2, \quad (\text{B-76})$$

where ρ^2 is the squared correlation between x and y . We conclude that the correlation coefficient (squared) is a measure of the proportion of the variance of y accounted for by variation in the mean of y given x . It is in this sense that correlation can be interpreted as a **measure of linear association** between two variables.

B.9 THE BIVARIATE NORMAL DISTRIBUTION

A bivariate distribution that embodies many of the features described earlier is the bivariate normal, which is the joint distribution of two normally distributed variables. The density is

$$f(x, y) = \frac{1}{2\pi\sigma_x\sigma_y\sqrt{1-\rho^2}} e^{-1/2[(\varepsilon_x^2 + \varepsilon_y^2 - 2\rho\varepsilon_x\varepsilon_y)/(1-\rho^2)]} \quad (\text{B-77})$$

$$\varepsilon_x = \frac{x - \mu_x}{\sigma_x}, \quad \varepsilon_y = \frac{y - \mu_y}{\sigma_y}.$$

868 APPENDIX B ♦ Probability and Distribution Theory

The parameters μ_x , σ_x , μ_y , and σ_y are the means and standard deviations of the marginal distributions of x and y , respectively. The additional parameter ρ is the correlation between x and y . The covariance is

$$\sigma_{xy} = \rho\sigma_x\sigma_y. \quad (\text{B-78})$$

The density is defined only if ρ is not 1 or -1 , which in turn requires that the two variables not be linearly related. If x and y have a bivariate normal distribution, denoted

$$(x, y) \sim N_2[\mu_x, \mu_y, \sigma_x^2, \sigma_y^2, \rho],$$

then

- The marginal distributions are normal:

$$\begin{aligned} f_x(x) &= N[\mu_x, \sigma_x^2], \\ f_y(y) &= N[\mu_y, \sigma_y^2]. \end{aligned} \quad (\text{B-79})$$

- The conditional distributions are normal:

$$\begin{aligned} f(y|x) &= N[\alpha + \beta x, \sigma_y^2(1 - \rho^2)] \\ \alpha &= \mu_y - \beta\mu_x, \quad \beta = \frac{\sigma_{xy}}{\sigma_x^2}, \end{aligned} \quad (\text{B-80})$$

and likewise for $f(x|y)$.

- x and y are independent if and only if $\rho = 0$. The density factors into the product of the two marginal normal distributions if $\rho = 0$.

Two things to note about the conditional distributions beyond their normality are their linear regression functions and their constant conditional variances. The conditional variance is less than the unconditional variance, which is consistent with the results of the previous section.

B.10 MULTIVARIATE DISTRIBUTIONS

The extension of the results for bivariate distributions to more than two variables is direct. It is made much more convenient by using matrices and vectors. The term **random vector** applies to a vector whose elements are random variables. The joint density is $f(\mathbf{x})$, whereas the cdf is

$$F(\mathbf{x}) = \int_{-\infty}^{x_n} \int_{-\infty}^{x_{n-1}} \cdots \int_{-\infty}^{x_1} f(\mathbf{t}) dt_1 \cdots dt_{n-1} dt_n. \quad (\text{B-81})$$

Note that the cdf is an n -fold integral. The marginal distribution of any one (or more) of the n variables is obtained by integrating or summing over the other variables.

B.10.1 MOMENTS

The expected value of a vector or matrix is the vector or matrix of expected values. A mean vector is defined as

$$\boldsymbol{\mu} = \begin{bmatrix} \mu_1 \\ \mu_2 \\ \vdots \\ \mu_n \end{bmatrix} = \begin{bmatrix} E[x_1] \\ E[x_2] \\ \vdots \\ E[x_n] \end{bmatrix} = E[\mathbf{x}]. \quad (\text{B-82})$$

Define the matrix

$$(\mathbf{x} - \boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu})' = \begin{bmatrix} (x_1 - \mu_1)(x_1 - \mu_1) & (x_1 - \mu_1)(x_2 - \mu_2) & \cdots & (x_1 - \mu_1)(x_n - \mu_n) \\ (x_2 - \mu_2)(x_1 - \mu_1) & (x_2 - \mu_2)(x_2 - \mu_2) & \cdots & (x_2 - \mu_2)(x_n - \mu_n) \\ \vdots & \vdots & \ddots & \vdots \\ (x_n - \mu_n)(x_1 - \mu_1) & (x_n - \mu_n)(x_2 - \mu_2) & \cdots & (x_n - \mu_n)(x_n - \mu_n) \end{bmatrix}$$

The expected value of each element in the matrix is the covariance of the two variables in the product. (The covariance of a variable with itself is its variance.) Thus,

$$E[(\mathbf{x} - \boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu})'] = \begin{bmatrix} \sigma_{11} & \sigma_{12} & \cdots & \sigma_{1n} \\ \sigma_{21} & \sigma_{22} & \cdots & \sigma_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_{n1} & \sigma_{n2} & \cdots & \sigma_{nn} \end{bmatrix} = E[\mathbf{xx}'] - \boldsymbol{\mu}\boldsymbol{\mu}', \quad (\text{B-83})$$

which is the **covariance matrix** of the random vector \mathbf{x} . Henceforth, we shall denote the covariance matrix of a random vector in boldface, as in

$$\text{Var}[\mathbf{x}] = \boldsymbol{\Sigma}.$$

By dividing σ_{ij} by $\sigma_i\sigma_j$, we obtain the **correlation matrix**:

$$\mathbf{R} = \begin{bmatrix} 1 & \rho_{12} & \rho_{13} & \cdots & \rho_{1n} \\ \rho_{21} & 1 & \rho_{23} & \cdots & \rho_{2n} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \rho_{n1} & \rho_{n2} & \rho_{n3} & \cdots & 1 \end{bmatrix}$$

B.10.2 SETS OF LINEAR FUNCTIONS

Our earlier results for the mean and variance of a linear function can be extended to the multivariate case. For the mean,

$$\begin{aligned} E[a_1x_1 + a_2x_2 + \cdots + a_nx_n] &= E[\mathbf{a}'\mathbf{x}] \\ &= a_1E[x_1] + a_2E[x_2] + \cdots + a_nE[x_n] \\ &= a_1\mu_1 + a_2\mu_2 + \cdots + a_n\mu_n \\ &= \mathbf{a}'\boldsymbol{\mu}. \end{aligned} \quad (\text{B-84})$$

For the variance,

$$\begin{aligned} \text{Var}[\mathbf{a}'\mathbf{x}] &= E[(\mathbf{a}'\mathbf{x} - E[\mathbf{a}'\mathbf{x}])^2] \\ &= E[\{\mathbf{a}'(\mathbf{x} - E[\mathbf{x}])\}^2] \\ &= E[\mathbf{a}'(\mathbf{x} - \boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu})'\mathbf{a}] \end{aligned}$$

as $E[\mathbf{x}] = \boldsymbol{\mu}$ and $\mathbf{a}'(\mathbf{x} - \boldsymbol{\mu}) = (\mathbf{x} - \boldsymbol{\mu})'\mathbf{a}$. Since \mathbf{a} is a vector of constants,

$$\text{Var}[\mathbf{a}'\mathbf{x}] = \mathbf{a}'E[(\mathbf{x} - \boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu})']\mathbf{a} = \mathbf{a}'\boldsymbol{\Sigma}\mathbf{a} = \sum_{i=1}^n \sum_{j=1}^n a_i a_j \sigma_{ij}. \quad (\text{B-85})$$

870 APPENDIX B ♦ Probability and Distribution Theory

Since it is the expected value of a square, we know that a variance cannot be negative. As such, the preceding quadratic form is nonnegative, and the symmetric matrix Σ must be nonnegative definite.

In the set of linear functions $\mathbf{y} = \mathbf{A}\mathbf{x}$, the i th element of \mathbf{y} is $y_i = \mathbf{a}_i\mathbf{x}$, where \mathbf{a}_i is the i th row of \mathbf{A} [see result (A-14)]. Therefore,

$$E[y_i] = \mathbf{a}_i\boldsymbol{\mu}.$$

Collecting the results in a vector, we have

$$E[\mathbf{A}\mathbf{x}] = \mathbf{A}\boldsymbol{\mu}. \quad (\text{B-86})$$

For two row vectors \mathbf{a}_i and \mathbf{a}_j ,

$$\text{Cov}[\mathbf{a}'_i\mathbf{x}, \mathbf{a}'_j\mathbf{x}] = \mathbf{a}'_i\Sigma\mathbf{a}_j.$$

Since $\mathbf{a}'_i\Sigma\mathbf{a}_j$ is the ij th element of $\mathbf{A}\Sigma\mathbf{A}'$,

$$\text{Var}[\mathbf{A}\mathbf{x}] = \mathbf{A}\Sigma\mathbf{A}'. \quad (\text{B-87})$$

This matrix will be either nonnegative definite or positive definite, depending on the column rank of \mathbf{A} .

B.10.3 NONLINEAR FUNCTIONS

Consider a set of possibly nonlinear functions of \mathbf{x} , $\mathbf{y} = \mathbf{g}(\mathbf{x})$. Each element of \mathbf{y} can be approximated with a linear Taylor series. Let \mathbf{j}^i be the row vector of partial derivatives of the i th function with respect to the n elements of \mathbf{x} :

$$\mathbf{j}^i(\mathbf{x}) = \frac{\partial g_i(\mathbf{x})}{\partial \mathbf{x}'} = \frac{\partial y_i}{\partial \mathbf{x}'}. \quad (\text{B-88})$$

Then, proceeding in the now familiar way, we use $\boldsymbol{\mu}$, the mean vector of \mathbf{x} , as the expansion point, so that $\mathbf{j}^i(\boldsymbol{\mu})$ is the row vector of partial derivatives evaluated at $\boldsymbol{\mu}$. Then

$$g_i(\mathbf{x}) \approx g_i(\boldsymbol{\mu}) + \mathbf{j}^i(\boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu}). \quad (\text{B-89})$$

From this we obtain

$$E[g_i(\mathbf{x})] \approx g_i(\boldsymbol{\mu}), \quad (\text{B-90})$$

$$\text{Var}[g_i(\mathbf{x})] \approx \mathbf{j}^i(\boldsymbol{\mu})\Sigma\mathbf{j}^i(\boldsymbol{\mu})', \quad (\text{B-91})$$

and

$$\text{Cov}[g_i(\mathbf{x}), g_j(\mathbf{x})] \approx \mathbf{j}^i(\boldsymbol{\mu})\Sigma\mathbf{j}^j(\boldsymbol{\mu})'. \quad (\text{B-92})$$

These results can be collected in a convenient form by arranging the row vectors $\mathbf{j}^i(\boldsymbol{\mu})$ in a matrix $\mathbf{J}(\boldsymbol{\mu})$. Then, corresponding to the preceding equations, we have

$$E[\mathbf{g}(\mathbf{x})] \simeq \mathbf{g}(\boldsymbol{\mu}), \quad (\text{B-93})$$

$$\text{Var}[\mathbf{g}(\mathbf{x})] \simeq \mathbf{J}(\boldsymbol{\mu})\Sigma\mathbf{J}(\boldsymbol{\mu})'. \quad (\text{B-94})$$

The matrix $\mathbf{J}(\boldsymbol{\mu})$ in the last preceding line is $\partial\mathbf{y}/\partial\mathbf{x}'$ evaluated at $\mathbf{x} = \boldsymbol{\mu}$.

B.11 THE MULTIVARIATE NORMAL DISTRIBUTION

The foundation of most multivariate analysis in econometrics is the multivariate normal distribution. Let the vector $(x_1, x_2, \dots, x_n)' = \mathbf{x}$ be the set of n random variables, $\boldsymbol{\mu}$ their mean vector, and $\boldsymbol{\Sigma}$ their covariance matrix. The general form of the joint density is

$$f(\mathbf{x}) = (2\pi)^{-n/2} |\boldsymbol{\Sigma}|^{-1/2} e^{-(1/2)(\mathbf{x}-\boldsymbol{\mu})' \boldsymbol{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu})}. \quad (\text{B-95})$$

If \mathbf{R} is the correlation matrix of the variables and $\mathbf{R}_{ij} = \sigma_{ij}/(\sigma_i\sigma_j)$, then

$$f(\mathbf{x}) = (2\pi)^{-n/2} (\sigma_1\sigma_2 \dots \sigma_n)^{-1} |\mathbf{R}|^{-1/2} e^{-(1/2)\boldsymbol{\varepsilon}' \mathbf{R}^{-1} \boldsymbol{\varepsilon}}, \quad (\text{B-96})$$

where $\varepsilon_i = (x_i - \mu_i)/\sigma_i$.⁸

Two special cases are of interest. If all the variables are uncorrelated, then $\rho_{ij} = 0$ for $i \neq j$. Thus, $\mathbf{R} = \mathbf{I}$, and the density becomes

$$\begin{aligned} f(\mathbf{x}) &= (2\pi)^{-n/2} (\sigma_1\sigma_2 \dots \sigma_n)^{-1} e^{-\boldsymbol{\varepsilon}' \boldsymbol{\varepsilon}/2} \\ &= f(x_1) f(x_2) \dots f(x_n) = \prod_{i=1}^n f(x_i). \end{aligned} \quad (\text{B-97})$$

As in the bivariate case, if normally distributed variables are uncorrelated, then they are independent. If $\sigma_i = \sigma$ and $\boldsymbol{\mu} = \mathbf{0}$, then $x_i \sim N[0, \sigma^2]$ and $\varepsilon_i = x_i/\sigma$, and the density becomes

$$f(\mathbf{x}) = (2\pi)^{-n/2} (\sigma^2)^{-n/2} e^{-\mathbf{x}' \mathbf{x}/(2\sigma^2)} \quad (\text{B-98})$$

Finally, if $\sigma = 1$,

$$f(\mathbf{x}) = (2\pi)^{-n/2} e^{-\mathbf{x}' \mathbf{x}/2}. \quad (\text{B-99})$$

This distribution is the **multivariate standard normal**, or **spherical normal distribution**.

B.11.1 MARGINAL AND CONDITIONAL NORMAL DISTRIBUTIONS

Let \mathbf{x}_1 be any subset of the variables, including a single variable, and let \mathbf{x}_2 be the remaining variables. Partition $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ likewise so that

$$\boldsymbol{\mu} = \begin{bmatrix} \mu_1 \\ \mu_2 \end{bmatrix} \quad \text{and} \quad \boldsymbol{\Sigma} = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{bmatrix}.$$

Then the marginal distributions are also normal. In particular, we have the following theorem.

THEOREM B.7 Marginal and Conditional Normal Distributions

If $[\mathbf{x}_1, \mathbf{x}_2]$ have a joint multivariate normal distribution, then the marginal distributions are

$$\mathbf{x}_1 \sim N(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_{11}) \quad (\text{B-100})$$

⁸This result is obtained by constructing $\boldsymbol{\Delta}$, the diagonal matrix with σ_i as its i th diagonal element. Then, $\mathbf{R} = \boldsymbol{\Delta}^{-1} \boldsymbol{\Sigma} \boldsymbol{\Delta}^{-1}$, which implies that $\boldsymbol{\Sigma}^{-1} = \boldsymbol{\Delta}^{-1} \mathbf{R}^{-1} \boldsymbol{\Delta}^{-1}$. Inserting this in (B-95) yields (B-96). Note that the i th element of $\boldsymbol{\Delta}^{-1}(\mathbf{x} - \boldsymbol{\mu})$ is $(x_i - \mu_i)/\sigma_i$.

THEOREM B.7 (Continued)

and

$$\mathbf{x}_2 \sim N(\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_{22}). \quad (\text{B-101})$$

The conditional distribution of \mathbf{x}_1 given \mathbf{x}_2 is normal as well:

$$\mathbf{x}_1 | \mathbf{x}_2 \sim N(\boldsymbol{\mu}_{1.2}, \boldsymbol{\Sigma}_{11.2}) \quad (\text{B-102})$$

where

$$\boldsymbol{\mu}_{1.2} = \boldsymbol{\mu}_1 + \boldsymbol{\Sigma}_{12} \boldsymbol{\Sigma}_{22}^{-1} (\mathbf{x}_2 - \boldsymbol{\mu}_2) \quad (\text{B-102a})$$

$$\boldsymbol{\Sigma}_{11.2} = \boldsymbol{\Sigma}_{11} - \boldsymbol{\Sigma}_{12} \boldsymbol{\Sigma}_{22}^{-1} \boldsymbol{\Sigma}_{21}. \quad (\text{B-102b})$$

Proof: We partition $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ as shown above and insert the parts in (B-95). To construct the density, we use (A-72) to partition the determinant,

$$|\boldsymbol{\Sigma}| = |\boldsymbol{\Sigma}_{22}| |\boldsymbol{\Sigma}_{11} - \boldsymbol{\Sigma}_{12} \boldsymbol{\Sigma}_{22}^{-1} \boldsymbol{\Sigma}_{21}|,$$

and (A-74) to partition the inverse,

$$\begin{bmatrix} \boldsymbol{\Sigma}_{11} & \boldsymbol{\Sigma}_{12} \\ \boldsymbol{\Sigma}_{21} & \boldsymbol{\Sigma}_{22} \end{bmatrix}^{-1} = \begin{bmatrix} \boldsymbol{\Sigma}_{11.2}^{-1} & -\boldsymbol{\Sigma}_{11.2}^{-1} \mathbf{B} \\ -\mathbf{B}' \boldsymbol{\Sigma}_{11.2}^{-1} & \boldsymbol{\Sigma}_{22}^{-1} + \mathbf{B}' \boldsymbol{\Sigma}_{11.2}^{-1} \mathbf{B} \end{bmatrix}.$$

For simplicity, we let

$$\mathbf{B} = \boldsymbol{\Sigma}_{12} \boldsymbol{\Sigma}_{22}^{-1}.$$

Inserting these in (B-95) and collecting terms produces the joint density as a product of two terms:

$$f(\mathbf{x}_1, \mathbf{x}_2) = f_{1.2}(\mathbf{x}_1 | \mathbf{x}_2) f_2(\mathbf{x}_2).$$

The first of these is a normal distribution with mean $\boldsymbol{\mu}_{1.2}$ and variance $\boldsymbol{\Sigma}_{11.2}$, whereas the second is the marginal distribution of \mathbf{x}_2 .

The conditional mean vector in the multivariate normal distribution is a linear function of the unconditional mean and the conditioning variables, and the conditional covariance matrix is constant and is smaller (in the sense discussed in Section A.7.3) than the unconditional covariance matrix. Notice that the conditional covariance matrix is the inverse of the upper left block of $\boldsymbol{\Sigma}^{-1}$; that is, this matrix is of the form shown in (A-74) for the partitioned inverse of a matrix.

B.11.2 THE CLASSICAL NORMAL LINEAR REGRESSION MODEL

An important special case of the preceding is that in which \mathbf{x}_1 is a single variable y and \mathbf{x}_2 is K variables, \mathbf{x} . Then the conditional distribution is a multivariate version of that in (B-80) with $\boldsymbol{\beta} = \boldsymbol{\Sigma}_{\mathbf{xx}}^{-1} \boldsymbol{\sigma}_{\mathbf{xy}}$, where $\boldsymbol{\sigma}_{\mathbf{xy}}$ is the vector of covariances of y with \mathbf{x}_2 . Recall that any random variable, y , can be written as its mean plus the deviation from the mean. If we apply this tautology to the multivariate normal, we obtain

$$y = E[y | \mathbf{x}] + (y - E[y | \mathbf{x}]) = \alpha + \boldsymbol{\beta}' \mathbf{x} + \varepsilon$$

where β is given above, $\alpha = \mu_y - \beta' \mu_x$, and ε has a normal distribution. We thus have, in this multivariate normal distribution, the **classical normal linear regression model**.

B.11.3 LINEAR FUNCTIONS OF A NORMAL VECTOR

Any linear function of a vector of joint normally distributed variables is also normally distributed. The mean vector and covariance matrix of $\mathbf{Ax} + \mathbf{b}$, where \mathbf{x} is normally distributed, follow the general pattern given earlier. Thus,

$$\text{If } \mathbf{x} \sim N[\boldsymbol{\mu}, \boldsymbol{\Sigma}], \text{ then } \mathbf{Ax} + \mathbf{b} \sim N[\mathbf{A}\boldsymbol{\mu} + \mathbf{b}, \mathbf{A}\boldsymbol{\Sigma}\mathbf{A}']. \tag{B-103}$$

If \mathbf{A} does not have full rank, then $\mathbf{A}\boldsymbol{\Sigma}\mathbf{A}'$ is singular and the density does not exist in the full dimensional space of \mathbf{x} though it does exist in the subspace of dimension equal to the rank of $\boldsymbol{\Sigma}$. Nonetheless, the individual elements of $\mathbf{Ax} + \mathbf{b}$ will still be normally distributed, and the joint *distribution* of the full vector is still a multivariate normal.

B.11.4 QUADRATIC FORMS IN A STANDARD NORMAL VECTOR

The earlier discussion of the chi-squared distribution gives the distribution of $\mathbf{x}'\mathbf{x}$ if \mathbf{x} has a standard normal distribution. It follows from (A-36) that

$$\mathbf{x}'\mathbf{x} = \sum_{i=1}^n x_i^2 = \sum_{i=1}^n (x_i - \bar{x})^2 + n\bar{x}^2. \tag{B-104}$$

We know from (B-32) that $\mathbf{x}'\mathbf{x}$ has a chi-squared distribution. It seems natural, therefore, to invoke (B-34) for the two parts on the right-hand side of (B-104). It is not yet obvious, however, that either of the two terms has a chi-squared distribution or that the two terms are independent, as required. To show these conditions, it is necessary to derive the distributions of **idempotent quadratic forms** and to show when they are independent.

To begin, the second term is the square of $\sqrt{n}\bar{x}$, which can easily be shown to have a standard normal distribution. Thus, the second term is the square of a standard normal variable and has chi-squared distribution with one degree of freedom. But the first term is the sum of n nonindependent variables, and it remains to be shown that the two terms are independent.

DEFINITION B.3 Orthonormal Quadratic Form
A particular case of (B-103) is the following:

If $\mathbf{x} \sim N[\mathbf{0}, \mathbf{I}]$ and \mathbf{C} is a square matrix such that $\mathbf{C}'\mathbf{C} = \mathbf{I}$, then $\mathbf{C}'\mathbf{x} \sim N[\mathbf{0}, \mathbf{I}]$.

Consider, then, a quadratic form in a standard normal vector \mathbf{x} with symmetric matrix \mathbf{A} :

$$q = \mathbf{x}'\mathbf{Ax}. \tag{B-105}$$

Let the characteristic roots and vectors of \mathbf{A} be arranged in a diagonal matrix $\boldsymbol{\Lambda}$ and an orthogonal matrix \mathbf{C} , as in Section A.6.3. Then

$$q = \mathbf{x}'\mathbf{C}\boldsymbol{\Lambda}\mathbf{C}'\mathbf{x}. \tag{B-106}$$

By definition, \mathbf{C} satisfies the requirement that $\mathbf{C}'\mathbf{C} = \mathbf{I}$. Thus, the vector $\mathbf{y} = \mathbf{C}'\mathbf{x}$ has a standard

normal distribution. Consequently,

$$q = \mathbf{y}' \Lambda \mathbf{y} = \sum_{i=1}^n \lambda_i y_i^2. \tag{B-107}$$

If λ_i is always one or zero, then

$$q = \sum_{j=1}^J y_j^2, \tag{B-108}$$

which has a chi-squared distribution. The sum is taken over the $j = 1, \dots, J$ elements associated with the roots that are equal to one. A matrix whose characteristic roots are all zero or one is idempotent. Therefore, we have proved the next theorem.

THEOREM B.8 Distribution of an Idempotent Quadratic Form in a Standard Normal Vector

If $\mathbf{x} \sim N[\mathbf{0}, \mathbf{I}]$ and \mathbf{A} is idempotent, then $\mathbf{x}' \mathbf{A} \mathbf{x}$ has a chi-squared distribution with degrees of freedom equal to the number of unit roots of \mathbf{A} , which is equal to the rank of \mathbf{A} .

The rank of a matrix is equal to the number of nonzero characteristic roots it has. Therefore, the degrees of freedom in the preceding chi-squared distribution equals J , the rank of \mathbf{A} .

We can apply this result to the earlier sum of squares. The first term is

$$\sum_{i=1}^n (x_i - \bar{x})^2 = \mathbf{x}' \mathbf{M}^0 \mathbf{x},$$

where \mathbf{M}^0 was defined in (A-34) as the matrix that transforms data to mean deviation form:

$$\mathbf{M}^0 = \mathbf{I} - \frac{1}{n} \mathbf{i} \mathbf{i}'.$$

Since \mathbf{M}^0 is idempotent, the sum of squared deviations from the mean has a chi-squared distribution. The degrees of freedom equals the rank \mathbf{M}^0 , which is not obvious except for the useful result in (A-108), that

- The rank of an idempotent matrix is equal to its trace. (B-109)

Each diagonal element of \mathbf{M}^0 is $1 - (1/n)$; hence, the trace is $n[1 - (1/n)] = n - 1$. Therefore, we have an application of Theorem B.8.

- If $\mathbf{x} \sim N(\mathbf{0}, \mathbf{I})$, $\sum_{i=1}^n (x_i - \bar{x})^2 \sim \chi^2[n - 1]$. (B-110)

We have already shown that the second term in (B-104) has a chi-squared distribution with one degree of freedom. It is instructive to set this up as a quadratic form as well:

$$n\bar{x}^2 = \mathbf{x}' \left[\frac{1}{n} \mathbf{i} \mathbf{i}' \right] \mathbf{x} = \mathbf{x}' [\mathbf{j} \mathbf{j}'] \mathbf{x}, \quad \text{where } \mathbf{j} = \left(\frac{1}{\sqrt{n}} \right) \mathbf{i}. \tag{B-111}$$

The matrix in brackets is the outer product of a nonzero vector, which always has rank one. You can verify that it is idempotent by multiplication. Thus, $\mathbf{x}' \mathbf{x}$ is the sum of two chi-squared variables.

one with $n - 1$ degrees of freedom and the other with one. It is now necessary to show that the two terms are independent. To do so, we will use the next theorem.

THEOREM B.9 Independence of Idempotent Quadratic Forms

If $\mathbf{x} \sim N[\mathbf{0}, \mathbf{I}]$ and $\mathbf{x}'\mathbf{A}\mathbf{x}$ and $\mathbf{x}'\mathbf{B}\mathbf{x}$ are two idempotent quadratic forms in \mathbf{x} , then $\mathbf{x}'\mathbf{A}\mathbf{x}$ and $\mathbf{x}'\mathbf{B}\mathbf{x}$ are independent if $\mathbf{A}\mathbf{B} = \mathbf{0}$. (B-112)

As before, we show the result for the general case and then specialize it for the example. Since both \mathbf{A} and \mathbf{B} are symmetric and idempotent, $\mathbf{A} = \mathbf{A}'\mathbf{A}$ and $\mathbf{B} = \mathbf{B}'\mathbf{B}$. The quadratic forms are therefore

$$\mathbf{x}'\mathbf{A}\mathbf{x} = \mathbf{x}'\mathbf{A}'\mathbf{A}\mathbf{x} = \mathbf{x}'_1\mathbf{x}_1, \quad \text{where } \mathbf{x}_1 = \mathbf{A}\mathbf{x}, \quad \text{and} \quad \mathbf{x}'\mathbf{B}\mathbf{x} = \mathbf{x}'_2\mathbf{x}_2, \quad \text{where } \mathbf{x}_2 = \mathbf{B}\mathbf{x}. \quad (\text{B-113})$$

Both vectors have zero mean vectors, so the covariance matrix of \mathbf{x}_1 and \mathbf{x}_2 is

$$E(\mathbf{x}_1\mathbf{x}'_2) = \mathbf{A}\mathbf{B}' = \mathbf{A}\mathbf{B} = \mathbf{0}.$$

Since $\mathbf{A}\mathbf{x}$ and $\mathbf{B}\mathbf{x}$ are linear functions of a normally distributed random vector, they are, in turn, normally distributed. Their zero covariance matrix implies that they are statistically independent,⁹ which establishes the independence of the two quadratic forms. For the case of $\mathbf{x}'\mathbf{x}$, the two matrices are \mathbf{M}^0 and $[\mathbf{I} - \mathbf{M}^0]$. You can show that $\mathbf{M}^0[\mathbf{I} - \mathbf{M}^0] = \mathbf{0}$ just by multiplying it out.

B.11.5 THE F DISTRIBUTION

The normal family of distributions (chi-squared, F , and t) can all be derived as functions of idempotent quadratic forms in a standard normal vector. The F distribution is the ratio of two independent chi-squared variables, each divided by its respective degrees of freedom. Let \mathbf{A} and \mathbf{B} be two idempotent matrices with ranks r_a and r_b , and let $\mathbf{A}\mathbf{B} = \mathbf{0}$. Then

$$\frac{\mathbf{x}'\mathbf{A}\mathbf{x}/r_a}{\mathbf{x}'\mathbf{B}\mathbf{x}/r_b} \sim F[r_a, r_b]. \quad (\text{B-114})$$

If $\text{Var}[\mathbf{x}] = \sigma^2\mathbf{I}$ instead, then this is modified to

$$\frac{(\mathbf{x}'\mathbf{A}\mathbf{x}/\sigma^2)/r_a}{(\mathbf{x}'\mathbf{B}\mathbf{x}/\sigma^2)/r_b} \sim F[r_a, r_b]. \quad (\text{B-115})$$

B.11.6 A FULL RANK QUADRATIC FORM

Finally, consider the general case,

$$\mathbf{x} \sim N[\boldsymbol{\mu}, \boldsymbol{\Sigma}].$$

We are interested in the distribution of

$$q = (\mathbf{x} - \boldsymbol{\mu})'\boldsymbol{\Sigma}^{-1}(\mathbf{x} - \boldsymbol{\mu}). \quad (\text{B-116})$$

⁹Note that both $\mathbf{x}_1 = \mathbf{A}\mathbf{x}$ and $\mathbf{x}_2 = \mathbf{B}\mathbf{x}$ have singular covariance matrices. Nonetheless, every element of \mathbf{x}_1 is independent of every element \mathbf{x}_2 , so the vectors are independent.

