3.7. The example showing that the simplex method can take an exponential number of iterations is due to Klee and Minty (1972). The Hirsch conjecture was made by Hirsch in 1957. The first results on the average case behavior of the simplex method were obtained by Borgwardt (1982) and Smale (1983). Schrijver (1986) contains an overview of the early research in this area, as well as proof of the $n/2$ bound on the number of pivots due to Haimovich (1983).

3.9. The results in Exercises 3.10 and 3.11, which deal with the smallest examples of cycling, are due to Marshall and Suurballe (1969). The matrix inversion lemma [Exercise 3.13(a)] is known as the Sherman-Morrison formula.

Chapter 4

Duality theory

Contents

In this chapter, we start with a linear programming problem, called the primal, and introduce another linear programming problem, called the dual. Duality theory deals with the relation between these two problems and uncovers the deeper structure of linear programming. It is a powerful theoretical tool that has numerous applications, provides new geometric insights, and leads to another algorithm for linear programming (the dual simplex method).

4.1 Motivation

Duality theory can be motivated as an outgrowth of the Lagrange multiplier method, often used in calculus to minimize a function subject to equality constraints. For example, in order to solve the problem

$$\begin{array}{ll} \text{minimize} & x^2 + y^2 \\ \text{subject to} & x + y = 1, \end{array}$$

we introduce a Lagrange multiplier p and form the Lagrangean $L(x, y, p)$ defined by

$$L(x, y, p) = x^2 + y^2 + p(1 - x - y).$$

While keeping p fixed, we minimize the Lagrangean over all x and y, subject to no constraints, which can be done by setting $\partial L/\partial x$ and $\partial L/\partial y$ to zero. The optimal solution to this unconstrained problem is

$$x = y = \frac{p}{2},$$

and depends on p. The constraint $x + y = 1$ gives us the additional relation $p = 1$, and the optimal solution to the original problem is $x = y = 1/2$.

The main idea in the above example is the following. Instead of enforcing the hard constraint $x + y = 1$, we allow it to be violated and associate a Lagrange multiplier, or *price*, p with the amount $1 - x - y$ by which it is violated. This leads to the unconstrained minimization of $x^2 + y^2 + p(1 - x - y)$. When the price is properly chosen ($p = 1$, in our example), the optimal solution to the constrained problem is also optimal for the unconstrained problem. In particular, under that specific value of p, the presence or absence of the hard constraint does not affect the optimal cost.

The situation in linear programming is similar: we associate a price variable with each constraint and start searching for prices under which the presence or absence of the constraints does not affect the optimal cost. It turns out that the right prices can be found by solving a new linear programming problem, called the *dual* of the original. We now motivate the form of the dual problem.

Consider the standard form problem

$$\begin{array}{lrl} \text{minimize} & \mathbf{c}'\mathbf{x} & \\ \text{subject to} & \mathbf{A}\mathbf{x} & = \mathbf{b} \\ & \mathbf{x} & \geq \mathbf{0}, \end{array}$$

which we call the *primal* problem, and let $\mathbf{x}^*$ be an optimal solution, assumed to exist. We introduce a *relaxed* problem in which the constraint $\mathbf{A}\mathbf{x} = \mathbf{b}$ is replaced by a penalty $\mathbf{p}'(\mathbf{b} - \mathbf{A}\mathbf{x})$, where $\mathbf{p}$ is a price vector of the same dimension as $\mathbf{b}$. We are then faced with the problem

$$\begin{array}{ll} \text{minimize} & \mathbf{c}'\mathbf{x} + \mathbf{p}'(\mathbf{b} - \mathbf{A}\mathbf{x}) \\ \text{subject to} & \mathbf{x} \geq \mathbf{0}. \end{array}$$

Let $g(\mathbf{p})$ be the optimal cost for the relaxed problem, as a function of the price vector $\mathbf{p}$. The relaxed problem allows for more options than those present in the primal problem, and we expect $g(\mathbf{p})$ to be no larger than the optimal primal cost $\mathbf{c}'\mathbf{x}^*$. Indeed,

$$g(\mathbf{p}) = \min_{\mathbf{x} \geq \mathbf{0}} \left[\mathbf{c}'\mathbf{x} + \mathbf{p}'(\mathbf{b} - \mathbf{A}\mathbf{x})\right] \leq \mathbf{c}'\mathbf{x}^* + \mathbf{p}'(\mathbf{b} - \mathbf{A}\mathbf{x}^*) = \mathbf{c}'\mathbf{x}^*,$$

where the last inequality follows from the fact that $\mathbf{x}^*$ is a feasible solution to the primal problem, and satisfies $\mathbf{A}\mathbf{x}^* = \mathbf{b}$. Thus, each $\mathbf{p}$ leads to a lower bound $g(\mathbf{p})$ for the optimal cost $\mathbf{c}'\mathbf{x}^*$. The problem

$$\begin{array}{ll} \text{maximize} & g(\mathbf{p}) \\ \text{subject to} & \text{no constraints} \end{array}$$

can be then interpreted as a search for the tightest possible lower bound of this type, and is known as the *dual* problem. The main result in duality theory asserts that the optimal cost in the dual problem is equal to the optimal cost $\mathbf{c}'\mathbf{x}^*$ in the primal. In other words, when the prices are chosen according to an optimal solution for the dual problem, the option of violating the constraints $\mathbf{A}\mathbf{x} = \mathbf{b}$ is of no value.

Using the definition of $g(\mathbf{p})$, we have

$$\begin{aligned} g(\mathbf{p}) &= \min_{\mathbf{x} \geq \mathbf{0}} \left[\mathbf{c}'\mathbf{x} + \mathbf{p}'(\mathbf{b} - \mathbf{A}\mathbf{x})\right] \\ &= \mathbf{p}'\mathbf{b} + \min_{\mathbf{x} \geq \mathbf{0}} (\mathbf{c}' - \mathbf{p}'\mathbf{A})\mathbf{x}. \end{aligned}$$

Note that

$$\min_{\mathbf{x} \geq \mathbf{0}} (\mathbf{c}' - \mathbf{p}'\mathbf{A})\mathbf{x} = \begin{cases} 0, & \text{if } \mathbf{c}' - \mathbf{p}'\mathbf{A} \geq \mathbf{0}', \\ -\infty, & \text{otherwise.} \end{cases}$$

In maximizing $g(\mathbf{p})$, we only need to consider those values of $\mathbf{p}$ for which $g(\mathbf{p})$ is not equal to $-\infty$. We therefore conclude that the dual problem is the same as the linear programming problem

$$\begin{array}{ll} \text{maximize} & \mathbf{p}'\mathbf{b} \\ \text{subject to} & \mathbf{p}'\mathbf{A} \leq \mathbf{c}'. \end{array}$$

In the preceding example, we started with the equality constraint $\mathbf{Ax} = \mathbf{b}$ and we ended up with no constraints on the sign of the price vector $\mathbf{p}$. If the primal problem had instead inequality constraints of the form $\mathbf{Ax} \geq \mathbf{b}$, they could be replaced by $\mathbf{Ax} - \mathbf{s} = \mathbf{b}$, $\mathbf{s} \geq \mathbf{0}$. The equality constraint can be written in the form

$$[\mathbf{A} \mid -\mathbf{I}] \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix} = \mathbf{0},$$

which leads to the dual constraints

$$\mathbf{p}'[\mathbf{A} \mid -\mathbf{I}] \leq [\mathbf{c}' \mid \mathbf{0}'],$$

or, equivalently,

$$\mathbf{p}'\mathbf{A} \leq \mathbf{c}', \quad \mathbf{p} \geq \mathbf{0}.$$

Also, if the vector $\mathbf{x}$ is free rather than sign-constrained, we use the fact

$$\min_{\mathbf{x}} (\mathbf{c}' - \mathbf{p}'\mathbf{A})\mathbf{x} = \begin{cases} 0, & \text{if } \mathbf{c}' - \mathbf{p}'\mathbf{A} = \mathbf{0}', \\ -\infty, & \text{otherwise,} \end{cases}$$

to end up with the constraint $\mathbf{p}'\mathbf{A} = \mathbf{c}'$ in the dual problem. These considerations motivate the general form of the dual problem which we introduce in the next section.

