

APPENDIX TO PART III

This Appendix, designated as **A3**, contains additional analytical results for Part III of the NOTEBOOK, and follows the notational conventions in Appendices A1 and A2.

A3.1. The Geometry of Linear Transformations

The ultimate objective of this section of the appendix is to develop the *Spectral Decomposition Theorem* for symmetric matrices, that illuminates many of the most important properties of covariance matrices. But to gain an intuitive understanding of this result, it is important to understand the geometry of linear transformations as represented by matrices. A *transformation*, T , on \mathbb{R}^n is simply a mapping that assigns every vector, $x \in \mathbb{R}^n$, to some other vector, $T(x) \in \mathbb{R}^n$, called the *image* of x under T . A transformation, T , is *linear* if and only if (*iff*) it preserves vector addition, i.e., *iff* for each pair of vectors, $x, y \in \mathbb{R}^n$, and scalars, $\alpha, \beta \in \mathbb{R}$,

$$(A3.1.1) \quad T(\alpha x + \beta y) = \alpha T(x) + \beta T(y)$$

The intimate connection between matrices and linear transformations is seen most readily in \mathbb{R}^2 . If we let $e_1 = (1, 0)' \in \mathbb{R}^2$ and $e_2 = (0, 1)' \in \mathbb{R}^2$ denote the so-called *identity basis* vectors in \mathbb{R}^2 (shown in Figure A3.1 below)¹,




Figure A3.1. Identity Basis




Figure A3.2. Basis Representation

then by definition any vector, $x = (x_1, x_2)' \in \mathbb{R}^2$ can be represented as:

$$(A3.1.2) \quad x = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = x_1 \begin{pmatrix} 1 \\ 0 \end{pmatrix} + x_2 \begin{pmatrix} 0 \\ 1 \end{pmatrix} = x_1 e_1 + x_2 e_2$$

¹ Note that we maintain the convention that all vectors are represented as *column* vectors, so that transpose notation is used for all inline representations [as in expression (1.1.2) of Part II].

This *basis representation* of x , shown in Figure A3.2, implies from (A3.1.1.) that the image of x under any linear transformation, T , can be represented as

$$(A3.1.3) \quad T(x) = T(x_1e_1 + x_2e_2) = x_1T(e_1) + x_2T(e_2)$$

So if we know where the identity basis vectors, (e_1, e_2) , are sent by T , then we can *construct the entire transformation*. In particular, if we now let

$$(A3.1.4) \quad T(e_1) = a_1 = \begin{pmatrix} a_{11} \\ a_{12} \end{pmatrix}, \quad T(e_2) = a_2 = \begin{pmatrix} a_{21} \\ a_{22} \end{pmatrix}$$

then this transformation can be represented for all $x = (x_1, x_2)' \in \mathbb{R}^2$ by

$$(A3.1.5) \quad T(x) = x_1T(e_1) + x_2T(e_2) = x_1 \begin{pmatrix} a_{11} \\ a_{12} \end{pmatrix} + x_2 \begin{pmatrix} a_{21} \\ a_{22} \end{pmatrix} \\ = \begin{pmatrix} a_{11}x_1 + a_{21}x_2 \\ a_{12}x_1 + a_{22}x_2 \end{pmatrix} = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = Ax$$

where the matrix,

$$(A3.1.6) \quad A = (a_1, a_2) = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix}$$

is designated as the *matrix representation* of transformation T . This is the fundamental relation between matrices and linear transformations. In fact, it is so fundamental that linear transformations are usually *defined* by their matrix representations as in (A3.1.5). So in the two-dimensional case, each linear transformation can be defined by its matrix representation, A , for all $x = (x_1, x_2) \in \mathbb{R}^2$ as in Figures A3.3 and A3.4 below:




Figure A3.3. Basis Image Vectors

Figure A3.4. General Image Vectors

More generally, if the *identity basis*² in \mathbb{R}^n is associated with the columns of the *identity matrix*, $I_n = (e_1, e_2, \dots, e_n)$, and if the images of these basis vectors under any linear transformation, T , are denoted by $T(e_i) = a_i = (a_{i1}, \dots, a_{in})' \in \mathbb{R}^n$, $i = 1, \dots, n$, then for all $x = (x_1, \dots, x_n)' \in \mathbb{R}^n$, T again has the matrix representation

$$(A3.1.7) \quad T(x) = \sum_{i=1}^n x_i T(e_i) = \sum_{i=1}^n x_i a_i = \sum_{i=1}^n x_i \begin{pmatrix} a_{i1} \\ \vdots \\ a_{in} \end{pmatrix} \\ = \begin{pmatrix} \sum_{j=1}^n a_{1j} x_j \\ \vdots \\ \sum_{j=1}^n a_{nj} x_j \end{pmatrix} = \begin{pmatrix} a_{11} & \cdots & a_{1n} \\ \vdots & \ddots & \vdots \\ a_{n1} & \cdots & a_{nn} \end{pmatrix} \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix} = Ax$$

In the analysis to follow, we shall use the terms *matrix* and *transformation* interchangeably. Note also that is this equivalence that motivates the basic *multiplication rules* of matrix algebra. So the *meaning* of these rules is often best understood in this way.

To examine some of the more important matrix properties, we begin by observing that every matrix can be written in two equivalent ways. First there is a *column representation* of A ,

$$(A3.1.8) \quad A = \begin{pmatrix} a_{11} & \cdots & a_{1n} \\ \vdots & \ddots & \vdots \\ a_{n1} & \cdots & a_{nn} \end{pmatrix} = \left(\begin{bmatrix} a_{11} \\ \vdots \\ a_{n1} \end{bmatrix} \cdots \begin{bmatrix} a_{1n} \\ \vdots \\ a_{nn} \end{bmatrix} \right) = (a_1, \dots, a_n)$$

where a_j denotes the j^{th} column of A . There is also a *row representation* of A ,

$$(A3.1.9) \quad A = \begin{pmatrix} a_{11} & \cdots & a_{1n} \\ \vdots & \ddots & \vdots \\ a_{n1} & \cdots & a_{nn} \end{pmatrix} = \begin{pmatrix} [a_{11} & \cdots & a_{1n}] \\ \vdots \\ [a_{n1} & \cdots & a_{nn}] \end{pmatrix} = \begin{pmatrix} a'_1 \\ \vdots \\ a'_n \end{pmatrix}$$

where a'_i denotes the i^{th} row of A .³ This in turn implies that matrix products, AB , can be written in two ways:

² A fuller discussion of *vectors bases* for linear spaces is given on page A3-16 below.

³ It is important to note, for example, that a'_i in (A3.1.9) is *not* the transpose of a_i in (A3.1.10). To be more precise here, one could use the “dot” notation, $a_{\cdot j}$, for columns and $a_{i \cdot}$, for rows. However, we choose not to add this notational complexity since the rows and columns of A will generally be clear in context.

$$(A3.1.10) \quad AB = (a_1, \dots, a_n) \begin{pmatrix} b'_1 \\ \vdots \\ b'_n \end{pmatrix} = \sum_{i=1}^n a_i b'_i$$

and

$$(A3.1.11) \quad AB = \begin{pmatrix} a'_1 \\ \vdots \\ a'_n \end{pmatrix} (b_1, \dots, b_n) = \begin{pmatrix} a'_1 b_1 & \cdots & a'_1 b_n \\ \vdots & \ddots & \vdots \\ a'_n b_1 & \cdots & a'_n b_n \end{pmatrix}$$

Both of these representations are very useful, and will be used throughout the analysis to follow. As one immediate application, it is important to note that for every matrix, $A = (a_{ij} : i, j = 1, \dots, n)$, the *transpose* matrix, $A' = (a_{ji} : i, j = 1, \dots, n)$, represents a linear transformation closely related to that of A . In particular, the rows of A are the columns of A' . So from a transformation viewpoint, A' represents the “row space” of A . Moreover, if for any matrices A and B we use the representations

$$(A3.1.12) \quad A = \begin{pmatrix} a'_1 \\ \vdots \\ a'_n \end{pmatrix} \Rightarrow A' = (a_1, \dots, a_n) \quad , \quad B = (b_1, \dots, b_n) \Rightarrow B' = \begin{pmatrix} b'_1 \\ \vdots \\ b'_n \end{pmatrix}$$

then (A3.1.11) together with the identity, $a'b = b'a$, imply that,

$$(A3.1.13) \quad B'A' = \begin{pmatrix} b'_1 \\ \vdots \\ b'_n \end{pmatrix} (a_1, \dots, a_n) = \begin{pmatrix} b'_1 a_1 & \cdots & b'_1 a_n \\ \vdots & \ddots & \vdots \\ b'_n a_1 & \cdots & b'_n a_n \end{pmatrix} = \begin{pmatrix} a'_1 b_1 & \cdots & a'_1 b_n \\ \vdots & \ddots & \vdots \\ a'_n b_1 & \cdots & a'_n b_n \end{pmatrix}' = (AB)'$$

and hence that the transpose of a product, AB , is the product of their transposes in the *reverse* order.

A3.1.1 Nonsingular Transformations and Inverses

Perhaps the single most important feature of a linear transformation is whether or not it has an “inverse”. In particular, a linear transformation A is said to be *nonsingular* iff there exists another linear transformation, A^{-1} , called the *inverse* of A such that

$$(A3.1.14) \quad A^{-1}A = I_n$$

This inverse transformation can be equivalently defined by the requirement that for all $x, y \in \mathbb{R}^n$,

$$(A3.1.15) \quad A^{-1}y = x \Leftrightarrow Ax = y$$

This version also shows that $AA^{-1} = I_n$. For if we let $X = (x_1, \dots, x_n)$ be defined by $Ax_i = e_i$, $i = 1, \dots, n$, so that $AX = I_n$, then by (A3.1.15), $A^{-1}e_i = x_i$, $i = 1, \dots, n$ implies that $A^{-1} = A^{-1}I_n = X$, and hence that $AA^{-1} = I_n$. Note also that since A^{-1} is well defined as a transformation (i.e., $A^{-1}y$ is uniquely defined), it must be true that A is a *one-to-one* transformation, i.e., for all $x_1, x_2 \in \mathbb{R}^n$,

$$(A3.1.16) \quad x_1 \neq x_2 \Rightarrow Ax_1 \neq Ax_2$$

For if $Ax_1 = y = Ax_2$ then we would have $\{x_1, x_2\} \subseteq A^{-1}y$, so that $A^{-1}y$ is not uniquely defined. As an additional consequence of (A3.1.14), note that for any pair of nonsingular transformations, A and B , we must have

$$(A3.1.17) \quad (B^{-1}A^{-1})AB = B^{-1}(A^{-1}A)B = B^{-1}I_nB = B^{-1}B = I_n$$

Since the same argument shows that $AB(B^{-1}A^{-1}) = I_n$, it then follows from (A3.1.14) that AB must also be nonsingular, and in particular, has a well defined inverse $(AB)^{-1}$ given by

$$(A3.1.18) \quad (AB)^{-1} = B^{-1}A^{-1}$$

A similar argument shows that transposes, A' , of nonsingular matrices, A , must also be nonsingular. To see this, observe that we may take transposes of the matrices in (A3.1.14) and use (A3.1.13) to obtain

$$(A3.1.19) \quad (AA^{-1})' = I_n = (A^{-1}A)' \Rightarrow (A^{-1})'A' = I_n = A'(A^{-1})'$$

So by again appealing to (A3.1.14), we see that A' has a well-defined inverse, $(A')^{-1}$, given by

$$(A3.1.20) \quad (A')^{-1} = (A^{-1})'$$

In other words, the inverse of A' is the transpose of A^{-1} (so that the operations of taking transposes and inverses are said to *commute*).

To examine some of the more geometric properties of nonsingular transformations, observe that if for any set, $S \subseteq \mathbb{R}^n$, we let

$$(A3.1.21) \quad A(S) = \{Ax : x \in \mathbb{R}^n\}$$

denote the *image* of S under transformation A , then nonsingular transformations A must map \mathbb{R}^n onto itself, i.e.,

$$(A3.1.22) \quad A(\mathbb{R}^n) = \mathbb{R}^n$$

Since $A(\mathbb{R}^n) \subseteq \mathbb{R}^n$ by definition, (A3.1.22) follows from the observation that for any $x \in \mathbb{R}^n$, $A(A^{-1}x) = x \Rightarrow x \in A(\mathbb{R}^n)$, so that $\mathbb{R}^n \subseteq A(\mathbb{R}^n)$. In summary, every nonsingular transformation is both *one-to-one* and *onto* as a mapping.

We next observe that for all transformations, A , the full image set, $A(\mathbb{R}^n)$, is of special importance since it is always a *linear subspace* of \mathbb{R}^n , i.e., it is contained in \mathbb{R}^n and is closed under linear combinations $[x, y \in A(\mathbb{R}^n) \Rightarrow \alpha x + \beta y \in A(\mathbb{R}^n)]$ for all scalars, α, β . In particular, since $A(\mathbb{R}^n)$ contains all vectors that are expressible as a linear combinations of the columns of $A = (a_1, \dots, a_n)$, it is said to be *spanned* by these columns, and is often written as:

$$(A3.1.23) \quad \text{span}(A) = A(\mathbb{R}^n) = \{Ax : x \in \mathbb{R}^n\} = \left\{ \sum_{i=1}^n x_i a_i : x = (x_1, \dots, x_n) \in \mathbb{R}^n \right\}$$

In these terms, we note that one final characterization of nonsingular transformations (and perhaps the most basic characterization) is in terms of linearly independent vectors. A set of vectors $\{z_1, \dots, z_k\} \subseteq \mathbb{R}^n$ is said to be *linearly independent* if and only if for all scalars, $(\alpha_1, \dots, \alpha_k)$,⁴

$$(A3.1.24) \quad \sum_{i=1}^k \alpha_i z_i = 0 \Rightarrow \alpha_i = 0, i = 1, \dots, k$$

In these terms, a matrix, $A = (a_1, \dots, a_n)$, is *nonsingular* iff its columns $\{a_1, \dots, a_n\}$ are linearly independent. So by replacing z_i with a_i and α_i with x_i in this general definition, we can write this nonsingularity condition for A in matrix form as follows. For all $x \in \mathbb{R}^n$,

$$(A3.1.25) \quad Ax = 0 \Rightarrow x = 0$$

This characterization of nonsingularity is essentially equivalent to the uniqueness condition in (A3.1.16) [since $Ax = 0$ for $x \neq 0$ would imply that $Ax = 0 = A0$].

⁴ Note that for convenience we drop the subscript notation on the n -vector of zeros by $0_n = (0, \dots, 0)'$. The dimension of 0 should always be clear in context. So in expression (A3.1.24) for example, the 0 on the left is k -dimensional and the 0's on the right are scalars (one-dimensional).

These general properties of nonsingular transformations are well illustrated by the transformation, $A = (a_1, a_2)$, in Figures A3.3 and A3.4 above. Here it is evident that every vector in \mathbb{R}^2 is representable as a linear combination of $Ae_1 = a_1$ and $Ae_2 = a_2$, so that $\text{span}(A) = \mathbb{R}^2$. Similarly, the only linear combination which is the zero vector is the pair of zero scalars, $0 = (0, 0)'$, so that (A3.1.13) holds. Hence, even without producing the inverse transformation, A^{-1} , it should be clear that A is nonsingular.

In these notes, we shall deal almost exclusively with nonsingular transformations. But to understand the full scope of the matrix decomposition theorems to follow, it is important to consider *all* linear transformations on \mathbb{R}^n . In particular, those linear transformations, A , for which *no* inverse exists are said to be *singular* transformations. In terms of (A3.1.16) above, this means is that there are distinct vectors, $x \neq y$ with $Ax = Ay$, so that the transformation A^{-1} is not well defined. In view of linearity, this in turn implies that there is a *nonzero* vector, namely $x - y$, with $A(x - y) = Ax - Ay = 0$. This observation shows that the characterizing property of singular transformations, A , is that there is a nontrivial set of vectors mapped into zero by A . This set is designated as the *null space* for A , written as

$$(A3.1.26) \quad \text{null}(A) = \{x \in \mathbb{R}^n : Ax = 0\}$$

As the term “space” implies, $\text{null}(A)$ is also a linear space, since

$$(A3.1.27) \quad \begin{aligned} x, y \in \text{null}(A) &\Rightarrow Ax = 0 = Ay \\ &\Rightarrow A(\alpha x + \beta y) = \alpha Ax + \beta Ay = 0 \\ &\Rightarrow \alpha x + \beta y \in \text{null}(A) \end{aligned}$$

For a nonsingular transformation this is trivially true, since $\text{null}(A) = \{0\}$ by (A3.1.13). But for singular transformations, $\text{null}(A)$ is a proper linear space. In fact, the two linear spaces, $\text{span}(A)$ and $\text{null}(A)$ completely characterize most of the geometric features of every linear transformation. A simple example of a singular transformation, A , in \mathbb{R}^2 is given by

$$(A3.1.28) \quad A = \begin{pmatrix} 2 & 1 \\ 2 & 1 \end{pmatrix}$$

where for the vector, $x = (-1, 2)' \neq 0$, we see from (A3.1.28) that $Ax = 0$. Here $\text{span}(A)$ and $\text{null}(A)$ are shown in Figure A3.5 below.




Figure A3.5. Singular Transformation

The image vectors $Ae_1 = (2, 2)'$ and $Ae_2 = (1, 1)'$ are seen to be collinear, so that $\text{span}(A)$ is reduced to a line, i.e., a one-dimensional subspace of \mathbb{R}^2 . Similarly, the point x above is also shown, and is seen to generate a one-dimensional subspace, $\text{null}(A)$, which is collapsed into 0 by A . [An example in \mathbb{R}^3 is given in Figure A3.16 below.] More generally, the dimensions of these two subspaces always add to n . To be more precise, for any linear subspace, $S \subseteq \mathbb{R}^n$, the *dimension* of S , denoted by $\dim(S)$ is the maximum number of linearly independent vectors in S . So by (A3.1.23), the dimension of $\text{span}(A)$ must be the maximum number of linearly independent columns (a_1, \dots, a_n) of A . Moreover, by (A3.1.26) the dimension of $\text{null}(A)$ must be the maximum number of linearly independent vectors mapped to zero by A . As seen in Figure A3.5

$$(A3.1.29) \quad \dim(\text{span}(A)) + \dim(\text{null}(A)) = n$$

where in this case, $n = 2$. It turns out that this is *always* true. Since its validity will be apparent from the Singular Value Decomposition Theorem below, we shall not offer a proof of this “rank-nullity” theorem here.⁵

For our later purposes, it is important to note that the maximum number of linearly independent columns of any matrix, A , is also called the *rank* of A , written as $\text{rank}(A)$. When matrices are not square, [as for example in the *Linear Invariance Theorem* for multi-normal random vectors, stated both in expression (3.2.22) of Part II and in expression (A3.2.121) below], then it is useful to distinguish between columns and rows of matrix A by designating the *column rank* (*row rank*) of A to be the maximum number of linearly independent columns (rows) of A . In these terms, matrix A is said to be of *full column rank* (*full row rank*) iff all its columns (rows) are linearly independent, i.e., iff its column rank (row rank) is equal to the number of columns (rows) of A . In terms of linear

⁵ For an elegant on line proof see http://en.wikipedia.org/wiki/Rank%E2%80%93nullity_theorem.

transformations, the row rank of A can also be viewed as the rank of the linear transformation represented by A' .

With this general discussion of linear transformations, we now consider several specific types of transformations that will play a central role in the decomposition theorems to follow.

A3.1.2 Scale Transformations

While there are many different types of linear transformations, it turns out that from a geometric view point there are essentially only two basic transformation types. The first, and by far the simplest, are *scale transformations* that simply rescale the identity basis vectors, as in Figures A3.6 and A3.7 below:




Figure A3.6. Positive Scalars




Figure A3.7. General Scalars

Figure A3.6 represents a *positive* scalar transformation in which all basis vectors are scaled by positive multiples. In many cases, such transformations result from simply changing the measurement units (dollars, meters, etc.) of the variables represented by each axis. However, some scale transformations may involve negative multiples, as in Figure A3.7. The matrix representations, A_1 and A_2 , of these respective transformations are given by the diagonal matrices (with zeros omitted for visual clarity),

$$(A3.1.30) \quad A_1 = \begin{pmatrix} 2 & \\ & 3 \end{pmatrix}, \quad A_2 = \begin{pmatrix} 2 & \\ & -3 \end{pmatrix}$$

More generally, every *diagonal matrix*,

$$(A3.1.31) \quad A = \text{diag}(a_{11}, \dots, a_{nn}) = \begin{pmatrix} a_{11} & & \\ & \ddots & \\ & & a_{nn} \end{pmatrix}$$

is the representation of a *scale transformation* on \mathbb{R}^n . A key feature of these simple matrices is that multiplication of diagonal matrices is simply multiplication of their corresponding diagonal elements. In (A3.1.30) for example,

$$(A3.1.32) \quad A_1 A_2 = \begin{pmatrix} 4 & & \\ & & \\ & & -9 \end{pmatrix} = A_2 A_1$$

So like real numbers themselves, multiplication of diagonal matrices is *commutative*, i.e., in a sequence of successive scale transformations, the ordering of these transformations make no difference. One other key feature is that matrix inversion can be done by inspection, since it is evident from (A3.1.14) that the inverse of A in (A3.1.31) must be:

$$(A3.1.33) \quad A^{-1} = \begin{pmatrix} 1/a_{11} & & \\ & \ddots & \\ & & 1/a_{nn} \end{pmatrix} = \begin{pmatrix} a_{11}^{-1} & & \\ & \ddots & \\ & & a_{nn}^{-1} \end{pmatrix}$$

In other words, undoing a scale transformation amounts to scaling by its reciprocals.

A3.1.3 Orthonormal Transformations

The second important class of linear transformations is far richer, and in fact, is given many different names, including *isometric transformations*, *orthonormal transformations* and *rigid motions*. From a geometric viewpoint the term “isometric” is perhaps most appropriate, since these transformations preserve both distances and angles (as we shall see below). But from a matrix viewpoint, the term “orthonormal” is most useful since it relates more directly to the corresponding matrix representations, $U = (u_1, \dots, u_n)$, of such transformations. In particular, if both distances and angles are preserved, then since the vectors in the identity basis, $I_n = (e_1, \dots, e_n)$, are mutually orthogonal and of unit length, it follows that their images

$$(A3.1.34) \quad U(e_1, \dots, e_n) = (Ue_1, \dots, Ue_n) = (u_1, \dots, u_n)$$

under U must necessarily have the same properties. More precisely, [recalling property (A2.4.4) in Appendix A2] it must be true that

$$(A3.1.35) \quad u_i' u_i = \|u_i\|^2 = 1, \quad i = 1, \dots, n, \quad \text{and}$$

$$(A3.1.36) \quad u_i' u_j = 0, \quad i \neq j = 1, \dots, n$$

These *defining conditions* for orthonormality can be written in equivalent matrix form as

$$(A3.1.37) \quad U'U = \begin{pmatrix} u'_1 \\ \vdots \\ u'_n \end{pmatrix} (u_1, \dots, u_n) = \begin{pmatrix} u'_1 u_1 & \cdots & u'_1 u_n \\ \vdots & \ddots & \vdots \\ u'_n u_1 & \cdots & u'_n u_n \end{pmatrix} = \begin{pmatrix} 1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & 1 \end{pmatrix} = I_n$$

Note also from (A3.1.26) that this condition implies that U must be nonsingular, since

$$(A3.1.38) \quad Ux = 0 \Rightarrow U'Ux = U'0 = 0 \Rightarrow I_n x = 0 \Rightarrow x = 0$$

Finally since this in turn implies that

$$(A3.1.39) \quad U' = U'(UU^{-1}) = (U'U)U^{-1} = U^{-1}$$

we see that the inverse of U is simply its transpose. This an equivalent form of the defining condition in (A3.1.37), though the geometric argument above is far more intuitive. All geometric and algebraic properties of such transformations are in turn readily established from these equivalent conditions. The most immediate result is that all *inner products* must be preserved, since for any vectors, $x, y \in \mathbb{R}^n$,

$$(A3.1.40) \quad (Ux)'(Uy) = x'(U'U)y = x'y$$

This in turn implies that all distances (lengths) are preserved, since

$$(A3.1.41) \quad \|Ux\|^2 = (Ux)'(Ux) = x'x = \|x\|^2 \Rightarrow \|Ux\| = \|x\|$$

Finally, if θ denotes the angle between any pair of vectors, x and y , as in Figure A3.8 below,




Figure A3.8. Vector Angles

then since the *Law of Cosines* asserts that

$$(A3.1.42) \quad \cos(\theta) = \frac{\|x\|^2 + \|y\|^2 - \|y-x\|^2}{2 \|x\| \|y\|}$$

it follows at once from (A3.1.41) that U must also preserve angles. In other words, all geometric figures are mapped into *congruent* copies by U .

Another natural consequence of the defining condition, $U'U = I_n$, is that compositions (products) of orthonormal transformations, U_1U_2 , must also be orthonormal since

$$(A3.1.43) \quad (U_1U_2)'(U_1U_2) = U_2'(U_1'U_1)U_2 = U_2'(I_n)U_2 = U_2'U_2 = I_n$$

This same argument obviously holds for any finite product, $U_1U_2 \cdots U_n$.

Rotations and Reflections

Such orthonormal transformations can be further classified into *rotations* and *reflections*, as illustrated in \mathbb{R}^2 by Figures A3.9 and A3.10 below:




Figure A3.9. Rotation




Figure A3.10. Reflection

In Figure A3.9, transformation U defines a counterclockwise rotation of the plane through an angle θ , and in Figure A3.10, transformation U reflects the plane about the dashed line shown, so that images of all points above this line are their reflections below the line, and *visa versa*. Clearly both distances and angles are preserved in both cases. But one important difference is that clockwise orderings (called “orientations”) are different. In particular, planar rotations are seen to preserve orientation, while reflections do not.

Another key difference from a practical viewpoint relates to the extendibility of these concepts to higher dimensions. In particular, while rotations are easily defined with respect to angles in \mathbb{R}^2 , the extension of this definition to \mathbb{R}^n is highly complex (to say the least). However, the extension of reflections is completely straightforward. For the case of \mathbb{R}^3 , the *reflection line* in Figure A3.10 is simply replaced by a *reflection plane* through the origin. For example, the transformation, $U = [e_1, e_2, -e_3]$, is seen to reflect all points in \mathbb{R}^3 about the (e_1, e_2) plane. More generally, every reflection in \mathbb{R}^n is uniquely defined by a $(n-1)$ -dimensional *reflection hyperplane* through the origin. In addition, such reflections can be given a unified matrix representation as we now show.

Householder Reflections

Observe that each $(n-1)$ -dimensional hyperplane is in fact the *orthogonal complement* of a single vector in \mathbb{R}^n . In the illustration above, the (e_1, e_2) plane can be characterized as the orthogonal complement of e_3 . More generally, if for any vector, $v \in \mathbb{R}^n - \{0\}$, we let

$$(A3.1.44) \quad v^\perp = \{x \in \mathbb{R}^n : x'v = 0\}$$

denote the *orthogonal complement* of v , then the reflection about this hyperplane through the origin is representable by the *Householder matrix*,

$$(A3.1.45) \quad H_v = I_n - \left(\frac{2}{v'v}\right)vv'$$

To see this, note first that

$$(A3.1.46) \quad H_v v = \left[I_n - \left(\frac{2}{v'v}\right)vv' \right] v = v - \left(\frac{2}{v'v}\right)v(v'v) = v - 2v = -v$$

so that the image of v is precisely its reflection (through the origin) about v^\perp . Moreover, for any $x \in v^\perp$ it also follows that

$$(A3.1.47) \quad H_v x = \left[I_n - \left(\frac{2}{v'v}\right)vv' \right] x = x - \left(\frac{2}{v'v}\right)v(v'x) = x - 0 = x$$

But since H_v is completely defined by this set of images, it then follows that H_v must be the unique reflection in \mathbb{R}^n about v^\perp . This is shown graphically by the \mathbb{R}^2 example in Figure A3.11 below:




Figure A3.11. Householder Reflection

Finally, since every reflection has such a representation, it follows that *all reflections* are representable by Householder matrices, as in (A3.1.45). So all reflections are easily computable in \mathbb{R}^n .

From a geometric viewpoint, the importance of this fact is that *all orthogonal transformations* \mathbb{R}^n are constructible as compositions of (at most n) reflections. Alternatively phrased, *every n -square orthonormal matrix is the product of at most n Householder matrices*. Since this fact will not actually be used in our subsequent analyses, we will not prove it here (see footnote 2 below). Rather we simply illustrate this general result by showing how all rotations in \mathbb{R}^2 (such as in Figure A3.9) are equivalent to (at most) a pair of reflections in \mathbb{R}^2 . For any given angle, θ , let the corresponding (counterclockwise) rotation be denoted by, R_θ , as in Figure A3.12.




Figure A3.12. Angular Rotation

This is clearly not a reflection, and moreover cannot be equivalent to any single reflection, since this would necessarily reverse the clockwise order of the basis vectors, e_1 and e_2 , as mentioned above. But it can be represented as a composition of two reflections as follows. Choose the first (Householder) reflection, $H_1 = H_{v_1}$, by setting $v_1 = R_\theta e_1 - e_1$, and observe that by construction it reflects $R_\theta e_1$ back into e_1 , as shown in Figure A3.13 below:




Figure A3.13. First Reflection

Notice also that since every reflection is an orthonormal transformation, the image of $R_\theta e_2$ under H_1 must continue to be orthogonal to that of $R_\theta e_1$. But in two dimensions, there are only two possibilities (with unit length), namely e_2 and $-e_2$. In this case, $H_1 R_\theta e_2 = -e_2$, as shown in the figure. Finally, this configuration is easily reflected back into (e_1, e_2) by simply choosing $H_2 = H_{v_2}$ with generating vector, $v_2 = e_2 - (-e_2) = 2e_2$, so that the orthogonal complement, v_2^\perp , in this case is simply the horizontal axis, as shown in Figure A3.14 below.




Figure A3.14. Second Reflection

Finally, since all Householder matrices in (A3.1.45) are seen to be *symmetric*, we may then conclude that:

$$(A3.1.48) \quad H_2 H_1 R_\theta [e_1, e_2] = [e_1, e_2] \Rightarrow H_2 H_1 R_\theta = I_2 \Rightarrow H_1 R_\theta = H_2'$$

$$\Rightarrow R_\theta = H_1' H_2' \Rightarrow R_\theta = H_1 H_2$$

Hence, each such rotation is seen to be equivalent to this particular pair of reflections.⁶ So in this sense, Householder reflections can be regarded as the fundamental “generator” of all orthonormal transformations.

⁶ The proof of the general representation of orthonormal matrices by products of Householder matrices is surprisingly difficult to find in standard references. But one can easily show this by extending the standard Householder construction of QR decompositions (see for example the nice discussion by Tom Lyche available on line at <http://heim.ifi.uio.no/~tom/ortrans.pdf>) which shows in particular that every orthonormal matrix, U , can be represented as, $U = H_1 H_2 \cdots H_n T$, for some choice of Householder matrices, $H_1 H_2 \cdots H_n$, together with an upper triangular matrix, T . But by successive multiplications of this expression by H_i' , $i = 1, \dots, n$, together with (A3.1.37), we obtain $T = H_n' \cdots H_2' H_1' U$, which implies from (A3.1.20) that T must also be orthonormal. Finally, since a simple inductive argument can be used to show that the only orthonormal triangular matrix is the identity matrix, I_n , it then follows that $U = H_1 H_2 \cdots H_n$.

Orthonormal Bases and Extensions

One final aspect of orthonormality is important to consider. Recall that we have often referred to the (orthogonal) columns of the identity matrix, $I_n = (e_1, \dots, e_n)$, as the *identity basis* for \mathbb{R}^n . So before proceeding to the Singular Decomposition Theorem, it is appropriate to formalize the more general concept of *orthonormal bases*. First we extend the notion of $\text{span}(A)$ in expression (A3.1.23) to any set of vectors, $z_1, \dots, z_k \in \mathbb{R}^n$ as follows:

$$(A3.1.49) \quad \text{span}(z_1, \dots, z_k) = \left\{ \sum_{i=1}^k \alpha_i z_i : \alpha_1, \dots, \alpha_k \in \mathbb{R} \right\}$$

Hence a vector, $x \in \mathbb{R}^n$, lies in $\text{span}(z_1, \dots, z_k)$ iff x can be expressed as a linear combination of (z_1, \dots, z_k) , i.e., iff $x = \sum_{i=1}^k \alpha_i z_i$ for some scalars, $\alpha_1, \dots, \alpha_k$. Next, recalling the definition of linear independence in expression (A3.1.24) above, we now say that a set of linearly independent vectors, z_1, \dots, z_k , forms a *basis* for a given linear subspace, L , of \mathbb{R}^n iff

$$(A3.1.50) \quad \text{span}(z_1, \dots, z_k) = L$$

The special feature of linear independence is that for each, $x \in L$, the α -coefficients in the representation, $x = \sum_{i=1}^k \alpha_i z_i$, must be *unique*.⁷ So in geometric terms, these coefficients $(\alpha_1, \dots, \alpha_k)$ yield a natural *coordinate system* for L . Notice also that if (z_1, \dots, z_k) is a basis for L , then *no larger* set $(z_1, \dots, z_k, z_{k+1})$ can be a basis since $z_{k+1} \in L$ implies that z_{k+1} must already be a linear combination of (z_1, \dots, z_k) , which would violate linear independence. So the size, k , of each basis is a unique characteristic of L , designated as the *dimension* of L , and often written as $\dim(L)$.

The single most important example of these concepts is of course the *identity basis* (e_1, \dots, e_n) for \mathbb{R}^n itself. But this basis has the important additional feature that its component vectors form an *orthonormal set*, i.e., they are each of unit length and are mutually orthogonal [as we have already seen for the columns of orthonormal matrices in (A3.1.35) and (A3.1.36) above]. Any basis with these properties is called an *orthonormal basis*. The key feature of such bases is that coordinates of any vector, $x \in \text{span}(z_1, \dots, z_k)$, are immediately constructible as *inner products* with the basis vectors, i.e., for each $i = 1, \dots, k$,

⁷ To see this, note simply that if $\sum_{i=1}^k \alpha_i z_i = x = \sum_{i=1}^k \theta_i z_i$ then $\sum_{i=1}^k (\alpha_i - \theta_i) z_i = 0$, so that by linear independence, $\alpha_i = \theta_i$, $i = 1, \dots, k$.

$$(A3.1.51) \quad z_i'x = z_i' \sum_{j=1}^k \alpha_j z_j = \sum_{j=1}^k \alpha_j z_i' z_j = (1)\alpha_i + \sum_{j \neq i} (0) = \alpha_i$$

This is why orthonormal bases provide such useful representations of linear spaces. So it is important to ask how such bases can be constructed.

In particular, for any given set of vectors, $z_1, \dots, z_k \in \mathbb{R}^n$, we next consider how to construct an orthonormal basis for $span(z_1, \dots, z_k)$. There is a remarkably simple procedure for doing so, known as the *Gram-Schmidt orthogonalization procedure*. Because the geometry of this procedure is of such fundamental importance, we begin by considering *orthogonal projections*. Given two vectors, $x, y \in \mathbb{R}^n$ (as illustrated for $n = 2$ in Figure A3.15 below), one may ask what vector in the span of y is “closest” to x , or equivalently, “best approximates” x ?




Figure A3.15. Simple Orthogonal Projection

If one were to imagine drawing circles around x , denoting points of equal distance from x , then the smallest circle touching the line, $span(y) = \{\alpha y : \alpha \in \mathbb{R}\}$, would be just tangent to this line, and would identify the desired closest point, y_x (shown as red in the figure). Formally, this amounts to finding the α which minimizes the distance, $\|x - \alpha y\|$, from x . But since minimizing distance is equivalent to minimizing *squared* distance, it follows that if we now write $\phi(\alpha) = \|x - \alpha y\|^2$, as a function of α , then we can identify this point by solving the “least squares” minimization problem:

$$(A3.1.52) \quad \min_{\alpha} \phi(\alpha) = \|x - \alpha y\|^2 = (x - \alpha y)'(x - \alpha y) = x'x - 2\alpha x'y + \alpha^2 y'y$$

Since the last equality is just a quadratic function in α , the desired “tangency” is given precisely by the first-order condition:

$$(A3.1.53) \quad 0 = \frac{d}{d\alpha} \phi(\alpha) = -2x'y + 2\alpha y'y \Rightarrow \alpha y'y = x'y \Rightarrow \alpha = \frac{x'y}{y'y}$$

So the vector closest to x in $span(y)$ is given by

$$(A3.1.54) \quad y_x = \left(\frac{x'y}{y'y} \right) y$$

and is designated as the *orthogonal projection* of x on y . The term “orthogonal” is of most importance for our present purposes, and is motivated by the fact that the difference vector, $x - y_x$ (shown by the red dashed line in the figure) is necessarily orthogonal to y , as can be seen by taking inner products:

$$(A3.1.55) \quad (x - y_x)'y = x'y - y_x'y = x'y - \left(\frac{x'y}{y'y} \right) y'y = x'y - x'y = 0$$

So if one starts with two vectors, (x, y) , and wishes to construct an *orthonormal basis* for $\text{span}(x, y)$, then this projection procedure yields a natural choice. In particular, since $x - y_x = x - (x'y/y'y)y$ is automatically a linear combination of (x, y) , it follows that $(y, x - y_x)$ yields a pair of orthogonal vectors in $\text{span}(x, y)$. Hence, by normalizing these, we have found an orthonormal basis for $\text{span}(x, y)$.

This argument implicitly assumes that x and y are linearly independent, so that the basis will consist of two orthonormal vectors. But notice also that if x and y were linearly dependent, so that x was already in $\text{span}(y)$ [i.e., $x - \alpha y = 0$ for some α], then the solution in (A3.1.54) would automatically yield $y_x = x$ so that $(x - y_x)$ is simply the zero vector. In other words, this procedure would *identify* this linear dependence, and tell us that by normalizing only y we would obtain a natural orthonormal basis for $\text{span}(x, y) = \text{span}(y)$.

This two-vector example defines the simplest possible instance of the Gram-Schmidt procedure. So all that remains is to be done is to show how this procedure can be extended to larger sets of vectors. This extension is extremely simple, and only uses the two-vector procedure detailed above. To see this, let us proceed to a three vector case. Suppose we given linearly independent vectors, $z_1, z_2, z_3 \in \mathbb{R}^n (n \geq 3)$, and wish to construct an orthonormal basis (u_1, u_2, u_3) for $\text{span}(z_1, z_2, z_3)$. To do so, we first construct an orthogonal basis (b_1, b_2, b_3) as follows:

Step 1. Start by setting

$$(A3.1.56) \quad b_1 = z_1.$$

Step 2. Project z_2 on b_1 and construct the difference vector,

$$(A3.1.57) \quad b_2 = z_2 - \frac{z_2' b_1}{b_1' b_1} b_1 \quad \left(= z_2 - \frac{z_2' z_1}{z_1' z_1} z_1 \right)$$

As in the example above, (b_1, b_2) , are now orthogonal, and are both in $\text{span}(z_1, z_2, z_3)$.

Step 3. Finally, project z_3 on b_1 and b_2 individually and take the vector difference:

$$(A3.1.58) \quad b_3 = z_3 - \left(\frac{z_3' b_1}{b_1' b_1} \right) b_1 - \left(\frac{z_3' b_2}{b_2' b_2} \right) b_2$$

Then by construction, $(b_1, b_2, b_3) \in \text{span}(z_1, z_2, z_3)$. Moreover, since b_1 and b_2 are already orthogonal, it follows that b_3 must necessarily be orthogonal to *both* b_1 and b_2 . To see this, note simply that for b_1 we have

$$(A3.1.59) \quad \begin{aligned} b_1' b_3 &= b_1' z_3 - \left(\frac{z_3' b_1}{b_1' b_1} \right) b_1' b_1 - \left(\frac{z_3' b_2}{b_2' b_2} \right) b_1' b_2 \\ &= b_1' z_3 - (z_3' b_1) - \left(\frac{z_3' b_2}{b_2' b_2} \right) (0) = b_1' z_3 - b_1' z_3 = 0 \quad , \end{aligned}$$

and similarly for b_2 . Given this orthogonal basis, it then follows by setting

$$(A3.1.60) \quad u_i = b_i / \|b_i\| \quad , \quad i = 1, 2, 3$$

that we must obtain an orthonormal basis (u_1, u_2, u_3) for $\text{span}(z_1, z_2, z_3)$. Again, if (z_1, z_2, z_3) are not linearly independent, then we need only normalize the nonzero vectors obtained. This will not only provide an orthonormal basis for $\text{span}(z_1, z_2, z_3)$, but will also indicate the *dimension* of this linear space.