876 APPENDIX B ♦ Probability and Distribution Theory

First, the vector can be written as $\mathbf{z} = \mathbf{x} - \boldsymbol{\mu}$, and $\boldsymbol{\Sigma}$ is the covariance matrix of \mathbf{z} as well as of \mathbf{x} . Therefore, we seek the distribution of

$$q = \mathbf{z}'\boldsymbol{\Sigma}^{-1}\mathbf{z} = \mathbf{z}'(\text{Var}[\mathbf{z}])^{-1}\mathbf{z}, \quad (\text{B-117})$$

where \mathbf{z} is normally distributed with mean $\mathbf{0}$. This equation is a quadratic form, but not necessarily in an idempotent matrix.¹⁰ Since $\boldsymbol{\Sigma}$ is positive definite, it has a square root. Define the symmetric matrix $\boldsymbol{\Sigma}^{1/2}$ so that $\boldsymbol{\Sigma}^{1/2}\boldsymbol{\Sigma}^{1/2} = \boldsymbol{\Sigma}$. Then

$$\boldsymbol{\Sigma}^{-1} = \boldsymbol{\Sigma}^{-1/2}\boldsymbol{\Sigma}^{-1/2}$$

and

$$\begin{aligned} \mathbf{z}'\boldsymbol{\Sigma}^{-1}\mathbf{z} &= \mathbf{z}'\boldsymbol{\Sigma}^{-1/2}\boldsymbol{\Sigma}^{-1/2}\mathbf{z} \\ &= (\boldsymbol{\Sigma}^{-1/2}\mathbf{z})'(\boldsymbol{\Sigma}^{-1/2}\mathbf{z}) \\ &= \mathbf{w}'\mathbf{w}. \end{aligned}$$

Now $\mathbf{w} = \mathbf{A}\mathbf{z}$, so

$$E(\mathbf{w}) = \mathbf{A}E[\mathbf{z}] = \mathbf{0}$$

and

$$\text{Var}[\mathbf{w}] = \mathbf{A}\boldsymbol{\Sigma}\mathbf{A}' = \boldsymbol{\Sigma}^{-1/2}\boldsymbol{\Sigma}\boldsymbol{\Sigma}^{-1/2} = \boldsymbol{\Sigma}^0 = \mathbf{I}.$$

This provides the following important result:

THEOREM B.10 **Distribution of a Standardized Normal Vector**
If $\mathbf{x} \sim N[\boldsymbol{\mu}, \boldsymbol{\Sigma}]$, then $\boldsymbol{\Sigma}^{-1/2}(\mathbf{x} - \boldsymbol{\mu}) \sim N[\mathbf{0}, \mathbf{I}]$.

The simplest special case is that in which \mathbf{x} has only one variable, so that the transformation is just $(x - \mu)/\sigma$. Combining this case with (B-32) concerning the sum of squares of standard normals, we have the following theorem.

THEOREM B.11 **Distribution of $\mathbf{x}'\boldsymbol{\Sigma}^{-1}\mathbf{x}$ When \mathbf{x} Is Normal**
If $\mathbf{x} \sim N[\boldsymbol{\mu}, \boldsymbol{\Sigma}]$, then $(\mathbf{x} - \boldsymbol{\mu})'\boldsymbol{\Sigma}^{-1}(\mathbf{x} - \boldsymbol{\mu}) \sim \chi^2[n]$.

B.11.7 INDEPENDENCE OF A LINEAR AND A QUADRATIC FORM

The t distribution is used in many forms of hypothesis tests. In some situations, it arises as the ratio of a linear to a quadratic form in a normal vector. To establish the distribution of these statistics, we use the following result.

¹⁰It will be idempotent only in the special case of $\boldsymbol{\Sigma} = \mathbf{I}$.

THEOREM B.12 Independence of a Linear and a Quadratic Form

A linear function \mathbf{Lx} and a symmetric idempotent quadratic form $\mathbf{x}'\mathbf{A}\mathbf{x}$ in a standard normal vector are statistically independent if $\mathbf{LA} = \mathbf{0}$.

The proof follows the same logic as that for two quadratic forms. Write $\mathbf{x}'\mathbf{A}\mathbf{x}$ as $\mathbf{x}'\mathbf{A}'\mathbf{A}\mathbf{x} = (\mathbf{Ax})'(\mathbf{Ax})$. The covariance matrix of the variables \mathbf{Lx} and \mathbf{Ax} is $\mathbf{LA} = \mathbf{0}$, which establishes the independence of these two random vectors. The independence of the linear function and the quadratic form follows since functions of independent random vectors are also independent.

The t distribution is defined as the ratio of a standard normal variable to the square root of a chi-squared variable divided by its degrees of freedom:

$$t[J] = \frac{N[0, 1]}{\{\chi^2[J]/J\}^{1/2}}$$

A particular case is

$$t[n - 1] = \frac{\sqrt{n}\bar{x}}{\left\{\left[\frac{1}{n-1}\right] \sum_{i=1}^n (x_i - \bar{x})^2\right\}^{1/2}} = \frac{\sqrt{n}\bar{x}}{s}$$

where s is the standard deviation of the values of \mathbf{x} . The distribution of the two variables in $t[n - 1]$ was shown earlier; we need only show that they are independent. But

$$\sqrt{n}\bar{x} = \frac{1}{\sqrt{n}}\mathbf{i}'\mathbf{x} = \mathbf{j}'\mathbf{x}$$

and

$$s^2 = \frac{\mathbf{x}'\mathbf{M}^0\mathbf{x}}{n - 1}$$

It suffices to show that $\mathbf{M}^0\mathbf{j} = \mathbf{0}$, which follows from

$$\mathbf{M}^0\mathbf{i} = [\mathbf{I} - \mathbf{i}(\mathbf{i}'\mathbf{i})^{-1}\mathbf{i}']\mathbf{i} = \mathbf{i} - \mathbf{i}(\mathbf{i}'\mathbf{i})^{-1}(\mathbf{i}'\mathbf{i}) = \mathbf{0}$$

APPENDIX C



ESTIMATION AND INFERENCE

C.1 INTRODUCTION

The probability distributions discussed in Appendix B serve as models for the underlying data generating processes that produce our observed data. The goal of statistical inference in econometrics is to use the principles of mathematical statistics to combine these theoretical distributions and the observed data into an empirical model of the economy. This analysis takes place in one of two frameworks, classical or Bayesian. The overwhelming majority of empirical study in econometrics

has been done in the classical framework. Our focus, therefore, will be on classical methods of inference. Bayesian methods are discussed in Chapter 16.¹

C.2 SAMPLES AND RANDOM SAMPLING

The classical theory of statistical inference centers on rules for using the sampled data effectively. These rules, in turn, are based on the properties of samples and sampling distributions.

A sample of n observations on one or more variables, denoted $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n$ is a **random sample** if the n observations are drawn independently from the same population, or probability distribution, $f(\mathbf{x}_i, \theta)$. The sample may be univariate if \mathbf{x}_i is a single random variable or multivariate if each observation contains several variables. A random sample of observations, denoted $[\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n]$ or $\{\mathbf{x}_i\}_{i=1, \dots, n}$, is said to be **independent, identically distributed**, which we denote *i.i.d.* The vector θ contains one or more unknown parameters. Data are generally drawn in one of two settings. A **cross section** is a sample of a number of observational units all drawn at the same point in time. A **time series** is a set of observations drawn on the same observational unit at a number of (usually evenly spaced) points in time. Many recent studies have been based on time series cross sections, which generally consist of the same cross sectional units observed at several points in time. Since the typical data set of this sort consists of a large number of cross-sectional units observed at a few points in time, the common term **panel data set** is usually more fitting for this sort of study.

C.3 DESCRIPTIVE STATISTICS

Before attempting to estimate parameters of a population or fit models to data, we normally examine the data themselves. In raw form, the sample data are a disorganized mass of information, so we will need some organizing principles to distill the information into something meaningful. Consider, first, examining the data on a single variable. In most cases, and particularly if the number of observations in the sample is large, we shall use some summary **statistics** to describe the sample data. Of most interest are measures of **location**—that is, the center of the data—and **scale**, or the dispersion of the data. A few measures of central tendency are as follows:

$$\text{mean: } \bar{x} = \frac{1}{n} \sum_{i=1}^n x_i,$$

$$\text{median: } M = \text{middle ranked observation}, \quad (\text{C-1})$$

$$\text{sample midrange: } \text{midrange} = \frac{\text{maximum} - \text{minimum}}{2}.$$

The dispersion of the sample observations is usually measured by the

$$\text{standard deviation: } s_x = \left[\frac{\sum_{i=1}^n (x_i - \bar{x})^2}{n-1} \right]^{1/2}. \quad (\text{C-2})$$

Other measures, such as the average absolute deviation from the sample mean, are also used, although less frequently than the standard deviation. The shape of the distribution of values is

¹An excellent reference is Leamer (1978). A summary of the results as they apply to econometrics is contained in Zellner (1971) and in Judge et al. (1985). See, as well, Poirier (1991). A recent textbook with a heavy Bayesian emphasis is Poirier (1995).

often of interest as well. Samples of income or expenditure data, for example, tend to be highly skewed while financial data such as asset returns and exchange rate movements are relatively more symmetrically distributed but are also more widely dispersed than other variables that might be observed. Two measures used to quantify these effects are the

$$\text{skewness} = \left[\frac{\sum_{i=1}^n (x_i - \bar{x})^3}{s_x^3 (n-1)} \right] \quad \text{and} \quad \text{kurtosis} = \left[\frac{\sum_{i=1}^n (x_i - \bar{x})^4}{s_x^4 (n-1)} \right]$$

(Benchmark values for these two measures are zero for a symmetric distribution, and three for one which is “normally” dispersed. The skewness coefficient has a bit less of the intuitive appeal of the mean and standard deviation, and the kurtosis measure has very little at all. The box and whisker plot is a graphical device which is often used to capture a large amount of information about the sample in a simple visual display. This plot shows in a figure the median, the range of values contained in the 25th and 75th percentile, some limits that show the normal range of values expected, such as the median plus and minus two standard deviations, and in isolation values that could be viewed as outliers. A box and whisker plot is shown in Figure C.1 for the income variable in Example C.1.

If the sample contains data on more than one variable, we will also be interested in measures of association among the variables. A **scatter diagram** is useful in a bivariate sample if the sample contains a reasonable number of observations. Figure C.1 shows an example for a small data set. If the sample is a multivariate one, then the degree of linear association among the variables can be measured by the pairwise measures

$$\text{covariance: } s_{xy} = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{n-1}, \tag{C-3}$$

$$\text{correlation: } r_{xy} = \frac{s_{xy}}{s_x s_y}.$$

If the sample contains data on several variables, then it is sometimes convenient to arrange the covariances or correlations in a

$$\text{covariance matrix: } \mathbf{S} = [s_{ij}] \tag{C-4}$$

or

$$\text{correlation matrix: } \mathbf{R} = [r_{ij}].$$

Some useful algebraic results for any two variables $(x_i, y_i), i = 1, \dots, n$, and constants a and b are

$$s_x^2 = \frac{\sum_{i=1}^n x_i^2 - n\bar{x}^2}{n-1}, \tag{C-5}$$

$$s_{xy} = \frac{\sum_{i=1}^n x_i y_i - n\bar{x}\bar{y}}{n-1}, \tag{C-6}$$

$$-1 \leq r_{xy} \leq 1,$$

$$r_{ax,by} = \frac{ab}{|ab|} r_{xy}, \quad a, b \neq 0, \tag{C-7}$$

$$s_{ax} = |a|s_x, \tag{C-8}$$

$$s_{ax,by} = abs_{xy}.$$

Note that these algebraic results parallel the theoretical results for bivariate probability distributions. (We note in passing, while the formulas in (C-2) and (C-5) are algebraically the same, (C-2) will generally be more accurate in practice, especially when the values in the sample are very widely dispersed.)

Example C.1 Descriptive Statistics for a Random Sample

Appendix Table FD1 contains a (hypothetical) sample of observations on income and education (The observations all appear in the calculations of the means below.) A scatter diagram appears in Figure C.1. It suggests a weak positive association between income and education in these data. The box and whisker plot for income at the left of the scatter plot shows the distribution of the income data as well.

$$\text{Means: } \bar{I} = \frac{1}{20} \left[\begin{array}{l} 20.5 + 31.5 + 47.7 + 26.2 + 44.0 + 8.28 + 30.8 + \\ 17.2 + 19.9 + 9.96 + 55.8 + 25.2 + 29.0 + 85.5 + \\ 15.1 + 28.5 + 21.4 + 17.7 + 6.42 + 84.9 \end{array} \right] = 31.278$$

$$\bar{E} = \frac{1}{20} \left[\begin{array}{l} 12 + 16 + 18 + 16 + 12 + 12 + 16 + 12 + 10 + 12 + \\ 16 + 20 + 12 + 16 + 10 + 18 + 16 + 20 + 12 + 16 \end{array} \right] = 14.600$$

Standard deviations:

$$s_I = \sqrt{\frac{1}{19} [(20.5 - 31.278)^2 + \dots + (84.9 - 31.278)^2]} = 22.376,$$

$$s_E = \sqrt{\frac{1}{19} [(12 - 14.6)^2 + \dots + (16 - 14.6)^2]} = 3.119.$$




FIGURE C.1 Box and Whisker Plot for Income and Scatter Diagram for Income and Education.

Covariance: $s_{IE} = \frac{1}{19}[20.5(12) + \dots + 84.9(16) - 20(31.28)(14.6)] = 23.597$.

Correlation: $r_{IE} = \frac{23.597}{(22.376)(3.119)} = 0.3382$.

The positive correlation is consistent with our observation in the scatter diagram.

The statistics described above will provide the analyst with a more concise description of the data than a raw tabulation. However, we have not, as yet, suggested that these measures correspond to some underlying characteristic of the process that generated the data. We do assume that there is an underlying mechanism, the data generating process, that produces the data in hand. Thus, these serve to do more than describe the data; they characterize that process, or population. Since we have assumed that there is an underlying probability distribution, it might be useful to produce a statistic that gives a broader view of the DGP. The **histogram** is a simple graphical device that produces this result—see Examples C.3 and C.4 for applications. For small samples or widely dispersed data, however, histograms tend to be rough and difficult to make informative. A burgeoning literature [see, e.g., Pagan and Ullah (1999)] has demonstrated the usefulness of the **kernel density estimator** as a substitute for the histogram as a descriptive tool for the underlying distribution that produced a sample of data. The underlying theory of the kernel density estimator is fairly complicated, but the computations are surprisingly simple. The estimator is computed using

$$\hat{f}(x^*) = \frac{1}{nh} \sum_{i=1}^n K \left[\frac{x_i - x^*}{h} \right]$$

where x_1, \dots, x_n are the n observations in the sample, $\hat{f}(x^*)$ denotes the estimated density function x^* is the value at which we wish to evaluate the density, and h and $K[\cdot]$ are the “bandwidth” and “kernel function” which we now consider. The density estimator is rather like a histogram, in which the bandwidth is the width of the intervals. The kernel function is a weight function which is generally chosen so that it takes large values when x^* is close to x_i and tapers off to zero in as they diverge in either direction. The weighting function used in the example below is the logistic density discussed in Section B.4.7. The bandwidth is chosen to be a function of $1/n$ so that the intervals can become narrower as the sample becomes larger (and richer). The one used below is $h = .9\text{Min}(s, \text{range}/3)/n^2$. (We will revisit this method of estimation in Chapter 16.) Example C.2 below illustrates the computation for the income data used in Example C.1.

Example C.2 Kernel Density Estimator for the Income Data

The following Figure C.2 suggests the large skew in the income data that is also suggested by the box and whisker plot (and the scatter plot) in Example C.1.

C.4 STATISTICS AS ESTIMATORS—SAMPLING DISTRIBUTIONS

The measures described in the preceding section summarize the data in a random sample. Each measure has a counterpart in the population, that is, the distribution from which the data were drawn. Sample quantities such as the means and the correlation coefficient correspond to population expectations, whereas the kernel density estimator and the values in Table C.1 parallel the population **pdf** and **cdf**. In the setting of a random sample, we expect these quantities to mimic




FIGURE C.2 Kernel Density Estimator.

TABLE C.1 Income Distribution

<i>Range</i>	<i>Relative Frequency</i>	<i>Cumulative Frequency</i>
<\$10,000	0.15	0.15
10,000–25,000	0.30	0.45
25,000–50,000	0.40	0.85
>50,000	0.15	1.00

the population, although not perfectly. The precise manner in which these quantities reflect the population values defines the sampling distribution of a sample statistic.

DEFINITION C.1 **Statistic**
A statistic is any function computed from the data in a sample.

If another sample were drawn under identical conditions, different values would be obtained for the observations, as each one is a random variable. Any statistic is a function of these random values, so it is also a random variable with a probability distribution called a **sampling distribution**. For example, the following shows an exact result for the sampling behavior of a widely used statistic.

THEOREM C.1 Sampling Distribution of the Sample Mean

If x_1, \dots, x_n are a random sample from a population with mean μ and variance σ^2 , then \bar{x} is a random variable with mean μ and variance σ^2/n .

Proof: $\bar{x} = (1/n)\sum_i x_i$. $E[\bar{x}] = (1/n)\sum_i \mu = \mu$. The observations are independent, so $\text{Var}[\bar{x}] = (1/n)^2 \text{Var}[\sum_i x_i] = (1/n^2)\sum_i \sigma^2 = \sigma^2/n$.

Example C.3 illustrates the behavior of the sample mean in samples of four observations drawn from a chi squared population with one degree of freedom. The crucial concepts illustrated in this example are, first, the mean and variance results in Theorem C.1 and, second, the phenomenon of **sampling variability**.

Notice that the fundamental result in Theorem C.1 does not assume a distribution for x_i . Indeed, looking back at Section C.3, nothing we have done so far has required any assumption about a particular distribution.

Example C.3 Sampling Distribution of a Sample Mean

Figure C.3 shows a frequency plot of the means of 1,000 random samples of four observations drawn from a chi-squared distribution with one degree of freedom, which has mean 1 and variance 2.

We are often interested in how a statistic behaves as the sample size increases. Example C.4 illustrates one such case. Figure C.4 shows two sampling distributions, one based on samples of three and a second, of the same statistic, but based on samples of six. The effect of increasing sample size in this figure is unmistakable. It is easy to visualize the behavior of this statistic if we extrapolate the experiment in Example C.4 to samples of, say, 100.

Example C.4 Sampling Distribution of the Sample Minimum

If x_1, \dots, x_n are a random sample from an exponential distribution with $f(x) = \theta e^{-\theta x}$, then the sampling distribution of the sample minimum in a sample of n observations, denoted $x_{(1)}$, is

$$f(x_{(1)}) = (n\theta)e^{-(n\theta)x_{(1)}}$$

Since $E[x] = 1/\theta$ and $\text{Var}[x] = 1/\theta^2$, by analogy $E[x_{(1)}] = 1/(n\theta)$ and $\text{Var}[x_{(1)}] = 1/(n\theta)^2$. Thus, in increasingly larger samples, the minimum will be arbitrarily close to 0. [The Chebychev inequality in Theorem D.2 can be used to prove this intuitively appealing result.]

Figure C.4 shows the results of a simple sampling experiment you can do to demonstrate this effect. It requires software that will allow you to produce pseudorandom numbers uniformly distributed in the range zero to one and that will let you plot a histogram and control the axes. (We used *EA/LimDep*. This can be done with *Stata*, *Excel*, or several other packages.) The experiment consists of drawing 1,000 sets of nine random values, U_{ij} , $i = 1, \dots, 1,000$, $j = 1, \dots, 9$. To transform these uniform draws to exponential with parameter θ —we used $\theta = 1.5$, use the inverse probability transform—see Section E.2.3. For an exponentially distributed variable, the transformation is $z_{ij} = -(1/\theta) \log(1 - U_{ij})$. We then created $z_{(1)}|3$ from the first three draws and $z_{(1)}|6$ from the other six. The two histograms show clearly the effect on the sampling distribution of increasing sample size from just 3 to 6.

Sampling distributions are used to make inferences about the population. To consider a perhaps obvious example, because the sampling distribution of the mean of a set of normally distributed observations has mean μ , the sample mean is a natural candidate for an estimate of μ . The observation that the sample “mimics” the population is a statement about the sampling




FIGURE C.3 Sampling Distribution of Means of 1,000 Samples of Size 4 from Chi-Squared [1].




FIGURE C.4 Histograms of the Sample Minimum of 3 and 6 Observations.

distributions of the sample statistics. Consider, for example, the sample data collected in Figure C.3. The sample mean of four observations clearly has a sampling distribution, which appears to have a mean roughly equal to the population mean. Our theory of parameter estimation departs from this point.

C.5 POINT ESTIMATION OF PARAMETERS

Our objective is to use the sample data to infer the value of a parameter or set of parameters, which we denote θ . A **point estimate** is a statistic computed from a sample that gives a single value for θ . The **standard error** of the estimate is the standard deviation of the sampling distribution of the statistic; the square of this quantity is the **sampling variance**. An **interval estimate** is a range of values that will contain the true parameter with a preassigned probability. There will be a connection between the two types of estimates; generally, if $\hat{\theta}$ is the point estimate, then the interval estimate will be $\hat{\theta} \pm$ a measure of sampling error.

An **estimator** is a rule or strategy for using the data to estimate the parameter. It is defined before the data are drawn. Obviously, some estimators are better than others. To take a simple example, your intuition should convince you that the sample mean would be a better estimator of the population mean than the sample minimum; the minimum is almost certain to underestimate the mean. Nonetheless, the minimum is not entirely without virtue; it is easy to compute, which is occasionally a relevant criterion. The search for good estimators constitutes much of econometrics. Estimators are compared on the basis of a variety of attributes. **Finite sample properties** of estimators are those attributes that can be compared regardless of the sample size. Some estimation problems involve characteristics that are not known in finite samples. In these instances, estimators are compared on the basis on their large sample, or **asymptotic properties**. We consider these in turn.

C.5.1 ESTIMATION IN A FINITE SAMPLE

The following are some finite sample estimation criteria for estimating a single parameter. The extensions to the multiparameter case are direct. We shall consider them in passing where necessary.

DEFINITION C.2 Unbiased Estimator

An estimator of a parameter θ is **unbiased** if the mean of its sampling distribution is θ . Formally,

$$E[\hat{\theta}] = \theta$$

or

$$E[\hat{\theta} - \theta] = \text{Bias}[\hat{\theta} | \theta] = 0$$

implies that $\hat{\theta}$ is unbiased. Note that this implies that the expected sampling error is zero. If θ is a vector of parameters, then the estimator is unbiased if the expected value of every element of $\hat{\theta}$ equals the corresponding element of θ .

If samples of size n are drawn repeatedly and $\hat{\theta}$ is computed for each one, then the average value of these estimates will tend to equal θ . For example, the average of the 1,000 sample means underlying Figure C.2 is 0.9038, which is reasonably close to the population mean of one. The sample minimum is clearly a biased estimator of the mean; it will almost always underestimate the mean, so it will do so on average as well.

Unbiasedness is a desirable attribute, but it is rarely used by itself as an estimation criterion. One reason is that there are many unbiased estimators that are poor uses of the data. For example, in a sample of size n , the first observation drawn is an unbiased estimator of the mean that clearly wastes a great deal of information. A second criterion used to choose among unbiased estimators is efficiency.

DEFINITION C.3 Efficient Unbiased Estimator

An unbiased estimator $\hat{\theta}_1$ is more **efficient** than another unbiased estimator $\hat{\theta}_2$ if the sampling variance of $\hat{\theta}_1$ is less than that of $\hat{\theta}_2$. That is,

$$\text{Var}[\hat{\theta}_1] < \text{Var}[\hat{\theta}_2].$$

In the multiparameter case, the comparison is based on the covariance matrices of the two estimators; $\hat{\theta}_1$ is more efficient than $\hat{\theta}_2$ if $\text{Var}[\hat{\theta}_2] - \text{Var}[\hat{\theta}_1]$ is a positive definite matrix.

By this criterion, the sample mean is obviously to be preferred to the first observation as an estimator of the population mean. If σ^2 is the population variance, then

$$\text{Var}[x_1] = \sigma^2 > \text{Var}[\bar{x}] = \frac{\sigma^2}{n}.$$

In discussing efficiency, we have restricted the discussion to unbiased estimators. Clearly, there are biased estimators that have smaller variances than the unbiased ones we have considered. Any constant has a variance of zero. Of course, using a constant as an estimator is not likely to be an effective use of the sample data. Focusing on unbiasedness may still preclude a tolerably biased estimator with a much smaller variance, however. A criterion that recognizes this possible tradeoff is the mean-squared error.

DEFINITION C.4 Mean-Squared Error

The mean-squared error of an estimator is

$$\begin{aligned} \text{MSE}[\hat{\theta} | \theta] &= E[(\hat{\theta} - \theta)^2] \\ &= \text{Var}[\hat{\theta}] + (\text{Bias}[\hat{\theta} | \theta])^2 && \text{if } \theta \text{ is a scalar,} && \text{(C-9)} \\ \text{MSE}[\hat{\theta} | \theta] &= \text{Var}[\hat{\theta}] + \text{Bias}[\hat{\theta} | \theta]\text{Bias}[\hat{\theta} | \theta]' && \text{if } \theta \text{ is a vector.} \end{aligned}$$

Figure C.5 illustrates the effect. On average, the biased estimator will be closer to the true parameter than will the unbiased estimator.

Which of these criteria should be used in a given situation depends on the particulars of that setting and our objectives in the study. Unfortunately, the MSE criterion is rarely operational; minimum mean-squared error estimators, when they exist at all, usually depend on unknown parameters. Thus, we are usually less demanding. A commonly used criterion is **minimum variance unbiasedness**.

Example C.5 Mean-Squared Error of the Sample Variance


In sampling from a normal distribution, the most frequently used estimator for σ^2 is

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

It is straightforward to show that s^2 is unbiased, so

$$\text{Var}[s^2] = \frac{2\sigma^4}{n - 1} = \text{MSE}[s^2 | \sigma^2].$$

FIGURE C.5 Sampling Distributions.



[A proof is based on the distribution of the idempotent quadratic form $(\mathbf{x} - \mathbf{i}\mu)' \mathbf{M}^0 (\mathbf{x} - \mathbf{i}\mu)$, which we discussed in Section B11.4.] A less frequently used estimator is

$$\hat{\sigma}^2 = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2 = [(n-1)/n]s^2.$$

This estimator is slightly biased downward:

$$E[\hat{\sigma}^2] = \frac{(n-1)E[s^2]}{n} = \frac{(n-1)\sigma^2}{n},$$

so its bias is

$$E[\hat{\sigma}^2 - \sigma^2] = \text{Bias}[\hat{\sigma}^2 | \sigma^2] = \frac{-1}{n}\sigma^2.$$

But it has a smaller variance than s^2 :

$$\text{Var}[\hat{\sigma}^2] = \left[\frac{n-1}{n} \right]^2 \left[\frac{2\sigma^4}{n-1} \right] < \text{Var}[s^2].$$

To compare the two estimators, we can use the difference in their mean-squared errors:

$$\text{MSE}[\hat{\sigma}^2 | \sigma^2] - \text{MSE}[s^2 | \sigma^2] = \sigma^4 \left[\frac{2n-1}{n^2} - \frac{2}{n-1} \right] < 0.$$

The biased estimator is a bit more precise. The difference will be negligible in a large sample, but, for example, it is about 1.2 percent in a sample of 16.

C.5.2 EFFICIENT UNBIASED ESTIMATION

In a random sample of n observations, the density of each observation is $f(x_i, \theta)$. Since the n observations are independent, their joint density is

$$\begin{aligned} f(x_1, x_2, \dots, x_n, \theta) &= f(x_1, \theta) f(x_2, \theta) \cdots f(x_n, \theta) \\ &= \prod_{i=1}^n f(x_i, \theta) = L(\theta | x_1, x_2, \dots, x_n). \end{aligned} \quad (\text{C-10})$$

This function, denoted $L(\theta | \mathbf{X})$, is called the likelihood function for θ given the data \mathbf{X} . It is frequently abbreviated to $L(\theta)$. Where no ambiguity can arise, we shall abbreviate it further to L .

Example C.6 Likelihood Functions for Exponential and Normal Distributions

If x_1, \dots, x_n are a sample of n observations from an exponential distribution with parameter θ , then

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

If x_1, \dots, x_n are a sample of n observations from a normal distribution with mean μ and standard deviation σ , then

$$\begin{aligned} L(\mu, \sigma) &= \prod_{i=1}^n (2\pi\sigma^2)^{-1/2} e^{-[1/(2\sigma^2)](x_i - \mu)^2} \\ &= (2\pi\sigma^2)^{-n/2} e^{-[1/(2\sigma^2)]\sum_i (x_i - \mu)^2}. \end{aligned} \quad (\text{C-11})$$

The likelihood function is the cornerstone for most of our theory of parameter estimation. An important result for efficient estimation is the following.

THEOREM C.2 Cramér–Rao Lower Bound

Assuming that the density of x satisfies certain regularity conditions, the variance of an unbiased estimator of a parameter θ will always be at least as large as

$$[I(\theta)]^{-1} = \left(-E \left[\frac{\partial^2 \ln L(\theta)}{\partial^2 \theta} \right] \right)^{-1} = \left(E \left[\left(\frac{\partial \ln L(\theta)}{\partial \theta} \right)^2 \right] \right)^{-1}. \quad (\text{C-12})$$

The quantity $I(\theta)$ is the information number for the sample. We will prove the result that the negative of the expected second derivative equals the expected square of the first derivative in Chapter 17. Proof of the main result of the theorem is quite involved. See, for example, Stuart and Ord (1989).

The regularity conditions are technical in nature. (See Section 17.4.1.) Loosely, they are conditions imposed on the density of the random variable that appears in the likelihood function; these conditions will ensure that the Lindberg–Levy central limit theorem will apply to the sample of observations on the random vector $\mathbf{y} = \partial \ln f(x | \boldsymbol{\theta}) / \partial \boldsymbol{\theta}$. Among the conditions are finite moments of x up to order 3. An additional condition normally included in the set is that the range of the random variable be independent of the parameters.

In some cases, the second derivative of the log likelihood is a constant, so the Cramér–Rao bound is simple to obtain. For instance, in sampling from an exponential distribution, from Example C.6,

$$\begin{aligned} \ln L &= n \ln \theta - \theta \sum_{i=1}^n x_i, \\ \frac{\partial \ln L}{\partial \theta} &= \frac{n}{\theta} - \sum_{i=1}^n x_i, \end{aligned}$$

so $\partial^2 \ln L / \partial \theta^2 = -n / \theta^2$ and the variance bound is $[I(\theta)]^{-1} = \theta^2 / n$. In most situations, the second derivative is a random variable with a distribution of its own. The following examples show two such cases.

Example C.7 Variance Bound for the Poisson Distribution

For the Poisson distribution,

$$\begin{aligned} f(x) &= \frac{e^{-\theta} \theta^x}{x!}, \\ \ln L &= -n\theta + \left(\sum_{i=1}^n x_i \right) \ln \theta - \sum_{i=1}^n \ln(x_i!), \\ \frac{\partial \ln L}{\partial \theta} &= -n + \frac{\sum_{i=1}^n x_i}{\theta}, \\ \frac{\partial^2 \ln L}{\partial \theta^2} &= \frac{-\sum_{i=1}^n x_i}{\theta^2}. \end{aligned}$$

The sum of n identical Poisson variables has a Poisson distribution with parameter equal to n times the parameter of the individual variables. Therefore, the actual distribution of the first derivative will be that of a linear function of a Poisson distributed variable. Since $E[\sum_{i=1}^n x_i] = nE[x_i] = n\theta$, the variance bound for the Poisson distribution is $[I(\theta)]^{-1} = \theta/n$. (Note also that the same result implies that $E[\partial \ln L / \partial \theta] = 0$, which is a result we will use in Chapter 17. The same result holds for the exponential distribution.)