In summary, the construction of the dual of a primal minimization problem can be viewed as follows. We have a vector of parameters (dual variables) $\mathbf{p}$, and for every $\mathbf{p}$ we have a method for obtaining a lower bound on the optimal primal cost. The dual problem is a maximization problem that looks for the tightest such lower bound. For some vectors $\mathbf{p}$, the corresponding lower bound is equal to $-\infty$, and does not carry any useful information. Thus, we only need to maximize over those $\mathbf{p}$ that lead to nontrivial lower bounds, and this is what gives rise to the dual constraints.

4.2 The dual problem

Let $\mathbf{A}$ be a matrix with rows $\mathbf{a}_i'$ and columns $\mathbf{A}_j$. Given a *primal* problem with the structure shown on the left, its *dual* is defined to be the maximization problem shown on the right:

$$\begin{array}{llll}
\text{minimize} & \mathbf{c}'\mathbf{x} & & \\
\text{subject to} & \mathbf{a}_i'\mathbf{x} \geq b_i, & i \in M_1, \\
& \mathbf{a}_i'\mathbf{x} \leq b_i, & i \in M_2, \\
& \mathbf{a}_i'\mathbf{x} = b_i, & i \in M_3, \\
& x_j \geq 0, & j \in N_1, \\
& x_j \leq 0, & j \in N_2, \\
& x_j \text{ free}, & j \in N_3,
\end{array}
\qquad
\begin{array}{lll}
\text{maximize} & \mathbf{p}'\mathbf{b} & \\
\text{subject to} & p_i \geq 0, & i \in M_1, \\
& p_i \leq 0, & i \in M_2, \\
& p_i \text{ free}, & i \in M_3, \\
& \mathbf{p}'\mathbf{A}_j \leq c_j, & j \in N_1, \\
& \mathbf{p}'\mathbf{A}_j \geq c_j, & j \in N_2, \\
& \mathbf{p}'\mathbf{A}_j = c_j, & j \in N_3.
\end{array}$$

Notice that for each constraint in the primal (other than the sign constraints), we introduce a variable in the dual problem; for each variable in the primal, we introduce a constraint in the dual. Depending on whether the primal constraint is an equality or inequality constraint, the corresponding dual variable is either free or sign-constrained, respectively. In addition, depending on whether a variable in the primal problem is free or sign-constrained, we have an equality or inequality constraint, respectively, in the dual problem. We summarize these relations in Table 4.1.

PRIMAL	minimize	maximize	DUAL
	$\geq b_i$	≥ 0	
constraints	$\leq b_i$	≤ 0	**variables**
	$= b_i$	free	
	≥ 0	$\leq c_j$	
variables	≤ 0	$\geq c_j$	**constraints**
	free	$= c_j$	

Table 4.1: Relation between primal and dual variables and constraints.

If we start with a maximization problem, we can always convert it into an equivalent minimization problem, and then form its dual according to the rules we have described. However, to avoid confusion, we will adhere to the convention that the primal is a minimization problem, and its dual is a maximization problem. Finally, we will keep referring to the objective function in the dual problem as a "cost" that is being maximized.

A problem and its dual can be stated more compactly, in matrix notation, if a particular form is assumed for the primal. We have, for example, the following pairs of primal and dual problems:

$$\begin{array}{ll}
\text{minimize} & \mathbf{c}'\mathbf{x} \\
\text{subject to} & \mathbf{Ax} = \mathbf{b} \\
& \mathbf{x} \geq \mathbf{0},
\end{array}
\qquad
\begin{array}{ll}
\text{maximize} & \mathbf{p}'\mathbf{b} \\
\text{subject to} & \mathbf{p}'\mathbf{A} \leq \mathbf{c}',
\end{array}$$

and

$$\begin{array}{ll}
\text{minimize} & \mathbf{c}'\mathbf{x} \\
\text{subject to} & \mathbf{Ax} \geq \mathbf{b},
\end{array}
\qquad
\begin{array}{ll}
\text{maximize} & \mathbf{p}'\mathbf{b} \\
\text{subject to} & \mathbf{p}'\mathbf{A} = \mathbf{c}' \\
& \mathbf{p} \geq \mathbf{0}.
\end{array}$$

Example 4.1 Consider the primal problem shown on the left and its dual shown on the right:

$$\begin{array}{llll} \text{minimize} & x_1 + 2x_2 + 3x_3 & \text{maximize} & 5p_1 + 6p_2 + 4p_3 \\ \text{subject to} & -x_1 + 3x_2 \quad = 5 & \text{subject to} & p_1 \text{ free} \\ & 2x_1 - x_2 + 3x_3 \geq 6 & & p_2 \geq 0 \\ & x_3 \leq 4 & & p_3 \leq 0 \\ & x_1 \geq 0 & & -p_1 + 2p_2 \leq 1 \\ & x_2 \leq 0 & & 3p_1 - p_2 \geq 2 \\ & x_3 \text{ free}, & & 3p_2 + p_3 = 3. \end{array}$$

We transform the dual into an equivalent minimization problem, rename the variables from p_1, p_2, p_3 to x_1, x_2, x_3, and multiply the three last constraints by -1. The resulting problem is shown on the left. Then, on the right, we show its dual:

$$\begin{array}{llll} \text{minimize} & -5x_1 - 6x_2 - 4x_3 & \text{maximize} & -p_1 - 2p_2 - 3p_3 \\ \text{subject to} & x_1 \text{ free} & \text{subject to} & p_1 - 3p_2 = -5 \\ & x_2 \geq 0 & & -2p_1 + p_2 - 3p_3 \leq -6 \\ & x_3 \leq 0 & & -p_3 \geq -4 \\ & x_1 - 2x_2 \geq -1 & & p_1 \geq 0 \\ & -3x_1 + x_2 \leq -2 & & p_2 \leq 0 \\ & -3x_2 - x_3 = -3, & & p_3 \text{ free}. \end{array}$$

We observe that the latter problem is equivalent to the primal problem we started with. (The first three constraints in the latter problem are the same as the first three constraints in the original problem, multiplied by -1. Also, if the maximization in the latter problem is changed to a minimization, by multiplying the objective function by -1, we obtain the cost function in the original problem.)

The first primal problem considered in Example 4.1 had all of the ingredients of a general linear programming problem. This suggests that the conclusion reached at the end of the example should hold in general. Indeed, we have the following result. Its proof needs nothing more than the steps followed in Example 4.1, with abstract symbols replacing specific numbers, and will therefore be omitted.

Theorem 4.1 *If we transform the dual into an equivalent minimization problem and then form its dual, we obtain a problem equivalent to the original problem.*

A compact statement that is often used to describe Theorem 4.1 is that "the dual of the dual is the primal."

Any linear programming problem can be manipulated into one of several equivalent forms, for example, by introducing slack variables or by using the difference of two nonnegative variables to replace a single free variable. Each equivalent form leads to a somewhat different form for the dual problem. Nevertheless, the examples that follow indicate that the duals of equivalent problems are equivalent.