The generalization of this stepwise procedure follows by simple induction. In particular, to obtain an orthogonal basis for $\text{span}(z_1, \dots, z_k)$, suppose we have already obtained an orthogonal set (b_1, b_2, \dots, b_m) in $\text{span}(z_1, \dots, z_k)$ with $3 \leq m < k$. To extend this orthogonal set, let

$$(A3.1.61) \quad b_{m+1} = z_{m+1} - \sum_{i=1}^m \left(\frac{z_{m+1}' b_i}{b_i' b_i} \right) b_i \in \text{span}(z_1, \dots, z_k)$$

Then the argument in (A3.1.59) again shows that b_{m+1} is orthogonal to each b_i , $i = 1, \dots, m$. So by induction, we thus obtain an orthogonal basis (b_1, \dots, b_k) for $\text{span}(z_1, \dots, z_k)$. This in turn yields an *orthonormal basis* (u_1, \dots, u_k) by normalizing all nonzero vectors in (b_1, \dots, b_k) as in (A3.1.60). Moreover, the number of such vectors will again identify the *dimension* of $\text{span}(z_1, \dots, z_k)$.

One final possibility is of interest. Suppose that we are given an orthogonal basis, (b_1, \dots, b_k) for some $\text{span}(z_1, \dots, z_k)$ with $k < n$, and wish to extend this to an orthogonal basis *for all of* \mathbb{R}^n . This is again quite simple, since we already have a basis for \mathbb{R}^n , namely the identity basis, (e_1, \dots, e_n) . So to extend (b_1, \dots, b_k) to a larger orthogonal basis, $(b_1, \dots, b_k, b_{k+1})$ we may proceed by setting $m = k$ in (A3.1.61) and then successively letting $z_{k+1} = e_i$ for each $i = 1, \dots, n$ until a nonzero difference vector, b_{k+1} , is found. There must be one, since not all e_i can lie in the lower dimensional space, $\text{span}(z_1, \dots, z_k)$. Once b_{k+1} is found, the procedure can be repeated by setting $m = k + 1$ in (A3.1.61) and continuing down the list of identity basis vectors, e_i , until a new nonzero difference vector, b_{k+2} , is found. Again by induction, this procedure must result in a full set of basis vectors, (b_1, \dots, b_n) , which yield the desired extension. These can in turn be normalized as in (A3.1.60) to obtain an orthonormal basis, (u_1, \dots, u_n) , for \mathbb{R}^n . Finally, if the original basis is already orthonormal, say (u_1, \dots, u_k) , then this procedure is designated as an *orthonormal extension* of (u_1, \dots, u_k) to all of \mathbb{R}^n .

A3.2 Singular Value Decomposition Theorem

While there are of course many special types of matrices that are of analytical interest [as for example the triangular Cholesky decompositions of symmetric matrices in (A2.7.44) of Appendix A2], our focus above on diagonal matrices and orthonormal matrices was for a reason. In the same way that orthonormal matrices have a simple decomposition into reflections, it turns out that *every* n -square matrix, A , is decomposable into a simple product of orthonormal and diagonal matrices as follows:

$$(A3.2.1) \quad A = U S V'$$

where U and V are *orthonormal* and where $S = \text{diag}(s_1, \dots, s_n)$ is a *nonnegative diagonal* matrix with diagonal entries, s_i , called the *singular values* of matrix A . In geometric terms, every linear transformation is constructible as a composition of a nonnegative scale transformation together with two orthonormal transformations. This fundamental result, known as the *Singular Value Decomposition (SVD) Theorem*, holds for all matrices (even rectangular matrices). At this level of generality, it has been designated by Gilbert Strang (1993,2009) as the *Fundamental Theorem of Linear Algebra*.

The main objective of the present section is to establish this theorem. By way of motivation, recall from the beginning of these notes that our ultimate objective is to establish the *Spectral Decomposition (SPD) Theorem* for symmetric matrices, which asserts that every symmetric matrix, A , can be represented in terms of a single orthonormal matrix, W , and diagonal matrix, $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$, as

$$(A3.2.2) \quad A = W \Lambda W'$$

where the diagonal entries, λ_i , are called the *eigenvalues* of A (see Section A3.3 below). So except for the nonnegativity of S in (A3.2.1), it would appear that this important result is simply a special case of the SVD Theorem with $W = U = V$. As we shall see below, this intuition is correct in many important cases. Moreover, it is essentially correct in *all* cases in the sense that an SPD can always be constructed from any given SVD. It is this relationship that provides the main motivation for our consideration of this more general result. But as emphasized by Strang's renaming of this result, anyone interested in understanding linear transformations should try to gain some understanding of (A3.2.1) in its own right.

While proofs of the SVD Theorem can be found in most standard texts on matrix algebra, the most common approach is to *start* with the SPD Theorem and then apply this result to the partitioned symmetric matrix,

$$(A3.2.3) \quad M_A = \begin{pmatrix} & A \\ A' & \end{pmatrix}$$

in order to establish the SVD Theorem. But this “trick” offers little insight into the geometric origins of either result. So the specific objectives of this section are to illustrate these origins with an easily visualized geometric example in \mathbb{R}^2 , and then use these insights to motivate a constructive proof of the SVD Theorem.¹

To develop our geometric argument, we require one further characterization of orthonormal transformations, V . Recall that all such transformations preserve distances. Conversely, to guarantee that V is orthonormal, it is enough to require that *all unit distances be preserved by V* , i.e., that for all $x \in \mathbb{R}^n$,

$$(A3.2.4) \quad \|x\| = 1 \Rightarrow \|Vx\| = 1$$

To see this, note first that since any vector, $x \in \mathbb{R}^n$, can be transformed to have unit length by the rescaling, $x \rightarrow \frac{1}{\|x\|}x$, it follows from (A3.2.4) that *all distances* must be preserved, since

$$(A3.2.5) \quad \left\| \frac{1}{\|x\|}x \right\| = \frac{1}{\|x\|}\|x\| = 1 \Rightarrow 1 = \left\| V \left(\frac{1}{\|x\|}x \right) \right\| = \frac{1}{\|x\|}\|Vx\| \\ \Rightarrow \|Vx\| = \|x\|$$

Moreover by observing from the identity

$$(A3.2.6) \quad \|x - y\|^2 = (x - y)'(x - y) = x'x - 2x'y + y'y = \|x\|^2 - 2x'y + \|y\|^2 \\ \Rightarrow x'y = \frac{1}{2}(\|x\|^2 + \|y\|^2 - \|x - y\|^2)$$

that inner products are entirely expressible in terms of distances, it then follows from (A3.2.5) that all inner products must be preserved as well. Hence the defining conditions for orthonormality in (A3.1.34) and (A3.1.35) must hold, and V is orthonormal.

Given this alternative characterization, we next observe that the product of matrices on the right hand side of (A3.2.1) can be directly interpreted geometrically as an orthonormal transformation, V' , followed by a rescaling, S , followed by a second orthonormal transformation, U .² But while this composite transformation is of course linear, the key question remains as to why *every* linear transformation, A , can be so represented. Assuming that A is nonsingular (so that its inverse exists), a more informative geometric approach is to start with transformation, A , and see how to “undo it” (i.e., invert it back to the identity) through a series of simple transformations. For the two dimensional case, this process can be illustrated by the four panels shown in Figure A3.16 below.

¹ Another excellent development of these ideas can be found in Section 5.12 of Meyer (2001).

² This is illustrated for example in Figure 6.8 of Strang (2009, p.366).

Starting from the upper left panel, suppose that a given transformation, A , maps the basis vectors (e_1, e_2) in \mathbb{R}^2 as shown in the upper right panel. In geometric terms, the key here is to consider not only how these basis vectors are transformed, but also how the entire *unit circle* (shown in blue) is transformed. In \mathbb{R}^2 the image of this circle is always some *ellipse*, as shown (in blue) in the upper right panel. Since the unit circle consists of all vectors of unit length, we see that some of these vectors will typically be “stretched” more than others by transformation A . In particular, since the *major axis* and *minor axis* of this ellipse (shown as thin blue lines) denote the directions of maximum and minimum distances from the origin, it follows that the vector on the unit circle which is “maximally stretched” by A must be the vector (not shown) that is mapped into the major axis of this ellipse. Similarly, the vector that is “minimally stretched” is mapped into the minor axis.

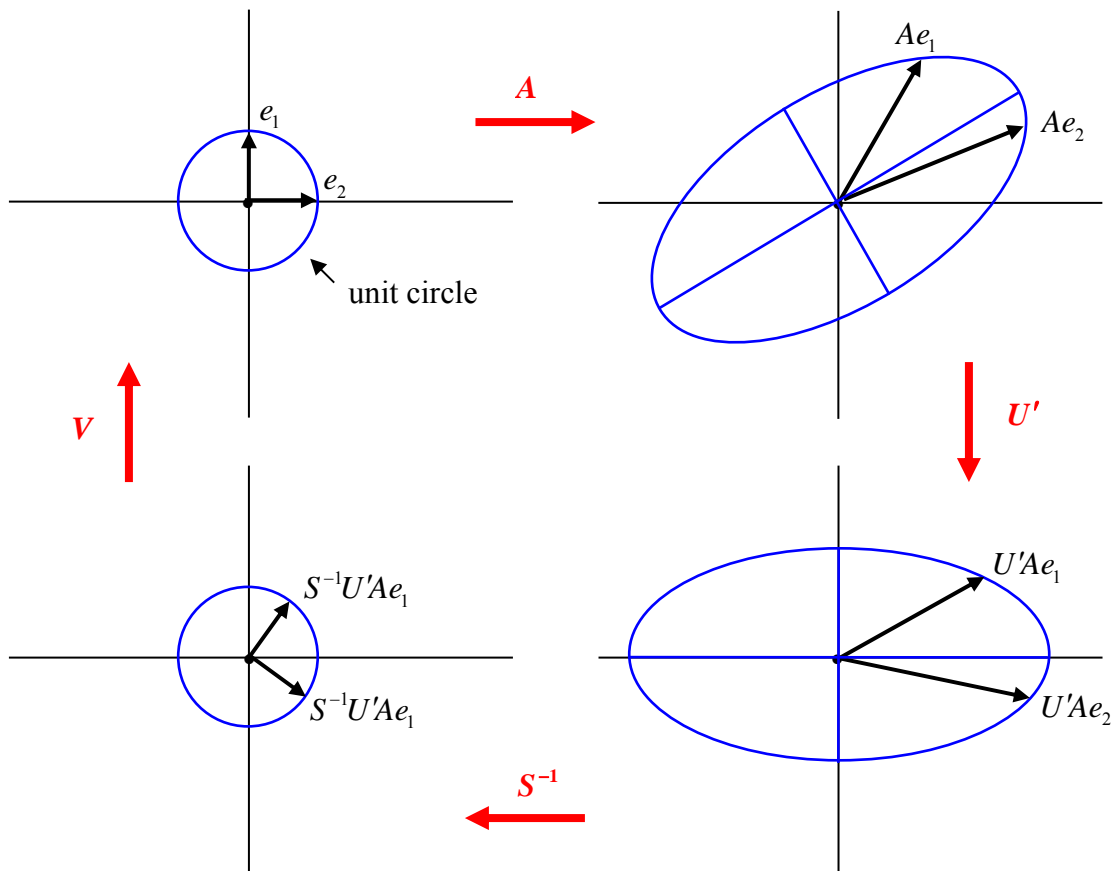


Figure A3.16. Geometry of SVD

So to remove all stretch effects, the simplest procedure is to rotate these (orthogonal) axes into the coordinate axes, and then rescale them back to unit lengths. The appropriate rotation is shown in the lower right panel, and is represented by an *orthonormal* matrix,

U' . The *rescaling* back to unit lengths is then shown in the lower left panel, and is represented by a *positive diagonal* matrix, S^{-1} . Notice also that by scaling the maximum and minimum lengths to unity, *all* intermediate lengths must also be scaled to unity.³ So the ellipse again becomes a unit circle. What this implies is that the transformation represented by the product, $S^{-1}U'A$, has actually *mapped the unit circle back into itself*. So if we now denote this product matrix by

$$(A3.2.7) \quad V' = S^{-1}U'A$$

then it follows that V' must satisfy (A3.2.4), and hence must be *orthonormal*. In particular, the images of e_1 and e_2 under this transformation (namely the two vectors, $S^{-1}U'Ae_1$ and $S^{-1}U'Ae_2$, shown in the lower left panel of Figure A3.16) must be *orthogonal*. So by construction we may use (A3.1.38) to conclude that

$$(A3.2.8) \quad S^{-1}U'A = V' \Rightarrow U'A = SV' \Rightarrow A = USV'$$

and thus that A is representable as in (A3.2.1) [where in this nonsingular case, S must be a positive diagonal matrix].

While this argument is quite transparent in \mathbb{R}^2 , it is more complex in higher dimensions. In particular, if the unit circle is now replaced by the *unit sphere* in \mathbb{R}^n ,

$$(A3.2.9) \quad \mathbb{S}_n = \{x \in \mathbb{R}^n : \|x\| = 1\}$$

then one can in principle construct similar arguments for the *ellipsoidal images*,

$$(A3.2.10) \quad A(\mathbb{S}_n) = \{Ax \in \mathbb{R}^n : x \in \mathbb{R}^n\}$$

of \mathbb{S}_n under linear transformations, A . The basic ideas can be illustrated for \mathbb{R}^3 as shown in Figure A3.17 below.

³ To show this formally, observe first that the equation of the ellipse in the lower right panel must be of the form, $a_1x_1^2 + a_2x_2^2 = c$ for some positive constants, a_1, a_2, c . So for the principle axes of this ellipse, say $(x_{01}, 0)$ and $(0, x_{02})$, it must be true that $a_1x_{01}^2 = c = a_2x_{02}^2$. But if the given scale transformation is denoted by, $S^{-1} = \text{diag}(s_1^{-1}, s_2^{-1})$, so that $S^{-1}x = (s_1^{-1}x_1, s_2^{-1}x_2)$, then for this unit scaling it must be also true that $s_1^{-1}x_{01} = 1 = s_2^{-1}x_{02}$, so that $x_{01} = s_1$ and $x_{02} = s_2$. These two relations together imply that $a_1 = s_1^{-2}c$ and $a_2 = s_2^{-2}c$ so that $c = a_1x_1^2 + a_2x_2^2 = (s_1^{-2}c)x_1^2 + (s_2^{-2}c)x_2^2$. Finally, by canceling c on both sides, we see that, $1 = (s_1^{-1}x_1)^2 + (s_2^{-1}x_2)^2 = \|(s_1^{-1}x_1, s_2^{-1}x_2)\|^2$, and may thus conclude that $\|(s_1^{-1}x_1, s_2^{-1}x_2)\| = 1$, i.e., that all transformed vectors $(s_1^{-1}x_1, s_2^{-1}x_2)$ have unit length.




Figure A3.17 Example in Three Dimensions

In this example, the unit sphere, \mathbb{S}_3 , shown (in red) on the left is mapped by the linear transformation,

$$(A3.2.11) \quad A = \begin{pmatrix} 0.7 & 0 & 0 \\ 0 & 1.8 & 0 \\ 0 & 0.7 & 0.7 \end{pmatrix}$$

into the ellipsoidal image set, $A(\mathbb{S}_3)$, shown (in red) on the right. The details of this example will be discussed further as we proceed. But for the moment, it should be clear that the first principle axis (major axis) of this ellipsoid is the line through the origin (not shown) that connects the two ends of this “football-shaped” set. So the point labeled, Av_1 , (to be discussed below) is the image of a point, $v_1 \in \mathbb{S}_3$, which is “maximally stretched” by transformation, A . The location of this particular point, v_1 , is shown on \mathbb{S}_3 (just below the x_2 axis). So by linearity, the other maximally stretched point in \mathbb{S}_3 (not shown) must be just opposite to v_1 on the line from v_1 through the origin. Note also that the second principle axis is a line through the origin which is orthogonal to the first principle axis and passes through the point labeled, Av_2 , in the figure.

While it is possible to construct an orthonormal transformation that rotates these axes into the coordinate axes, and then rescale the ellipsoidal image back to a sphere as in Figure A3.16 above, the details of such a construction are extremely tedious (especially in higher dimensions). Hence the two most important features of the argument in Figure A3.16 are (i) its graphical simplicity in \mathbb{R}^2 , and (ii) its role in suggesting a more tractable approach to the SVD Theorem in \mathbb{R}^n . In particular, this approach is motivated by the observation that the critical task in the above argument is to *identify* those unit vectors in

\mathbb{S}_n that are mapped by A into the principle axes of the ellipsoidal image, $A(\mathbb{S}_n)$, so that the appropriate rotations can be defined. Note in particular that the vector mapped into the major axis of the ellipse in Figure A3.16 (or ellipsoid in Figure A3.17) is by definition that unit vector, v_1 , with *maximal image length*, $\|Av_1\|$. So the most natural procedure for identifying v_1 is to solve the maximization problem

$$(A3.2.12) \quad \text{maximize: } \|Av\| \quad \text{subject to: } v \in \mathbb{S}_n$$

There will of course be two solutions, corresponding to each end of the ellipse (or ellipsoid). But this vector is essentially unique up to a choice of direction. The second key point established for the case of \mathbb{R}^2 was that the vector mapped into the minor axis is necessarily determined (up to a choice of direction) as one *orthogonal* to v . In the case of \mathbb{R}^2 , this was established by verifying that the transformation, V' , in (A3.2.7) was *orthonormal*. In higher dimensions, a direct proof of this fact is much more difficult. So our approach will be to start by *assuming* that this is the case, and use this assumption to construct a sequence of maximization problems similar to (A3.2.12). The final solutions to these problems will be seen to yield precisely desired representation in (A3.2.1), and thus show (among other things) that V' in (A3.2.7) is indeed orthonormal in all cases.

Before developing this sequential maximization procedure, it is appropriate to make a few preliminary remarks. First of all, this approach to establishing the SVD Theorem is known in the literature as the “variational” approach, and is in fact one of the oldest approaches to this problem.⁴ Second, it turns out that there is a more useful way of representing image lengths, $\|Av\|$, that will be seen to have added benefits in the following analysis. In particular, if for any vector, $v \in \mathbb{S}_n$, the image vector, Av , is simply rescaled to a vector of unit length, $u \in \mathbb{S}_n$, as shown (in red) for $n = 2$ in Figure A3.18 below,




Figure A3.18. Rescaling Convention

then by construction

⁴ See Stewart (1993) for an interesting historical discussion of this variational approach, which goes back to the work of Jordan in the 1870's.

$$(A3.2.13) \quad Av = su$$

for some scalar, s . Note that if $Av = 0$ then (A3.2.13) will hold trivially for $s = 0$. While we shall eventually need to deal with this degenerate case, we focus for the present on vectors, $v \in \mathbb{R}^n$, with $Av \neq 0$ [i.e., $v \notin \text{null}(A)$] so that $s \neq 0$. Moreover, by replacing u with $-u$ if necessary, we can always ensure that $s > 0$, so that by construction, $\|Av\| = s\|u\| = s > 0$. Thus, as an alternative to (A3.2.12), one can find the direction, v , of maximal stretch by solving the associated maximization problem:

$$(A3.2.14) \quad \text{maximize: } s = s(v, u) \quad \text{subject to: } Av = su, \quad \|u\| = 1, \quad \|v\| = 1$$

Note also that since $u'u = \|u\|^2$ for any vector, u , it follows from the first constraint that

$$(A3.2.15) \quad u'Av = su'u = s\|u\|^2 = s$$

Hence (A3.2.14) can be simplified to

$$(A3.2.16) \quad \text{maximize: } u'Av \quad \text{subject to: } u'u = 1, \quad v'v = 1$$

As we shall see below, the advantage of this alternative formulation is that will allow us to solve simultaneously for all three matrices, U , S , and V in (A3.2.1), where u and v will turn out to be column vectors of U and V respectively, and where $s (= u'Av)$ will be the diagonal elements of S . This constrained maximization problem thus constitutes the center piece of the present analysis, and will be used recursively to construction the full SVD representation for arbitrary linear transformations.

Before doing so, it is important to note finally that (A3.2.16) must *always* have a solution. While this may seem obvious in our original two dimensional problem, it is less so in higher dimensions. In particular, since the objective function, $u'Av$, in (A3.2.16) is a *bilinear form* in u and v (i.e., it is linear in u for each fixed v , and linear in v for each fixed u) there are no natural maxima or minima for this function. But the existence of such solutions follows from what is usually called the *Extreme Value Theorem*. The one-dimensional version simply states that every continuous function, $f(x)$, on a closed bounded interval, $[a, b] \subset \mathbb{R}$, has both a maximum and minimum value. This can be seen intuitively as in Figure A3.19 below:




Figure A3.19. Extreme Values

The generalized version simply shows that same is true for continuous functions on nonempty closed bounded sets in any finite-dimensional space, \mathbb{R}^N .⁵ In the present case, the bilinear form, $f(u, v) = u'Av$ is a continuous function on \mathbb{R}^{2n} constrained to the product of unit spheres, $\mathbb{S}_n \times \mathbb{S}_n = \{u \in \mathbb{R}^n : \|u\| = 1\} \times \{v \in \mathbb{R}^n : \|v\| = 1\} \subset \mathbb{R}^n \times \mathbb{R}^n = \mathbb{R}^{2n}$, which is easily seen to be a nonempty closed bounded set in \mathbb{R}^{2n} . Hence there always exists a maximum solution to (4). Moreover, since both the objective function, $f(u, v)$, and constraint functions, $u'u$ and $v'v$, are *continuously differentiable* on \mathbb{R}^{2n} , this maximum can be characterized by the first order conditions of the associated Lagrangian function [recall expression (A2.8.38) in Section 8 of the Appendix, A2, to Part II of these notes]:

$$(A3.2.17) \quad L(u, v, s, \theta) = u'Av + \frac{1}{2} [s(1 - u'u) + \theta(1 - v'v)]$$

(where the factor of $\frac{1}{2}$ is introduced for notational convenience only). By using expressions (A2.7.7) and (A2.7.11) [with $A = I_n$] in Appendix A2, we see that the first order conditions for u and v are given respectively by

$$(A3.2.18) \quad 0 = \nabla_u L = Av + \frac{1}{2}[-2su + 0] = Av - su \Rightarrow Av = su$$

$$(A3.2.19) \quad 0 = \nabla_v L = A'u + \frac{1}{2}[0 - 2\theta v] = A'u - \theta v \Rightarrow A'u = \theta v$$

where (A3.2.19) also uses the identity, $\nabla_v(u'Av) = (u'A)' = A'u$. Similarly, the first order conditions for s and θ reduce to the constraints

$$(A3.2.20) \quad u'u = 1 = v'v$$

At this point, notice that conditions (A3.2.18) and (A3.2.20) are simply the constraints in (A3.2.14) that originally motivated this formulation. In particular, transformation A must achieve its maximum stretch, s , at vector v . Hence the most important new information provided by this solution is condition (A3.2.19), which shows that there is a parallel relation for the *transpose*, A' , of A . In particular, the same argument leading to (A3.2.14) shows that maximum stretch, θ , of the transpose transformation, A' , must be achieved at vector u , so that there is a clear duality between these two transformations. Moreover, by the symmetry of inner products, it follows from (A3.2.18), (A3.2.19) and (A3.2.20) that

$$(A3.2.21) \quad s = s(u'u) = u'(su) = u'(Av) = v'(A'u) = v'(\theta v) = \theta(v'v) = \theta$$

So in fact this maximum stretch value must be the *same* for both A and A' .

⁵ See for example Theorem 16.6 in Bartle (1964). More general versions can found in Murphy (2008).

Before extending this argument to obtain the SVD representation (A3.2.1) for all matrices, A , it is essential to distinguish between the nonsingular and singular cases. Figure A3.17 above illustrates a typical nonsingular case, which is by far the most important case for all applications that we consider in these notes. However, since this same representation also holds for singular matrices, it is instructive to see what this means for the geometry of linear transformations. To illustrate the basic differences between these two cases, we now consider the following modification of transformation, A , in expression (A3.2.11) above

$$(A3.2.22) \quad A_0 = \begin{pmatrix} 0.7 & 0 & 0 \\ 0 & 1.8 & 0 \\ 0 & 0.7 & 0 \end{pmatrix}$$

Here the matrix, A_0 , differs from A in only the third column, which is now the zero vector. This of course implies that $A_0 e_3 = 0$, and hence that A_0 is singular. The corresponding modification of Figure A3.17 is shown in Figure A3.20 below.




Figure A3.20 A Singular Example in Three Dimensions

The key difference here is that $\text{span}(A_0)$ is now a *two-dimensional* plane (shown in blue). So the ellipsoid in Figure A3.17 has now been collapsed into an ellipse on this plane. Notice also that even though \mathbb{S}_3 is only the surface of an ellipsoidal solid in \mathbb{R}^3 , the image set, $A_0(\mathbb{S}_3)$, consists of the full area inside the ellipse on the right (including the origin). But the initial maximization problem in (A3.2.12) above is still well defined, and is seen to have a solution very similar to the full-dimensional case in Figure A3.17. Notice in particular that the analysis of this ellipse in $\text{span}(A_0)$ is qualitatively the same as that for the ellipse in upper right panel of Figure A3.17 for the \mathbb{R}^2 case. More generally, it will turn out that for any singular matrix, A , one proceeds by first analyzing

the ellipsoid in $\text{span}(A)$, and then extending this analysis to the collapsed dimensions in $\text{null}(A)$ in order to complete the SVD representation.

This extension process is most transparent in \mathbb{R}^2 . So before proceeding with the formal argument, it is instructive to reconsider the singular example expression (A3.1.28) together with Figure A3.5. This figure is reproduced in Figure A3.21 below, where the unit circle, \mathbb{S}_2 , is now included. The image set, $A(\mathbb{S}_2)$, is given by the red line segment shown, which by definition lies in $\text{span}(A)$. So the possible solution vectors, $v \in \mathbb{S}_2$, in (A3.2.12) are seen to be either v_1 or \tilde{v}_1 , with images, Av_1 and $A\tilde{v}_1$, corresponding to the end points of interval, $A(\mathbb{S}_2)$, as shown in the figure. For purposes of discussion, we now focus on v_1 . In this case, the full solution to this maximization problem is given by the triple, (v_1, s_1, u_1) , where u_1 is the unit-scaled version of Av_1 (shown by the red point), with scale factor, s_1 , denoting the maximum-stretch value, i.e., $Av_1 = s_1 u_1$.




Figure A3.21 Singular Example in Two Dimensions

To complete the desired Singular Value Decomposition of matrix A in (A3.1.28), we would like to find a unit vector, $v_2 \in \mathbb{S}_2$, that is mapped by A in the minor axis of this one-dimensional “ellipse”. But while the major axis is well defined, there appears to be no meaningful minor axis. Here is where we use our assumption above that the end points of this “axis” must be the images of unit vectors *orthogonal* to v_1 . If so, then there are seen to be only two possible choices, namely the points v_2 and \tilde{v}_2 shown in blue.

Moreover, since these points both lie in $\text{null}(A)$, it follows by definition that

$Av_2 = A\tilde{v}_2 = 0$. So under this assumption, the origin must constitute the relevant “minor axis”. In addition, $\|Av_2\| = \|A\tilde{v}_2\| = 0$ implies that both points are equally good

solutions. So if we now focus on v_2 , then the solution value must be given by

$s_2 = \|Av_2\| = 0$. Finally, to complete this solution, observe that if we choose any unit

vector, u_2 , orthogonal to u_1 , (such as the point, u_2 , just to the right of v_2 in the figure), then it is automatically true that $Av_2 = 0 = \lambda_2 u_2$. So this degenerate “maximal stretch” solution is summarized by the triple (v_2, s_2, u_2) where $s_2 = 0$. Notice that when taken together, these two solutions can be written as

$$(A3.2.23) \quad \left. \begin{array}{l} Av_1 = s_1 u_1 \\ Av_2 = s_2 u_2 \end{array} \right\} \Rightarrow A(v_1, v_2) = (u_1, u_2) \begin{pmatrix} s_1 & \\ & s_2 \end{pmatrix} \Rightarrow AV = US$$

where $V = (v_1, v_2)$ and $U = (u_1, u_2)$ are *orthonormal* matrices by construction. So this in turn implies that

$$(A3.2.24) \quad A = USV'$$

and thus that (A3.2.1) holds for this choice of matrices. In the present case, it can readily be verified that these matrices have the exact form:

$$(A3.2.25) \quad USV' = \begin{pmatrix} 1/\sqrt{2} & 1/\sqrt{2} \\ 1/\sqrt{2} & -1/\sqrt{2} \end{pmatrix} \begin{pmatrix} \sqrt{10} & \\ & 0 \end{pmatrix} \begin{pmatrix} 2/\sqrt{5} & 1/\sqrt{5} \\ 1/\sqrt{5} & -2/\sqrt{5} \end{pmatrix} = \begin{pmatrix} 2 & 1 \\ 2 & 1 \end{pmatrix} = A$$

where for example, $s_1 = \sqrt{10} \approx 3.16$, is the length of the major axis vector,

$Av_1 = (\sqrt{10}, \sqrt{10})$ in Figure A3.21. The most important feature of this singular example is that all analysis of the collapsed “minor axis” in $null(A)$ is formally identical to that of the positive “major axis” in $span(A)$. The only difference is that the solutions in the collapsed case are nonunique, so that *any* choice of a unit vector, u_2 , orthogonal to u_1 will work.

Note finally that nonuniqueness of solutions is also possible for positive axes of the ellipsoid in $span(A)$. A simple example is provided by any *orthonormal* transformation, $A \equiv U$, where $U(\mathbb{S}_n) = \mathbb{S}_n$ implies that all “axes” of this spherical image must have the same length. In this extreme case, there are infinitely many SVD representations of U , including the trivial one, $U = U I_n I_n$. A more interesting example is based on the matrix, A , in (A3.2.11) with SVD given by⁶

$$(A3.2.26) \quad USV' = \begin{pmatrix} 0 & 1 & 0 \\ -0.9131 & 0 & -0.4076 \\ -0.4076 & 0 & 0.9131 \end{pmatrix} \begin{pmatrix} 1.95 & & \\ & 0.7 & \\ & & 0.6462 \end{pmatrix} \begin{pmatrix} 0 & -0.98924 & -0.14633 \\ 1 & 0 & 0 \\ 0 & -0.14633 & 0.98924 \end{pmatrix}$$

⁶ This solution was obtained numerically with the MATLAB program, **svd.m**.

$$= \begin{pmatrix} 0.7 & 0 & 0 \\ 0 & 1.8 & 0 \\ 0 & 0.7 & 0.7 \end{pmatrix} = A$$

While all principle axes in this example are distinct, notice that the lengths of the second and third axes (0.7 and 0.6462) are almost the same. Geometrically, this implies that the intersection of the surface of the ellipsoid in the right panel of Figure A3.17 with the plane orthogonal to the major axis vector, Av_1 , must be almost circular (as shown by the blue curve in the figure). So one can imagine that this intersection can be made exactly circular by an appropriately small modification of the matrix, A .⁷ In this circular case, it should be clear that while the principle axis vector, Av_1 , is still unique (up to a choice of sign) there is no unique choice of the second principle axis vector, Av_2 , shown in Figure A3.17. Any selection of a unit vector, v_2 , orthogonal to v_1 will do. But, as we shall see below, the actual maximization problem for identifying this principle axis is still well defined, and most importantly, all such choices of v_2 must satisfy the corresponding Lagrangian first-order conditions.

With these preliminary observations, we are now ready to extend the maximization problem in (A3.2.16) in order to obtain a full singular value decomposition (SVD) of matrices, A (which will further clarify the natural duality between A and A').⁸

Singular Value Decomposition Theorem. *For any n -square matrix, A , there exist orthonormal matrices, $U = (u_i : i = 1, \dots, n)$, $V = (v_i : i = 1, \dots, n)$ and a nonnegative diagonal matrix, $S = \text{diag}(s_i : i = 1, \dots, n)$, such that*

$$(A3.2.27) \quad A = USV'$$

Proof: To establish this result, we begin by observing that if (A3.2.27) holds, then [as an extension of (A3.2.23) above] it follows by definition that,

$$(A3.2.28) \quad A = USV' \Rightarrow AV = US$$

⁷ One such modification, A° , is obtained by simply replacing A with $A^\circ = \text{diag}(1.95, 0.7, 0.7)$ and using U and V in (A3.2.27) to define $A^\circ = US^\circ V'$.

⁸ As mentioned earlier, more compact versions of this SVD Theorem can be obtained by appealing to the Spectral Decomposition Theorem and employing the symmetric-matrix device in (A3.2.3) above. [For a “variational” version of this proof see Theorem 7.3.10 in Horn and Johnson (1985).] However, it should be emphasized that essentially all direct proofs of the Spectral Decomposition Theorem implicitly embed \mathbb{R} in the complex plane, \mathbb{C} , to ensure existence of such decompositions. Hence one of the objectives of the present approach is to avoid any appeal to complex number theory whatsoever.

$$\begin{aligned} \Rightarrow (Av_1, \dots, Av_n) &= (u_1, \dots, u_n) \begin{pmatrix} s_1 & & \\ & \ddots & \\ & & s_n \end{pmatrix} = (s_1 u_1, \dots, s_n u_n) \\ \Rightarrow Av_i &= s_i u_i, \quad i = 1, \dots, n \end{aligned}$$

where the last line is seen to have exactly the same form as (A3.2.18) above. Hence if we now denote the solution to (A3.2.17) by (u_1, s_1, v_1) , so that conditions (A3.2.18) through (A3.2.20) imply

$$(A3.2.29) \quad Av_1 = s_1 u_1, \quad A'u_1 = s_1 v_1, \quad u_1' u_1 = 1 = v_1' v_1$$

then our objective is to extend this relation to a full SVD as in (A3.2.28) by generating the successive triplets, (u_i, s_i, v_i) , one at a time. Here it is instructive to generate the first triplet, (u_2, s_2, v_2) , in full detail, and then proceed by induction for the rest. To do so, we begin by observing that in order for U and V to be orthonormal, we must require that (u_2, v_2) satisfy the orthogonality conditions, $u_2' u_1 = 0 = v_2' v_1$. So if we now let $\mathbb{S}(u_1) = \{u \in \mathbb{S}_n : u' u_1 = 0\}$ and $\mathbb{S}(v_1) = \{v \in \mathbb{S}_n : v' v_1 = 0\}$ denote the vectors of unit length orthogonal to u_1 and v_1 respectively, then in geometric terms, the task is to find “maximal stretch” vectors, $(u_2, v_2) \in \mathbb{S}(u_1) \times \mathbb{S}(v_1)$, for transformation A which generate the “second principle axes” of the ellipsoids, $A'(\mathbb{S}_n)$ and $A(\mathbb{S}_n)$, respectively. [For example, the set $\mathbb{S}(v_1)$ for the nonsingular illustration in Figure A3.17 above is shown by the blue circle on \mathbb{S}_3 in the left panel, with corresponding image, $A[\mathbb{S}(v_1)]$, shown by the blue circle in the right panel. Similarly, for the singular illustration in Figure A3.20, the set $\mathbb{S}(v_1)$ is again shown on the left as a (different) blue circle on \mathbb{S}_3 , with associated image now corresponding to interval shown in dark blue in the right panel. [Note that (for sake of visual clarity) neither the vector, u_1 , or its orthogonal set, $\mathbb{S}(u_1)$, are shown in these figures.] As a natural extension of (A3.2.17), the appropriate maximization problem for determining (u_2, v_2) is given by

$$(A3.2.30) \quad \text{maximize: } u_2' A v_2 \quad \text{subject to: } (u_2, v_2) \in \mathbb{S}(u_1) \times \mathbb{S}(v_1)$$

Moreover, $\mathbb{S}(u_1)$ and $\mathbb{S}(v_1)$ are again a nonempty closed bounded subsets of \mathbb{R}^n for $n \geq 2$ [implying that $\mathbb{S}(u_1) \times \mathbb{S}(v_1)$ must be a nonempty closed bounded subset of \mathbb{R}^{2n}]. So the same argument using the Generalized Extreme Value Theorem again shows that a solution to (A3.2.30) must exist. Since the above constraint conditions for (u_2, v_2) can be equivalently stated as

$$(A3.2.31) \quad u_2' u_2 = 1, \quad v_2' v_2 = 1, \quad u_2' u_1 = 0, \quad v_2' v_1 = 0$$

it follows that the appropriate Lagrangian function for this problem takes the form:

$$(A3.2.32) \quad L(u_2, v_2, s_2, \theta_2, \alpha_2, \beta_2) = u_2' A v_2 + \frac{1}{2} [s_2 (1 - u_2' u_2) + \theta_2 (1 - v_2' v_2)] \\ + \alpha_2 (u_2' u_1) + \beta_2 (v_2' v_1)$$

Here the first order conditions for u_2 and v_2 are given respectively by

$$(A3.2.33) \quad 0 = \nabla_{u_2} L = A v_2 - s_2 u_2 + \alpha_2 u_1 \Rightarrow A v_2 = s_2 u_2 - \alpha_2 u_1$$

$$(A3.2.34) \quad 0 = \nabla_{v_2} L = A' u_2 - \theta_2 v_2 + \beta_2 v_1 \Rightarrow A' u_2 = \theta_2 v_2 - \beta_2 v_1$$

with corresponding first order conditions for $(s_2, \theta_2, \alpha_2, \beta_2)$ given precisely by the conditions in (A3.2.31) above. At this point it is important to recall from the discussion in Section 8 of the Appendix to Part II that the validity of this Lagrangian formulation requires that the constraint gradient vectors be *linearly independent* [recall expression (A3.1.24) above]. But this is automatically guaranteed by the mutual orthogonality of (u_2, u_1) and (v_2, v_1) . Hence the task remaining in this second step is to show that $\alpha_2 = 0 = \beta_2$, so that (A3.2.33) and (A3.2.34) will have the same form as (A3.2.18) and (A3.2.19). But since the solution in (A3.2.32) is assumed to satisfy (A3.2.18) and (A3.2.19), together with (A3.2.31), it follows by premultiplying (A3.2.33) by u_1' that

$$(A3.2.35) \quad u_1' A v_2 = s_2 u_1' u_2 - \alpha_2 u_1' u_1 = 0 - \alpha_2 = -\alpha_2 \\ \Rightarrow \alpha_2 = -u_1' A v_2 = -(A' u_1)' v_2 = -(s_1 v_1)' v_2 = -s_1 (v_1' v_2) = 0$$

Similarly, premultiplying (A3.2.34) by v_1' , we see that

$$(A3.2.36) \quad v_1' A' u_2 = s_2 v_1' v_2 - \beta_2 v_1' v_1 = 0 - \beta_2 = -\beta_2 \\ \Rightarrow \beta_2 = -(v_1' A') u_2 = -(A v_1)' u_2 = -(s_1 u_1)' u_2 = -s_1 (u_1' u_2) = 0$$

Hence $\alpha_2 = 0 = \beta_2$, and conditions (A3.2.33) and (A3.2.34) reduce to

$$(A3.2.37) \quad A v_2 = s_2 u_2$$

$$(A3.2.38) \quad A' u_2 = \theta_2 v_2$$

Moreover, exactly the same argument in (A3.2.21) with (u_2, s_2, v_2) replacing (u, s, v) now shows that $\theta_2 = s_2$, so that (A3.2.38) becomes

$$(A3.2.39) \quad A'u_2 = s_2v_2$$

Hence the maximal stretch, s_2 , for transformation A among vectors in $\mathbb{S}(v_1)$ is achieved at v_2 , and similarly, the same maximal stretch for transformation A' among vectors in $\mathbb{S}(u_1)$ is achieved at u_2 . Most importantly for our present purposes, expression (A3.2.37) shows that (u_2, s_2, v_2) yields the desired second row for the SVD in expression (A3.2.28). Note finally that this solution (u_2, s_2, v_2) may not be unique, even when (u_1, s_1, v_1) is unique [such as in the modification of example (A3.2.11) illustrated above]. But all such solutions must necessarily satisfy conditions (A3.2.31), (A3.2.37) and (A3.2.39).