Consider, finally, a multivariate case. If θ is a vector of parameters, then $\mathbf{I}(\theta)$ is the **information matrix**. The Cramér–Rao theorem states that the difference between the covariance matrix of any unbiased estimator and the inverse of the information matrix,

$$[\mathbf{I}(\theta)]^{-1} = \left(-E \left[\frac{\partial^2 \ln L(\theta)}{\partial \theta \partial \theta'} \right] \right)^{-1} = \left\{ E \left[\left(\frac{\partial \ln L(\theta)}{\partial \theta} \right) \left(\frac{\partial \ln L(\theta)}{\partial \theta'} \right) \right] \right\}^{-1}, \quad (\text{C-13})$$

will be a nonnegative definite matrix.

In most settings, numerous estimators are available for the parameters of a distribution. The usefulness of the Cramér–Rao bound is that if one of these is known to attain the variance bound, then there is no need to consider any other to seek a more efficient estimator. Regarding the use of the variance bound, we emphasize that if an unbiased estimator attains it, then that estimator is efficient. If a given estimator does not attain the variance bound, however, then we do not know, except in a few special cases, whether this estimator is efficient or not. It may be that no unbiased estimator can attain the Cramér–Rao bound, which can leave the question of whether a given unbiased estimator is efficient or not unanswered.

We note, finally, that in some cases we further restrict the set of estimators to linear functions of the data.

DEFINITION C.5 Minimum Variance Linear Unbiased Estimator (MVLUE)

An estimator is the minimum variance linear unbiased estimator or best linear unbiased estimator (BLUE) if it is a linear function of the data and has minimum variance among linear unbiased estimators.

In a few instances, such as the normal mean, there will be an efficient linear unbiased estimator; \bar{x} is efficient among all unbiased estimators, both linear and nonlinear. In other cases, such as the normal variance, there is no linear unbiased estimator. This criterion is useful because we can sometimes find an MVLUE without having to specify the distribution at all. Thus, by limiting ourselves to a somewhat restricted class of estimators, we free ourselves from having to assume a particular distribution.

C.6 INTERVAL ESTIMATION

Regardless of the properties of an estimator, the estimate obtained will vary from sample to sample, and there is some probability that it will be quite erroneous. A point estimate will not provide any information on the likely range of error. The logic behind an **interval estimate** is that we use the sample data to construct an interval, [lower (\mathbf{X}), upper (\mathbf{X})], such that we can expect this interval to contain the true parameter in some specified proportion of samples, or

equivalently, with some desired level of confidence. Clearly, the wider the interval, the more confident we can be that it will, in any given sample, contain the parameter being estimated.

The theory of interval estimation is based on a **pivotal quantity**, which is a function of both the parameter and a point estimate that has a known distribution. Consider the following examples.

Example C.8 Confidence Intervals for the Normal Mean

In sampling from a normal distribution with mean μ and standard deviation σ ,

$$z = \frac{\sqrt{n}(\bar{x} - \mu)}{s} \sim t[n - 1]$$

and

$$c = \frac{(n - 1)s^2}{\sigma^2} \sim \chi^2[n - 1].$$

Given the pivotal quantity, we can make probability statements about events involving the parameter and the estimate. Let $p(g, \theta)$ be the constructed random variable, for example, z or c . Given a prespecified **confidence level**, $1 - \alpha$, we can state that

$$\text{Prob}(\text{lower} \leq p(g, \theta) \leq \text{upper}) = 1 - \alpha, \tag{C-14}$$

where lower and upper are obtained from the appropriate table. This statement is then manipulated to make equivalent statements about the endpoints of the intervals. For example, the following statements are equivalent:

$$\text{Prob}\left(-z \leq \frac{\sqrt{n}(\bar{x} - \mu)}{s} \leq z\right) = 1 - \alpha,$$

$$\text{Prob}\left(\bar{x} - \frac{zs}{\sqrt{n}} \leq \mu \leq \bar{x} + \frac{zs}{\sqrt{n}}\right) = 1 - \alpha.$$

The second of these is a statement about the interval, not the parameter; that is, it is the interval that is random, not the parameter. We attach a probability, or $100(1 - \alpha)$ percent confidence level, to the interval itself; in repeated sampling, an interval constructed in this fashion will contain the true parameter $100(1 - \alpha)$ percent of the time.

In general, the interval constructed by this method will be of the form

$$\text{lower}(\mathbf{X}) = \hat{\theta} - e_1,$$

$$\text{upper}(\mathbf{X}) = \hat{\theta} + e_2,$$

where \mathbf{X} is the sample data, e_1 and e_2 are sampling errors, and $\hat{\theta}$ is a point estimate of θ . It is clear from the preceding example that if the sampling distribution of the pivotal quantity is either t or standard normal, which will be true in the vast majority of cases we encounter in practice, then the confidence interval will be

$$\hat{\theta} \pm C_{1-\alpha/2}[\text{se}(\hat{\theta})], \tag{C-15}$$

where $\text{se}(\cdot)$ is the (known or estimated) standard error of the parameter estimate and $C_{1-\alpha/2}$ is the value from the t or standard normal distribution that is exceeded with probability $1 - \alpha/2$. The usual values for α are 0.10, 0.05, or 0.01. The theory does not prescribe exactly how to choose the endpoints for the confidence interval. An obvious criterion is to minimize the width of the interval. If the sampling distribution is symmetric, then the symmetric interval is the best one. If the sampling distribution is not symmetric, however, then this procedure will not be optimal.

Example C.9 Estimated Confidence Intervals for a Normal Mean and Variance

In a sample of 25, $\bar{x} = 1.63$ and $s = 0.51$. Construct a 95 percent confidence interval for μ . Assuming that the sample of 25 is from a normal distribution,

$$\text{Prob}\left(-2.064 \leq \frac{5(\bar{x} - \mu)}{s} \leq 2.064\right) = 0.95,$$

where 2.064 is the critical value from a t distribution with 24 degrees of freedom. Thus, the confidence interval is $1.63 \pm [2.064(0.51)/5]$ or $[1.4195, 1.8405]$.

Remark: Had the parent distribution not been specified, it would have been natural to use the standard normal distribution instead, perhaps relying on the central limit theorem. But a sample size of 25 is small enough that the more conservative t distribution might still be preferable.

The chi-squared distribution is used to construct a confidence interval for the variance of a normal distribution. Using the data from Example 4.29, we find that the usual procedure would use

$$\text{Prob}\left(12.4 \leq \frac{24s^2}{\sigma^2} \leq 39.4\right) = 0.95,$$

where 12.4 and 39.4 are the 0.025 and 0.975 cutoff points from the chi-squared (24) distribution. This procedure leads to the 95 percent confidence interval $[0.1581, 0.5032]$. By making use of the asymmetry of the distribution, a narrower interval can be constructed. Allocating 4 percent to the left-hand tail and 1 percent to the right instead of 2.5 percent to each, the two cutoff points are 13.4 and 42.9, and the resulting 95 percent confidence interval is $[0.1455, 0.4659]$.

Finally, the confidence interval can be manipulated to obtain a confidence interval for a function of a parameter. For example, based on the preceding, a 95 percent confidence interval for σ would be $[\sqrt{0.1581}, \sqrt{0.5032}] = [0.3976, 0.7094]$.

C.7 HYPOTHESIS TESTING

The second major group of statistical inference procedures is hypothesis tests. The classical testing procedures are based on constructing a statistic from a random sample that will enable the analyst to decide, with reasonable confidence, whether or not the data in the sample would have been generated by a hypothesized population. The formal procedure involves a statement of the hypothesis, usually in terms of a “null” or maintained hypothesis and an “alternative,” conventionally denoted H_0 and H_1 , respectively. The procedure itself is a rule, stated in terms of the data, that dictates whether the null hypothesis should be rejected or not. For example, the hypothesis might state a parameter is equal to a specified value. The decision rule might state that the hypothesis should be rejected if a sample estimate of that parameter is too far away from that value (where “far” remains to be defined). The classical, or Neyman–Pearson, methodology involves partitioning the sample space into two regions. If the observed data (i.e., the test statistic) fall in the **rejection region** (sometimes called the **critical region**), then the null hypothesis is rejected; if they fall in the **acceptance region**, then it is not.

C.7.1 CLASSICAL TESTING PROCEDURES

Since the sample is random, the test statistic, however defined, is also random. The same test procedure can lead to different conclusions in different samples. As such, there are two ways such a procedure can be in error:

1. **Type I error.** The procedure may lead to rejection of the null hypothesis when it is true.
2. **Type II error.** The procedure may fail to reject the null hypothesis when it is false.

To continue the previous example, there is some probability that the estimate of the parameter will be quite far from the hypothesized value, even if the hypothesis is true. This outcome might cause a type I error.

DEFINITION C.6 Size of a Test

The probability of a type I error is the *size* of the test. This is conventionally denoted α and is also called the *significance level*.

The size of the test is under the control of the analyst. It can be changed just by changing the decision rule. Indeed, the type I error could be eliminated altogether just by making the rejection region very small, but this would come at a cost. By eliminating the probability of a type I error—that is, by making it unlikely that the hypothesis is rejected—we must increase the probability of a type II error. Ideally, we would like both probabilities to be as small as possible. It is clear, however, that there is a tradeoff between the two. The best we can hope for is that for a given probability of type I error, the procedure we choose will have as small a probability of type II error as possible.

DEFINITION C.7 Power of a Test

The *power* of a test is the probability that it will correctly lead to rejection of a false null hypothesis:

$$\text{power} = 1 - \beta = 1 - \text{Prob}(\text{type II error}). \quad (\text{C-16})$$

For a given significance level α , we would like β to be as small as possible. Since β is defined in terms of the alternative hypothesis, it depends on the value of the parameter.

Example C.10 Testing a Hypothesis About a Mean

For testing $H_0: \mu = \mu^0$ in a normal distribution with known variance σ^2 , the decision rule is to reject the hypothesis if the absolute value of the z statistic, $\sqrt{n}(\bar{x} - \mu^0)/\sigma$, exceeds the predetermined critical value. For a test at the 5 percent significance level, we set the critical value at 1.96. The power of the test, therefore, is the probability that the absolute value of the test statistic will exceed 1.96 given that the true value of μ is, in fact, not μ^0 . This value depends on the alternative value of μ , as shown in Figure C.6. Notice that for this test the power is equal to the size at the point where μ equals μ^0 . As might be expected, the test becomes more powerful the farther the true mean is from the hypothesized value.

Testing procedures, like estimators, can be compared using a number of criteria.

DEFINITION C.8 Most Powerful Test

A test is *most powerful* if it has greater power than any other test of the same size.




FIGURE C.6 Power Function for a Test.

This requirement is very strong. Since the power depends on the alternative hypothesis, we might require that the test be **uniformly most powerful (UMP)**, that is, have greater power than any other test of the same size for all admissible values of the parameter. There are few situations in which a UMP test is available. We usually must be less stringent in our requirements. Nonetheless, the criteria for comparing hypothesis testing procedures are generally based on their respective power functions. A common and very modest requirement is that the test be **unbiased**.

DEFINITION C.9 Unbiased Test

A test is *unbiased* if its power $(1 - \beta)$ is greater than or equal to its size α for all values of the parameter.

If a test is **biased**, then, for some values of the parameter, we are more likely to accept the null hypothesis when it is false than when it is true.

The use of the term *unbiased* here is unrelated to the concept of an unbiased estimator. Fortunately, there is little chance of confusion. Tests and estimators are clearly connected, however. The following criterion derives, in general, from the corresponding attribute of a parameter estimate.

DEFINITION C.10 Consistent Test

A test is *consistent* if its power goes to one as the sample size grows to infinity.

Example C.11 Consistent Test About a Mean

A confidence interval for the mean of a normal distribution is $\bar{x} \pm t_{1-\alpha/2}(s/\sqrt{n})$, where \bar{x} and s are the usual consistent estimators for μ and σ , n is the sample size, and $t_{1-\alpha/2}$ is the correct critical value from the t distribution with $n - 1$ degrees of freedom. For testing $H_0: \mu = \mu_0$ versus $H_1: \mu \neq \mu_0$, let the procedure be to reject H_0 if the confidence interval does not contain μ_0 . Since \bar{x} is consistent for μ , one can discern if H_0 is false as $n \rightarrow \infty$, with probability 1, because \bar{x} will be arbitrarily close to the true μ . Therefore, this test is consistent.

As a general rule, a test will be consistent if it is based on a consistent estimator of the parameter.

C.7.2 TESTS BASED ON CONFIDENCE INTERVALS

There is an obvious link between interval estimation and the sorts of hypothesis tests we have been discussing here. The confidence interval gives a range of plausible values for the parameter. Therefore, it stands to reason that if a hypothesized value of the parameter does not fall in this range of plausible values, then the data are not consistent with the hypothesis, and it should be rejected. Consider, then, testing

$$H_0: \theta = \theta_0,$$

$$H_1: \theta \neq \theta_0.$$

We form a confidence interval based on $\hat{\theta}$ as described earlier:

$$\hat{\theta} - C_{1-\alpha/2}[\text{se}(\hat{\theta})] < \theta < \hat{\theta} + C_{1-\alpha/2}[\text{se}(\hat{\theta})].$$

H_0 is rejected if θ_0 exceeds the upper limit or is less than the lower limit. Equivalently, H_0 is rejected if

$$\left| \frac{\hat{\theta} - \theta_0}{\text{se}(\hat{\theta})} \right| > C_{1-\alpha/2}.$$

In words, the hypothesis is rejected if the estimate is too far from θ_0 , where the distance is measured in standard error units. The critical value is taken from the t or standard normal distribution, whichever is appropriate.

Example C.12 Testing a Hypothesis About a Mean with a Confidence Interval

For the results in Example C.8, test $H_0: \mu = 1.98$ versus $H_1: \mu \neq 1.98$, assuming sampling from a normal distribution:

$$t = \left| \frac{\bar{x} - 1.98}{s/\sqrt{n}} \right| = \left| \frac{1.63 - 1.98}{0.102} \right| = 3.43.$$

The 95 percent critical value for $t(24)$ is 2.064. Therefore, reject H_0 . If the critical value for the standard normal table of 1.96 is used instead, then the same result is obtained.

If the test is one-sided, as in

$$H_0: \theta \geq \theta_0,$$

$$H_1: \theta < \theta_0,$$

then the critical region must be adjusted. Thus, for this test, H_0 will be rejected if a point estimate of θ falls sufficiently below θ_0 . (Tests can usually be set up by departing from the decision criterion, "What sample results are inconsistent with the hypothesis?")

Example C.13 One-Sided Test About a Mean

A sample of 25 from a normal distribution yields $\bar{x} = 1.63$ and $s = 0.51$. Test

$$H_0: \mu \leq 1.5,$$

$$H_1: \mu > 1.5.$$

Clearly, no observed \bar{x} less than or equal to 1.5 will lead to rejection of H_0 . Using the borderline value of 1.5 for μ , we obtain

$$\text{Prob}\left(\frac{\sqrt{n}(\bar{x} - 1.5)}{s} > \frac{5(1.63 - 1.5)}{0.51}\right) = \text{Prob}(t_{24} > 1.27).$$

This is approximately 0.11. This value is not unlikely by the usual standards. Hence, at a significant level of 0.11, we would not reject the hypothesis.

C.7.3 SPECIFICATION TESTS

The hypothesis testing procedures just described are known as “classical” testing procedures. In each case, the null hypothesis tested came in the form of a restriction on the alternative. You can verify that in each application we examined, the parameter space assumed under the null hypothesis is a subspace of that described by the alternative. For that reason, the models implied are said to be “nested.” The null hypothesis is contained within the alternative. This approach suffices for most of the testing situations encountered in practice, but there are common situations in which two competing models cannot be viewed in these terms. For example, consider a case in which there are two completely different, competing theories to explain the same observed data. Many models for censoring and truncation discussed in Chapter 22 rest upon a fragile assumption of normality, for example. Testing of this nature requires a different approach from the classical procedures discussed here. These are discussed at various points throughout the book, for example, in Chapter 13, where we study the difference between fixed and random effects models.

APPENDIX D**LARGE SAMPLE DISTRIBUTION THEORY****D.1 INTRODUCTION**

Most of this book is about parameter estimation. In studying that subject, we will usually be interested in determining how best to use the observed data when choosing among competing estimators. That, in turn, requires us to examine the sampling behavior of estimators. In a few cases, such as those presented in Appendix C and the least squares estimator considered in Chapter 4, we can make broad statements about sampling distributions that will apply regardless of the size of the sample. But, in most situations, it will only be possible to make approximate statements about estimators, such as whether they improve as the sample size increases and what can be said about their sampling distributions in large samples as an approximation to the finite samples we actually observe. This appendix will collect most of the formal, fundamental theorems

and results needed for this analysis. A few additional results will be developed in the discussion of time series analysis later in the book.

D.2 LARGE-SAMPLE DISTRIBUTION THEORY¹

In most cases, whether an estimator is exactly unbiased or what its exact sampling variance is in samples of a given size will be unknown. But we may be able to obtain approximate results about the behavior of the distribution of an estimator as the sample becomes large. For example, it is well known that the distribution of the mean of a sample tends to approximate normality as the sample size grows, regardless of the distribution of the individual observations. Knowledge about the limiting behavior of the distribution of an estimator can be used to infer an approximate distribution for the estimator in a finite sample. To describe how this is done, it is necessary, first, to present some results on convergence of random variables.

D.2.1 CONVERGENCE IN PROBABILITY

Limiting arguments in this discussion will be with respect to the sample size n . Let x_n be a sequence random variable indexed by the sample size.

DEFINITION D.1 Convergence in Probability

The random variable x_n converges in probability to a constant c if $\lim_{n \rightarrow \infty} \text{Prob}(|x_n - c| > \varepsilon) = 0$ for any positive ε .

Convergence in probability implies that the values that the variable may take that are not close to c become increasingly unlikely as n increases. To consider one example, suppose that the random variable x_n takes two values, zero and n , with probabilities $1 - (1/n)$ and $(1/n)$, respectively. As n increases, the second point will become ever more remote from any constant but, at the same time, will become increasingly less probable. In this example, x_n converges in probability to zero. The crux of this form of convergence is that all the mass of the probability distribution becomes concentrated at points close to c . If x_n converges in probability to c , then we write

$$\text{plim } x_n = c. \quad (\text{D-1})$$

We will make frequent use of a special case of convergence in probability, **convergence in mean square** or **convergence in quadratic mean**.

THEOREM D.1 Convergence in Quadratic Mean

If x_n has mean μ_n and variance σ_n^2 such that the ordinary limits of μ_n and σ_n^2 are c and 0 , respectively, then x_n converges in mean square to c , and

$$\text{plim } x_n = c.$$

¹A comprehensive summary of many results in large-sample theory appears in White (2001). The results discussed here will apply to samples of independent observations. Time series cases in which observations are correlated are analyzed in Chapters 19 and 20.

A proof of Theorem D.1 can be based on another useful theorem.

THEOREM D.2 Chebychev's Inequality

If x_n is a random variable and c and ϵ are constants, then $\text{Prob}(|x_n - c| > \epsilon) \leq E[(x_n - c)^2]/\epsilon^2$.

To establish the Chebychev inequality, we use another result [see Goldberger (1991, p. 31)].

THEOREM D.3 Markov's Inequality

If y_n is a nonnegative random variable and δ is a positive constant, then $\text{Prob}[y_n \geq \delta] \leq E[y_n]/\delta$.

Proof: $E[y_n] = \text{Prob}[y_n < \delta]E[y_n | y_n < \delta] + \text{Prob}[y_n \geq \delta]E[y_n | y_n \geq \delta]$. Since y_n is nonnegative, both terms must be nonnegative, so $E[y_n] \geq \text{Prob}[y_n \geq \delta]E[y_n | y_n \geq \delta]$. Since $E[y_n | y_n \geq \delta]$ must be greater than or equal to δ , $E[y_n] \geq \text{Prob}[y_n \geq \delta]\delta$, which is the result.

Now, to prove Theorem D.1., let y_n be $(x_n - c)^2$ and δ be ϵ^2 in Theorem D.3. Then, $(x_n - c)^2 > \delta$ implies that $|x_n - c| > \epsilon$. Finally, we will use a special case of the Chebychev inequality, where $c = \mu_n$, so that we have

$$\text{Prob}(|x_n - \mu_n| > \epsilon) \leq \sigma_n^2/\epsilon^2. \tag{D-2}$$

Taking the limits of μ_n and σ_n^2 in (D-2), we see that if

$$\lim_{n \rightarrow \infty} E[x_n] = c \quad \text{and} \quad \lim_{n \rightarrow \infty} \text{Var}[x_n] = 0, \tag{D-3}$$

then

$$\text{plim } x_n = c.$$

We have shown that convergence in mean square implies convergence in probability. Mean-square convergence implies that the distribution of x_n collapses to a spike at $\text{plim } x_n$, as shown in Figure D.1.

Example D.1 Mean Square Convergence of the Sample Minimum in Exponential Sampling

As noted in Example C.4, in sampling of n observations from an exponential distribution, for the sample minimum $x_{(1)}$,

$$\lim_{n \rightarrow \infty} E[x_{(1)}] = \lim_{n \rightarrow \infty} \frac{1}{n\theta} = 0$$

and

$$\lim_{n \rightarrow \infty} \text{Var}[x_{(1)}] = \lim_{n \rightarrow \infty} \frac{1}{(n\theta)^2} = 0.$$

Therefore,

$$\text{plim } x_{(1)} = 0.$$

Note, in particular, that the variance is divided by n^2 . Thus, this estimator converges very rapidly to 0.




FIGURE D.1 Quadratic Convergence to a Constant, θ .

Convergence in probability does not imply convergence in mean square. Consider the simple example given earlier in which x_n equals either zero or n with probabilities $1 - (1/n)$ and $(1/n)$. The exact expected value of x_n is 1 for all n , which is not the probability limit. Indeed, if we let $\text{Prob}(x_n = n^2) = (1/n)$ instead, the mean of the distribution explodes, but the probability limit is still zero. Again, the point $x_n = n^2$ becomes ever more extreme but, at the same time, becomes ever less likely.

The conditions for convergence in mean square are usually easier to verify than those for the more general form. Fortunately, we shall rarely encounter circumstances in which it will be necessary to show convergence in probability in which we cannot rely upon convergence in mean square. Our most frequent use of this concept will be in formulating consistent estimators.

DEFINITION D.2 Consistent Estimator

An estimator $\hat{\theta}_n$ of a parameter θ is a **consistent** estimator of θ if and only if

$$\text{plim } \hat{\theta}_n = \theta. \tag{D-4}$$

THEOREM D.4 Consistency of the Sample Mean

The mean of a random sample from any population with finite mean μ and finite variance σ^2 is a consistent estimator of μ .

Proof: $E[\bar{x}_n] = \mu$ and $\text{Var}[\bar{x}_n] = \sigma^2/n$. Therefore, \bar{x}_n converges in mean square to μ , or $\text{plim } \bar{x}_n = \mu$.

Theorem D.4 is broader than it might appear at first.

COROLLARY TO THEOREM D.4 Consistency of a Mean of Functions

In random sampling, for any function $g(x)$, if $E[g(x)]$ and $\text{Var}[g(x)]$ are finite constants, then

$$\text{plim} \frac{1}{n} \sum_{i=1}^n g(x_i) = E[g(x)]. \tag{D-5}$$

Proof: Define $y_i = g(x_i)$ and use Theorem D.4.

Example D.2 Estimating a Function of the Mean

In sampling from a normal distribution with mean μ and variance 1, $E[e^x] = e^{\mu+1/2}$ and $\text{Var}[e^x] = e^{2\mu+2} - e^{2\mu+1}$. (See Section B.4.4 on the lognormal distribution.) Hence,

$$\text{plim} \frac{1}{n} \sum_{i=1}^n e^{x_i} = e^{\mu+1/2}.$$

D.2.2 OTHER FORMS OF CONVERGENCE AND LAWS OF LARGE NUMBERS

Theorem D.4 and the corollary given above are particularly narrow forms of a set of results known as **laws of large numbers** that are fundamental to the theory of parameter estimation. Laws of large numbers come in two forms depending on the type of convergence considered. The simpler of these are “weak laws of large numbers” which rely on convergence in probability as we defined it above. “Strong laws” rely on a broader type of convergence called **almost sure convergence**. Overall, the law of large numbers is a statement about the behavior of an average of a large number of random variables.

THEOREM D.5 Khinchine’s Weak Law of Large Numbers

If $x_i, i = 1, \dots, n$ is a random (i.i.d.) sample from a distribution with finite mean $E[x_i] = \mu$, then

$$\text{plim} \bar{x}_n = \mu.$$

Proofs of this and the theorem below are fairly intricate. Rao (1973) provides one.

Notice that this is already broader than Theorem D.4, as it does not require that the variance of the distribution be finite. On the other hand, it is not broad enough, since most of the situations we encounter where we will need a result such as this will not involve i.i.d. random sampling. A broader result is

THEOREM D.6 Chebychev's Weak Law of Large Numbers

If $x_i, i = 1, \dots, n$ is a sample of observations such that $E[x_i] = \mu_i < \infty$ and $\text{Var}[x_i] = \sigma_i^2 < \infty$ such that $\bar{\sigma}_n^2/n = (1/n^2)\sum_i \sigma_i^2 \rightarrow 0$ as $n \rightarrow \infty$ then $\text{plim}(\bar{x}_n - \bar{\mu}_n) = 0$.

There is a subtle distinction between these two theorems that you should notice. The Chebychev theorem does not state that \bar{x}_n converges to $\bar{\mu}_n$, or even that it converges to a constant at all. That would require a precise statement about the behavior of $\bar{\mu}_n$. The theorem states that as n increases without bound, these two quantities will be arbitrarily close to each other—that is, the difference between them converges to a constant, zero. This is an important notion that enters the derivation when we consider statistics that converge to random variables, instead of to constants. What we do have with these two theorems is extremely broad conditions under which a sample mean will converge in probability to its population counterpart. The more important difference between the Khinchine and Chebychev theorems is that the second allows for heterogeneity in the distributions of the random variables that enter the mean.

In analyzing time series data, the sequence of outcomes is itself viewed as a random event. Consider, then, the sample mean, \bar{x}_n . The preceding results concern the behavior of this statistic as $n \rightarrow \infty$ for a particular realization of the sequence $\bar{x}_1, \dots, \bar{x}_n$. But, if the sequence, itself, is viewed as a random event, then limit to which \bar{x}_n converges may be also. The stronger notion of almost sure convergence relates to this possibility.

DEFINITION D.3 Almost Sure Convergence

The random variable x_n converges almost surely to the constant c if and only if

$$\lim_{n \rightarrow \infty} \text{Prob}(|x_i - c| > \varepsilon \text{ for all } i \geq n) = 0 \text{ for all } \varepsilon > 0.$$

Almost sure convergence differs from convergence in probability in an important respect. Note that the index in the probability statement is “ i ,” not “ n .” The definition states that if a sequence converges almost surely, then there is an n large enough such that for any positive ε the probability that the sequence will not converge to c goes to zero. This is denoted $x_n \xrightarrow{a.s.} x$. Again, it states that the probability of observing a sequence that does not converge to c ultimately vanishes. Intuitively, it states that once the sequence x_n becomes close to c , it stays close to c .

From the two definitions, it is clear that almost sure convergence is a stronger form of convergence. Almost sure convergence implies convergence in probability. The proof is obvious given the statements of the definitions. The event described in the definition of almost sure convergence, for any $i \geq n$, includes $i = n$, which is the condition for convergence in probability.

Almost sure convergence is used in a stronger form of the law of large numbers:

THEOREM D.7 Kolmogorov's Strong Law of Large Numbers

If $x_i, i = 1, \dots, n$ is a sequence of independently distributed random variables such that $E[x_i] = \mu_i < \infty$ and $\text{Var}[x_i] = \sigma_i^2 < \infty$ such that $\sum_{i=1}^{\infty} \sigma_i^2/i^2 < \infty$ as $n \rightarrow \infty$ then $x_n - \bar{\mu}_n \xrightarrow{a.s.} 0$.

THEOREM D.8 Markov's Strong Law of Large Numbers

If $\{z_i\}$ is a sequence of independent random variables with $E[z_i] = \mu_i < \infty$ and if for some $\delta > 0$, $\sum_{i=1}^{\infty} E[|z_i - \mu_i|^{1+\delta}] / i^{1+\delta} < \infty$, then $\bar{z}_n - \bar{\mu}_n$ converges almost surely to 0, which we denote $\bar{z}_n - \bar{\mu}_n \xrightarrow{a.s.} 0$.²

The variance condition is satisfied if every variance in the sequence is finite, but this is not strictly required; it only requires that the variances in the sequence increase at a slow enough rate that the sequence of variances as defined is bounded. The theorem allows for heterogeneity in the means and variances. If we return to the conditions of the Khinchine theorem, i.i.d. sampling, we have a corollary:

COROLLARY TO THEOREM D.8 (Kolmogorov)

If $x_i, i = 1, \dots, n$ is a sequence of independent and identically distributed random variables such that $E[x_i] = \mu < \infty$ and $E[|x_i|] < \infty$ then $x_n - \mu \xrightarrow{a.s.} 0$.

Note that the corollary requires identically distributed observations while the theorem only requires independence. Finally, another form of convergence encountered in the analysis of time series data is convergence in r th mean:

DEFINITION D.4 Convergence in r th Mean

If x_n is a sequence of random variables such that $E[|x_n|^r] < \infty$ and $\lim_{n \rightarrow \infty} E[|x_n - c|^r] = 0$, then x_n converges in r th mean to c . This is denoted $x_n \xrightarrow{r.m.} c$.

Surely the most common application is the one we met earlier, convergence in means square, which is convergence in the second mean. Some useful results follow from this definition:

THEOREM D.9 Convergence in Lower Powers

If x_n converges in r th mean to c then x_n converges in s th mean to c for any $s < r$. The proof uses Jensen's Inequality, Theorem D.13. Write $E[|x_n - c|^s] = E[(|x_n - c|^r)^{s/r}] \leq \{E[|x_n - c|^r]\}^{s/r}$ and the inner term converges to zero so the full function must also.

²The use of the expected absolute deviation differs a bit from the expected squared deviation that we have used heretofore to characterize the spread of a distribution. Consider two examples. If $z \sim N[0, \sigma^2]$, then $E[|z|] = \text{Prob}[z < 0]E[z | z < 0] + \text{Prob}[z \geq 0]E[z | z \geq 0] = 0.7989\sigma$. (See Theorem 22.2.) So, finite expected absolute value is the same as finite second moment for the normal distribution. But if z takes values $[0, n]$ with probabilities $[1 - 1/n, 1/n]$, then the variance of z is $(n - 1)$, but $E[|z - \mu_z|]$ is $2 - 2/n$. For this case, finite expected absolute value occurs without finite expected second moment. These are different characterizations of the spread of the distribution.

