Lecture 25

Quadratic Forms

This lecture provides a short introduction to quadratic forms, which have important applications in mathematical social science. Quadratic forms are central to the study of multivariate optimization, where they help classify stationary points.

25.1 Introduction to Quadratic Forms

Recall that a function is k-homogeneous if scaling its arguments by λ scales the function value by λ^k. Homogeneity of degree 2 is a core property of quadratic forms. When $f[\lambda \boldsymbol{x}] = \lambda^2 f[\boldsymbol{x}]$, it follows that f is even (i.e., $f[-\boldsymbol{x}] = f[\boldsymbol{x}]$) and that $f[\mathbf{0}] = 0$.

Definition 25.1 (Quadratic Form) An ***algebraic form*** is a homogeneous polynomial; each term is a monomial of identical degree. A ***quadratic form*** is a polynomial that is homogeneous of degree two; each term is a monomial of degree two.

A quadratic form $\mathbb{R}^N \xrightarrow{q} \mathbb{R}$ has the following representation, which is often written as $q = \boldsymbol{x} \mapsto \boldsymbol{x}^\top \cdot \boldsymbol{A} \cdot \boldsymbol{x}$.

$$q = \boldsymbol{x} \mapsto \sum_{r=1}^{N} \sum_{k=1}^{N} a_{r,k} x_r x_k \tag{25.1}$$

Example 25.1 Let f be the squaring function on the real numbers: $f = x \mapsto x^2$. Then for any scalar λ, $f[\lambda x] = (\lambda x)^2 = \lambda^2(x^2) = \lambda^2 f[x]$. So f is a quadratic form on $\mathbb{R}$. The upper plot of Figure 25.1 illustrates this case.
Let f be a binary real polynomial: $F = \langle x_1, x_2 \rangle \mapsto x_1^2 + x_2^2$. Then $f[\lambda \boldsymbol{x}] = (\lambda x_1)^2 + (\lambda x_2)^2 = \lambda^2(x_1^2 + x_2^2) = \lambda^2 f[\boldsymbol{x}]$. So f is a quadratic form on $\mathbb{R}^2$.

25.1.1 Extrema of Quadratic Forms

Theorem 25.1 Let $\mathbb{R}^N \xrightarrow{q} \mathbb{R}$ be a quadratic form. If q has an extremum, then $\mathbf{0}$ is an extreme point of q.

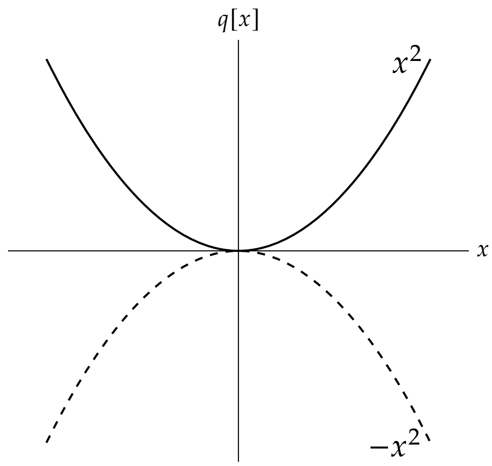

Figure 25.1: Two Unary Quadratic Forms

Proof: Suppose $\boldsymbol{x}$ maximizes q. Then $q[\boldsymbol{x}] \geq 0$, since $q[\mathbf{0}] = 0$. (A quadratic form always maps the origin to zero.) If $q[\boldsymbol{x}] = k > 0$, then $q[2\boldsymbol{x}] = 4k > k$, contradicting $\boldsymbol{x}$ being a maximizer. So if $\boldsymbol{x}$ maximizes q, then $f[\boldsymbol{x}] = 0 = f[\mathbf{0}]$. The proof for a minimizer is symmetrical. ■

Example 25.2 Figure 25.1 displays the graph of two functions, $x \mapsto x^2$ and $x \mapsto -x^2$. The first has no maximum but has a minimum at 0. The second has no minimum but has a maximum at 0.

Real quadratic forms that are positive definite or negative definite prove particularly useful in mathematical social science. Recall that a function is positive definite if every nonzero input produces a positive output, and a function is negative definite if every nonzero input produces a negative output. From theorem 25.1, if q is positive definite, then $\mathbf{0}$ is the unique minimizer. Similarly, if q is negative definite, then $\mathbf{0}$ is the unique maximizer.

Example 25.3 (Extrema of a Unary Quadratic Form) The unary quadratic form $x \mapsto ax^2$ has an extremum at 0, which is unique for $a \neq 0$. The sign of the parameter a determines the definiteness of the function, which correspondingly determines whether the extremum is a minimum or a maximum. If $a < 0$, the function is negative definite, and 0 is the unique maximizer. If $a > 0$, the function is positive definite, and 0 is the unique minimizer. Figure 25.1 illustrates this with two plots, for $a = 1$ and $a = -1$.

Example 25.4 (Extrema of a Binary Quadratic Form) The binary quadratic form $\langle x, y\rangle \mapsto a(x^2 + y^2)$ has an extremum at $\mathbf{0}$. The sign of the parameter a determines the definiteness of the function, which correspondingly determines if the extremum is a minimum or a maximum. If $a > 0$, this function is positive definite, and $\mathbf{0}$ is the unique minimizer. If $a < 0$, this function is negative definite, and $\mathbf{0}$ is the unique maximizer.

Compare this to the binary real function $\langle x, y\rangle \mapsto (x + y)^2$. In this case, $\mathbf{0}$ is still a minimizer of q, but so is any point where $x = -y$. This function is nonnegative definite, but it is not positive definite.

25.1.2 Extrema of Quadratic Functions

Theorem 25.1 tells us that unary quadratic forms have their extrema (if any) at 0. This result provides insight to the extrema of any unary real quadratic function. For ease of discussion, parameterize the family of such functions as

$$f = x \mapsto ax^2 + bx + c \tag{25.2}$$

where $a \neq 0$. A function in the family is quadratic but, unless $b = 0$ and $c = 0$, is not a quadratic form. Nevertheless, it implies a similar simple relation between changes in the argument and changes in the function value. Starting with a value x, consider a change in the argument of size h.

$$\begin{aligned} f[x + h] &= a(x + h)^2 + b(x + h) + c \\ &= ax^2 + 2axh + ah^2 + bx + bh + c \\ &= (ax^2 + bx + c) + (2ax + b)h + ah^2 \end{aligned} \tag{25.3}$$

Since the first-order derivative is $f' = x \mapsto 2ax + b$ and the second-order derivative is $f'' = x \mapsto 2a$, write this as

$$f[x + h] = f[x] + f'[x]h + \frac{1}{2}f''[x]\,h^2 \tag{25.4}$$

Find the stationary point $\hat{x}$ by solving $f'[\hat{x}] = 0$, so that $\hat{x} = -b/(2a)$. Then

$$f[\hat{x} + h] - f[\hat{x}] = \frac{1}{2}f''[\hat{x}]\,h^2 \tag{25.5}$$

If $\hat{x}$ is a maximizer, then $f[\hat{x} + h] - f[\hat{x}] \leq 0$ regardless of the value of h. Looking at (25.5), this means that at a maximum, any unary quadratic function must satisfy

$$\frac{1}{2}f''[\hat{x}]\,h^2 \leq 0 \tag{25.6}$$

Since $h^2 > 0$, this in turn requires that $f''[\hat{x}] \leq 0$. Furthermore, $f''[\hat{x}] < 0$ ensures a strict maximum. So $a < 0$ ensures that $\hat{x} = -b/2a$ is a strict global maximizer, regardless of the values of b and c. (The reasoning for a minimizer is analogous.) This is a *global* result for quadratic functions, not just quadratic forms.