The task remaining is to extend this argument by induction to all rows of (A3.2.28). To do so, we start with the *inductive hypothesis* that for a given $k \leq n$, the first $k-1$ rows have been filled with triplets, (u_i, s_i, v_i) , $i = 1, \dots, k-1$, satisfying

$$(A3.2.40) \quad Av_i = s_iu_i, \quad i = 1, \dots, k-1$$

$$(A3.2.41) \quad A'u_i = s_iv_i, \quad i = 1, \dots, k-1$$

$$(A3.2.42) \quad u'_iu_i = 1 = v'_iv_i, \quad i = 1, \dots, k-1$$

$$(A3.2.43) \quad u'_iu_j = 0 = v'_iv_j, \quad i, j = 1, \dots, k-1, i \neq j$$

If we now let $\mathbb{S}(u_1, \dots, u_{k-1}) = \{u \in \mathbb{S}_n : u'_iu_i = 0, i = 1, \dots, k-1\}$ denote the set of vectors in \mathbb{S}_n orthogonal to (u_1, \dots, u_{k-1}) , and similarly let $\mathbb{S}(v_1, \dots, v_{k-1}) = \{v \in \mathbb{S}_n : v'_iv_i = 0, i = 1, \dots, k-1\}$ denote the vectors in \mathbb{S}_n orthogonal to (v_1, \dots, v_{k-1}) , then since these nonempty sets are again closed and bounded, one final application of the Generalized Extreme Value Theorem shows that the maximization problem

$$(A3.2.44) \quad \text{maximize: } u'_kAv_k \quad \text{subject to: } (u_k, v_k) \in \mathbb{S}(u_1, \dots, u_{k-1}) \times \mathbb{S}(v_1, \dots, v_{k-1})$$

must have a solution. Moreover, as an extension of (A3.2.31) and (A3.2.32), it follows that if the constraint conditions on (u_k, v_k) are written explicitly as

$$(A3.2.45) \quad u'_ku_k = 1 = v'_kv_k, \quad u'_iu_i = 0 = v'_iv_i, \quad i = 1, \dots, k-1$$

then the appropriate Lagrangian function for (A3.2.44) is seen to have the form:

$$(A3.2.46) \quad L(u_k, v_k, s_k, \theta_k, \alpha_1, \dots, \alpha_{k-1}, \beta_1, \dots, \beta_{k-1}) = u'_kAv_k + \frac{1}{2} [s_k(1 - u'_ku_k) + \theta_k(1 - v'_kv_k)] \\ + \sum_{i=1}^{k-1} \alpha_i(u'_iu_i) + \sum_{i=1}^{k-1} \beta_i(v'_iv_i)$$

Here the first order conditions for u_k and v_k have the respective forms

$$(A3.2.47) \quad 0 = \nabla_{u_k} L = Av_k - s_k u_k + \sum_{i=1}^{k-1} \alpha_i u_i \Rightarrow Av_k = s_k u_k - \sum_{i=1}^{k-1} \alpha_i u_i$$

$$(A3.2.48) \quad 0 = \nabla_{v_k} L = A'u_k - \theta_k v_k + \sum_{i=1}^{k-1} \beta_i v_i \Rightarrow A'u_k = \theta_k v_k - \sum_{i=1}^{k-1} \beta_i v_i$$

and the remaining first order conditions are now given by (A3.2.45). [Note again from the orthogonality conditions in (A3.2.45) that the constraint gradient vectors in both (A3.2.47) and (A3.2.48) are linearly independent, so that this Lagrangian formulation of (A3.2.46) is indeed valid.] Next, to show that $\alpha_j = 0 = \beta_j$, $j = 1, \dots, k-1$, we again premultiply (A3.2.47) by u'_j and use the inductive hypotheses (A3.2.40) through (A3.2.43) together with (A3.2.45) to conclude that

$$(A3.2.49) \quad u'_j Av_k = s_k u'_j u_k - \sum_{i=1}^{k-1} \alpha_i u'_j u_i = 0 - [\alpha_j (u'_j u_j) + 0] = -\alpha_j \\ \Rightarrow \alpha_j = -u'_j Av_k = -(A'u_j)' v_k = -(s_j v_j)' v_k = -s_j (v'_j v_k) = 0$$

Similarly, by premultiplying (A3.2.48) by v'_j , we see that

$$(A3.2.50) \quad v'_j A'u_k = \theta_k v'_j v_k - \sum_{i=1}^{k-1} \beta_i v'_j v_i = 0 - [\beta_j (v'_j v_j) + 0] = -\beta_j \\ \Rightarrow \beta_j = -v'_j A'u_k = -(Av_j)' u_k = -(\theta_j u_j)' u_k = -\theta_j (u'_j u_k) = 0$$

Hence (A3.2.47) and (A3.2.48) reduce to

$$(A3.2.51) \quad Av_k = s_k u_k$$

$$(A3.2.52) \quad A'u_k = \theta_k v_k$$

Finally, since the argument in (A3.2.21) with (u_k, s_k, v_k) replacing (u, s, v) again shows that $\theta_k = s_k$, we see that (A3.2.52) becomes

$$(A3.2.53) \quad A'u_k = s_k v_k$$

Thus the conditions in (A3.2.40) through (A3.2.43) hypothesized for $i = 1, \dots, k-1$ are seen to hold for k as well, and it follows by induction that they must hold for *all* $i = 1, \dots, n$. Most importantly for our purposes, conditions (A3.2.40) together with (A3.2.42) and (A3.2.43) are now seen to yield a full SVD for A as in expression (A3.2.28). \square

This particular proof of the SVD Theorem has a number of additional geometric consequences. Note first from (A3.2.45) and (A3.2.51) that

$$(A3.2.54) \quad u'_k A v_k = s_k (u'_k u_k) = s_k, \quad k = 1, \dots, n$$

so that the stretch values, s_k , are indeed the maximum values of the objective function, $u'_k A v_k$, at each step, k . Moreover, since this objective function is formally the same at each step, and since the constraints sets form a nested decreasing sequence of sets, i.e.,

$$(A3.2.55) \quad \mathbb{S}(u_1, \dots, u_k) \times \mathbb{S}(v_1, \dots, v_k) \subset \mathbb{S}(u_1, \dots, u_{k-1}) \times \mathbb{S}(v_1, \dots, v_{k-1}), \quad k = 1, \dots, n$$

it follows that these maximal values must necessarily form a *non-increasing* sequence, so that

$$(A3.2.56) \quad s_1 \geq s_2 \geq \dots \geq s_n$$

In geometric terms, these singular values thus yield the *successive lengths of the principle axes* corresponding to the ellipsoidal image, $A(\mathbb{S}_n)$, of the unit sphere, \mathbb{S}_n , under the linear transformation, A . In particular, if $s_1 > s_2 > \dots > s_n > 0$, then A is *nonsingular* and the n -dimensional ellipsoid, $A(\mathbb{S}_n)$, has a well defined set of principle axes.

However, if there are say k repetitions of a positive singular value, such as in the modified version of Figure A3.17 illustrated above with $k = 2$, then a k -dimensional “slice” through this ellipsoid will be spherical. Similarly, if the last k singular values are zero, then A is *singular* and its null space, $null(A)$, has exactly dimension k . So a great deal of information about A is conveyed by these singular values.

However, it should also be emphasized that the programming formulation of this proof is not meant to provide a method for *computing* the SVD of a matrix. This is particularly evident when there are repeated positive singular values (either positive or zero). Here there are infinitely many programming solutions, and procedures such as Gram-Schmidt orthogonalization must be used to construct appropriate orthonormal sets of solution vectors. While there exist very efficient methods for constructing such decompositions (often based on Householder representations in section A3.1.2 above), such procedures are beyond the scope of these notes.⁹

We now consider some of the more useful consequences of the SVD Theorem for our purposes. As already mentioned, one direct consequence is to clarify the geometric relation between A and A' . In particular, it follows at once from (A3.2.27) together with (A3.1.13) that

$$(A3.2.57) \quad A = USV' \Leftrightarrow A' = VSU'$$

⁹ For a discussion of such methods as used by MATLAB, see Chapter 10 of Mohler (2004).

So the singular values of A and A' must always be the same. More the above proof shows their respective ellipsoidal images, $A(\mathbb{S}_n)$ and $A'(\mathbb{S}_n)$, of the unit sphere, \mathbb{S}_n , must essentially be rotations of one another, where the roles of the orthonormal matrices, U and V , are exactly reversed. A simple illustration of this relationship is given in Figure A3.22 below, where the unit circle, \mathbb{S}_2 , is shown in black, and the elliptical images, $A(\mathbb{S}_2)$ and $A'(\mathbb{S}_2)$, for a given matrix, A ,¹⁰ are shown in blue and red, respectively.

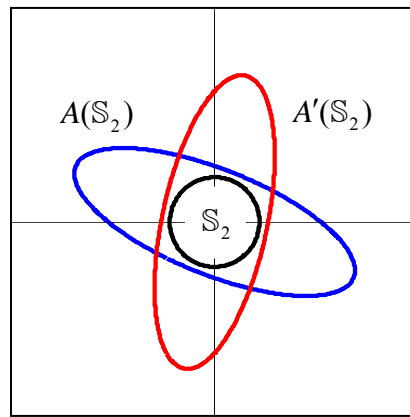


Figure A3.22. Ellipsoidal Relations for Transposes

Next we consider a number of SVD consequences that will be used in our subsequent analyses.

A3.2.1. Inverses and Pseudoinverses.

Note first that since the inverse of an orthonormal matrix is simply its transpose, it follows at once from the SVD Theorem that for any nonsingular matrix, A ,

$$(A3.2.58) \quad A = U S V' \Rightarrow A^{-1} = V S^{-1} U' = (v_1, \dots, v_n) \begin{pmatrix} \frac{1}{s_1} u'_1 \\ \vdots \\ \frac{1}{s_n} u'_n \end{pmatrix}$$

Thus, by recalling (A3.1.11), we see that the inverse, A^{-1} , can be determined from the SVD of A almost by inspection. While this of course assumes that this SVD has already been calculated, it nonetheless provides a powerful analytical tool in many contexts. For example, it now reveals the behavior of “almost nonsingular” matrices, which by

¹⁰ The particular matrix used here was $A = [1.0689, 2.9443 ; 0.8095, -1.4384]$.

definition have at least one singular value, s_i , very close to zero. But since this in turn implies that $1/s_i$ must be very large, it can be seen from the last equality in (A3.2.58) that vectors in the u_i direction are being stretched enormously. So this shows not only that A^{-1} is becoming unstable, but also the *directions* in which this instability is worst.

Even more important is the fact that this SVD shows how to construct *generalized inverses* for singular matrices. In particular, when no inverse exists for A , this SVD representation suggest a very natural “best approximation” to such an inverse. The idea is seen most clearly in trying to solve the associated linear equation system, $Ax = b$. If A^{-1} exists, then there is an exact solution, $x = A^{-1}b$. But if A is singular, one would like to find x so that Ax is as “close” to b as possible, i.e. so that $\|Ax - b\|$ is *minimized*. But by (A3.2.54),¹¹

$$(A3.2.59) \quad Ax - b = (USV')x - b = USV'x - UU'b = U(SV'x - U'b) = U(S\tilde{x} - \tilde{b})$$

where $\tilde{x} = V'x$ and $\tilde{b} = U'b$. But since U is orthonormal and hence preserves distances, it follows that

$$(A3.2.60) \quad \|Ax - b\| = \|S\tilde{x} - \tilde{b}\| = \left\| \begin{pmatrix} s_1 & & \\ & \ddots & \\ & & s_n \end{pmatrix} \begin{pmatrix} \tilde{x}_1 \\ \vdots \\ \tilde{x}_n \end{pmatrix} - \begin{pmatrix} \tilde{b}_1 \\ \vdots \\ \tilde{b}_n \end{pmatrix} \right\| = \left\| \begin{pmatrix} \tilde{x}_1 s_1 - \tilde{b}_1 \\ \vdots \\ \tilde{x}_n s_n - \tilde{b}_n \end{pmatrix} \right\|$$

So this approximation problem has now been reduced to a diagonal form for which the solution is seen to be trivial, namely, set

$$(A3.2.61) \quad \tilde{x}_i = \begin{cases} \tilde{b}_i / s_i & , s_i > 0 \\ 0 & , s_i = 0 \end{cases}$$

Finally, if we assume (for convenience) that the first k components of S are the positive ones, and set

$$(A3.2.62) \quad S^- = \text{diag}(s_1^{-1}, \dots, s_k^{-1}, 0, \dots, 0)$$

then it follows from (A3.2.61) together with the definitions of \tilde{x} and \tilde{b} that

$$(A3.2.63) \quad \tilde{x} = S^- \tilde{b} \Rightarrow V'x = S^- U'b \Rightarrow x = (V S^- U') b$$

Finally, since this argument is completely *independent* of the choice of b , it follows by setting

¹¹ The following argument is base on the excellent discussion of SVD properties in Kalman (1996).

$$(A3.2.64) \quad A^- = V S^- U'$$

that A^- yields a natural generalization of (A3.2.58) which is designated as the *pseudoinverse* (or *Moore-Penrose inverse*) of A . Moreover, since minimizing distance is the same as minimizing *squared* distance, this pseudoinverse, A^- , always provides the *least squares solution* to linear equation systems.

These observations are immediately applicable to OLS regression. In particular, recall from expression (A2.7.69) that the least squares solution, $\hat{\beta}$, for estimating β satisfies the linear equation system:

$$(A3.2.65) \quad (X'X)\hat{\beta} = X'y$$

So if $(X'X)^{-1}$ exists (as was assumed) then $\hat{\beta} = (X'X)^{-1}X'y$. But in cases where $X'X$ is singular, one can still determine a least squares solution by setting

$$(A3.2.66) \quad \hat{\beta} = (X'X)^- X'y$$

Moreover, even in cases where $X'X$ is technically nonsingular but is in fact “almost singular” (i.e., exhibits strong multicollinearities), one can often obtain a more stable estimate by using (A3.2.66). So the SVD Theorem is seen to have very practical applications in such cases.

A3.2.2. Determinants and Volumes

Recall from expression (3.2.11) in Part II of this NOTEBOOK that we encountered determinants in the density function of the multi-normal distribution. The main objective of this section is to clarify the role of determinants in such densities, and to emphasize their broader role in describing the volume changes associated with linear transformations. To do so, we require some preliminary facts about matrix determinants. For the simple case of a 2×2 matrix, A , recall that the *determinant* of A is given by

$$(A3.2.67) \quad A = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix} \Rightarrow |A| = a_{11}a_{22} - a_{12}a_{21},$$

which in turn plays a critical role in calculating the *inverse* of A :

$$(A3.2.68) \quad A^{-1} = \frac{1}{|A|} \begin{pmatrix} a_{22} & -a_{12} \\ -a_{21} & a_{11} \end{pmatrix}$$

[In fact, the determinant itself originated as part of the first general solution of linear equations (Cramer's Rule, 1750).] Note in particular from (A3.2.68) that such solutions exist iff $|A| \neq 0$. The geometric meaning of this relationship will become clear below.

But for the present, we simply note that the formula in (A3.2.67) offers little insight by itself. Indeed, the general formula for determinants (in terms of alternating-signed sums of products of matrix elements)¹² is even more obtuse. But one important observation about this formula can be made in terms of the following instance of a Householder reflection, $A = H_v$, in \mathbb{R}^2 [recall expression (A3.1.44) above], where in this case $v = (1, -1)'$ [with $v'v = 2$], so that:

$$(A3.2.69) \quad A = I_2 - \frac{2}{v'v} vv' = I_2 - \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

By expression (A3.2.67) this matrix has a *negative* determinant, $|A| = -1$. To interpret the meaning of this negative sign, note from Figure A3.23 below that this transformation simply *reverses* the basis vectors, (e_1, e_2) , so that $A(e_1, e_2) = (e_2, e_1)$:




Figure A3.23. Order Reversal

More generally, negative values of determinants are always associated with such reversals of orientation. But this “sign” property of determinants is not of direct interest for our purposes (even though the present Householder example will prove useful later). Rather, we are primarily interested in the *absolute value* of determinants. As mentioned above, these absolute values tell us exactly how volumes are transformed under linear transformations. The standard example which is often shown in the literature is illustrated in Figure A3.24 below:




Figure A3.24. Volume Transformation

¹² See for example section 0.3 in Horn and Johnson (1985)

In terms of Figure A3.3, we here set $e_1 = (1, 0)'$, $e_2 = (0, 1)'$, $Ae_1 = (a_{11}, a_{21})'$ and $Ae_2 = (a_{12}, a_{22})'$ in order to emphasize the role of each matrix element. The key point is that the unit area of the unit square on the left is transformed by A into a parallelogram on the right with area given precisely by $|A|$, which in the case is seen to be positive (no reversal of orientation). This in turn implies (from linearity) that every area on the left is transformed by A into an area scaled by a factor of $|A|$. But even in this simple case, it is not obvious that the parallelogram area should be given by $a_{11}a_{22} - a_{12}a_{21}$. While the geometric proof in this case is not difficult, its generalization to linear transformations, A , in \mathbb{R}^n is tedious, to say the least. So our first objective is to show that this relation between volume and absolute determinant values can be made completely transparent in terms of the SVD of A .

To do so, we must first deal with the (unfortunate) notational fact that the symbol, $|\cdot|$, is used both for determinants and absolute values. This is often resolved by using “ $\det(A)$ ” for the determinant of A , so that its absolute value can be directly represented by $|\det(A)|$. But since the relevant determinants for our purposes will almost always be nonnegative, we choose to stay with the simpler notation, $|A|$. Where it is essential to specify absolute values of determinants (such as in the present section) we shall simply write, $|A|_+$.

Aside from this notational convention, the only algebraic properties of determinants that we require are the *product rule*,

$$(A3.2.70) \quad |AB| = |A||B|$$

the *symmetry rule*,

$$(A3.2.71) \quad |A'| = |A|$$

and the *diagonal rule*,

$$(A3.2.72) \quad |\text{diag}(a_1, a_2, \dots, a_n)| = a_1 a_2 \cdots a_n$$

Note in particular that for absolute values, the product rule implies

$$(A3.2.73) \quad |AB|_+ = |A|_+ |B|_+$$

Together with the SVD Theorem, these properties of determinants imply that *the absolute determinant of any matrix is the product of its singular values*, i.e., that for all transformations, A , in (A3.2.28)

$$(A3.2.74) \quad A = U \text{diag}(s_1, \dots, s_n) V' \Rightarrow |A|_+ = \prod_{i=1}^n s_i$$

To see this, note first from (A3.2.72) that $|I_n| = 1$, so that by the defining property of orthonormal transformations, U ,

$$(A3.2.75) \quad 1 = |I_n| = |U'U| = |U'| |U| = |U|^2 \Rightarrow |U| = \pm 1$$

Hence the absolute determinant of U must be unity, i.e.,

$$(A3.2.76) \quad U \text{ orthonormal} \Rightarrow |U|_+ = 1$$

Hence it follows from (A3.2.70) and (A3.2.72) that

$$(A3.2.77) \quad A = U \text{diag}(s_1, \dots, s_n) V' \Rightarrow |A|_+ = |U|_+ |\text{diag}(s_1, \dots, s_n)|_+ |V'|_+ \\ = (1) |\text{diag}(s_1, \dots, s_n)| (1) \\ = \prod_{i=1}^n s_i$$

Using this result, it is a simple matter to show that for any linear transformation, A , on \mathbb{R}^n , volumes are transformed by a factor of $|A|_+$. To do so, observe that if the *unit cube* in \mathbb{R}^n is denoted by

$$(A3.2.78) \quad C_n = [0, 1]^n = \{x = (x_1, \dots, x_n) \in \mathbb{R}^n : 0 \leq x_i \leq 1, i = 1, \dots, n\}$$

and if we denote the volume of any set, $T \subset \mathbb{R}^n$ by $\text{vol}(T)$,¹³ then clearly $\text{vol}(C_n) = 1$. So if the *image* of C_n under transformation A is denoted by

$$(A3.2.79) \quad A(C_n) = \{Ax : x \in C_n\}$$

then it suffices to show that $\text{vol}[A(C_n)]$ is always given by $|A|_+$. But since each *linear* transformation scales all volumes by the same amount, if we now denote this common *scale factor* by $s(A) = \text{vol}[A(C_n)]$,¹⁴ then for all $T \subset \mathbb{R}^n$,

$$(A3.2.80) \quad \frac{\text{vol}[A(T)]}{\text{vol}(T)} = \frac{\text{vol}[A(C_n)]}{\text{vol}(C_n)} = s(A) \Rightarrow \text{vol}[A(T)] = s(A) \text{vol}(T)$$

In these terms, our objective is to show that for any linear transformation, A , on \mathbb{R}^n ,

¹³ The knowledgeable reader will note that technically we here refer to any *measurable* set, $T \subset \mathbb{R}^n$.

¹⁴ Be careful not to confuse scale factors, $s(A)$, with singular values, s .

$$(A3.2.81) \quad s(A) = |A|_+$$

Here we need only appeal to certain elementary properties of volume itself. The most fundamental property concerns scale transformations of individual coordinates. For example, if a transformation scales all coordinate axes by 2, then volumes increase by a factor of 2^n . More generally, since positive diagonal matrices, $D = \text{diag}(d_1, \dots, d_n)$, scale each coordinate, x_i , by a factor of d_i , i.e., since

$$(A3.2.82) \quad Dx = D(x_1, \dots, x_n) = (d_1x_1, \dots, d_nx_n),$$

it follows that $\text{vol}[D(C_n)] = d_1d_2 \cdots d_n$, so that by definition

$$(A3.2.83) \quad D = \text{diag}(d_1, d_2, \dots, d_n) \Rightarrow s(D) = \prod_{i=1}^n d_i$$

In fact, this is how volumes of n -dimensional “boxes” are computed. Note also that if coordinates are scaled by factors, d_i , *one at a time*, then since the composition of these transformations is precisely D , the cumulative effect of these scale changes is necessarily *multiplicative*. More generally, the cumulative scale effect of any successive transformations, say A followed by B , is always multiplicative. For example, if A doubles volumes and B triples volumes, then the composite transformation, BA , increases volumes by a factor of $(2)(3) = 6$. More generally, for all transformations, A and B ,

$$(A3.2.84) \quad s(BA) = s(B)s(A)$$

The only other property of volume that we require is one we have already seen, namely that *orthonormal transformations preserve volumes*. So by definition,

$$(A3.2.85) \quad U \text{ orthonormal} \Rightarrow s(U) = 1$$

Given these volume properties, it follows at once from the SVD Theorem together with (A3.2.84) that

$$\begin{aligned} (A3.2.86) \quad A = U \text{diag}(s_1, \dots, s_n) V' &\Rightarrow s(A) = s(U) s[\text{diag}(s_1, \dots, s_n)] s(V') \\ &= (1) \left(\prod_{i=1}^n s_i \right) (1) \\ &= |A|_+ \end{aligned}$$

This result has far reaching consequences for determinants, and shows why they play such a fundamental role in linear algebra. With respect to matrix inverses in particular, note that if $|A| = 0$ (so that $|A|_+ = 0$) then $s(A) = 0$ implies that all volumes are

collapsed to zero. So from a geometric viewpoint, A must collapse the space into a lower dimensional subspace, such as the examples in Figures A3.20 and A3.21 above.

A3.2.3 Linear Transformations of Random Vectors

The final objective of this section is to illustrate the consequences of these results for linear transformations of random vectors. In particular, our objective is to complete the derivation of the multi-normal distribution sketched in Section 3.2.1 of Part II in this NOTEBOOK, and to show how the multi-normal density in (3.2.11) is derived. The key element we focus on is the role of the determinant, $|\Sigma|$, of the covariance matrix, Σ . In fact, this determinant reflects the volume transformation associated with a particular linear transformation, as we now show. To do so, we start by considering the *standard normal random vector*, $X = (X_1, \dots, X_n)'$, of independent standard normal variates,

$X_i \sim N(0,1)$, $i = 1, \dots, n$. Recall from the *Linear Invariance Theorem* of Section 3.2.2 of Part II that if for some nonsingular matrix, A , the random vector, Y , is defined by

$$(A3.2.87) \quad Y = AX + \mu \quad ,$$

then since $X \sim N(0, I_n)$, this theorem asserts that $Y \sim N(\mu, \Sigma)$ with $\Sigma = AA'$. Moreover, since *all* covariance matrices, Σ , are of this form for some A [as we have already seen from the Cholesky Theorem in Appendix A2, below expression (A2.7.45)], it follows that *all multi-normal random vectors, Y , are derivable as linear transformations of the standard normal vector, X* . In fact, this is precisely how the general multi-normal distribution is *defined*.

Our goal is to establish this result by starting with the probability density of the standard normal random vector, X , and show how this density is transformed under (A3.2.87). To do so, we first recall from the argument in (3.2.7) of Part II [with $(\mu_i, \sigma_i) = (0, 1)$, $i = 1, \dots, n$] that the probability density, $f(x) = f(x_1, \dots, x_n)$, of the standard normal random vector, X , is necessarily of the form:

$$(A3.2.88) \quad f(x) = f(x_1) \cdots f(x_n) = \left(\frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}x_1^2} \right) \cdots \left(\frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}x_n^2} \right) \\ = \left(\frac{1}{\sqrt{2\pi}} \right)^n e^{-\frac{1}{2}(x_1^2 + \cdots + x_n^2)} = (2\pi)^{-n/2} e^{-\frac{1}{2}x'x}$$

where $x' = (x_1, \dots, x_n)$. So to obtain the desired distribution of Y , it suffices to show that this standard normal density is transformed by (A3.2.87) into a probability density, $g(y) = g(y_1, \dots, y_n)$, of the form (3.2.11) in Part II, i.e., that:

$$(A3.2.89) \quad g(y) = (2\pi)^{-n/2} |\Sigma|^{-1/2} e^{-\frac{1}{2}(y-\mu)'\Sigma^{-1}(y-\mu)}$$

But before doing so, it is important to emphasize that even though expressions like (A3.2.87) are usually referred to as “linear transformations”, they technically involve *linear transformations, A, plus translation terms, μ* (and are properly classified as *affine transformations*). Only in the case, $\mu = 0$, is this a linear transformation [as defined in (A3.1.1) above]. So to simplify the present development further, we start with the case, $\mu = 0$, where (A3.2.87) reduces to a proper *linear transformation*,

$$(A3.2.90) \quad Y = AX$$

It will be seen later that adding a nonzero translation term, μ , is then a simple matter.

To begin this development, we start by observing the role of the determinant, $|\Sigma|$, in (A3.2.89) in fact has nothing to do with “normality” itself. So clarify this role, it is more convenient to regard X in (A3.2.90) as a general continuous random vector with density, $f(x)$. To derive the associated density, $g(y)$, of Y , we begin by recalling that all probability densities are by definition representations of event probabilities in terms of *volumes*. In particular, if for any selected value, $y_0 = (y_{01}, \dots, y_{0n})$, of Y we consider small intervals, $\Delta_{0i}(\varepsilon) = [y_{0i} - \varepsilon, y_{0i} + \varepsilon]$ about each component value, y_{0i} , for some $\varepsilon > 0$, and denote the n -cube defined by these intervals as,

$$(A3.2.91) \quad \Delta_0(\varepsilon) = \Delta_{01}(\varepsilon) \times \Delta_{02}(\varepsilon) \times \dots \times \Delta_{0n}(\varepsilon) \subset \mathbb{R}^n,$$

then the probability, $\Pr[Y \in \Delta_0(\varepsilon)]$, of event $\Delta_0(\varepsilon)$ is represented by the integral of density g over this region of \mathbb{R}^n , i.e.,

$$(A3.2.92) \quad \Pr[Y \in \Delta_0(\varepsilon)] = \int_{\Delta_0(\varepsilon)} g(y) dy = \int_{\Delta_{0n}(\varepsilon)} \dots \int_{\Delta_{01}(\varepsilon)} g(y_1, \dots, y_n) dy_1 \dots dy_n$$

This is illustrated for the case of $n = 2$ on the right-hand side of Figure A3.25 below, where the 2-cube, $\Delta_0(\varepsilon)$, is seen to be a square (shown in blue) about point, y_0 :




Figure A3.25. Linear Transformation of Variables

So the probability integral in (A3.2.92) is simply the volume under that portion of density, g , above this square (also shown in blue). The key point here is that if the value of ε is sufficiently small, then this volume is well approximated by the box with *base*, $\Delta_0(\varepsilon)$, and *height*, $g(y_0)$. More precisely, if the *area* (more generally, *volume*) of this base is denoted by $vol[\Delta_0(\varepsilon)]$, so that the volume of the box (height \times base) is given by, $g(y_0)vol[\Delta_0(\varepsilon)]$, then we obtain the approximation:

$$(A3.2.93) \quad \Pr[Y \in \Delta_0(\varepsilon)] = g(y_0)vol[\Delta_0(\varepsilon)] + e_Y(\varepsilon)$$

where the magnitude of error term, $e_Y(\varepsilon)$, is assumed to be much smaller than $vol[\Delta_0(\varepsilon)]$, so that as ε approaches zero,

$$(A3.2.94) \quad \lim_{\varepsilon \rightarrow 0} \frac{e_Y(\varepsilon)}{vol[\Delta_0(\varepsilon)]} = 0$$

To gain some feeling for such error representations, observe that if we divide both sides of (A3.2.93) by $vol[\Delta_0(\varepsilon)]$, and let $\varepsilon \rightarrow 0$ then we obtain

$$(A3.2.95) \quad \lim_{\varepsilon \rightarrow 0} \frac{\Pr[Y \in \Delta_0(\varepsilon)]}{vol[\Delta_0(\varepsilon)]} = g(y_0) ,$$

which is essentially definition of probability density, $g(y_0)$, at y_0 .

In order to associate these quantities with the random vector, X , observe from (A3.2.90) that since, $y = Ax \Leftrightarrow x = A^{-1}y$, it follows that Y -outcome, y , occurs iff X -outcome, $A^{-1}y$, occurs. So by using the same image notation in (A3.2.10) to write

$$(A3.2.96) \quad A^{-1}[\Delta_0(\varepsilon)] = \{A^{-1}y: y \in \Delta_0(\varepsilon)\}$$

we obtain the following probability identity

$$(A3.2.97) \quad \Pr[Y \in \Delta_0(\varepsilon)] = \Pr(A^{-1}Y \in A^{-1}[\Delta_0(\varepsilon)]) = \Pr(X \in A^{-1}[\Delta_0(\varepsilon)])$$

This provides the *key link* between the X and Y distributions. The X -event in the last equality is shown (for the $n = 2$ case) on the left-hand side of Figure A3.25 as a parallelogram (in red) representing the image of $\Delta_0(\varepsilon)$ under A^{-1} . (Note also that the bold red arrow shows the direction of this inverse relationship.) But since X is continuously distributed with density, f , this probability again has a “box” approximation with base, $A^{-1}[\Delta_0(\varepsilon)]$, and height, $f(A^{-1}y_0)$, i.e.,

$$(A3.2.98) \quad \Pr(X \in A^{-1}[\Delta_0(\varepsilon)]) = f(A^{-1}y) \text{vol}(A^{-1}[\Delta_0(\varepsilon)]) + e_x(\varepsilon)$$

where the error, $e_x(\varepsilon)$, again satisfies

$$(A3.2.99) \quad \lim_{\varepsilon \rightarrow 0} \frac{e_x(\varepsilon)}{\text{vol}(A^{-1}[\Delta_0(\varepsilon)])} = 0$$

But now we are in a position to simplify (A3.2.98) by using (A3.2.86), together with (A3.2.80) to obtain

$$(A3.2.100) \quad \text{vol}(A^{-1}[\Delta_0(\varepsilon)]) = s(A^{-1}) \text{vol}[\Delta_0(\varepsilon)] = |A^{-1}|_+ \text{vol}[\Delta_0(\varepsilon)]$$

This can be further simplified by recalling from the same argument as (A3.2.75) that

$$(A3.2.101) \quad 1 = |I_n| = |AA^{-1}| = |A||A^{-1}| \Rightarrow |A^{-1}| = |A|^{-1} \Rightarrow |A^{-1}|_+ = |A|_+^{-1}$$

so that (A3.2.100) becomes:

$$(A3.2.102) \quad \text{vol}(A^{-1}[\Delta_0(\varepsilon)]) = |A|_+^{-1} \text{vol}[\Delta_0(\varepsilon)]$$

By combining the results in (A3.2.93), (A3.2.97), (A3.2.98) and (A3.2.102), we obtain the identity

$$\begin{aligned} (A3.2.103) \quad g(y_0) \text{vol}[\Delta_0(\varepsilon)] + e_y(\varepsilon) &= \Pr[Y \in \Delta_0(\varepsilon)] \\ &= \Pr(X \in A^{-1}[\Delta_0(\varepsilon)]) \\ &= f(A^{-1}y_0) |A|_+^{-1} \text{vol}[\Delta_0(\varepsilon)] + e_x(\varepsilon) \end{aligned}$$

which after dividing by $\text{vol}[\Delta_0(\varepsilon)]$ and again using (A3.2.102) yields

$$\begin{aligned} (A3.2.104) \quad g(y_0) + \frac{e_y(\varepsilon)}{\text{vol}[\Delta_0(\varepsilon)]} &= f(A^{-1}y_0) |A|_+^{-1} + \frac{e_x(\varepsilon)}{\text{vol}[\Delta_0(\varepsilon)]} \\ &= |A|_+^{-1} \left\{ f(A^{-1}y_0) + \frac{e_x(\varepsilon)}{\text{vol}(A^{-1}[\Delta_0(\varepsilon)])} \right\} \end{aligned}$$

Finally, by letting $\varepsilon \rightarrow 0$ and using (A3.2.95) and (A3.2.99), we obtain the density relation:

$$(A3.2.105) \quad g(y_0) + \lim_{\varepsilon \rightarrow 0} \frac{e_Y(\varepsilon)}{\text{vol}[\Delta_0(\varepsilon)]} = |A|_+^{-1} \left\{ f(A^{-1}y_0) + \lim_{\varepsilon \rightarrow 0} \frac{e_X(\varepsilon)}{\text{vol}(A^{-1}[\Delta_0(\varepsilon)])} \right\}$$

$$\Rightarrow g(y_0) = f(A^{-1}y_0) |A|_+^{-1}$$

But since this is an identity for all choices of y_0 , we can now replace y_0 by y and write

$$(A3.2.106) \quad g(y) = f(A^{-1}y) |A|_+^{-1}$$

This is the key result for constructing density, $g(y)$, from $f(x)$ under linear transformations, $Y = AX$, as in (A3.2.88). Essentially it asserts that the desired density, $g(y)$, at y is obtained by evaluating f at $A^{-1}y$ and rescaling to adjust for the volume changes created by A .

But before applying this result to the multi-normal case, we first extend (A3.2.102) to include *translations* as in (A3.2.87). To do so, observe that if we now let $Z = Y - \mu$ so that

$$(A3.2.107) \quad Y = AX + \mu \Rightarrow Y - \mu = AX \Rightarrow Z = AX$$

then Z is seen to be related to X by a *linear transformation*. Hence if $h(z)$ denotes the density of Z , then it follows from (A3.2.106) that

$$(A3.2.108) \quad h(z) = f(A^{-1}z) |A|_+^{-1}$$

But since Y is related to Z by a simple *translation* operator, T , defined by

$$(A3.2.109) \quad Y = T(Z) = Z + \mu$$

with associated inverse,

$$(A3.2.110) \quad Z = T^{-1}(Y) = Y - \mu$$

we can now use $h(z)$ to obtain $g(y)$ from these relations. Here the key point to note is that translations on \mathbb{R}^n simply shift locations, and involve *no rescaling of volumes*.¹⁵ So in fact, the relation between h and g in this case reduces simply to:

$$(A3.2.111) \quad g(y) = h(T^{-1}y) = h(y - \mu)$$

¹⁵ Here it is worth noting that the terms *isometric transformations* and *rigid motions* mentioned at the beginning of Section A3.1.2 formally include translations as well as rotations and reflections, since all such transformations preserve both distances and angles.

Finally, by combining (A3.2.108) and (A3.2.111), we obtain the desired general relation

$$(A3.2.112) \quad g(y) = f[A^{-1}(y - \mu)] |A|_+^{-1}$$

between densities g and f for linear transformations plus translations in (A3.2.87).

Before applying this to the multi-normal case, we can make one additional observation about covariances that is independent of normality. Recall from expression (3.2.21) in Part II of the NOTEBOOK that

$$(A3.2.113) \quad Y = AX + \mu \Rightarrow \text{cov}(Y) = A \text{cov}(X) A'$$

So if we let $\text{cov}(Y) = \Sigma$ and assume that $\text{cov}(X) = I_n$, then as in the standard normal case of (A3.2.87) we obtain the formal identity:

$$(A3.2.114) \quad \Sigma = A I_n A' = A A'$$

But by the determinantal identities in (A3.2.70) and (A3.2.71), this in turn implies that

$$(A3.2.115) \quad |\Sigma| = |A| |A'| = |A|^2 > 0$$

So (as we have already seen in Sylvester's Condition leading to the Cholesky Theorem in Appendix A2) the determinant of every (nonsingular) covariance matrix is *positive*. This means that "plus" subscripts can be dropped for determinants of covariance matrices. In particular, by letting $|\Sigma|^{1/2}$ denote the positive square root of $|\Sigma|$, it follows that

$$(A3.2.116) \quad |A|_+ = |\Sigma|^{1/2}$$

and hence that (A3.2.112) can also be written as

$$(A3.2.117) \quad g(y) = f[A^{-1}(y - \mu)] |\Sigma|^{-1/2}$$

[with the understanding that $\Sigma = \text{cov}(Y)$].