Example 4.2 Consider the primal problem shown on the left and its dual shown on the right:

$$\begin{array}{llll} \text{minimize} & \mathbf{c}'\mathbf{x} & \text{maximize} & \mathbf{p}'\mathbf{b} \\ \text{subject to} & \mathbf{A}\mathbf{x} \geq \mathbf{b} & \text{subject to} & \mathbf{p} \geq \mathbf{0} \\ & \mathbf{x} \text{ free}, & & \mathbf{p}'\mathbf{A} = \mathbf{c}'. \end{array}$$

We transform the primal problem by introducing surplus variables and then obtain its dual:

$$\begin{array}{llll} \text{minimize} & \mathbf{c}'\mathbf{x} + \mathbf{0}'\mathbf{s} & \text{maximize} & \mathbf{p}'\mathbf{b} \\ \text{subject to} & \mathbf{A}\mathbf{x} - \mathbf{s} = \mathbf{b} & \text{subject to} & \mathbf{p} \text{ free} \\ & \mathbf{x} \text{ free} & & \mathbf{p}'\mathbf{A} = \mathbf{c}' \\ & \mathbf{s} \geq \mathbf{0}, & & -\mathbf{p} \leq \mathbf{0}. \end{array}$$

Alternatively, if we take the original primal problem and replace $\mathbf{x}$ by sign-constrained variables, we obtain the following pair of problems:

$$\begin{array}{llll} \text{minimize} & \mathbf{c}'\mathbf{x}^+ - \mathbf{c}'\mathbf{x}^- & \text{maximize} & \mathbf{p}'\mathbf{b} \\ \text{subject to} & \mathbf{A}\mathbf{x}^+ - \mathbf{A}\mathbf{x}^- \geq \mathbf{b} & \text{subject to} & \mathbf{p} \geq \mathbf{0} \\ & \mathbf{x}^+ \geq \mathbf{0} & & \mathbf{p}'\mathbf{A} \leq \mathbf{c}' \\ & \mathbf{x}^- \geq \mathbf{0}, & & -\mathbf{p}'\mathbf{A} \leq -\mathbf{c}'. \end{array}$$

Note that we have three equivalent forms of the primal. We observe that the constraint $\mathbf{p} \geq \mathbf{0}$ is equivalent to the constraint $-\mathbf{p} \leq \mathbf{0}$. Furthermore, the constraint $\mathbf{p}'\mathbf{A} = \mathbf{c}'$ is equivalent to the two constraints $\mathbf{p}'\mathbf{A} \leq \mathbf{c}'$ and $-\mathbf{p}'\mathbf{A} \leq -\mathbf{c}'$. Thus, the duals of the three variants of the primal problem are also equivalent.

The next example is in the same spirit and examines the effect of removing redundant equality constraints in a standard form problem.

Example 4.3 Consider a standard form problem, assumed feasible, and its dual:

$$\begin{array}{llll} \text{minimize} & \mathbf{c}'\mathbf{x} & \text{maximize} & \mathbf{p}'\mathbf{b} \\ \text{subject to} & \mathbf{A}\mathbf{x} = \mathbf{b} & \text{subject to} & \mathbf{p}'\mathbf{A} \leq \mathbf{c}'. \\ & \mathbf{x} \geq \mathbf{0}, & & \end{array}$$

Let $\mathbf{a}'_1, \ldots, \mathbf{a}'_m$ be the rows of $\mathbf{A}$ and suppose that $\mathbf{a}_m = \sum_{i=1}^{m-1} \gamma_i \mathbf{a}_i$ for some scalars $\gamma_1, \ldots, \gamma_{m-1}$. In particular, the last equality constraint is redundant and can be eliminated. By considering an arbitrary feasible solution $\mathbf{x}$, we obtain

$$b_m = \mathbf{a}'_m \mathbf{x} = \sum_{i=1}^{m-1} \gamma_i \mathbf{a}'_i \mathbf{x} = \sum_{i=1}^{m-1} \gamma_i b_i. \tag{4.1}$$

Note that the dual constraints are of the form $\sum_{i=1}^{m} p_i \mathbf{a}'_i \leq \mathbf{c}'$ and can be rewritten as

$$\sum_{i=1}^{m-1} (p_i + \gamma_i p_m) \mathbf{a}'_i \leq \mathbf{c}'.$$

Furthermore, using Eq. (4.1), the dual cost $\sum_{i=1}^{m} p_i b_i$ is equal to

$$\sum_{i=1}^{m-1} (p_i + \gamma_i p_m) b_i.$$

If we now let $q_i = p_i + \gamma_i p_m$, we see that the dual problem is equivalent to

$$\begin{array}{ll} \text{maximize} & \sum_{i=1}^{m-1} q_i b_i \\ \text{subject to} & \sum_{i=1}^{m-1} q_i \mathbf{a}_i' \le \mathbf{c}'. \end{array}$$

We observe that this is the exact same dual that we would have obtained if we had eliminated the last (and redundant) constraint in the primal problem, before forming the dual.

The conclusions of the preceding two examples are summarized and generalized by the following result.

Theorem 4.2 *Suppose that we have transformed a linear programming problem Π_1 to another linear programming problem Π_2, by a sequence of transformations of the following types:*

(a) *Replace a free variable with the difference of two nonnegative variables.*

(b) *Replace an inequality constraint by an equality constraint involving a nonnegative slack variable.*

(c) *If some row of the matrix $\mathbf{A}$ in a feasible standard form problem is a linear combination of the other rows, eliminate the corresponding equality constraint.*

Then, the duals of Π_1 and Π_2 are equivalent, i.e., they are either both infeasible, or they have the same optimal cost.

The proof of Theorem 4.2 involves a combination of the various steps in Examples 4.2 and 4.3, and is left to the reader.

4.3 The duality theorem

We saw in Section 4.1 that for problems in standard form, the cost $g(\mathbf{p})$ of any dual solution provides a lower bound for the optimal cost. We now show that this property is true in general.

Theorem 4.3 (Weak duality) *If $\mathbf{x}$ is a feasible solution to the primal problem and $\mathbf{p}$ is a feasible solution to the dual problem, then*

$$\mathbf{p}'\mathbf{b} \le \mathbf{c}'\mathbf{x}.$$

Proof. For any vectors $\mathbf{x}$ and $\mathbf{p}$, we define

$$\begin{aligned} u_i &= p_i(\mathbf{a}_i'\mathbf{x} - b_i), \\ v_j &= (c_j - \mathbf{p}'\mathbf{A}_j)x_j. \end{aligned}$$

Suppose that $\mathbf{x}$ and $\mathbf{p}$ are primal and dual feasible, respectively. The definition of the dual problem requires the sign of p_i to be the same as the sign of $\mathbf{a}_i'\mathbf{x} - b_i$, and the sign of $c_j - \mathbf{p}'\mathbf{A}_j$ to be the same as the sign of x_j. Thus, primal and dual feasibility imply that

$$u_i \ge 0, \qquad \forall\, i,$$

and

$$v_j \ge 0, \qquad \forall\, j.$$

Notice that

$$\sum_i u_i = \mathbf{p}'\mathbf{A}\mathbf{x} - \mathbf{p}'\mathbf{b},$$

and

$$\sum_j v_j = \mathbf{c}'\mathbf{x} - \mathbf{p}'\mathbf{A}\mathbf{x}.$$

We add these two equalities and use the nonnegativity of u_i, v_j, to obtain

$$0 \le \sum_i u_i + \sum_j v_j = \mathbf{c}'\mathbf{x} - \mathbf{p}'\mathbf{b}.$$ □

The weak duality theorem is not a deep result, yet it does provide some useful information about the relation between the primal and the dual. We have, for example, the following corollary.