Relation to Completing the Square

Recall that a unary polynomial is called ***monic*** if its leading coefficient if its leading coefficient is 1. Parameterize monic quadratic functions as $x \mapsto x^2 + bx + c$ and note that

$$x^2 + bx + c = (x + b/2)^2 + (c - b^2/4) \tag{25.7}$$

Re-expressing the function body in this way is called ***completing the square***. Let $y = x+b/2$ and $k = c - b^2/4$, the defining expression becomes $y^2 + k$, which is clearly minimized at $y = 0$ to produce a value of k. Equivalently, f is minimized at $-b/2$ to produce a value of $c - b^2/4$.

Completing the square of the general quadratic function $x \mapsto ax^2 + bx + c$ requires attention to the sign of a. Note that

$$ax^2 + bx + c = a\big((x + b/2a)^2 + (c/a - (b/2a)^2\big) \tag{25.8}$$

This time, define $y = x + b/(2a)$ and $k = c - (b/2a)^2$, so that the expression $(y^2 + k)$ is minimized at $y = 0$ to produce a value of k. Regardless of the sign of a, the extreme value of the function is $c/a - (b/2a)^2$. But if $a > 0$ then the function is minimized at $-b/2a$, while if $a < 0$ then it is maximized at $-b/2a$.

Example 25.5 (Laffer Curve) The *Laffer curve* postulates that tax revenue is zero if the tax rate is 0% or 100% but is positive in the real interval (0 .. 1). Represent the tax-revenue function $r : [0, 1] \to \mathbb{R}_{\geq 0}$ by the quadratic function $r = t \mapsto \alpha(t - t^2)$, where $\alpha > 0$. Solve for the tax maximizing t by completing the square. That is, rewrite this function as $r = t \mapsto -\alpha\left((t - 1/2)^2 - 1/4\right)$, so the extremum evidently occurs at $t = 1/2$. Since $-\alpha < 0$, this is the unique maximizer of the tax-revenue function.

25.1.3 Generalizing to Other Functions

This observation about the role of the second-order derivative generalizes even further, when a local result suffices. As long as f is continuously twice differentiable, the following *second-order Taylor approximation* of the function provides a good local approximation of the function near x.

$$f[x + h] \approx f[x] + f'[x]h + \frac{1}{2}f''[x]h^2 \tag{25.9}$$

In the quadratic case, as shown above, this is an equality. More generally, this approximiation becomes very good as h becomes small. The standard first-order necessary condition for an extremum, $f'[\hat{x}] = 0$, can thereby be supplemented with second-order conditions. Setting $f'[x] = 0$,

$$f[\hat{x} + h] - f[\hat{x}] \approx \frac{1}{2}f''[\hat{x}]h^2 \tag{25.10}$$

A second-order necessary condition for a local maximum is therefore $f''[x] \leq 0$, and a second-order sufficient condition for a strict local maximum $f''[x] < 0$.

Example 25.6 Let $f = x \mapsto \sin[x]$ so that $f' = x \mapsto \cos[x]$ and $f'' = x \mapsto -\sin[x]$. Let $\hat{x} = \pi/2$ radians, and note that $\cos[\pi/2] = 0$ and $-\sin[\pi/2] = -1$. So $f'[\hat{x}] = 0$ and $f''[\hat{x}] < 0$, satisfying the necessary and sufficient conditions for strict local maximum.

Example 25.7 Let $f = x \mapsto x^4$ so that $f' = x \mapsto 4x^3$ and $f'' = x \mapsto 12x^2$. Let $\hat{x} = 0$, and note that $f'[\hat{x}] = 0$ and $f''[\hat{x}] = 0$. This satisfies the second-order necessary condition for strict local minimum. In fact, 0 is a *global* minimizer of this function. Nevertheless, the second-order sufficient condition is not satisfied.

25.1.4 Matrix Representation

Recall from definition 25.1 that $q[\boldsymbol{x}] = \sum_{r=1}^{N} \sum_{k=1}^{N} a_{r,k} x_r x_k$ describes any quadratic form $\mathbb{R}^N \xrightarrow{q} \mathbb{R}$. Give this the matrix representation $\boldsymbol{x}^\top \boldsymbol{A} \boldsymbol{x}$, where the r,k-th element of $\boldsymbol{A}$ is $a_{r,k}$ and the k-th element of the column vector $\boldsymbol{x}$ is x_k.[1]

Consider the binary quadratic form, $\mathbb{R}^2 \xrightarrow{q} \mathbb{R}$ represented by the matrix $\boldsymbol{A}$.

$$\boldsymbol{A} = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \tag{25.11}$$

If $\boldsymbol{A}$ is a diagonal matrix, so that the off-diagonal elements a_{12} and a_{21} are zero, $q[\boldsymbol{x}]$ is the weighted sum of squares $a_{11}x_1^2 + a_{22}x_2^2$. In the special case where $\boldsymbol{A} = \boldsymbol{I}_2$, $q[\boldsymbol{x}]$ is the simple sum of squares $x_1^2 + x_2^2$, which can never be negative and will only be zero when $\boldsymbol{x} = \boldsymbol{0}$. This is an example of a positive definite quadratic form.

Each quadratic form has a unique representation as a symmetric matrix. This implies that the properties of the quadratic form can be studied as properties of its associated symmetric matrix. For example, as shown later in this lecture, the second-order conditions for extrema have natural matrix characterizations.

[1] By convention, in this context, the result of this matrix multiplication is a scalar instead of a 1×1 matrix.

Example 25.8 Consider any real, symmetric 2×2 matrix

$$\boldsymbol{A} = \begin{bmatrix} a_{11} & a_{12} \\ a_{12} & a_{22} \end{bmatrix}$$

The associated quadratic form is

$$\begin{aligned} q[\boldsymbol{x}] &= \boldsymbol{x}^\top \boldsymbol{A} \boldsymbol{x} \\ &= a_{11} x_1^2 + 2 a_{12} x_1 x_2 + a_{22} x_2^2 \end{aligned}$$

If $\boldsymbol{A} = \boldsymbol{I}_2$ then q is positive at all nonzero vectors and has a unique global minimum at 0. Since q is positive definite in this case, $\boldsymbol{I}_2$ is a positive definite matrix. However if $\boldsymbol{A} = -\boldsymbol{I}_2$, then q is negative at all nonzero vectors and has a unique global maximum at 0. Since q is negative definite in this case, say that $-\boldsymbol{I}_2$ is a negative definite matrix. Figure 25.2 illustrates these two cases.

