Exercise 8.9.4 Consider the equation $ax^2 + bxy + cy^2 = d$, where $b \neq 0$. Introduce new variables x_1 and y_1 by rotating the axes counterclockwise through an angle θ. Show that the resulting equation has no x_1y_1-term if θ is given by

$$\begin{aligned} \cos 2\theta &= \frac{a-c}{\sqrt{b^2+(a-c)^2}} \\ \sin 2\theta &= \frac{b}{\sqrt{b^2+(a-c)^2}} \end{aligned}$$

[*Hint*: Use equation (8.8) preceding Theorem 8.9.2 to get x and y in terms of x_1 and y_1, and substitute.]

Exercise 8.9.5 Prove properties (1)–(5) preceding Example 8.9.4.

Exercise 8.9.6 If $A \overset{c}{\sim} B$ show that A is invertible if and only if B is invertible.

Exercise 8.9.7 If $\mathbf{x} = (x_1, \dots, x_n)^T$ is a column of variables, $A = A^T$ is $n \times n$, B is $1 \times n$, and c is a constant, $\mathbf{x}^T A\mathbf{x} + B\mathbf{x} = c$ is called a **quadratic equation** in the variables x_i.

a. Show that new variables $y_1, \dots, y_n$ can be found such that the equation takes the form

$$\lambda_1 y_1^2 + \cdots + \lambda_r y_r^2 + k_1y_1 + \cdots + k_ny_n = c$$

b. Put $x_1^2 + 3x_2^2 + 3x_3^2 + 4x_1x_2 - 4x_1x_3 + 5x_1 - 6x_3 = 7$ in this form and find variables y_1, y_2, y_3 as in (a).

Exercise 8.9.8 Given a symmetric matrix A, define $q_A(\mathbf{x}) = \mathbf{x}^T A\mathbf{x}$. Show that $B \overset{c}{\sim} A$ if and only if B is symmetric and there is an invertible matrix U such that $q_B(\mathbf{x}) = q_A(U\mathbf{x})$ for all $\mathbf{x}$. [*Hint*: Theorem 8.9.3.]

Exercise 8.9.9 Let $q(\mathbf{x}) = \mathbf{x}^T A\mathbf{x}$ be a quadratic form where $A = A^T$.

a. Show that $q(\mathbf{x}) > 0$ for all $\mathbf{x} \neq \mathbf{0}$, if and only if A is positive definite (all eigenvalues are positive). In this case, q is called **positive definite**.

b. Show that new variables $\mathbf{y}$ can be found such that $q = \|\mathbf{y}\|^2$ and $\mathbf{y} = U\mathbf{x}$ where U is upper triangular with positive diagonal entries. [*Hint*: Theorem 8.3.3.]

Exercise 8.9.10 A **bilinear form** β on $\mathbb{R}^n$ is a function that assigns to every pair $\mathbf{x}$, $\mathbf{y}$ of columns in $\mathbb{R}^n$ a number $\beta(\mathbf{x}, \mathbf{y})$ in such a way that

$$\begin{aligned} \beta(r\mathbf{x} + s\mathbf{y}, \mathbf{z}) &= r\beta(\mathbf{x}, \mathbf{z}) + s\beta(\mathbf{y}, \mathbf{z}) \\ \beta(\mathbf{x}, r\mathbf{y} + s\mathbf{z}) &= r\beta(\mathbf{x}, \mathbf{z}) + s\beta(\mathbf{x}, \mathbf{z}) \end{aligned}$$

for all $\mathbf{x}$, $\mathbf{y}$, $\mathbf{z}$ in $\mathbb{R}^n$ and r, s in $\mathbb{R}$. If $\beta(\mathbf{x}, \mathbf{y}) = \beta(\mathbf{y}, \mathbf{x})$ for all $\mathbf{x}$, $\mathbf{y}$, β is called **symmetric**.

a. If β is a bilinear form, show that an $n \times n$ matrix A exists such that $\beta(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T A\mathbf{y}$ for all $\mathbf{x}$, $\mathbf{y}$.

b. Show that A is uniquely determined by β.

c. Show that β is symmetric if and only if $A = A^T$.