Corollary 4.1

(a) *If the optimal cost in the primal is $-\infty$, then the dual problem must be infeasible.*

(b) *If the optimal cost in the dual is $+\infty$, then the primal problem must be infeasible.*

Proof. Suppose that the optimal cost in the primal problem is $-\infty$ and that the dual problem has a feasible solution $\mathbf{p}$. By weak duality, $\mathbf{p}$ satisfies $\mathbf{p}'\mathbf{b} \le \mathbf{c}'\mathbf{x}$ for every primal feasible $\mathbf{x}$. Taking the minimum over all primal feasible $\mathbf{x}$, we conclude that $\mathbf{p}'\mathbf{b} \le -\infty$. This is impossible and shows that the dual cannot have a feasible solution, thus establishing part (a). Part (b) follows by a symmetrical argument. □

Another important corollary of the weak duality theorem is the following.

Corollary 4.2 *Let* $\mathbf{x}$ *and* $\mathbf{p}$ *be feasible solutions to the primal and the dual, respectively, and suppose that* $\mathbf{p}'\mathbf{b} = \mathbf{c}'\mathbf{x}$. *Then,* $\mathbf{x}$ *and* $\mathbf{p}$ *are optimal solutions to the primal and the dual, respectively.*

Proof. Let $\mathbf{x}$ and $\mathbf{p}$ be as in the statement of the corollary. For every primal feasible solution $\mathbf{y}$, the weak duality theorem yields $\mathbf{c}'\mathbf{x} = \mathbf{p}'\mathbf{b} \leq \mathbf{c}'\mathbf{y}$, which proves that $\mathbf{x}$ is optimal. The proof of optimality of $\mathbf{p}$ is similar. □

The next theorem is the central result on linear programming duality.

Theorem 4.4 (Strong duality) *If a linear programming problem has an optimal solution, so does its dual, and the respective optimal costs are equal.*

Proof. Consider the standard form problem

$$\begin{aligned} \text{minimize} \quad & \mathbf{c}'\mathbf{x} \\ \text{subject to} \quad & \mathbf{A}\mathbf{x} = \mathbf{b} \\ & \mathbf{x} \geq \mathbf{0}. \end{aligned}$$

Let us assume temporarily that the rows of $\mathbf{A}$ are linearly independent and that there exists an optimal solution. Let us apply the simplex method to this problem. As long as cycling is avoided, e.g., by using the lexicographic pivoting rule, the simplex method terminates with an optimal solution $\mathbf{x}$ and an optimal basis $\mathbf{B}$. Let $\mathbf{x}_B = \mathbf{B}^{-1}\mathbf{b}$ be the corresponding vector of basic variables. When the simplex method terminates, the reduced costs must be nonnegative and we obtain

$$\mathbf{c}' - \mathbf{c}_B'\mathbf{B}^{-1}\mathbf{A} \geq \mathbf{0}',$$

where $\mathbf{c}_B'$ is the vector with the costs of the basic variables. Let us define a vector $\mathbf{p}$ by letting $\mathbf{p}' = \mathbf{c}_B'\mathbf{B}^{-1}$. We then have $\mathbf{p}'\mathbf{A} \leq \mathbf{c}'$, which shows that $\mathbf{p}$ is a feasible solution to the dual problem

$$\begin{aligned} \text{maximize} \quad & \mathbf{p}'\mathbf{b} \\ \text{subject to} \quad & \mathbf{p}'\mathbf{A} \leq \mathbf{c}'. \end{aligned}$$

In addition,

$$\mathbf{p}'\mathbf{b} = \mathbf{c}_B'\mathbf{B}^{-1}\mathbf{b} = \mathbf{c}_B'\mathbf{x}_B = \mathbf{c}'\mathbf{x}.$$

It follows that $\mathbf{p}$ is an optimal solution to the dual (cf. Corollary 4.2), and the optimal dual cost is equal to the optimal primal cost.

If we are dealing with a general linear programming problem Π_1 that has an optimal solution, we first transform it into an equivalent standard

form problem Π_2, with the same optimal cost, and in which the rows of the matrix $\mathbf{A}$ are linearly independent. Let D_1 and D_2 be the duals of Π_1 and Π_2, respectively. By Theorem 4.2, the dual problems D_1 and D_2 have the same optimal cost. We have already proved that Π_2 and D_2 have the same optimal cost. It follows that Π_1 and D_1 have the same optimal cost (see Figure 4.1). □

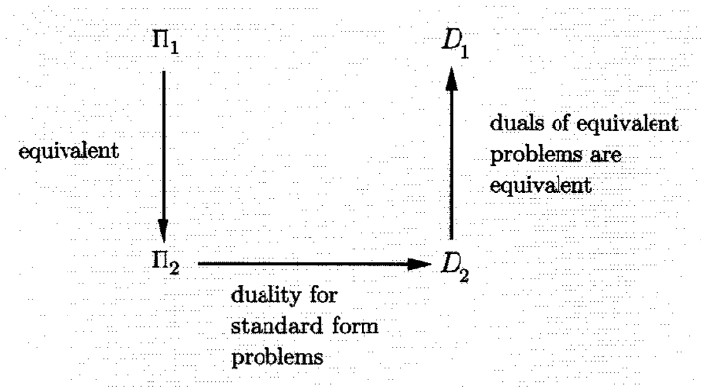

Figure 4.1: Proof of the duality theorem for general linear programming problems.

The preceding proof shows that an optimal solution to the dual problem is obtained as a byproduct of the simplex method as applied to a primal problem in standard form. It is based on the fact that the simplex method is guaranteed to terminate and this, in turn, depends on the existence of pivoting rules that prevent cycling. There is an alternative derivation of the duality theorem, which provides a geometric, algorithm-independent view of the subject, and which is developed in Section 4.7. At this point, we provide an illustration that conveys most of the content of the geometric proof.

Example 4.4 Consider a solid ball constrained to lie in a polyhedron defined by inequality constraints of the form $\mathbf{a}_i'\mathbf{x} \geq b_i$. If left under the influence of gravity, this ball reaches equilibrium at the lowest corner $\mathbf{x}^*$ of the polyhedron; see Figure 4.2. This corner is an optimal solution to the problem

$$\begin{aligned} \text{minimize} \quad & \mathbf{c}'\mathbf{x} \\ \text{subject to} \quad & \mathbf{a}_i'\mathbf{x} \geq b_i, \qquad \forall\, i, \end{aligned}$$

where $\mathbf{c}$ is a vertical vector pointing upwards. At equilibrium, gravity is counterbalanced by the forces exerted on the ball by the "walls" of the polyhedron. The latter forces are normal to the walls, that is, they are aligned with the vectors $\mathbf{a}_i$. We conclude that $\mathbf{c} = \sum_i p_i \mathbf{a}_i$, for some nonnegative coefficients p_i; in particular, the vector $\mathbf{p}$ is a feasible solution to the dual problem

$$\begin{aligned} \text{maximize} \quad & \mathbf{p}'\mathbf{b} \\ \text{subject to} \quad & \mathbf{p}'\mathbf{A} = \mathbf{c}' \\ & \mathbf{p} \geq \mathbf{0}. \end{aligned}$$

Given that forces can only be exerted by the walls that touch the ball, we must have $p_i = 0$, whenever $\mathbf{a}_i'\mathbf{x}^* > b_i$. Consequently, $p_i(b_i - \mathbf{a}_i'\mathbf{x}^*) = 0$ for all i. We therefore have $\mathbf{p}'\mathbf{b} = \sum_i p_i b_i = \sum_i p_i \mathbf{a}_i'\mathbf{x}^* = \mathbf{c}'\mathbf{x}^*$. It follows (Corollary 4.2) that $\mathbf{p}$ is an optimal solution to the dual, and the optimal dual cost is equal to the optimal primal cost.