Finally, to apply these results to the multi-normal case, we need only observe that if X is standard normal, $X \sim N(0, I_n)$, with density in (A3.2.88), then $g(y)$ in (A3.2.117) takes the form:

$$(A3.2.118) \quad g(y) = f[A^{-1}(y - \mu)] |\Sigma|^{-1/2} \\ = \left\{ (2\pi)^{-n/2} e^{-\frac{1}{2}[A^{-1}(y-\mu)]'[A^{-1}(y-\mu)]} \right\} |\Sigma|^{-1/2}$$

$$= (2\pi)^{-n/2} |\Sigma|^{-1/2} e^{-\frac{1}{2}(y-\mu)'(A^{-1})'(A^{-1})(y-\mu)}$$

But by (A3.2.114) together with the matrix identities in (A3.1.18) and (A3.1.20) we see that

$$(A3.2.119) \quad (A^{-1})'(A^{-1}) = (A')^{-1}A^{-1} = (AA')^{-1} = \Sigma^{-1}$$

so that (A3.2.118) becomes

$$(A3.2.120) \quad g(y) = (2\pi)^{-n/2} |\Sigma|^{-1/2} e^{-\frac{1}{2}(y-\mu)'\Sigma^{-1}(y-\mu)}$$

Thus the resulting probability density is precisely that in (A3.2.89), and the multi-normal case is established. In particular, the family of multi-normal random vectors, $Y \sim N(\mu, \Sigma)$, is seen to be generated by transformations, $Y = AX + \mu$, of the standard normal random vector, X , satisfying $\Sigma = AA'$, with A nonsingular. As an immediate consequence of this, we have the following simple proof of the *Linear Invariance Theorem* of Section 3.2.2 of Part II, which we now restate for convenience as:

Linear Invariance Theorem. For any multi-normal random vector, $X \sim N(\mu, \Sigma)$, and affine transformation, $Y = AX + b$, of X with A of full row rank, Y is also multi-normally distributed as

$$(A3.2.121) \quad Y \sim N(A\mu + b, A\Sigma A')$$

Proof: If C denotes the Cholesky decomposition of Σ so that $\Sigma = CC'$, and if we define the random vector, Z , by

$$(A3.2.122) \quad Z = C^{-1}(X - \mu)$$

so that by construction,

$$(A3.2.123) \quad X = CZ + \mu$$

then the argument above shows that $Z \sim N(0, I_n)$. But since

$$(A3.2.124) \quad Y = AX + b = A(CZ + \mu) + b = (AC)Z + (A\mu + b)$$

shows that Y is an affine transformation of Z , the same argument shows that Y is multi-normally distributed. Moreover, from expressions (3.2.18) and (3.2.21) in Part II we see that the mean and covariance of Y are given respectively by

$$(A3.2.125) \quad E(y) = AC E(Z) + (A\mu + b) = A\mu + b, \text{ and}$$

$$(A3.2.126) \quad \begin{aligned} \text{cov}(Y) &= \text{cov}[ACZ + (A\mu + b)] = (AC)\text{cov}(Z)(AC)' \\ &= AC(I_n)C'A' = A(CC')A' = A\Sigma A' \end{aligned}$$

Thus we must have $Y \sim N(A\mu + b, A\Sigma A')$, and the result is established. \square

Finally it is important to clarify the above requirement that A be of *full row rank*. Note in particular that if A has fewer rows than columns, say $m < n$, then the random vector, $Y = AX + b$, must be of dimension m (where b must also be m -dimensional so that vector addition is well defined). So it is implicitly assumed that $N(A\mu + b, A\Sigma A')$ is a multi-normal distribution on \mathbb{R}^m with density given by replacing μ and Σ in (A3.2.120) with $A\mu + b$ and $A\Sigma A'$, respectively. With this in mind, it should be clear from (A3.2.120) that such a density is only defined if $A\Sigma A'$ is a *nonsingular* covariance matrix (i.e., with a well defined inverse). As shown in Corollary 3 of Section A3.4.3 below, the condition that A be of full row rank, insures that the m -square covariance matrix, $A\Sigma A'$, will indeed be nonsingular.

A3.3 Eigenvalues and Eigenvectors

As stated earlier, the single most important application of the Singular Value Decomposition Theorem for our purposes is to provide a simple proof of the *Spectral Decomposition Theorem* for symmetric matrices. Recall from (A3.2.2) that this theorem asserts that if matrix A is symmetric (i.e., $A' = A$) then there exists an orthonormal matrix, U , and diagonal matrix, Λ , such that

$$(A3.3.1) \quad A = U\Lambda U'$$

The elements of $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$ are called the *eigenvalues* of A and the columns of $U = (u_1, \dots, u_n)$ are the associated *eigenvectors*. However, these concepts are much more general, and indeed, provide additional geometric intuition about linear transformations in general. In particular, it is useful to consider eigenvalues and eigenvectors for *nonnegative* spatial weight matrices, W , which may possibly be non-symmetric (as for example in the case of nearest-neighbor matrices). So it is convenient to start with a broader consideration of these concepts, and then focus in on symmetric matrices.

For any given n -square matrix, A , and *nonzero* vector, $x \in \mathbb{R}^n$, if A maps x into a scalar multiple of itself, i.e., if

$$(A3.3.2) \quad Ax = \lambda x$$

for some scalar, $\lambda \in \mathbb{R}$, then λ is designated as a *eigenvalue* of A with associated *eigenvector*, x .¹ In geometric terms, A simply “stretches” or “shrinks” each eigenvector, x , by a factor, λ , with a reversal in direction if $\lambda < 0$.

Before analyzing these concepts, it is important to reiterate that our present view of (real-valued) matrices, $A \in \mathbb{R}^{n \times n}$, is as representations of linear transformations on \mathbb{R}^n . So our focus is naturally on the geometric properties of such transformations on the *real vector space*, \mathbb{R}^n . But such matrices can also be viewed as representing a class of linear transformations on the *complex vector space*, \mathbb{C}^n . This distinction is important for the present discussion because the general theory of eigenvalues and eigenvectors treats such matrices as linear transformations on \mathbb{C}^n . The reason for this can be seen by the following equivalent view of eigenvalues. If we rewrite (A3.3.2) as,

$$(A3.3.3) \quad Ax - \lambda x = 0 \Rightarrow (A - \lambda I_n)x = 0$$

then it becomes clear that the eigenvalues of A are precisely those values for which the matrix, $A - \lambda I_n$, is *singular*. So, as was observed following expression (A3.2.86) above, this is equivalent to the condition that

¹ The word “eigen” is German for “own” as in belonging to. So the eigenvalues of A are also referred to as its “own” values or “characteristic” values.

$$(A3.3.4) \quad |A - \lambda I_n| = 0$$

But by the definition of determinants, this is simply a polynomial equation in λ , called the *characteristic equation* for A . For example, if $n = 2$, then by (A3.2.67) above expression (A3.3.4) takes the form,

$$(A3.3.5) \quad 0 = \begin{vmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{vmatrix} - \begin{vmatrix} \lambda & \\ & \lambda \end{vmatrix} = (a_{11} - \lambda)(a_{22} - \lambda) - a_{12}a_{21}$$

So the eigenvalues of 2×2 matrices are thus seen to be the roots of a quadratic equation. More generally they are the roots of an n^{th} -degree polynomial called the *characteristic polynomial* for A . In this setting, the key result here is of course the *Fundamental Theorem of Algebra*, which tells us that there are always exactly n roots to this equation (counting repetitions) if we allow *complex-valued* roots. So if A is regarded as a linear transformation on \mathbb{C}^n , where both λ and x in (A3.3.2) can be complex-valued, then one obtains a very elegant and powerful theory of eigenvalues and eigenvectors. But from a *geometric* viewpoint, there is a fundamental difference between the simple scaling of real vectors in \mathbb{R}^n and the corresponding interpretation of expression (A3.3.2) in \mathbb{C}^n . In particular, multiplication of complex numbers involves *rotation* as well as scaling. We shall return to these issues in Section A3.5.4 below, where the geometric meaning of such rotations will be interpreted in \mathbb{R}^n . But for the present our attention is restricted to the case of real eigenvalues. Indeed, one major objective of these notes is to show that the *eigenvalues of symmetric matrices are always real* – without any appeal to complex numbers whatsoever. Hence, unless otherwise stated, we implicitly assume that the relevant eigenvalues and eigenvectors for matrices, A , in this section are real valued, i.e., are meaningful for A as a linear transformation on \mathbb{R}^n .

In this setting, we begin by noting from (A3.3.2) that each eigenvector, x , for λ is inherently nonunique. In particular, every nonzero scalar multiple, αx , of x is also an eigenvector for λ , since

$$(A3.3.6) \quad A(\alpha x) = \alpha Ax = \alpha \lambda x = \lambda(\alpha x)$$

To remove such obvious redundancies, representative eigenvectors are by convention normalized to have unit length, $\|x\| = 1$.

With this normalization, the next question concerns the relation between eigenvectors for *distinct* eigenvalues. Our objective is to show that such eigenvectors must always be *linearly independent*. Here some geometric intuition can be gained by considering several examples. We start with the simplest and most transparent example of eigenvalues and associated eigenvectors, namely those for diagonal matrices, $A = \text{diag}(a_{11}, \dots, a_{mm})$. Here it is obvious that

$$(A3.3.7) \quad AI_n = I_n A \Rightarrow A(e_1, \dots, e_n) = (e_1, \dots, e_n) \begin{pmatrix} a_{11} & & \\ & \ddots & \\ & & a_{nn} \end{pmatrix}$$

$$\Rightarrow Ae_i = a_{ii} e_i, \quad i = 1, \dots, n$$

So if we now denote the *set of eigenvalues* for any matrix, A , by $Eig(A)$, then for diagonal matrices in (A3.3.7) it is clear that $Eig(A) = \{a_{ii} : i = 1, \dots, n\}$ with associated eigenvectors, e_i , $i = 1, \dots, n$. This example shows that n -square matrices can indeed have n distinct eigenvalues. Notice also that all eigenvectors in this case are in fact *orthogonal*, and hence are necessarily linearly independent even if their eigenvalues are not distinct. We shall see below that this property is shared by *all symmetric matrices* (of which diagonal matrices are the simplest example).

Of course, the orthogonal basis in (A3.3.7) is a very special case. A more typical example with a full set of eigenvalues is given by the following simple matrix

$$(A3.3.8) \quad A = \begin{pmatrix} 3 & 1 \\ 0 & 2 \end{pmatrix}$$

for which it can easily be verified that the eigenvalues of A are $Eig(A) = \{\lambda_1, \lambda_2\} = \{3, 2\}$ with associated eigenvectors given respectively by $x_1 = e_1 = (1, 0)'$ and

$x_2 = (-1/\sqrt{2}, 1/\sqrt{2})'$, as shown in Figure A3.26 below.




Figure A3.26. Non-orthogonal Example

In both these examples, the eigenvectors associated with distinct eigenvalues are indeed linearly independent. But the question remains as to whether this is always true. To see that it is, we now consider a general matrix, A , and suppose that λ_1 and λ_2 are two

distinct eigenvalues of A with associated eigenvectors, x_1 and x_2 . Clearly, x_1 and x_2 cannot themselves be linearly dependent since this would mean that $x_1 = \alpha x_2$ for some $\alpha \neq 0$. But the normalization condition, $\|x_1\| = \|x_2\| = 1$, would then imply that $\alpha = 1$ and hence that $x_1 = x_2$, which is not possible for distinct eigenvalues λ_1 and λ_2 . However, we can still ask whether there could possibly be another vector, x_3 , in $\text{span}(x_1, x_2)$ which is also an eigenvector of A with distinct eigenvalue, $\lambda_3 \neq \lambda_i, i = 1, 2$. A representation of $\text{span}(x_1, x_2)$ is shown in Figure A3.27 below where it is assumed for sake of illustration that $0 < \lambda_1 < \lambda_2$ and that x_3 (shown in blue) is a positive linear combination, $x_3 = a x_1 + b x_2$, of x_1 and x_2 .




Figure A3.27. Linear Independence Example

Now if x_3 were an eigenvector of A with eigenvalue, λ_3 , then by definition,

$$(A3.3.9) \quad Ax_3 = \lambda_3 x_3 = \lambda_3(ax_1 + bx_2) = (\lambda_3 a)x_1 + (\lambda_3 b)x_2$$

So the coefficients $(\lambda_3 a, \lambda_3 b)$ of this new linear combination of x_1 and x_2 would necessarily be proportional to the original coefficients (a, b) , as shown by all points on the blue line in the figure. But by hypothesis,

$$(A3.3.10) \quad Ax_3 = A(ax_1 + bx_2) = a(Ax_1) + b(Ax_2) = (a\lambda_1)x_1 + (b\lambda_2)x_2$$

which together with $0 < \lambda_1 < \lambda_2$, shows that in fact more weight is now placed on the maximal eigenvector, x_2 , and thus that proportionality *cannot* hold. More generally, the same argument shows that the image of any vector, $x_3 \in \text{span}(x_1, x_2)$ [not collinear with either x_1 or x_2] is necessarily “pulled toward” this maximal eigenvector (shown by the arrow in the figure), and cannot itself be an eigenvector. So we may conclude that no

eigenvector with eigenvalue distinct from λ_1 and λ_2 can be collinear with (x_1, x_2) , i.e., can lie in $\text{span}(x_1, x_2)$.

While this illustration involves only triples of distinct eigenvalues, the argument is in fact quite general, and can be used to show that *eigenvectors for distinct eigenvalues must always be linearly independent*.² But since our main interest is in symmetric matrices, where the argument will be seen to be even more transparent, the above example suffices for our purposes.

A final property of eigenvalues relates to their possible repetitions, and can again be illustrated most easily by diagonal matrices, $A = \text{diag}(a_{11}, \dots, a_{nn})$. Notice in particular that this is the one case where the characteristic equation in (A3.3.4) is completely transparent, since

$$(A3.3.11) \quad 0 = |A - \lambda I_n| = \left| \begin{pmatrix} a_{11} - \lambda & & \\ & \ddots & \\ & & a_{nn} - \lambda \end{pmatrix} \right| = (a_{11} - \lambda) \cdots (a_{nn} - \lambda)$$

This implies at once that the diagonal elements of A are indeed the roots of its characteristic equation. If some of these diagonal elements are the same, then such repeated roots are designated as *algebraic multiplicities*. For example, the matrix $A = \text{diag}(1, 1, 3)$ has only two distinct eigenvalues, $\text{Eig}(A) = \{1, 3\}$, but since $(1 - \lambda)$, appears twice in (A3.3.11), this eigenvalue is said to have an *algebraic multiplicity* of two. Notice also from (A3.3.7) that since there are two linearly independent eigenvectors for this eigenvalue, namely e_1 and e_2 , its *geometric multiplicity* (i.e., the maximum number of its linearly independent eigenvectors) is also two. More generally, it follows at once from (A3.3.7) that algebraic and geometric multiplicities of eigenvalues are *always identical* for diagonal matrices.

But for general matrices, even when eigenvalues do exist, these two multiplicities need not be the same. For example, while the algebraic and geometric multiplicities of $\lambda = 2$ in the diagonal matrix, $A = \text{diag}(2, 2)$, are both equal to two, consider the following (modest) variation of this matrix:

$$(A3.3.12) \quad A = \begin{pmatrix} 2 & 1 \\ 0 & 2 \end{pmatrix}$$

This matrix is still nonsingular, and moreover, has the same characteristic equation, since $0 = |A - \lambda I_2| = (2 - \lambda)(2 - \lambda)$. So the algebraic multiplicity of $\lambda = 2$ is two. But observe that if $x = (x_1, x_2)'$ is any associated eigenvector, then

² A simpler and more elegant proof of this fact is given in Lemma 1.3.8 of Horn and Johnson (1985). The advantage of the present argument is that it provides some geometric intuition as to *why* this is true.

$$(A3.3.13) \quad Ax = 2x \Rightarrow \begin{pmatrix} 2 & 1 \\ 0 & 2 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = 2 \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \Rightarrow \begin{pmatrix} 2x_1 + x_2 = 2x_1 \\ 2x_2 = 2x_2 \end{pmatrix} \Rightarrow x_2 = 0$$

Moreover, there is only *one* eigenvector (up to a choice of sign) with this property, namely, $x = (1, 0)'$. So the geometric multiplicity this eigenvalue is *one*. Such matrices are usually designated as *defective matrices* in the literature. The reason for this “defective” property can be seen by plotting the transformation, as in Figure A3.28 below:




Figure A3.28. A “Defective” Transformation

Here we have used the vector notation, $v = (x, y)'$, for points in the plane, and have displayed the unique eigenvector by $v_0 = (1, 0)'$, with associated image,

$Av_0 = 2v_0 = (0, 2)'$. To show where all other points are sent, we have fixed the y-coordinate value at $y = 1$, and have plotted the four points, $v_1 = (-2, 1)'$, $v_2 = (-1, 1)'$, $v_3 = (0, 1)'$, and $v_4 = (1, 1)'$, as shown in blue. Multiplying each of these four vectors by matrix A in (A3.3.12), we obtain the corresponding image vectors, $Av_1 = (-3, 2)'$,

$Av_2 = (-1, 2)'$, $Av_3 = (1, 2)'$, and $Av_4 = (3, 2)'$, shown in red. The key point to notice is that all these image vectors are *to the right* of the original vectors, indicating that (along with a certain amount of stretching) each vector has been *rotated clockwise* toward the eigenvector, v_0 . Similarly, by extending all vector arrows in the opposite direction through the origin, it is clear that the vectors, $-v_1, -v_2, -v_3, -v_4$, are also rotated clockwise toward the negative eigenvector, $-v_0$. This shows that all nonzero vectors other than these unique eigenvectors are rotated clockwise to some degree, and thus cannot be eigenvectors. So essentially, such matrices involve some form of *non-rigid*

rotations that can reduce the number of linearly independent eigenvectors associated with repeated eigenvalues.

Given these general properties of real eigenvalues and eigenvectors, our objective is to apply these concepts to symmetric matrices in particular. But before doing so, it is important to reiterate that not all matrices have a full complement of real eigenvalues. The following simple orthonormal matrix will turn out to be a particularly important case in point:

$$(A3.3.14) \quad U = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$$

Geometrically, this matrix rotates the plane counterclockwise through an angle of 90° , as shown in Figure A3.29 below. Clearly no vector can possibly be mapped by this transformation into a scalar times itself.




Figure A3.28. Example with No Eigenvalues

In algebraic terms, the characteristic equation of U takes the form,

$$(A3.3.15) \quad 0 = |U - \lambda I_2| = \lambda^2 + 1$$

which is seen to have only the “imaginary” solutions, $\lambda = \pm\sqrt{-1}$. We shall return to to examples of this type in Section A3.5.4 below.

A3.4 Spectral Decomposition Theorem

We begin by recalling from the very beginning of Section A3.2 that there seems to be an obvious relation between the Spectral Decomposition (SPD) Theorem for symmetric matrices and the Singular Value Decomposition (SVD) Theorem for general matrices. Since the SVD Theorem shows that for *every* matrix, A , there exist orthonormal matrices, U , V , and a diagonal matrix of singular values, S , such that

$$(A3.4.1) \quad A = U S V'$$

it follows that once for symmetric matrices, A , we must have

$$(A3.4.2) \quad U S V' = A = A' = V S U'$$

So at first glance, this identity would appear to suggest that $U = V$, and thus that (A3.3.1) must hold with $\Lambda = S$. To see that this intuition is wrong, recall that $|U| = \pm 1$, which together with the nonnegativity of the singular values, S , must imply that

$$(A3.4.3) \quad |A| = |U| |S| |U'| = |U| |S| |U| = |U|^2 |S| = |S| \geq 0$$

and thus that the determinant of *every* symmetric matrix is *nonnegative*! But we have already seen from (A3.2.69) that the symmetric (orthonormal) matrix

$$(A3.4.4) \quad A = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

has a *negative* determinant, $|A| = -1$. More generally, the fact that singular values are by construction nonnegative, shows that the relation between singular values and eigenvalues for symmetric matrices is not immediately obvious.

This is made even more clear by a closer examination of this particular counterexample. Here one can verify (by direct multiplication) that (A3.3.2) holds for this matrix A with

$$(A3.4.5) \quad U = \begin{pmatrix} 0 & -1 \\ -1 & 0 \end{pmatrix}, \quad S = \begin{pmatrix} 1 & \\ & 1 \end{pmatrix}, \quad V = \begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix}$$

Moreover, since U and V are easily seen to be orthonormal, this is indeed a singular value decomposition of A with $U \neq V$. Here it can also be verified by direction computation that

$$(A3.4.6) \quad U S U' = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix} \neq A \neq \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} = V S V'$$

so that neither U nor V yield spectral decompositions of A with $\Lambda = S$. But it turns out that A does indeed have a *unique* spectral decomposition:

$$(A3.4.7) \quad A = W \Lambda W'$$

with orthonormal matrix, W , and diagonal matrix, Λ , given by

$$(A3.4.8) \quad W = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -1 \\ -1 & -1 \end{pmatrix}, \quad \Lambda = \begin{pmatrix} -1 & \\ & 1 \end{pmatrix}$$

So at first glance, there would seem to be little relation between the decompositions if A in (A3.4.5) and (A3.4.8). But closer inspection show that the *absolute value* of Λ is precisely S . As we shall see below, this relationship is fundamental.

A3.4.1 Eigenvalues and Eigenvectors of Symmetric Matrices

To gain further insight, it is convenient for the moment to *suppose* that the SPD Theorem is true, and to examine its geometric consequences. To do so, note first that if (A3.3.1) holds for a symmetric matrix, A , then since $U^{-1} = U'$, it follow at once that

$$(A3.4.9) \quad A = U\Lambda U' \Rightarrow AU = U\Lambda(UU') = U\Lambda$$

$$\Rightarrow A(u_1, \dots, u_n) = (u_1, \dots, u_n) \begin{pmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{pmatrix}$$

$$\Rightarrow Au_i = \lambda_i u_i, \quad i = 1, \dots, n$$

Thus, as an extension of the diagonal-matrix case in (A3.3.7), we see that the diagonal elements $(\lambda_1, \dots, \lambda_n)$ of Λ must indeed be the *eigenvalues* of A with associated *orthonormal eigenvectors* (u_1, \dots, u_n) , as asserted at the beginning of Section A3.3. Note also that by definition this decomposition implies that all eigenvalues must be *real*.

Moreover, these eigenvalues and eigenvectors together imply that such matrices (like diagonal matrices) are in fact representations of *scale transformations with respect to some appropriate coordinate system*. This can be illustrated in two dimensions by the symmetric matrix,

$$(A3.4.10) \quad A = \begin{pmatrix} 3 & 1 \\ 1 & 3 \end{pmatrix}$$

Here A does indeed have a spectral decomposition as in (A3.3.1) with $\Lambda = \text{diag}(2, 4)$ and orthonormal matrix,

$$(A3.4.11) \quad U = \frac{1}{\sqrt{2}} \begin{pmatrix} -1 & 1 \\ 1 & 1 \end{pmatrix} = \left[\begin{pmatrix} \frac{-1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{pmatrix} \begin{pmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{pmatrix} \right] = (u_1, u_2)$$

[which is precisely W in (A3.4.8) above]. So the eigenvectors for this matrix are the two diagonal vectors, u_1 and u_2 , with corresponding eigenvalues, $(\lambda_1 = 2, \lambda_2 = 4)$, as shown in Figure A3.29 below:




Figure A3.29. Positive Eigenvalue Case

So if (u_1, u_2) are regarded as the coordinate axes, then A is seen to be a pure scale transformation with respect to this coordinate system. More generally, if the spectral decomposition of A is regarded as a composition of the respective transformations, U' , A and U , then we obtain a diagram very reminiscent of that in Figure A3.16, with V in the last step replaced by U . In particular, the eigenvectors (i.e., columns of U) correspond precisely to the *principle axes* of the ellipsoidal image, $A(\mathbb{S}_n)$, of the unit circle, \mathbb{S}_n , as seen for $n = 2$ in the figure. So this shows geometrically that there must be an intimate connection between the *singular value decomposition* (SVD) and *spectral decomposition* (SPD) of symmetric matrices.

In fact for the present matrix, A , in (A3.4.10), these two decompositions are *identical*. The special feature of this symmetric matrix that leads to this identity is that its eigenvalues are all *positive*. What this implies is that these eigenvalues play exactly the same role as singular values, i.e., they measure the lengths of these axes from the origin. More generally, this suggests that the lengths of such axes for symmetric matrices, A , should be precisely the *absolute values* of their eigenvalues. In other words, the eigenvalues of A should differ only in *sign* from the associated singular values of A .

All these conjectures will be shown to be true in the following sections. But for the moment, we continue with our illustrations by reconsidering the counterexample in expression (A3.4.4) above. As mentioned already, the eigenvectors here are precisely the

same as in (A3.4.11) above. So only the eigenvalues are different, as shown in Figure A3.30 below.




Figure A3.29. Mixed Signs Case

The key point to notice [as was evident in the SVD of this matrix in expression (A3.4.5) above] is that the unit circle is mapped onto itself, i.e., $A(\mathbb{S}_2) = \mathbb{S}_2$. So *any* set of orthonormal axes can be used for an SVD. However, this is not true for the SPD. In fact there is *exactly one* eigenvector, u_1 , associated with eigenvalue, $\lambda_1 = -1$, and exactly one eigenvector, u_2 , associated with eigenvalue, $\lambda_2 = 1$.³ So in contrast to the SVD, this SPD is essentially unique. But it will be shown below that in spite of its nonuniqueness, the SVD for A still contains enough information to allow the SPD to be constructed explicitly. The essential reason for this is that relations between U and V implicit in the identity (A3.4.2) will yield additional analytical information.

Before proceeding to these analytical results, it should be noted that there is one additional complication that cannot be illustrated in two dimensions. Consider the following 4-dimensional version of the matrix in (A3.4.4) above:

$$(A3.4.12) \quad A = \begin{pmatrix} & & & 1 \\ & & 1 & \\ & 1 & & \\ 1 & & & \end{pmatrix}$$

which can be seen (by direct multiplication) to have eigenvalues, $A = \text{diag}(-1, -1, 1, 1)$ with associated eigenvectors:

³ Again, remember that if u_i is an eigenvector for λ_i , then so is $-u_i$. So we are implicitly ignoring this trivial form of nonuniqueness.

$$(A3.4.13) \quad U = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 & 1 & 1 & 0 \\ -1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 \\ 0 & -1 & 1 & 0 \end{pmatrix} = (u_1, u_2, u_3, u_4)$$

Note in particular that the unit sphere is again mapped onto itself, i.e., $A(\mathbb{S}_4) = \mathbb{S}_4$, so that any set of four mutually orthonormal axes can again be used to define an SVD. But now, the two distinct eigenvalues, 1 and -1, both have two-dimensional spaces of eigenvectors, namely $\text{span}(u_1, u_2)$ and $\text{span}(u_3, u_4)$, respectively. So even the SPD is nonunique in this example. Such cases require more effort to construct an admissible SPD from any given SVD. So this general case will be treated by itself.

With these examples in mind, we now proceed to establish the Spectral Decomposition (SPD) Theorem in stages. The first task is to establish some general consequences of singular value decompositions (SVD) for symmetric matrices. This will provide a general foundation for the SPD results to follow.

A3.4.2 Some Consequences of SVD for Symmetric Matrices

Here we focus on the additional information contained in the identity (A3.4.2) for symmetric matrices. These equalities can be rewritten in the following way:

$$(A3.4.14) \quad A = U S V' \Rightarrow A V = U S$$

$$(A3.4.15) \quad A = A' = V S U' \Rightarrow A U = V S$$

By adding the right hand sides we obtain

$$(A3.4.16) \quad A(V + U) = (U + V)S = (V + U)S$$

and similarly by subtracting the right hand sides we have

$$(A3.4.17) \quad A(V - U) = (U - V)S = (V - U)(-S)$$

But if we now let

$$(A3.4.18) \quad X = U + V$$

$$(A3.4.19) \quad Y = V - U$$

then (A3.4.16) and (A3.4.17) yield the associated sets of *eigenvalue equations*:

$$(A3.4.20) \quad A X = X S$$

$$(A3.4.21) \quad AY = Y(-S)$$

Note that neither matrix, $X = (x_1, \dots, x_n)$ or $Y = (y_1, \dots, y_n)$, is orthonormal, or even orthogonal. But nonetheless, the respective rows of each relation (A3.4.20) and (A3.4.21) yield well defined *eigenvalue relations*:

$$(A3.4.22) \quad Ax_i = s_i x_i, \quad i = 1, \dots, n$$

$$(A3.4.23) \quad Ay_i = (-s_i) y_i, \quad i = 1, \dots, n$$

This shows us that each *nonzero* column of X and Y , namely each $x_i = u_i + v_i \neq 0$ and $y_i = u_i - v_i \neq 0$, respectively, must yield corresponding *real* eigenvalues, s_i or $-s_i$ for symmetric matrix, A . As we shall see below, many columns of X and/or Y must be zero. But the key point to notice is that for all $i = 1, \dots, n$, the column pair (x_i, y_i) *cannot both be zero*. For if so then

$$(A3.4.24) \quad u_i + v_i = x_i = 0 = y_i = u_i - v_i \\ \Rightarrow v_i = -v_i \Rightarrow v_i = 0$$

which contradicts the normalization condition, $\|v_i\| = 1$. So conditions (A3.4.22) and (A3.4.23) together will provide us with a full complement of *real eigenvalues* for A in every case.

Thus the first major consequence of these observations is that without loss of generality we can focus our attention on *real* eigenvalues for symmetric matrices. This is of sufficient importance to be stated formally. If the set of *distinct eigenvalues* for any matrix, A , is denoted by $Eig(A)$, and if we define symmetric matrices by the condition that $A = A'$, then this first consequence of SVD can be stated as follows:⁴

$$(A3.4.25) \quad A = A' \Rightarrow Eig(A) \subset \mathbb{R}$$

A second consequence (as suggested by the examples above) is that all eigenvalues in (A3.4.22) and (A3.4.23) are either the singular values of A or their negatives. So the absolute magnitudes of all eigenvalues can be determined by the SVD of A . To state this more formally, let the set of *distinct singular values* of any matrix, A , be denoted by $Sing(A)$, and let the negatives of these values be denoted by $-Sing(A)$. Then, in a manner paralleling (A3.4.25), this second consequence of SVD can be stated as follows:

⁴ The standard proof of this fact is to show that eigenvalues of symmetric matrices must always be equal to their complex conjugates, and hence must be real (see for example Theorem 4.1.3 in Horn and Johnson, 1985).

$$(A3.4.26) \quad A = A' \Rightarrow \text{Eig}(A) \subseteq \text{Sing}(A) \cup [-\text{Sing}(A)]$$

There is a third important consequence that relates to the eigenvectors associated with distinct eigenvalues of symmetric matrices. Recall that in Figure A3.27 above a geometric argument was sketched showing that eigenvectors for distinct (real) eigenvalues are always linearly independent. For symmetric matrices we have the stronger property that such eigenvectors must actually be *orthogonal*. This can be demonstrated as follows:⁵

Orthogonality of Distinct Eigenvectors. For any symmetric matrix, A , and eigenvectors, x_i, x_j , associated with distinct eigenvalues, $\lambda_i, \lambda_j \in \text{Eig}(A)$,

$$(A3.4.27) \quad \lambda_i \neq \lambda_j \Rightarrow x_i'x_j = 0$$

Proof: By definition we must have,

$$(A3.4.28) \quad Ax_i = \lambda_i x_i \quad , \quad \text{and}$$

$$(A3.4.29) \quad Ax_j = \lambda_j x_j$$

But premultiplying (A3.4.28) by x_j' and employing the symmetry of A , we see that,

$$(A3.4.30) \quad \lambda_i x_j'x_i = x_j'Ax_i = (Ax_i)'x_j = x_i'Ax_j = x_i'(\lambda_j x_j) = \lambda_j x_i'x_j = \lambda_j x_j'x_i \\ \Rightarrow (x_j'x_i)(\lambda_i - \lambda_j) = 0$$

So if $\lambda_i \neq \lambda_j$ then we may conclude that $x_j'x_i = 0 = x_i'x_j$. \square

Given these properties of eigenvalues and eigenvectors for symmetric matrices, the key questions remaining are (i) how to identify which of the values on the right hand side of (A3.4.26) are relevant in any particular case, and (ii) how to construct their associated eigenvectors in terms of the SVD of A . To answer these questions, we shall proceed on a case-by-case basis from the simplest to the most general cases.

⁵ It should be noted that both the statement and proof of this result make constant use of property (A3.4.25), since eigenvectors for real eigenvalues can always be restricted to \mathbb{R}^n . This allows orthogonality (and indeed all inner products) to be defined solely on \mathbb{R}^n . While this same analysis can of course be carried on \mathbb{C}^n using complex inner products, property (A3.4.25) shows that this is not necessary for real symmetric matrices.

A3.4.3 Spectral Decomposition for Symmetric Definite and Semidefinite Matrices

The simplest and by far the most important cases for our purposes all involve symmetric definite or semidefinite matrices. So this is the best place to begin. Recall from expressions (A2.7.36) and (A2.7.67) in Appendix A2 that an n -square matrix, A , is *positive semidefinite* iff for all $x \in \mathbb{R}^n$,

$$(A3.4.31) \quad x \neq 0 \Rightarrow x'Ax \geq 0$$

and is *positive definite* iff this inequality is strict, i.e., iff

$$(A3.4.32) \quad x \neq 0 \Rightarrow x'Ax > 0$$

Moreover, A is *negative definite (semidefinite)* iff $-A$ is positive definite (semidefinite). Since all results for symmetric positive definite and semidefinite matrices are immediately extendable to their negative counterparts by just reversing signs, we focus only on (A3.4.31) and (A3.4.32). Hence our first result is to show that for symmetric positive semidefinite matrices, A , the SVD and SPD of A are essentially identical. In particular, the eigenvalues of A are *precisely* its singular values, and their associated eigenvectors can be taken directly from the SVD of A . Moreover, if A is *positive definite* then all eigenvalues of A are positive, and each SVD for A is *precisely* an SPD for A . These results can be stated more formally as follows:⁶

Spectral Decomposition of Symmetric Positive Semidefinite Matrices

(i) If A is a symmetric positive semidefinite matrix with SVD,

$$(A3.4.33) \quad A = U S V' ,$$

then it must be true that

$$(A3.4.34) \quad A = U S U' = V S V'$$

(ii) If in addition A is positive definite, then $\text{diag}(S) > 0$ and $U = V$.

Proof: (i) To establish the first equality, it must be shown that

$$(A3.4.35) \quad A u_i = s_i u_i$$

⁶ It is of interest to note that a direct proof for this case follows from the standard construction of principle components in multivariate analysis, which in fact closely parallels the above proof of the Singular Decomposition Theorem. See for example the classic treatment in Anderson (1958, pp.273-275).

for all $i = 1, \dots, n$. But by applying the same column decomposition in (A3.2.28) to (A3.4.15) for the SVD in (A3.4.33), it follows that

$$(A3.4.36) \quad Au_i = s_i v_i, \quad i = 1, \dots, n$$

Given this representation, there are two cases to consider. First if $s_i = 0$, then it follows at once from (A3.4.36) that,

$$(A3.4.37) \quad Au_i = s_i v_i = 0 = s_i u_i$$

On the other hand if $s_i > 0$, then observe from (A3.4.23) that we must have $y_i = 0$. For if not, then since $y_i \neq 0 \Rightarrow y_i' y_i > 0$, it would follow from (A3.4.23) that

$$(A3.4.38) \quad Ay_i = -s_i y_i \Rightarrow y_i' Ay_i = -s_i y_i' y_i < 0$$

which contradicts the positive semidefiniteness of A . Thus we must have

$$(A3.4.39) \quad 0 = y_i = u_i - v_i \Rightarrow u_i = v_i$$

and may conclude again from (A3.4.36) that

$$(A3.4.40) \quad Au_i = s_i v_i = s_i u_i$$

So (A3.4.35) must hold in all cases, and the first equality (A3.4.34) is established. The second equality follows in exactly the same way by replacing (A3.4.15) with (A3.4.14) and thus switching the roles of u_i and v_i in (A3.4.36).

(ii) Finally, if A is positive definite, then since $u_i' u_i = \|u_i\|^2 = 1$ for each $i = 1, \dots, n$, it follows from (A3.4.35) that

$$(A3.4.41) \quad Au_i = s_i u_i \Rightarrow u_i' Au_i = s_i u_i' u_i = s_i$$

and hence from positive definiteness that $s_i > 0$. Thus $\text{diag}(S) = \text{diag}(s_1, \dots, s_n) > 0$. Moreover since the argument in (A3.4.38) and (A3.4.39) now holds for all $i = 1, \dots, n$, it also follows that $U = V$. \square

For symmetric positive definite matrices, A , the above theorem (now referred to as *SPD Theorem 1*), shows that the two decompositions, SVD and SPD, of A exhibit a *one-to-one* correspondence. As a direct consequence of this correspondence, we now have the following additional characterizations of positive definiteness:

Corollary 1. For any symmetric positive semidefinite matrix, A , the following three properties are equivalent:

- (A3.4.42) A is positive definite.
 (A3.4.43) A has all positive eigenvalues.
 (A3.4.44) A is nonsingular.