8.10 An Application to Constrained Optimization

It is a frequent occurrence in applications that a function $q = q(x_1, x_2, \dots, x_n)$ of n variables, called an **objective function**, is to be made as large or as small as possible among all vectors $\mathbf{x} = (x_1, x_2, \dots, x_n)$ lying in a certain region of $\mathbb{R}^n$ called the **feasible region**. A wide variety of objective functions q arise in practice; our primary concern here is to examine one important situation where q is a quadratic form. The next example gives some indication of how such problems arise.

Example 8.10.1

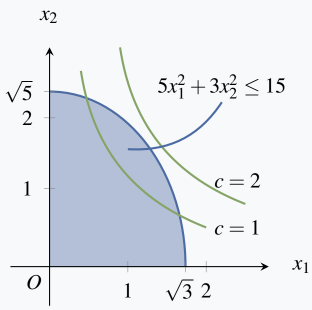

A politician proposes to spend x_1 dollars annually on health care and x_2 dollars annually on education. She is constrained in her spending by various budget pressures, and one model of this is that the expenditures x_1 and x_2 should satisfy a constraint like

$$5x_1^2 + 3x_2^2 \leq 15$$

Since $x_i \geq 0$ for each i, the feasible region is the shaded area shown in the diagram. Any choice of feasible point (x_1, x_2) in this region will satisfy the budget constraints. However, these choices have different effects on voters, and the politician wants to choose $\mathbf{x} = (x_1, x_2)$ to maximize some measure $q = q(x_1, x_2)$ of voter satisfaction. Thus the assumption is that, for any value of c, all points on the graph of $q(x_1, x_2) = c$ have the same appeal to voters. Hence the goal is to find the largest value of c for which the graph of $q(x_1, x_2) = c$ contains a feasible point.

The choice of the function q depends upon many factors; we will show how to solve the problem for any quadratic form q (even with more than two variables). In the diagram the function q is given by

$$q(x_1, x_2) = x_1 x_2$$

and the graphs of $q(x_1, x_2) = c$ are shown for $c = 1$ and $c = 2$. As c increases the graph of $q(x_1, x_2) = c$ moves up and to the right. From this it is clear that there will be a solution for some value of c between 1 and 2 (in fact the largest value is $c = \frac{1}{2}\sqrt{15} = 1.94$ to two decimal places).

The constraint $5x_1^2 + 3x_2^2 \leq 15$ in Example 8.10.1 can be put in a standard form. If we divide through by 15, it becomes $\left(\frac{x_1}{\sqrt{3}}\right)^2 + \left(\frac{x_2}{\sqrt{5}}\right)^2 \leq 1$. This suggests that we introduce new variables $\mathbf{y} = (y_1, y_2)$ where $y_1 = \frac{x_1}{\sqrt{3}}$ and $y_2 = \frac{x_2}{\sqrt{5}}$. Then the constraint becomes $\|\mathbf{y}\|^2 \leq 1$, equivalently $\|\mathbf{y}\| \leq 1$. In terms of these new variables, the objective function is $q = \sqrt{15} y_1 y_2$, and we want to maximize this subject to $\|\mathbf{y}\| \leq 1$. When this is done, the maximizing values of x_1 and x_2 are obtained from $x_1 = \sqrt{3} y_1$ and $x_2 = \sqrt{5} y_2$.

Hence, for constraints like that in Example 8.10.1, there is no real loss in generality in assuming that the constraint takes the form $\|\mathbf{x}\| \leq 1$. In this case the principal axes theorem solves the problem. Recall that a vector in $\mathbb{R}^n$ of length 1 is called a *unit vector*.