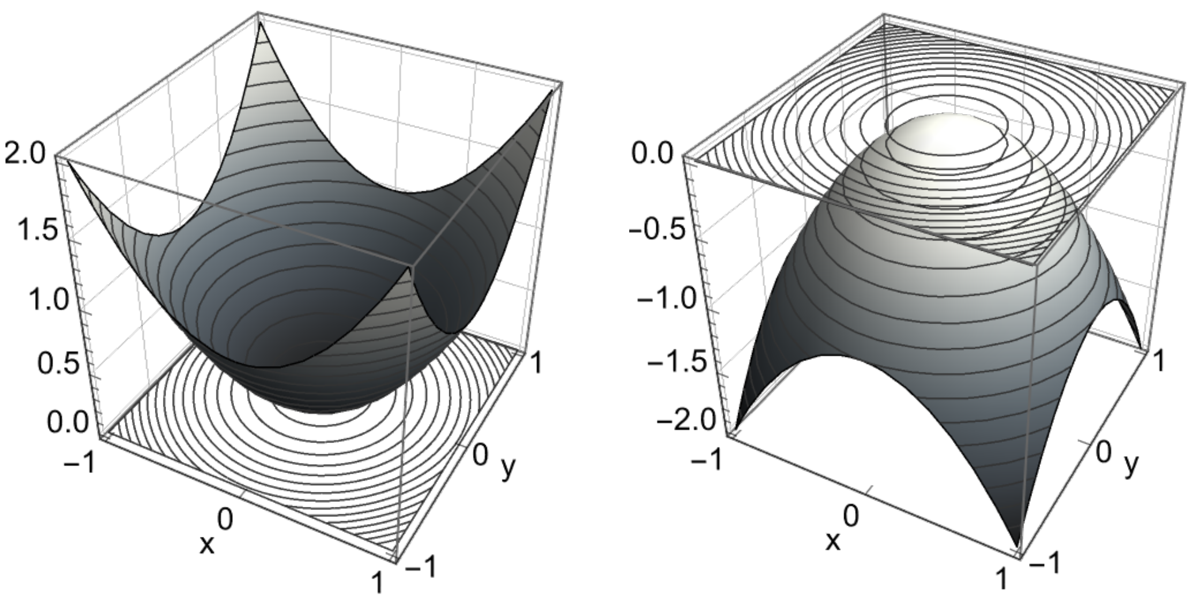

Figure 25.2: Definite Quadratic Forms

Since symmetric matrices have particularly convenient properties, it is conventional to use the symmetric matrix representation of any quadratic form q. To see that such a matrix exists, note that the properties of tranposition imply that $\boldsymbol{x}^\top \boldsymbol{A} \boldsymbol{x} = \boldsymbol{x}^\top \boldsymbol{A}^\top \boldsymbol{x}$. This means that if q can be represented by $\boldsymbol{A}$, it can also be represented by $\boldsymbol{A}^\top$. It also means that any convex combination of these two matrices also provides a matrix representation. Of particular interest is the symmetric representation.

$$q[\boldsymbol{x}] = \boldsymbol{x}^\top \left(\frac{1}{2} \boldsymbol{A} + \frac{1}{2} \boldsymbol{A}^\top \right) \boldsymbol{x} \tag{25.12}$$

If the coefficient matrix $\boldsymbol{A}$ is not symmetric in some setting, it is therefore conventional to represent the quadratic form by the unique *symmetric* matrix $\boldsymbol{Q} = (1/2)(\boldsymbol{A} + \boldsymbol{A}^\top)$.

Example 25.9 Suppose we have

$$\boldsymbol{A} = \begin{bmatrix} 3 & 4 \\ 5 & 6 \end{bmatrix} \qquad \boldsymbol{Q} = (1/2)(\boldsymbol{A} + \boldsymbol{A}^\top) = \begin{bmatrix} 3 & 4.5 \\ 4.5 & 6 \end{bmatrix}$$

Note that $\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x} = \boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x}$. In more detail, let $\boldsymbol{x}^\top = \langle x_1, x_2 \rangle$. Then

$$\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x} = 3x_1^2 + 9x_1x_2 + 6x_2^2 = \boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x}$$

Example 25.10 As a more abstract example, suppose $q[\boldsymbol{x}] = a_{11}x_1^2 + a_{12}x_1x_2 + a_{22}x_2^2$. Let

$$\boldsymbol{A} = \begin{bmatrix} a_{11} & a_{12} \\ 0 & a_{22} \end{bmatrix} \qquad \boldsymbol{Q} = \begin{bmatrix} a_{11} & a_{12}/2 \\ a_{12}/2 & a_{22} \end{bmatrix} \tag{25.13}$$

then $q[\boldsymbol{x}] = \boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x} = \boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x}$, where $\boldsymbol{Q} = (1/2)\big(\boldsymbol{A} + \boldsymbol{A}^\top\big)$.

Recall that from definition 25.1, a quadratic form $q : \mathbb{R}^N \to \mathbb{R}$ can be represented as $q[\boldsymbol{x}] = \sum_{r=1}^N \sum_{k=1}^N a_{r,k}x_r x_k$. Note that the inner summation involves the dot product of the r-th row of $\boldsymbol{A}$ and the vector $\boldsymbol{x}$. That is,

$$\begin{aligned} q[\boldsymbol{x}] &= \sum_{r=1}^N x_r \sum_{k=1}^N a_{r,k}x_k \\ &= \sum_{r=1}^N x_r(\boldsymbol{A}_{r,:} \cdot \boldsymbol{x}) \end{aligned} \tag{25.14}$$

So if $x_r = 0$, the r-th row of $\boldsymbol{A}$ makes no contribution to the result. This makes intuitive sense, because $\boldsymbol{x}^\top \boldsymbol{A}$ is a weighted sum of the rows of $\boldsymbol{A}$, and the r-th row receives zero weight. However, the order of summation is arbitrary, so we could equally well write this as

$$\begin{aligned} q[\boldsymbol{x}] &= \sum_{k=1}^N \big(\sum_{r=1}^N a_{r,k}x_r\big)x_k \\ &= \sum_{k=1}^N (\boldsymbol{x}^\top \cdot \boldsymbol{A}_{:,k})x_k \end{aligned} \tag{25.15}$$

So it is equally true that if the $x_k = 0$, then the k-th column of $\boldsymbol{A}$ contributes nothing to the result. This makes intuitive sense, because $\boldsymbol{A}\boldsymbol{x}$ is a weighted sum of the columms of

$\boldsymbol{A}$, and the k-th column receives zero weight. In short, if $x_i = 0$, then the i-th row and i-th column of $\boldsymbol{A}$ contribute nothing to the result.

Here is another way to make the same observation. As in Lecture 21, let $\boldsymbol{A}_{r,k}$ be the submatrix produced from $\boldsymbol{A}$ by discarding its r-th row and k-th column. Similarly, let $\boldsymbol{x}_{\hat{\imath}}$ be the vector $\boldsymbol{x}$ but with its i-th entry deleted. Then when $x_i = 0$,

$$\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x} = \boldsymbol{x}_{\hat{\imath}}^\top \boldsymbol{A}_{i,i}\boldsymbol{x}_{\hat{\imath}} \tag{25.16}$$

Example 25.11 Consider the following matrix $\boldsymbol{A}$ and vector $\boldsymbol{b}$.

$$\boldsymbol{A} = \begin{bmatrix} 4 & 5 & 9 & 4 \\ 10 & 8 & 7 & 9 \\ 6 & 5 & 1 & 5 \\ 5 & 5 & 7 & 4 \end{bmatrix} \qquad \boldsymbol{x} = \begin{bmatrix} 4 \\ 5 \\ 0 \\ 2 \end{bmatrix}$$

Note that $x_3 = 0$. Therefore produce

$$\boldsymbol{A}_{3,3} = \begin{bmatrix} 4 & 5 & 4 \\ 10 & 8 & 9 \\ 5 & 5 & 4 \end{bmatrix} \qquad \boldsymbol{x}_{\hat{3}} = \begin{bmatrix} 4 \\ 5 \\ 2 \end{bmatrix}$$

Compute $\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x} = 792 = \boldsymbol{x}_{\hat{3}}^\top \boldsymbol{A}_{3,3}\boldsymbol{x}_{\hat{3}}$.