THEOREM D.10 Generalized Chebychev's Inequality

If x_n is a random variable and c is a constant such that with $E[|x_n - c|^r] < \infty$ and ε is a positive constant, then $\text{Prob}(|x_n - c| > \varepsilon) \leq E[|x_n - c|^r]/\varepsilon^r$.

We have considered two cases of this result already, when $r = 1$ which is the Markov inequality, Theorem D.3 and $r = 2$, which is the Chebychev inequality we looked at first in Theorem D.2.

THEOREM D.11 Convergence in r th mean and Convergence in Probability

If $x_n \xrightarrow{r.m.} c$, for some $r > 0$, then $x_n \xrightarrow{p} c$. The proof relies on Theorem D.9. By assumption, $\lim_{n \rightarrow \infty} E[|x_n - c|^r] = 0$ so for some n sufficiently large, $E[|x_n - c|^r] < \infty$. By Theorem D.9, then, $\text{Prob}(|x_n - c| \geq \varepsilon) \leq E[|x_n - c|^r]/\varepsilon^r$ for any $\varepsilon > 0$. The denominator of the fraction is a fixed constant and the numerator converges to zero by our initial assumption, so $\lim_{n \rightarrow \infty} \text{Prob}(|x_n - c| > \varepsilon) = 0$, which completes the proof.

One implication of Theorem D.11 is that although convergence in mean square is a convenient way to prove convergence in probability, it is actually stronger than necessary, as we get the same result for any positive r .

Finally, we note that we have now shown that both almost sure convergence and convergence in r th mean are stronger than convergence in probability; each implies the latter. But they, themselves, are different notions of convergence, and neither implies the other.

DEFINITION D.5 Convergence of a Random Vector or Matrix

Let \mathbf{x}_n denote a random vector and \mathbf{X}_n a random matrix, and \mathbf{c} and \mathbf{C} denote a vector and matrix of constants with the same dimensions as \mathbf{x}_n and \mathbf{X}_n , respectively. All of the preceding notions of convergence can be extended to $(\mathbf{x}_n, \mathbf{c})$ and $(\mathbf{X}_n, \mathbf{C})$ by applying the results to the respective corresponding elements.

D.2.3 CONVERGENCE OF FUNCTIONS

A particularly convenient result is the following.

THEOREM D.12 Slutsky Theorem

For a continuous function $g(x_n)$ that is not a function of n ,

$$\text{plim } g(x_n) = g(\text{plim } x_n). \quad (\text{D-6})$$

The generalization of Theorem D.12 to a function of several random variables is direct, as illustrated in the next example.

Example D.3 Probability Limit of a Function of \bar{x} and s^2

In random sampling from a population with mean μ and variance σ^2 , the exact expected value of \bar{x}_n^2/s_n^2 will be difficult, if not impossible, to derive. But, by the Slutsky theorem,

$$\text{plim} \frac{\bar{x}_n^2}{s_n^2} = \frac{\mu^2}{\sigma^2}.$$

An application that highlights the difference between expectation and probability is suggested by the following useful relationships.

THEOREM D.13 Inequalities for Expectations
Jensen's Inequality. If $g(x_n)$ is a concave function of x_n , then $g(E[x_n]) \geq E[g(x_n)]$.
Cauchy-Schwartz Inequality. For two random variables, $E[|xy|] \leq \{E[x^2]\}^{1/2} \{E[y^2]\}^{1/2}$.

Although the expected value of a function of x_n may not equal the function of the expected value—it exceeds it if the function is concave—the probability limit of the function is equal to the function of the probability limit.

The Slutsky theorem highlights a comparison between the expectation of a random variable and its probability limit. Theorem D.12 extends directly in two important directions. First, though stated in terms of convergence in probability, the same set of results applies to convergence in r th mean and almost sure convergence. Second, so long as the functions are continuous, the Slutsky Theorem can be extended to vector or matrix valued functions of random scalars, vectors, or matrices. The following describe some specific applications. Some implications of the Slutsky theorem are now summarized.

THEOREM D.14 Rules for Probability Limits
 If x_n and y_n are random variables with $\text{plim} x_n = c$ and $\text{plim} y_n = d$, then

$\text{plim}(x_n + y_n) = c + d,$ (sum rule) (D-7)

$\text{plim} x_n y_n = cd,$ (product rule) (D-8)

$\text{plim} x_n/y_n = c/d$ if $d \neq 0.$ (ratio rule) (D-9)

If \mathbf{W}_n is a matrix whose elements are random variables and if $\text{plim} \mathbf{W}_n = \mathbf{\Omega}$, then

$\text{plim} \mathbf{W}_n^{-1} = \mathbf{\Omega}^{-1}.$ (matrix inverse rule) (D-10)

If \mathbf{X}_n and \mathbf{Y}_n are random matrices with $\text{plim} \mathbf{X}_n = \mathbf{A}$ and $\text{plim} \mathbf{Y}_n = \mathbf{B}$, then

$\text{plim} \mathbf{X}_n \mathbf{Y}_n = \mathbf{AB}.$ (matrix product rule) (D-11)

D.2.4 CONVERGENCE TO A RANDOM VARIABLE

The preceding has dealt with conditions under which a random variable converges to a constant, for example, the way that a sample mean converges to the population mean. In order to develop

a theory for the behavior of estimators, as a prelude to the discussion of limiting distributions, we now consider cases in which a random variable converges not to a constant, but to another random variable. These results will actually subsume those in the preceding section, as a constant may always be viewed as a degenerate random variable, that is one with zero variance.

DEFINITION D.6 Convergence in Probability to a Random Variable

The random variable x_n converges in probability to the random variable x if $\lim_{n \rightarrow \infty} \text{Prob}(|x_n - x| > \varepsilon) = 0$ for any positive ε .

As before, we write $\text{plim } x_n = x$ to denote this case. The interpretation (at least the intuition) of this type of convergence is different when x is a random variable. The notion of closeness defined here relates not to the concentration of the mass of the probability mechanism generating x_n at a point c , but to the closeness of that probability mechanism to that of x . One can think of this as a convergence of the CDF of x_n to that of x .

DEFINITION D.7 Almost Sure Convergence to a Random Variable

The random variable x_n converges almost surely to the random variable x if and only if $\lim_{n \rightarrow \infty} \text{Prob}(|x_i - x| > \varepsilon \text{ for all } i \geq n) = 0$ for all $\varepsilon > 0$.

DEFINITION D.8 Convergence in r th mean to a Random Variable

The random variable x_n converges in r th mean to the random variable x if and only if $\lim_{n \rightarrow \infty} E[|x_n - x|^r] = 0$. This is labeled $x_n \xrightarrow{r.m.} x$. As before, the case $r = 2$ is labeled convergence in mean square.

Once again, we have to revise our understanding of convergence when convergence is to a random variable.

THEOREM D.15 Convergence of Moments

Suppose $x_n \xrightarrow{r.m.} x$ and $E[|x|^r]$ is finite. Then, $\lim_{n \rightarrow \infty} E[|x_n|^r] = E[|x|^r]$.

Theorem D.15 raises an interesting question. Suppose we let r grow, and suppose that $x_n \xrightarrow{r.m.} x$ and, in addition, all moments are finite. If this holds for any r , do we conclude that these random variables have the same distribution? The answer to this longstanding problem in probability theory—the problem of the sequence of moments—is no. The sequence of moments does not uniquely determine the distribution. Although convergence in r th mean and almost surely still both imply convergence in probability, it remains true, even with convergence to a random variable instead of a constant, that these are different forms of convergence.

**D.2.5 CONVERGENCE IN DISTRIBUTION:
LIMITING DISTRIBUTIONS**

A second form of convergence is **convergence in distribution**. Let x_n be a sequence of random variables indexed by the sample size, and assume that x_n has cdf $F_n(x)$.

DEFINITION D.9 Convergence in Distribution
 x_n converges in distribution to a random variable x with cdf $F(x)$ if $\lim_{n \rightarrow \infty} |F_n(x_n) - F(x)| = 0$ at all continuity points of $F(x)$.

This statement is about the probability distribution associated with x_n ; it does not imply that x_n converges at all. To take a trivial example, suppose that the exact distribution of the random variable x_n is

$$\text{Prob}(x_n = 1) = \frac{1}{2} + \frac{1}{n+1}, \quad \text{Prob}(x_n = 2) = \frac{1}{2} - \frac{1}{n+1}.$$

As n increases without bound, the two probabilities converge to $\frac{1}{2}$, but x_n does not converge to a constant.

DEFINITION D.10 Limiting Distribution
*If x_n converges in distribution to x , where $F(x_n)$ is the cdf of x_n , then $F(x)$ is the **limiting distribution** of x . This is written*

$$x_n \xrightarrow{d} x.$$

The limiting distribution is often given in terms of the pdf, or simply the parametric family. For example, “the limiting distribution of x_n is standard normal.”

Convergence in distribution can be extended to random vectors and matrices, though not in the element by element manner that we extended the earlier convergence forms. The reason is that convergence in distribution is a property of the CDF of the random variable, not the variable itself. Thus, we can obtain a convergence result analogous to that in Definition D.9 for vectors or matrices by applying definition to the joint CDF for the elements of the vector or matrices. Thus, $\mathbf{x}_n \xrightarrow{d} \mathbf{x}$ if $\lim_{n \rightarrow \infty} |F_n(\mathbf{x}_n) - F(\mathbf{x})| = 0$ and likewise for a random matrix.

Example D.4 Limiting Distribution of t_{n-1}

Consider a sample of size n from a standard normal distribution. A familiar inference problem is the test of the hypothesis that the population mean is zero. The test statistic usually used is the t statistic:

$$t_{n-1} = \frac{\bar{x}_n}{S_n/\sqrt{n}},$$

where

$$S_n^2 = \frac{\sum_{i=1}^n (x_i - \bar{x}_n)^2}{n-1}.$$

The exact distribution of the random variable t_{n-1} is t with $n - 1$ degrees of freedom. The density is different for every n :

$$f(t_{n-1}) = \frac{\Gamma(n/2)}{\Gamma[(n-1)/2]} [(n-1)\pi]^{-1/2} \left[1 + \frac{t_{n-1}^2}{n-1} \right]^{-n/2} \quad (\text{D-12})$$

as is the cdf, $F_{n-1}(t) = \int_{-\infty}^t f_{n-1}(x) dx$. This distribution has mean zero and variance $(n-1)/(n-3)$. As n grows to infinity, t_{n-1} converges to the standard normal, which is written

$$t_{n-1} \xrightarrow{d} N[0, 1].$$

DEFINITION D.11 Limiting Mean and Variance

The *limiting mean* and *variance* of a random variable are the mean and variance of the limiting distribution, assuming that the limiting distribution and its moments exist.

For the random variable with $t[n]$ distribution, the exact mean and variance are zero and $n/(n-2)$, whereas the limiting mean and variance are zero and one. The example might suggest that the limiting mean and variance are zero and one; that is, that the moments of the limiting distribution are the ordinary limits of the moments of the finite sample distributions. This situation is almost always true, but it need not be. It is possible to construct examples in which the exact moments do not even exist, even though the moments of the limiting distribution are well defined.³ Even in such cases, we can usually derive the mean and variance of the limiting distribution.

Limiting distributions, like probability limits, can greatly simplify the analysis of a problem. Some results that combine the two concepts are as follows.⁴

THEOREM D.16 Rules for Limiting Distributions

1. If $x_n \xrightarrow{d} x$ and $\text{plim } y_n = c$, then

$$x_n y_n \xrightarrow{d} cx, \quad (\text{D-13})$$

which means that the limiting distribution of $x_n y_n$ is the distribution of cx . Also,

$$x_n + y_n \xrightarrow{d} x + c, \quad (\text{D-14})$$

$$x_n / y_n \xrightarrow{d} x/c, \quad \text{if } c \neq 0. \quad (\text{D-15})$$

2. If $x_n \xrightarrow{d} x$ and $g(x_n)$ is a continuous function, then

$$g(x_n) \xrightarrow{d} g(x). \quad (\text{D-16})$$

This result is analogous to the Slutsky theorem for probability limits. For an example, consider the t_n random variable discussed earlier. The exact distribution of t_n^2 is $F[1, n]$. But as $n \rightarrow \infty$, t_n converges to a standard normal variable. According to this result, the limiting distribution of t_n^2 will be that of the square of a standard normal, which is chi-squared with one

³See, for example, Maddala (1977a, p. 150).

⁴For proofs and further discussion, see, for example, Greenberg and Webster (1983).

THEOREM D.16 (Continued)

degree of freedom. We conclude, therefore, that

$$F[1, n] \xrightarrow{d} \text{chi-squared}[1]. \tag{D-17}$$

We encountered this result in our earlier discussion of limiting forms of the standard normal family of distributions.

3. *If y_n has a limiting distribution and $\text{plim}(x_n - y_n) = 0$, then x_n has the same limiting distribution as y_n .*

The third result in Theorem D.16 combines convergence in distribution and in probability. The second result can be extended to vectors and matrices.

Example D.5 The F Distribution

Suppose that $\mathbf{t}_{1,n}$ and $\mathbf{t}_{2,n}$ are a $K \times 1$ and an $M \times 1$ random vector of variables whose components are independent with each distributed as t with n degrees of freedom. Then, as we saw in the preceding, for any component in either random vector, the limiting distribution is standard normal, so for the entire vector, $\mathbf{t}_{j,n} \xrightarrow{d} \mathbf{z}_n$, a vector of independent standard normally distributed variables. The results so far show that $\frac{(\mathbf{t}_{1,n}^* \mathbf{t}_{1,n})/K}{(\mathbf{t}_{2,n}^* \mathbf{t}_{2,n})/M} \xrightarrow{d} F[K, M]$.

Finally, a specific case of result 2 in Theorem D.16 produces a tool known as the Cramér-Wold device.

THEOREM D.17 Cramer-Wold Device

If $\mathbf{x}_n \xrightarrow{d} \mathbf{x}$, then $\mathbf{c}'\mathbf{x}_n \xrightarrow{d} \mathbf{c}'\mathbf{x}$ for all conformable vectors \mathbf{c} with real valued elements.

By allowing \mathbf{c} to be a vector with just a one in a particular position and zeros elsewhere, we see that convergence in distribution of a random vector \mathbf{x}_n to \mathbf{x} does imply that each component does likewise.

D.2.6 CENTRAL LIMIT THEOREMS

We are ultimately interested in finding a way to describe the statistical properties of estimators when their exact distributions are unknown. The concepts of consistency and convergence in probability are important. But the theory of limiting distributions given earlier is not yet adequate. We rarely deal with estimators that are not consistent for something, though perhaps not always the parameter we are trying to estimate. As such,

$$\text{if } \text{plim } \hat{\theta}_n = \theta, \text{ then } \hat{\theta}_n \xrightarrow{d} \theta.$$

That is, the limiting distribution of $\hat{\theta}_n$ is a spike. This is not very informative, nor is it at all what we have in mind when we speak of the statistical properties of an estimator. (To endow our finite sample estimator $\hat{\theta}_n$ with the zero sampling variance of the spike at θ would be optimistic in the extreme.)

As an intermediate step, then, to a more reasonable description of the statistical properties of an estimator, we use a **stabilizing transformation** of the random variable to one that does have

a well-defined limiting distribution. To jump to the most common application, whereas

$$\text{plim } \hat{\theta}_n = \theta,$$

we often find that

$$z_n = \sqrt{n}(\hat{\theta}_n - \theta) \xrightarrow{d} f(z),$$

where $f(z)$ is a well-defined distribution with a mean and a positive variance. An estimator which has this property is said to be **root- n consistent**. The single most important theorem in econometrics provides an application of this proposition. A basic form of the theorem is as follows.

THEOREM D.18 Lindberg–Levy Central Limit Theorem (Univariate)

If x_1, \dots, x_n are a random sample from a probability distribution with finite mean μ and finite variance σ^2 and $\bar{x}_n = (1/n) \sum_{i=1}^n x_i$, then

$$\sqrt{n}(\bar{x}_n - \mu) \xrightarrow{d} N[0, \sigma^2],$$

A proof appears in Rao (1973, p. 127).

The result is quite remarkable as it holds regardless of the form of the parent distribution. For a striking example, return to Figure C.2. The distribution from which the data were drawn in that figure does not even remotely resemble a normal distribution. In samples of only four observations the force of the central limit theorem is clearly visible in the sampling distribution of the means. The sampling experiment Example D.5 shows the effect in a systematic demonstration of the result.

The Lindberg–Levy theorem is one of several forms of this extremely powerful result. For our purposes, an important extension allows us to relax the assumption of equal variances. The Lindberg–Feller form of the central limit theorem is the centerpiece of most of our analysis in econometrics.

THEOREM D.19 Lindberg–Feller Central Limit Theorem (with Unequal Variances)

Suppose that $\{x_i\}$, $i = 1, \dots, n$, is a sequence of independent random variables with finite means μ_i and finite positive variances σ_i^2 . Let

$$\bar{\mu}_n = \frac{1}{n}(\mu_1 + \mu_2 + \dots + \mu_n) \quad \text{and} \quad \bar{\sigma}_n^2 = \frac{1}{n}(\sigma_1^2 + \sigma_2^2 + \dots).$$

If no single term dominates this average variance, which we could state as $\lim_{n \rightarrow \infty} \max(\sigma_i) / (n\bar{\sigma}_n) = 0$, and if the average variance converges to a finite constant, $\bar{\sigma}^2 = \lim_{n \rightarrow \infty} \bar{\sigma}_n^2$, then

$$\sqrt{n}(\bar{x}_n - \bar{\mu}_n) \xrightarrow{d} N[0, \bar{\sigma}^2].$$




FIGURE D.2 The Exponential Distribution.

In practical terms, the theorem states that sums of random variables, regardless of their form, will tend to be normally distributed. The result is yet more remarkable in that *it does not require the variables in the sum to come from the same underlying distribution. It requires, essentially, only that the mean be a mixture of many random variables, none of which is large compared with their sum.* Since nearly all the estimators we construct in econometrics fall under the purview of the central limit theorem, it is obviously an important result.

Example D.6 The Lindberg–Levy Central Limit Theorem

We'll use a sampling experiment to demonstrate the operation of the central limit theorem. Consider random sampling from the exponential distribution with mean 1.5—this is the setting used in Example C.4. The density is shown in Figure D.2.

We've drawn 1,000 samples of 3, 6, and 20 observations from this population and computed the sample means for each. For each mean, we then computed $z_{in} = \sqrt{n}(\bar{x}_{in} - \mu)$ where $i = 1, \dots, 1,000$ and n is 3, 6 or 20. The three rows of figures in Figure D.3 show histograms of the observed samples of sample means and kernel density estimates of the underlying distributions for the three samples of transformed means.

Proof of the Lindberg–Feller theorem requires some quite intricate mathematics [see Loeve (1977), for example] that are well beyond the scope of our work here. We do note an important consideration in this theorem. The result rests on a condition known as the Lindberg condition. The sample mean computed in the theorem is a mixture of random variables from possibly different distributions. The Lindberg condition, in words, states that the contribution of the tail areas of these underlying distributions to the variance of the sum must be negligible in the limit. The condition formalizes the assumption in Theorem D.12 that the average variance be positive and not be dominated by any single term. (For an intuitively crafted mathematical discussion of this condition, see White (2001, pp. 117–118.) The condition is essentially impossible to verify in practice, so it is useful to have a simpler version of the theorem which encompasses it.





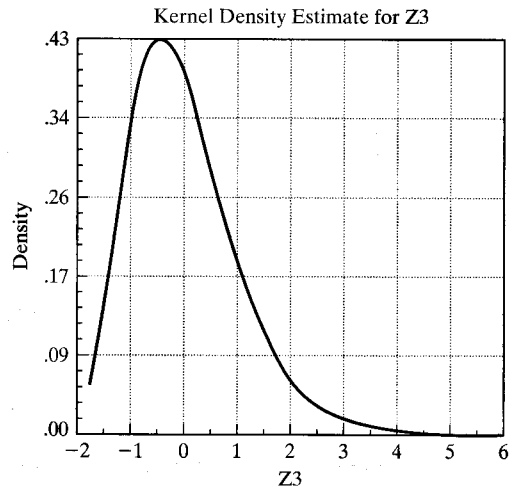



FIGURE D.3 THE CENTRAL LIMIT THEOREM.

THEOREM D.20 Liapounov Central Limit Theorem

Suppose that $\{x_i\}$, is a sequence of independent random variables with finite means μ_i and finite positive variances σ_i^2 such that $E[|x_i - \mu_i|^{2+\delta}]$ is finite for some $\delta > 0$. If $\bar{\sigma}_n$ is positive and finite for all n sufficiently large, then

$$\sqrt{n}(\bar{x}_n - \bar{\mu}_n)/\bar{\sigma}_n \xrightarrow{d} N[0, 1].$$

This version of the central limit theorem requires only that moments slightly larger than two be finite.

Note the distinction between the laws of large numbers in Theorems D.5 and D.6 and the central limit theorems. Neither assert that sample means tend to normality. Sample means (that is, the distributions of them) converge to spikes at the true mean. It is the transformation of the mean, $\sqrt{n}(\bar{x}_n - \mu)/\sigma$, that converges to standard normality. To see this at work, if you have access to the necessary software, you might try reproducing Example D.5 using the raw means, \bar{x}_{in} . What do you expect to observe?

For later purposes, we will require multivariate versions of these theorems. Proofs of the following may be found, for example, in Greenberg and Webster (1983) or Rao (1973) and references cited there.

THEOREM D.18A Multivariate Lindberg–Levy Central Limit Theorem

If $\mathbf{x}_1, \dots, \mathbf{x}_n$ are a random sample from a multivariate distribution with finite mean vector $\boldsymbol{\mu}$ and finite positive definite covariance matrix \mathbf{Q} , then

$$\sqrt{n}(\bar{\mathbf{x}}_n - \boldsymbol{\mu}) \xrightarrow{d} N[\mathbf{0}, \mathbf{Q}],$$

where

$$\bar{\mathbf{x}}_n = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i.$$

In order to get from D.18 to D.18A (and D.19 to D.19A) we need to add a step. Theorem D.18 applies to the individual elements of the vector. A vector has a multivariate normal distribution if the individual elements are normally distributed and if every linear combination is normally distributed. We can use Theorem D.18 (D.19) for the individual terms and Theorem D.17 to establish that linear combinations behave likewise. This establishes the extensions.

The extension of the Lindberg–Feller theorem to unequal covariance matrices requires some intricate mathematics. The following is an informal statement of the relevant conditions. Further discussion and references appear in Fomby, Hill, and Johnson (1984) and Greenberg and Webster (1983).

THEOREM D.19A Multivariate Lindberg–Feller Central Limit Theorem

Suppose that $\mathbf{x}_1, \dots, \mathbf{x}_n$ are a sample of random vectors such that $E[\mathbf{x}_i] = \boldsymbol{\mu}_i$, $\text{Var}[\mathbf{x}_i] = \mathbf{Q}_i$, and all mixed third moments of the multivariate distribution are finite. Let

$$\bar{\boldsymbol{\mu}}_n = \frac{1}{n} \sum_{i=1}^n \boldsymbol{\mu}_i,$$

$$\bar{\mathbf{Q}}_n = \frac{1}{n} \sum_{i=1}^n \mathbf{Q}_i.$$

We assume that

$$\lim_{n \rightarrow \infty} \bar{\mathbf{Q}}_n = \mathbf{Q},$$

where \mathbf{Q} is a finite, positive definite matrix, and that for every i ,

$$\lim_{n \rightarrow \infty} (n\bar{\mathbf{Q}}_n)^{-1} \mathbf{Q}_i = \lim_{n \rightarrow \infty} \left(\sum_{i=1}^n \mathbf{Q}_i \right)^{-1} \mathbf{Q}_i = \mathbf{0}.$$

We allow the means of the random vectors to differ, although in the cases that we will analyze, they will generally be identical. The second assumption states that individual components of the sum must be finite and diminish in significance. There is also an implicit assumption that the sum of matrices is nonsingular. Since the limiting matrix is nonsingular, the assumption must hold for large enough n , which is all that concerns us here. With these in place, the result is

$$\sqrt{n}(\bar{\mathbf{x}}_n - \bar{\boldsymbol{\mu}}_n) \xrightarrow{d} N[\mathbf{0}, \mathbf{Q}].$$

D.2.7 THE DELTA METHOD

At several points in Appendix C, we used a linear Taylor series approximation to analyze the distribution and moments of a random variable. We are now able to justify this usage. We complete the development of Theorem D.12 (probability limit of a function of a random variable), Theorem D.16 (2) (limiting distribution of a function of a random variable), and the central limit theorems, with a useful result that is known as “the **delta method**.” For a single random variable (sample mean or otherwise), we have the following theorem.

THEOREM D.21 Limiting Normal Distribution of a Function

If $\sqrt{n}(z_n - \mu) \xrightarrow{d} N[0, \sigma^2]$ and if $g(z_n)$ is a continuous function not involving n , then

$$\sqrt{n}[g(z_n) - g(\mu)] \xrightarrow{d} N[0, \{g'(\mu)\}^2 \sigma^2]. \quad (\text{D-18})$$

Notice that the mean and variance of the limiting distribution are the mean and variance of the linear Taylor series approximation:

$$g(z_n) \simeq g(\mu) + g'(\mu)(z_n - \mu).$$

The multivariate version of this theorem will be used at many points in the text.

THEOREM D.21A Limiting Normal Distribution of a Set of Functions

If \mathbf{z}_n is a $K \times 1$ sequence of vector-valued random variables such that $\sqrt{n}(\mathbf{z}_n - \mu) \xrightarrow{d} N[\mathbf{0}, \Sigma]$ and if $\mathbf{c}(\mathbf{z}_n)$ is a set of J continuous functions of \mathbf{z}_n not involving n , then

$$\sqrt{n}[\mathbf{c}(\mathbf{z}_n) - \mathbf{c}(\mu)] \xrightarrow{d} N[\mathbf{0}, \mathbf{C}(\mu)\Sigma\mathbf{C}(\mu)'], \quad (\text{D-19})$$

where $\mathbf{C}(\mu)$ is the $J \times K$ matrix $\partial\mathbf{c}(\mu)/\partial\mu'$. The j th row of $\mathbf{C}(\mu)$ is the vector of partial derivatives of the j th function with respect to μ' .

D.3 ASYMPTOTIC DISTRIBUTIONS

The theory of limiting distributions is only a means to an end. We are interested in the behavior of the estimators themselves. The limiting distributions obtained through the central limit theorem all involve unknown parameters, generally the ones we are trying to estimate. Moreover, our samples are always finite. Thus, we depart from the limiting distributions to derive the asymptotic distributions of the estimators.


DEFINITION D.12 Asymptotic Distribution

An asymptotic distribution is a distribution that is used to approximate the true finite sample distribution of a random variable.⁵

By far the most common means of formulating an asymptotic distribution (at least by econometricians) is to construct it from the known limiting distribution of a function of the random variable. If

$$\sqrt{n}[(\bar{x}_n - \mu)/\sigma] \xrightarrow{d} N[0, 1],$$

⁵We depart somewhat from some other treatments [e.g., White (2001), Hayashi (2000, p. 90)] at this point, because they make no distinction between an asymptotic distribution and the limiting distribution, although the treatments are largely along the lines discussed here. In the interest of maintaining consistency of the discussion, we prefer to retain the sharp distinction and derive the asymptotic distribution of an estimator, \mathbf{t} by first obtaining the limiting distribution of $\sqrt{n}(\mathbf{t} - \theta)$. By our construction, the limiting distribution of \mathbf{t} is degenerate, whereas the asymptotic distribution of $\sqrt{n}(\mathbf{t} - \theta)$ is not useful.


FIGURE D.4 True Versus Asymptotic Distribution.

then approximately, or asymptotically, $\bar{x}_n \sim N[\mu, \sigma^2/n]$, which we write as

$$\bar{x}_n \stackrel{a}{\sim} N[\mu, \sigma^2/n].$$

The statement “ \bar{x}_n is asymptotically normally distributed with mean μ and variance σ^2/n ” says only that this normal distribution provides an approximation to the true distribution, not that the true distribution is exactly normal.

Example D.7 Asymptotic Distribution of the Mean of an Exponential Sample

In sampling from an exponential distribution with parameter θ , the exact distribution of \bar{x}_n is that of $\theta/(2n)$ times a chi-squared variable with $2n$ degrees of freedom. The asymptotic distribution is $N[\theta, \theta^2/n]$. The exact and asymptotic distributions are shown in Figure D.4 for the case of $\theta = 1$ and $n = 16$.

Extending the definition, suppose that $\hat{\theta}_n$ is an estimator of the parameter vector θ . The asymptotic distribution of the vector $\hat{\theta}_n$ is obtained from the limiting distribution:

$$\sqrt{n}(\hat{\theta}_n - \theta) \xrightarrow{d} N[\mathbf{0}, \mathbf{V}] \quad (\text{D-20})$$

implies that

$$\hat{\theta}_n \stackrel{a}{\sim} N\left[\theta, \frac{1}{n}\mathbf{V}\right]. \quad (\text{D-21})$$

This notation is read “ $\hat{\theta}_n$ is asymptotically normally distributed, with mean vector θ and covariance matrix $(1/n)\mathbf{V}$.” The covariance matrix of the asymptotic distribution is the **asymptotic covariance matrix** and is denoted

$$\text{Asy. Var}[\hat{\theta}_n] = \frac{1}{n}\mathbf{V}.$$

Note, once again, the logic used to reach the result; (D-20) holds exactly as $n \rightarrow \infty$. We assume that it holds approximately for finite n , which leads to (D-21).

DEFINITION D.13 Asymptotic Normality and Asymptotic Efficiency

An estimator $\hat{\theta}_n$ is asymptotically normal if (D-20) holds. The estimator is asymptotically efficient if the covariance matrix of any other consistent, asymptotically normally distributed estimator exceeds $(1/n)\mathbf{V}$ by a nonnegative definite matrix.

For most estimation problems, these are the criteria used to choose an estimator.

Example D.8 Asymptotic Inefficiency of the Median in Normal Sampling

In sampling from a normal distribution with mean μ and variance σ^2 , both the mean \bar{x}_n and the median M_n of the sample are consistent estimators of μ . Since the limiting distributions of both estimators are spikes at μ , they can only be compared on the basis of their asymptotic properties. The necessary results are

$$\bar{x}_n \overset{a}{\sim} N[\mu, \sigma^2/n] \quad \text{and} \quad M_n \overset{a}{\sim} N[\mu, (\pi/2)\sigma^2/n]. \tag{D-22}$$

Therefore, the mean is more efficient by a factor of $\pi/2$. (But, see Examples E.1 and E.2 for a finite sample result.)

D.3.1 ASYMPTOTIC DISTRIBUTION OF A NONLINEAR FUNCTION

Theorems D.12 and D.14 for functions of a random variable have counterparts in asymptotic distributions.