Proof: To establish this equivalence it suffices to show that (A3.4.42) \Rightarrow (A3.4.43) \Rightarrow (A3.4.44) \Rightarrow (A3.4.42). But if A is positive definite and λ is any eigenvalue of A with eigenvector, x , then by (A3.4.32) together with $x'x > 0$ it follows that,

$$(A3.4.45) \quad Ax = \lambda x \Rightarrow x'Ax = \lambda x'x \Rightarrow \lambda = \frac{x'Ax}{x'x} > 0$$

and thus that all eigenvalues must be positive. Next, to show that positive eigenvalues imply nonsingularity, observe since the symmetric positive semidefiniteness of A implies from part (i) of SPD Theorem 1 that A has a spectral decomposition,

$$(A3.4.46) \quad A = U S U'$$

[given by the first equality in (A3.4.34)], it follows that if all eigenvalues are positive, then the positive diagonal matrix, S , in (A3.4.46) must have a well defined inverse, S^{-1} . But this together with the orthonormality of matrix, U , implies that

$$(A3.4.47) \quad U S^{-1} U' = (U S U')^{-1} = A^{-1}$$

and thus that A is nonsingular. Finally, to show that nonsingularity implies positive definiteness, note first from the nonnegativity of the diagonal matrix, S , in (A3.4.46) that S has a well defined *square root*,

$$(A3.4.48) \quad S^{1/2} = \text{diag}(s_1^{1/2}, \dots, s_n^{1/2})$$

satisfying $S = S^{1/2} S^{1/2}$. So for any $x \in \mathbb{R}^n$ it follows that,

$$(A3.4.49) \quad x'Ax = x'U S U'x = x'U S^{1/2} S^{1/2} U'x = (S^{1/2} U'x)'(S^{1/2} U'x) = \|S^{1/2} U'x\|^2$$

But since for any vector z , $\|z\|^2 = 0 \Rightarrow \|z\| = 0 \Rightarrow z = 0$, we see from (A3.4.49) that

$$(A3.4.50) \quad x'Ax = 0 \Rightarrow S^{1/2} U'x = 0 \Rightarrow U S^{1/2} (S^{1/2} U'x) = 0 \Rightarrow Ax = 0$$

So if $x'Ax = 0$ for any $x \neq 0$, then it would also be true that $Ax = 0$, which contradicts the nonsingularity of A . Thus, nonsingularity together with the positive semidefiniteness

of A imply that $x'Ax > 0$ must hold whenever $x \neq 0$, and it follows that A is positive definite. \square

For our present purposes, the single most important application of these results is to characterize the spectral properties of covariance matrices. To begin with, we can now give a more complete statement of the Positive Definiteness Property for nonsingular covariance matrices stated in Appendix A2 (page A2-27):

Corollary 2. *Every nonsingular covariance matrix, Σ , is positive definite with spectral decomposition:*

$$(A3.4.51) \quad \Sigma = U\Lambda U' \quad , \quad \text{diag}(\Lambda) > 0$$

Proof: For convenience we start by repeating the argument in Appendix A2. First recall that if the *covariance matrix* of a random vector, $X = (X_1, \dots, X_n)'$, is denoted by $\Sigma = \text{cov}(X)$, then the symmetry of covariances, $\sigma_{ij} = \text{cov}(X_i, X_j) = \text{cov}(X_j, X_i) = \sigma_{ji}$ implies that Σ is symmetric. Moreover, since for any coefficient vector, $a \neq 0$, we must have

$$(A3.4.52) \quad a'\Sigma a = \text{var}(a'X) \geq 0$$

it follows that Σ is positive semidefinite. Hence by SPD Theorem 1 together with Corollary 1 above, it follows that once nonsingularity of Σ implies both positive definiteness and (A3.4.51). \square

In addition, recall from the discussion following the Linear Invariance Theorem in Section A.2.3 that *reduced* covariance matrices of the form, $A\Sigma A'$, were asserted to be nonsingular whenever A is of full row rank. We are now in a position to establish this result:

Corollary 3. *For any nonsingular n -square covariance matrix, Σ , and any $m \times n$ matrix, A , with $1 \leq m \leq n$, if A is of full row rank then $A\Sigma A'$ is also a nonsingular (m -square) covariance matrix.*

Proof: The matrix, $A\Sigma A'$, has already been shown to be an m -square covariance matrix in expression (3.2.21) of Part II of these notes. So it remains to be shown that $A\Sigma A'$ is nonsingular. To do so, recall first (from the end of Section A3.1.1) that A is of full row rank iff its rows are linearly independent. But since these rows are precisely the columns of $A' = (a_1, \dots, a_m)$, it follows from the definition of linear independence [expression (A3.1.24)], that for any $x = (x_1, \dots, x_m) \in \mathbb{R}^m$,

$$(A3.4.53) \quad A'x = 0 \Rightarrow \sum_{i=1}^m x_i a_i = 0 \Rightarrow x_i = 0, i = 1, \dots, m \Rightarrow x = 0$$

Moreover, if Σ is a nonsingular covariance matrix, then by (A3.4.51),

$$(A3.4.54) \quad \Sigma = U \Lambda U' = U \Lambda^{1/2} \Lambda^{1/2} U'$$

where again $\Lambda^{1/2} = \text{diag}(\lambda_1^{1/2}, \dots, \lambda_n^{1/2})$. So by essentially the same argument as in (A3.4.49) and (A3.4.50) it follows that for any $x \in \mathbb{R}^m$,

$$(A3.4.55) \quad \begin{aligned} x' A \Sigma A' x = 0 &\Rightarrow x' A U \Lambda^{1/2} \Lambda^{1/2} U' A' x = 0 \\ &\Rightarrow (\Lambda^{1/2} U' A' x)' \Lambda^{1/2} U' A' x = 0 \\ &\Rightarrow \|\Lambda^{1/2} U' A' x\|^2 = 0 \\ &\Rightarrow \Lambda^{1/2} U' A' x = 0 \\ &\Rightarrow U \Lambda^{1/2} (\Lambda^{1/2} U' A' x) = 0 \\ &\Rightarrow \Sigma A' x = 0 \end{aligned}$$

But this together with (A3.4.53) and the nonsingularity of Σ then shows that

$$(A3.4.56) \quad x' A \Sigma A' x = 0 \Rightarrow \Sigma (A' x) = 0 \Rightarrow A' x = 0 \Rightarrow x = 0$$

Finally, since the covariance matrix, $A \Sigma A'$, is symmetric and positive semidefinite by (A3.4.51), it must then be true that

$$(A3.4.57) \quad x \neq 0 \Rightarrow x' (A \Sigma A') x > 0$$

for all $x \in \mathbb{R}^m$. Thus $A \Sigma A'$ is positive definite, and we may conclude from (A3.4.44) that $A \Sigma A'$ is also nonsingular. \square

A3.4.4 Spectral Decompositions with Distinct Eigenvalues

There is a second class of symmetric matrices for which each SVD directly yields a unique SPD, namely those symmetric matrices for which all eigenvalues are distinct. Here it is of interest to recall the example given in expression (A3.4.4) above, i.e.,

$$(A3.4.58) \quad A = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

with distinct eigenvalues, $\text{Eig}(A) = \{-1, 1\}$, but with (necessarily) repeating singular values given by the absolute values of $\text{Eig}(A)$, so that A has only one distinct singular value, namely $\text{Sing}(A) = \{1\}$. Here the SVD in (A3.4.5) appeared to exhibit little direct

relation to the SPD in (A3.4.8). Indeed, the situation is even worse for this matrix. In particular, since the unit circle is mapped onto itself by A , it follows that every pair of orthogonal unit vectors can serve as principle axes for this “ellipse”. This nonuniqueness can be seen algebraically by observing that since A is itself orthonormal, it follows that for *any* other orthonormal matrix, V , the product, $U = AV$, must also be orthonormal. But since the singular values of A are given by the identity matrix, $S = I_2$, we may then use V to construct a distinct SVD for A by the product:

$$(A3.4.59) \quad A = AVV' = (AV)(I_2)V' = USV'$$

Thus there are seen to be *infinitely* many SVDs for A . On the other hand, since the eigenvalues of A are distinct, we have already seen from (A3.4.27) that their corresponding eigenvectors must be orthogonal, and thus must form a basis for \mathbb{R}^2 . So these eigenvectors [in (A3.4.8)] must in fact be *unique* (up to a choice of signs). Given this stark contrast, it would appear that there is little hope of constructing the unique SPD for A from its highly nonunique SVDs. But as we now show, this can indeed be done so long as eigenvalues are *distinct* in the sense that each has a geometric multiplicity of one [as in the case of (A3.4.58)]. Note also from the orthogonality of eigenvectors for distinct eigenvalues in (A3.4.27) that this in turn implies that the SPD for such symmetric matrices must be *unique*. With these observations, we now show that:

Spectral Decomposition of Symmetric Matrices with Distinct Eigenvalues

If A is a symmetric matrix with distinct eigenvalues, then each SVD,

$$(A3.4.60) \quad A = USV' = (u_1, \dots, u_n) \begin{pmatrix} s_1 & & \\ & \ddots & \\ & & s_n \end{pmatrix} \begin{pmatrix} v_1' \\ \vdots \\ v_n' \end{pmatrix},$$

of A yields exactly the same SPD,

$$(A3.4.61) \quad A = W\Lambda W' = (w_1, \dots, w_n) \begin{pmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{pmatrix} \begin{pmatrix} w_1' \\ \vdots \\ w_n' \end{pmatrix}$$

in terms of the relations in (A3.4.22) and (A3.4.23).

Proof: Our objective is to give an explicit construction of (A3.4.61) in terms of (A3.4.60). To do so, we note first from the assumed distinctness of eigenvalues that there can be at most one zero eigenvalue for A , say $\lambda_i = 0$ with eigenvector satisfying,

$Aw_i = 0$. But since this implies that A is singular there must at least one zero singular value [for otherwise, we would have $|A|_+ > 0$ by (A3.2.78), which contradicts the singularity of A]. Moreover, if there were more than one, the argument in (A3.4.27) shows that there would be more than one zero eigenvalue for A , which contradicts the distinctness assumption. So there is exactly one u_i with $Au_i = 0$ as in (A3.4.37). Thus we may then set $w_i = u_i$, and conclude from (A3.4.27) that this will always form an admissible entry in the orthonormal matrix, W , of (A3.4.61). For the positive singular values, $s_i \in \text{diag}(S)$, we now consider the possible distinct eigenvalues they can generate by (A3.4.22) and (A3.4.23). In view of the distinctness assumption, either exactly one of the values $(s_i, -s_i)$ belongs to $\text{Eig}(A)$ or both do. The first case is the simplest, and is equivalent to the condition derived from (A3.4.60) that s_i only appear in one equation of the equation systems (A3.4.22) and (A3.4.23) with a nonzero eigenvector. If for notational simplicity we let (λ_i, w_i) denote the associated eigenvalue-eigenvector pair to be constructed in (A3.4.61)⁷ then by using the definitions of $x_i = u_i + v_i$ and $y_i = u_i - v_i$ in (A3.4.22) and (A3.4.23), respectively, (and recalling that at least one of these vectors must be nonzero) we may set

$$(A3.4.62) \quad w_i = \begin{cases} \frac{u_i + v_i}{\|u_i + v_i\|} & , \text{ if } u_i + v_i \neq 0 \\ \frac{u_i - v_i}{\|u_i - v_i\|} & , \text{ if } u_i + v_i = 0 \end{cases}$$

and similarly, set

$$(A3.4.63) \quad \lambda_i = \begin{cases} s_i & , \text{ if } u_i + v_i \neq 0 \\ -s_i & , \text{ if } u_i + v_i = 0 \end{cases}$$

Again the orthogonality of eigenvectors for distinct eigenvalues in (A3.4.27) guarantees that these normalized vectors are automatically admissible components of W .

Turning to the second case, where both $(s_i, -s_i)$ appear in equation systems (A3.4.22) and (A3.4.23) with nonzero eigenvectors, observe that s_i must appear twice in $\text{diag}(S)$, say in positions i and j . If we consider the values of x and y in columns i and j of both (A3.4.22) and (A3.4.23), namely,

⁷ More formally, the rows and columns of A can always be permuted to satisfy this relation. The standard convention in the literature is thus to say that “by relabeling if necessary” we can use i for both s_i and its associated eigenvalue, λ_i .

$$(A3.4.64) \quad \begin{aligned} x_i &= u_i + v_i & y_i &= u_i - v_i \\ x_j &= u_j + v_j & y_j &= u_j - v_j \end{aligned}$$

then it must be true that either x_i or x_j is nonzero, and similarly that either y_i or y_j are nonzero. But if *both* x_i and x_j are nonzero, then they must be scalar multiples of one another. For otherwise, columns i and j of equation system (A3.4.22) would yield two linearly independent solutions, $(Ax_i = s_i x_i, Ax_j = s_j x_j)$ with $s_i = s_j$, and it would follow that eigenvalue, s_i , has a multiplicity of two. But since this contradicts the assumption of distinct eigenvalues, x_i and x_j must be linearly dependent, i.e., scalar multiples of one another. This in turn implies that they must have the same normalizations, which can be written terms of u and v by:

$$(A3.4.65) \quad \frac{u_i + v_i}{\|u_i + v_i\|} = \frac{u_j + v_j}{\|u_j + v_j\|}$$

Moreover, since exactly the same argument for y_i and y_j shows that if both are nonzero then

$$(A3.4.66) \quad \frac{u_i - v_i}{\|u_i - v_i\|} = \frac{u_j - v_j}{\|u_j - v_j\|}$$

With these observations, if we now denote the eigenvalues for s_i and $s_j (= -s_i)$ in (A3.4.59) by λ_i and λ_j (again using the convention in footnote 6), then by definition, $\lambda_i = s_i > 0$ and $\lambda_j = -s_i < 0$, with associated eigenvectors given respectively by

$$(A3.4.67) \quad w_i = \begin{cases} \frac{u_i + v_i}{\|u_i + v_i\|} & , \text{ if } u_i + v_i \neq 0 \\ \frac{u_j + v_j}{\|u_j + v_j\|} & , \text{ if } u_i + v_i = 0 \end{cases}$$

$$(A3.4.68) \quad w_j = \begin{cases} \frac{u_j - v_j}{\|u_j - v_j\|} & , \text{ if } u_j - v_j \neq 0 \\ \frac{u_i - v_i}{\|u_i - v_i\|} & , \text{ if } u_j - v_j = 0 \end{cases}$$

Again, it follows from (A3.4.65) and (A3.4.66) that these choices of w_i and w_j are insensitive to whether the i^{th} or j^{th} quantities are used first on the right hand sides of

(A3.4.65) and (A3.4.66). Note also from the orthogonality of eigenvectors for distinct eigenvalues that these normalized vectors will always yield admissible components of W .

Finally, since the multiplicity of each singular value, $s \in \text{diag}(S)$ determines exactly the number of eigenvalues generated by s (including the $s = 0$ case), it follows that this procedure must generate precisely n eigenvalues $(\lambda_1, \dots, \lambda_n)$ with corresponding orthonormal eigenvectors (w_1, \dots, w_n) generating a basis for \mathbb{R}^n . So by construction, this procedure must yield a complete representation of A as in (A3.4.61). \square

So for the case of distinct eigenvalues, we see that the unique SPD for symmetric matrix A can be explicitly constructed from any of its possible SVDs. Here it is instructive to see how this procedure works for the example in (A3.4.58). In this case, the SVD produced in (A3.4.5) yields

$$(A3.4.69) \quad X = U + V = (u_1 + v_1, u_2 + v_2) = \left[\begin{pmatrix} -1 \\ -1 \end{pmatrix} \begin{pmatrix} -1 \\ -1 \end{pmatrix} \right]$$

$$(A3.4.70) \quad Y = U - V = (u_1 - v_1, u_2 - v_2) = \left[\begin{pmatrix} 1 \\ -1 \end{pmatrix} \begin{pmatrix} 1 \\ -1 \end{pmatrix} \right]$$

So this is case where all four elements of (A3.4.64) are nonzero, and thus where the identities in (A3.4.65) and (A3.4.66) are seen to hold. Moreover, since the norms of all these vectors are seen to be $\sqrt{2}$, it follows that they yield precisely the pair of normalized eigenvectors in (A3.4.8). Moreover, one can verify by direct computation that any choice of an orthonormal matrix, V , in (A3.4.59) will always produce vectors that are scalar multiples of those in (A3.4.69) and (A3.4.70), and thus will yield the same eigenvectors for W .

Finally it is important to note that the case of distinct eigenvalues is overwhelmingly the most common case observed in practice. Indeed, it is a simple matter to show that within the space of all n -square symmetric matrices, the subset possessing two or more common eigenvalues must have *zero volume*. So if one were to choose a symmetric matrix at random, then *with probability one*, this matrix will have all distinct eigenvalues.

A3.4.5 General Spectral Decomposition Theorem

Nonetheless, it is clear that in numerous modeling contexts, theoretical considerations can often lead to symmetric matrices with additional structure yielding repeated eigenvalues. In particular, for singular matrices it is clear that zero eigenvalue may be repeated many times. So it is of practical interest to show that the information contained in SVDs for such matrices can still be used to construct their SPDs. The major difference in this general case with repeated eigenvalues is that the SPD itself is *not unique*. So there will necessarily be some degree of nonuniqueness in the SVD construction of SPDs.

Here we begin with one preliminary result that will enable us to verify dimensional consistency for all subspaces of repeated eigenvalues. In particular, note that for any pair of $n \times k$ matrices, A and B , the matrix *sum*, $A + B$, is well defined. In addition, A and B are said to be *mutually orthogonal* iff their columns are orthogonal, i.e., $A'B = O_k$.⁸ Hence, recalling that the *rank* of a matrix is by definition the dimension of its span [$\text{rank}(A) = \dim(\text{span}(A))$], we have the following useful rank equality:⁹

Rank Lemma. For any mutually orthogonal $n \times k$ matrices, A and B ,

$$(A3.4.71) \quad \text{rank}(A + B) = \text{rank}(A) + \text{rank}(B)$$

Proof: By definition it suffices to show that

$$(A3.4.72) \quad \dim(\text{span}(A + B)) = \dim(\text{span}(A)) + \dim(\text{span}(B))$$

But if we choose any bases $[x_1, \dots, x_k]$ and $[y_1, \dots, y_h]$ for $\text{span}(A)$ and $\text{span}(B)$, respectively, then by mutual orthogonality it follows that $[x_1, \dots, x_k, y_1, \dots, y_h]$ must constitute a *linearly independent set*. To see this, note that since $x_i \in \text{span}(A) \Rightarrow x_i = Az$ for some $z \in \mathbb{R}^k$, and similarly that $y_j \in \text{span}(B) \Rightarrow y_j = Bw$ for some $w \in \mathbb{R}^k$, this together with the mutual orthogonality condition, $A'B = O_k$, implies that

$$(A3.4.73) \quad x_i' y_j = (Az)'(Bw) = z'(A'B)w = 0$$

and hence that $[x_1, \dots, x_k]$ and $[y_1, \dots, y_h]$ are mutually orthogonal sets of vectors. This together with the linear independence of basis vectors implied that the full set of vectors, $[x_1, \dots, x_k, y_1, \dots, y_h]$, is linearly independent. Finally since for any vector, $v \in \mathbb{R}^n$,

$$(A3.4.74) \quad v \in \text{span}(A + B) \Rightarrow v = (A + B)u \quad \text{for some } u \in \mathbb{R}^k$$

$$\Rightarrow v = Au + Bu \in \text{span}(A) + \text{span}(B)$$

$$\Rightarrow v = \sum_{i=1}^k \alpha_i x_i + \sum_{j=1}^h \beta_j y_j$$

for some coefficients $(\alpha_1, \dots, \alpha_k)$ and $(\beta_1, \dots, \beta_h)$, it then follows that $[x_1, \dots, x_k, y_1, \dots, y_h]$ must be a basis for $\text{span}(A + B)$. Thus we may have

$$(A3.4.75) \quad \dim(\text{span}(A + B)) = k + h = \dim(\text{span}(A)) + \dim(\text{span}(B))$$

⁸ As with the n -square identity matrix, I_n , we here denote the n -square zero matrix by O_n .

⁹ A detailed development of other rank properties can be found in Chapter 6 of Searle (1982).

and may conclude that condition (A3.4.72) holds. \square

With this preliminary result, we are now ready to establish the following general form of the Spectral Decomposition Theorem:

Spectral Decomposition Theorem: For any symmetric matrix, A , there exists a diagonal matrix, Λ , and an orthonormal matrix, W , such that

$$(A3.4.76) \quad A = W \Lambda W'$$

Proof: Our approach is again to start with any SVD,

$$(A3.4.77) \quad A = U S V'$$

for the n -square symmetric matrix, A , and to construct an SPD for A as in (A3.4.76). To do so, we first note that by relabeling the rows and columns of A if necessary, we may assume that the sets of common singular values (including singleton sets) are grouped into blocks, $S_i = \text{diag}(s_{i1}, \dots, s_{in_i})$, $i = 1, \dots, m (\leq n)$, along the diagonal of matrix, S , where each block has common value, $s_i \equiv s_{ij}$, $j = 1, \dots, n_i (\geq 1)$, and has associated orthonormal sets of column vectors, $U_i = (u_{i1}, \dots, u_{in_i})$ and $V_i = (v_{i1}, \dots, v_{in_i})$. With this grouping, expression (A3.4.77) can be written as

$$(A3.4.78) \quad A = (U_1, \dots, U_m) \begin{pmatrix} S_1 & & \\ & \ddots & \\ & & S_m \end{pmatrix} \begin{pmatrix} V_1' \\ \vdots \\ V_m' \end{pmatrix} \Rightarrow AV_i = U_i S_i, \quad i = 1, \dots, m$$

Our objective is then to construct an SPD of A with corresponding block form:

$$(A3.4.78) \quad A = W \Lambda W' = (W_1, \dots, W_m) \begin{pmatrix} \Lambda_1 & & \\ & \ddots & \\ & & \Lambda_m \end{pmatrix} \begin{pmatrix} W_1' \\ \vdots \\ W_m' \end{pmatrix} \Rightarrow AW_i = W_i \Lambda_i, \quad i = 1, \dots, m$$

As with U and V above, the key conditions to be satisfied by W are that each block, W_i , have orthonormal columns, and that the columns in different blocks be mutually orthogonal. To construct (A3.4.79), we start by observing there is one special case that can be handled without further analysis. In particular, if matrix A is *singular*, then exactly one block, S_i , in (A3.4.78), will have $s_i = 0$. But since the vectors in this block must satisfy,

$$(A3.4.80) \quad AV_i = O_{n_i} \Rightarrow Av_{ij} = 0 = 0v_{ij}, \quad j = 1, \dots, n_i$$

it follows [as in (A3.4.37) above] that these are automatically eigenvectors for $s_i = 0$, and in addition, must be *orthonormal* by the properties of SVDs. Moreover, since all other eigenvalues of A must be nonzero, it also follows from (A3.4.27) that the columns of V_i will automatically be *orthogonal* to all other eigenvectors of A . Thus by setting

$$(A3.4.81) \quad A_i = S_i = O_{n_i} \quad \text{and} \quad W_i = V_i,$$

we are guaranteed to obtain an admissible “zero” block in (A3.4.76). So we may henceforth assume (from the nonnegativity of singular values) that $s_i > 0$. For these cases, recall from (A3.4.18) through (A3.4.21) that if we set

$$(A3.4.82) \quad X_i = U_i + V_i, \quad i = 1, \dots, m$$

$$(A3.4.83) \quad Y_i = U_i - V_i, \quad i = 1, \dots, m$$

that by construction,

$$(A3.4.84) \quad AX_i = s_i X_i, \quad i = 1, \dots, m$$

$$(A3.4.85) \quad AY_i = -s_i Y_i, \quad i = 1, \dots, m$$

So all columns in X_i and Y_i are potential eigenvectors for A . Here there three possible cases to be considered, namely (i) $U_i = V_i \Rightarrow Y_i = O$, (ii) $U_i = -V_i \Rightarrow X_i = O$, or (iii) $X_i \neq O$ and $Y_i \neq O$.¹⁰ In case (i), all eigenvalues in block i are positive, and governed by (A3.4.83). But since (A3.4.82) and (A3.4.83) together imply that $X_i + Y_i = 2U_i$, it follows that

$$(A3.4.86) \quad X_i = X_i + O = X_i + Y_i = 2U_i$$

Thus, except for a factor of 2, the columns of X_i automatically form a set of n_i admissible eigenvectors for block i . So here we may set

$$(A3.4.87) \quad A_i = S_i = s_i I_{n_i} \quad \text{and} \quad W_i = \frac{1}{2} X_i = U_i$$

to obtain the desired block i in (A3.4.79). The construction for case (ii) is essentially identical except that now

¹⁰ For notational simplicity, we here take the common dimension of these zero matrices (namely $n \times n_i$) to be understood.

$$(A3.4.88) \quad Y_i = O + Y_i = X_i + Y_i = 2U_i$$

with all eigenvalues given by $-s_i$. So in this case, we can construct the desired block i by setting

$$(A3.4.89) \quad A_i = -S_i = (-s_i)I_{n_i} \quad \text{and} \quad W_i = \frac{1}{2}Y_i = U_i$$

This leaves case (iii), in which both s_i and $-s_i$ are eigenvalues for A .

This is by far the most complex case, and requires additional analysis. Here we start by observing from the distinctness of the eigenvalues, s_i and $-s_i$, together with (A3.4.84), (A3.4.85) and the orthogonality condition (A3.4.27), that the matrices X_i and Y_i must now be mutually orthogonal (i.e., $X_i'Y_i = O_{n_i}$). So by the Rank Lemma above, we must have

$$(A3.4.90) \quad \text{rank}(X_i + Y_i) = \text{rank}(X_i) + \text{rank}(Y_i)$$

Moreover, since it continues to be true that $X_i + Y_i = 2U_i$, it then follows that

$$(A3.4.91) \quad \begin{aligned} \dim(\text{span}(X_i)) + \dim(\text{span}(Y_i)) &= \text{rank}(X_i) + \text{rank}(Y_i) \\ &= \text{rank}(X_i + Y_i) = \text{rank}(2U_i) = n_i \end{aligned}$$

So if we now use the *Gram-Schmidt orthogonalization procedure* [summarized by expression (A3.1.60) above] to construct orthogonal bases $[b_1, \dots, b_{k_i}]$ and $[c_1, \dots, c_{h_i}]$ for $\text{span}(X_i)$ and $\text{span}(Y_i)$, respectively, then these basis vectors will constitute the desired eigenvectors for this case. To verify this, observe first from (A3.4.91) that

$$(A3.4.92) \quad k_i + h_i = n_i$$

and hence that there are again exactly n_i of these basis vectors. Moreover, since $b_j \in \text{span}(X_i) \Rightarrow b_j = X_i z_j$ for some $z_j \in \mathbb{R}^{n_i}$, it follows from (A3.4.84) that

$$(A3.4.93) \quad A X_i = s_i X_i \Rightarrow A X_i z_j = s_i X_i z_j \Rightarrow A b_j = s_i b_j, \quad j = 1, \dots, k_i$$

and thus that the basis vectors $[b_1, \dots, b_{k_i}]$ form an orthogonal set of eigenvectors for eigenvalue, s_i . Similarly, since $c_j \in \text{span}(Y_i) \Rightarrow c_j = Y_i u_j$ for some $u_j \in \mathbb{R}^{n_i}$, it follows from (A3.4.85) that

$$(A3.4.94) \quad AY_i = (-s_i)Y_i \Rightarrow AY_i u_j = (-s_i)X_i u_j \Rightarrow Ac_j = (-s_i)c_j, \quad j = 1, \dots, h_i$$

and thus that the basis vectors $[c_1, \dots, c_{h_i}]$ form an orthogonal set of eigenvectors for eigenvalue, $-s_i$. Finally, since the distinctness of s_i and $-s_i$ again implies that $[b_1, \dots, b_{k_i}]$ and $[c_1, \dots, c_{h_i}]$ are mutually orthogonal, and also orthogonal to the eigenvectors for all other distinct eigenvalues of A , we may conclude that the normalizations of these eigenvectors yield an admissible choice for W_i . Hence, if we let $1'_m = (1, \dots, 1)$ denote the unit vector of length m , then an admissible choice for block i in (A3.4.79) is now given by:

$$(A3.4.95) \quad W_i = \left(\frac{b_1}{\|b_1\|}, \dots, \frac{b_{k_i}}{\|b_{k_i}\|}, \frac{c_1}{\|c_1\|}, \dots, \frac{c_{h_i}}{\|c_{h_i}\|} \right) \quad \text{and} \quad A_i = \text{diag}(s_i 1'_{k_i}, -s_i 1'_{h_i})$$

By way of summary, the blocks defined respectively by (A3.4.81), (A3.4.87), (A3.4.89) and (A3.4.95) yield a full specification of expression (A3.4.79), and thus the desired SPD for A is established. \square

As one final comment, we begin by reiterating that our main objective in this section has been to show that the spectral decomposition (SPD) of any symmetric matrix, A , can be constructed from its singular value decomposition (SVD). However, this appears to leave open the converse question of how to construct SVDs of symmetric matrices from their SPDs. But since the singular values of A are simply the absolute values of its eigenvalues, it turns out to be a simple matter to transform each SPD into a corresponding SVD. To do so, recall from (A3.1.10) that the SPD in (A3.4.76) can be rewritten as:

$$(A3.4.96) \quad A = W \Lambda W' = (w_1, \dots, w_n) \begin{pmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{pmatrix} \begin{pmatrix} w'_1 \\ \vdots \\ w'_n \end{pmatrix} = \sum_{i=1}^n \lambda_i w_i w'_i$$

To convert these eigenvalues to absolute form, observe that if $\text{sgn}(\lambda)$ denotes the *sign* of any number, λ , then by definition, $\lambda = |\lambda| \text{sgn}(\lambda)$, so that (A3.4.96) can be written as,

$$(A3.4.97) \quad A = \sum_{i=1}^n |\lambda_i| \text{sgn}(\lambda_i) w_i w'_i$$

But if we define, $U = (u_1, \dots, u_n)$, $S = \text{diag}(s_1, \dots, s_n)$, and $V = (v_1, \dots, v_n)$ by

$$(A3.4.98) \quad \begin{aligned} u_i &= \text{sgn}(\lambda_i) w_i, & i &= 1, \dots, n \\ s_i &= |\lambda_i|, & i &= 1, \dots, n \\ v_i &= w_i, & i &= 1, \dots, n \end{aligned}$$

then by definition,

$$(A3.4.99) \quad A = \sum_{i=1}^n s_i u_i v_i' = (u_1, \dots, u_n) \begin{pmatrix} s_1 & & \\ & \ddots & \\ & & s_n \end{pmatrix} \begin{pmatrix} v_1' \\ \vdots \\ v_n' \end{pmatrix} = USV'$$

where S is nonnegative diagonal matrix and where $V (= W)$ is orthonormal. Moreover, since

$$(A3.4.100) \quad \begin{aligned} u_i' u_i &= \text{sgn}(\lambda_i)^2 w_i' w_i = 1 & , i = 1, \dots, n \\ u_i' u_j &= \text{sgn}(\lambda_i) \text{sgn}(\lambda_j) w_i' w_j = 0 & , i = 1, \dots, n \end{aligned}$$

it also follows that U is orthonormal, and thus that (A3.4.99) is automatically an SVD for A . However, the SPDs of symmetric matrices clearly contain more information, and turn out to be far more useful than their corresponding SVDs. So this final result only serves to complete the full correspondence between the two.

A3.5 Nonnegative Matrices

Recall that every spatial weights matrix, $W = (w_{ij} : i, j = 1, \dots, n)$, is by definition a *nonnegative* matrix. In this section we develop some of the key properties of such matrices that will prove to be useful in our subsequent analyses. To do so, we begin with a consideration of additional matrix properties which are particularly appropriate for spatial weights matrices.

A3.5.1 Strongly Connected Matrices

Consider the following spatial weights matrix,

$$(A3.5.1) \quad W = \begin{pmatrix} 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \end{pmatrix}$$

that represents (queen) *contiguity* relations among the five regions depicted in Figure A3.30 below, where $w_{ij} = 1$ if and only if regions i and j share a common boundary:




Figure A3.30. Disconnected Regions




Figure A3.31. Relabeling of Regions

It is not obvious from matrix, W , that these regions actually form two *disconnected* regional subsystems, $\{R_1, R_3, R_4\}$ and $\{R_2, R_5\}$, which share no common boundaries. However, this relation can be seen more clearly by *relabeling* these regions as shown in Figure A3.31, so that W now takes the form:

$$(A3.5.2) \quad W = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \end{pmatrix} = \begin{pmatrix} W_1 & \\ & W_2 \end{pmatrix}$$

With this relabeling, W is seen to be a *block-diagonal* matrix, where each block, $W_i, i = 1, 2$, corresponds to a single connected subsystem of regions. Note that for each $k > 0$ we must also have,

$$(A3.5.3) \quad W^k = \begin{pmatrix} W_1^k & \\ & W_2^k \end{pmatrix}$$

Moreover, recall from Section 3.3.3 that these k^{th} powers of W effectively count the *influence paths* of length k between each pair of regions. So the block-diagonal form of (A3.5.3) inherited from (A3.5.2) implies that there can be *no influence paths* (of any length) between these two subsystems of regions. In other words, these subsystems are completely *disconnected* from one another in terms of possible spatial influences. More specifically, with respect to the types of “ripple effects” described in Section 3.3.3, this block-diagonal property implies that no effects originating in one subsystem can ever influence regions of the other subsystem. Thus in terms of spatial dependencies, these subsystems can effectively be analyzed separately from one another.

While this seems obvious for such physically disjoint subsystems, there are in fact more subtle types of spatial disconnectedness that must also be considered. As one example, suppose we consider *centroid distances* between members of the 5-region system in Figure A3.5.4 below:




Figure A3.32. Centroid Distance

As an instance of *multiple nearest-neighbor relations*, we now consider only the first two nearest neighbors for each region within this small system of regions. Such two-neighbor relations are represented by the spatial weights matrix in expression (A3.5.4) below, where for example, the first two nearest neighbors of R_1 are seen from Figure A3.32 to be R_2 and R_3 .

$$(A3.5.4) \quad W = \begin{pmatrix} 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 1 & 0 \end{pmatrix} = \begin{pmatrix} W_{11} & O_{3 \times 2} \\ W_{21} & W_{22} \end{pmatrix}$$

Here it is clear that the subsystem of regions $\{R_1, R_2, R_3\}$ (represented by submatrix, W_{11}) is self-contained in the sense that all nearest neighbors are in the same subsystem. But while the complementary subsystem, $\{R_4, R_5\}$ (represented by W_{22}) also contains all of its first nearest neighbors, this is not true of second nearest neighbors, which are seen to be R_2 and R_3 . Thus there is an *asymmetry* between these two subsystems, with W_{12} containing all zeros while W_{21} is nonzero. Moreover, this (lower) *block triangularity* property of W is again seen to be preserved by all powers of W , as illustrated by W^2 ,

$$(A3.5.5) \quad W^2 = \begin{pmatrix} W_{11} & O_{3 \times 2} \\ W_{21} & W_{22} \end{pmatrix}^2 = \begin{pmatrix} W_{11}^2 & O_{3 \times 2} \\ W_{21}W_{11} + W_{22}W_{21} & W_{22}^2 \end{pmatrix}$$

which continues to be (lower) block triangular. So in terms of spatial influence paths, this block triangular structure implies that there are *no influence paths* from subsystem $\{R_4, R_5\}$ to $\{R_1, R_2, R_3\}$. In other words, no spatial effects originating in $\{R_4, R_5\}$ can ever influence $\{R_1, R_2, R_3\}$. Thus while these two subsystems are not fully disconnected with respect to such (second-nearest-neighbor) influences, there is clearly some lack of connectivity here.

With these observations, we now focus on spatial weights matrices, $W = (w_{ij} : i, j = 1, \dots, n)$, that are *strongly connected* in the sense that for every pair of distinct regions, i and j , there exists *at least one influence path* from j to i , i.e., at least one sequence coefficients $(w_{i,m_1}, w_{m_1,m_2}, \dots, w_{m_{k-1},m_k}, w_{m_k,j})$ of positive coefficients in W . For our later purposes, it is important to note that it is enough to focus on influence paths of length $k \leq n - 1$. For since each longer path must repeat some index, m_i , at least once, we can always remove the path cycles created by such repetitions. Under this condition, the underlying system of regions, $i = 1, \dots, n$, is then said to be *strongly connected with respect to W* . Some consequences of failures of this condition will be discussed further in Section ?? below. But for the moment we consider the analytical of strong connectivity.

A3.5.2 Perron-Frobenius Theorem

The single most important properties of such strongly connected weights matrices relate to their *eigenstructure*. In particular, their *maximum* eigenvalue is always a *positive real* number with a *unique positive* eigenvector of unit length. To establish this fundamental result, known as the *Perron-Frobenius (P-F) Theorem* it is convenient to focus on general *n-square matrices*, A , and to state all results in this context.

To do so, we begin with some instructive examples. First recall from expression (3.3.13) in Section 3 that the matrix

$$(A3.5.6) \quad A = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$$

is a pure (90°) rotation of the plane, and thus has *no real* eigenvalues. Thus one consequence of the P-F Theorem is that nonnegative matrices, A , are guaranteed to have at least one *real* eigenvalue. To illustrate this in the context of spatial weights matrices, recall from expression (3.3.12) that one of the simplest possible examples is provided by the contiguity matrix for two (connected) regions

$$(A3.5.7) \quad W = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

Here the Perron-Frobenius properties above are almost immediate from the *row normalized* structure of W , which by definition ensures that $W1_2 = 1_2$, and thus that W has a real positive eigenvalue, $\lambda_1 = 1$, with positive eigenvector, $x_1 = 1_2$. The full eigenstructure of W is given in Figure 3.4, and shows that the normalization, $x_1 = (1/\sqrt{2})1_2$, yields the unique positive eigenvector with unit length. A somewhat more interesting example is provided by the contiguity matrix,

$$(A3.5.8) \quad W = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}$$

for the regional subsystem $\{R_1, R_2, R_3\}$ in Figure A3.31 above. Here it can be verified by direct multiplication that [as in expression (3.3.14)] the eigenstructure,

$$(A3.5.9) \quad WX = X\Lambda$$

of W is given by

$$(A3.5.10) \quad \Lambda = \begin{pmatrix} \sqrt{2} & & \\ & 0 & \\ & & -\sqrt{2} \end{pmatrix} \quad X = \begin{pmatrix} \frac{1}{2} & -\frac{1}{\sqrt{2}} & \frac{1}{2} \\ \frac{1}{\sqrt{2}} & 0 & -\frac{1}{\sqrt{2}} \\ \frac{1}{2} & \frac{1}{\sqrt{2}} & \frac{1}{2} \end{pmatrix}$$

So here the maximum positive eigenvalue is $\lambda_1 = \sqrt{2}$ with unique positive eigenvector, $x_1 = (\frac{1}{2}, \frac{1}{\sqrt{2}}, \frac{1}{2})$, of unit length. Note also that while this matrix is strongly connected, it is also a *singular* matrix (which can be seen directly from the fact that the first and last columns of W are the same). So strongly connectivity says nothing about nonsingularity of matrices.

Finally, just to illustrate that these properties have nothing to do with the zero-one nature of contiguity matrices, it is of interest to consider *centroid distance* between these three

regions, as shown in Figure A3.32 below (where by construction these centroids form a “3-4-5” triangle”):




Figure A3.32. Centroid Distance

Here the corresponding spatial weights matrix is given by:

$$(A3.5.11) \quad W = \begin{pmatrix} 0 & 3 & 5 \\ 3 & 0 & 4 \\ 5 & 4 & 0 \end{pmatrix}$$

with eigenstructure, $WX = XD$, that can be calculated numerically in MATLAB as

$$(A3.5.12) \quad D = \begin{pmatrix} 8.0558 & & \\ & -2.8755 & \\ & & -5.1803 \end{pmatrix}, \quad X = \begin{pmatrix} 0.58133 & 0.55396 & -0.59598 \\ 0.52511 & -0.81493 & -0.24527 \\ 0.62155 & 0.17037 & -0.76463 \end{pmatrix}$$

by using the command:

```
>> [X,D] = eig(W);
```

Here we see again that there is exactly one real positive eigenvalue, **8.0558**, with corresponding positive eigenvector, **[0.58133, 0.52511, 0.52511]'**.

With these examples of nonnegative strongly connected matrices, our objective is to establish the Perron-Frobenius Theorem for such matrices, which for our purposes can be stated as follows:

Perron-Frobenius Theorem. *For any strongly connected nonnegative matrix, A , the eigenvalue of A with maximum absolute value is positive and has a unique positive eigenvector of unit length.*

The only technicality here is that the *absolute value*, $|\lambda|$, of an eigenvalue, $\lambda \in \text{Eig}(A)$, is taken to include *complex* eigenvalues as well, where $|\lambda|$ is usually called the *modulus* of λ and is always a *real positive* number. Since the arguments below will only require the use of such absolute values, there is little harm in assuming for the present that all eigenvalues, $\lambda \in \text{Eig}(A)$, are real. The interpretation of complex eigenvalues is deferred to Section A3.5.4 below.

There are several approaches to proving this theorem. We start with the two-dimensional case, where a simple “angle” approach can be applied that offers some intuition as to why the P-F theorem is true. We then develop a proof for the general case.

Proof Sketch for Two Dimensions

Consider the linear transformation in Figure A3.33 below, with strongly connected nonnegative matrix representation, A .




Figure A3.33. Angle Approach




Figure A3.34. Zero Point

By nonnegativity, the images, Ae_1 and Ae_2 , of the basis vectors, e_1 and e_2 , must be in the nonnegative quadrant. Moreover, by strong connectivity, neither e_1 and e_2 can be an eigenvector. For example, if e_2 were an eigenvector, then we would have

$$(A3.5.13) \quad Ae_2 = \lambda e_2 \Rightarrow \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix} \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ \lambda \end{pmatrix} \Rightarrow A = \begin{pmatrix} a_{11} & 0 \\ a_{21} & \lambda \end{pmatrix}$$

which in turn implies that there can be no path from node $j = 2$ to node $i = 1$. A similar argument holds for e_1 . So in Figure A3.33 it follows that Ae_2 must be to the right of e_2 , yielding a *positive* (clockwise) angle, $\theta_2 = \theta(e_2, Ae_2)$, between these vectors. Similarly, Ae_1 must be above e_1 , yielding a *negative* (counter-clockwise) angle, $\theta_1 = \theta(e_1, Ae_1)$, as in the figure. So one can see already that there “should” be some vector, x_0 , between e_1 and e_2 for which this angle is zero, and thus for which x_0 is a scalar multiple of Ax_0 , i.e.,

an eigenvector of A (as shown in the figure). To verify this, we first connect e_1 and e_2 by the line segment shown in blue (usually called the *unit simplex* in \mathbb{R}^2), and express all points on this line as linear (convex) combinations of e_1 and e_2 :

$$(A3.5.14) \quad x(\alpha) = \alpha e_2 + (1 - \alpha)e_1, \quad \alpha \in [0,1]$$

Next, we define the corresponding angles for all these points as follows:

$$(A3.5.15) \quad \theta(\alpha) = \theta[x(\alpha), Ax(\alpha)], \quad \alpha \in [0,1]$$

where in particular, $\theta(0) = \theta(e_1, Ae_1) = \theta_1 < 0$, and $\theta(1) = \theta(e_2, Ae_2) = \theta_2 > 0$. Then it is easy to see that (A3.5.15) must yield a *continuous* function, $\theta(\alpha)$, on the unit interval, as shown schematically in Figure A3.34. But since $\theta(0) < 0$ and $\theta(1) > 0$, it then follows from the *Intermediate Value Theorem*¹ that there must be some $\alpha_0 \in (0,1)$ for which $\theta(\alpha_0) = 0$. Thus if we now let $x_0 = x(\alpha_0)$, and set $\lambda_0 = \|Ax_0\|/\|x_0\| > 0$, then by definition this implies that:

$$(A3.5.16) \quad 0 = \theta(x_0, Ax_0) \Rightarrow \frac{Ax_0}{\|Ax_0\|} = \frac{x_0}{\|x_0\|} \Rightarrow Ax_0 = \lambda_0 x_0$$

and thus that x_0 is a positive eigenvector of A with positive eigenvalue, λ_0 . So all that remains to be shown is that λ_0 is the largest eigenvalue for A . Again, this can be seen intuitively for \mathbb{R}^2 by simply observing that the image, $A(C) = \{Az : z \in C\}$ of the unit circle, $C = \{z : \|z\| = 1\}$, under any transformation such as A in Figure A3.33 must be “elongated” in the direction of this eigenvector, as seen schematically in Figure A3.35 below (where x_0 has been rescaled to $z_0 = x_0/\|x_0\| \in C$),




Figure A3.35. Direction of Elongation

¹ See for example Theorem 16.4 in Bartle (1975).