25.2 Definite Matrices

A quadratic form $\mathbb{R}^N \xrightarrow{q} \mathbb{R}$ may have the special property that determing the sign of function value does not require any information about the argument vector. Consider a symmetric matrix $\boldsymbol{Q}$ representing such a quadratic form. Such a matrix will correspondingly have the special property that determing the sign of $\boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x}$ does not require any information about $\boldsymbol{x}$. There, it is conventional to assign to any real symmetric matrix the definiteness of its quadratic form.[2]

- $\boldsymbol{Q}$ is ***nonnegative definite*** iff $\boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x} \geq 0 \quad \forall \boldsymbol{x} \in \mathbb{R}^n$.
- $\boldsymbol{Q}$ is ***nonpositive definite*** iff $\boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x} \leq 0 \quad \forall \boldsymbol{x} \in \mathbb{R}^n$.

If neither of these conditions apply, the matrix is *indefinite*.

A quadratic form q will always have $q[\boldsymbol{0}] = 0$, but some quadratic forms definitely have a strict sign everywhere else. These cases produce refinements of the definiteness definitions, which are substantially easier to test for.

- $\boldsymbol{Q}$ is ***positive definite*** iff $\boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x} > 0 \quad \forall \boldsymbol{x} \in \mathbb{R}^n - \{\boldsymbol{0}\}$.

[2]Some authors use the term ***positive semidefinite*** instead of nonnegative definite and the term ***negative semidefinite*** instead of nonpositive definite.

- $\boldsymbol{Q}$ is ***negative definite*** iff $\boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x} < 0 \quad \forall \boldsymbol{x} \in \mathbb{R}^n - \{\boldsymbol{0}\}$.

A quadratic form is negative definite iff it has a unique maximum at zero, and it is positive definite iff it has a unique minimum at zero.

Theorem 25.2 (Definiteness and Nonsingularity) Positive definite and negative definite matrices are nonsingular.

Proof : Show this by showing that the columns of a positive definite matrix $\boldsymbol{Q}$ are linearly independent. Suppose $\boldsymbol{Q}\boldsymbol{x} = \boldsymbol{0}$ for some $\boldsymbol{x} \neq \boldsymbol{0}$; then $\boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x} = 0$ for some $\boldsymbol{x} \neq \boldsymbol{0}$ and $\boldsymbol{Q}$ is not positive definite. The proof for negative definite matrices is the same. ■

Theorem 25.3 (Definite Inverse) A real symmetric matrix is positive (negative) definite iff it has a positive (negative) definite inverse.

Proof : To show this for a symmetric positive definite matrix $\boldsymbol{Q}$, note that $\boldsymbol{Q}^{-1} = \boldsymbol{Q}^{-1}\boldsymbol{Q}\boldsymbol{Q}^{-1}$. So

$$\boldsymbol{x}^\top \boldsymbol{Q}^{-1}\boldsymbol{x} = \boldsymbol{x}^\top \boldsymbol{Q}^{-1}\boldsymbol{Q}\boldsymbol{Q}^{-1}\boldsymbol{x} = \boldsymbol{y}^\top \boldsymbol{Q}\boldsymbol{y}$$

where $\boldsymbol{y} = \boldsymbol{Q}\boldsymbol{x}$. ■

Recall from Lecture 21 that nonsingular matrices have nonzero determinants. In addition, we can sign the determinate of definite matrices. One way to see this is to observe that a weighted average of positive definite matrices must be positive definite. So for any positive definite matrix $\boldsymbol{Q}$, any weighted average of the identity matrix and $\boldsymbol{Q}$ must be positive definite.

Theorem 25.4 A positive definite matrix has a positive determinant.

Proof : Define the function $f : [0 \,..\, 1] \to \mathbb{R}$ by $f = \lambda \mapsto \det[(1-\lambda)\boldsymbol{I} + \lambda\boldsymbol{Q}]$. By definition of the determinate, f is a polynomial in λ, so it is a continuous function. Note $f[0] = \det \boldsymbol{I} = 1$. By the intermediate value theorem, for $f[1] = \det \boldsymbol{Q}$ to be negative, there would have to be some value of λ in $[0 \,..\, 1]$ that produces $f[\lambda] = 0$. However, this is impossible, because a weighted average of positive definite matrices is positive definite, and a positive definite matrix is nonsingular and therefore must have a nonzero determinant. ■

Example 25.12 If $\boldsymbol{A} = \begin{bmatrix} 2 & 1 \\ 1 & 1 \end{bmatrix}$ then $(1-\lambda)\boldsymbol{I} + \lambda\boldsymbol{A} = \begin{bmatrix} 1+\lambda & \lambda \\ \lambda & 1 \end{bmatrix}$, which has the determinant $1 + \lambda - \lambda^2$. This polynomial in λ is positive for $\lambda \in [0 \,..\, 1]$.

Exercise 25.1 If $\boldsymbol{Q}$ is negative definite, what is the sign of its determinant?

25.2.1 Testing for Definiteness

It is evidently not practical to test a matrix $\boldsymbol{Q}$ for definiteness by systematically computing the sign of $\boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x}$ for all the possible values of $\boldsymbol{x}$. Fortunately, there are other ways to test for definiteness. This is especially simple for diagonal matrices. A diagonal matrix is positive definite iff all its diagonal elements are positive, and it is negative definite iff all its diagonal elements are negative.

If $\boldsymbol{A}$ is a *diagonal* matrix, then $\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x} = \sum_{n=1}^{N} a_{n,n} x_n^2$. This is a weighted sum of squares. For a definite matrix, the sign of this expression does not depend on the value of $\boldsymbol{x}$. So to be positive definite, a diagonal matrix must have only positive diagonal elements. This is also sufficient for the matrix to be positive definite, since it assures all positive weights in the weighted sum of squares. Similarly, to be negative definite, the diagonal matrix must have only negative diagonal elements, and this is sufficent for the matrix to be negative definite.

Example 25.13 Consider the quadratic form represented by the identity matrix, $\boldsymbol{I}_N$. In this case $\boldsymbol{x}^\top \boldsymbol{I}\boldsymbol{x} = \boldsymbol{x}^\top \boldsymbol{x}$ is the sum of squared elements, which is always positive for nonzero $\boldsymbol{x}$.

Extend this line of argument to any real symmetric matrix $\boldsymbol{A}$. Let $q = \boldsymbol{x} \mapsto \boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x}$ be the associated quadratic form, and let $\boldsymbol{e}^i$ be the i-th standard unit vector. Note that $q[\boldsymbol{e}^i] = a_{i,i}$. Therefore $a_{i,i} > 0$ is necessary for $\boldsymbol{A}$ to be positive definite. Similarly, all diagonal elements must be negative for $\boldsymbol{A}$ to be negative definite. A constant sign on elements on the diagonal is a necessary condition for definiteness. A quick examination of the diagonal is therefore helpful when determining the definiteness of matrices: any sign variation along the diagonal implies that the matrix is indefinite. However, as illustrated by the following example, sign consistency along the diagonal is not sufficient for definiteness.

Example 25.14

$$\begin{bmatrix} 1 & 1 \end{bmatrix} \begin{bmatrix} 1 & 2 \\ 2 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 1 \end{bmatrix} = 6 \qquad \begin{bmatrix} -1 & 1 \end{bmatrix} \begin{bmatrix} 1 & 2 \\ 2 & 1 \end{bmatrix} \begin{bmatrix} -1 \\ 1 \end{bmatrix} = -2$$

Aside from diagonal matrices, finding sufficient conditions for definiteness therefore requires more work. Equation (25.16) implies that $\boldsymbol{A}_{i,i}$ is positive definite whenever $\boldsymbol{A}$ is positive definite. This in turn implies that every every principal submatrix of a positive definite matrix is itself positive definite. Since positive definite matrices have positive determinants, all of the principal minors of a positive definite matrix are positive. In fact, positive leading principal minors also ensures positive definiteness.