THEOREM D.22 Asymptotic Distribution of a Nonlinear Function

If $\sqrt{n}(\hat{\theta}_n - \theta) \xrightarrow{d} N[0, \sigma^2]$ and if $g(\theta)$ is a continuous function not involving n , then $g(\hat{\theta}_n) \overset{a}{\sim} N[g(\theta), (1/n)\{g'(\theta)\}^2\sigma^2]$. If $\hat{\theta}_n$ is a vector of parameter estimators such that $\hat{\theta}_n \overset{a}{\sim} N[\theta, (1/n)\mathbf{V}]$ and if $\mathbf{c}(\theta)$ is a set of J continuous functions not involving n , then, $\mathbf{c}(\hat{\theta}_n) \overset{a}{\sim} N[\mathbf{c}(\theta), (1/n)\mathbf{C}(\theta)\mathbf{V}\mathbf{C}(\theta)']$, where $\mathbf{C}(\theta) = \partial\mathbf{c}(\theta)/\partial\theta'$.

Example D.9 Asymptotic Distribution of a Function of Two Estimators

Suppose that b_n and t_n are estimators of parameters β and θ such that

$$\begin{bmatrix} b_n \\ t_n \end{bmatrix} \overset{a}{\sim} N \left[\begin{pmatrix} \beta \\ \theta \end{pmatrix}, \begin{pmatrix} \sigma_{\beta\beta} & \sigma_{\beta\theta} \\ \sigma_{\theta\beta} & \sigma_{\theta\theta} \end{pmatrix} \right].$$

Find the asymptotic distribution of $c_n = b_n/(1 - t_n)$. Let $\gamma = \beta/(1 - \theta)$. By the Slutsky theorem, c_n is consistent for γ . We shall require

$$\frac{\partial\gamma}{\partial\beta} = \frac{1}{1 - \theta} = \gamma_\beta, \quad \frac{\partial\gamma}{\partial\theta} = \frac{\beta}{(1 - \theta)^2} = \gamma_\theta.$$

Let Σ be the 2×2 asymptotic covariance matrix given previously. Then the asymptotic

variance of c_n is

$$\text{Asy. Var}[c_n] = (\gamma_\beta \gamma_\theta) \Sigma \begin{pmatrix} \gamma_\beta \\ \gamma_\theta \end{pmatrix} = \gamma_\beta^2 \sigma_{\beta\beta} + \gamma_\theta^2 \sigma_{\theta\theta} + 2\gamma_\beta \gamma_\theta \sigma_{\beta\theta},$$

which is the variance of the linear Taylor series approximation:

$$\hat{y}_n \simeq \gamma + \gamma_\beta(b_n - \beta) + \gamma_\theta(t_n - \theta).$$

D.3.2 ASYMPTOTIC EXPECTATIONS

The asymptotic mean and variance of a random variable are usually the mean and variance of the asymptotic distribution. Thus, for an estimator with the limiting distribution defined in

$$\sqrt{n}(\hat{\theta}_n - \theta) \xrightarrow{d} N[\mathbf{0}, \mathbf{V}],$$

the asymptotic expectation is θ and the asymptotic variance is $(1/n)\mathbf{V}$. This statement implies, among other things, that the estimator is “asymptotically unbiased.”

At the risk of clouding the issue a bit, it is necessary to reconsider one aspect of the previous description. We have deliberately avoided the use of consistency even though, in most instances, that is what we have in mind. The description thus far might suggest that consistency and asymptotic unbiasedness are the same. Unfortunately (because it is a source of some confusion), they are not. They are if the estimator is consistent and asymptotically normally distributed, or CAN. They may differ in other settings, however. There are at least three possible definitions of asymptotic unbiasedness:

1. The mean of the limiting distribution of $\sqrt{n}(\hat{\theta}_n - \theta)$ is 0.
2. $\lim_{n \rightarrow \infty} E[\hat{\theta}_n] = \theta$. (D-23)
3. $\text{plim } \hat{\theta}_n = \theta$.

In most cases encountered in practice, the estimator in hand will have all three properties, so there is no ambiguity. It is not difficult to construct cases in which the left-hand sides of all three definitions are different, however.⁶ There is no general agreement among authors as to the precise meaning of asymptotic unbiasedness, perhaps because the term is misleading at the outset; *asymptotic* refers to an approximation, whereas *unbiasedness* is an exact result.⁷ Nonetheless, the majority view seems to be that (2) is the proper definition of asymptotic unbiasedness.⁸ Note, though, that this definition relies on quantities that are generally unknown and that may not exist.

A similar problem arises in the definition of the asymptotic variance of an estimator. One common definition is

$$\text{Asy. Var}[\hat{\theta}_n] = \frac{1}{n} \lim_{n \rightarrow \infty} E \left[\left\{ \sqrt{n}(\hat{\theta}_n - \lim_{n \rightarrow \infty} E[\hat{\theta}_n]) \right\}^2 \right].^9 \quad (\text{D-24})$$

This result is a **leading term approximation**, and it will be sufficient for nearly all applications.

⁶See, for example, Maddala (1977a, p. 150).

⁷See, for example, Theil (1971, p. 377).

⁸Many studies of estimators analyze the “asymptotic bias” of, say, $\hat{\theta}_n$ as an estimator of a parameter θ . In most cases, the quantity of interest is actually $\text{plim } [\hat{\theta}_n - \theta]$. See, for example, Greene (1980b) and another example in Johnston (1984, p. 312).

⁹Kmenta (1986, p.165).

Note, however, that like definition 2 of asymptotic unbiasedness, it relies on unknown and possibly nonexistent quantities.

Example D.10 Asymptotic Moments of the Sample Variance

The exact expected value and variance of the variance estimator

$$m_2 = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2 \tag{D-25}$$

are

$$E[m_2] = \frac{(n-1)\sigma^2}{n} \tag{D-26}$$

and

$$\text{Var}[m_2] = \frac{\mu_4 - \sigma^4}{n} - \frac{2(\mu_4 - 2\sigma^4)}{n^2} + \frac{\mu_4 - 3\sigma^4}{n^3}, \tag{D-27}$$

where $\mu_4 = E[(x - \mu)^4]$. [See Goldberger (1964, pp. 97–99).] The leading term approximation would be

$$\text{Asy. Var}[m_2] = \frac{1}{n}(\mu_4 - \sigma^4).$$

D.4 SEQUENCES AND THE ORDER OF A SEQUENCE

This section has been concerned with sequences of constants, denoted, for example c_n , and random variables, such as x_n , that are indexed by a sample size, n . An important characteristic of a sequence is the rate at which it converges (or diverges). For example, as we have seen, the mean of a random sample of n observations from a distribution with finite mean, μ , and finite variance, σ^2 , is itself a random variable with variance $\gamma_n^2 = \sigma^2/n$. We see that as long as σ^2 is a finite constant, γ_n^2 is a sequence of constants that converges to zero. Another example is the random variable $x_{(1),n}$, the minimum value in a random sample of n observations from the exponential distribution with mean $1/\theta$ defined in Example C.4. It turns out that $x_{(1),n}$ has variance $1/(n\theta)^2$. Clearly, this variance also converges to zero, but, intuition suggests, faster than σ^2/n does. On the other hand, the sum of the integers from one to n , $S_n = n(n+1)/2$, obviously diverges as $n \rightarrow \infty$, albeit faster (one might expect) than the log of the likelihood function for the exponential distribution in Example 4.6, which is $\ln L(\theta) = n(\ln \theta - \theta \bar{x}_n)$. As a final example, consider the downward bias of the maximum likelihood estimator of the variance of the normal distribution, $c_n = (n-1)/n$, which is a constant that converges to one. (See Examples C.5.)

We will define the rate at which a sequence converges or diverges in terms of the order of the sequence.

DEFINITION D.14 Order n^δ
 A sequence c_n is of order n^δ , denoted $O(n^\delta)$, if and only if $\text{plim}(1/n^\delta)c_n$ is a finite nonzero constant.

DEFINITION D.15 Order less than n^δ

A sequence c_n is of order less than n^δ , denoted $o(n^\delta)$, if and only if $\text{plim}(1/n^\delta)c_n$ equals zero.

Thus, in our examples, γ_n^2 is $O(n^{-1})$, $\text{Var}[x_{(1),n}]$ is $O(n^{-2})$ and $o(n^{-1})$, S_n is $O(n^2)$ (δ equals +2 in this case), $\ln L(\theta)$ is $O(n)$ (δ equals +1), and c_n is $O(1)$ ($\delta = 0$). Important particular cases that we will encounter repeatedly in our work are sequences for which $\delta = 1$ or -1 .

The notion of order of a sequence is often of interest in econometrics in the context of the variance of an estimator. Thus, we see in Section C.3 that an important element of our strategy for forming an asymptotic distribution is that the variance of the limiting distribution of $\sqrt{n}(\bar{x}_n - \mu)/\sigma$ is $O(1)$. In Example D.10 the variance of m_2 is the sum of three terms that are $O(n^{-1})$, $O(n^{-2})$, and $O(n^{-3})$. The sum is $O(n^{-1})$, because $n \text{Var}[m_2]$ converges to $\mu_4 - \sigma^4$, the numerator of the first, or *leading term*, whereas the second and third terms converge to zero. This term is also the *dominant term* of the sequence. Finally, consider the two divergent examples in the preceding list. S_n is simply a deterministic function of n that explodes. However, $\ln L(\theta) = n \ln \theta - \theta \sum_i x_i$ is the sum of a constant that is $O(n)$ and a random variable with variance equal to n/θ . The random variable “diverges” in the sense that its variance grows without bound as n increases.

APPENDIX E



COMPUTATION AND OPTIMIZATION

E.1 INTRODUCTION

The computation of empirical estimates by econometricians involves using digital computers and software written either by the researchers themselves or by others.¹ It is also a surprisingly balanced mix of art and science. It is important for software users to be aware of how results are obtained, not only to understand routine computations, but also to be able to explain the occasional strange and contradictory results that do arise. This appendix will describe some of the basic elements of computing and a number of tools that are used by econometricians.² Sections E.2

¹It is one of the interesting aspects of the development of econometric methodology that the adoption of certain classes of techniques has proceeded in discrete jumps with the development of software. Noteworthy examples include the appearance, both around 1970, of G. K. Joreskog's LISREL [Joreskog and Sorbom (1981)] program, which spawned a still-growing industry in linear structural modeling, and TSP [Hall (1982)], which was among the first computer programs to accept symbolic representations of econometric models and which provided a significant advance in econometric practice with its LSQ procedure for systems of equations.

²This discussion is not intended to teach the reader how to write computer programs. For those who expect to do so, there are whole libraries of useful sources. Three very useful works are Kennedy and Gentle (1980), Abramovitz and Stegun (1971), and especially Press et al. (1986). The third of these provides a wealth of expertly written programs and a large amount of information about how to do computation efficiently and accurately. A recent survey of many areas of computation is Judd (1998).

and E.3 present issues that arise in generation of artificial data using Monte Carlo methods. Section E.4 describes bootstrapping, which is a method often used for estimating variances when analytical expressions cannot be obtained. Section E.5 then describes some techniques for computing certain integrals and derivatives that are recurrent in econometric applications. Section E.6 presents methods of optimization of functions. Some examples are also given in Section E.6.

E.2 DATA INPUT AND GENERATION

The data used in an econometric study can be broadly characterized as either real or simulated. “Real” data consist of actual measurements on some physical phenomenon such as the level of activity of an economy or the behavior of real consumers. For present purposes, the defining characteristic of such data is that they are generated outside the context of the empirical study and are gathered for the purpose of measuring some aspect of their real-world counterpart, such as an elasticity of some aspect of consumer behavior. The alternative is simulated data, produced by the analyst with a random number generator, usually for the purpose of studying the behavior of econometric estimators for which the statistical properties are unknown or impossible to derive. This section will consider a few aspects of the manipulation of data with a computer.

E.2.1 GENERATING PSEUDO-RANDOM NUMBERS

Monte Carlo methods and Monte Carlo studies of estimators are enjoying a flowering in the econometrics literature. In these studies, data are generated internally in the computer using **pseudo-random number generators**. These computer programs generate sequences of values that appear to be strings of draws from a specified probability distribution. There are many types of random number generators, but most take advantage of the inherent inaccuracy of the digital representation of real numbers. The method of generation is usually by the following steps:

0. Set a seed.
1. Update the seed by $\text{seed}_j = \text{seed}_{j-1} \times s$ value.
2. $x_j = \text{seed}_j \times x$ value.
3. Transform x_j if necessary, then move x_j to desired place in memory.
4. Return to Step 1, or exit if no additional values are needed.

Random number generators produce sequences of values that resemble strings of random draws from the specified distribution. In fact, the sequence of values produced by the preceding method is not truly random at all; it is a deterministic (Markov) chain of values. The set of 32 bits in the random value only appear random when subjected to certain tests. [See Press et al. (1986).] Since the series is, in fact, deterministic, at any point that a generator produces a value it has produced before, it must thereafter replicate the entire sequence. Since modern digital computers typically use 32-bit double precision variables to represent numbers, it follows that the longest string of values that this kind of generator can produce is $2^{32} - 1$ (about 2.1 billion). This length is the *period* of a random number generator. (A generator with a shorter period than this would be inefficient, since it is possible to achieve this period with some fairly simple algorithms.) Some improvements in the periodicity of a generator can be achieved by the method of *shuffling*. By this method, a set of, say, 128 values is maintained in an array. The random draw is used to select one of these 128 positions from which the draw is taken and then the value in the array is replaced with a draw from the generator. The period of the generator can also be increased by combining several generators. [See L’Ecuyer (1998) and Greene (2001).]

The deterministic nature of pseudo-random number generators is both a flaw and a virtue. Since many Monte Carlo studies require billions of draws, the finite period of any generator represents a nontrivial consideration. On the other hand, being able to reproduce a sequence of

values just by resetting the seed to its initial value allows the researcher to replicate a study.³ The seed itself can be a problem. It is known that certain seeds in particular generators will produce shorter series or series that do not pass randomness tests. For example, *congruential* generators of the sort discussed above should be started from odd seeds.

E.2.2 SAMPLING FROM A STANDARD UNIFORM POPULATION

When sampling from a standard uniform, $U[0, 1]$ population, the sequence is a kind of difference equation, since given the initial seed, x_j is ultimately a function of x_{j-1} . In most cases, the result at step 2 is a pseudodraw from the continuous uniform distribution in the range zero to one, which can then be transformed to a draw from another distribution by using the fundamental probability transformation.

E.2.3 SAMPLING FROM CONTINUOUS DISTRIBUTIONS

As soon as the sequence of $U[0, 1]$ values is obtained, there are several ways to transform them to a sample from the desired distribution. A common approach is to use the fundamental probability transform. For continuous distributions, this is done by treating the draw, F , as if F were $F(x)$, where F is the cdf of x . For example, if we desire draws from the exponential distribution with known θ , then $F(x) = 1 - \exp(-\theta x)$. The inverse transform is $x = (-1/\theta) \ln(1 - F)$. For example, for a draw of 0.4 with $\theta = 5$, the associated x would be 0.1022. One of the most common applications is the draws from the standard normal distribution, which is complicated because there is no closed form for $\Phi^{-1}(F)$. There are several ways to proceed. One is to approximate the inverse function. One well-known approximation is given in Abramovitz and Stegun (1971):

$$\Phi^{-1}(F) = x \approx T - \frac{c_0 + c_1 T + c_2 T^2}{1 + d_1 T + d_2 T^2 + d_3 T^3},$$

where $T = [\ln(1/H^2)]^{1/2}$ and $H = F$ if $F > 0.5$ and $1 - F$ otherwise. The sign is then reversed if $F < 0.5$. A second method is to transform the $U[0, 1]$ values directly to a standard normal value. The Box-Muller (1958) method is $z = (-2 \ln x_1)^{1/2} \cos(2\pi x_2)$, where x_1 and x_2 are two independent $U[0, 1]$ draws. A second $N[0, 1]$ draw can be obtained from the same two values by replacing \cos with \sin in the transformation. The Marsaglia-Bray (1964) generator, $z_i = x_i(-2/v) \ln v)^{1/2}$, where $x_i = 2w_i - 1$, w_i is a random draw from $U[0, 1]$ and $v = x_1^2 + x_2^2$, $i = 1, 2$, is often used as well. (The pair of draws must be rejected and redrawn if $v \geq 1$.) Sequences of draws from the standard normal distribution can be transformed easily into draws from other distributions by making use of the results in Section B.4. The square of a standard normal has chi-squared [1], and the sum of K chi-squareds is chi-squared [K]. From this relationship, it is possible to produce samples from the chi-squared, t , F , and beta distributions. A related problem is obtaining draws from the truncated normal distribution. An obviously inefficient (albeit effective) method of drawing values from the truncated normal $[\mu, \sigma^2]$ distribution in the range $[L, U]$ is simply to draw F from the $U[0, 1]$ distribution and transform it first to a standard normal variate as discussed previously and then to the $N[\mu, \sigma^2]$ variate by using $x = \mu + \sigma \Phi^{-1}(F)$. Finally, the value x is retained if it falls in the range $[L, U]$ and discarded otherwise. This method will require, on average, $1/[\Phi(U) - \Phi(L)]$ draws per observation, which could be substantial. A direct transformation that requires only one draw is as follows. Let $P_j = \Phi[(j - \mu)/\sigma]$, $j = L, U$. Then

$$x = \mu + \sigma \Phi^{-1}[P_L + F \times (P_U - P_L)]. \quad (\text{E-1})$$

³Current trends in the econometrics literature dictate that readers of empirical studies be able to replicate applied work. In Monte Carlo studies, at least in principle, data can be replicated efficiently merely by providing the random number generator and the seed.

E.2.4 SAMPLING FROM A MULTIVARIATE NORMAL POPULATION

A common application involves draws from a multivariate normal distribution with specified mean μ and covariance matrix Σ . To sample from this K -variate distribution, we begin with a draw, \mathbf{z} , from the K -variate standard normal distribution just by stacking K independent draws from the univariate standard normal distribution. Let \mathbf{T} be the square root of Σ such that $\mathbf{T}\mathbf{T}' = \Sigma$.⁴ The desired draw is then just $\mathbf{x} = \mu + \mathbf{Tz}$. A draw from a Wishart distribution of order K (which is a multivariate generalization of the chi-squared distribution) can be produced by computing $\mathbf{X}'\mathbf{M}\mathbf{X}$, where each row of \mathbf{X} is a draw from the multivariate normal distribution. Note that the Wishart is a matrix-variate random variable and that a sample of M draws from the Wishart distribution ultimately requires $M \times N \times K$ draws from the standard normal distribution, however generated.

E.2.5 SAMPLING FROM A DISCRETE POPULATION

Discrete distributions, such as the Poisson, present a different problem. There is no obvious inverse transformation for most of these. One inefficient, albeit unfortunately, unavoidable method for some distributions is to draw the F and then search sequentially for the discrete value that has cdf equal to or greater than F . This procedure makes intuitive sense, but it can involve a lot of computation. The **rejection method** described by Press et al. (1986, pp. 203–209) will be more efficient (although not more accurate) for some distributions.

E.2.6 THE GIBBS SAMPLER

The following problem is pervasive in Bayesian statistics and econometrics, although it has many applications in classical problems as well. (See Chapter 16 for an application.) We are given a joint density $f(x, y_1, y_2, \dots, y_K)$. We are interested in studying the characteristics, such as the mean, of the marginal distribution,

$$f(x) = \int_{y_K} \cdots \int_{y_1} f(x, y_1, y_2, \dots, y_K) dy_1 \cdots dy_K.$$

The direct approach, actually doing the integration to obtain the marginal density, may be infeasible or at least complicated enough to seem so. But the **Gibbs sampler**, a technique that has begun to enjoy a surge of activity in the econometrics literature, allows one to generate random draws from the marginal density $f(x)$ without having to compute it.^{5,6} The theory is presented in Casella and George (1992), among others. We will briefly sketch the mechanics of the technique and examine an application to a bivariate distribution.

Consider a two-variable case, $f(x, y)$ in which $f(x|y)$ and $f(y|x)$ are known. A “Gibbs sequence” of draws, $y_0, x_0, y_1, x_1, y_2, \dots, y_M, x_M$, is generated as follows. First, y_0 is specified “manually.” Then x_0 is obtained as a random draw from the population $f(x|y_0)$. Then y_1 is drawn

⁴In practice, this is usually done with a Cholesky decomposition in which \mathbf{T} is a lower triangular matrix. See Section B.7.11.

⁵A very readable introduction to the technique, on which we have based most of this discussion, is Casella and George (1992).

⁶The technique lends itself naturally to Bayesian applications, which is where most of the applications are to be found. See, for example, Albert and Chib (1993a,b), Chib (1992), Chib and Greenberg (1996), and Carlin and Chib (1995). There are classical applications as well, as surveyed in Tanner (1993) and Gelfand and Smith (1990).

from $f(y | x_0)$, and so on. The iteration is, generically, as follows.

1. Draw x_j from $f(x | y_j)$.
2. Draw y_{j+1} from $f(y | x_j)$.
3. Exit or return to step 1.

Note that the sequence of values are not independent; they are a **Markov chain**. If enough iterations are completed, the final observation x_M is a draw from $f(x)$ and likewise for y_M .⁷ Characteristics of the marginal distributions, such as the means, variances, or values of the densities, can then be studied just by using corresponding averages of the functions of the observations.

E.3 MONTE CARLO STUDIES

Simulated data generated by the methods of the previous section have various uses in econometrics. One of the more common applications is the derivation of the properties of estimators or in obtaining comparisons of the properties of estimators. For example, in time-series settings, most of the known results for characterizing the sampling distributions of estimators are asymptotic, large-sample results. But the typical time series is not very long, and descriptions that rely on T , the number of observations, going to infinity may not be very accurate. Exact, finite sample properties are usually intractable, however, which leaves the analyst with only the choice of learning about the behavior of the estimators experimentally.

In the typical application, one would either compare the properties of two or more estimators while holding the sampling conditions fixed or study how the properties of an estimator are affected by changing conditions such as the sample size or the value of an underlying parameter.

Example E.1 Monte Carlo Study of the Mean Versus the Median

In Example D.7, we compared the asymptotic distributions of the sample mean and the sample median in random sampling from the normal distribution. The basic result is that both estimators are consistent, but the mean is asymptotically more efficient by a factor of

$$\frac{\text{Asy. Var}[\text{Median}]}{\text{Asy. Var}[\text{Mean}]} = \frac{\pi}{2} = 1.5708.$$

This result is useful, but it does not tell which is the better estimator in small samples, nor does it suggest how the estimators would behave in some other distribution. It is known that the mean is affected by outlying observations whereas the median is not. The effect is averaged out in large samples, but the small sample behavior might be very different. To investigate the issue, we constructed the following experiment: We sampled 500 observations from the t distribution with d degrees of freedom by sampling $d + 1$ values from the standard normal distribution and then computing

$$t_{ir} = \frac{z_{ir,d+1}}{\sqrt{\frac{1}{d} \sum_{l=1}^d z_{ir,l}^2}}, \quad i = 1, \dots, 500, \quad r = 1, \dots, 100.$$

The t distribution with a low value of d was chosen because it has very thick tails and because large, outlying values have high probability. For each value of d , we generated $R = 100$ replications. For each of the 100 replications, we obtained the mean and median. Since both are unbiased, we compared the mean squared errors around the true expectations using

$$M_d = \frac{(1/R) \sum_{r=1}^R (\text{median}_r - 0)^2}{(1/R) \sum_{r=1}^R (\bar{x}_r - 0)^2}.$$

⁷Determining when to stop the sequence is an interesting and yet unsolved problem. See Casella and George (1992, pp. 172–173), Raftery and Lewis (1992), Roberts (1992), and Zellner and Min (1995).

We obtained ratios of 0.6761, 1.2779, and 1.3765 for $d = 3, 6,$ and $10,$ respectively. (You might want to repeat this experiment with different degrees of freedom.) These results agree with what intuition would suggest. As the degrees of freedom parameter increases, which brings the distribution closer to the normal distribution, the sample mean becomes more efficient—the ratio should approach its limiting value of 1.5708 as d increases. What might be surprising is the apparent overwhelming advantage of the median when the distribution is very nonnormal.

The preceding is a very small, straightforward application of the technique. In a typical study, there are many more parameters to be varied and more dimensions upon which the results are to be studied. One of the practical problems in this setting is how to organize the results. There is a tendency in Monte Carlo work to proliferate tables indiscriminately. It is incumbent on the analyst to collect the results in a fashion that is useful to the reader. For example, this requires some judgment on how finely one should vary the parameters of interest. One useful possibility that will often mimic the thought process of the reader is to collect the results of bivariate tables in carefully designed contour plots.

There are any number of situations in which Monte Carlo simulation offers the only method of learning about finite sample properties of estimators. Still, there are a number of problems with Monte Carlo studies. To achieve any level of generality, the number of parameters that must be varied and hence the amount of information that must be distilled can become enormous. Second, they are limited by the design of the experiments, so the results they produce are rarely generalizable. For our example, we may have learned something about the t distribution. But the results that would apply in other distributions remain to be described. And, unfortunately, real data will rarely conform to any specific distribution, so no matter how many other distributions we analyze, our results would still only be suggestive. In more general terms, this problem of **specificity** [Hendry (1984)] limits most Monte Carlo studies to quite narrow ranges of applicability. There are very few that have proved general enough to have provided a widely cited result.⁸

E.4 BOOTSTRAPPING AND THE JACKKNIFE

The technique of **bootstrapping** is used to obtain a description of the sampling properties of empirical estimators using the sample data themselves, rather than broad theoretical results.⁹ Suppose that $\hat{\theta}_n$ is an estimator of a parameter vector θ based on a sample $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$. An approximation to the statistical properties of $\hat{\theta}_n$ can be obtained by studying a sample of bootstrap estimators $\hat{\theta}^{(b)}_m, b = 1, \dots, B,$ obtained by sampling m observations, *with replacement*, from \mathbf{X} and recomputing $\hat{\theta}$ with each sample. After a total of B times, the desired sampling characteristic is computed from

$$\hat{\Theta} = [\hat{\theta}^{(1)}_m, \dots, \hat{\theta}^{(B)}_m].$$

For example, if it were known that the estimator were consistent and if n were reasonably large, then one might approximate the asymptotic covariance matrix of the estimator $\hat{\theta}$ by using

$$\text{Est. Asy. Var}[\hat{\theta}] = \frac{1}{B} \sum_{b=1}^B [\hat{\theta}^{(b)}_m - \hat{\theta}_n][\hat{\theta}^{(b)}_m - \hat{\theta}_n]'$$

⁸Two that have withstood the test of time are Griliches and Rao (1969) and Kmenta and Gilbert (1968).

⁹See Efron (1979) and Efron and Tibshirani (1993).

This technique was developed by Efron (1979) and has been appearing with increasing frequency in the applied econometrics literature. [See, for example, Veall (1987, 1992), Vinod (1993), and Vinod and Raj (1994).] An application of this technique to the least absolute deviations in the linear model is shown in the example below and in Chapter 5, and to a model of binary choice in Section 21.5.3.

Example E.2 Bootstrapping the Variance of the Median

As discussed earlier, there are few cases in which an exact expression for the sampling variance of the median are known. In the previous example, we examined the case of the median of a sample of 500 observations from the t distribution with 10 degrees of freedom. This is one of those cases in which there is no exact formula for the asymptotic variance of the median. However, we can use the bootstrap technique to estimate one empirically. (You might want to replicate this experiment. It is revisited in the exercises.) To demonstrate, consider the same data as used in the preceding example. We have a sample of 500 observations, for which we have computed the median, $-.00786$. We drew 100 samples of 500 with replacement from this sample and recomputed the median with each of these samples. The empirical square root of the mean squared deviation around this estimate of $-.00786$ was 0.056. In contrast, consider the same calculation for the mean. The sample mean is -0.07247 . The sample standard deviation is 1.08469, so the standard error of the mean is 0.04657. (The bootstrap estimate of the standard error of the mean was 0.052.) This agrees with our expectation in that the sample mean should generally be a more efficient estimator of the mean of the distribution in a large sample.

There is another approach we might take in this situation. Consider the regression model

$$y_i = \alpha + \varepsilon_i$$

where ε_i has a symmetric distribution with finite variance. As discussed in Chapter 8, the least absolute deviations estimator of the coefficient in this model is an estimator of the median (which equals the mean) of the distribution. So, this presents another estimator. Once again, the bootstrap estimator must be used to estimate the asymptotic variance of the estimator. Using the same data, we fit this regression model using the LAD estimator. The coefficient estimate is $-.05397$ with a bootstrap estimated standard error of 0.05872. The estimated standard error agrees with the earlier one. The difference in the estimated coefficient stems from the different computations—the regression estimate is the solution to a linear programming problem while the earlier estimate is the actual sample median.

The jackknife estimator is similar to the bootstrap estimator—Efron and Tibshirani argue that it is an approximation to the bootstrap. The jackknife procedure is carried out by redoing the estimation for $i = 1, \dots, n$ times, in each case leaving out the i th observation. The remaining computations are the same as for the bootstrap. The comparison of the two procedures is inconclusive, but Efron and Tibshirani suggest that by several criteria, the bootstrap is likely to be preferable. For a large sample, the simple advantage of the bootstrap estimator in terms of the amount of computation is likely to be compelling.

E.5 COMPUTATION IN ECONOMETRICS

This section will discuss some aspects of data manipulation.

E.5.1 COMPUTING INTEGRALS

One advantage of computers is their ability rapidly to compute approximations to complex functions such as logs and exponents. The basic functions, such as these, trigonometric functions, and so forth, are standard parts of the libraries of programs that accompany all scientific computing installations.¹⁰ But one of the very common applications that often requires some high-level creativity by econometricians is the evaluation of integrals that do not have simple closed forms and that do not typically exist in “system libraries.” We will consider several of these in this section. We will not go into detail on the nuts and bolts of how to compute integrals with a computer; rather, we will turn directly to the most common applications in econometrics.