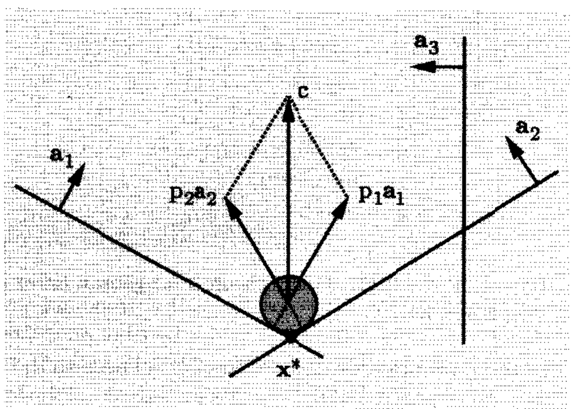

Figure 4.2: A mechanical analogy of the duality theorem.

Recall that in a linear programming problem, exactly one of the following three possibilities will occur:

(a) There is an optimal solution.

(b) The problem is "unbounded"; that is, the optimal cost is $-\infty$ (for minimization problems), or $+\infty$ (for maximization problems).

(c) The problem is infeasible.

This leads to nine possible combinations for the primal and the dual, which are shown in Table 4.2. By the strong duality theorem, if one problem has an optimal solution, so does the other. Furthermore, as discussed earlier, the weak duality theorem implies that if one problem is unbounded, the other must be infeasible. This allows us to mark some of the entries in Table 4.2 as "impossible."

	Finite optimum	**Unbounded**	**Infeasible**
Finite optimum	Possible	Impossible	Impossible
Unbounded	Impossible	Impossible	Possible
Infeasible	Impossible	Possible	Possible

Table 4.2: The different possibilities for the primal and the dual.

The case where both problems are infeasible can indeed occur, as shown by the following example.

Example 4.5 Consider the infeasible primal

$$\begin{aligned} \text{minimize} \quad & x_1 + 2x_2 \\ \text{subject to} \quad & x_1 + x_2 = 1 \\ & 2x_1 + 2x_2 = 3. \end{aligned}$$

Its dual is

$$\begin{aligned} \text{maximize} \quad & p_1 + 3p_2 \\ \text{subject to} \quad & p_1 + 2p_2 = 1 \\ & p_1 + 2p_2 = 2, \end{aligned}$$

which is also infeasible.

There is another interesting relation between the primal and the dual which is known as Clark's theorem (Clark, 1961). It asserts that unless both problems are infeasible, at least one of them must have an unbounded feasible set (Exercise 4.21).

Complementary slackness

An important relation between primal and dual optimal solutions is provided by the *complementary slackness* conditions, which we present next.

Theorem 4.5 (Complementary slackness) *Let $\mathbf{x}$ and $\mathbf{p}$ be feasible solutions to the primal and the dual problem, respectively. The vectors $\mathbf{x}$ and $\mathbf{p}$ are optimal solutions for the two respective problems if and only if:*

$$\begin{aligned} p_i(\mathbf{a}_i'\mathbf{x} - b_i) &= 0, \qquad \forall\, i, \\ (c_j - \mathbf{p}'\mathbf{A}_j)x_j &= 0, \qquad \forall\, j. \end{aligned}$$

Proof. In the proof of Theorem 4.3, we defined $u_i = p_i(\mathbf{a}_i'\mathbf{x} - b_i)$ and $v_j = (c_j - \mathbf{p}'\mathbf{A}_j)x_j$, and noted that for $\mathbf{x}$ primal feasible and $\mathbf{p}$ dual feasible, we have $u_i \geq 0$ and $v_j \geq 0$ for all i and j. In addition, we showed that

$$\mathbf{c}'\mathbf{x} - \mathbf{p}'\mathbf{b} = \sum_i u_i + \sum_j v_j.$$

By the strong duality theorem, if $\mathbf{x}$ and $\mathbf{p}$ are optimal, then $\mathbf{c}'\mathbf{x} = \mathbf{p}'\mathbf{b}$, which implies that $u_i = v_j = 0$ for all i, j. Conversely, if $u_i = v_j = 0$ for all i, j, then $\mathbf{c}'\mathbf{x} = \mathbf{p}'\mathbf{b}$, and Corollary 4.2 implies that $\mathbf{x}$ and $\mathbf{p}$ are optimal. □

The first complementary slackness condition is automatically satisfied by every feasible solution to a problem in standard form. If the primal problem is not in standard form and has a constraint like $\mathbf{a}_i'\mathbf{x} \geq b_i$, the corresponding complementary slackness condition asserts that the dual variable p_i is zero unless the constraint is active. An intuitive explanation is that a constraint which is not active at an optimal solution can be removed from the problem without affecting the optimal cost, and there is no point in associating a nonzero price with such a constraint. Note also the analogy with Example 4.4, where "forces" were only exerted by the active constraints.

If the primal problem is in standard form and a nondegenerate optimal basic feasible solution is known, the complementary slackness conditions determine a unique solution to the dual problem. We illustrate this fact in the next example.

Example 4.6 Consider a problem in standard form and its dual:

$$\begin{array}{llcl}
\text{minimize} & 13x_1 + 10x_2 + 6x_3 & & \\
\text{subject to} & 5x_1 + x_2 + 3x_3 & = & 8 \\
& 3x_1 + x_2 & = & 3 \\
& x_1, x_2, x_3 \geq 0, & &
\end{array}
\qquad
\begin{array}{llcl}
\text{maximize} & 8p_1 + 3p_2 & & \\
\text{subject to} & 5p_1 + 3p_2 & \leq & 13 \\
& p_1 + p_2 & \leq & 10 \\
& 3p_1 & \leq & 6.
\end{array}$$

As will be verified shortly, the vector $\mathbf{x}^* = (1, 0, 1)$ is a nondegenerate optimal solution to the primal problem. Assuming this to be the case, we use the complementary slackness conditions to construct the optimal solution to the dual. The condition $p_i(\mathbf{a}_i'\mathbf{x}^* - b_i) = 0$ is automatically satisfied for each i, since the primal is in standard form. The condition $(c_j - \mathbf{p}'\mathbf{A}_j)x_j^* = 0$ is clearly satisfied for $j = 2$, because $x_2^* = 0$. However, since $x_1^* > 0$ and $x_3^* > 0$, we obtain

$$5p_1 + 3p_2 = 13,$$

and

$$3p_1 = 6,$$

which we can solve to obtain $p_1 = 2$ and $p_2 = 1$. Note that this is a dual feasible solution whose cost is equal to 19, which is the same as the cost of $\mathbf{x}^*$. This verifies that $\mathbf{x}^*$ is indeed an optimal solution as claimed earlier.

We now generalize the above example. Suppose that x_j is a basic variable in a nondegenerate optimal basic feasible solution to a primal

problem in standard form. Then, the complementary slackness condition $(c_j - \mathbf{p}'\mathbf{A}_j)x_j = 0$ yields $\mathbf{p}'\mathbf{A}_j = c_j$ for every such j. Since the basic columns $\mathbf{A}_j$ are linearly independent, we obtain a system of equations for $\mathbf{p}$ which has a unique solution, namely, $\mathbf{p}' = \mathbf{c}_B'\mathbf{B}^{-1}$. A similar conclusion can also be drawn for problems not in standard form (Exercise 4.12). On the other hand, if we are given a degenerate optimal basic feasible solution to the primal, complementary slackness may be of very little help in determining an optimal solution to the dual problem (Exercise 4.17).

We finally mention that if the primal constraints are of the form $\mathbf{A}\mathbf{x} \geq \mathbf{b}$, $\mathbf{x} \geq \mathbf{0}$, and the primal problem has an optimal solution, then there exist optimal solutions to the primal and the dual which satisfy *strict complementary slackness*; that is, a variable in one problem is nonzero if and only if the corresponding constraint in the other problem is active (Exercise 4.20). This result has some interesting applications in discrete optimization, but these lie outside the scope of this book.