Definition 25.2 A k-th order ***principal submatrix*** of the matrix $\boldsymbol{Q}_{N\times N}$ is any $k \times k$ submatrix of $\boldsymbol{Q}$ created by deleting all but k rows and the k *corresponding* columns. (That is, retaining row i implies retaining column i.) A k-th order ***principal minor*** is the determinant of a k-th order principal submatrix. A k-th order ***leading*** principal submatrix is formed by keeping only the *first* k rows and columns, and its determinant is called the k-th order ***leading*** principal minor.[3]

[3]Among economists this terminology is not completely standard. For example Chiang and Wainwright (2004, p.322) use the term *principal minor* to denote only leading principal minors.

Remark 25.1 An $N \times N$ matrix will have $C[N,k]$ principal submatrices of order k.

Example 25.15 A 3×3 matrix has one principle submatrix of order 1. three principle submatrices of order 2, and three principle submatrices of order 1.

Theorem 25.5 A real symmetric matrix $\boldsymbol{Q}$ is positive definite iff its leading principal minors are positive. A real symmetric matrix $\boldsymbol{Q}$ is negative definite iff $a_{11} < 0$ and its leading principal minors alternate sign.

Proof: See Simon and Blume (1994, ch.16.4). ■

Testing for semidefiniteness is unfortunately more complicated: it involves looking at all the principal minors.

Theorem 25.6 A real, symmetric matrix $\boldsymbol{Q}$ is nonnegative-definite iff each of its principal minors is nonnegative.

Proof: Necessity follows from the earlier discussion. For sufficiency, see Simon and Blume (1994, ch.16.4). ■

It follows that a real, symmetric matrix $\boldsymbol{Q}$ is nonpositive-definite iff each of its k-th order principal minors has sign $(-1)^k$ or is zero.

Exercise 25.2 Consider the following two symmetric matrices.

$$\boldsymbol{A} = \begin{bmatrix} 0 & 0 \\ 0 & a_{22} \end{bmatrix} \qquad \boldsymbol{B} = \begin{bmatrix} a & a \\ a & a \end{bmatrix}$$

Can $\boldsymbol{A}$ be positive definite? Can $\boldsymbol{A}$ be nonnegative definite? Can $\boldsymbol{B}$ be positive definite? Can $\boldsymbol{B}$ be nonnegative definite?

Diagonal matrices once again assist with insight. Let $\boldsymbol{A}$ be a diagonal matrix. Then

$$q[\boldsymbol{x}] = \boldsymbol{x}^\top \boldsymbol{Q} \boldsymbol{x} = \sum_{i=1}^{N} a_{i,i} x_i^2 \tag{25.17}$$

That is, $q[\boldsymbol{x}]$ is a weighted sum of squares. Note that $q[e^i] = a_{i,i}$, therefore a necessary condition for $\boldsymbol{A}$ to be nonnegative definite is that $a_{i,i} \geq 0$. This is also sufficient, for a positively weighted sum of squares must be positive.

Nonnegative diagonal elements are also the necessary and sufficient condition for all principal minors to be either positive or zero, since each of these evaluate to a product of k elements from the diagonal of $\boldsymbol{A}$. (Remember, $\boldsymbol{A}$ is diagonal.) This ensures that the leading principal minors are nonnegative, but it does not conversely ensure that examining them is sufficient to determine definiteness. For example, if $a_{11} = 0$ then all the leading principal minors evaluate to zero. To be sure no element on the diagonal is negative, we may need to examine all the first order principal minors (i.e., all the elements on the diagonal).

The reasoning for nonpositive definite diagonal matrices is parallel. Since $q[\boldsymbol{e}^i]a_{i,i}$, a necessary condition for the diagonal matrix $\boldsymbol{A}$ to be nonpositive definite is that $a_{i,i} \leq 0$. This is also sufficient, for a negatively weighted sum of squares must be negative. This is also the necessary and sufficient condition for all principal minors to be either negative or zero, since each of these evaluate to a product of k elements from the diagonal of $\boldsymbol{Q}$. But again, while this condition ensures that the leading principal minors are nonpositive, it does not conversely ensure that examining them is sufficient to determine definiteness. For example, if $a_{1,1} = 0$ then all the leading principal minors evaluate to zero. To be sure no element on the diagonal is positive, we may need to examine all the first order principal minors (i.e., all the elements on the diagonal).

25.3 Linear Constraints

The previous section focuses on the characteristics of quadratic forms on $\mathbb{R}^N$. The characteristics on a subspace of $\mathbb{R}^N$ can differ in useful ways. This section characterizes the subspace by a linear constraint: for some matrix $\boldsymbol{G}_{R\times N}$, the constraint is that $\boldsymbol{G}\boldsymbol{x} = \boldsymbol{0}$. (Note that $\boldsymbol{x} = \boldsymbol{0}$ always satisfies the constraint, so the constraint set is not empty.) While the previous section demonstrates how to characterize the definiteness of a quadratic form on $\mathbb{R}^N$, it can be a bit more work to characterize its behavior on the subspace $\{\boldsymbol{x}|\boldsymbol{G}\boldsymbol{x} = \boldsymbol{0}\}$.

Of course if the quadratic form is definite on a vector space, it is also definite on any subspace. But a quadratic form that is indefinite on a vector space may nevertheless be definite on a subspace. For example, $x \mapsto x^3$ is positive definite on $\mathbb{R}_{\geq 0}$.

Example 25.16 Consider the quadratic form $\mathbb{X} \xrightarrow{q} \mathbb{R}$ defined by $q = \langle x_1, x_2\rangle \mapsto x_1^2 - x_2^2$, which is indefinite on the domain $\mathbb{R}^2$. Correspondingly $\boldsymbol{Q} = \left[\begin{smallmatrix}1 & 0\\ 0 & -1\end{smallmatrix}\right]$ is indefinite. Define the constraint set to be $\mathbb{X} = \{\boldsymbol{x} \in \mathbb{R}^2 \mid \langle 1, 0\rangle \cdot \boldsymbol{x} = 0\}$. On this restricted domain, $x_1 = 0$. Note that $q[0, x_2] = -x_2^2$ is negative definite on the constraint set.

25.3.1 Testing for Definiteness

Suppose the quadratic form q, characterized by the real symmetric matrix $\boldsymbol{Q}$, is defined on the subspace $\mathbb{X} = \{\boldsymbol{x} \in \mathbb{R}^N \mid \boldsymbol{G}\boldsymbol{x} = \boldsymbol{0}\}$, where G is an $R \times N$ matrix with rank R. Construct the bordered matrix $\overline{\boldsymbol{Q}}$ as follows:

$$\overline{\boldsymbol{Q}} = \begin{bmatrix} \boldsymbol{0} & \boldsymbol{G} \\ \boldsymbol{G}^\top & \boldsymbol{Q} \end{bmatrix} \tag{25.18}$$

The bordered matrix $\overline{\boldsymbol{Q}}$ will characterize the definiteness of the quadratic form q on this subspace (Hassell and Rees, 1993). Once again the test for definiteness examines the leading principal minors. Obviously the first R leading principal minors of $\overline{\boldsymbol{Q}}$ evaluate to 0, so any information will be contained in the rest. It turns out that to test for definiteness we can discard the next R leading principal minors as well, so that we only need to consider the last $N - R$ leading principal minors.