E.5.2 THE STANDARD NORMAL CUMULATIVE DISTRIBUTION FUNCTION

The standard normal cumulative distribution function (cdf) is ubiquitous in econometric models. Yet this most homely of applications must be computed by approximation. There are a number of ways to do so.¹¹ Recall that what we desire is

$$\Phi(x) = \int_{-\infty}^x \phi(t) dt, \quad \text{where } \phi(t) = \frac{1}{\sqrt{2\pi}} e^{-t^2/2}.$$

One way to proceed is to use a Taylor series:

$$\Phi(x) \approx \sum_{i=0}^M \frac{1}{i!} \frac{d^i \Phi(x_0)}{dx_0^i} (x - x_0)^i.$$

The normal cdf has some advantages for this approach. First, the derivatives are simple and not integrals. Second, the function is **analytic**; as $M \rightarrow \infty$, the approximation converges to the true value. Third, the derivatives have a simple form; they are the **Hermite polynomials** and they can be computed by a simple recursion. The 0th term in the expansion above is $\Phi(x)$ evaluated at the expansion point. The first derivative of the cdf is the pdf, so the terms from 2 onward are the derivatives of $\phi(x)$, once again evaluated at x_0 . The derivatives of the standard normal pdf obey the recursion

$$\phi^i / \phi(x) = -x\phi^{i-1} / \phi(x) - (i-1)\phi^{i-2} / \phi(x),$$


where ϕ^i is $d^i \phi(x) / dx^i$. The zero and one terms in the sequence are one and $-x$. The next term is $x^2 - 1$, followed by $3x - x^3$ and $x^4 - 6x^2 + 3$, and so on. The approximation can be made more accurate by adding terms. Consider using a fifth-order Taylor series approximation around the point $x = 0$, where $\Phi(0) = 0.5$ and $\phi(0) = 0.3989423$. Evaluating the derivatives at zero and assembling the terms produces the approximation

$$\Phi(x) \approx \frac{1}{2} + 0.3989423[x - x^3/6 + x^5/40].$$

[Some of the terms (every other one, in fact) will conveniently drop out.] Figure E.1 shows the actual values (F) and approximate values (FA) over the range -2 to 2 . The figure shows two important points. First, the approximation is remarkably good over most of the range. Second, as is usually true for Taylor series approximations, the quality of the approximation deteriorates as one gets far from the expansion point.

¹⁰Of course, at some level, these must have been programmed as approximations by someone.

¹¹Many system libraries provide a related function, the *error function*, $\text{erf}(x) = (2/\sqrt{\pi}) \int_0^x e^{-t^2} dt$. If this is available, then the normal cdf can be obtained from $\Phi(x) = \frac{1}{2} + \frac{1}{2} \text{erf}(x/\sqrt{2})$, $x \geq 0$ and $\Phi(x) = 1 - \Phi(-x)$, $x \leq 0$.


FIGURE E.1 Approximation to Normal cdf.

Unfortunately, it is the tail areas of the standard normal distribution that are usually of interest, so the preceding is likely to be problematic. An alternative approach that is used much more often is a polynomial approximation reported by Abramovitz and Stegun (1971, p. 932):

$$\Phi(-|x|) = \phi(x) \sum_{i=1}^5 a_i t^i + \varepsilon(x), \quad \text{where } t = 1/[1 + a_0|x|].$$

(The complement is taken if x is positive.) The error of approximation is less than $\pm 7.5 \times 10^{-8}$ for all x . (Note that the error exceeds the function value at $|x| > 5.7$, so this is the operational limit of this approximation.)

E.5.3 THE GAMMA AND RELATED FUNCTIONS

The standard normal cdf is probably the most common application of numerical integration of a function in econometrics. Another very common application is the class of gamma functions. For positive constant P , the gamma function is

$$\Gamma(P) = \int_0^{\infty} t^{P-1} e^{-t} dt.$$

The gamma function obeys the recursion $\Gamma(P) = (P-1)\Gamma(P-1)$, so for integer values of P , $\Gamma(P) = (P-1)!$ This result suggests that the gamma function can be viewed as a generalization of the factorial function for noninteger values. Another convenient value is $\Gamma(\frac{1}{2}) = \sqrt{\pi}$. By making a change of variable, it can be shown that for positive constants a , c , and P ,

$$\int_0^{\infty} t^{P-1} e^{-at^c} dt = \int_0^{\infty} t^{-(P+1)} e^{-a/t^c} dt = \left(\frac{1}{c}\right) a^{-P/c} \Gamma\left(\frac{P}{c}\right).$$

As a generalization of the factorial function, the gamma function will usually overflow for the sorts of values of P that normally appear in applications. The log of the function should

normally be used instead. The function $\ln \Gamma(P)$ can be approximated remarkably accurately with only a handful of terms and is very easy to program. A number of approximations appear in the literature; they are generally modifications of **Sterling's approximation** to the factorial function $P! \approx (2\pi P)^{1/2} P^P e^{-P}$, so

$$\ln \Gamma(P) \approx (P - 0.5)\ln P - P + 0.5 \ln(2\pi) + C + \varepsilon(P),$$

where C is the correction term [see, e.g., Abramovitz and Stegun (1971, p. 257), Press et al. (1986, p. 157), or Rao (1973, p. 59)] and $\varepsilon(P)$ is the approximation error.¹²

The derivatives of the gamma function are

$$\frac{d^r \Gamma(P)}{dP^r} = \int_0^\infty (\ln P)^r t^{P-1} e^{-t} dt.$$

The first two derivatives of $\ln \Gamma(P)$ are denoted $\Psi(P) = \Gamma'/\Gamma$ and $\Psi'(P) = (\Gamma'' - \Gamma'^2)/\Gamma^2$ and are known as the **digamma** and **trigamma** functions.¹³ The beta function, denoted $\beta(a, b)$,

$$\beta(a, b) = \int_0^1 t^{a-1} (1-t)^{b-1} dt = \frac{\Gamma(a)\Gamma(b)}{\Gamma(a+b)},$$

is related.

E.5.4 APPROXIMATING INTEGRALS BY QUADRATURE

The digamma and trigamma functions, and the gamma function for noninteger values of P and values that are not integers plus $\frac{1}{2}$, do not exist in closed form and must be approximated. Most other applications will also involve integrals for which no simple computing function exists. The simplest approach to approximating

$$F(x) = \int_{L(x)}^{U(x)} f(t) dt$$

is likely to be a variant of Simpson's rule, or the trapezoid rule. For example, one approximation [see Press et al. (1986, p. 108)] is

$$F(x) \approx \Delta \left[\frac{1}{3} f_1 + \frac{4}{3} f_2 + \frac{2}{3} f_3 + \frac{4}{3} f_4 + \dots + \frac{2}{3} f_{N-2} + \frac{4}{3} f_{N-1} + \frac{1}{3} f_N \right],$$

where f_j is the function evaluated at N equally spaced points in $[L, U]$ including the endpoints and $\Delta = (U - L)/(N - 1)$. There are a number of problems with this method, most notably that it is difficult to obtain satisfactory accuracy with a moderate number of points.

Gaussian quadrature is a popular method of computing integrals. The general approach is to use an approximation of the form

$$\int_L^U W(x) f(x) dx \approx \sum_{j=1}^M w_j f(a_j),$$

where $W(x)$ is viewed as a "weighting" function for integrating $f(x)$, w_j is the **quadrature weight**, and a_j is the **quadrature abscissa**. Different weights and abscissas have been derived for several

¹²For example, one widely used formula is $C = z^{-1}/12 - z^{-3}/360 - z^{-5}/1260 + z^{-7}/1680 - q$, where $z = P$ and $q = 0$ if $P > 18$, or $z = P + J$ and $q = \ln[P(P+1)(P+2) \dots (P+J-1)]$, where $J = 18 - \text{INT}(P)$, if not. Note, in the approximation, we write $\Gamma(P) = (P!)/P +$ a correction.

¹³Tables of specific values for the gamma, digamma, and trigamma functions appear in Abramovitz and Stegun (1971). Most contemporary econometric programs have built-in functions for these common integrals, so the tables are not generally needed.

weighting functions. Two weighting functions common in econometrics are

$$W(x) = x^c e^{-x}, \quad x \in [0, \infty),$$

for which the computation is called **Gauss-Laguerre quadrature**, and

$$W(x) = e^{-x^2}, \quad x \in (-\infty, \infty),$$

for which the computation is called **Gauss-Hermite quadrature**. The theory for deriving weights and abscissas is given in Press et al. (1986, pp. 121-125). Tables of weights and abscissas for many values of M are given by Abramovitz and Stegun (1971).

E.5.5 MONTE CARLO INTEGRATION

The quadrature methods have proved very useful in empirical research and are surprisingly accurate even for a small number of points. There are integrals that defy treatment in this form, however. Recent work has brought many advances in techniques for evaluating complex integrals by using Monte Carlo methods rather than direct numerical approximations.

Example E.3 Fractional Moments of the Truncated Normal Distribution

The following function appeared in Greene's (1990) study of the stochastic frontier model:

$$h(r, \varepsilon) = \frac{\int_0^\infty z^r \frac{1}{\sigma} \phi \left[\frac{z - (-\varepsilon - \theta\sigma^2)}{\sigma} \right] dz}{\int_0^\infty \frac{1}{\sigma} \phi \left[\frac{z - (-\varepsilon - \theta\sigma^2)}{\sigma} \right] dz}$$

If we let $\mu = -(\varepsilon + \theta\sigma^2)$, we can show that the denominator is just $1 - \Phi(-\mu/\sigma)$, which is a value from the standard normal cdf. But the numerator is complex. It does not fit into a form that lends itself to Gauss-Laguerre integration because the exponential function involves both z and z^2 . An alternative form that has potential is obtained by making the change of variable to $w = (z - \mu)/\sigma$, which produces the right weighting function. But now the range of integration is not $-\infty$ to $+\infty$; it is $-\mu/\sigma$ to $+\infty$. There is another approach. Suppose that z is a random variable with $N[\mu, \sigma^2]$ distribution. Then the density of the truncated normal (at zero) distribution for z is

$$f(z|z > 0) = \frac{f(z)}{\text{Prob}[z > 0]} = \frac{\frac{1}{\sigma} \phi \left[\frac{z - \mu}{\sigma} \right]}{\int_0^\infty \frac{1}{\sigma} \phi \left[\frac{z - \mu}{\sigma} \right] dz}$$

This result is exactly the weighting function that appears in $h(r, \varepsilon)$, and the function being weighted is z^r . Therefore, $h(r, \varepsilon)$ is the expected value of z^r given that z is greater than zero. That is, $h(r, \varepsilon)$ is a possibly fractional moment—we do not restrict r to integer values—from the truncated (at zero) normal distribution when the untruncated variable has mean $-(\varepsilon + \theta\sigma^2)$ and variance σ^2 .

Now that we have identified the function, how do we compute it? We have already concluded that the familiar quadrature methods will not suffice. (And, no one has previously derived closed forms for the fractional moments of the normal distribution, truncated or not.) But, if we can draw a random sample of observations from this truncated normal distribution $\{z_i\}$, then the sample mean of $w_i = z_i^r$ will converge in probability (mean square) to its population counterpart. [The remaining detail is to establish that this expectation is finite, which it is for the truncated normal distribution; see Amemiya (1973).] Since we showed earlier how to draw observations from a truncated normal distribution, this remaining step is simple.

The preceding is a fairly straightforward application of **Monte Carlo integration**. In certain cases, an integral can be approximated by computing the sample average of a set of function values. The approach taken here was to interpret the integral as an expected value. We then had to establish that the mean we were computing was finite. Our basic statistical result for the behavior of sample means implies that with a large enough sample, we can approximate the integral as closely as we like.

The general approach is widely applicable in Bayesian econometrics and has begun to appear in classical statistics and econometrics as well.¹⁴ For direct application in a straightforward class of problems, we consider the general computation

$$F(x) = \int_L^U f(x)g(x) dx,$$

where $g(x)$ is a continuous function in the range $[L, U]$. (We could achieve greater generality by allowing more complicated functions, but for current purposes, we limit ourselves to straightforward cases.) Now, suppose that $g(x)$ is nonnegative in the entire range $[L, U]$. To normalize the weighting function, we suppose, as well, that

$$K = \int_L^U g(x) dx$$

is a known constant. Then $h(x) = g(x)/K$ is a probability density function in the range because it satisfies the axioms of probability.¹⁵ Let

$$H(x) = \int_L^x h(t) dt.$$

Then $H(L) = 0$, $H(U) = 1$, $H'(x) = h(x) > 0$, and so on. Then

$$\int_L^U f(x)g(x) dx = K \int_L^U f(x) \frac{g(x)}{K} dx = KE_{h(x)}[f(x)],$$

where we use the notation $E_{h(x)}[f(x)]$ to denote the expected value of the function $f(x)$ when x is drawn from the population with probability density function $h(x)$. We assume that this expected value is a finite constant. This set of results defines the computation. We now assume that we are able to draw (pseudo)random samples from the population $h(x)$. Since K is a known constant and the means of random samples are unbiased and consistent estimators of their population counterparts, the sample mean of the functions,

$$\hat{F}(x) = K \times \frac{1}{n} \sum_{i=1}^n f(x_i^h)$$

where x_i^h is a random draw from $h(\cdot)$, is a consistent estimator of the integral. [The claim is based on the Corollary to Theorem D.4, as the integral is equal to the expected value of the function $f(x)$.]

Suppose that the problem is well defined as above but that it is not possible to draw random samples from the population $h(\cdot)$. If there is another probability density function that resembles $h(\cdot)$, say, $I(x)$, then there may be an alternative strategy. We can rewrite our computation in the

¹⁴See Geweke (1986, 1988, 1989) for discussion and applications. A number of other references are given in Poirier (1995, p. 654).

¹⁵In many applications, K will already be part of the desired integral.

form

$$F(x) = K \int_L^U f(x) h(x) dx = K \int_L^U \left[\frac{f(x) h(x)}{I(x)} \right] I(x) dx.$$

Then we can interpret our integral as the expected value of $[f(x)h(x)]/I(x)$ when the population has density $I(x)$. The new density $I(x)$ is called an **importance function**. The same strategy works if certain fairly benign conditions are imposed on the importance function. [See Geweke (1989).] The range of variation is an important consideration, for example. If the range of x is, say, $(-\infty, +\infty)$ and we choose an importance function that is nonzero only in the range $(0, +\infty)$, then our strategy is likely to come to some difficulties.

Example E.4 Mean of a Lognormal Distribution

Consider computing the mean of a lognormal distribution. If $x \sim N[0, 1]$, then e^x is distributed as lognormal and has density

$$g(x) = \frac{1}{x\sqrt{2\pi}} e^{-1/2(\ln x)^2}, \quad x > 0$$

(see Section B.4.4). The expected value is $\int xg(x) dx$, so $f(x) = x$. Suppose that we did not know how to draw a random sample from this lognormal distribution (by just exponentiating our draws from the standard normal distribution) [or that the true mean is $\exp(0.5) = 1.649$]. Consider using a $\chi^2[1]$ as an importance function, instead. This chi-squared distribution is a gamma distribution with parameters $P = \lambda = \frac{1}{2}$ [see (B-39)], so

$$I(x) = \frac{1^{1/2}}{\Gamma\left(\frac{1}{2}\right)} x^{-1/2} e^{-(1/2)x}.$$

After a bit of manipulation, we find that

$$\frac{f(x)g(x)}{I(x)} = q(x) = e^{(1/2)[x - (\ln x)^2]} x^{1/2}.$$

Therefore, to estimate the mean of this lognormal distribution, we can draw a random sample of values x_i from the $\chi^2[1]$ distribution, which we can do by squaring the draws in a sample from the standard normal distribution, then computing the average of the sample of values, $q(x_i)$.

We carried out this experiment with 1000 draws from a standard normal distribution. The mean of our sample was 1.6974, compared with a true mean of 1.649, so the error was less than 3 percent.

E.5.6 MULTIVARIATE NORMAL PROBABILITIES AND SIMULATED MOMENTS

The computation of bivariate normal probabilities requires a large amount of computing effort. Quadrature methods have been developed for trivariate probabilities as well, but the amount of computing effort needed at this level is enormous. For integrals of level greater than three, satisfactory (in terms of speed and accuracy) direct approximations remain to be developed. Our work thus far does suggest an alternative approach. Suppose that \mathbf{x} has a K -variate normal distribution with mean vector $\mathbf{0}$ and covariance matrix Σ . (No generality is sacrificed by the assumption of a zero mean, since we could just subtract a nonzero mean from the random vector wherever it appears in any result.) We wish to compute the K -variate probability, $\text{Prob}[a_1 < x_1 < b_1, a_2 < x_2 < b_2, \dots, a_K < x_K < b_K]$. Our Monte Carlo integration technique is well suited for this well-defined problem. As a first approach, consider sampling R observations, $\mathbf{x}_r, r = 1, \dots, R$,

from this multivariate normal distribution, using the method described in Section E.2. Now, define

$$d_r = \mathbf{1}[a_1 < x_{r1} < b_1, a_2 < x_{r2} < b_2, \dots, a_K < x_{rK} < b_K].$$

(That is, $d_r = 1$ if the condition is true and 0 otherwise.) Based on our earlier results, it follows that

$$\text{plim } \bar{d} = \text{plim } \frac{1}{R} \sum_{r=1}^R d_r = \text{Prob}[a_1 < x_1 < b_1, a_2 < x_2 < b_2, \dots, a_K < x_K < b_K].^{16}$$

This method is valid in principle, but in practice it has proved to be unsatisfactory for several reasons. For large-order problems, it requires an enormous number of draws from the distribution to give reasonable accuracy. Also, even with large numbers of draws, it appears to be problematic when the desired tail area is very small. Nonetheless, the idea is sound, and recent research has built on this idea to produce some quite accurate and efficient simulation methods for this computation. A survey of the methods is given in McFadden and Ruud (1994).¹⁷

Among the simulation methods examined in the survey, the GHK smooth recursive simulator appears to be the most accurate.¹⁸ The method is surprisingly simple. The general approach uses

$$\text{Prob}[a_1 < x_1 < b_1, a_2 < x_2 < b_2, \dots, a_K < x_K < b_K] \approx \frac{1}{R} \sum_{r=1}^R \prod_{k=1}^K Q_{rk},$$

where Q_{rk} are easily computed univariate probabilities. The probabilities Q_{rk} are computed according to the following recursion: We first factor Σ using the Cholesky factorization $\Sigma = \mathbf{L}\mathbf{L}'$, where \mathbf{L} is a lower triangular matrix (see Section 2.7.11). The elements of \mathbf{L} are l_{km} , where $l_{km} = 0$ if $m > k$. Then we begin the recursion with

$$Q_{r1} = \Phi(b_1/l_{11}) - \Phi(a_1/l_{11}).$$

Note that $l_{11} = \sigma_{11}$, so this is just the marginal probability, $\text{Prob}[a_1 < x_1 < b_1]$. Now, generate a random observation ε_{r1} from the truncated standard normal distribution in the range

$$A_{r1} \text{ to } B_{r1} = a_1/l_{11} \text{ to } b_1/l_{11}.$$

(Note, again, that the range is standardized since $l_{11} = \sigma_{11}$.) The draw can be obtained from a $U[0, 1]$ observation using (E-1). For steps $k = 2, \dots, K$, compute

$$A_{rk} = \left[a_k - \sum_{m=1}^{k-1} l_{km} \varepsilon_{rm} \right] / l_{kk},$$

$$B_{rk} = \left[b_k - \sum_{m=1}^{k-1} l_{km} \varepsilon_{rm} \right] / l_{kk}.$$

Then

$$Q_{rk} = \Phi(B_{rk}) - \Phi(A_{rk}),$$

¹⁶This method was suggested by Lerman and Manski (1981).

¹⁷A symposium on the topic of simulation methods appears in *Review of Economic Statistics*, Vol. 76, November 1994. See, especially, McFadden and Ruud (1994), Stern (1994), Geweke, Keane, and Runkle (1994), and Breslaw (1994). See, as well, Gourieroux and Monfort (1996).

¹⁸See Geweke (1989), Hajivassiliou (1990), and Keane (1994). Details on the properties of the simulator are given in Börsch-Supan and Hajivassiliou (1990).

and in preparation for the next step in the recursion, we generate a random draw from the truncated standard normal distribution in the range A_{rk} to B_{rk} . This process is replicated R times, and the estimated probability is the sample average of the simulated probabilities.

The GHK simulator has been found to be impressively fast and accurate for fairly moderate numbers of replications. Its main usage has been in computing functions and derivatives for maximum likelihood estimation of models that involve multivariate normal integrals. We will revisit this in the context of the method of simulated moments when we examine the probit model in Chapter 21.

E.5.7 COMPUTING DERIVATIVES

For certain functions, the programming of derivatives may be quite difficult. Numeric approximations can be used, although it should be borne in mind that analytic derivatives obtained by formally differentiating the functions involved are to be preferred. First derivatives can be approximated by using

$$\frac{\partial F(\theta)}{\partial \theta_i} \approx \frac{F(\dots \theta_i + \varepsilon \dots) - F(\dots \theta_i - \varepsilon \dots)}{2\varepsilon}$$

The choice of ε is a remaining problem. Extensive discussion may be found in Quandt (1983).

There are three drawbacks to this means of computing derivatives compared with using the analytic derivatives. A possible major consideration is that it may substantially increase the amount of computation needed to obtain a function and its gradient. In particular, $K + 1$ function evaluations (the criterion and K derivatives) are replaced with $2K + 1$ functions. The latter may be more burdensome than the former, depending on the complexity of the partial derivatives compared with the function itself. The comparison will depend on the application. But in most settings, careful programming that avoids superfluous or redundant calculation can make the advantage of the analytic derivatives substantial. Second, the choice of ε can be problematic. If it is chosen too large, then the approximation will be inaccurate. If it is chosen too small, then there may be insufficient variation in the function to produce a good estimate of the derivative. A compromise that is likely to be effective is to compute ε_i separately for each parameter, as in

$$\varepsilon_i = \text{Max}[\alpha|\theta_i|, \gamma]$$

[see Goldfeld and Quandt (1971)]. The values α and γ should be relatively small, such as 10^{-5} . Third, although numeric derivatives computed in this fashion are likely to be reasonably accurate, in a sum of a large number of terms, say several thousand, enough approximation error can accumulate to cause the numerical derivatives to differ significantly from their analytic counterparts. Second derivatives can also be computed numerically. In addition to the preceding problems, however, it is generally not possible to ensure negative definiteness of a Hessian computed in this manner. Unless the choice of ε is made extremely carefully, an indefinite matrix is a possibility. In general, the use of numeric derivatives should be avoided if the analytic derivatives are available.

E.6 OPTIMIZATION

Nonlinear optimization (e.g., maximizing log-likelihood functions) is an intriguing practical problem. Theory provides few hard and fast rules, and there are relatively few cases in which it is obvious how to proceed. This section introduces some of the terminology and underlying theory

of nonlinear optimization.¹⁹ We begin with a general discussion on how to search for a solution to a nonlinear optimization problem and describe some specific commonly used methods. We then consider some practical problems that arise in optimization. An example is given in the final section.

Consider maximizing the quadratic function

$$F(\boldsymbol{\theta}) = a + \mathbf{b}'\boldsymbol{\theta} - \frac{1}{2}\boldsymbol{\theta}'\mathbf{C}\boldsymbol{\theta},$$

where \mathbf{C} is a positive definite matrix. The first-order condition for a maximum is

$$\frac{\partial F(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} = \mathbf{b} - \mathbf{C}\boldsymbol{\theta} = \mathbf{0}. \quad (\text{E-2})$$

This set of *linear* equations has the unique solution

$$\boldsymbol{\theta} = \mathbf{C}^{-1}\mathbf{b}. \quad (\text{E-3})$$

This is a linear optimization problem. Note that it has a **closed-form solution**; for any a , \mathbf{b} , and \mathbf{C} , the solution can be computed directly.²⁰ In the more typical situation,

$$\frac{\partial F(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} = \mathbf{0} \quad (\text{E-4})$$

is a set of nonlinear equations that cannot be solved explicitly for $\boldsymbol{\theta}$.²¹ The techniques considered in this section provide systematic means of searching for a solution.

We now consider the general problem of maximizing a function of several variables:

$$\text{maximize}_{\boldsymbol{\theta}} F(\boldsymbol{\theta}), \quad (\text{E-5})$$

where $F(\boldsymbol{\theta})$ may be a log-likelihood or some other function. Minimization of $F(\boldsymbol{\theta})$ is handled by maximizing $-F(\boldsymbol{\theta})$. Two special cases are

$$F(\boldsymbol{\theta}) = \sum_{i=1}^n f_i(\boldsymbol{\theta}), \quad (\text{E-6})$$

which is typical for maximum likelihood problems, and the **least squares problem**,²²

$$f_i(\boldsymbol{\theta}) = -(y_i - f(\mathbf{x}_i, \boldsymbol{\theta}))^2. \quad (\text{E-7})$$

We treated the nonlinear least squares problem in detail in Chapter 9. An obvious way to search for the $\boldsymbol{\theta}$ that maximizes $F(\boldsymbol{\theta})$ is by trial and error. If $\boldsymbol{\theta}$ has only a single element and it is known approximately where the optimum will be found, then a **grid search** will be a feasible strategy. An example is a common time-series problem in which a one-dimensional search for a correlation coefficient is made in the interval $(-1, 1)$. The grid search can proceed in the obvious fashion—that is, $\dots, -0.1, 0, 0.1, 0.2, \dots$, then $\hat{\theta}_{\max} - 0.1$ to $\hat{\theta}_{\max} + 0.1$ in increments of 0.01, and so on—until the desired precision is achieved.²³ If $\boldsymbol{\theta}$ contains more than one parameter, then a grid search

¹⁹There are numerous excellent references that offer a more complete exposition. Among these are Quandt (1983), Bazzara and Shetty (1979), Fletcher (1980), and Judd (1998).

²⁰Notice that the constant a is irrelevant to the solution. Many maximum likelihood problems are presented with the preface “neglecting an irrelevant constant.” For example, the log-likelihood for the normal linear regression model contains a term $-(n/2) \ln(2\pi)$ —that can be discarded.

²¹See, for example, the normal equations for the nonlinear least squares estimators of Chapter 9.

²²Least squares is, of course, a minimization problem. The negative of the criterion is used to maintain consistency with the general formulation.

²³There are more efficient methods of carrying out a one-dimensional search, for example, the **golden section** method. See Press et al. (1986, Chap. 10).

is likely to be extremely costly, particularly if little is known about the parameter vector at the outset. Nonetheless, relatively efficient methods have been devised. Quandt (1983) and Fletcher (1980) contain further details.

There are also systematic, derivative-free methods of searching for a function optimum that resemble in some respects the algorithms that we will examine in the next section. The **downhill simplex** (and other simplex) methods²⁴ have been found to be very fast and effective for some problems. A recent entry in the econometric literature is the method of **simulated annealing**.²⁵ These derivative-free methods, particularly the latter, are often very effective in problems with many variables in the objective function, but they usually require far more function evaluations than the methods based on derivatives that are considered below. Since the problems typically analyzed in econometrics involve relatively few parameters but often quite complex functions involving large numbers of terms in a summation, on balance, the gradient methods are usually going to be preferable.²⁶

E.6.1 ALGORITHMS

A more effective means of solving most nonlinear maximization problems is by an **iterative algorithm**:

Beginning from initial value θ_0 , at entry to iteration t , if θ_t is not the optimal value for θ , compute direction vector Δ_t , step size λ_t , then

$$\theta_{t+1} = \theta_t + \lambda_t \Delta_t. \quad (\text{E-8})$$

Figure E.2 illustrates the structure of an iteration for a hypothetical function of two variables. The direction vector Δ_t is shown in the figure with θ_t . The dashed line is the set of points $\theta_t + \lambda_t \Delta_t$. Different values of λ_t lead to different contours; for this θ_t and Δ_t , the best value of λ_t is about 0.5.

Notice in Figure E.2 that for a given direction vector Δ_t and current parameter vector θ_t , a secondary optimization is required to find the best λ_t . Translating from Figure E.2, we obtain the form of this problem as shown in Figure E.3. This subsidiary search is called a **line search**, as we search along the line $\theta_t + \lambda_t \Delta_t$ for the optimal value of $F(\cdot)$. The formal solution to the line search problem would be the λ_t that satisfies

$$\frac{\partial F(\theta_t + \lambda_t \Delta_t)}{\partial \lambda_t} = \mathbf{g}(\theta_t + \lambda_t \Delta_t)' \Delta_t = 0, \quad (\text{E-9})$$

where \mathbf{g} is the vector of partial derivatives of $F(\cdot)$ evaluated at $\theta_t + \lambda_t \Delta_t$. In general, this problem will also be a nonlinear one. In most cases, adding a formal search for λ_t will be too expensive, as well as unnecessary. Some approximate or ad hoc method will usually be chosen. It is worth emphasizing that finding the λ_t that maximizes $F(\theta_t + \lambda_t \Delta_t)$ at a given iteration does not generally lead to the overall solution in that iteration. This situation is clear in Figure E.3, where the optimal value of λ_t leads to $F(\cdot) = 2.0$, at which point we reenter the iteration.

E.6.2 GRADIENT METHODS

The most commonly used algorithms are **gradient methods**, in which

$$\Delta_t = \mathbf{W}_t \mathbf{g}_t, \quad (\text{E-10})$$

²⁴See Nelder and Mead (1965) and Press et al. (1986).

²⁵See Goffe, Ferrier, and Rodgers (1994) and Press et al. (1986, pp. 326–334).

²⁶Goffe, Ferrier, and Rodgers (1994) did find that the method of simulated annealing was quite adept at finding the best among multiple solutions. This problem is frequent for derivative-based methods, because they usually have no method of distinguishing between a local optimum and a global one.




FIGURE E.2 Iteration.




FIGURE E.3 Line Search.

where \mathbf{W}_t is a positive definite matrix and \mathbf{g}_t is the **gradient** of $F(\boldsymbol{\theta}_t)$:

$$\mathbf{g}_t = \mathbf{g}(\boldsymbol{\theta}_t) = \frac{\partial F(\boldsymbol{\theta}_t)}{\partial \boldsymbol{\theta}_t}. \tag{E-11}$$

These methods are motivated partly by the following. Consider a linear Taylor series approximation to $F(\boldsymbol{\theta}_t + \lambda_t \boldsymbol{\Delta}_t)$ around $\lambda_t = 0$:

$$F(\boldsymbol{\theta}_t + \lambda_t \boldsymbol{\Delta}_t) \simeq F(\boldsymbol{\theta}_t) + \lambda_t \mathbf{g}(\boldsymbol{\theta}_t)' \boldsymbol{\Delta}_t. \tag{E-12}$$

Let $F(\boldsymbol{\theta}_t + \lambda_t \boldsymbol{\Delta}_t)$ equal F_{t+1} . Then,

$$F_{t+1} - F_t \simeq \lambda_t \mathbf{g}'_t \boldsymbol{\Delta}_t.$$

If $\boldsymbol{\Delta}_t = \mathbf{W}_t \mathbf{g}_t$, then

$$F_{t+1} - F_t \simeq \lambda_t \mathbf{g}'_t \mathbf{W}_t \mathbf{g}_t.$$

If \mathbf{g}_t is not $\mathbf{0}$ and λ_t is small enough, then $F_{t+1} - F_t$ must be positive. Thus, if $F(\boldsymbol{\theta})$ is not already at its maximum, then we can always find a step size such that a gradient-type iteration will lead to an increase in the function. (Recall that \mathbf{W}_t is assumed to be positive definite.)