A geometric view

We now develop a geometric view that allows us to visualize pairs of primal and dual vectors without having to draw the dual feasible set.

We consider the primal problem

$$\begin{array}{ll}
\text{minimize} & \mathbf{c}'\mathbf{x} \\
\text{subject to} & \mathbf{a}_i'\mathbf{x} \geq b_i, \qquad i = 1, \ldots, m,
\end{array}$$

where the dimension of $\mathbf{x}$ is equal to n. We assume that the vectors $\mathbf{a}_i$ span $\Re^n$. The corresponding dual problem is

$$\begin{array}{ll}
\text{maximize} & \mathbf{p}'\mathbf{b} \\
\text{subject to} & \displaystyle\sum_{i=1}^{m} p_i\mathbf{a}_i = \mathbf{c} \\
& \mathbf{p} \geq \mathbf{0}.
\end{array}$$

Let I be a subset of $\{1, \ldots, m\}$ of cardinality n, such that the vectors $\mathbf{a}_i$, $i \in I$, are linearly independent. The system $\mathbf{a}_i'\mathbf{x} = b_i$, $i \in I$, has a unique solution, denoted by $\mathbf{x}^I$, which is a basic solution to the primal problem (cf. Definition 2.9 in Section 2.2). We assume, that $\mathbf{x}^I$ is nondegenerate, that is, $\mathbf{a}_i'\mathbf{x} \neq b_i$ for $i \notin I$.

Let $\mathbf{p} \in \Re^m$ be a dual vector (not necessarily dual feasible), and let us consider what is required for $\mathbf{x}^I$ and $\mathbf{p}$ to be optimal solutions to the primal and the dual problem, respectively. We need:

(a) $\mathbf{a}_i'\mathbf{x}^I \geq b_i$, for all i, (primal feasibility),

(b) $p_i = 0$, for all $i \notin I$, (complementary slackness),

(c) $\sum_{i=1}^{m} p_i\mathbf{a}_i = \mathbf{c}$, (dual feasibility),

(d) $\mathbf{p} \geq \mathbf{0}$, (dual feasibility).

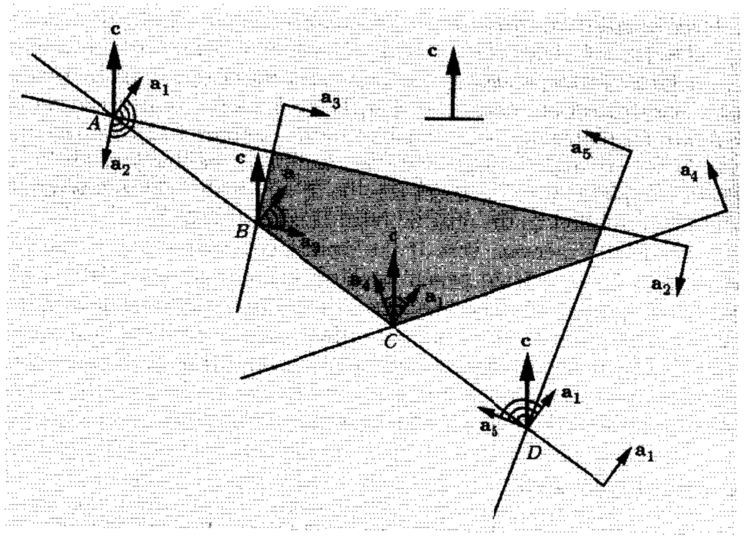

Figure 4.3: Consider a primal problem with two variables and five inequality constraints ($n = 2$, $m = 5$), and suppose that no two of the vectors $\mathbf{a}_i$ are collinear. Every two-element subset I of $\{1, 2, 3, 4, 5\}$ determines basic solutions $\mathbf{x}^I$ and $\mathbf{p}^I$ of the primal and the dual, respectively.
If $I = \{1, 2\}$, $\mathbf{x}^I$ is primal infeasible (point A) and $\mathbf{p}^I$ is dual infeasible, because $\mathbf{c}$ cannot be expressed as a nonnegative linear combination of the vectors $\mathbf{a}_1$ and $\mathbf{a}_2$.
If $I = \{1, 3\}$, $\mathbf{x}^I$ is primal feasible (point B) and $\mathbf{p}^I$ is dual infeasible.
If $I = \{1, 4\}$, $\mathbf{x}^I$ is primal feasible (point C) and $\mathbf{p}^I$ is dual feasible, because $\mathbf{c}$ can be expressed as a nonnegative linear combination of the vectors $\mathbf{a}_1$ and $\mathbf{a}_4$. In particular, $\mathbf{x}^I$ and $\mathbf{p}^I$ are optimal.
If $I = \{1, 5\}$, $\mathbf{x}^I$ is primal infeasible (point D) and $\mathbf{p}^I$ is dual feasible.

Given the complementary slackness condition (b), condition (c) becomes

$$\sum_{i \in I} p_i \mathbf{a}_i = \mathbf{c}.$$

Since the vectors $\mathbf{a}_i$, $i \in I$, are linearly independent, the latter equation has a unique solution that we denote by $\mathbf{p}^I$. In fact, it is readily seen that the vectors $\mathbf{a}_i$, $i \in I$, form a basis for the dual problem (which is in standard form) and $\mathbf{p}^I$ is the associated basic solution. For the vector $\mathbf{p}^I$ to be dual feasible, we also need it to be nonnegative. We conclude that once the complementary slackness condition (b) is enforced, feasibility of

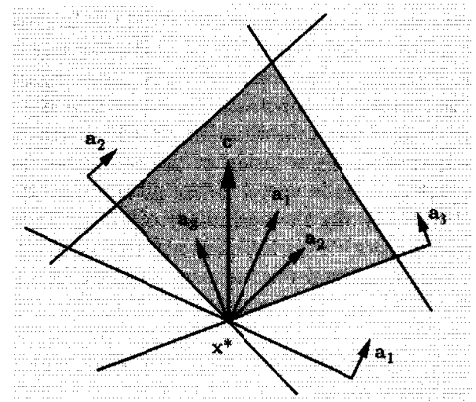

Figure 4.4: The vector $\mathbf{x}^*$ is a degenerate basic feasible solution of the primal. If we choose $I = \{1, 2\}$, the corresponding dual basic solution $\mathbf{p}^I$ is infeasible, because $\mathbf{c}$ is not a nonnegative linear combination of $\mathbf{a}_1$, $\mathbf{a}_2$. On the other hand, if we choose $I = \{1, 3\}$ or $I = \{2, 3\}$, the resulting dual basic solution $\mathbf{p}^I$ is feasible and, therefore, optimal.

the resulting dual vector $\mathbf{p}^I$ is equivalent to $\mathbf{c}$ being a nonnegative linear combination of the vectors $\mathbf{a}_i$, $i \in I$, associated with the active primal constraints. This allows us to visualize dual feasibility without having to draw the dual feasible set; see Figure 4.3.

If $\mathbf{x}^*$ is a degenerate basic solution to the primal, there can be several subsets I such that $\mathbf{x}^I = \mathbf{x}^*$. Using different choices for I, and by solving the system $\sum_{i \in I} p_i \mathbf{a}_i = \mathbf{c}$, we may obtain several dual basic solutions $\mathbf{p}^I$. It may then well be the case that some of them are dual feasible and some are not; see Figure 4.4. Still, if $\mathbf{p}^I$ is dual feasible (i.e., all p_i are nonnegative) and if $\mathbf{x}^*$ is primal feasible, then they are both optimal, because we have been enforcing complementary slackness and Theorem 4.5 applies.

4.4 Optimal dual variables as marginal costs

In this section, we elaborate on the interpretation of the dual variables as prices. This theme will be revisited, in more depth, in Chapter 5.