Since this two-dimensional argument is meant only to provide some intuition for the P-F Theorem, we shall not attempt a more detailed proof of this “elongation” property. Rather, we turn now to a general proof, based on the version given in Chapter 9 of Sternberg (2010)². Before doing so however, it should be mentioned that the argument above for \mathbb{R}^2 can be directly generalized by replacing the simple “angle” argument with a more general fixed-point argument on the unit simplex in \mathbb{R}^n [as shown by Debreu and Herstein (1953)].³ However, the proof below requires only the existence of maxima, as we have already seen in the discussion of Figure A3.19 above. In addition, it provides some degree of geometric intuition that can be displayed graphically.

This proof makes extensive use of the follow characterization of *strong connectivity* which is useful in its own right. In particular, a nonnegative n -square matrix, A , is *strongly connected* if and only if the associated matrix,

$$(A3.5.17) \quad P = (I_n + A)^{n-1}$$

is *positive* (i.e., has all positive elements). First recall from the *Binomial Theorem* that for any number, a ,

$$(A3.5.18) \quad (1+a)^m = 1 + \binom{m}{1}a + \binom{m}{2}a^2 + \cdots + \binom{m}{m}a^m = 1 + \sum_{k=1}^m \binom{m}{k}a^k$$

Exactly the same result applies to matrices, and shows in particular that

$$(A3.5.19) \quad P = (I_n + A)^{n-1} = I_n + \sum_{k=1}^{n-1} \binom{n-1}{k} A^k$$

Since A is nonnegative and since both the diagonal elements of I_n and combinatorial coefficients on the right hand side are positive, it follows that P is positive if and only if for each distinct ij -pair, at least one of the matrices, A^k , $k = 1, \dots, n-1$, has a positive component, $(A^k)_{ij}$. But since $(A^k)_{ij}$ is the number of influence paths of length k from j to i , this says that P is positive if and only if there is at least one influence path of length $k \leq n-1$ for every distinct ij -pair – which is precisely *strong connectedness*.

This result also provides a simple method for verifying strong connectedness. In fact, by simply taking recursive (geometric) powers

$$(A3.5.20) \quad (I_n + A), (I_n + A)^2, ((I_n + A)^2)^2 = (I_n + A)^4, \dots$$

one will typically reach a positive matrix rather quickly when A is strongly connected, and will otherwise identify persistent zeros indicating which nodes fail to be connected.

² This proof is also available online at <http://www.math.harvard.edu/library/sternberg/slides/1180912pf.pdf>.

³ Another more sophisticated (and powerful) geometric approach is Birkhoff’s contraction-mapping argument, which is developed and discussed in detail by Kohlberg and Pratt (1982).

A final property of matrix, P , which is crucial for the analysis to follow is that it *commutes with A* , i.e.,

$$(A3.5.21) \quad PA = (I_n + A)^{n-1}A = A + A^n = A(I_n + A^{n-1}) = AP$$

With these preliminaries we now proceed to the proof.

Proof for the General Case:

If the *nonnegative orthant* in \mathbb{R}^n is denoted by $Q = \{x \in \mathbb{R}^n : x \geq 0\}$, then our analysis focuses on the *positive orthant* in \mathbb{R}^n ,

$$(A3.5.22) \quad Q_+ = \{x \in Q : x \neq 0\}$$

consisting of all nonnegative vectors with at least one positive component. We also denote the *strictly positive orthant* in \mathbb{R}^n by

$$(A3.5.23) \quad Q_{++} = \{x \in \mathbb{R}^n : x > 0\}$$

In this setting, if $A = (a_{ij} : i, j = 1, \dots, n)$ denotes a given *strongly connected nonnegative n -square matrix*, with associated positive matrix P in (A3.5.17), then these two matrices exhibit the following mapping properties. First, observe from the positivity of P that for each $x \in Q_+$ it must be true that $Px > 0$, since any positive component of x is enough to make all components of Px positive. So we must have

$$(A3.5.24) \quad P(Q_+) \subseteq Q_{++}$$

In particular, this implies that for any $x, y \in \mathbb{R}^n$,

$$(A3.5.25) \quad x - y \in Q_+ \Rightarrow P(x - y) \in Q_{++} \Rightarrow Px > Py$$

Turning next to A , observe first for any $x \in Q_+$ we must have $Ax \in Q_+$. For if not, then $Ax = 0$ would imply $PAx = P0 = 0$, which together with (A3.5.21) yields $A(Px) = 0$. But since Px is strictly positive by (A3.5.24), this would in turn require that $A = O_{n \times n}$ and thus would violate strong connectivity. So we must have

$$(A3.5.26) \quad A(Q_+) \subseteq Q_+$$

It is this property of strong connectedness underscores our interest in Q_+ . For our later purposes, we note also that

$$(A3.5.27) \quad A(Q_{++}) \subseteq Q_{++}$$

For if $x > 0$ and for any i it is true that $\sum_j a_{ij}x_j = 0$, then $a_{ij} = 0$ must hold for all j , and there can be no influence path from any j to i . Thus again by strong connectedness we must have $Ax > 0$.

Given these general mapping properties, the following proof consists of three parts. First we establish the *existence* of a positive eigenvalue, λ^+ , for A with associated positive eigenvector, x^+ , with unit length. We then show that λ^+ has *maximum* absolute value among all eigenvalues of A . Finally it is shown that the eigenvector, x^+ , is *unique*.

(i) Existence. The proof of existence focuses on a scaling relation between vectors, $x \in Q_+$, and their images, $Ax \in Q_+$, that can be used to characterize eigenvectors in a useful way. In particular, if the *scale function*, $s(x)$, for A is defined for all $x = (x_1, \dots, x_n)' \in Q_+$ by:

$$(A3.5.28) \quad s(x) = \max\{\alpha \in \mathbb{R}_+ : \alpha x \leq Ax\}$$

then the nature of this function can be seen graphically for $n = 2$ as in Figure A3.36 below:




Figure A3.36. Scale Function

Here the scaled version, $s(x)x$, of x is seen to lie on the boundary of the box,

$$(A3.5.29) \quad B_x = \{y \in \mathbb{R}^n : 0 \leq y \leq Ax\}$$

which for $n = 2$ is represented by the dashed lines in Figure A3.36. In this case, x is shown to be inside B_x , so that $s(x) \geq 1$. But if x is outside B_x , then $s(x)x$ is still on the boundary of B_x with $0 \leq s(x) < 1$. For our later purposes, it is useful to write $s(x)$ more explicitly as follows. Note first by definition that for each component of $x = (x_i : i = 1, \dots, n)$ we must have,

$$(A3.5.30) \quad s(x)x_i \leq (Ax)_i \left[= \sum_j a_{ij}x_j \right], \quad i = 1, \dots, n$$

But the nonnegativity condition, $(Ax)_i \geq 0$, implies that (A3.5.30) must hold automatically when $x_i = 0$, and thus that we need only consider *positive* components, $x_i > 0$. Moreover, since the largest scalar, α , satisfying

$$(A3.5.31) \quad \alpha \leq \frac{(Ax)_i}{x_i} \text{ for all } i \text{ with } x_i > 0$$

is simply the *minimum* of these ratios, it then follows from (A3.5.28) that $s(x)$ must be given explicitly by,⁴

$$(A3.5.32) \quad s(x) = \min_{\{i: x_i > 0\}} \frac{(Ax)_i}{x_i}$$

Note in particular that while Figure A3.36 illustrates only the strictly positive case, $x \in Q_{++}$, we now see from (A3.5.32) that this scale function, $s(x)$ is indeed well defined for all $x \in Q_+$.

But our interest in this scale function is seen most clearly by noting first from definition (A3.5.28) that $s(x)x$ never exceeds Ax in any component, i.e., that

$$(A3.5.33) \quad s(x)x \leq Ax, \quad x \in Q_+$$

Moreover, in the extreme case, where $s(x)x = Ax$, it also follows by definition that $s(x)$ must be an *eigenvalue* of A with *eigenvector*, x . So our object is to find conditions under which this extreme will be achieved.

To do so, we now focus on the relation between this scale function and the positive matrix, P . In particular, starting with any nonnegative scalar, $\alpha \leq s(x)$, observe from (A3.5.33) together with the positivity of P and the commutativity property, $PA = AP$, in (A3.5.21) that

$$(A3.5.34) \quad \alpha \leq s(x) \Rightarrow \alpha x \leq s(x)x \leq Ax \Rightarrow Ax - \alpha x \geq 0 \Rightarrow P(Ax - \alpha x) \geq 0 \\ \Rightarrow P(Ax) \geq P(\alpha x) \Rightarrow \alpha Px \leq PAx = APx$$

So by applying the scale function, s , to Px , we now see that

$$(A3.5.35) \quad \alpha \leq s(x) \Rightarrow \alpha(Px) \leq A(Px) \Rightarrow \alpha \leq s(Px)$$

⁴ This is known as the *Collatz-Wielandt Formula*, as discussed for example in Chapter 8 of Meyer (2001)

and may conclude [by setting $\alpha = s(x)$] that for all $x \in Q_+$

$$(A3.5.36) \quad s(x) \leq s(Px)$$

Moreover if $s(x)x \neq Ax$, then by almost the same argument, we see from (A3.5.24) that

$$(A3.5.37) \quad s(x)x \leq Ax \neq s(x)x \Rightarrow Ax - s(x)x \in Q_+ \\ \Rightarrow P(Ax - s(x)x) \in Q_{++} \Rightarrow s(x)Px < PAx = APx$$

and thus, by again applying the scale function, s , to Px , that

$$(A3.5.38) \quad s(x)x \neq Ax \Rightarrow s(x) < s(Px)$$

But this last result shows that if this scale function achieves a *maximum* in Q_+ , say at $z \in Q_+$, then it must be true that

$$(A3.5.39) \quad Az = s(z)z$$

For if not, then by (A3.5.38) we would have $s(z) < s(Pz)$. But since $z \in Q_+ \Rightarrow Pz \in Q_+$, this would contradict the maximality of $s(z)$. Thus to establish the existence of an eigenvector for A in Q_+ , it suffices to show that s does achieve a maximum on Q_+ .

To do so, we start by observing that this scale function is itself invariant under positive scalar transformations of x . In particular, for any positive scalar, r , it must be true for any $\alpha \in \mathbb{R}_+$ that

$$(A3.5.40) \quad \alpha(rx) \leq A(rx) \Leftrightarrow r(\alpha x) \leq r(Ax) \Leftrightarrow \alpha x \leq Ax$$

and thus from (A3.5.28) that for all $x \in Q_+$

$$(A3.5.41) \quad r > 0 \Rightarrow s(rx) = s(x)$$

In particular, since every $x \in Q_+$ has positive length, $\|x\| > 0$, it suffices to set $r = 1/\|x\|$, and consider the *normalization*, $u_x = x/\|x\|$, of each $x \in Q_+$, i.e., the subset,

$$(A3.5.42) \quad U_+ = \{u_x : x \in Q_+\} \subset Q_+$$

which in geometric terms is simply the intersection of Q_+ with the *unit sphere*, $U = \{u \in \mathbb{R}^n : \|u\| = 1\}$ in \mathbb{R}^n , as shown for $n = 2$ by the quarter circle in Figure A3.37 below:




Figure A3.37. Construction for Existence Proof

The key point here is that U_+ is a *compact* (closed and bounded) subset of Q_+ , and all values of s are necessarily achieved in U_+ . So s achieves a maximum in Q_+ if and only if it achieves a maximum in U_+ . But we can do even better than this by considering the image

$$(A3.5.43) \quad P(U_+) = \{Pu : u \in U_+\} \subset Q_{++}$$

of U_+ under P , as shown for $n = 2$ in Figure A3.37 above. For the dominance property in (A3.5.36) then shows that the maximum value of s in Q_+ (and thus in U_+) must also be achieved on $P(U_+)$. But here it is quite simple to show that maximum values of s must *always exist*. To do so, observe first that since all linear mappings are continuous, $P(U_+)$ must be the image of a compact set, U_+ , under a continuous mapping, P , and thus must itself be compact.⁵ But as we have already seen in Section A3.2, the Extreme Value Theorem⁶ ensures the existence of maxima for every continuous function on $P(U_+)$. In particular, the scale function, s , is easily seen to be continuous on $P(U_+)$. This is particularly evident from expression (A3.5.32), which reduces to the following simple form for *positive* vectors,

$$(A3.5.44) \quad s(x) = \min_{i=1, \dots, n} \frac{(Ax)_i}{x_i}$$

Since the continuity of A implies that each ratio, $(Ax)_i / x_i$, must be continuous on all of Q_{++} [and thus on $P(U_+)$], and since the minimum of a finite set of continuous functions

⁵ See for example Theorem 16.5 in Bartle (1964).

⁶ Again this is demonstrated in Theorem 16.6 of Bartle (1964).

is continuous, it then follows from (A3.5.44) that s is continuous on $P(U_+)$. Thus by the Extreme Value Theorem there must be some *positive* vector, $z^+ \in P(U_+) \subset Q_+$, satisfying

$$(A3.5.45) \quad s(z^+) \geq s(x) \quad \text{for all } x \in Q_+$$

(as again illustrated for $n = 2$ in Figure A3.37). But we have already seen from (A3.5.35) that this implies $Az^+ = s(z^+)z^+$. Moreover, since $Az^+ > 0$ by (A3.5.27) it follows that $s(z^+) > 0$. Finally, by setting

$$(A3.5.46) \quad \lambda^+ = s(z^+)$$

and taking x^+ to be the normalization of z , i.e.,

$$(A3.5.47) \quad x^+ = u_z = \frac{z}{\|z\|} > 0$$

(as in Figure A3.37), it then follows that

$$(A3.5.48) \quad Az = s(z)z \Rightarrow \frac{1}{\|z\|}Az = \frac{s(z)}{\|z\|}z \Rightarrow A\left(\frac{z}{\|z\|}\right) = s(z)\left(\frac{z}{\|z\|}\right) \\ \Rightarrow \boxed{Ax^+ = \lambda^+x^+}$$

and thus that the desired existence result is established.

(ii) Maximality. One added bonus of the argument above is that can also be used to show that no eigenvalue, $\lambda \in \text{Eig}(A)$, has greater absolute value than λ^+ , i.e., that

$$(A3.5.49) \quad |\lambda| \leq \lambda^+ \quad \text{for all } \lambda \in \text{Eig}(A)$$

Here we start by observing that for any $\lambda \in \text{Eig}(A)$ and any associated eigenvector, $y = (y_i : i = 1, \dots, n)'$, it follows by definition that,

$$(A3.5.50) \quad Ay = \lambda y \Rightarrow \lambda y_i = \sum_{j=1}^n a_{ij}y_j, \quad i = 1, \dots, n$$

It is important to note that both λ and y in (A3.5.50) may be *complex valued*. But in any case, their absolute values, $|\lambda|$ and $|y| = (|y_i| : i = 1, \dots, n)$, are *real*, and behave in the same way as for real valued λ and y .⁷ In particular, for any (possibly complex)

⁷ The *absolute value (modulus)* of a complex number, $x = a + bi$, is defined by $|x| = \sqrt{a^2 + b^2}$.

numbers, x and z , it is always true that $|xz| = |x||z|$ and that $|x+z| \leq |x|+|z|$, so that by (A3.5.50) together with the nonnegativity of A ,

$$(A3.5.51) \quad |\lambda \|y_i| = |\lambda y_i| = \left| \sum_{j=1}^n a_{ij} y_j \right| \leq \sum_{j=1}^n |a_{ij} y_j| = \sum_{j=1}^n a_{ij} |y_j|, \quad i=1, \dots, n$$

$$\Rightarrow |\lambda \|y| \leq A|y|$$

But since $y \neq 0 \Rightarrow |y| \in \mathcal{Q}_+$, it then follows from the definition of the scale function, s , in (A3.5.28) together with z^+ in (A3.5.45) and (A3.5.46) that

$$(A3.5.52) \quad |\lambda| \leq s(|y|) \leq s(z^+) = \lambda^+$$

and thus that (A3.5.49) must hold.

(iii) Uniqueness. To establish uniqueness of the eigenvector, x^+ , it suffices to show that that the equality,

$$(A3.5.53) \quad Ay = \lambda^+ y$$

cannot hold for any nonzero vector, $y \in \mathbb{R}^n$, which is not a scalar multiple of x^+ .⁸ To formulate this more precisely, we start by observing that every pair of distinct points, $x, z \in \mathbb{R}^n$, define a unique *line* in \mathbb{R}^n , denoted by

$$(A3.5.54) \quad L(x, z) = \{\alpha x + (1-\alpha)z : \alpha \in \mathbb{R}\}$$

In these terms, all linear multiples of x^+ are precisely those points on the line, $L(0, x^+)$, shown for $n=2$ in Figure A3.38 below (as a thin red line). So our objective is to show that *all* vectors, y , satisfying (A3.5.53) must lie on this line, i.e., that for all $y \in \mathbb{R}^n$,

$$(A3.5.55) \quad Ay = \lambda^+ y \Rightarrow y \in L(0, x^+)$$

However, the argument is made simpler by observing that all eigenvectors of A are shared by the *positive* matrix, P . In particular, if (A3.3.53) holds for y then,

$$(A3.5.56) \quad Ay = \lambda^+ y \Rightarrow (I+A)y = y + Ay = y + \lambda^+ y = (1+\lambda^+)y$$

⁸ Because λ^+ is real, we may focus on *real* eigenvectors. Here it is worth noting that if $Ay = \lambda^+ y$ then it is trivially true that $A(y+iy) = Ay + iAy = \lambda^+ y + i\lambda^+ y = \lambda^+(y+iy)$. But as we shall see in Section A3.5.3 below, complex eigenvectors are *only* of interest when their associated eigenvalues are complex.

$$\Rightarrow (I + A)^2 y = (I + A)[(1 + \lambda^+)y] = (1 + \lambda^+)(I + A)y = (1 + \lambda^+)^2 y$$

$$\vdots$$

$$\Rightarrow (I + A)^{n-1} y = (1 + \lambda^+)^{n-1} y$$

$$\Rightarrow Py = \theta^+ y$$

where $\theta^+ = (1 + \lambda^+)^{n-1} > 0$. So each eigenvector of A for eigenvalue, λ^+ , is necessarily an eigenvector of P with associated positive eigenvalue, θ^+ . So in particular,

$$(A3.5.57) \quad Px^+ = \theta^+ x^+$$

Moreover, if $Ay = \lambda^+ y$ for any $y \notin L(0, x^+)$, then it must also be true that $Py = \theta^+ y$. Thus to establish uniqueness of x^+ for eigenvalue, $\lambda^+ \in \text{Eig}(A)$, it is enough to establish uniqueness of x^+ for eigenvalue, $\theta^+ \in \text{Eig}(P)$. So our objective in (A3.5.55) will be achieved if it can be shown that

$$(A3.5.58) \quad \boxed{Py = \theta^+ y \Rightarrow y \in L(0, x^+)}$$

The advantage of this reformulation is that the *positivity* of P makes it simpler to show that violations of (A3.5.58) lead to a contradiction. In this context, the remaining argument is essentially a (more geometric) version of the uniqueness proof for the case of positive matrices in Lemma 2.2 of Cheng et al. (2012).




Figure A3.38. Boundary Eigenvectors

To begin with, suppose to the contrary that $Py = \theta^+ y$ for some $y \notin L(0, x^+)$. Then $y \neq x^+$ implies that the line, $L(y, x^+)$, passing through these two points is well defined. Such a line is shown (in blue) for $n = 2$ in Figure A3.38. Moreover, if both x^+ and y are eigenvectors for θ^+ , then it follows that *every* point on the line, $L(y, x^+)$, must also be an eigenvector for θ^+ . To see this, observe that each point on this line can be written as a linear (affine) combination of x^+ and y as follows,

$$(A3.5.59) \quad y^\alpha = \alpha y + (1 - \alpha)x^+$$

But by simple linearity, we see that

$$(A3.5.60) \quad \begin{aligned} Py^\alpha &= P[\alpha y + (1 - \alpha)x^+] = \alpha Py + (1 - \alpha)Px^+ \\ &= \alpha(\theta^+ y) + (1 - \alpha)(\theta^+ x^+) = \theta^+ [\alpha y + (1 - \alpha)x^+] = \theta^+ y^\alpha \end{aligned}$$

and thus that y^α is also an eigenvector θ^+ . The next key point to observe is that since $x^+ \in Q_{++}$, it follows that (regardless of where y is located) the line, $L(y, x^+)$, must intersect the *boundary* of Q_{++} , which is precisely the set of nonnegative y with at least one zero component (as shown for $n = 2$ by the point y^0 in Figure A3.38). To verify this observation, note first from the representation in (A3.5.59) that for each component, y_i , of y with $y_i \neq x_i^+$, there is a unique α value, say α_i , at which $y_i^\alpha = 0$, as can be seen as follows:

$$(A3.5.61) \quad 0 = y_i^\alpha = \alpha y_i + (1 - \alpha)x_i^+ \Rightarrow \alpha_i = \frac{x_i^+}{x_i^+ - y_i}$$

If among these α values, we choose α_m to be one with *smallest absolute value*, i.e.,

$$(A3.5.62) \quad |\alpha_m| = \min \{|\alpha_i| : y_i \neq x_i^+\}$$

then the desired point, $y^0 \in L(y, x^+)$, is given by

$$(A3.5.63) \quad y^0 = y^{\alpha_m} = \alpha_m y + (1 - \alpha_m)x^+$$

To see this intuitively, note that since points, y^α , with smaller values of $|\alpha|$ are by construction closer to $x^+ \in Q_{++}$, it then follows that y^0 must be the closest point to x^+ on the line $L(y, x^+)$ which has at least one zero component. Thus it is natural to expect this point to be in the boundary of Q_{++} . To verify this, recall first that by construction,

$y_m^0 = 0$. To show that $y^0 \geq 0$, suppose to the contrary that $y_i^0 < 0$ for some i . Then this must in turn imply that $y_i \neq x_i^+$. For if $y_i = x_i^+$, then the positivity of x^+ would imply

$$(A3.5.64) \quad y_i^0 = y_i^{\alpha_m} = \alpha_m y_i + (1 - \alpha_m)x_i^+ = x_i^+ - \alpha_m(x_i^+ - y_i) = x_i^+ > 0$$

which contradicts $y_i^0 < 0$. So by employing (A3.5.62) we see on the one hand that

$$(A3.5.65) \quad y_i^0 < 0 \Rightarrow y_i \neq x_i^+ \Rightarrow |\alpha_m| \leq |\alpha_i|$$

But on the other hand, since

$$(A3.5.66) \quad 0 > y_i^0 = y_i^{\alpha_m} = \alpha_m y_i + (1 - \alpha_m)x_i^+ = x_i^+ - \alpha_m(x_i^+ - y_i) \\ \Rightarrow \alpha_m(x_i^+ - y_i) > x_i^+ > 0$$

and since every positive product can be rewritten as,

$$(A3.5.67) \quad \alpha_m(x_i^+ - y_i) = |\alpha_m(x_i^+ - y_i)| = |\alpha_m| |(x_i^+ - y_i)|,$$

we see from (A3.5.66) that,

$$(A3.5.68) \quad y_i^0 < 0 \Rightarrow |\alpha_m| |x_i^+ - y_i| > x_i^+ \Rightarrow |\alpha_m| > \frac{x_i^+}{|x_i^+ - y_i|} = |\alpha_i|$$

which contradicts (A3.5.65). Thus we may conclude that $y^0 \geq 0$ with at least one zero component, $y_m^0 = 0$, and thus that y^0 is indeed in the boundary of Q_{++} . In addition, y^0 must have at least one *positive* component. For if not then $y^0 = 0 \in L(0, x^+)$ would contradict our hypothesis that $y^0 \notin L(0, x^+)$. Thus it also follows that $y^0 \in Q_+$. Finally, to show that no such y^0 can exist, recall from (A3.5.24) that $P(Q_+) \subset Q_{++}$, and thus that $Py^0 > 0$. But if y^0 were an eigenvector for θ^+ , then the positivity of θ^+ would imply that

$$(A3.5.69) \quad Py^0 = \theta^+ y^0 \Rightarrow y^0 = \frac{1}{\theta^+} Py^0 > 0$$

which contradicts our choice of y^0 with $y_m^0 = 0$. Thus no such y^0 exists, and we may conclude that all eigenvectors of θ^+ lie in $L(0, x^+)$, i.e., that (A3.5.62) must hold. \square

For our analyses below, we shall employ the notational conventions of Section 3.3.1 in the text by making the matrix dependencies of maximum eigenvalues and eigenvectors explicit. For each strongly connected nonnegative matrix, A , let the *maximum eigenvalue*

of A be denoted by $\lambda_A > 0$, with associated unique positive eigenvector, x_A , of unit length, so that (A3.5.48) now becomes

$$(A3.5.70) \quad Ax_A = \lambda_A x_A$$

Similarly, if we denote the maximum absolute value of all eigenvalues of A by⁹

$$(A3.5.71) \quad |\lambda|_A = \max_{\lambda \in \text{Eig}(A)} |\lambda|$$

then the maximality result in (A3.5.49) can now be restated as

$$(A3.5.72) \quad \lambda_A = |\lambda|_A$$

A3.5.3 Application to Spatial Autoregressive Kernels

The objective of this section is to apply the above results to the power expansion of spatial autoregressive kernels, as developed in Section 3.3 of the text. To do so, we start by establishing expression (3.3.7), which asserts that for *any* n -square matrix, A ,

$$(A3.5.73) \quad (I_n - A)^{-1} = \sum_{k=0}^{\infty} A^k \Leftrightarrow \lim_{k \rightarrow \infty} A^k = O_n$$

[The infinite sum on the left is known as the *Neumann series*, and has a long history in mathematics, as developed for example in Meyer (2000, Section 7.10).] Because this series can only converge if A^k shrinks to zero, necessity of the limit condition on the right hand side of (A3.5.73) should be obvious. To prove sufficiency of this condition, note simply that for any positive integer, k , we must have¹⁰

$$\begin{aligned} (A3.5.74) \quad (I_n - A)(I_n + A + A^2 + \cdots + A^{k-1}) \\ &= (I_n + A + A^2 + \cdots + A^{k-1}) - (A + A^2 + \cdots + A^k) \\ &= I_n - A^k \end{aligned}$$

and thus may conclude that

$$(A3.5.75) \quad \lim_{k \rightarrow \infty} A^k = O_n \Rightarrow \lim_{k \rightarrow \infty} (I_n - A^k) = I_n$$

⁹ This maximum absolute value is almost always referred to as the *spectral radius*, $\rho(A)$, of A . However, we choose the present notation and terminology both to avoid confusion with the spatial dependence parameter, ρ , and more importantly, to emphasize its relation to eigenvalues explicitly.

¹⁰ The cancelation of almost all terms here is known as a “telescoping sum”.

$$\Rightarrow \lim_{k \rightarrow \infty} (I_n - A)(I_n + A + A^2 + \cdots + A^{k-1}) = I_n$$

$$\Rightarrow (I_n - A) \sum_{k=0}^{\infty} A^k = I_n$$

$$\Rightarrow \sum_{k=0}^{\infty} A^k = (I_n - A)^{-1}$$

Given these general results, it should be clear that the instance of (A3.5.73) which is of most importance for our purposes is the assertion that for any spatial weights matrix, W , and spatial dependency parameter, ρ , the “ripple” expansion,

$$(A3.5.76) \quad (I_n - \rho W)^{-1} = \sum_{k=0}^{\infty} (\rho W)^k = \sum_{k=0}^{\infty} \rho^k W^k$$

holds if and only if it is true that

$$(A3.5.77) \quad \lim_{k \rightarrow \infty} \rho^k W^k = O_n$$

In this context, our main objective is to show that

$$(A3.5.78) \quad \lim_{k \rightarrow \infty} \rho^k W^k = O_n \Leftrightarrow |\rho| < 1/\lambda_W$$

For this, together with (A3.5.73), will yield the main result of this section, namely that

$$(A3.5.79) \quad (I_n - \rho W)^{-1} = \sum_{k=0}^{\infty} \rho^k W^k \Leftrightarrow |\rho| < 1/\lambda_W$$

As suggested in Section 3.3.1, the intuition behind (A3.5.78) is rather simple. In fact, powers (A^k) of any matrix, A , should shrink to zero whenever all vectors are “shrunk” by A . This in turn should require that all eigenvalues be less than one in absolute value, i.e., that $|\lambda|_A < 1$. But since $|\lambda|_{\rho W} = |\rho| \lambda_W$, it follows that the right hand side of (A3.5.78) simply asserts that $|\rho| \lambda_W < 1$. To see that this intuition is basically correct, we first establish necessity of this condition by observing that for all k ,

$$(A3.5.80) \quad Wx_W = \lambda_W x_W \Rightarrow W^k x_W = \lambda_W^k x_W$$

$$\Rightarrow |\rho|^k W^k x_W = |\rho|^k \lambda_W^k x_W = (|\rho| \lambda_W)^k x_W$$

and thus from the positivity of both λ_W and x_W that

$$(A3.5.81) \quad \rho^k W^k \rightarrow O_n \Rightarrow |\rho|^k W^k \rightarrow O_n \Rightarrow |\rho|^k W^k x_W \rightarrow 0$$

$$\Rightarrow (|\rho| \lambda_W)^k x_W \rightarrow 0$$

$$\Rightarrow (|\rho| \lambda_w)^k \rightarrow 0$$

$$\Rightarrow |\rho| \lambda_w < 1 \Rightarrow |\rho| < 1 / \lambda_w$$

However, sufficiency of this condition is somewhat more subtle, and requires an additional result that is useful in its own right. To see the difficulty here, consider the follows weight matrix, W , that has been partitioned to show that it is *not* strongly connected:

$$(A3.5.82) \quad W = \begin{pmatrix} 0 & 1 & | & 0 & 0 \\ 1 & 0 & | & 0 & 0 \\ \hline 1 & 1 & | & 0 & 1 \\ 1 & 1 & | & 1 & 0 \end{pmatrix}$$

Here it can be verified by direct calculation that *all* eigenvalues of W are real with unit absolute value (i.e., 1 and -1). In particular, there is still a positive maximum eigenvalue, $\lambda_w = 1$, with unique *nonnegative* eigenvector, $x_w = (0, 0, 2^{-1/2}, 2^{-1/2})'$ of unit length. The surprising feature of this matrix is that even though its maximum eigenvalue is unity, the powers of this matrix diverge without bound. In particular it may be verified by direct calculation that for each power, $k = 1, 2, \dots$

$$(A3.5.83) \quad W^k = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ k & k & 1 & 0 \\ k & k & 0 & 1 \end{pmatrix}$$

So in such cases it clear that having a maximum absolute eigenvalue of one ($|\lambda|_w = 1$) need *not* bound the growth of W powers. However, if W is *strongly connected*, then such powers are indeed bounded. In particular [by using the argument of Corollary 8.1.33 in Horn and Johnson (1985)], we now show that for any strongly connected weight matrix, A , with maximum eigenvalue, λ_A , the scaled matrix powers, $(A / \lambda_A)^k$, are *uniformly bounded*, i.e., that there exists a constant matrix, C_A , such that for all positive powers, k ,

$$(A3.5.84) \quad O_n \leq \frac{1}{\lambda_A^k} A^k \leq C_A$$

To construct C_A , we first observe that if the components of A^k are denoted by $A^k = [a_{ij}^{(k)} : i, j = 1, \dots, n]$, then [as in (A3.5.80) above],

$$(A3.5.85) \quad Ax_A = \lambda_A x_A \Rightarrow A^k x_A = \lambda_A^k x_A \Rightarrow \frac{1}{\lambda_A^k} A^k x_A = x_A \\ \Rightarrow \frac{1}{\lambda_A^k} \sum_j a_{ij}^{(k)} x_{A,j} = x_{A,i} \quad , \quad i = 1, \dots, n$$

So if the maximum and minimum components of eigenvector, $x_A > 0$, are denoted by x_A^{\min} and x_A^{\max} , respectively, then it follows at once that for all components, $a_{ij}^{(k)}$,

$$(A3.5.86) \quad \frac{1}{\lambda_A^k} a_{ij}^{(k)} x_A^{\min} \leq \frac{1}{\lambda_A^k} a_{ij}^{(k)} x_{A,j} \leq \frac{1}{\lambda_A^k} \sum_j a_{ij}^{(k)} x_{A,j} = x_{A,i} \leq x_A^{\max}$$

$$\Rightarrow \frac{1}{\lambda_A^k} a_{ij}^{(k)} \leq \frac{x_A^{\max}}{x_A^{\min}}$$

But since the right hand side is *independent* of k , this in fact yields the desired uniform bound. In particular, if the bounding matrix, C_A , is defined simply to have constant component values, x_A^{\max} / x_A^{\min} , i.e.,

$$(A3.5.87) \quad C_A = \left(\frac{x_A^{\max}}{x_A^{\min}} \right) \mathbf{1}_n \mathbf{1}'_n$$

then it follows at once from (A3.5.86) that (A3.5.84) must hold for this choice of C_A .

With this result, the desired sufficiency argument for (A3.5.78) is almost immediate. By letting $A = W$ in (A3.5.84), it is enough to observe that

$$(A3.5.88) \quad |\rho| < 1/\lambda_W \Rightarrow |\rho| \lambda_W \in (0,1) \Rightarrow (|\rho| \lambda_W)^k \rightarrow 0$$

$$\Rightarrow (|\rho| \lambda_W)^k C_W \rightarrow O_n \Rightarrow (|\rho| \lambda_W)^k \left(\frac{1}{\lambda_W^k} W^k \right) \rightarrow O_n$$

$$\Rightarrow |\rho|^k \lambda_W^k \left(\frac{1}{\lambda_W^k} W^k \right) \rightarrow O_n \Rightarrow |\rho|^k W^k \rightarrow O_n$$

$$\Rightarrow \rho^k W^k \rightarrow O_n$$

So (A3.5.78) is established, and it follows that the key convergence condition in (A3.5.79) must hold.

Having established necessary and sufficient condition for the existence of this power expansion for spatial autoregressive kernels, it is important to consider some of the general implications and possible generalizations of this result. We do so in the following two subsections.

The Importance of Power Representations

Recall from the closing discussion of Section 3.3.3 that spatial autoregressive kernels, $(I_n - \rho W)^{-1}$, can exist even when the power representation

$$(A3.5.89) \quad (I_n - \rho W)^{-1} = \sum_{k=0}^{\infty} \rho^k W^k$$

does not. However, it was also asserted that for the most important case of positive spatial dependencies, $\rho > 0$, such kernels must necessarily have *negative* elements when ρ exceeds $1/\lambda_w$. In particular, we now establish the following result:

For any strongly connected spatial weights matrix, W , and nonnegative spatial dependency parameter, $\rho \geq 0$, if $(I_n - \rho W)^{-1}$ exists then,

$$(A3.5.90) \quad (I_n - \rho W)^{-1} \geq O_n \Leftrightarrow \rho < 1/\lambda_w$$

To establish (A3.5.90), note first that since this assertion is trivially true when $\rho = 0$ (i.e., since $I_n \geq O_n$ and $0 < 1/\lambda_w$ by the positivity of λ_w), we need only consider the case of positive spatial dependence, $\rho > 0$. For the sufficiency part (\Leftarrow) of assertion (A3.5.90), we then obtain the following sharper result. Here strong connectedness together with the argument in (A3.5.19) above shows that for each ij pair, we must have $(W^k)_{ij} > 0$ for some $k \leq n-1$. Thus if $\rho > 0$ then it follows from (A3.5.79) that,

$$(A3.5.91) \quad \rho < 1/\lambda_w \Rightarrow (I_n - \rho W)^{-1} = \sum_{k=0}^{\infty} \rho^k W^k \geq \sum_{k=0}^{n-1} \rho^k W^k > O_n$$

and thus that $(I_n - \rho W)^{-1}$ is *strictly positive* in this case. To establish necessity (\Rightarrow), note first that if $\rho = 1/\lambda_w$ that since,

$$(A3.5.92) \quad (I_n - \rho W)x_w = x_w - \left(\frac{1}{\lambda_w}\right)Wx_w = x_w - \left(\frac{1}{\lambda_w}\right)(\lambda_w x_w) = x_w - x_w = 0$$

it follows from the positivity of x_w that $(I_n - \rho W)$ must be singular, and thus that $(I_n - \rho W)^{-1}$ cannot exist. So it is enough to show that $(I_n - \rho W)^{-1}$ has negative elements whenever $\rho > 1/\lambda_w$. To do so, note first that x_w is also an eigenvector of $(I_n - \rho W)^{-1}$ with eigenvalue, $(1 - \rho\lambda_w)^{-1}$, since

$$(A3.5.93) \quad \begin{aligned} (I_n - \rho W)x_w &= x_w - \rho Wx_w = x_w - \rho(\lambda_w x_w) = (1 - \rho\lambda_w)x_w \\ &\Rightarrow x_w = (I_n - \rho W)^{-1}(1 - \rho\lambda_w)x_w \\ &\Rightarrow (I_n - \rho W)^{-1}x_w = \left(\frac{1}{1 - \rho\lambda_w}\right)x_w \end{aligned}$$

But if $\rho > 1/\lambda_w$, then this eigenvalue is *negative*,

$$(A3.5.94) \quad \rho > 1/\lambda_w \Rightarrow \rho\lambda_w > 1 \Rightarrow 1 - \rho\lambda_w < 0$$

which together with the positivity of x_w that

$$(A3.5.95) \quad \rho > 1/\lambda_w \Rightarrow (I_n - \rho W)^{-1} x_w < 0$$

So whenever $\rho > 1/\lambda_w$, we see that $(I_n - \rho W)^{-1}$ maps the positive vector, x_w , into a negative vector, and must therefore have *negative elements*. Thus the assertion in (A3.5.90) is established. \square

It should also be noted from (A3.5.93) that for $\rho > 1/\lambda_w$, the autoregressive kernel, $(I_n - \rho W)^{-1}$, must in fact have negative elements in *every row*. So we see from the form of the spatial autoregressive model itself,

$$(A3.5.96) \quad u = \rho W u + \varepsilon \Rightarrow u = (I_n - \rho W)^{-1} \varepsilon$$

that each component of u must be a *decreasing function* of at least one element of ε [as in the example of expression (3.3.37) in the text]. Thus for strongly connected weight matrices, W , with spatial dependency parameters, $\rho > 1/\lambda_w$, there must always be widespread *negative dependencies* between u and ε in spite of the positivity of ρ . So while this spatial autoregressive model is formally well defined for a broader range of ρ values, the behavioral meaning of this model loses much of its appeal. Similar observations can be made for the strong connectivity condition, as we now show.

Possible Relaxations of Strong Connectedness

Given the above results for strongly connected weight matrices, it must be emphasized that many of these results are in fact far more general. Here we simply summarize such generalizations, and then reexamine strong connectedness from this perspective.

First, certain parts of the Perron-Frobenius Theorem itself are extendable to *all nonnegative matrices*, A . Of most importance for our purposes is the fact that the eigenvalue, λ_A , with maximum absolute value, $|\lambda_A|$, is *always nonnegative*, so that

$$(A3.5.97) \quad \lambda_A = |\lambda_A|$$

holds identically for all nonnegative matrices, as demonstrated, for example, in expression (8.3.1) of Meyer (2000).¹¹ [It is worth noting however, that this extension is only made possible by first establishing the Perron Theorem (i.e., the original case of

¹¹ In Meyer (see footnote 9 above) the maximum absolute value $|\lambda_A|$ is designated as the spectral radius, $\rho(A)$.

positive matrices) and then using the fact that all nonnegative matrices are appropriate limits of positive matrices.]