Theorem 8.10.1

Consider the quadratic form $q = q(\mathbf{x}) = \mathbf{x}^T A\mathbf{x}$ where A is an $n \times n$ symmetric matrix, and let λ_1 and λ_n denote the largest and smallest eigenvalues of A, respectively. Then:

1. $\max\{q(\mathbf{x}) \mid \|\mathbf{x}\| \leq 1\} = \lambda_1$*, and* $q(\mathbf{f}_1) = \lambda_1$ *where* $\mathbf{f}_1$ *is any unit* λ_1*-eigenvector.*

2. $\min\{q(\mathbf{x}) \mid \|\mathbf{x}\| \leq 1\} = \lambda_n$*, and* $q(\mathbf{f}_n) = \lambda_n$ *where* $\mathbf{f}_n$ *is any unit* λ_n*-eigenvector.*

Proof. Since A is symmetric, let the (real) eigenvalues λ_i of A be ordered as to size as follows:

$$\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n$$

By the principal axes theorem, let P be an orthogonal matrix such that $P^T AP = D = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_n)$. Define $\mathbf{y} = P^T\mathbf{x}$, equivalently $\mathbf{x} = P\mathbf{y}$, and note $\|\mathbf{y}\| = \|\mathbf{x}\|$ because $\|\mathbf{y}\|^2 = \mathbf{y}^T\mathbf{y} = \mathbf{x}^T(PP^T)\mathbf{x} = \mathbf{x}^T\mathbf{x} = \|\mathbf{x}\|^2$. If we write $\mathbf{y} = (y_1, y_2, \dots, y_n)^T$, then

$$\begin{aligned} q(\mathbf{x}) &= q(P\mathbf{y}) = (P\mathbf{y})^T A(P\mathbf{y}) \\ &= \mathbf{y}^T(P^T AP)\mathbf{y} = \mathbf{y}^T D\mathbf{y} \\ &= \lambda_1 y_1^2 + \lambda_2 y_2^2 + \cdots + \lambda_n y_n^2 \end{aligned} \tag{8.9}$$

Now assume that $\|\mathbf{x}\| \le 1$. Since $\lambda_i \le \lambda_1$ for each i, (8.9) gives

$$q(\mathbf{x}) = \lambda_1 y_1^2 + \lambda_2 y_2^2 + \cdots + \lambda_n y_n^2 \le \lambda_1 y_1^2 + \lambda_1 y_2^2 + \cdots + \lambda_1 y_n^2 = \lambda_1\|\mathbf{y}\|^2 \le \lambda_1$$

because $\|\mathbf{y}\| = \|\mathbf{x}\| \le 1$. This shows that $q(\mathbf{x})$ cannot exceed λ_1 when $\|\mathbf{x}\| \le 1$. To see that this maximum is actually achieved, let $\mathbf{f}_1$ be a unit eigenvector corresponding to λ_1. Then

$$q(\mathbf{f}_1) = \mathbf{f}_1^T A\mathbf{f}_1 = \mathbf{f}_1^T(\lambda_1\mathbf{f}_1) = \lambda_1(\mathbf{f}_1^T\mathbf{f}_1) = \lambda_1\|\mathbf{f}_1\|^2 = \lambda_1$$

Hence λ_1 is the maximum value of $q(\mathbf{x})$ when $\|\mathbf{x}\| \le 1$, proving (1). The proof of (2) is analogous. □

The set of all vectors $\mathbf{x}$ in $\mathbb{R}^n$ such that $\|\mathbf{x}\| \le 1$ is called the **unit ball**. If $n = 2$, it is often called the unit disk and consists of the unit circle and its interior; if $n = 3$, it is the unit sphere and its interior. It is worth noting that the maximum value of a quadratic form $q(\mathbf{x})$ as $\mathbf{x}$ ranges *throughout* the unit ball is (by Theorem 8.10.1) actually attained for a unit vector $\mathbf{x}$ on the *boundary* of the unit ball.

Theorem 8.10.1 is important for applications involving vibrations in areas as diverse as aerodynamics and particle physics, and the maximum and minimum values in the theorem are often found using advanced calculus to minimize the quadratic form on the unit ball. The algebraic approach using the principal axes theorem gives a geometrical interpretation of the optimal values because they are eigenvalues.

Example 8.10.2

Maximize and minimize the form $q(\mathbf{x}) = 3x_1^2 + 14x_1x_2 + 3x_2^2$ subject to $\|\mathbf{x}\| \le 1$.