The quadratic form q is postive definite on the restricted domain iff $(-1)^R \left|\overline{\boldsymbol{Q}}\right| > 0$, and the last $N - R$ leading principal minors all share the sign of $\overline{\boldsymbol{Q}}$. (Here R is the number of constraints, so for a single constraint, this means these leading principal minors must all be negative.) The quadratic form q is negative definite on its domain if $(-1)^N \left|\overline{\boldsymbol{Q}}\right| > 0$, and the last $N - R$ leading principal minors alternate in sign. (Here N is the number of variables.) For a single constraint, $\left|\overline{\boldsymbol{Q}}\right|$ is the $N + 1$-th leading principal minor, so in this case these leading principal minors must have the following characteristic: a k-th order leading principal minors must be positive when k is odd and negative when k is even.)

This section focuses on on the single-constraint case. (For more generality, see Debreu (1952), which also gives conditions for semidefiniteness.) Let $\mathbb{R}^2 \xrightarrow{q} \mathbb{R}$ be a quadratic form represented by the symmetric matrix $\boldsymbol{A}_{2\times2}$. Given a vector $\boldsymbol{g} = \langle g_1, g_2\rangle$, construct the constraint set $\mathbb{X}$ as the the nullspace of this vector (null[$\boldsymbol{g}$]). That is, $\mathbb{X} = \{\langle x_1, x_2\rangle \in \mathbb{R}^2 \mid g_1x_1 + g_2x_2 = 0\}$, or equivalently, $\mathbb{X} = \{\boldsymbol{x} \in \mathbb{R}^2 \mid \boldsymbol{g} \cdot \boldsymbol{x} = 0\}$. As long as $\boldsymbol{g} \neq \boldsymbol{0}$, the constraint set comprises the points along a line. Characterize these points parametrically as

$$\mathbb{X} = \{t \cdot \langle g_2, -g_1\rangle \mid t \in \mathbb{R}\} \tag{25.19}$$

When restricted to this constraint set, the quadratic form now involves the single variable t.

$$q_r = t^2\langle g_2, -g_1\rangle^\top \boldsymbol{A}\langle g_2, -g_1\rangle \tag{25.20}$$

Naturally enough, the constrained problem has a unary expression, since any $\boldsymbol{x} \in \mathbb{X}$ lies on a line through the origin. The result may be considered as an unconstrained unary quadratic form, subject to the earlier analysis of this form. For example, this unary quadratic form is positive definite iff $\langle g_2, -g_1\rangle^\top \boldsymbol{A}\langle g_2, -g_1\rangle > 0$. Expanding this expression and then imposing symmetry produces

$$\begin{aligned}\langle g_2, -g_1\rangle^\top \boldsymbol{A}\langle g_2, -g_1\rangle &= g_2^2 a_{1,1} - g_1g_2a_{1,2} - g_1g_2a_{2,1} + g_1^2 a_{2,2} \\ &= g_2^2 a_{1,1} - 2g_1g_2a_{1,2} + g_1^2 a_{2,2}\end{aligned} \tag{25.21}$$

Algebraically, this expression negates the determinant of the following bordered matrix.

$$\overline{\boldsymbol{Q}} = \begin{bmatrix} 0 & g_1 & g_2 \\ g_1 & a_{1,1} & a_{1,2} \\ g_2 & a_{1,2} & a_{2,2} \end{bmatrix}$$

That is

$$\begin{aligned}\left|\overline{\boldsymbol{Q}}\right| &= -g_1(g_1 a_{2,2} - a_{1,2} g_2) + g_2(g_1 a_{1,2} - a_{1,1} g_2)\\ &= -(a_{1,1} g_2^2 - 2a_{1,2} g_1 g_2 + a_{2,2} g_1^2)\\ &= -\langle g_2, -g_1\rangle^\top \boldsymbol{A}\langle g_2, -g_1\rangle\end{aligned}$$

So if $\left|\overline{\boldsymbol{Q}}\right| < 0$ then our quadratic form is positive definite, whereas if $\left|\overline{\boldsymbol{Q}}\right| > 0$ then our quadratic form is negative definite. Checking the value of this determinant is therefore a common and convenient check of the second-order conditions for an extremum.

Perhaps surprisingly, this procedure works in higher dimeions. For example, let $\mathbb{R}^3 \xrightarrow{q} \mathbb{R}$ be a quadratic form represented by the symmetric matrix $\boldsymbol{A}_{3\times 3}$. Given a vector $\boldsymbol{g} = \langle g_1, g_2, g_3\rangle$, once again construct the constraint set X as the nullspace of the vector (null[$\boldsymbol{g}$]). That is, $\mathrm{X} = \{\langle x_1, x_2, x_3\rangle \in \mathbb{R}^3 \mid g_1 x_1 + g_2 x_2 + g_3 x_3 = 0\}$, or equivalently, $\mathrm{X} = \{\boldsymbol{x} \in \mathbb{R}^2 \mid \boldsymbol{g} \cdot \boldsymbol{x} = 0\}$. As long as $\boldsymbol{g} \neq \boldsymbol{0}$, the constraint set comprises the points of a plane. Consider the determinant of a bordered symmetric matrix.

$$\overline{\boldsymbol{Q}} = \begin{bmatrix} 0 & g_1 & g_2 & g_3 \\ g_1 & a_{1,1} & a_{1,2} & a_{1,3} \\ g_2 & a_{2,1} & a_{2,2} & a_{2,3} \\ g_3 & a_{3,1} & a_{3,2} & a_{3,3} \end{bmatrix}$$

Expand across the first row to get

$$\begin{aligned}\left|\overline{\boldsymbol{Q}}\right| =& g_1^2(a_{2,3}a_{3,2} - a_{2,2}a_{3,3}) + g_2^2(a_{1,3}a_{3,1} - a_{1,1}a_{3,3}) + g_3^2(a_{1,2}a_{2,1} - a_{1,1}a_{2,2})\\ &+ g_3 g_1(a_{1,3}a_{2,2} - a_{1,2}a_{2,3}) + g_3 g_1(a_{2,2}a_{3,1} - a_{2,1}a_{3,2})\\ &+ g_2 g_1(a_{1,2}a_{3,3} - a_{1,3}a_{3,2}) + g_2 g_1(a_{2,1}a_{3,3} - a_{2,3}a_{3,1})\\ &+ g_2 g_3(a_{1,1}a_{2,3} - a_{1,3}a_{2,1}) + g_2 g_3(a_{1,1}a_{3,2} - a_{1,2}a_{3,1})\end{aligned}$$

Assume $g_1 \neq 0$, which means our constraint matrix has rank 1. We need to consider the last $N - R = 2 - 1 = 1$ leading principal minors of $\overline{\boldsymbol{Q}}$. Therefore compute

We can readily see that this result makes good sense. Solve our constraint for $x_1 = -(g_2/g_1)x_2$ and substitute this into our quadratic form to get

$$\begin{aligned}q[\boldsymbol{x}] &= \boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x}\\ &= a_{1,1}x_1^2 + 2a_{1,2}x_1 x_2 + a_{2,2}x_2^2\\ &= a_{1,1}\big((g_2/g_1)x_2\big)^2 - 2a_{1,2}(g_2/g_1)x_2 x_2 + a_{2,2}x_2^2\\ &= \big(a_{1,1}(g_2/g_1)^2 - 2a_{1,2}(g_2/g_1) + a_{2,2}\big)x_2^2\\ &= \big(a_{1,1}g_2^2 - 2a_{1,2}g_2 g_1 + a_{2,2}g_1^2\big)x_2^2/g_1^2\end{aligned}$$

The expression in parentheses is opposite in sign to our condition for definiteness.