Consider the standard form problem

$$\begin{aligned} \text{minimize} \quad & \mathbf{c}'\mathbf{x} \\ \text{subject to} \quad & \mathbf{A}\mathbf{x} = \mathbf{b} \\ & \mathbf{x} \geq \mathbf{0}. \end{aligned}$$

We assume that the rows of $\mathbf{A}$ are linearly independent and that there is a nondegenerate basic feasible solution $\mathbf{x}^*$ which is optimal. Let $\mathbf{B}$ be the corresponding basis matrix and let $\mathbf{x}_B = \mathbf{B}^{-1}\mathbf{b}$ be the vector of basic variables, which is positive, by nondegeneracy. Let us now replace $\mathbf{b}$ by $\mathbf{b} + \mathbf{d}$, where $\mathbf{d}$ is a small perturbation vector. Since $\mathbf{B}^{-1}\mathbf{b} > \mathbf{0}$, we also have $\mathbf{B}^{-1}(\mathbf{b} + \mathbf{d}) > \mathbf{0}$, as long as $\mathbf{d}$ is small. This implies that the same basis leads to a basic feasible solution of the perturbed problem as well. Perturbing the right-hand side vector $\mathbf{b}$ has no effect on the reduced costs associated with this basis. By the optimality of $\mathbf{x}^*$ in the original problem, the vector of reduced costs $\mathbf{c}' - \mathbf{c}_B'\mathbf{B}^{-1}\mathbf{A}$ is nonnegative and this establishes that the same basis is optimal for the perturbed problem as well. Thus, the optimal cost in the perturbed problem is

$$\mathbf{c}_B'\mathbf{B}^{-1}(\mathbf{b} + \mathbf{d}) = \mathbf{p}'(\mathbf{b} + \mathbf{d}),$$

where $\mathbf{p}' = \mathbf{c}_B'\mathbf{B}^{-1}$ is an optimal solution to the dual problem. Therefore, a small change of $\mathbf{d}$ in the right-hand side vector $\mathbf{b}$ results in a change of $\mathbf{p}'\mathbf{d}$ in the optimal cost. We conclude that each component p_i of the optimal dual vector can be interpreted as the *marginal cost* (or *shadow price*) per unit increase of the ith requirement b_i.

We conclude with yet another interpretation of duality, for standard form problems. In order to develop some concrete intuition, we phrase our discussion in terms of the diet problem (Example 1.3 in Section 1.1). We interpret each vector $\mathbf{A}_j$ as the nutritional content of the jth available food, and view $\mathbf{b}$ as the nutritional content of an ideal food that we wish to synthesize. Let us interpret p_i as the "fair" price per unit of the ith nutrient. A unit of the jth food has a value of c_j at the food market, but it also has a value of $\mathbf{p}'\mathbf{A}_j$ if priced at the nutrient market. Complementary slackness asserts that every food which is used (at a nonzero level) to synthesize the ideal food, should be consistently priced at the two markets. Thus, duality is concerned with two alternative ways of cost accounting. The value of the ideal food, as computed in the food market, is $\mathbf{c}'\mathbf{x}^*$, where $\mathbf{x}^*$ is an optimal solution to the primal problem; the value of the ideal food, as computed in the nutrient market, is $\mathbf{p}'\mathbf{b}$. The duality relation $\mathbf{c}'\mathbf{x}^* = \mathbf{p}'\mathbf{b}$ states that when prices are chosen appropriately, the two accounting methods should give the same results.

4.5 Standard form problems and the dual simplex method

In this section, we concentrate on the case where the primal problem is in standard form. We develop the *dual simplex method*, which is an alternative to the simplex method of Chapter 3. We also comment on the relation between the basic feasible solutions to the primal and the dual, including a discussion of dual degeneracy.

In the proof of the strong duality theorem, we considered the simplex method applied to a primal problem in standard form and defined a dual vector $\mathbf{p}$ by letting $\mathbf{p}' = \mathbf{c}_B'\mathbf{B}^{-1}$. We then noted that the primal optimality condition $\mathbf{c}' - \mathbf{c}_B'\mathbf{B}^{-1}\mathbf{A} \geq \mathbf{0}'$ is the same as the dual feasibility condition $\mathbf{p}'\mathbf{A} \leq \mathbf{c}'$. We can thus think of the simplex method as an algorithm that maintains primal feasibility and works towards dual feasibility. A method with this property is generally called a *primal* algorithm. An alternative is to start with a dual feasible solution and work towards primal feasibility. A method of this type is called a *dual* algorithm. In this section, we present a dual simplex method, implemented in terms of the full tableau. We argue that it does indeed solve the dual problem, and we show that it moves from one basic feasible solution of the dual problem to another. An alternative implementation that only keeps track of the matrix $\mathbf{B}^{-1}$, instead of the entire tableau, is called a *revised dual simplex* method (Exercise 4.23).

The dual simplex method

Let us consider a problem in standard form, under the usual assumption that the rows of the matrix $\mathbf{A}$ are linearly independent. Let $\mathbf{B}$ be a basis matrix, consisting of m linearly independent columns of $\mathbf{A}$, and consider the corresponding tableau

$-\mathbf{c}_B'\mathbf{B}^{-1}\mathbf{b}$	$\bar{\mathbf{c}}'$
$\mathbf{B}^{-1}\mathbf{b}$	$\mathbf{B}^{-1}\mathbf{A}$

or, in more detail,

$-\mathbf{c}_B'\mathbf{x}_B$	$\bar{c}_1$	$\cdots$	$\bar{c}_n$
$x_{B(1)}$	$\vert$		$\vert$
$\vdots$	$\mathbf{B}^{-1}\mathbf{A}_1$	$\cdots$	$\mathbf{B}^{-1}\mathbf{A}_n$
$x_{B(m)}$	$\vert$		$\vert$

We do not require $\mathbf{B}^{-1}\mathbf{b}$ to be nonnegative, which means that we have a basic, but not necessarily feasible solution to the primal problem. However, we assume that $\bar{\mathbf{c}} \geq \mathbf{0}$; equivalently, the vector $\mathbf{p}' = \mathbf{c}_B'\mathbf{B}^{-1}$ satisfies $\mathbf{p}'\mathbf{A} \leq \mathbf{c}'$, and we have a feasible solution to the dual problem. The cost of this dual feasible solution is $\mathbf{p}'\mathbf{b} = \mathbf{c}_B'\mathbf{B}^{-1}\mathbf{b} = \mathbf{c}_B'\mathbf{x}_B$, which is the negative of the entry at the upper left corner of the tableau. If the inequality $\mathbf{B}^{-1}\mathbf{b} \geq \mathbf{0}$ happens to hold, we also have a primal feasible solution with the same cost, and optimal solutions to both problems have been found. If the inequality $\mathbf{B}^{-1}\mathbf{b} \geq \mathbf{0}$ fails to hold, we perform a change of basis in a manner we describe next.

We find some ℓ such that $x_{B(\ell)} < 0$ and consider the ℓth row of the tableau, called the *pivot row*; this row is of the form $(x_{B(\ell)}, v_1, \ldots, v_n)$, where v_i is the ℓth component of $\mathbf{B}^{-1}\mathbf{A}_i$. For each i with $v_i < 0$ (if such i exist), we form the ratio $\bar{c}_i/|v_i|$ and let j be an index for which this ratio is smallest; that is, $v_j < 0$ and

$$\frac{\bar{c}_j}{|v_j|} = \min_{\{i \mid v_i < 0\}} \frac{\bar{c}_i}{|v_i|}. \tag{4.2}$$

(We call the corresponding entry v_j the *pivot element*. Note that x_j must be a nonbasic variable, since the jth column in the tableau contains the negative element v_j.) We then perform a change of basis: column $\mathbf{A}_j$ enters the basis and column $\mathbf{A}_{B(\ell)}$ exits. This change of basis (or *pivot*) is effected exactly as in the primal simplex method: we add to each row of the tableau a multiple of the pivot row so that all entries in the pivot column are set to zero, with the exception of the pivot element which is set to 1. In particular, in order to set the reduced cost in the pivot column to zero, we multiply the pivot row by $\bar{c}_j/|v_j|$ and add it to the zeroth row. For every i, the new value of $\bar{c}_i$ is equal to

$$\bar{c}_i + v_i \frac{\bar{c}_j}{|v_j|},$$

which is nonnegative because of the way that j was selected [cf. Eq. (4.2)]. We conclude that the reduced costs in the new tableau will also be nonnegative and dual feasibility has been maintained.