In the following, we will omit the iteration index t , except where it is necessary to distinguish one vector from another. The following are some commonly used algorithms.²⁷

Steepest Ascent The simplest algorithm to employ is the **steepest ascent** method, which uses

$$\mathbf{W} = \mathbf{I} \text{ so that } \boldsymbol{\Delta} = \mathbf{g}. \tag{E-13}$$

As its name implies, the direction is the one of greatest increase of $F(\cdot)$. Another virtue is that the line search has a straightforward solution; at least near the maximum, the optimal λ is

$$\lambda = \frac{-\mathbf{g}'\mathbf{g}}{\mathbf{g}'\mathbf{H}\mathbf{g}}, \tag{E-14}$$

where

$$\mathbf{H} = \frac{\partial^2 F(\boldsymbol{\theta})}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}'}$$

Therefore, the steepest ascent iteration is

$$\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t - \frac{\mathbf{g}'_t \mathbf{g}_t}{\mathbf{g}'_t \mathbf{H}_t \mathbf{g}_t} \mathbf{g}_t. \tag{E-15}$$

Computation of the second derivatives matrix may be extremely burdensome. Also, if \mathbf{H}_t is not negative definite, which is likely if $\boldsymbol{\theta}_t$ is far from the maximum, the iteration may diverge. A systematic line search can bypass this problem. This algorithm usually converges very slowly, however, so other techniques are usually used.

Newton's Method The template for most gradient methods in common use is Newton's method. The basis for **Newton's method** is a linear Taylor series approximation. Expanding the first-order conditions,

$$\frac{\partial F(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} = \mathbf{0},$$

²⁷ A more extensive catalog may be found in Judge et al. (1985, Appendix B). Those mentioned here are some of the more commonly used ones and are chosen primarily because they illustrate many of the important aspects of nonlinear optimization.

equation by equation, in a linear Taylor series around an arbitrary θ^0 yields

$$\frac{\partial F(\theta)}{\partial \theta} \simeq \mathbf{g}^0 + \mathbf{H}^0(\theta - \theta^0) = \mathbf{0}, \tag{E-16}$$

where the superscript indicates that the term is evaluated at θ^0 . Solving for θ and then equating θ to θ_{t+1} and θ^0 to θ_t , we obtain the iteration

$$\theta_{t+1} = \theta_t - \mathbf{H}_t^{-1} \mathbf{g}_t. \tag{E-17}$$

Thus, for Newton's method,

$$\mathbf{W} = -\mathbf{H}^{-1}, \quad \Delta = -\mathbf{H}^{-1} \mathbf{g}, \quad \lambda = 1. \tag{E-18}$$

Newton's method will converge very rapidly in many problems. If the function is quadratic, then this method will reach the optimum in one iteration from any starting point. If the criterion function is globally concave, as it is in a number of problems that we shall examine in this text, then it is probably the best algorithm available. This method is very well suited to maximum likelihood estimation.

Alternatives to Newton's Method Newton's method is very effective in some settings, but it can perform very poorly in others. If the function is not approximately quadratic or if the current estimate is very far from the maximum, then it can cause wide swings in the estimates and even fail to converge at all. A number of algorithms have been devised to improve upon Newton's method. An obvious one is to include a line search at each iteration rather than use $\lambda = 1$. Two problems remain, however. At points distant from the optimum, the second derivatives matrix may not be negative definite, and, in any event, the computational burden of computing \mathbf{H} may be excessive.

The **quadratic hill-climbing method** proposed by Goldfeld, Quandt, and Trotter (1966) deals directly with the first of these problems. In any iteration, if \mathbf{H} is not negative definite, then it is replaced with

$$\mathbf{H}_\alpha = \mathbf{H} - \alpha \mathbf{I}, \tag{E-19}$$

where α is a positive number chosen large enough to ensure the negative definiteness of \mathbf{H}_α . Another suggestion is that of Greenstadt (1967), which uses, at every iteration,

$$\mathbf{H}_\pi = - \sum_{i=1}^n |\pi_i| \mathbf{c}_i \mathbf{c}_i', \tag{E-20}$$

where π_i is the i th characteristic root of \mathbf{H} and \mathbf{c}_i is its associated characteristic vector. Other proposals have been made to ensure the negative definiteness of the required matrix at each iteration.²⁸

Quasi-Newton Methods: Davidon-Fletcher-Powell A very effective class of algorithms has been developed that eliminates second derivatives altogether and has excellent convergence properties, even for ill-behaved problems. These are the **quasi-Newton methods**, which form

$$\mathbf{W}_{t+1} = \mathbf{W}_t + \mathbf{E}_t,$$

where \mathbf{E}_t is a positive definite matrix.²⁹ As long as \mathbf{W}_0 is positive definite— \mathbf{I} is commonly used— \mathbf{W}_t will be positive definite at every iteration. In the **Davidon-Fletcher-Powell (DFP) method**,

²⁸See, for example, Goldfeld and Quandt (1971).

²⁹See Fletcher (1980).

after a sufficient number of iterations, \mathbf{W}_{t+1} will be an approximation to $-\mathbf{H}^{-1}$. Let

$$\delta_t = \lambda_t \Delta_t \quad \text{and} \quad \gamma_t = \mathbf{g}(\theta_{t+1}) - \mathbf{g}(\theta_t). \quad (\text{E-21})$$

The DFP **variable metric algorithm** uses

$$\mathbf{W}_{t+1} = \mathbf{W}_t + \frac{\delta_t \delta_t'}{\delta_t' \gamma_t} + \frac{\mathbf{W}_t \gamma_t \gamma_t' \mathbf{W}_t}{\gamma_t' \mathbf{W}_t \gamma_t}. \quad (\text{E-22})$$

Notice that in the DFP algorithm, the change in the first derivative vector is used in \mathbf{W} ; an estimate of the inverse of the second derivatives matrix is being accumulated.

The variable metric algorithms are those that update \mathbf{W} at each iteration while preserving its definiteness. For the DFP method, the accumulation of \mathbf{W}_{t+1} is of the form

$$\mathbf{W}_{t+1} = \mathbf{W}_t + \mathbf{a}\mathbf{a}' + \mathbf{b}\mathbf{b}' = \mathbf{W}_t + [\mathbf{a} \quad \mathbf{b}][\mathbf{a} \quad \mathbf{b}]'$$

The two-column matrix $[\mathbf{a} \quad \mathbf{b}]$ will have rank two; hence, DFP is called a **rank two update** or **rank two correction**. The **Broyden-Fletcher-Goldfarb-Shanno (BFGS)** method is a rank three correction that subtracts $v\mathbf{d}\mathbf{d}'$ from the DFP update, where $v = (\gamma_t' \mathbf{W}_t \gamma_t)$ and

$$\mathbf{d}_t = \left(\frac{1}{\delta_t' \gamma_t} \right) \delta_t - \left(\frac{1}{\gamma_t' \mathbf{W}_t \gamma_t} \right) \mathbf{W}_t \gamma_t.$$

There is some evidence that this method is more efficient than DFP. Other methods, such as **Broyden's method**, involve a rank one correction instead. Any method that is of the form

$$\mathbf{W}_{t+1} = \mathbf{W}_t + \mathbf{Q}\mathbf{Q}'$$

will preserve the definiteness of \mathbf{W} , regardless of the number of columns in \mathbf{Q} .

The DFP and BFGS algorithms are extremely effective and are among the most widely used of the gradient methods. An important practical consideration to keep in mind is that although \mathbf{W}_t accumulates an estimate of the negative inverse of the second derivatives matrix for both algorithms, in maximum likelihood problems it rarely converges to a very good estimate of the covariance matrix of the estimator and should generally not be used as one.

E.6.3 ASPECTS OF MAXIMUM LIKELIHOOD ESTIMATION

Newton's method is often used for maximum likelihood problems. For solving a maximum likelihood problem, the **method of scoring** replaces \mathbf{H} with

$$\hat{\mathbf{H}} = E[\mathbf{H}(\theta)], \quad (\text{E-23})$$

which will be recognized as the asymptotic covariance of the maximum likelihood estimator. There is some evidence that where it can be used, this method performs better than Newton's method. The exact form of the expectation of the Hessian of the log likelihood is rarely known, however.³⁰ Newton's method, which uses actual instead of expected second derivatives, is generally used instead.

One-Step Estimation A convenient variant of Newton's method is the **one-step maximum likelihood estimator**. It has been shown that if θ^0 is any consistent initial estimator of θ and \mathbf{H}^* is \mathbf{H} , $\hat{\mathbf{H}}$, or any other asymptotically equivalent estimator of $\text{Var}[\mathbf{g}(\hat{\theta}_{\text{MLE}})]$, then

$$\theta^1 = \theta^0 - (\mathbf{H}^*)^{-1} \mathbf{g}^0 \quad (\text{E-24})$$

³⁰Amemiya (1981) provides a number of examples.

is an estimator of θ that has the same asymptotic properties as the maximum likelihood estimator.³¹ (Note that it is *not* the maximum likelihood estimator. As such, for example, it should not be used as the basis for likelihood ratio tests.)

Covariance Matrix Estimation In computing maximum likelihood estimators, a commonly used method of estimating \mathbf{H} simultaneously simplifies the calculation of \mathbf{W} and solves the occasional problem of indefiniteness of the Hessian. The method of Berndt et al. (1974) replaces \mathbf{W} with

$$\hat{\mathbf{W}} = \left[\sum_{i=1}^n \mathbf{g}_i \mathbf{g}'_i \right]^{-1} = (\mathbf{G}'\mathbf{G})^{-1}, \tag{E-25}$$

where

$$\mathbf{g}_i = \frac{\partial \ln f(y_i | \mathbf{x}_i, \theta)}{\partial \theta}. \tag{E-26}$$

Then, \mathbf{G} is the $n \times K$ matrix with i th row equal to \mathbf{g}'_i . Although $\hat{\mathbf{W}}$ and other suggested estimators of $(-\mathbf{H})^{-1}$ are asymptotically equivalent, $\hat{\mathbf{W}}$ has the additional virtues that it is always nonnegative definite, and it is only necessary to differentiate the log-likelihood once to compute it.

The Lagrange Multiplier Statistic The use of $\hat{\mathbf{W}}$ as an estimator of $(-\mathbf{H})^{-1}$ brings another intriguing convenience in maximum likelihood estimation. When testing restrictions on parameters estimated by maximum likelihood, one approach is to use the **Lagrange multiplier** statistic. We will examine this test at length at various points in this book, so we need only sketch it briefly here. The logic of the LM test is as follows. The gradient $\mathbf{g}(\theta)$ of the log-likelihood function equals $\mathbf{0}$ at the unrestricted maximum likelihood estimators (that is, at least to within the precision of the computer program in use). If $\hat{\theta}_r$ is an MLE that is computed subject to some restrictions on θ , then we know that $\mathbf{g}(\hat{\theta}_r) \neq \mathbf{0}$. The LM test is used to test whether, at $\hat{\theta}_r$, \mathbf{g}_r is *significantly* different from $\mathbf{0}$ or whether the deviation of \mathbf{g}_r from $\mathbf{0}$ can be viewed as sampling variation. The covariance matrix of the gradient of the log-likelihood is $-\mathbf{H}$, so the Wald statistic for testing this hypothesis is $W = \mathbf{g}'(-\mathbf{H})^{-1}\mathbf{g}$. Now, suppose that we use $\hat{\mathbf{W}}$ to estimate $-\mathbf{H}^{-1}$. Let \mathbf{G} be the $n \times K$ matrix with i th row equal to \mathbf{g}'_i , and let \mathbf{i} denote an $n \times 1$ column of ones. Then the LM statistic can be computed as

$$\text{LM} = \mathbf{i}'\mathbf{G}(\mathbf{G}'\mathbf{G})^{-1}\mathbf{G}'\mathbf{i}.$$

Since $\mathbf{i}'\mathbf{i} = n$,

$$\text{LM} = n[\mathbf{i}'\mathbf{G}(\mathbf{G}'\mathbf{G})^{-1}\mathbf{G}'\mathbf{i}/n] = nR_r^2,$$

where R_r^2 is the *uncentered* R^2 in a regression of a column of ones on the derivatives of the log-likelihood function.

The Concentrated Log-Likelihood Many problems in maximum likelihood estimation can be formulated in terms of a partitioning of the parameter vector $\theta = [\theta_1, \theta_2]$ such that at the solution to the optimization problem, $\theta_{2,ML}$, can be written as an explicit function of $\theta_{1,ML}$. When the solution to the likelihood equation for θ_2 produces

$$\theta_{2,ML} = \mathbf{t}(\theta_{1,ML}),$$

then, if it is convenient, we may “concentrate” the log-likelihood function by writing

$$F^*(\theta_1, \theta_2) = F[\theta_1, \mathbf{t}(\theta_1)] = F_c(\theta_1).$$

³¹See, for example, Rao (1973).

The unrestricted solution to the problem $\text{Max}_{\theta_1} F_c(\theta_1)$ provides the full solution to the optimization problem. Once the optimizing value of θ_1 is obtained, the optimizing value of θ_2 is simply $\mathbf{t}(\hat{\theta}_{1,ML})$. Note that $F^*(\theta_1, \theta_2)$ is a subset of the set of values of the log-likelihood function, namely those values at which the second parameter vector satisfies the first-order conditions.³²

E.6.4 OPTIMIZATION WITH CONSTRAINTS

Occasionally, some of or all the parameters of a model are constrained, for example, to be positive in the case of a variance or to be in a certain range, such as a correlation coefficient. Optimization subject to constraints is often yet another art form. The elaborate literature on the general problem provides some guidance—see, for example, Appendix B in Judge et al. (1985)—but applications still, as often as not, require some creativity on the part of the analyst. In this section, we will examine a few of the most common forms of constrained optimization as they arise in econometrics.

Parametric constraints typically come in two forms, which may occur simultaneously in a problem. Equality constraints can be written $\mathbf{c}(\theta) = \mathbf{0}$, where $c_j(\theta)$ is a continuous and differentiable function. Typical applications include linear constraints on slope vectors, such as a requirement that a set of elasticities in a log-linear model add to one; exclusion restrictions, which are often cast in the form of interesting hypotheses about whether or not a variable should appear in a model (i.e., whether a coefficient is zero or not); and equality restrictions, such as the symmetry restrictions in a translog model, which require that parameters in two different equations be equal to each other. Inequality constraints, in general, will be of the form $a_j \leq c_j(\theta) \leq b_j$, where a_j and b_j are known constants (either of which may be infinite). Once again, the typical application in econometrics involves a restriction on a single parameter, such as $\sigma > 0$ for a variance parameter, $-1 \leq \rho \leq 1$ for a correlation coefficient, or $\beta_j \geq 0$ for a particular slope coefficient in a model. We will consider the two cases separately.

In the case of equality constraints, for practical purposes of optimization, there are usually two strategies available. One can use a Lagrangean multiplier approach. The new optimization problem is

$$\text{Max}_{\theta, \lambda} L(\theta, \lambda) = F(\theta) + \lambda' \mathbf{c}(\theta).$$

The necessary conditions for an optimum are

$$\begin{aligned} \frac{\partial L(\theta, \lambda)}{\partial \theta} &= \mathbf{g}(\theta) + \mathbf{C}(\theta)' \lambda = \mathbf{0}, \\ \frac{\partial L(\theta, \lambda)}{\partial \lambda} &= \mathbf{c}(\theta) = \mathbf{0}, \end{aligned}$$

where $\mathbf{g}(\theta)$ is the familiar gradient of $F(\theta)$ and $\mathbf{C}(\theta)$ is a $J \times K$ matrix of derivatives with j th row equal to $\partial c_j / \partial \theta'$. The joint solution will provide the constrained optimizer, as well as the Lagrange multipliers, which are often interesting in their own right. The disadvantage of this approach is that it increases the dimensionality of the optimization problem. An alternative strategy is to eliminate some of the parameters by either imposing the constraints directly on the function or by solving out the constraints. For exclusion restrictions, which are usually of the form $\theta_j = 0$, this step usually means dropping a variable from a model. Other restrictions can often be imposed

³²A formal proof that this is a valid way to proceed is given by Amemiya (1985, pp. 125–127).

just by building them into the model. For example, in a function of θ_1 , θ_2 , and θ_3 , if the restriction is of the form $\theta_3 = \theta_1\theta_2$, then θ_3 can be eliminated from the model by a direct substitution.

Inequality constraints are more difficult. For the general case, one suggestion is to transform the constrained problem into an unconstrained one by imposing some sort of penalty function into the optimization criterion that will cause a parameter vector that violates the constraints, or nearly does so, to be an unattractive choice. For example, to force a parameter θ_j to be nonzero, one might maximize the augmented function $F(\theta) - |1/\theta_j|$. This approach is feasible, but it has the disadvantage that because the penalty is a function of the parameters, different penalty functions will lead to different solutions of the optimization problem. For the most common problems in econometrics, a simpler approach will usually suffice. One can often reparameterize a function so that the new parameter is unconstrained. For example, the “method of squaring” is sometimes used to force a parameter to be positive. If we require θ_j to be positive, then we can define $\theta_j = \alpha^2$ and substitute α^2 for θ_j wherever it appears in the model. Then an unconstrained solution for α is obtained. An alternative reparameterization for a parameter that must be positive that is often used is $\theta_j = \exp(\alpha)$. To force a parameter to be between zero and one, we can use the function $\theta_j = 1/[1 + \exp(\alpha)]$. The range of α is now unrestricted. Experience suggests that a third, less orthodox approach works very well for many problems. When the constrained optimization is begun, there is a starting value θ^0 that begins the iterations. Presumably, θ^0 obeys the restrictions. (If not, and none can be found, then the optimization process must be terminated immediately.) The next iterate, θ^1 , is a step away from θ^0 , by $\theta^1 = \theta^0 + \lambda_0\delta^0$. Suppose that θ^1 violates the constraints. By construction, we know that there is some value θ_*^1 between θ^0 and θ^1 that does not violate the constraint, where “between” means only that a shorter step is taken. Therefore, the next value for the iteration can be θ_*^1 . The logic is true at every iteration, so a way to proceed is to alter the iteration so that the step length is shortened when necessary when a parameter violates the constraints.

E.6.5 SOME PRACTICAL CONSIDERATIONS

The reasons for the good performance of many algorithms, including DFP, are unknown. Moreover, different algorithms may perform differently in given settings. Indeed, for some problems, one algorithm may fail to converge whereas another will succeed in finding a solution without great difficulty. In view of this, computer programs such as GQOPT³³ and Gauss that offer a menu of different preprogrammed algorithms can be particularly useful. It is sometimes worth the effort to try more than one algorithm on a given problem.

Step Sizes Except for the steepest ascent case, an optimal line search is likely to be infeasible or to require more effort than it is worth in view of the potentially large number of function evaluations required. In most cases, the choice of a step size is likely to be rather ad hoc. But within limits, the most widely used algorithms appear to be robust to inaccurate line searches. For example, one method employed by the widely used TSP computer program³⁴ is the method of *squeezing*, which tries $\lambda = 1, \frac{1}{2}, \frac{1}{4}$, and so on until an improvement in the function results. Although this approach is obviously a bit unorthodox, it appears to be quite effective when used with the Gauss–Newton method for nonlinear least squares problems. (See Chapter 9.) A somewhat more elaborate rule is suggested by Berndt et al. (1974). Choose an ε between 0 and $\frac{1}{2}$, and then find a λ such that

$$\varepsilon < \frac{F(\theta + \lambda\Delta) - F(\theta)}{\lambda g' \Delta} < 1 - \varepsilon. \quad (\text{E-27})$$

³³Goldfeld and Quandt (1972).

³⁴Hall (1982, p. 147).

Of course, which value of ε to choose is still open, so the choice of λ remains ad hoc. Moreover, in neither of these cases is there any optimality to the choice; we merely find a λ that leads to a function improvement. Other authors have devised relatively efficient means of searching for a step size without doing the full optimization at each iteration.³⁵

Assessing Convergence Ideally, the iterative procedure should terminate when the gradient is zero. In practice, this step will not be possible, primarily because of accumulated rounding error in the computation of the function and its derivatives. Therefore, a number of alternative convergence criteria are used. Most of them are based on the relative changes in the function or the parameters. There is considerable variation in those used in different computer programs, and there are some pitfalls that should be avoided. A critical absolute value for the elements of the gradient or its norm will be affected by any scaling of the function, such as normalizing it by the sample size. Similarly, stopping on the basis of small absolute changes in the parameters can lead to premature convergence when the parameter vector approaches the maximizer. It is probably best to use several criteria simultaneously, such as the proportional change in both the function and the parameters. Belsley (1980) discusses a number of possible stopping rules. One that has proved useful and is immune to the scaling problem is to base convergence on $\mathbf{g}'\mathbf{H}^{-1}\mathbf{g}$.

Multiple Solutions It is possible for a function to have several local extrema. It is difficult to know a priori whether this is true of the one at hand. But if the function is not globally concave, then it may be a good idea to attempt to maximize it from several starting points to ensure that the maximum obtained is the global one. Ideally, a starting value near the optimum can facilitate matters; in some settings, this can be obtained by using a consistent estimate of the parameter for the starting point. The method of moments, if available, is sometimes a convenient device for doing so.

No Solution Finally, it should be noted that in a nonlinear setting the iterative algorithm can break down, even in the absence of constraints, for at least two reasons. The first possibility is that the problem being solved may be so numerically complex as to defy solution. The second possibility, which is often neglected, is that the proposed model may simply be inappropriate for the data. In a linear setting, a low R^2 or some other diagnostic test may suggest that the model and data are mismatched, but as long as the full rank condition is met by the regressor matrix, a linear regression can *always* be computed. Nonlinear models are not so forgiving. The failure of an iterative algorithm to find a maximum of the criterion function may be a warning that the model is not appropriate for this body of data.

E.6.6 Examples

To illustrate the use of gradient methods, we consider some simple problems.

E.6.6.a Function of One Parameter

First, consider maximizing a function of a single variable, $f(\theta) = \ln(\theta) - 0.1\theta^2$. The function is shown in Figure E.4. The first and second derivatives are

$$f'(\theta) = \frac{1}{\theta} - 0.2\theta,$$

$$f''(\theta) = \frac{-1}{\theta^2} - 0.2.$$

³⁵See, for example, Joreskog and Gruvaeus (1970), Powell (1964), Quandt (1983), and Hall (1982).




FIGURE E.4 Function of One Variable Parameter.

TABLE E.1 Iterations for Newton's Method

Iteration	θ	f	f'	f''
0	5.00000	-0.890562	-0.800000	-0.240000
1	1.66667	0.233048	0.266667	-0.560000
2	2.14286	0.302956	0.030952	-0.417778
3	2.23404	0.304718	0.000811	-0.400363
4	2.23607	0.304719	0.0000004	-0.400000

Equating f' to zero yields the solution $\theta = \sqrt{5} = 2.236$. At the solution, $f'' = -0.4$, so this equation is indeed a maximum. To demonstrate the use of an iterative method, we solve this problem using Newton's method. Observe, first, that the second derivative is always negative for any admissible (positive) θ .³⁶ Therefore, it should not matter where we start the iterations; we shall eventually find the maximum. For a single parameter, Newton's method is

$$\theta_{t+1} = \theta_t - [f'_t / f''_t].$$

The sequence of values that results when 5 is used as the starting value is given in Table E.1. The path of the iterations is also shown in the table.

E.6.6.b Function of Two Parameters: The Gamma Distribution

For random sampling from the gamma distribution,

$$f(y_i, \beta, \rho) = \frac{\beta^\rho}{\Gamma(\rho)} e^{-\beta y_i} y_i^{\rho-1}.$$

³⁶In this problem, an inequality restriction, $\theta > 0$, is required. As is common, however, for our first attempt we shall neglect the constraint.

TABLE E.2 Iterative Solutions to $\text{Max}(\rho, \beta) \rho \ln \beta - \ln \Gamma(\rho) - 3\beta + \rho - 1$

Iter.	Trial 1				Trial 2				Trial 3			
	DFP		Newton		DFP		Newton		DFP		Newton	
	ρ	β	ρ	β	ρ	β	ρ	β	ρ	β	ρ	β
0	4.000	1.000	4.000	1.000	8.000	3.000	8.000	3.000	2.000	7.000	2.000	7.000
1	3.981	1.345	3.812	1.203	7.117	2.518	2.640	0.615	6.663	2.027	-47.7	-233.
2	4.005	1.324	4.795	1.577	7.144	2.372	3.203	0.931	6.195	2.075	—	—
3	5.217	1.743	5.190	1.728	7.045	2.389	4.257	1.357	5.239	1.731	—	—
4	5.233	1.744	5.231	1.744	5.114	1.710	5.011	1.656	5.251	1.754	—	—
5	—	—	—	—	5.239	1.747	5.219	1.740	5.233	1.744	—	—
6	—	—	—	—	5.233	1.744	5.233	1.744	—	—	—	—

The log-likelihood is $\ln L(\beta, \rho) = n\rho \ln \beta - n \ln \Gamma(\rho) - \beta \sum_{i=1}^n y_i + (\rho - 1) \sum_{i=1}^n \ln y_i$. (See Section 17.5.4 and Example 18.5.) It is often convenient to scale the log-likelihood by the sample size. Suppose, as well, that we have a sample with $\bar{y} = 3$ and $\ln \bar{y} = 1$. Then the function to be maximized is $F(\beta, \rho) = \rho \ln \beta - \ln \Gamma(\rho) - 3\beta + \rho - 1$. The derivatives are

$$\begin{aligned} \frac{\partial F}{\partial \beta} &= \frac{\rho}{\beta} - 3, & \frac{\partial F}{\partial \rho} &= \ln \beta - \frac{\Gamma'}{\Gamma} + 1 = \ln \beta - \Psi(\rho) + 1, \\ \frac{\partial^2 F}{\partial \beta^2} &= -\frac{\rho}{\beta^2}, & \frac{\partial^2 F}{\partial \rho^2} &= \frac{-(\Gamma\Gamma'' - \Gamma'^2)}{\Gamma^2} = -\Psi'(\rho), & \frac{\partial^2 F}{\partial \beta \partial \rho} &= \frac{1}{\beta}. \end{aligned}$$

Finding a good set of starting values is often a difficult problem. Here we choose three starting points somewhat arbitrarily: $(\rho^0, \beta^0) = (4, 1)$, $(8, 3)$, and $(2, 7)$. The solution to the problem is $(5.233, 1.7438)$. We used Newton's method and DFP with a line search to maximize this function.³⁷ For Newton's method, $\lambda = 1$. The results are shown in Table E.2. The two methods were essentially the same when starting from a good starting point (trial 1), but they differed substantially when starting from a poorer one (trial 2). Note that DFP and Newton approached the solution from different directions in trial 2. The third starting point shows the value of a line search. At this starting value, the Hessian is extremely large, and the second value for the parameter vector with Newton's method is $(-47.671, -233.35)$, at which point F cannot be computed and this method must be abandoned. Beginning with $\mathbf{H} = \mathbf{I}$ and using a line search, DFP reaches the point $(6.63, 2.03)$ at the first iteration, after which convergence occurs routinely in three more iterations. At the solution, the Hessian is $[[-1.72038, 0.191153], [0.191153, -0.210579]]$. The diagonal elements of the Hessian are negative and its determinant is 0.32574, so it is negative definite. (The two characteristic roots are -1.7442 and -0.18675). Therefore, this result is indeed the maximum of the function.

E.6.6.c A Concentrated Log-Likelihood Function

There is another way that the preceding problem might have been solved. The first of the necessary conditions implies that at the joint solution for (β, ρ) , β will equal $\rho/3$. Suppose that we impose this requirement on the function we are maximizing. The **concentrated** (over β) **log-likelihood function** is then produced:

$$\begin{aligned} F_c(\rho) &= \rho \ln(\rho/3) - \ln \Gamma(\rho) - 3(\rho/3) + \rho - 1 \\ &= \rho \ln(\rho/3) - \ln \Gamma(\rho) - 1. \end{aligned}$$

³⁷The one used is described in Joreskog and Gruvaeus (1970).




FIGURE E.5 Concentrated Log-Likelihood.

This function could be maximized by an iterative search or by a simple one-dimensional grid search. Figure E.5 shows the behavior of the function. As expected, the maximum occurs at $\rho = 5.233$. The value of β is found as $5.23/3 = 1.743$.

The concentrated log-likelihood is a useful device in many problems. Note the interpretation of the function plotted in Figure E.5. The original function of ρ and β is a surface in three dimensions. The curve in Figure E.5 is a projection of that function; it is a plot of the function values above the line $\beta = \rho/3$. By virtue of the first-order condition, we know that one of these points will be the maximizer of the function. Therefore, we may restrict our search for the overall maximum of $F(\beta, \rho)$ to the points on this line.

APPENDIX F



DATA SETS USED IN APPLICATIONS

The following tables list the variables in the data sets used in the applications in the text. The data sets, themselves, can be downloaded from the website for the text.

TABLE F1.1 Consumption and Income, 10 Yearly Observations, 1970–1979

C = Consumption,
Y = Disposable Income.

Source: Council of Economic Advisors, *Economic Report of the President* (Washington, D.C.: U.S. Government Printing Office, 1987).

TABLE F2.1 Consumption and Income, 11 Yearly Observations, 1940–1950

Year = Date,

X = Disposable Income,

C = Consumption,

W = War years dummy variable, one in 1942–1945, zero other years.

Source: *Economic Report of the President* (Washington, D.C.: U.S. Government Printing Office, 1983).

TABLE F2.2 The U.S. Gasoline Market, 36 Yearly Observations, 1960–1995

G = Total U.S. gasoline consumption, computed as total expenditure divided by price index,

P_g = Price index for gasoline,

Y = Per capita disposable income,

P_{nc} = Price index for new cars,

P_{uc} = Price index for used cars,

P_p = Price index for public transportation,

P_d = Aggregate price index for consumer durables,

P_n = Aggregate price index for consumer nondurables,

P_s = Aggregate price index for consumer services,

Pop = U.S. total population in millions.

Source: Council of Economic Advisors, *Economic Report of the President: 1996* (Washington, D.C.: U.S. Government Printing Office, 1996).

TABLE F3.1 Investment, 15 Yearly Observations, 1968–1982

Year = Date,

GNP = Nominal GNP,

Invest = Nominal Investment,

CPI = Consumer price index,

Interest = Interest rate.

Source: *Economic Report of the President* (Washington, D.C.: U.S. Government Printing Office, 1983). CPI 1967 is 79.06. The interest rate is the average yearly discount rate at the New York Federal Reserve Bank.

TABLE F4.1 Labor Supply Data from Mroz (1987)

LFP = A dummy variable = 1 if woman worked in 1975, else 0,

WHRS = Wife's hours of work in 1975,

KL6 = Number of children less than 6 years old in household,

K618 = Number of children between ages 6 and 18 in household,

WA = Wife's age,

WE = Wife's educational attainment, in years,

WW = Wife's average hourly earnings, in 1975 dollars,

RPWG = Wife's wage reported at the time of the 1976 interview (not = 1975 estimated wage),

HHRS = Husband's hours worked in 1975,

HA = Husband's age,

HE = Husband's educational attainment, in years,

HW = Husband's wage, in 1975 dollars,

FAMINC = Family income, in 1975 dollars,

WMED = Wife's mother's educational attainment, in years,

WFED = Wife's father's educational attainment, in years,

UN = Unemployment rate in county of residence, in percentage points,

CIT = Dummy variable = one if live in large city (SMSA), else zero,

AX = Actual years of wife's previous labor market experience.