Exercise 25.3 Consider a real symmetric matrix $\boldsymbol{Q}_{N\times N}$ and a linear constraint set $S = \{x \in \mathbb{R}^N | g^\top x = 0\}$. Recall that S is a subspace of $\mathbb{R}^N$. Let the matrix $\overline{\boldsymbol{Q}}_{N\times N}$ map $\mathbb{R}^N$ to that subspace, and assume the quadratic form represented by $\boldsymbol{Q}$ is positive definite on this subspace. Let $x \in \mathbb{R}^N$ be any nonzero vector. Find a problem in the following reasoning.

For any nonzero vector $\boldsymbol{y} \in S$ we know $\boldsymbol{y}^\top \boldsymbol{Q}\boldsymbol{y} > 0$. For any vector $x \in \mathbb{R}$ we know $\overline{\boldsymbol{Q}}\boldsymbol{x} \in S$. Consider $x \neq \boldsymbol{0}$, so that $(\overline{\boldsymbol{Q}}\boldsymbol{x})^\top \boldsymbol{Q}\overline{\boldsymbol{Q}}\boldsymbol{x} > 0$. But $(\overline{\boldsymbol{Q}}\boldsymbol{x})^\top \boldsymbol{Q}\overline{\boldsymbol{Q}}\boldsymbol{x} = \boldsymbol{x}^\top \overline{\boldsymbol{Q}}^\top \boldsymbol{Q}\overline{\boldsymbol{Q}}\boldsymbol{x}$. Define $\boldsymbol{C} = \overline{\boldsymbol{Q}}^\top \boldsymbol{Q}\overline{\boldsymbol{Q}}$. Since for any nonzero $\boldsymbol{x}$ we have $\boldsymbol{x}^\top \boldsymbol{C}\boldsymbol{x} > 0$, we can conclude that $\boldsymbol{C}$ is positive definite.

25.4 Derivatives

Recall that definition 25.1 implies that a quadratic form $q : \mathbb{R}^N \to \mathbb{R}$ can be represented as $q[\boldsymbol{x}] = \sum_{r=1}^N \sum_{k=1}^N a_{r,k} x_r x_k$. Produce the first-order partial derivative with respect to x_k, as usual.

$$\frac{\partial q[\boldsymbol{x}]}{\partial x_k} = \sum_{r=1}^N a_{kr} x_r + \sum_{r=1}^N a_{rk} x_r \tag{25.22}$$

Based on (25.22), a column vector of first-order partial derivatives can be written as follows.

$$\frac{\partial q[\boldsymbol{x}]}{\partial \boldsymbol{x}} = (\boldsymbol{A} + \boldsymbol{A}^\top)\boldsymbol{x} \tag{25.23}$$

Similarly, a row vector of first-order partial derivatives can be written as

$$\frac{\partial q[\boldsymbol{x}]}{\partial \boldsymbol{x}^\top} = \boldsymbol{x}^\top(\boldsymbol{A}^\top + \boldsymbol{A}) \tag{25.24}$$

This row vector is the ***gradient*** vector of the quadratic form. When $\boldsymbol{A}$ is symmetric, the vectors of partial derivatives simplify to

$$\frac{\partial q[\boldsymbol{x}]}{\partial \boldsymbol{x}} = 2\boldsymbol{A}\boldsymbol{x} \qquad \frac{\partial q[\boldsymbol{x}]}{\partial \boldsymbol{x}^\top} = 2\boldsymbol{x}^\top\boldsymbol{A} \tag{25.25}$$

Next consider the second-order partial derivatives. For aribtrary variables x_r and x_k,

$$\frac{\partial^2 q[\boldsymbol{x}]}{\partial x_r \partial x_k} = a_{rk} + a_{kr} \tag{25.26}$$

By inspection, a matrix of second-order partial derivatives can be written as

$$\frac{\partial^2 q[\boldsymbol{x}]}{\partial \boldsymbol{x}^\top \partial \boldsymbol{x}} = \boldsymbol{A}^\top + \boldsymbol{A} \tag{25.27}$$

This is the ***hessian matrix*** of the quadratic form. When $\boldsymbol{A}$ is symmetric this simplifies to

$$\frac{\partial^2 q[\boldsymbol{x}]}{\partial \boldsymbol{x}^\top \partial \boldsymbol{x}} = 2\boldsymbol{A} \tag{25.28}$$

25.4.1 Derivatives of the Binary Quadratic Form

If $\boldsymbol{A}_{2\times 2}$ represents the quadratic form q, then the first-order derivatives are

$$\frac{\partial q}{\partial \boldsymbol{x}} = \begin{bmatrix} \frac{\partial q}{\partial x_1} \\ \frac{\partial q}{\partial x_2} \end{bmatrix} = \begin{bmatrix} 2a_{11}x_1 + (a_{12} + a_{21})x_2 \\ 2a_{22}x_2 + (a_{12} + a_{21})x_1 \end{bmatrix} = (\boldsymbol{A} + \boldsymbol{A}^\top)\boldsymbol{x} \tag{25.29}$$

When $\boldsymbol{A}$ is symmetric, as we typically assume, this becomes

$$\frac{\partial q}{\partial \boldsymbol{x}} = 2\begin{bmatrix} a_{11}x_1 + a_{12}x_2 \\ a_{22}x_2 + a_{12}x_1 \end{bmatrix} = 2\boldsymbol{A}\boldsymbol{x} \tag{25.30}$$

The second-order derivatives are

$$\frac{\partial^2 q}{\partial \boldsymbol{x}^\top \partial \boldsymbol{x}} = \begin{bmatrix} \frac{\partial^2 q}{\partial x_1^2} & \frac{\partial^2 q}{\partial x_1 \partial x_2} \\ \frac{\partial^2 q}{\partial x_2 \partial x_1} & \frac{\partial^2 q}{\partial x_2^2} \end{bmatrix} = \begin{bmatrix} 2a_{11} & (a_{12} + a_{21}) \\ (a_{12} + a_{21}) & 2a_{22} \end{bmatrix} = (\boldsymbol{A}^\top + \boldsymbol{A}) \tag{25.31}$$

When $\boldsymbol{A}$ is symmetric, this becomes

$$\frac{\partial^2 q}{\partial \boldsymbol{x}^\top \partial \boldsymbol{x}} = 2\begin{bmatrix} a_{11} & a_{12} \\ a_{12} & a_{22} \end{bmatrix} = 2\boldsymbol{A} \tag{25.32}$$

25.4.2 Derivatives of a General Quadratic Functional: Bivariate Case

A general representation of a quadratic polynomial is $\boldsymbol{x} \mapsto \boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x} + \boldsymbol{b}^\top \boldsymbol{x} + c$. In the bivariate case,

$$\boldsymbol{A} = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \quad \boldsymbol{b} = \begin{bmatrix} b_1 \\ b_2 \end{bmatrix} \quad \boldsymbol{x} = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

Let f be this bivariate quadratic function:

$$f = \langle x_1, x_2 \rangle \mapsto a_{11}x_1^2 + (a_{12} + a_{21})x_1x_2 + a_{22}x_2^2 + b_1x_1 + b_2x_2 + c \tag{25.33}$$