In addition, it can be shown by using this more general result that our main convergence condition (A3.5.79) for spatial autoregressive kernels continues to hold for *all* spatial weights matrices, W . In fact, since the maximum absolute value (spectral radius), $|\lambda|_A$, is perfectly well defined for arbitrary matrices, A , (even when no real eigenvalues exist), this convergence condition turns out to be meaningful in all cases. More specifically, it can be shown [as for example in expressions (7.10.8) through (7.10.11) of Meyers (2000)] that for every n -square matrix, A ,

$$(A3.5.98) \quad (I_n - A)^{-1} = \sum_{k=0}^{\infty} A^k \Leftrightarrow |\lambda|_A < 1$$

In this light, expression (A3.5.79) is simply the special case with $A = \rho W$ and $|\lambda|_{\rho W} = |\rho \lambda_W| = |\rho| |\lambda_W|$.

In light of these more general results, it is appropriate to ask why strong connectivity of spatial weights matrices, W , is of so much interest. Here the argument parallels that for positive spatial dependency parameters, where values of ρ were restricted to ensure meaningful (nonnegative) spatial influences. For the case of strong connectivity, the present argument is most transparent for the important case of *symmetric* weight matrices, W , such as the (queen) contiguity in expression (A3.5.2) above. This example was used to motivate strong connectivity precisely because the underlying spatial system in Figure A3.31 consisted of physically separated subsystems that shared no boundaries whatsoever. More generally, it is easily shown that failures of strong connectivity for symmetric weight matrices, W , are *always* of this type, namely that W must have a *block diagonal* form reflecting subsystems that share no linkages whatsoever.

When analyzing spatial autocorrelation effects in particular, there are strong statistical grounds for treating such blocks separately. For example, suppose we consider a *Spatial Errors Model*,

$$(A3.5.99) \quad Y = X\beta + u, \quad u = \rho Wu + \varepsilon, \quad \varepsilon \sim N(0, \sigma^2 I_n)$$

for the two subsystems in Figure A3.31, where Y is some relevant regional variable, say per capita income. Then by partitioning vectors in terms of these two subsystems, we see that all spatial autoregressive dependencies can be written as:

$$(A3.5.100) \quad \begin{pmatrix} u_1 \\ u_2 \end{pmatrix} = \rho \begin{pmatrix} W_1 & \\ & W_2 \end{pmatrix} \begin{pmatrix} u_1 \\ u_2 \end{pmatrix} + \begin{pmatrix} \varepsilon_1 \\ \varepsilon_2 \end{pmatrix}$$

$$\Rightarrow u_i = \rho W_i u_i + \varepsilon_i, \quad i = 1, 2$$

$$\Rightarrow u_i = (I_{n_i} - \rho W_i)^{-1} \varepsilon_i = B_i^{-1} \varepsilon_i, \quad i = 1, 2$$

But since ε_1 and ε_2 are *independent* by hypothesis, it then follows that

$$\begin{aligned} \text{(A3.5.101)} \quad \text{cov}(u_1, u_2) &= \text{cov}(B_1^{-1}\varepsilon_1, B_2^{-1}\varepsilon_2) \\ &= B_1^{-1} \text{cov}(\varepsilon_1, \varepsilon_2) (B_2')^{-1} = B_1^{-1} O_{n_1 \times n_2} (B_2')^{-1} \\ &\Rightarrow \text{cov}(u_1, u_2) = O_{n_1 \times n_2} \end{aligned}$$

and thus (by multi-normality) that u_1 and u_2 are also *independent*. So we see that in terms of unobserved spatial errors, these regional subsystems are *statistically independent*. In other words, all “ripple effects” generated in a given subsystem (as in Section 3.3.3) must stay in that subsystem, and never influence other subsystems.

When W is not symmetric then some linkages are possible, and the argument becomes more complex. But again the most important ideas are best illustrated by examples. A particularly dramatic example was given in Figure 7.6 of the text, where (asymmetric) first nearest neighbors were plotted. Here it was shown that the spatial linkages provided by this weight matrix were so sparse that *no spatial autocorrelation* was detectable by standard tests. Even if this neighbor relation is symmetrized, as in Figure A3.39 below, such relations are still so sparse that Eire is seen to be partitioned into six disjoint subsystems. So unless there is some reason to believe that spatial correlation effects are only possible within these six subregions, such sparse linkage structures are seen to impose strong prior restrictions on possible correlation effects.

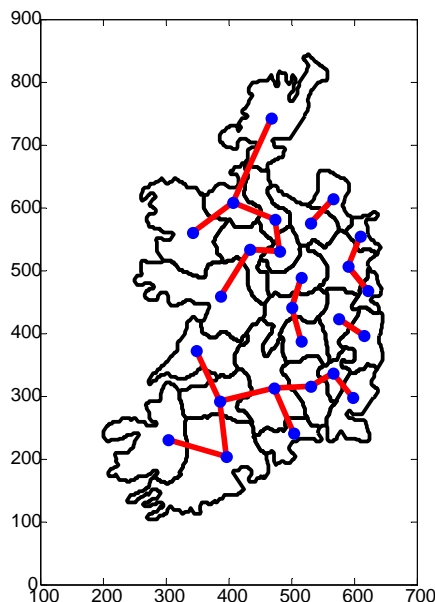


Figure A3.39. Symmetric Nearest-Neighbor Relations in Eire

More generally, unless there are compelling prior reasons to believe that certain subregions cannot be influenced by others in terms of spatial autocorrelation, it is most reasonable to consider spatial weight structures, W , that at least allow the possibility of such influences. The most direct way to ensure this is to require that W be strongly connected. In the case of Eire, for example, the two matrices (W and W_{shares}) exhibiting by far the most significant spatial autocorrelation in Table 7.2 are both strongly connected.¹²

A3.5.4 Geometry of Complex Eigenvalues

Before leaving the topic of nonnegative matrices it is important to note that while their maximum eigenvalues are always nonnegative – and hence real – their other eigenvalues need *not be real*. For spatial weights matrices in particular, recall from the Spectral Decomposition Theorem in Section A3.4.5 above, that this can never happen for symmetric spatial relations. But if such relations are asymmetric, then this is quite possible, as illustrated by the following simple example. Suppose that distances within a given city are measured by shortest driving distances, and that asymmetries are here created by the presence of one-way streets. In particular, suppose we consider nearest neighbors among the three street-corner locations (a , b , c) shown in Figure A3.40 below:




Figure A3.40. Asymmetric Urban Neighbors

While location b is the nearest neighbor of a , and c is the nearest neighbor of b , the street configuration shows that the nearest neighbor of c is location a . So the relevant nearest-neighbor weights matrix for these three locations is given by

$$(A3.5.102) \quad W = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix}$$

(where the top row is “ a ”, etc.) This spatial weights matrix is strongly connected, so that its maximum eigenvalue must be positive, with a unique positive eigenvector of unit length. But this is the *only* real eigenvalue of W . It particular, it can be verified by simple

¹² Here the queen contiguity matrix, W_{queen} , is also strongly connected, while W_{nn5} is not. So the fact that spatial autocorrelation is slightly “more significant” for the latter shows that strong connectivity by itself is certainly not the *only* consideration in constructing relevant weight matrices.

multiplication (using the rules for complex numbers)¹³ that the full eigenstructure, $W = V\Lambda V^{-1}$, is given by

$$(A3.5.103) \quad \Lambda = \begin{pmatrix} 1 & & \\ & -\frac{1}{2} + i\frac{\sqrt{3}}{2} & \\ & & -\frac{1}{2} - i\frac{\sqrt{3}}{2} \end{pmatrix}, \quad V = \left(\frac{1}{\sqrt{3}}\right) \begin{pmatrix} 1 & -\frac{1}{2} + i\frac{\sqrt{3}}{2} & -\frac{1}{2} - i\frac{\sqrt{3}}{2} \\ 1 & -\frac{1}{2} - i\frac{\sqrt{3}}{2} & -\frac{1}{2} + i\frac{\sqrt{3}}{2} \\ 1 & 1 & 1 \end{pmatrix}$$

So while $\lambda_w = 1 > 0$ with $x_w = (1, 1, 1)'$, the other two eigenvalues are seen to be complex valued with associated complex eigenvectors.¹⁴ To understand the geometric meaning of this eigenstructure, it is convenient to ignore the specific numerical values in (A3.5.95) for the moment and suppose that $W = V\Lambda V^{-1}$ with

$$(A3.5.104) \quad \Lambda = \begin{pmatrix} \lambda & & \\ & \alpha + i\beta & \\ & & \alpha - i\beta \end{pmatrix}, \quad V = (v, x + iy, x - iy)$$

where v, x, y are vectors in \mathbb{R}^3 . Then by definition,

$$(A3.5.105) \quad WV = V\Lambda \Rightarrow W(v, x + iy, x - iy) = (v, x + iy, x - iy) \begin{pmatrix} \lambda & & \\ & \alpha + i\beta & \\ & & \alpha - i\beta \end{pmatrix}$$

$$\Rightarrow \begin{bmatrix} Wv = \lambda v \\ W(x + iy) = (\alpha + i\beta)(x + iy) \\ W(x - iy) = (\alpha - i\beta)(x - iy) \end{bmatrix}$$

The first line simply shows that λ is a real eigenvalue with eigenvector, v . Moreover, (by using vector versions of the rules in footnote 13), it follows that the second line can be expanded as follows:

$$(A3.5.106) \quad W(x + iy) = (\alpha + i\beta)(x + iy) = (\alpha x + i\beta y + ibx - by)$$

$$\Rightarrow Wx + iWy = (\alpha x - by) + i(\beta x + \alpha y)$$

$$\Rightarrow [Wx - (\alpha x - \beta y)] + i[Wy - (\beta x + \alpha y)] = 0$$

¹³ The only rules required here are complex addition, $(a + ib) + (c + id) = (a + c) + i(b + d)$, and multiplication, $(a + ib)(c + id) = (ac - bd) + i(ad + bc)$.

¹⁴ Here it should be noted that complex eigenvalues of real-valued matrices must *necessarily* have complex eigenvectors (so that all imaginary parts can cancel out in multiplications).

But since the complex vector in the last line of (A3.5.106) can be zero if and only if *both* its real part and imaginary parts are zero, it follows that both vectors $Wx - (\alpha x - \beta y)$ and $Wy - (\beta x + \alpha y)$ must be zero. We thus obtain the following pair of *real-valued* equalities:

$$(A3.5.107) \quad \begin{aligned} Wx &= \alpha x - \beta y \\ Wy &= \beta x + \alpha y \end{aligned}$$

that allow us to interpret this (A3.5.106) in more familiar terms. Before doing so, it is important to note that an expansion of the third line of (A3.5.106) yields exactly the same pairs of real-valued equality. In particular, this shows that the complex-valued relations in the second and third lines are strongly coupled, and we shall return to this coupling below. But for the present we focus the two equations in (A3.5.107), and observe that they can be combined as follows:

$$(A3.5.108) \quad \begin{aligned} W(x, y) &= [Wx, Wy] = [\alpha x - \beta y, \beta x + \alpha y] \\ &= \left[(x, y) \begin{pmatrix} \alpha \\ -\beta \end{pmatrix}, (x, y) \begin{pmatrix} \beta \\ \alpha \end{pmatrix} \right] = (x, y) \begin{pmatrix} \alpha & \beta \\ -\beta & \alpha \end{pmatrix} \end{aligned}$$

Note moreover that the 2×2 matrix on the right hand side has very special properties. In particular, its columns are *orthogonal* with the same (squared) length, $\alpha^2 + \beta^2$, as can be seen from the matrix product:

$$(A3.5.109) \quad \begin{pmatrix} \alpha & \beta \\ -\beta & \alpha \end{pmatrix}' \begin{pmatrix} \alpha & \beta \\ -\beta & \alpha \end{pmatrix} = \begin{pmatrix} \alpha^2 + \beta^2 & 0 \\ 0 & \alpha^2 + \beta^2 \end{pmatrix} = (\alpha^2 + \beta^2) I_2$$

In geometric terms, this matrix involves both *rotation* and *dilation*, which can be made more explicit by rewriting the matrix as follows,

$$(A3.5.110) \quad \begin{pmatrix} \alpha & \beta \\ -\beta & \alpha \end{pmatrix} = \sqrt{\alpha^2 + \beta^2} \begin{pmatrix} \frac{\alpha}{\sqrt{\alpha^2 + \beta^2}} & \frac{\beta}{\sqrt{\alpha^2 + \beta^2}} \\ -\frac{\beta}{\sqrt{\alpha^2 + \beta^2}} & \frac{\alpha}{\sqrt{\alpha^2 + \beta^2}} \end{pmatrix} = \|(\alpha, \beta)\| R_{\alpha, \beta}$$

In this form, the matrix, $R_{\alpha, \beta}$, on the right hand side is now seen to be an *orthonormal* matrix (with columns of unit length), and thus is a pure *rotation* (as discussed in Section A3.1.2 above). By plotting this transformation for typical values (say $\alpha = \sqrt{3}$ and $\beta = 1$), we can see that the basis vectors, (e_1, e_2) , are both rotated by the same angle, θ ($= 30^\circ$), in \mathbb{R}^2 , where $\cos(\theta) = \alpha / \sqrt{\alpha^2 + \beta^2}$ and $\sin(\theta) = \beta / \sqrt{\alpha^2 + \beta^2}$ (see also Figure A3.9 above):




Figure A3.41. Rotation Operator

In these terms, the rotation matrix, $R_{\alpha,\beta}$ can be expressed entirely in terms of angle, θ , so that (A3.5.110) becomes

$$(A3.5.111) \quad \begin{pmatrix} \alpha & \beta \\ -\beta & \alpha \end{pmatrix} = \|(\alpha, \beta)\| \begin{pmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{pmatrix}$$

This serves to emphasize the geometric properties of this matrix in terms of *rotation* by angle, θ , and *dilation* by scaling factor, $\|(\alpha, \beta)\|$.

Given this geometric interpretation, we now put all of these results together in the following way. If the three vectors, v, x, y , in (A3.5.105) are now designated as *basis vectors*, b_1, b_2, b_3 , then by (A3.5.105) together with (A3.5.108), we can express W equivalently as follows:

$$(A3.5.112) \quad W(b_1, b_2, b_3) = [Wb_1, Wb_2, Wb_3] = [\lambda b_1, \alpha b_2 - \beta b_3, \beta b_2 + \alpha b_3]$$

$$= (b_1, b_2, b_3) \begin{pmatrix} \lambda & 0 & 0 \\ 0 & \alpha & \beta \\ 0 & -\beta & \alpha \end{pmatrix}$$

Note that this looks “almost” like the first line of (A3.5.105) – except that there are no longer any complex numbers! In fact, by writing these vectors in matrix form as $B = [b_1, b_2, b_3]$, we obtain the following equivalent *real-valued* decomposition of W :

$$(A3.5.113) \quad WB = B \begin{pmatrix} \lambda & 0 & 0 \\ 0 & \alpha & \beta \\ 0 & -\beta & \alpha \end{pmatrix} \Rightarrow W = B \begin{pmatrix} \lambda & 0 & 0 \\ 0 & \alpha & \beta \\ 0 & -\beta & \alpha \end{pmatrix} B^{-1}$$

From an algebraic viewpoint, we have thus re-expressed linear transformation, W , in terms of a new set of basis vectors, B , which best reveal the intrinsic structure this transformation. If all eigenvalues of W were *real* (and distinct), then this new version would of course be precisely the eigenstructure of W (i.e., $W = V\Lambda V^{-1}$). So in all cases, (A3.5.113) yields a new representation of W that is expressible entirely in real terms (known as the *real Jordan form* of W)¹⁵. But even more important for our present purposes is the geometric intuition that this representation provides about the linear transformation itself. In the present case of (A3.5.102) and (A3.5.103), where $\alpha = -\frac{1}{2}$, $\beta = \frac{\sqrt{3}}{2}$, and

$$(A3.5.114) \quad b_1 = \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}, b_2 = \begin{pmatrix} -\frac{1}{2} \\ -\frac{1}{2} \\ 1 \end{pmatrix}, b_3 = \begin{pmatrix} \frac{\sqrt{3}}{2} \\ -\frac{\sqrt{3}}{2} \\ 0 \end{pmatrix}$$

this transformation, W , can be depicted as in Figure A3.42 below:




Figure A3.42. Geometry of the W Transformation

Here the unit vector, $b_1 = 1_3$, is fixed under W , and serves as the axis of rotation for a rotation in the plain defined by vectors b_2 and b_3 .¹⁶ In this case, W is seen from (A3.5.102) to be *orthonormal*, so that there is no dilation associated with this rotation ($b_2^2 + b_3^2 = 1/4 + 3/4 = 1$). But in all cases, complex eigenvalues indicate that there is some degree of rotation in the plane defined by the corresponding pairs of basis vectors (such as b_2 and b_3).

¹⁵ A particularly lucid treatment of real Jordan forms can be found in the lecture notes by K.R. Matthews, available online at <http://www.numbertheory.org/courses/MP274/realjord.pdf>.

¹⁶ For ease of visual representation, the angle of rotation (120°) for these vectors has been shortened.

Here it is worth noting that even in \mathbb{R}^n , rotations are fundamentally *two dimensional* in nature. This classic result (known as *Euler's Theorem*)¹⁷, shows for example in that in the present case of \mathbb{R}^3 , every possible way of turning the unit sphere about the origin (no matter how complex) is always equivalent to a simple rotation about some fixed axis, as in Figure A3.42. This in fact is why complex eigenvalues must always appear in *pairs* ($\alpha + i\beta, \alpha - i\beta$) [called “complex conjugates”]. The vectors for their real and imaginary parts then define the plane in which this rotation occurs. As a consequence, every matrix of odd dimension ($n = 3, 5, \dots$) must have *at least one real eigenvalue*. For the case of even dimensions ($n = 2, 4, \dots$), matrices can have all complex eigenvalues. An example matrix, A , was given for $n = 2$ in (3.3.13) of the text [and can be extended to all higher (even) dimensions by block-diagonal repetitions of A].

This geometric interpretation of complex eigenvalues is no accident. In fact, it is a simple matter to show that the 2×2 orthogonal matrix in (A3.5.111) yields a *completely general representation* of complex numbers, $\alpha + i\beta$. In these terms, complex addition becomes matrix addition, and complex multiplication,

$$(A3.5.115) \quad (a + ib)(c + id) = (ac - bd) + i(ad + cb)$$

becomes matrix multiplication

$$(A3.5.116) \quad \begin{pmatrix} a & b \\ -b & a \end{pmatrix} \begin{pmatrix} c & d \\ -d & c \end{pmatrix} = \begin{pmatrix} ac - bd & ad + cb \\ -(ad + cb) & ac - bd \end{pmatrix}$$

Real numbers, a , are then simply the special case of diagonal matrices, aI_2 . Of fundamental importance is the fact that the real number, -1 , in this matrix representation now has a “square root” since

$$(A3.5.117) \quad \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}^2 = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} = \begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix}$$

So if “negation” is interpreted as “reversing direction” on the real line, then by extending this idea to *two dimensions*, the identity in (A3.5.117) shows that reversing directions (i.e., turning 180°) can be achieved by two successive 90° rotations in the plane. The key point here is that by introducing planar rotations together with dilations, one produces an extension of the real numbers in which negative roots are now well defined. Of course, such matrices are far more cumbersome to manipulate than the powerful notation of complex numbers. Thus, their main advantage (as seen in the present case of complex eigenvalues) is to provide some intuition as to the geometry underlying these mathematical constructs.

¹⁷ A nice discussion and analysis of *Euler's Theorem* can be found in Palais, et al. (2009).

A3.6 Geometry of Correlation in Regression

Recall that the geometry of squared correlation was developed as a goodness-of-fit measure for spatial regression in Section 9.3.3 of the text. Here we take an in-depth look at this geometry. To do so, we begin by noting from the development of the Regression Dual in Section 9.1.1 that data vectors can readily be depicted in *variable space*, as shown in Figure 9.4. Moreover, since correlation is interpretable as angle cosines, it should in principle be a simple matter to display such angles in variable space. The difficulty here is that standard sample correlation is in fact the angle cosine between data vectors in *deviation form*. What this means geometrically is that for any vectors, $x, y \in \mathbb{R}^n$, the *sample correlation* in (9.3.20),

$$(A3.6.1) \quad r(x, y) = \cos[\theta(Dx, Dy)]$$

is the angle cosine between the images, Dx and Dy , of these vectors under the linear transformation defined by the *deviation matrix*,

$$(A3.6.2) \quad D = I_n - \frac{1}{n}(1_n 1_n')$$

in (9.1.17). Recall also from (9.1.9) that $D1_n = 0$ and thus that for any vector, x , which is orthogonal to 1_n ,

$$(A3.6.3) \quad 1_n'x = 0 \Rightarrow Dx = x - \frac{1}{n}1_n(1_n'x) = x - 0 = x$$

A3.6.1 Deviation Space

In geometric terms, this implies that D in fact *projects* \mathbb{R}^n onto the $(n-1)$ -dimensional subspace of all vectors orthogonal to the unit vector, 1_n , here denoted by

$$(A3.6.4) \quad \mathbb{D} = \{x \in \mathbb{R}^n : 1_n'x = 0\} = \{Dx : x \in \mathbb{R}^n\} = D(\mathbb{R}^n)$$

and designated as the *deviation space* for \mathbb{R}^n .¹ For the case of $n=3$, the associated 2-dimensional deviation space, \mathbb{D} , is depicted in Figure A3.43 below (where the projected images, Dx and Dy , of two vectors, x and y , in \mathbb{R}^3 are also shown).

However, it is difficult to analyze this space directly in terms of the standard coordinate system in \mathbb{R}^n . So it is necessary to construct a new orthonormal basis for \mathbb{R}^n which yields explicit coordinates for \mathbb{D} . To do so, we start by letting $b_0 = (1/\sqrt{n})1_n$ denote the

¹ In more standard terminology, this subspace constitutes the *orthogonal complement* to the unit vector in \mathbb{R}^n .

normalization of 1_n with unit length ($\|1_n\| = \sqrt{1_n'1_n} = \sqrt{n} \Rightarrow \|b_0\| = 1$). Our objective is to extend b_0 to an orthonormal basis,

$$(A3.6.5) \quad M = [b_0, b_1, \dots, b_{n-1}] = [b_0, B]$$

for all of \mathbb{R}^n . To motivate this extension, notice first that since the deviation space, \mathbb{D} , consists of all vectors orthogonal to 1_n (and hence b_0) it then must follow that all columns of B will lie in \mathbb{D} . Moreover, since the number of these columns ($n-1$) is precisely the dimension of \mathbb{D} , it then follows that must automatically constitute an *orthonormal basis* for \mathbb{D} . This is shown in Figure A3.43 by the pair of red vectors, b_1 and b_2 .




Figure A3.43 Deviation Space in \mathbb{R}^3




Figure A3.44 Regression in \mathbb{D}

To construct such a basis, we note first that there are in principle infinitely many choices. Moreover, the *Gram-Schmidt orthogonalization procedure* outlined in expressions (A3.1.55) through (A3.1.57) above provides an operational method for constructing such a basis explicitly for any given dimension, n . However, for the case of $n = 4$ that will be used in our illustrations below, it is instructive to give a simple concrete example, as shown below:

$$(A3.6.6) \quad M = [b_0, B] = (1/2) \begin{pmatrix} 1 & \sqrt{2} & & 1 \\ 1 & -\sqrt{2} & & 1 \\ 1 & & \sqrt{2} & -1 \\ 1 & & -\sqrt{2} & -1 \end{pmatrix}$$

(where in this case, $\sqrt{n} = 2$). Here it can be seen by inspection that all columns are orthogonal with unit norms.

Turning now to the more general properties of such matrices in (A3.6.5), note first from the orthonormality of M that,

$$(A3.6.7) \quad M'M = MM' = I_n$$

so that in particular,

$$(A3.6.8) \quad \begin{pmatrix} b'_0 \\ B' \end{pmatrix} (b_0 B) = \begin{pmatrix} b'_0 b_0 & b'_0 B \\ B' b_0 & B' B \end{pmatrix} = \begin{pmatrix} 1 & \\ & I_{n-1} \end{pmatrix} \Rightarrow B'B = I_{n-1}$$

This shows that B is indeed an orthonormal basis for the deviation space, \mathbb{D} .

A3.6.2 Regression in Deviation Space

Given such a basis, our main objective is to show that OLS continues to be an *orthogonal projection* inside this subspace. Here we start by recalling the following definitions and relations developed in Section 9.1 of the text. For any given any data (y, X) with

$X = (1_n, x_1, \dots, x_k) = (1_n, X_v)$, let the OLS beta estimates, $\hat{\beta}$, and predictions, \hat{y} be defined by

$$(A3.6.9) \quad \hat{\beta} = (X'X)^{-1} X'y = \begin{pmatrix} \hat{\beta}_0 \\ \hat{\beta}_v \end{pmatrix}$$

$$(A3.6.10) \quad \hat{y} = X \hat{\beta} = \hat{\beta}_0 1_n + X_v \hat{\beta}_v$$

If the corresponding OLS residuals, \hat{e} , are defined by

$$(A3.6.11) \quad \hat{e} = y - \hat{y}$$

then by construction, these residuals satisfy

$$(A3.6.12) \quad D\hat{e} = \hat{e}$$

$$(A3.6.13) \quad \hat{y}'\hat{e} = 0$$

where the last condition describes the key orthogonality relation between OLS predictions and residuals. Our objective is to show that the same relation holds for their images in deviation space. In particular, if for any $z \in \mathbb{R}^n$ with projected image, $Dz \in \mathbb{D}$, we denote the coordinate representation of this image in terms of basis, B , for \mathbb{D} by

$$(A3.6.14) \quad \langle z \rangle = B'Dz = \begin{pmatrix} b'_1 Dz \\ \vdots \\ b'_{n-1} Dz \end{pmatrix}$$

then our objective is to show that

$$(A3.6.15) \quad \langle \hat{y} \rangle' \langle \hat{e} \rangle = 0$$

To do so, observe from (A3.6.7), (A3.6.12) and (A3.6.13) that

$$(A3.5.16) \quad (M'D\hat{y})'(M'D\hat{e}) = \hat{y}'D(MM')D\hat{e} = \hat{y}'D(I_n)D\hat{e} = \hat{y}'D\hat{e} = \hat{y}'\hat{e} = 0$$

But since $b_0'Dz = (1/\sqrt{n})1_n'Dz = 0$ for all z , it then follows from (A3.6.7), (A3.6.12) and (A3.6.16) that

$$(A3.6.17) \quad 0 = \begin{pmatrix} b_0'D\hat{y} \\ B'D\hat{y} \end{pmatrix}' \begin{pmatrix} b_0'D\hat{e} \\ B'D\hat{e} \end{pmatrix} = \begin{pmatrix} 0 \\ B'D\hat{y} \end{pmatrix}' \begin{pmatrix} 0 \\ B'D\hat{e} \end{pmatrix} = 0 + (B'D\hat{y})'(B'D\hat{e}) = \langle \hat{y} \rangle' \langle \hat{e} \rangle$$

and thus that (A3.6.15) must hold.

Finally, we show that this implies that all slope estimates, $\hat{\beta}_v$, in (A3.6.9) can be obtained by orthogonal projection in the deviation space. To do so, observe simply that by multiplying both sides of (A3.6.10) by $B'D$ and applying (A3.6.14), we obtain

$$(A3.6.18) \quad \begin{aligned} \hat{y} = \hat{\beta}_0 1_n + \sum_{j=1}^k \hat{\beta}_j x_j &\Rightarrow B'D\hat{y} = \hat{\beta}_0 B'D1_n + \sum_{j=1}^k \hat{\beta}_j B'Dx_j = \sum_{j=1}^k \hat{\beta}_j B'Dx_j \\ &\Rightarrow B'D\hat{y} = 0 + \sum_{j=1}^k \hat{\beta}_j B'Dx_j = \sum_{j=1}^k \hat{\beta}_j B'Dx_j \\ &\Rightarrow \langle \hat{y} \rangle = \sum_{j=1}^k \hat{\beta}_j \langle x_j \rangle \end{aligned}$$

as depicted for the $n = 3$ case in Figure A3.44 above. Moreover, by applying the same procedure to (A3.6.11) we also have,

$$(A3.6.19) \quad y = \hat{y} + \hat{e} \Rightarrow B'Dy = B'D\hat{y} + B'D\hat{e} \Rightarrow \langle y \rangle = \langle \hat{y} \rangle + \langle \hat{e} \rangle$$

So by (A3.6.18) and (A3.6.19) together with (A3.6.15), we see that $\langle \hat{y} \rangle$ must necessarily be the *orthogonal projection* of $\langle y \rangle$ into the span of $(\langle x_j \rangle : j = 1, \dots, k)$ -- with exactly the same beta estimates (see also Figure A3.45 below).

A3.6.3 Application to Squared Correlation for OLS and GLS

Given these regression results for deviation space, \mathbb{D} , we now return to our study of *squared correlation*, which in the present setting is directly expressible in terms of (squared) *angle cosines* in \mathbb{D} . In particular, for OLS regressions in \mathbb{R}^4 , the relevant angles between deviation images in $\mathbb{D} \subset \mathbb{R}^3$ can be illustrated graphically, as in Figure A3.45 below. To do so, it is instructive to consider a regression with $k = 2$ explanatory variables, (x_1, x_2) , and with $n = 4$ samples, so that (A3.6.18) takes the specific form:

$$(A3.6.20) \quad \hat{y} = \hat{\beta}_0 1_4 + \hat{\beta}_1 x_1 + \hat{\beta}_2 x_2 \in \mathbb{R}^4 \Rightarrow \langle \hat{y} \rangle = \hat{\beta}_1 \langle x_1 \rangle + \hat{\beta}_2 \langle x_2 \rangle \in \mathbb{D} \subset \mathbb{R}^3$$

These two explanatory variables yield a well-defined *regression plane* in \mathbb{D} that allows regression in this space to be depicted in a manner paralleling *regression dual* representation in Figure 9.6 of the text [where the unit (intercept) vector, 1_3 , and

explanatory-variable vector, x , in *variable space*, \mathbb{R}^3 , are now replaced by the projected images, $\langle x_1 \rangle$ and $\langle x_2 \rangle$, of these two explanatory variables in *deviation space*, \mathbb{D}].





Figure A3.45. OLS in Deviation Space **Figure A3.46. GLS in Deviation Space**

Here the angle, θ , between $\langle y \rangle$ and its projection, $\langle \hat{y} \rangle$, into this regression plane provides some indication of how “close” $\langle \hat{y} \rangle$ is to $\langle y \rangle$. As mentioned in Section 9.1, the advantage of using angles to measure closeness (rather than direct Euclidean distance) is that angles are independent of the particular units of $\langle y \rangle$ and $\langle \hat{y} \rangle$, and thus allow direct comparisons between different regressions. This is perhaps the single most fundamental reason why *squared angle cosines*,

$$(A3.6.21) \quad r^2(y, \hat{y}) = \cos^2 [\theta(\langle y \rangle, \langle \hat{y} \rangle)] = \cos^2 [\theta(Dy, D\hat{y})] = \frac{(y'D\hat{y})^2}{(y'Dy)(\hat{y}'D\hat{y})}$$

are used as general *goodness-of-fit* measure.

Finally, the graphic representations of these angles in Figures A3.45 and A3.46 allow one additional property of this measure to be illustrated. First observe from Figure A3.45 that the OLS projection, $\langle \hat{y} \rangle$, of $\langle y \rangle$ into regression plane constitutes an *orthogonal* projection [by (A3.6.15)]. This in turn implies that among all points on this regression plane, the angle, θ , between $\langle y \rangle$ and $\langle \hat{y} \rangle$ is as *small as possible*. But this is *not* true of the broader class of GLS regressions. In particular, we have already seen from the argument Section 7.2.2 that for any GLS model,

$$(A3.6.22) \quad Y = X\beta + u, \quad u \sim N(0, V)$$

where V may possibly involve additional unknown parameters [such as SEM with $V = \sigma^2 V_p$ in (7.3.1)], the predicted value of $E(Y | X)$ for data (y, X) is of the form

$$\begin{aligned}
 \text{(A3.6.23)} \quad \hat{y} &= X\hat{\beta} = X[(X'\hat{V}^{-1}X)^{-1}X'\hat{V}^{-1}y] \\
 &\Rightarrow B'D\hat{y} = B'DX\hat{\beta} = \sum_{j=1}^k B'Dx_j\hat{\beta}_j \\
 &\Rightarrow \langle \hat{y} \rangle = \sum_{j=1}^k \hat{\beta}_j \langle x_j \rangle
 \end{aligned}$$

So in terms of the ($k = 2$, $n = 4$) example in Figure A3.46, we see that while $\langle \hat{y} \rangle$ is still lies in the span of $\langle x_1 \rangle$ and $\langle x_2 \rangle$, it may no longer be the orthogonal projection of $\langle y \rangle$ (or any scalar multiple of this projection). So as shown, the angle between $\langle y \rangle$ and $\langle \hat{y} \rangle$ is almost always larger than for the OLS case. What this means is that for GLS models like the SEM model, the “goodness of fit” in terms of squared correlations can almost never be as large as for OLS. So there is an implicit bias toward OLS with respect to this measure.

However, it should also be added that for the SLM model, this bias no longer holds. For as was seen in expression (9.3.19) of the text, the explanatory variables (including the intercept) are in fact transformed by this model. So the span of these transformed variables in deviation space is almost never the same as that for OLS. In fact, as was seen in the Eire example, $\langle \hat{y} \rangle$ for SLM was much closer to $\langle y \rangle$ than was $\langle \hat{y} \rangle$ for OLS.

A3.7 Large Sample Properties of Maximum Likelihood Estimators

Here we develop the two most important large sample properties of Maximum Likelihood Estimators (discussed in Section 8 of the text), namely *consistency* and *asymptotic normality*. To simplify the development as much as possible, we focus on random variables with *continuous densities* governed by only a single *scalar parameter*. This approach, which allows the use of many graphical explanations, is largely inspired by the lecture notes of Dmitry Panchenko (2006).¹ For a more rigorous treatment of the vector case, see for example the lecture notes of Constantine E. Frangakis (2006).²

Here it must also be stressed that, like the development in Section 8, we consider only the classic case of *independent* random samples. Note in particular that this restriction appears to exclude the types of *spatial autocorrelation* effects that are of central interest for our purposes. However, the large-sample results that have been obtained for spatially autocorrelated samples are almost always based on an underlying assumption of “asymptotic independence”. In particular, it is assumed that larger sample sizes are obtained by considering larger spatial domains in which samples sufficiently far apart are approximately independent (in an appropriate sense). This approach, typically referred to as *expansion asymptotics*, is technically demanding and is well beyond the scope of this appendix.³

A3.7.1 Some Useful Preliminary Results

Here we gather together a number of results that will be used in subsequent sections. The first is the *Mean Value Theorem* for differentiable functions, which asserts that

Mean Value Theorem. *For any scalars, $a < b$, and continuous function, $f : [a, b] \rightarrow \mathbb{R}$, which is differentiable on (a, b) , there exists at least one interior point, $c \in (a, b)$, with derivative, $f'(c)$, given by*

$$(A3.7.1) \quad f'(c) = \frac{f(b) - f(a)}{b - a}$$

While the proof of this result is not difficult, its validity is seen most easily in graphical terms by Figure A3.47 below:

¹ These notes can be found online at <http://ocw.mit.edu/courses/mathematics/18-443-statistics-for-applications-fall-2006/lecture-notes/lecture3.pdf>.

² These notes can be found online at www.biostat.jhsph.edu/bstcourse/bio771.

³ For a detailed development of these ideas, see Lee (2002, 2004), and especially the online appendix to the latter reference. For spatial autoregressive models in particular, see also the development in Kelejian and Prucha (2010).




Figure A3.47. Mean Value Theorem

Since the right hand side of (A3.7.1) is seen to be the slope of the line from point $[a, f(a)]$ to $[b, f(b)]$ in the figure, we can find such an interior point by simply shifting this line up or down until a point of tangency is reached, like the point c shown. The slope, $f'(c)$ at this point must then satisfy (A3.7.1). Note that such tangencies are guaranteed by the smoothness of function f . Note also that there are two points of tangency in this example, illustrating that c is not always unique.

We now record (without proofs) some useful results for sequences of random variables. Because our later applications will only involve continuous random variables (i.e., with continuous density functions), we focus only on this simple case. First, a sequence of continuous random variables, X_n , is said to *converge in distribution* to a (continuous) random variable, X , and written as $X_n \rightarrow_d X$, if and only if their cumulative distribution functions (*cdf*'s), $F_n(x) = \Pr(X_n \leq x)$, converge pointwise to the *cdf* of X , i.e., if and only if for all $x \in \mathbb{R}$,

$$(A3.7.2) \quad \lim_{n \rightarrow \infty} \Pr(X_n \leq x) = \Pr(X \leq x) = F(x)$$

Notice that unlike convergence in probability, this says nothing about X_n getting “close” to any particular random variable, X . Indeed, if Y has the same distribution as X , then by definition it must also be true that $X_n \rightarrow_d Y$. For this reason, it is often more convenient to say that X_n converges in distribution to F , i.e., $X_n \rightarrow_d F$. This is was already apparent in our statement of the *Central Limit Theorem*, where the even more general notation, $X_n \rightarrow_d N(0,1)$, was used (rather than the normal *cdf*, $X_n \rightarrow_d \Phi$). While all these definitions are equivalent, each is useful in certain contexts. As we shall see, the random variable definition, $X_n \rightarrow_d X$, is most useful for stating the following properties of convergence in distribution. The first property, known as the *Continuous Mapping Theorem*, says essentially that the property of “convergence in distribution” is preserved under continuous transformations of random variables. In particular, this

theorem asserts that for any sequence of (continuous) random variables, (X_n) , converging in distribution to X , and any *continuous* function, $g(x)$,

$$(A3.7.3) \quad X_n \rightarrow_d X \Rightarrow g(X_n) \rightarrow_d g(X)$$

For example, if $g(x) = x^2$, then the theorem asserts that whenever $X_n \rightarrow_d X$, it must also be true that $X_n^2 \rightarrow_d X^2$. Again, this is a very compact way of saying that if the cdf's of X_n converge to that of X , then the cdf's of X_n^2 must converge to the cdf associated with the square of X (or any random variable with the same distribution).

The second useful result for our purposes is *Slutsky's Theorem*, which establishes convergence in distribution for certain functions of pairs of random variables. In addition, it establishes certain relations between convergence in distribution and convergence in probability. Here we state only that part of the theorem which we shall use. In particular, for any random variables, $(X_n), (Y_n), X$ and nonzero constant, $c \neq 0$,

$$(A3.7.4) \quad (X_n \rightarrow_d X) \text{ and } (Y_n \xrightarrow{prob} c) \Rightarrow \frac{X_n}{Y_n} \rightarrow_d \frac{X}{c}$$

The key to this result is the fact that the denominator Y_n converges in probability to a (nonzero) *constant* value. Indeed, the result is false if c is a random variable. Intuitively the constancy of c implies that for large n the denominator contributes little to the overall fluctuations in the ratio. [For a fuller statement of both the Continuous Mapping Theorem and Slutsky's Theorem, together with the applications of most interest for our purposes, see for example the lecture notes by James L. Powell at Berkeley, (pp.9-11) available online at http://eml.berkeley.edu/~powell/e240b_sp10/asynotes.pdf.⁴

As one application of these results, we now develop an equivalent version of the Central Limit Theorem (CLT) that will be more useful for our present purposes. Recall from Section 3.1.3 that for any independent samples, X_1, \dots, X_n of a random variable, X , with mean, μ , and variance, σ^2 , the classical version of the CLT asserts that the standardized sample mean,

$$(A3.7.5) \quad Z_n = \frac{\bar{X}_n - \mu}{\sigma / \sqrt{n}}$$

converges in distribution to the standard normal distribution, i.e.,

$$(A3.7.6) \quad Z_n \rightarrow_d N(0,1)$$

⁴ Another good discussion is given in the online lecture notes by Charles J. Geyer at the University of Minnesota (<http://www.stat.umn.edu/geyer/5101/notes/>).