Source: 1976 Panel Study of Income Dynamics, Mroz (1987).

TABLE F4.2 The Longley Data, 15 Yearly Observations, 1947–1962

Employ = Employment (1,000s),
Price = GNP deflator,
GNP = Nominal GNP (millions),
Armed = Armed forces,
Year = Date.

Source: Longley (1967).

TABLE F5.1 Macroeconomics Data Set, Quarterly, 1950I to 2000IV

Year = Date,
Qtr = Quarter,
Realgdp = Real GDP (\$bil),
Realcons = Real consumption expenditures,
Realinvs = Real investment by private sector,
Realgovt = Real government expenditures,
Realdpi = Real disposable personal income,
CPI.U = Consumer price index,
MI = Nominal money stock,
Tbilrate = Quarterly average of month end 90 day *T*-bill rate,
Unemp = Unemployment rate,
Pop = Population, mil. interpolate of year end figures using constant growth rate per quarter,
Infl = Rate of inflation (first observation is missing),
Realint = Ex post real interest rate = *Tbilrate*—*Infl*. (First observation missing).

Source: Department of Commerce, BEA website and www.economagic.com.

TABLE F5.2 Cost Function, 123 1970 Cross-section Firm Level Observations

Id = Observation,
Year = 1970 for all observations,
Cost = Total cost,
Q = Total output,
Pl = Wage rate,
Sl = Cost share for labor,
Pk = Capital price index,
Sk = Cost share for capital,
Pf = Fuel price,
Sf = Cost share for fuel.

Source: Christensen and Greene (1976). Note the file contains some extra observations. These are the holding companies. Use only the first 123 observations to replicate Christensen and Greene.

TABLE F6.1 Production for SIC 33: Primary Metals, 27 Statewide Observations

Obs = Observation number,
Valueadd = Value added,
Labor = Labor input,
Capital = Capital stock.

Note: Data are per establishment, labor is a measure of labor input, and capital is the gross value of plant and equipment. A scale factor used to normalize the capital figure in the original study has been omitted. Further details on construction of the data are given in Aigner et al. (1977) and in Hildebrand and Liu (1957).

TABLE F7.1 Costs for U.S. Airlines, 90 Observations on 6 Firms for 1970–1984

I = Airline,
 T = Year,
 Q = Output, in revenue passenger miles, index number,
 C = Total cost, in \$1,000,
 PF = Fuel price,
 LF = Load factor, the average capacity utilization of the fleet.

Source: These data are a subset of a larger data set provided to the author by Professor Moshe Kim. They were originally constructed by Christensen Associates of Madison, Wisconsin.

TABLE F7.2 Solow's Technological Change Data, 41 Yearly Observations, 1909–1949

$Year$ = Date,
 Q = Output,
 K = Capital/labor ratio,
 A = Index of technology.

Source: Solow (1957, p. 314). Several Variables are omitted.

TABLE F8.1 Pesaran and Hall Inflation Data

Pa = Actual inflation,
 Pe = Expected inflation,

Source: Pesaran and Hall (1988).

TABLE F9.1 Income and Expenditure Data. 100 Cross Section Observations

MDR = Number of derogatory reports,
 Acc = Credit card application accepted (1 = yes),
 Age = Age in years + 12ths of a year,
 $Income$ = Income, divided by 10,000,
 $Avgexp$ = Avg. monthly credit card expenditure,
 $Ownrent$ = OwnRent, individual owns (1) or rents (0) home,
 $Selfempl$ = Self employed (1 = yes, 0 = no).

Source: Greene (1992).

TABLE F9.2 Statewide Data on Transportation Equipment Manufacturing, 25 Observations

$State$ = Observation,
 $ValueAdd$ = Output,
 $Capita$ = Capital input,
 $Labor$ = Labor input,
 $Nfirm$ = Number of firms.

Source: A Zellner and N. Revankar (1970, p. 249). *Note:* "Value added," "Capital," and "Labor" are in millions of 1957 dollars. Data used for regression examples are per establishment. Raw data are used for the stochastic frontier application in Chapter 16.

TABLE F11.1 Bollerslev and Ghysels Exchange Rate Data, 1974 Daily Observations

Y = Nominal return on Mark/Pound exchange rate, daily.

Source: Bollerslev (1986).

TABLE F13.1 Grunfeld Investment Data, 100 Yearly Observations on 5 Firms for 1935–1954

I = Gross investment, from *Moody's Industrial Manual* and annual reports of corporations,
 F = Value of the firm from *Bank and Quotation Record* and *Moody's Industrial Manual*,
 C = Stock of plant and equipment, from *Survey of Current Business*.

Source: *Moody's Industrial Manual*, *Survey of Current Business*.

TABLE F14.1 Manufacturing Costs, U.S. Economy, 25 Yearly Observations, 1947–1971

$Year$ = Date,
 $Cost$ = Cost index,
 K = Capital cost share,
 L = Labor cost share,
 E = Energy cost share,
 M = Materials cost share,
 Pk = Capital price,
 Pl = Labor price,
 Pe = Energy price,
 Pm = Materials price.

Source: Berndt and Wood (1975).

TABLE F14.2 Cost Function, 145 U.S. Electricity Producers, 1955 Data–Nerlove

$Firm$ = Observation,
 $Year$ = 1955 for all observations,
 $Cost$ = Total cost,
 $Output$ = Total output,
 Pl = Wage rate,
 Sl = Cost share for labor,
 Pk = Capital price index,
 Sk = Cost share for capital,
 Pf = Fuel price,
 Sf = Cost share for fuel.

Source: Nerlove (1963) and Christensen and Greene (1976).

Note: The data file contains several extra observations that are aggregates of commonly owned firms. Use only the first 145 for analysis.

TABLE F15.1 Klein's Model I, 22 Yearly Observations, 1920–1941

$Year$ = Date,
 C = Consumption,
 P = Corporate profits,
 Wp = Private wage bill,
 I = Investment,
 KI = previous year's capital stock,
 X = GNP,
 Wg = Government wage bill,
 G = Government spending,
 T = Taxes.

Source: Klein (1950).

TABLE F16.1 Bertschek and Lechner Binary Choice Data

y_{it} = One if firm i realized a product innovation in year t and zero if not,
 x_{it2} = Log of sales,
 x_{it3} = Relative size = Ratio of employment in business unit to employment in the industry,
 x_{it4} = Ratio of industry imports to (industry sales + imports),
 x_{it5} = Ratio of industry foreign direct investment to (industry sales + imports),
 x_{it6} = Productivity = Ratio of industry value added to industry employment,
 x_{it7} = Dummy variable indicating firm is in the raw materials sector,
 x_{it8} = Dummy variable indicating firm is in the investment goods sector.

Note: These data are proprietary.

Source: Bertschek and Lechner (1998).

TABLE F18.1 Dahlberg and Johanssen–Municipal Expenditure Data

ID = Identification,
 $Year$ = Date,
 $Expend$ = Expenditure,
 $Revenue$ = Revenue from taxes and fees,
 $Grants$ = Grants from Central Government.

Source: Dahlberg and Johanssen (2000), *Journal of Applied Econometrics*, data archive.

TABLE F20.1 Bond Yield on a Moody's Aaa Rated, Monthly, 60 Monthly Observations, 1990–1994

$Date$ = Year.Month,
 Y = Corporate bond rate in percent/year.

Source: *National Income and Product Accounts*, U.S. Department of Commerce, Bureau of Economic Analysis, *Survey of Current Business*: Business Statistics.

TABLE F20.2 Money, Output, and Price Deflator Data, 136 Quarterly Observations, 1950–1983

Y = Nominal GNP,
 $M1$ = M1 measure of money stock,
 P = Implicit price deflator for GNP.

Source: *National Income and Product Accounts*, U.S. Department of Commerce, Bureau of Economic Analysis, *Survey of Current Business*: Business Statistics.

TABLE F21.1 Program Effectiveness, 32 Cross Section Observations

Obs = Observation,
 $TUCE$ = Test score on economics test,
 PSI = Participation in program,
 $GRADE$ = Grade increase (1) or decrease (0) indicator.

Source: Spector and Mazzeo (1980).

TABLE F21.2 Data Used to Study Travel Mode Choice, 840 Observations, on 4 Modes for 210 Individuals

$Mode$ = Choice: Air, Train, Bus, or Car,
 $Ttime$ = Terminal waiting time, 0 for car,
 $Invc$ = In vehicle cost—cost component,
 $Invt$ = Travel time, in vehicle,
 GC = Generalized cost measure,
 $Hinc$ = Household income,
 $Psize$ = Party size in mode chosen.

Source: Greene and Hensher (1997).

TABLE F21.3 Ship Accidents, 40 Observations on 5 Types in 4 Vintages and 2 Service Periods

Type = Ship type,
TA, TB, TC, TD, TE = Type indicators,
Y6064, Y6569, Y7074, Y7579 = Year constructed indicators,
O6064, O7579 = Years operated indicators,
Months = Measure of service amount,
Acc = Accidents.

Source: McCullagh and Nelder (1983).

TABLE F21.4 Expenditure and Default Data, 1,319 Observations

Cardhldr = Dummy variable, one if application for credit card accepted, zero if not,
Majordrg = Number of major derogatory reports,
Age = Age *n* years plus twelfths of a year,
Income = Yearly income (divided by 10,000),
Exp_Inc = Ratio of monthly credit card expenditure to yearly income,
Avgexp = Average monthly credit card expenditure,
Ownrent = 1 if owns their home, 0 if rent,
Selfempl = 1 if self employed, 0 if not,
Depndt = 1 + number of dependents,
Inc_per = Income divided by number of dependents,
Cur_add = Months living at current address,
Major = Number of major credit cards held,
Active = Number of active credit accounts.

Source: Greene (1992).

TABLE F22.1 Strike Duration Data, 63 Observations in 9 Years, 1968–1976

Year = Date,
T = Strike duration in days,
PROD = Unanticipated output.

Source: Kennan (1985).

TABLE F22.2 Fair's (1977) Extramarital Affairs Data, 601 Cross-section Observations

y = Number of affairs in the past year,
z1 = Sex,
z2 = Age,
z3 = Number of years married,
z4 = Children,
z5 = Religiousness,
z6 = Education,
z7 = Occupation,
z8 = Self rating of marriage,
 Several variables not used are denoted *X1, . . . , X5*.

Source: Fair (1977) and <http://fairmodel.econ.yale.edu/rayfair/pdf/1978ADAT.ZIP>.

TABLE G.2 Percentiles of the Student's *t* Distribution. Table Entry Is *x* Such that $\text{Prob}[t_n \leq x] = P$

<i>n</i>	.750	.900	.950	.975	.990	.995
1	1.000	3.078	6.314	12.706	31.821	63.657
2	.816	1.886	2.920	4.303	6.965	9.925
3	.765	1.638	2.353	3.182	4.541	5.841
4	.741	1.533	2.132	2.776	3.747	4.604
5	.727	1.476	2.015	2.571	3.365	4.032
6	.718	1.440	1.943	2.447	3.143	3.707
7	.711	1.415	1.895	2.365	2.998	3.499
8	.706	1.397	1.860	2.306	2.896	3.355
9	.703	1.383	1.833	2.262	2.821	3.250
10	.700	1.372	1.812	2.228	2.764	3.169
11	.697	1.363	1.796	2.201	2.718	3.106
12	.695	1.356	1.782	2.179	2.681	3.055
13	.694	1.350	1.771	2.160	2.650	3.012
14	.692	1.345	1.761	2.145	2.624	2.977
15	.691	1.341	1.753	2.131	2.602	2.947
16	.690	1.337	1.746	2.120	2.583	2.921
17	.689	1.333	1.740	2.110	2.567	2.898
18	.688	1.330	1.734	2.101	2.552	2.878
19	.688	1.328	1.729	2.093	2.539	2.861
20	.687	1.325	1.725	2.086	2.528	2.845
21	.686	1.323	1.721	2.080	2.518	2.831
22	.686	1.321	1.717	2.074	2.508	2.819
23	.685	1.319	1.714	2.069	2.500	2.807
24	.685	1.318	1.711	2.064	2.492	2.797
25	.684	1.316	1.708	2.060	2.485	2.787
26	.684	1.315	1.706	2.056	2.479	2.779
27	.684	1.314	1.703	2.052	2.473	2.771
28	.683	1.313	1.701	2.048	2.467	2.763
29	.683	1.311	1.699	2.045	2.462	2.756
30	.683	1.310	1.697	2.042	2.457	2.750
35	.682	1.306	1.690	2.030	2.438	2.724
40	.681	1.303	1.684	2.021	2.423	2.704
45	.680	1.301	1.679	2.014	2.412	2.690
50	.679	1.299	1.676	2.009	2.403	2.678
60	.679	1.296	1.671	2.000	2.390	2.660
70	.678	1.294	1.667	1.994	2.381	2.648
80	.678	1.292	1.664	1.990	2.374	2.639
90	.677	1.291	1.662	1.987	2.368	2.632
100	.677	1.290	1.660	1.984	2.364	2.626
∞	.674	1.282	1.645	1.960	2.326	2.576

TABLE G.3 Percentiles of the Chi-Squared Distribution. Table Entry Is c such that $\text{Prob}[\chi_n^2 \leq c] = P$

n	.005	.010	.025	.050	.100	.250	.500	.750	.900	.950	.975	.990	.995
1	.00004	.0002	.001	.004	.02	.10	.45	1.32	2.71	3.84	5.02	6.63	7.88
2	.01	.02	.05	.10	.21	.58	1.39	2.77	4.61	5.99	7.38	9.21	10.60
3	.07	.11	.22	.35	.58	1.21	2.37	4.11	6.25	7.81	9.35	11.34	12.84
4	.21	.30	.48	.71	1.06	1.92	3.36	5.39	7.78	9.49	11.14	13.28	14.86
5	.41	.55	.83	1.15	1.61	2.67	4.35	6.63	9.24	11.07	12.83	15.09	16.75
6	.68	.87	1.24	1.64	2.20	3.45	5.35	7.84	10.64	12.59	14.45	16.81	18.55
7	.99	1.24	1.69	2.17	2.83	4.25	6.35	9.04	12.02	14.07	16.01	18.48	20.28
8	1.34	1.65	2.18	2.73	3.49	5.07	7.34	10.22	13.36	15.51	17.53	20.09	21.95
9	1.73	2.09	2.70	3.33	4.17	5.90	8.34	11.39	14.68	16.92	19.02	21.67	23.59
10	2.16	2.56	3.25	3.94	4.87	6.74	9.34	12.55	15.99	18.31	20.48	23.21	25.19
11	2.60	3.05	3.82	4.57	5.58	7.58	10.34	13.70	17.28	19.68	21.92	24.72	26.76
12	3.07	3.57	4.40	5.23	6.30	8.44	11.34	14.85	18.55	21.03	23.34	26.22	28.30
13	3.57	4.11	5.01	5.89	7.04	9.30	12.34	15.98	19.81	22.36	24.74	27.69	29.82
14	4.07	4.66	5.63	6.57	7.79	10.17	13.34	17.12	21.06	23.68	26.12	29.14	31.32
15	4.60	5.23	6.26	7.26	8.55	11.04	14.34	18.25	22.31	25.00	27.49	30.58	32.80
16	5.14	5.81	6.91	7.96	9.31	11.91	15.34	19.37	23.54	26.30	28.85	32.00	34.27
17	5.70	6.41	7.56	8.67	10.09	12.79	16.34	20.49	24.77	27.59	30.19	33.41	35.72
18	6.26	7.01	8.23	9.39	10.86	13.68	17.34	21.60	25.99	28.87	31.53	34.81	37.16
19	6.84	7.63	8.91	10.12	11.65	14.56	18.34	22.72	27.20	30.14	32.85	36.19	38.58
20	7.43	8.26	9.59	10.85	12.44	15.45	19.34	23.83	28.41	31.41	34.17	37.57	40.00
21	8.03	8.90	10.28	11.59	13.24	16.34	20.34	24.93	29.62	32.67	35.48	38.93	41.40
22	8.64	9.54	10.98	12.34	14.04	17.24	21.34	26.04	30.81	33.92	36.78	40.29	42.80
23	9.26	10.20	11.69	13.09	14.85	18.14	22.34	27.14	32.01	35.17	38.08	41.64	44.18
24	9.89	10.86	12.40	13.85	15.66	19.04	23.34	28.24	33.20	36.42	39.36	42.98	45.56
25	10.52	11.52	13.12	14.61	16.47	19.94	24.34	29.34	34.38	37.65	40.65	44.31	46.93
30	13.79	14.95	16.79	18.49	20.60	24.48	29.34	34.80	40.26	43.77	46.98	50.89	53.67
35	17.19	18.51	20.57	22.47	24.80	29.05	34.34	40.22	46.06	49.80	53.20	57.34	60.27
40	20.71	22.16	24.43	26.51	29.05	33.66	39.34	45.62	51.81	55.76	59.34	63.69	66.77
45	24.31	25.90	28.37	30.61	33.35	38.29	44.34	50.98	57.51	61.66	65.41	69.96	73.17
50	27.99	29.71	32.36	34.76	37.69	42.94	49.33	56.33	63.17	67.50	71.42	76.15	79.49

TABLE G.4 95th Percentiles of the *F* Distribution. Table Entry is *f* such that $\text{Prob}[F_{n_1, n_2} \leq f] = .95$

<i>n</i> ₁ = Degrees of Freedom for the Numerator									
<i>n</i> ₂	1	2	3	4	5	6	7	8	9
1	161.45	199.50	215.71	224.58	230.16	233.99	236.77	238.88	240.54
2	18.51	19.00	19.16	19.25	19.30	19.33	19.35	19.37	19.38
3	10.13	9.55	9.28	9.12	9.01	8.94	8.89	8.85	8.81
4	7.71	6.94	6.59	6.39	6.26	6.16	6.09	6.04	6.00
5	6.61	5.79	5.41	5.19	5.05	4.95	4.88	4.82	4.77
6	5.99	5.14	4.76	4.53	4.39	4.28	4.21	4.15	4.10
7	5.59	4.74	4.35	4.12	3.97	3.87	3.79	3.73	3.68
8	5.32	4.46	4.07	3.84	3.69	3.58	3.50	3.44	3.39
9	5.12	4.26	3.86	3.63	3.48	3.37	3.29	3.23	3.18
10	4.96	4.10	3.71	3.48	3.33	3.22	3.14	3.07	3.02
15	4.54	3.68	3.29	3.06	2.90	2.79	2.71	2.64	2.59
20	4.35	3.49	3.10	2.87	2.71	2.60	2.51	2.45	2.39
25	4.24	3.39	2.99	2.76	2.60	2.49	2.40	2.34	2.28
30	4.17	3.32	2.92	2.69	2.53	2.42	2.33	2.27	2.21
40	4.08	3.23	2.84	2.61	2.45	2.34	2.25	2.18	2.12
50	4.03	3.18	2.79	2.56	2.40	2.29	2.20	2.13	2.07
70	3.98	3.13	2.74	2.50	2.35	2.23	2.14	2.07	2.02
100	3.94	3.09	2.70	2.46	2.31	2.19	2.10	2.03	1.97
∞	3.84	3.00	2.60	2.37	2.21	2.10	2.01	1.94	1.88

<i>n</i> ₂	10	12	15	20	30	40	50	60	∞
1	241.88	243.91	245.95	248.01	250.10	251.14	252.20	252.20	254.19
2	19.40	19.41	19.43	19.45	19.46	19.47	19.48	19.48	19.49
3	8.79	8.74	8.70	8.66	8.62	8.59	8.57	8.57	8.53
4	5.96	5.91	5.86	5.80	5.75	5.72	5.69	5.69	5.63
5	4.74	4.68	4.62	4.56	4.50	4.46	4.43	4.43	4.37
6	4.06	4.00	3.94	3.87	3.81	3.77	3.74	3.74	3.67
7	3.64	3.57	3.51	3.44	3.38	3.34	3.30	3.30	3.23
8	3.35	3.28	3.22	3.15	3.08	3.04	3.01	3.01	2.93
9	3.14	3.07	3.01	2.94	2.86	2.83	2.79	2.79	2.71
10	2.98	2.91	2.85	2.77	2.70	2.66	2.62	2.62	2.54
15	2.54	2.48	2.40	2.33	2.25	2.20	2.16	2.16	2.07
20	2.35	2.28	2.20	2.12	2.04	1.99	1.95	1.95	1.85
25	2.24	2.16	2.09	2.01	1.92	1.87	1.82	1.82	1.72
30	2.16	2.09	2.01	1.93	1.84	1.79	1.74	1.74	1.63
40	2.08	2.00	1.92	1.84	1.74	1.69	1.64	1.64	1.52
50	2.03	1.95	1.87	1.78	1.69	1.63	1.58	1.58	1.45
70	1.97	1.89	1.81	1.72	1.62	1.57	1.50	1.50	1.36
100	1.93	1.85	1.77	1.68	1.57	1.52	1.45	1.45	1.30
∞	1.83	1.75	1.67	1.57	1.46	1.39	1.34	1.31	1.30

TABLE G.5 99th Percentiles of the *F* Distribution. Table Entry is *f* such that $\text{Prob}[F_{n_1, n_2} \leq f] = .99$

		<i>n</i> ₁ = Degrees of Freedom for the Numerator								
<i>n</i> ₂	1	2	3	4	5	6	7	8	9	
1	4052.18	4999.50	5403.35	5624.58	5763.65	5858.99	5928.36	5981.07	6022.47	
2	98.50	99.00	99.17	99.25	99.30	99.33	99.36	99.37	99.39	
3	34.12	30.82	29.46	28.71	28.24	27.91	27.67	27.49	27.35	
4	21.20	18.00	16.69	15.98	15.52	15.21	14.98	14.80	14.66	
5	16.26	13.27	12.06	11.39	10.97	10.67	10.46	10.29	10.16	
6	13.75	10.92	9.78	9.15	8.75	8.47	8.26	8.10	7.98	
7	12.25	9.55	8.45	7.85	7.46	7.19	6.99	6.84	6.72	
8	11.26	8.65	7.59	7.01	6.63	6.37	6.18	6.03	5.91	
9	10.56	8.02	6.99	6.42	6.06	5.80	5.61	5.47	5.35	
10	10.04	7.56	6.55	5.99	5.64	5.39	5.20	5.06	4.94	
15	8.68	6.36	5.42	4.89	4.56	4.32	4.14	4.00	3.89	
20	8.10	5.85	4.94	4.43	4.10	3.87	3.70	3.56	3.46	
25	7.77	5.57	4.68	4.18	3.85	3.63	3.46	3.32	3.22	
30	7.56	5.39	4.51	4.02	3.70	3.47	3.30	3.17	3.07	
40	7.31	5.18	4.31	3.83	3.51	3.29	3.12	2.99	2.89	
50	7.17	5.06	4.20	3.72	3.41	3.19	3.02	2.89	2.78	
70	7.01	4.92	4.07	3.60	3.29	3.07	2.91	2.78	2.67	
100	6.90	4.82	3.98	3.51	3.21	2.99	2.82	2.69	2.59	
∞	6.66	4.63	3.80	3.34	3.04	2.82	2.66	2.53	2.43	

<i>n</i> ₂	10	12	15	20	30	40	50	60	∞
1	6055.85	6106.32	6157.28	6208.73	6260.65	6286.78	6313.03	6313.03	6362.68
2	99.40	99.42	99.43	99.45	99.47	99.47	99.48	99.48	99.50
3	27.23	27.05	26.87	26.69	26.50	26.41	26.32	26.32	26.14
4	14.55	14.37	14.20	14.02	13.84	13.75	13.65	13.65	13.47
5	10.05	9.89	9.72	9.55	9.38	9.29	9.20	9.20	9.03
6	7.87	7.72	7.56	7.40	7.23	7.14	7.06	7.06	6.89
7	6.62	6.47	6.31	6.16	5.99	5.91	5.82	5.82	5.66
8	5.81	5.67	5.52	5.36	5.20	5.12	5.03	5.03	4.87
9	5.26	5.11	4.96	4.81	4.65	4.57	4.48	4.48	4.32
10	4.85	4.71	4.56	4.41	4.25	4.17	4.08	4.08	3.92
15	3.80	3.67	3.52	3.37	3.21	3.13	3.05	3.05	2.88
20	3.37	3.23	3.09	2.94	2.78	2.69	2.61	2.61	2.43
25	3.13	2.99	2.85	2.70	2.54	2.45	2.36	2.36	2.18
30	2.98	2.84	2.70	2.55	2.39	2.30	2.21	2.21	2.02
40	2.80	2.66	2.52	2.37	2.20	2.11	2.02	2.02	1.82
50	2.70	2.56	2.42	2.27	2.10	2.01	1.91	1.91	1.70
70	2.59	2.45	2.31	2.15	1.98	1.89	1.78	1.78	1.56
100	2.50	2.37	2.22	2.07	1.89	1.80	1.69	1.69	1.45
∞	2.34	2.20	2.06	1.90	1.72	1.61	1.50	1.50	1.16

TABLE G.6 Durbin–Watson Statistic: 5 Percent Significance Points of *dL* and *dU*

<i>n</i>	<i>k</i> = 1		<i>k</i> = 2		<i>k</i> = 3		<i>k</i> = 4		<i>k</i> = 5		<i>k</i> = 10		<i>k</i> = 15	
	<i>dL</i>	<i>dU</i>	<i>dL</i>	<i>dU</i>	<i>dL</i>	<i>dU</i>	<i>dL</i>	<i>dU</i>	<i>dL</i>	<i>dU</i>	<i>dL</i>	<i>dU</i>	<i>dL</i>	<i>dU</i>
15	1.08	1.36	.95	1.54	.82	1.75	.69	1.97	.56	2.21				
16	1.10	1.37	.98	1.54	.86	1.73	.74	1.93	.62	2.15	.16	3.30		
17	1.13	1.38	1.02	1.54	.90	1.71	.78	1.90	.67	2.10	.20	3.18		
18	1.16	1.39	1.05	1.53	.93	1.69	.82	1.87	.71	2.06	.24	3.07		
19	1.18	1.40	1.08	1.53	.97	1.68	.86	1.85	.75	2.02	.29	2.97		
20	1.20	1.41	1.10	1.54	1.00	1.68	.90	1.83	.79	1.99	.34	2.89	.06	3.68
21	1.22	1.42	1.13	1.54	1.03	1.67	.93	1.81	.83	1.96	.38	2.81	.09	3.58
22	1.24	1.43	1.15	1.54	1.05	1.66	.96	1.80	.86	1.94	.42	2.73	.12	3.55
23	1.26	1.44	1.17	1.54	1.08	1.66	.99	1.79	.90	1.92	.47	2.67	.15	3.41
24	1.27	1.45	1.19	1.55	1.10	1.66	1.01	1.78	.93	1.90	.51	2.61	.19	3.33
25	1.29	1.45	1.21	1.55	1.12	1.66	1.04	1.77	.95	1.89	.54	2.57	.22	3.25
26	1.30	1.46	1.22	1.55	1.14	1.65	1.06	1.76	.98	1.88	.58	2.51	.26	3.18
27	1.32	1.47	1.24	1.56	1.16	1.65	1.08	1.76	1.01	1.86	.62	2.47	.29	3.11
28	1.33	1.48	1.26	1.56	1.18	1.65	1.10	1.75	1.03	1.85	.65	2.43	.33	3.05
29	1.34	1.48	1.27	1.56	1.20	1.65	1.12	1.74	1.05	1.84	.68	2.40	.36	2.99
30	1.35	1.49	1.28	1.57	1.21	1.65	1.14	1.74	1.07	1.83	.71	2.36	.39	2.94
31	1.36	1.50	1.30	1.57	1.23	1.65	1.16	1.74	1.09	1.83	.74	2.33	.43	2.99
32	1.37	1.50	1.31	1.57	1.24	1.65	1.18	1.73	1.11	1.82	.77	2.31	.46	2.84
33	1.38	1.51	1.32	1.58	1.26	1.65	1.19	1.73	1.13	1.81	.80	2.28	.49	2.80
34	1.39	1.51	1.33	1.58	1.27	1.65	1.21	1.73	1.15	1.81	.82	2.26	.52	2.75
35	1.40	1.52	1.34	1.53	1.28	1.65	1.22	1.73	1.16	1.80	.85	2.24	.55	2.72
36	1.41	1.52	1.35	1.59	1.29	1.65	1.24	1.73	1.18	1.80	.87	2.22	.58	2.68
37	1.42	1.53	1.36	1.59	1.31	1.66	1.25	1.72	1.19	1.80	.89	2.20	.60	2.65
38	1.43	1.54	1.37	1.59	1.32	1.66	1.26	1.72	1.21	1.79	.91	2.18	.63	2.61
39	1.43	1.54	1.38	1.60	1.33	1.66	1.27	1.72	1.22	1.79	.93	2.16	.65	2.59
40	1.44	1.54	1.39	1.60	1.34	1.66	1.29	1.72	1.23	1.79	.95	2.15	.68	2.56
45	1.48	1.57	1.43	1.62	1.38	1.67	1.34	1.72	1.29	1.78	1.04	2.09	.79	2.44
50	1.50	1.59	1.46	1.63	1.42	1.67	1.38	1.72	1.34	1.77	1.11	2.04	.88	2.35
55	1.53	1.60	1.49	1.64	1.45	1.68	1.41	1.72	1.38	1.77	1.17	2.01	.96	2.28
60	1.55	1.62	1.51	1.65	1.48	1.69	1.44	1.73	1.41	1.77	1.22	1.98	1.03	2.23
65	1.57	1.63	1.54	1.66	1.50	1.70	1.47	1.73	1.44	1.77	1.27	1.96	1.09	2.18
70	1.58	1.64	1.55	1.67	1.52	1.70	1.49	1.74	1.46	1.77	1.30	1.95	1.14	2.15
75	1.60	1.65	1.57	1.68	1.54	1.71	1.51	1.74	1.49	1.77	1.34	1.94	1.18	2.12
80	1.61	1.66	1.59	1.69	1.56	1.72	1.53	1.74	1.51	1.77	1.37	1.93	1.22	2.09
85	1.62	1.67	1.60	1.70	1.57	1.72	1.55	1.75	1.52	1.77	1.40	1.92	1.26	2.07
90	1.63	1.68	1.61	1.70	1.59	1.73	1.57	1.75	1.54	1.78	1.42	1.91	1.29	2.06
95	1.64	1.69	1.62	1.71	1.60	1.73	1.58	1.75	1.56	1.78	1.44	1.90	1.32	2.04
100	1.65	1.69	1.63	1.72	1.61	1.74	1.59	1.76	1.57	1.78	1.46	1.90	1.35	2.03

Source: Extracted from N.E. Savin and K.J. White, "The Durbin–Watson Test for Serial Correlation with Extreme Sample Sizes and Many Regressors," *Econometrica*, 45 (8), Nov. 1977, pp. 1992–1995.

Note: *k* is the number of regressors excluding the intercept.