Write the column vector of first-order partial derivatives as

$$\begin{aligned} f'[\boldsymbol{x}] = \frac{\partial f[\boldsymbol{x}]}{\partial \boldsymbol{x}} &= \begin{bmatrix} \partial f[\boldsymbol{x}]/\partial x_1 \\ \partial f[\boldsymbol{x}]/\partial x_2 \end{bmatrix} \\ &= \begin{bmatrix} 2a_{11}x_1 + (a_{12} + a_{21})x_2 + b_1 \\ 2a_{22}x_2 + (a_{12} + a_{21})x_1 + b_2 \end{bmatrix} \\ &= (\boldsymbol{A} + \boldsymbol{A}^\top)\boldsymbol{x} + \boldsymbol{b} \end{aligned} \tag{25.34}$$

Imposing symmetry on $\boldsymbol{A}$ produces $f' = \boldsymbol{x} \mapsto 2\boldsymbol{A}\boldsymbol{x} + \boldsymbol{b}$. In other words, finding the stationary points is a matter of solving a linear system of equations. When $\boldsymbol{A}$ is symmetrica, this system has the form $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}'$, where $\boldsymbol{b}' = -(1/2)\boldsymbol{b}$. This is familiar territory.

Similarly

$$\begin{aligned} f''[\boldsymbol{x}] \overset{\text{def}}{=} \frac{\partial^2 f[\boldsymbol{x}]}{\partial \boldsymbol{x} \partial \boldsymbol{x}^\top} &= \begin{bmatrix} \partial^2 f[\boldsymbol{x}]/\partial x_1^2 & \partial^2 f[\boldsymbol{x}]/\partial x_1 \partial x_2 \\ \partial^2 f[\boldsymbol{x}]/\partial x_2 \partial x_1 & \partial^2 f[\boldsymbol{x}]/\partial x_2^2 \partial f[\boldsymbol{x}]/\partial x_2 \end{bmatrix} \\ &= \begin{bmatrix} 2a_{11} & a_{12} + a_{21} \\ a_{12} + a_{21} & 2a_{22} \end{bmatrix} \\ &= (\boldsymbol{A}^\top + \boldsymbol{A}) \end{aligned} \tag{25.35}$$

Imposing symmetry on $\boldsymbol{A}$ produces $f'' = \boldsymbol{x} \mapsto 2\boldsymbol{A}$.

Now suppose we have a change in the argument of size h.

$$\begin{aligned} f[\boldsymbol{x} + \boldsymbol{h}] =& a_{11}(x_1 + h_1)^2 + (a_{12} + a_{21})(x_1 + h_1)(x_2 + h_2) \\ &+ a_{22}(x_2 + h_2)^2 + b_1(x_1 + h_1) + b_2(x_2 + h_2) + c \end{aligned} \tag{25.36}$$

so that

$$\begin{aligned} f[\boldsymbol{x} + \boldsymbol{h}] - f[\boldsymbol{x}] =& a_{11}(2x_1 h_1 + h_1^2) + (a_{12} + a_{21})(x_1 h_2 + x_2 h_1 + h_1 h_2) \\ &+ a_{22}(2x_2 h_2 + h_2^2) + b_1 h_1 + b_2 h_2 \\ =& \begin{bmatrix} x_1 & x_2 \end{bmatrix} \begin{bmatrix} 2a_{11} & (a_{12} + a_{21}) \\ (a_{12} + a_{21}) & 2a_{22} \end{bmatrix} \begin{bmatrix} h_1 \\ h_2 \end{bmatrix} + \begin{bmatrix} b_1 & b_2 \end{bmatrix} \begin{bmatrix} h_1 \\ h_2 \end{bmatrix} \\ &+ \frac{1}{2} \begin{bmatrix} h_1 & h_2 \end{bmatrix} \begin{bmatrix} 2a_{11} & a_{12} + a_{21} \\ a_{12} + a_{21} & 2a_{22} \end{bmatrix} \begin{bmatrix} h_1 \\ h_2 \end{bmatrix} \\ =& \boldsymbol{x}^\top (\boldsymbol{A}^\top + \boldsymbol{A}) \boldsymbol{h} + \boldsymbol{b}^\top \boldsymbol{h} + \frac{1}{2} \boldsymbol{h}^\top (\boldsymbol{A}^\top + \boldsymbol{A}) \boldsymbol{h} \\ =& \frac{\partial f}{\partial \boldsymbol{x}^\top} \boldsymbol{h} + \frac{1}{2} \boldsymbol{h}^\top \frac{\partial^2 f}{\partial \boldsymbol{x} \partial \boldsymbol{x}^\top} \boldsymbol{h} \end{aligned} \tag{25.37}$$

Once again we can address the question of extrema. Consider the case of a maximum. Regardless of the direction $\boldsymbol{h}$, we must get a no increase in f. This requires the gradient vector be 0. This is a first-order necessary conditition. It also requires the hessian matrix be nonpositive definite. This is a second-order necessary conditition. If in addition the hessian matrix is negative definite, we have a global maximum of this quadratic function.

Why Do the Interactions Matter?

The first order conditions appear natural: they just carry order the intuition from the unary case. That is, they require that in each considered direction, the function is (locally) flat. This is obviously a necessary condition to be at an extremum.

Unfortunately, similar reasoning for the curvature does not take us far enough. It is true for example that if a function f is be convex at a minimizer, it must be convex in each variable individually. However, the way the variables interact matters as well and cannot be ignored.

To illustrate, consider the quadratic form associated with the matrix $\begin{bmatrix} 1.0 & -1.5 \\ -1.5 & 1.0 \end{bmatrix}$. That is, $q = \langle x, y\rangle \mapsto x^2 + y^2 - 3xy$. The first order necessary condition of a zero gradient is satisfied at $\langle 0, 0\rangle$. Additionally, note that $q[x, 0] = x^2$ and $q[0, y] = y^2$, so at the stationary point q is convex in each variable individually. This implies that each associated unary function is minimized at the origin, with a value of 0. Nevertheless, $q[1, 1] = -1 < 0$, so q is not minimized at the origin. In fact, $q[x, x] = -x^2$, so we can clearly make the function value as negative as desired. This additional exploration shows that 0 is a saddle point.

Recommended Reading

Simon and Blume (1994, ch.16)

Problems for Review

Exercise 25.4 Suppose we have the matrices

$$\boldsymbol{x} = \begin{bmatrix} 2 \\ 4 \end{bmatrix} \qquad \boldsymbol{A} = \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix}$$

Calculate $\boldsymbol{x}^\top \boldsymbol{A}\boldsymbol{x}$.

Exercise 25.5 Consider a 3×3 matrix $\boldsymbol{A} = [a_{ij}]$. Construct all the principal minors of each order. Which are leading principal minors?

Exercise 25.6 Find the definiteness of each of the following matrices.

$$\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \quad \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} \quad \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$$

$$\begin{bmatrix} -1 & 0 \\ 0 & -1 \end{bmatrix} \quad \begin{bmatrix} -1 & 0 \\ 0 & 0 \end{bmatrix} \quad \begin{bmatrix} 0 & 0 \\ 0 & -1 \end{bmatrix} \quad \begin{bmatrix} -1 & 0 \\ 0 & 1 \end{bmatrix}$$

Computational Exercise 25.1 Create a function `innerprod[x,y]` that computes the dot product of two N-vectors by iterating over the elements. Compare with the built in functionality for producing dot products by matrix multiplication.