Solution. The matrix of q is $A = \begin{bmatrix} 3 & 7 \\ 7 & 3 \end{bmatrix}$, with eigenvalues $\lambda_1 = 10$ and $\lambda_2 = -4$, and corresponding unit eigenvectors $\mathbf{f}_1 = \frac{1}{\sqrt{2}}(1, 1)$ and $\mathbf{f}_2 = \frac{1}{\sqrt{2}}(1, -1)$. Hence, among all unit vectors $\mathbf{x}$ in $\mathbb{R}^2$, $q(\mathbf{x})$ takes its maximal value 10 at $\mathbf{x} = \mathbf{f}_1$, and the minimum value of $q(\mathbf{x})$ is -4 when $\mathbf{x} = \mathbf{f}_2$.

As noted above, the objective function in a constrained optimization problem need not be a quadratic form. We conclude with an example where the objective function is linear, and the feasible region is determined by linear constraints.

Example 8.10.3

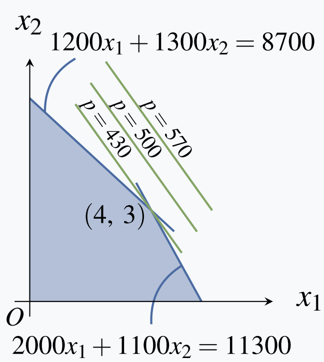

A manufacturer makes x_1 units of product 1, and x_2 units of product 2, at a profit of \$70 and \$50 per unit respectively, and wants to choose x_1 and x_2 to maximize the total profit $p(x_1, x_2) = 70x_1 + 50x_2$. However x_1 and x_2 are not arbitrary; for example, $x_1 \geq 0$ and $x_2 \geq 0$. Other conditions also come into play. Each unit of product 1 costs \$1200 to produce and requires 2000 square feet of warehouse space; each unit of product 2 costs \$1300 to produce and requires 1100 square feet of space. If the total warehouse space is 11 300 square feet, and if the total production budget is \$8700, x_1 and x_2 must also satisfy the conditions

$$\begin{aligned} 2000x_1 + 1100x_2 &\leq 11300 \\ 1200x_1 + 1300x_2 &\leq 8700 \end{aligned}$$

The feasible region in the plane satisfying these constraints (and $x_1 \geq 0$, $x_2 \geq 0$) is shaded in the diagram. If the profit equation $70x_1 + 50x_2 = p$ is plotted for various values of p, the resulting lines are parallel, with p increasing with distance from the origin. Hence the best choice occurs for the line $70x_1 + 50x_2 = 430$ that touches the shaded region at the point $(4, 3)$. So the profit p has a maximum of $p = 430$ for $x_1 = 4$ units and $x_2 = 3$ units.

Example 8.10.3 is a simple case of the general **linear programming** problem[23] which arises in economic, management, network, and scheduling applications. Here the objective function is a linear combination $q = a_1x_1 + a_2x_2 + \cdots + a_nx_n$ of the variables, and the feasible region consists of the vectors $\mathbf{x} = (x_1, x_2, \dots, x_n)^T$ in $\mathbb{R}^n$ which satisfy a set of linear inequalities of the form $b_1x_1 + b_2x_2 + \cdots + b_nx_n \leq b$. There is a good method (an extension of the gaussian algorithm) called the **simplex algorithm** for finding the maximum and minimum values of q when $\mathbf{x}$ ranges over such a feasible set. As Example 8.10.3 suggests, the optimal values turn out to be vertices of the feasible set. In particular, they are on the boundary of the feasible region, as is the case in Theorem 8.10.1.

8.11 An Application to Statistical Principal Component Analysis

Linear algebra is important in multivariate analysis in statistics, and we conclude with a very short look at one application of diagonalization in this area. A main feature of probability and statistics is the idea of a *random variable* X, that is a real-valued function which takes its values according to a probability law (called its *distribution*). Random variables occur in a wide variety of contexts; examples include the number of meteors falling per square kilometre in a given region, the price of a share of a stock, or the duration of a long distance telephone call from a certain city.

The values of a random variable X are distributed about a central number μ, called the *mean* of X. The mean can be calculated from the distribution as the *expectation* $E(X) = \mu$ of the random variable X.

[23] More information is available in "Linear Programming and Extensions" by N. Wu and R. Coppins, McGraw-Hill, 1981.