So if we now consider the (continuous) linear transformation, $Z_n \rightarrow \sigma Z_n$, then by the Continuous Mapping Theorem it follows that $\sigma Z_n \rightarrow_d \sigma Z$. But recall from the Linear Invariance Theorem [expression (3.2.22)] that for any standard normal random variable, $Z \sim N(0,1)$, the transformed variable, aZ , [with variance, $a^2 \text{var}(Z) = a^2$] must be distributed as $N(0, a^2)$. So by transforming each standardized mean, Z_n , in (A2.7.6) to σZ_n , it then follows that

$$(A3.7.7) \quad \sigma Z_n \rightarrow_d N(0, \sigma^2)$$

Moreover, since (A3.7.5) implies that

$$(A3.7.8) \quad \sigma Z_n = \sigma \left(\frac{\bar{X}_n - \mu}{\sigma / \sqrt{n}} \right) = \sqrt{n} (\bar{X}_n - \mu)$$

we may now dispense with standardized means altogether and write (A3.7.7) as

$$(A3.7.9) \quad \sqrt{n} (\bar{X}_n - \mu) \rightarrow_d N(0, \sigma^2)$$

This equivalent version of the CLT is useful in cases where the variance of the original samples needs to be preserved. Recall finally from the discussion in Section 3.1.3 that from a practical viewpoint, this implies that for “sufficiently large” sample sizes, n , the left hand side of (A3.7.9) is *approximately distributed* (\approx_d) as $N(0, \sigma^2)$, i.e.,

$$(A3.7.10) \quad \sqrt{n} (\bar{X}_n - \mu) \approx_d N(0, \sigma^2)$$

A final result is *Jensen's Inequality* for convex or concave functions of random variables. Since we only require this inequality for differentiable strictly concave functions of continuous random variables, it suffices to establish the result for this case. Here we start by observing that if $g(x)$ is a *strictly concave differentiable* function, then by definition all tangents are *above* the function as shown in Figure A3.48 below. In particular, since the slope of the tangent at x is simply the derivative, $g'(x)$, it follows that for every other point, y , on the line, we must have

$$(A3.7.11) \quad g(y) < g(x) + g'(x)(y - x)$$

as seen in the figure. But since this strict inequality holds for *all* $y \neq x$, if we now consider these values as the realizations of a *continuous random variable*, X , with mean $E(X) = \mu$, and set $x = \mu$ in (A3.7.11), then it follows at once that

$$(A3.7.12) \quad g(X) < g(\mu) + g'(\mu)(X - \mu), \quad X \neq \mu$$




Figure A3.48 A Smooth Concave Function

But since $\text{prob}(X = \mu) = 0$ implies that this single point has no effect on expectations, it follows by taking expectations of both sides [and using the linearity property, $E(a + bY) = a + bE(Y)$] that,

$$\begin{aligned}
 \text{(A3.7.13)} \quad E[g(X)] &< E[g(\mu) + g'(\mu)(X - \mu)] \\
 &= g(\mu) + g'(\mu)[E(X) - \mu] = g(\mu) + (0) \\
 &= g[E(X)]
 \end{aligned}$$

Thus we obtain the most common form of *Jensen's Inequality*, which asserts that for differentiable strictly concave functions, g , of continuous random variables, X , it must always be true that

$$\text{(A3.7.14)} \quad E[g(X)] < g[E(X)]$$

A3.7.2 Consistency of Maximum Likelihood Estimators

The objective of this section is to sharpen the informal consistency argument for the sample-mean case in Section 8.1.2. To do so, let Y_1, \dots, Y_n be independent samples of a random variable, Y , with continuous probability density, $f(y|\theta)$, governed by an (unknown) scalar parameter, θ . We also assume for simplicity that $f(y|\theta)$ is *positive* on all of \mathbb{R} (like the normal density), and that θ *uniquely identifies* this density in the sense that distinct values of θ always yield distinct densities [such as the normal density, $\phi(x|\mu)$, in the example of Section 8.1]. For any realized sample values, (y_1, \dots, y_n) , recall (from Section 7.1) that the *log likelihood* if θ given (y_1, \dots, y_n) is defined by

$$(A3.7.15) \quad L_n(\theta) = L_n(\theta | y_1, \dots, y_n) = \log[f_n(y_1, \dots, y_n | \theta)] \\ = \log\left[\prod_{i=1}^n f(y_i | \theta)\right] = \sum_{i=1}^n \log[f(y_i | \theta)]$$

So if the *standardized likelihood function*, $\bar{L}_n(\theta)$, is given by

$$(A3.7.16) \quad \bar{L}_n(\theta) = \bar{L}_n(\theta | y_1, \dots, y_n) = \frac{1}{n} L_n(\theta | y_1, \dots, y_n) = \frac{1}{n} \sum_{i=1}^n \log[f(y_i | \theta)]$$

and if the expected value of $\log[f(Y | \theta)]$ under the *true value*, θ_0 , of parameter θ is denoted by⁵

$$(A3.7.17) \quad \bar{L}(\theta) = E_{\theta_0}(\log[f(Y | \theta)])$$

then (as a generalization of Section 8.1.2) it follows from the *Law of Large Numbers* (together the independence of samples Y_1, \dots, Y_n) that $\bar{L}_n(\theta)$ must converge in probability to $\bar{L}(\theta)$, i.e., that for each value of θ ,

$$(A3.7.18) \quad \bar{L}_n(\theta) \xrightarrow[\text{prob}]{} \bar{L}(\theta)$$

If the *maximum likelihood estimator*, $\hat{\theta}_n$, of θ given (y_1, \dots, y_n) is now defined (equivalently) in terms of the *standardized likelihood* in (A3.7.16) by

$$(A3.7.19) \quad \bar{L}_n(\hat{\theta}_n) = \max_{\theta} \bar{L}_n(\theta)$$

then by (A3.7.18) it is not surprising that one can infer many large sample properties of $\hat{\theta}_n$ from comparable properties of the limit function, $\bar{L}(\theta)$. Perhaps the single most important property of this limit function is that it virtually always achieves a unique maximum at the *true value*, θ_0 , of θ i.e.,

Lemma (Limit-Function Maximum). *If θ_0 is the true value of θ , then for all other values, $\theta \neq \theta_0$,*

$$(A3.7.20) \quad \bar{L}(\theta) < \bar{L}(\theta_0)$$

⁵ The subscript, θ_0 , denotes that expectation is taken with respect to the true density, $f(y | \theta_0)$. Note also that since $\log[f(y | \theta)] \rightarrow -\infty$ as $f(y | \theta) \rightarrow 0$, it is possible that for positive densities, $f(y | \theta)$, with “sufficiently thin tails”, the expected value of $\log[f(Y | \theta)]$ may not be finite. So (as in the normal density case of Section 8.1) we shall always assume that $|E_{\theta_0}(\log[f(Y | \theta)])| < \infty$.

Proof: Observe that by definition

$$\begin{aligned}
 \text{(A3.7.21)} \quad \bar{L}(\theta) - \bar{L}(\theta_0) &= E_{\theta_0}(\log[f(Y|\theta)]) - E_{\theta_0}(\log[f(Y|\theta_0)]) \\
 &= E_{\theta_0}(\log[f(Y|\theta)] - \log[f(Y|\theta_0)]) \\
 &= E_{\theta_0}\left(\log\left[\frac{f(Y|\theta)}{f(Y|\theta_0)}\right]\right)
 \end{aligned}$$

But since $\log(x)$ is a differentiable strictly concave function, and since the positivity and continuity of $f(y|\theta)$ for all θ implies that the ratio, $f(Y|\theta)/f(Y|\theta_0)$, yields a well-defined random variable with continuous density, it follows from Jensen's Inequality in (A3.7.14) above that

$$\text{(A3.7.22)} \quad E_{\theta_0}\left(\log\left[\frac{f(Y|\theta)}{f(Y|\theta_0)}\right]\right) < \log\left[E_{\theta_0}\left(\frac{f(Y|\theta)}{f(Y|\theta_0)}\right)\right]$$

Moreover, by writing the expectation on the right hand side in integral form, we see that

$$\text{(A3.7.23)} \quad E_{\theta_0}\left(\frac{f(Y|\theta)}{f(Y|\theta_0)}\right) = \int_y \left(\frac{f(y|\theta)}{f(y|\theta_0)}\right) f(y|\theta_0) dy = \int_y f(y|\theta) dy = 1$$

and thus may rewrite (A3.7.22) as

$$\text{(A3.7.24)} \quad E_{\theta_0}\left(\log\left[\frac{f(Y|\theta)}{f(Y|\theta_0)}\right]\right) < \log(1) = 0$$

Finally, by combining (A3.7.21) and (A3.7.24), we may conclude that

$$\text{(A3.7.25)} \quad \bar{L}(\theta) - \bar{L}(\theta_0) = E_{\theta_0}\left(\log\left[\frac{f(Y|\theta)}{f(Y|\theta_0)}\right]\right) < 0$$

and thus that (A3.7.20) must hold. \square

Given this maximum result, the key idea of the consistency proof to follow can be illustrated graphically as follows. First recall from (A3.7.19) that $\hat{\theta}_n$ maximizes the (continuous) standardized likelihood function, $\bar{L}_n(\theta) = \bar{L}_n(\theta | y_1, \dots, y_n)$, in (A3.7.13). For simplicity, we now assume that this maximum is always *unique*, as illustrated by the two representative standardized likelihood functions, \bar{L}_1 and \bar{L}_n , in Figure A3.49 below. The corresponding limit function, \bar{L} , is shown in red.




Figure A3.49 Convergence of Estimators

But since $\bar{L}_n(\theta)$ converges in probability to $\bar{L}(\theta)$ for each θ , it seems reasonable to suppose that the maxima of these functions should also converge in probability to the maximum of the limit function, as depicted in Figure A3.49. This is of course precisely the desired consistency result, namely, that $\hat{\theta}_n$ converge in probability to θ_0 .

Unfortunately, such convergence fails to hold without further assumptions. More generally, even for bounded continuous functions $\{f_n(x)\}$ converging pointwise to a limit function, $f(x)$, with all maxima unique, it need not be true that the maxima of f_n converge to the maximum of f . This is well illustrated by the example shown in Figure A3.50 below.




Figure A3.50 Failure of Maxima Convergence

Here the (piecewise linear) functions, $f_n(x)$, [illustrated for $n = 4, 8, 16$] are seen to be bounded and continuous, and are defined for each $x \in [0, 1]$ by:

$$(A3.7.26) \quad f_n(x) = \begin{cases} nx & , \quad 0 \leq x \leq 1/n \\ 2-nx & , \quad 1/n \leq x \leq 2/n \\ (x/2)-1/n & , \quad 2/n \leq x \leq 1 \end{cases}$$

By construction, the maximum of $f_n(x)$ is uniquely achieved at $x_n = 1/n$ with value, $f_n(x_n) = 1$ for all n . In particular, these maximum points clearly converge to $x = 0$, as shown in the figure. But, it is also clear [from the bottom line of (A3.7.26)] that these functions converge pointwise to the linear function (shown in red),

$$(A3.7.27) \quad f(x) = x/2 \quad , \quad x \in [0,1]$$

with unique maximum $f(1) = 1/2$ at $x = 1$. So even though f_n converges pointwise to f , neither the maxima of these functions nor their associated maximum values converge to those of f .

If one looks more closely at Figure A3.50, it can be seen that while $f_n(x)$ does converge to $f(x)$ at every point, x , there are always points where f_n is far away from f . In this particular case, one need look no further than the maximum points, x_n , of these functions, where $f_n(x_n) - f(x_n) = (1) - (1/2n) \geq 1/2$ for all n . More formally, while the functions f_n converge *pointwise* to f , they fail to converge *uniformly* to f . In other words, the overall “shape” of f_n never gets close to that of the limit function, f . It is precisely this this failure that creates the problem.

So it is necessary to impose a stronger “uniform convergence” condition here. Moreover, unlike the deterministic functions in the example above, standardized likelihood functions, $\bar{L}_n(\theta) = \bar{L}_n(\theta | Y_1, \dots, Y_n)$, are in fact *random* functions depending on the underlying random sample (Y_1, \dots, Y_n) . Thus it is necessary to impose a “uniform probabilistic convergence” condition on such functions. This simplest way to state this condition is to require that the “biggest” difference, $|\bar{L}_n(\theta | Y_1, \dots, Y_n) - \bar{L}(\theta)|$, between these random functions and \bar{L} converges in probability to zero as n becomes large. But since an infinite collection of differences may fail to have a distinct largest value, we focus on the *supremum* (least upper bound) of this collection. More formally, if we now denote the set of *feasible values* for parameter θ by $\Theta \subseteq \mathbb{R}$ (where in particular, $\theta_0 \in \Theta$), then this supremum can be written as

$$(A3.7.28) \quad \sup_{\theta \in \Theta} |\bar{L}_n(\theta) - \bar{L}(\theta)| = \sup_{\theta \in \Theta} |\bar{L}_n(\theta | Y_1, \dots, Y_n) - \bar{L}(\theta)|$$

For example, if θ can be any value between 0 and 1 [i.e., $\Theta = (0,1)$], and if for a given sample (y_1, \dots, y_n) it were true for each $\theta \in \Theta$ that $|\bar{L}_n(\theta | y_1, \dots, y_n) - \bar{L}(\theta)| = \theta/2$, then

while there is no maximum value of $\theta/2$ in the open interval, Θ , it is clear that $\sup_{\theta \in \Theta} |\bar{L}_n(\theta | y_1, \dots, y_n) - \bar{L}(\theta)| = 1/2$.

With this notation, we now require that sequence of standardized likelihood functions, \bar{L}_n converge uniformly in probability to \bar{L} on Θ , written as

$$(A3.7.29) \quad \sup_{\theta \in \Theta} |\bar{L}_n(\theta) - \bar{L}(\theta)| \xrightarrow{prob} 0$$

Finally, even when functions converge uniformly to a limit function, there is one additional problem with convergence of maxima that must be dealt with. This can again be illustrated most easily with deterministic functions, as in Figure A3.51 below. Here there is little need to formalize these functions explicitly. The key point is that when the function domain is *unbounded* (such as the nonnegative real line shown), the sequence of function maxima need not converge at all. In the present case, each function, f_n (shown for f_1, f_2, f_n) can be thought of as a positive constant plus a scaled version of a standard normal density centered at n , say $\alpha_n + \beta_n \phi(x - n)$, with centers, $x_n = n$, and with both $\alpha_n \downarrow 0$ and $\beta_n \downarrow 0$.




Figure A3.51 Divergence of Maxima

As seen in the figure, these functions *converge uniformly* to the limit function, $f(x) \equiv 0$, (shown in red), while their maxima diverge toward infinity. So it should be clear we must require some boundedness conditions to ensure that convergence of maxima is guaranteed in an appropriate sense. For the present scalar case, we simply require that Θ be a *compact* (closed and bounded) interval.

Given these conditions, we can now establish the desired consistency result, where all conditions above are summarized as follows:

Consistency Theorem. Given continuous standardized likelihood functions, $\bar{L}_n(\theta)$, together with limiting likelihood, $\bar{L}(\theta)$, defined on $\Theta \subseteq \mathbb{R}$, if

- (i) Θ is compact,
- (ii) \bar{L} is continuous on Θ ,
- (iii) $\hat{\theta}_n$ uniquely maximizes \bar{L}_n on Θ ,
- (iv) θ_0 uniquely maximizes \bar{L} on Θ , and
- (v) \bar{L}_n converges uniformly in probability to \bar{L} on Θ ,

then $\hat{\theta}_n$ converges in probability to θ_0 , i.e.,

$$(A3.7.30) \quad \hat{\theta}_n \xrightarrow{prob} \theta_0$$

Comments on the Consistency Theorem. The present proof follows the treatment in Frangakis (2006, Theorem 8.2), and is easily seen to be extendable to parameter vectors, $\theta = (\theta_1, \dots, \theta_k)$. Before proceeding with the proof, it should be noted that condition (iv) is simply a restatement of the above lemma on limit function maxima. But for completeness we include it as one of the key conditions.

Proof: First choose any small $\varepsilon > 0$ and let the ε -neighborhood of θ_0 in Θ be denoted by

$$(A3.7.31) \quad \Theta_\varepsilon = \{\theta \in \Theta : |\theta - \theta_0| < \varepsilon\} = (\theta_0 - \varepsilon, \theta_0 + \varepsilon)$$

as shown in Figure A3.52 below.




Figure A3.52. Limit Function




Figure A3.53. Value Dominance

Our objective is to show that $\Pr(\hat{\theta}_n \in \Theta_\varepsilon) \rightarrow 1$ for all $\varepsilon > 0$, which by definition will imply that (A3.7.30) must hold. To do so, we start by considering the values of θ *not* in Θ_ε , i.e., in the *complement* of Θ_ε in Θ , as denoted by

$$(A3.7.32) \quad \Theta_\varepsilon^c = \Theta - \Theta_\varepsilon = \{\theta \in \Theta : |\theta - \theta_0| \geq \varepsilon\}$$

Note from the compactness of Θ [condition (i)] together with the weak inequality in (A3.7.32) that Θ_ε^c is a closed bounded set, and thus is also compact. Moreover, since \bar{L} is continuous on all of Θ , it must be continuous on Θ_ε^c . Thus by the *Extreme Value Theorem* (recall Figure A3.19), it follows that \bar{L} must achieve a *maximum* in Θ_ε^c at some point, $\theta_\varepsilon \in \Theta_\varepsilon^c$. For the example shown in Figure A3.53, the point, θ_ε , happens to lie in the boundary between Θ_ε^c and Θ_ε . However this need not be the case, and more generally, θ_ε may not even be unique. But by the uniqueness of the global maximum of \bar{L} at θ_0 , it must always be true that

$$(A3.7.33) \quad \bar{L}(\theta_0) > \bar{L}(\theta_\varepsilon) \geq \bar{L}(\theta) \text{ for all } \theta \in \Theta_\varepsilon^c$$

If we now denote the positive difference between $\bar{L}(\theta_0)$ and $\bar{L}(\theta_\varepsilon)$ by

$$(A3.7.34) \quad \bar{L}(\theta_0) - \bar{L}(\theta_\varepsilon) = \delta > 0$$

(as shown in Figure A3.52), and define a band (or “sleeve”) of width, $\delta/2$, about the \bar{L} (as shown by the dotted lines in Figure A3.53), then it follows from the uniform probabilistic convergence of the standardized likelihood functions, \bar{L}_n , to \bar{L} [condition (v)] that for sufficiently large n these functions will lie entirely inside this band (with probability approaching one). This already suggests that the maxima, $\hat{\theta}_n$, of these functions cannot be very far from θ_0 . To make this more precise, we now restate definition (A3.7.29) in terms of probability events as follows. If for each n we let A_n denote the event that the random likelihood function \bar{L}_n is within distance $\delta/2$ of \bar{L} , i.e.,

$$(A3.7.35) \quad A_n \Leftrightarrow \sup_{\theta \in \Theta} |\bar{L}_n(\theta) - \bar{L}(\theta)| < \delta/2 \\ \Leftrightarrow \bar{L}_n(\theta) \in (\bar{L}(\theta) - \delta/2, \bar{L}(\theta) + \delta/2) \text{ for all } \theta \in \Theta$$

then by (A3.7.29), \bar{L}_n converges uniformly in probability to \bar{L} on Θ if and only if

$$(A3.7.36) \quad \lim_{n \rightarrow \infty} \Pr(A_n) = 1$$

But if we now consider the sub-event

$$(A3.7.37) \quad A_n^1 \Leftrightarrow \bar{L}_n(\theta) < \bar{L}(\theta) + \delta/2 \text{ for all } \theta \in \Theta_\varepsilon^c$$

then by definition, event A_n implies A_n^1 . Moreover, if we consider a second sub-event

$$(A3.7.38) \quad A_n^2 \Leftrightarrow \bar{L}_n(\theta_0) > \bar{L}(\theta_0) - \delta/2$$

then since A_n also implies A_n^2 , it follows that

$$(A3.7.39) \quad A_n \Rightarrow A_n^1 \cap A_n^2$$

But by (A3.7.31) together with (A3.7.33) we also see that,

$$(A3.7.40) \quad \begin{aligned} \bar{L}(\theta) &\leq \bar{L}(\theta_\varepsilon) = \bar{L}(\theta_0) - \delta \\ &\Rightarrow \bar{L}(\theta) + \delta/2 \leq (\bar{L}(\theta_0) - \delta) + \delta/2 = \bar{L}(\theta_0) - \delta/2 \end{aligned}$$

which together with the definition of A_n^1 yields the implication,

$$(A3.7.41) \quad A_n^1 \Rightarrow \bar{L}_n(\theta) < \bar{L}(\theta_0) - \delta/2 \text{ for all } \theta \in \Theta_\varepsilon^c$$

In other words, event A_n^1 implies that all $\bar{L}_n(\theta)$ values for $\theta \in \Theta_\varepsilon^c$ must be *less* than $\bar{L}(\theta_0) - \delta/2$. This is illustrated for the point θ shown in blue in Figure A3.53, with value, $\bar{L}_n(\theta)$, below the red horizontal line denoting value, $\bar{L}(\theta_0) - \delta/2$. But since event A_n^2 implies that $\bar{L}_n(\theta_0)$ is above this line, and since by construction, $\theta_0 \in \Theta_\varepsilon = \Theta - \Theta_\varepsilon^c$, it follows that if both events A_n^1 and A_n^2 occur, then \bar{L}_n *cannot* achieve its maximum in Θ_ε^c . Moreover, since \bar{L}_n does achieve a unique maximum on Θ at $\hat{\theta}_n$ [condition (iii)], it must be true that $\hat{\theta}_n$ lies in Θ_ε . Thus we see that

$$(A3.7.42) \quad A_n^1 \cap A_n^2 \Rightarrow \hat{\theta}_n \in \Theta_\varepsilon,$$

which together with (A3.7.42) shows that

$$(A3.7.43) \quad A_n \Rightarrow \hat{\theta}_n \in \Theta_\varepsilon.$$

But since the probability inequality, $\Pr(A) \leq \Pr(B)$, holds whenever $A \Rightarrow B$, it then follows that for all n we must have

$$(A3.7.44) \quad \Pr(A_n) \leq \Pr(\hat{\theta}_n \in \Theta_\varepsilon) \leq 1$$

Finally, since this in turn implies from (A3.7.36) that

$$(A3.7.45) \quad \lim_{n \rightarrow \infty} \Pr(\hat{\theta}_n \in \Theta_\varepsilon) = 1$$

we see that (A3.7.31) must hold. \square

A3.7.3 Asymptotic Normality of Maximum Likelihood Estimators

The purpose of this final section is to sharpen the asymptotic normality argument given for $\hat{\mu}_n$ in Section 8.1.2. So here again we replace $\hat{\mu}_n$ with a general (scalar) estimator, $\hat{\theta}_n$. To start with, recall from the argument for $\hat{\mu}_n$ that the asymptotic variance was related to the second derivative of the standardized likelihood [as in expressions (8.1.14) through (8.1.16)]. So for this to even be meaningful, the basic likelihood function,

$$(A3.7.46) \quad L(\theta | y) = \log f(y | \theta)$$

must itself be at least twice differentiable. More generally, we will now require that this likelihood function be *continuously twice differentiable*, i.e., that for all realizations, y , both functions,

$$(A3.7.47) \quad L'(\theta | y) = \frac{d}{d\theta} L(\theta | y)$$

$$(A3.7.48) \quad L''(\theta | y) = \frac{d}{d\theta} L'(\theta | y) = \frac{d^2}{d\theta^2} L(\theta | y)$$

are well defined and continuous on Θ .⁶

In terms of this notation, *Fisher Information* about the true value, θ_0 , of θ in distribution $f(y | \theta)$ is defined to be

$$(A3.7.49) \quad \mathcal{I}(\theta_0) = -E_{\theta_0} [L''(\theta_0 | Y)] = -E_{\theta_0} \left(\frac{d^2}{d\theta^2} [\log f(Y | \theta)]_{\theta=\theta_0} \right)$$

For our present purposes, Fisher information is most simply expressed as the second derivative of the limiting standardized likelihood function, \bar{L} , in (A3.7.17). This can be seen by noting from the continuity of the second derivative that order of differentiation and expectation (integration) in (A3.7.49) can be reversed to yield:⁷

⁶ Technically assertions like this need only hold for “almost all” realizations, y . But for simplicity we shall ignore such measure-theoretic considerations whenever possible.

⁷ This important result (which we shall use several times below) is known as *Leibniz rule* for integrals.

$$\begin{aligned}
\text{(A3.7.50)} \quad \mathcal{I}(\theta_0) &= -E_{\theta_0} \left(\frac{d^2}{d\theta^2} [\log f(Y | \theta)]_{\theta=\theta_0} \right) \\
&= -\frac{d^2}{d\theta^2} \left(E_{\theta_0} [\log f(Y | \theta)] \right)_{\theta=\theta_0} \\
&= -\frac{d^2}{d\theta^2} [\bar{L}(\theta)]_{\theta=\theta_0} = -\bar{L}''(\theta_0)
\end{aligned}$$

For purposes of maximum likelihood estimation, this version shows that Fisher information is precisely the *negative curvature* of \bar{L} at its maximum value, $\bar{L}(\theta_0)$ (as illustrated for μ_0 in Figures 8.2 through 8.5).

Before analyzing this quantity further, it should be emphasized that the version, $\mathcal{I}_n(\theta_0)$, used in the text [expression (8.1.21)] describes the Fisher information about parameter θ_0 in *sampling distributions of size n* generated by density f . The sampling version is more convenient for estimating large-sample variance, and is related to $\mathcal{I}(\theta_0)$ in the following simple way. First, by noting from the additivity of log products and linearity of differentiation operations that

$$\begin{aligned}
\text{(A3.7.51)} \quad L_n''(\theta_0 | y_1, \dots, y_n) &= \frac{d^2}{d\theta^2} \left[\log \left(\prod_{i=1}^n f(y_i | \theta) \right) \right]_{\theta=\theta_0} \\
&= \frac{d^2}{d\theta^2} \left[\sum_{i=1}^n \log f(y_i | \theta) \right]_{\theta=\theta_0} \\
&= \sum_{i=1}^n \frac{d^2}{d\theta^2} [\log f(y_i | \theta)]_{\theta=\theta_0}
\end{aligned}$$

we see from (A3.7.49) together with (8.1.21) that (again by interchanging differentiation and expectation):

$$\begin{aligned}
\text{(A3.7.52)} \quad \mathcal{I}_n(\theta_0) &= -E_{\theta_0} [L_n''(\theta_0 | Y_1, \dots, Y_n)] \\
&= -E_{\theta_0} \left(\sum_{i=1}^n \frac{d^2}{d\theta^2} [\log f(Y_i | \theta)]_{\theta=\theta_0} \right) \\
&= -\sum_{i=1}^n \frac{d^2}{d\theta^2} \left(E_{\theta_0} [\log f(Y_i | \theta)] \right)_{\theta=\theta_0} \\
&= -\sum_{i=1}^n \frac{d^2}{d\theta^2} [\bar{L}(\theta)]_{\theta=\theta_0} \\
&= \sum_{i=1}^n [-\bar{L}''(\theta_0)] = \sum_{i=1}^n \mathcal{I}(\theta_0)
\end{aligned}$$

$$\Rightarrow \mathcal{I}_n(\theta_0) = n\mathcal{I}(\theta_0)$$

This makes it clear that the most fundamental form of Fisher information for our purposes is given by $\mathcal{I}(\theta_0)$.

With these observations, our next objective is to show that $\mathcal{I}(\theta_0)$ can also be expressed in terms of the *first derivative* of likelihood, L , in (A3.7.47). To do so, we start by expanding the derivatives in (A3.7.47) and (A3.7.48), respectively, as

$$(A3.7.53) \quad L'(y|\theta) = \frac{d}{d\theta}[\log f(y|\theta)] = \frac{f'(y|\theta)}{f(y|\theta)}$$

and,

$$(A3.7.54) \quad L''(y|\theta) = \frac{d}{d\theta} \left[\frac{f'(y|\theta)}{f(y|\theta)} \right] = \frac{f''(y|\theta)}{f(y|\theta)} - \left(\frac{f'(y|\theta)}{f(y|\theta)} \right)^2$$

$$= \frac{f''(y|\theta)}{f(y|\theta)} - L'(\theta|y)^2$$

Next we observe from the simple probability identity

$$(A3.7.55) \quad \int_y f(y|\theta) dy = 1, \quad \theta \in \Theta$$

(and in particular from the constancy of the right hand side) that

$$(A3.7.56) \quad 0 = \frac{d}{d\theta} \int_y f(y|\theta) dy = \int_y \frac{d}{d\theta} [f(y|\theta)] dy = \int_y f'(y|\theta) dy, \quad \theta \in \Theta$$

and similarly, that

$$(A3.7.57) \quad 0 = \frac{d^2}{d\theta^2} \int_y f(y|\theta) dy = \int_y \frac{d^2}{d\theta^2} [f(y|\theta)] dy = \int_y f''(y|\theta) dy, \quad \theta \in \Theta$$

By using these results [and in particular, (A3.7.54)], we may rewrite Fisher information in (A3.7.49) as follows:

$$(A3.7.58) \quad \mathcal{I}(\theta_0) = -E_{\theta_0} [L''(\theta_0|Y)] = -\int_y L''(\theta_0|y) f(y|\theta_0) dy$$

$$= -\int_y \left[\frac{f''(y|\theta_0)}{f(y|\theta_0)} - L'(\theta_0|y)^2 \right] f(y|\theta_0) dy$$

$$\begin{aligned}
&= -\int_y \frac{f''(y|\theta_0)}{f(y|\theta_0)} f(y|\theta_0) dy + \int_y L'(\theta_0|y)^2 f(y|\theta_0) dy \\
&= -\int_y f''(y|\theta_0) dy + E_{\theta_0} [L'(\theta_0|Y)^2]
\end{aligned}$$

But since the first term on the right is zero by (A3.7.57), it then follows that

$$(A3.7.59) \quad \mathcal{I}(\theta_0) = E_{\theta_0} [L'(\theta_0|Y)^2]$$

and thus that Fisher information is also given by the second moment of the first derivative of L evaluated at θ_0 . It will turn out that this random variable, $L'(\theta_0|Y)$, exhibits a close asymptotic relation to the maximum likelihood estimator, $\hat{\theta}_n$, itself.

To gain some feeling for this relationship, note that since $\hat{\theta}_n$ is by definition the maximum value of the differentiable function, $L_n(\theta|y_1, \dots, y_n)$, it follows that if this maximum is achieved anywhere *interior* to Θ , then $\hat{\theta}_n$ must be uniquely identified by the first-order condition:

$$(A3.7.60) \quad 0 = L'_n(\hat{\theta}_n|y_1, \dots, y_n) = \sum_{i=1}^n L'(\hat{\theta}_n|y_i)$$

But since consistency of the estimator, $\hat{\theta}_n$, implies that $\hat{\theta}_n \approx \theta_0$ for large n (with probability approaching one), this first-order condition is itself characterized by the behavior of the random variable, $L'(\theta|Y)$, near the true value $\theta = \theta_0$. In view of this relation, $L'(\theta|Y)$, is usually designated as the *score function* for the maximum-likelihood problem.

As a first key property of this score function, note from (A3.7.56) that

$$(A3.7.61) \quad E_{\theta_0} [L'(\theta_0|Y)] = \int_y f'(y|\theta_0) dy = 0$$

Thus $L'(\theta_0|Y)$ has mean zero, and it follows from (A3.7.59) that Fisher information is precisely the *variance* of the score function at θ_0 . This yields a direct *variance interpretation* of $\mathcal{I}(\theta_0)$ -- which we shall return to in the proof below.

With these preliminaries, we are now ready to state the main result of this section, where we again summarize the key assumptions required for asymptotic normality:

Asymptotic Normality Theorem. *If all conditions of the above Consistency Theorem are met, and if in addition,*

(i) *The underlying density function, $f(y|\theta)$, is continuously twice differentiable on Θ for each value of y ,*

(ii) *The unique maximum likelihood estimator, $\hat{\theta}_n$, is interior to Θ , and*

(iii) *Fisher information, $\mathcal{I}(\theta_0)$, at the true value, θ_0 , is positive and finite,*

then $\hat{\theta}_n$ is asymptotically normally distributed, and in particular,

$$(A3.7.62) \quad \sqrt{n}(\hat{\theta}_n - \theta_0) \rightarrow_d N(0, \mathcal{I}(\theta_0)^{-1})$$

Proof: We begin by recalling from (3.7.62) together with conditions (i) and (ii) above that the maximum likelihood estimator, $\hat{\theta}_n = \hat{\theta}_n(y_1, \dots, y_n)$, for any given sample (y_1, \dots, y_n) is obtained by evaluating the derivative,

$$(A3.7.63) \quad L'_n(\theta) = L'_n(\theta | y_1, \dots, y_n) = \sum_{i=1}^n L'(\theta | y_i)$$

of the likelihood function at zero. But since this function is itself continuous and differentiable on all of Θ (and thus on the interval between $\hat{\theta}_n$ and θ_0), it follows from the Mean Value Theorem [expression (A3.7.1)] that there must be some intermediate point, θ_n , between $\hat{\theta}_n$ and θ_0 such that ⁸

$$(A3.7.64) \quad L''_n(\theta_n) = \frac{L'_n(\hat{\theta}_n) - L'_n(\theta_0)}{\hat{\theta}_n - \theta_0}$$

But since $L'_n(\hat{\theta}_n) = 0$ by definition, it then follows that

$$(A3.7.65) \quad L''_n(\theta_n) = \frac{-L'_n(\theta_0)}{\hat{\theta}_n - \theta_0} \Rightarrow \hat{\theta}_n - \theta_0 = -\frac{L'_n(\theta_0)}{L''_n(\theta_n)} = -\frac{L'_n(\theta_0)/n}{L''_n(\theta_n)/n}$$

$$\Rightarrow \sqrt{n}(\hat{\theta}_n - \theta_0) = \frac{\sqrt{n}[L'_n(\theta_0)/n]}{-[L''_n(\theta_n)/n]}$$

⁸ Here we ignore the (zero probability) case, $\hat{\theta}_n = \theta_0$, and for simplicity assume that the open interval between these values is entirely contained in Θ (at least for n sufficiently large).

So it is immediate from (A3.7.62) that the asymptotic distribution we seek must be the asymptotic distribution of the ratio on the right hand side. At first glance this would seem to be an even more difficult task. But fortunately the denominator has a simple probability limit that enables us to employ Slutsky's Theorem. To see this, note that for any choice of $\theta \in \Theta$

$$(A3.7.66) \quad L_n''(\theta) / n = \frac{1}{n} (L_n''(\theta | Y_1, \dots, Y_n)) = \frac{1}{n} \sum_{i=1}^n L''(\theta | Y_i)$$

But since the random variables, $L''(\theta | Y_i), i=1, \dots, n$, are independent and identically distributed, it follows at once from the Law of Large numbers that

$$(A3.7.67) \quad \frac{1}{n} \sum_{i=1}^n L''(\theta | Y_i) \xrightarrow{prob} E_{\theta_0} [L''(\theta | Y)]$$

Moreover by the consistency of $\hat{\theta}_n$ we also see that

$$(A3.7.68) \quad |\theta_n - \theta_0| < |\hat{\theta}_n - \theta_0| \xrightarrow{prob} 0 \Rightarrow \theta_n \xrightarrow{prob} \theta_0$$

So by combining these results, we obtain the probability limit,

$$(A3.7.69) \quad \frac{1}{n} \sum_{i=1}^n L''(\theta_n | Y_i) \xrightarrow{prob} E_{\theta_0} [L''(\theta_0 | Y)]$$

and may conclude from (A3.7.66) and (A3.7.69) together with (A3.7.49) that

$$(A3.7.70) \quad -L_n''(\theta_n) / n = -\frac{1}{n} \sum_{i=1}^n L''(\theta_n | Y_i) \xrightarrow{prob} -E_{\theta_0} [L''(\theta_0 | Y)] = \mathcal{I}(\theta_0)$$

Since this limit is constant, it then follows from Slutsky's Theorem [expression (A3.7.4) above] that the limiting distribution of the ratio in (A3.7.65) can be obtained by simply determining the limiting distribution of its numerator. But the numerator is again of the form

$$(A3.7.71) \quad L_n'(\theta_0) / n = \frac{1}{n} L_n'(\theta_0 | Y_1, \dots, Y_n) = \frac{1}{n} \sum_{i=1}^n L'(\theta_0 | Y_i)$$

where the random variables, $L'(\theta_0 | Y_i), i=1, \dots, n$, are independently and identically distributed. So this expression is simply a sample mean of the random variable, $L'(\theta_0 | Y)$, i.e., for the *score function* evaluated at θ_0 . Thus (as promised) the limiting distribution of $\hat{\theta}_n$ is seen to be determined entirely by its associated score function. To make this relation precise, we first recall from (A3.7.61) that the mean of $L'(\theta_0 | Y)$ is zero, so that the full numerator is seen to be of the form

$$\begin{aligned}
\text{(A3.7.72)} \quad \sqrt{n} [L'_n(\theta_0) / n] &= \sqrt{n} \left(\frac{1}{n} \sum_{i=1}^n L'(\theta_0 | Y_i) - E_{\theta_0} [L'(\theta_0 | Y)] \right) \\
&= \sqrt{n} \left(\frac{1}{n} \sum_{i=1}^n X_i - E(X) \right) \\
&= \sqrt{n} (\bar{X} - \mu)
\end{aligned}$$

where $X_i \equiv L'(\theta_0 | Y_i)$ and $\mu = E(X) = E_{\theta_0} [L'(\theta_0 | Y)] (= 0)$. So it follows at once from the form of the Central Limit Theorem in (A3.7.9) that

$$\text{(A3.7.73)} \quad \sqrt{n} [L'_n(\theta_0) / n] \rightarrow_d N(0, \text{var}_{\theta_0} [L'(\theta_0 | Y)])$$

But as we have already seen, the zero-mean property of $L'(\theta_0 | Y)$ implies that its variance, $\text{var}_{\theta_0} [L'(\theta_0 | Y)]$, is precisely *Fisher information* in (A3.7.59):

$$\text{(A3.7.74)} \quad \text{var}_{\theta_0} [L'(\theta_0 | Y)] = E_{\theta_0} [L'(\theta_0 | Y)^2] = \mathcal{I}(\theta_0)$$

So (A3.7.73) takes the more explicit form:

$$\text{(A3.7.75)} \quad \sqrt{n} [L'_n(\theta_0) / n] \rightarrow_d N[0, \mathcal{I}(\theta_0)]$$

and it follows from Slutsky's Theorem together with (A3.7.65), (A3.7.70) and (A3.7.73) that

$$\text{(A3.7.76)} \quad \sqrt{n} (\hat{\theta}_n - \theta_0) \rightarrow_d \frac{Z}{\mathcal{I}(\theta_0)} \quad \text{where } Z \sim N[0, \mathcal{I}(\theta_0)]$$

Finally, by recalling from the Invariance Theorem for normal random variables that

$$\text{(A3.7.77)} \quad Z \sim N[0, \mathcal{I}(\theta_0)] \Rightarrow \frac{Z}{\mathcal{I}(\theta_0)} \sim N \left[0, \frac{\mathcal{I}(\theta_0)}{\mathcal{I}(\theta_0)^2} \right] = N(0, \mathcal{I}(\theta_0)^{-1})$$

we may conclude from (A3.7.76) that

$$\text{(A3.7.78)} \quad \sqrt{n} (\hat{\theta}_n - \theta_0) \rightarrow_d N(0, \mathcal{I}(\theta_0)^{-1})$$

and thus that the desired result in (A3.7.62) is established. \square