Example 4.7 Consider the tableau

		x_1	x_2	x_3	x_4	x_5
	0	2	6	10	0	0
$x_4 =$	2	-2	4	1	1	0
$x_5 =$	-1	4	-2^*	-3	0	1

Since $x_{B(2)} < 0$, we choose the second row to be the pivot row. Negative entries of the pivot row are found in the second and third column. We compare the corresponding ratios $6/|-2|$ and $10/|-3|$. The smallest ratio is $6/|-2|$ and, therefore, the second column enters the basis. (The pivot element is indicated by an asterisk.) We multiply the pivot row by 3 and add it to the zeroth row. We multiply the pivot row by 2 and add it to the first row. We then divide the pivot row by -2. The new tableau is

		x_1	x_2	x_3	x_4	x_5
	-3	14	0	1	0	3
$x_4 =$	0	6	0	-5	1	2
$x_2 =$	1/2	-2	1	3/2	0	$-1/2$

The cost has increased to 3. Furthermore, we now have $\mathbf{B}^{-1}\mathbf{b} \geq \mathbf{0}$, and an optimal solution has been found.

Note that the pivot element v_j is always chosen to be negative, whereas the corresponding reduced cost $\bar{c}_j$ is nonnegative. Let us temporarily assume that $\bar{c}_j$ is in fact positive. Then, in order to replace $\bar{c}_j$ by zero, we need to add a positive multiple of the pivot row to the zeroth row. Since $x_{B(\ell)}$ is negative, this has the effect of adding a negative quantity to the upper left corner. Equivalently, the dual cost increases. Thus, as long as the reduced cost of every nonbasic variable is nonzero, the dual cost increases with each basis change, and no basis will ever be repeated in the course of the algorithm. It follows that the algorithm must eventually terminate and this can happen in one of two ways:

(a) We have $\mathbf{B}^{-1}\mathbf{b} \geq \mathbf{0}$ and an optimal solution.

(b) All of the entries $v_1, \ldots, v_n$ in the pivot row are nonnegative and we are therefore unable to locate a pivot element. In full analogy with the primal simplex method, this implies that the optimal dual cost is equal to $+\infty$ and the primal problem is infeasible; the proof is left as an exercise (Exercise 4.22).

We now provide a summary of the algorithm.

An iteration of the dual simplex method

1. A typical iteration starts with the tableau associated with a basis matrix $\mathbf{B}$ and with all reduced costs nonnegative.
2. Examine the components of the vector $\mathbf{B}^{-1}\mathbf{b}$ in the zeroth column of the tableau. If they are all nonnegative, we have an optimal basic feasible solution and the algorithm terminates; else, choose some ℓ such that $x_{B(\ell)} < 0$.
3. Consider the ℓth row of the tableau, with elements $x_{B(\ell)}, v_1, \ldots, v_n$ (the pivot row). If $v_i \geq 0$ for all i, then the optimal dual cost is $+\infty$ and the algorithm terminates.
4. For each i such that $v_i < 0$, compute the ratio $\bar{c}_i/|v_i|$ and let j be the index of a column that corresponds to the smallest ratio. The column $\mathbf{A}_{B(\ell)}$ exits the basis and the column $\mathbf{A}_j$ takes its place.
5. Add to each row of the tableau a multiple of the ℓth row (the pivot row) so that v_j (the pivot element) becomes 1 and all other entries of the pivot column become 0.

Let us now consider the possibility that the reduced cost $\bar{c}_j$ in the pivot column is zero. In this case, the zeroth row of the tableau does not change and the dual cost $\mathbf{c}_B'\mathbf{B}^{-1}\mathbf{b}$ remains the same. The proof of termination given earlier does not apply and the algorithm can cycle. This can be avoided by employing a suitable anticycling rule, such as the following.

Lexicographic pivoting rule for the dual simplex method

1. Choose any row ℓ such that $x_{B(\ell)} < 0$, to be the pivot row.
2. Determine the index j of the entering column as follows. For each column with $v_i < 0$, divide all entries by $|v_i|$, and then choose the lexicographically smallest column. If there is a tie between several lexicographically smallest columns, choose the one with the smallest index.

If the dual simplex method is initialized so that every column of the tableau $[\text{that is, each vector } (\bar{c}_j, \mathbf{B}^{-1}\mathbf{A}_j)]$ is lexicographically positive, and if the above lexicographic pivoting rule is used, the method terminates in a finite number of steps. The proof is similar to the proof of the corresponding result for the primal simplex method (Theorem 3.4) and is left as an exercise (Exercise 4.24).

When should we use the dual simplex method

At this point, it is natural to ask when the dual simplex method should be used. One such case arises when a basic feasible solution of the dual problem is readily available. Suppose, for example, that we already have an optimal basis for some linear programming problem, and that we wish to solve the same problem for a different choice of the right-hand side vector $\mathbf{b}$. The optimal basis for the original problem may be primal infeasible under the new value of $\mathbf{b}$. On the other hand, a change in $\mathbf{b}$ does not affect the reduced costs and we still have a dual feasible solution. Thus, instead of solving the new problem from scratch, it may be preferable to apply the dual simplex algorithm starting from the optimal basis for the original problem. This idea will be considered in more detail in Chapter 5.

The geometry of the dual simplex method

Our development of the dual simplex method was based entirely on tableau manipulations and algebraic arguments. We now present an alternative viewpoint based on geometric considerations.

We continue assuming that we are dealing with a problem in standard form and that the matrix $\mathbf{A}$ has linearly independent rows. Let $\mathbf{B}$ be a basis matrix with columns $\mathbf{A}_{B(1)}, \ldots, \mathbf{A}_{B(m)}$. This basis matrix determines a basic solution to the primal problem with $\mathbf{x}_B = \mathbf{B}^{-1}\mathbf{b}$. The same basis can also be used to determine a dual vector $\mathbf{p}$ by means of the equations

$$\mathbf{p}'\mathbf{A}_{B(i)} = c_{B(i)}, \qquad i = 1, \ldots, m.$$

These are m equations in m unknowns; since the columns $\mathbf{A}_{B(1)}, \ldots, \mathbf{A}_{B(m)}$ are linearly independent, there is a unique solution $\mathbf{p}$. For such a vector $\mathbf{p}$, the number of linearly independent active dual constraints is equal to the dimension of the dual vector, and it follows that we have a basic solution to the dual problem. In matrix notation, the dual basic solution p satisfies $\mathbf{p}'\mathbf{B} = \mathbf{c}'_B$, or $\mathbf{p}' = \mathbf{c}'_B\mathbf{B}^{-1}$, which was referred to as the vector of simplex multipliers in Chapter 3. If $\mathbf{p}$ is also dual feasible, that is, if $\mathbf{p}'\mathbf{A} \le \mathbf{c}'$, then $\mathbf{p}$ is a basic feasible solution of the dual problem.

To summarize, a basis matrix $\mathbf{B}$ is associated with a basic solution to the primal problem and also with a basic solution to the dual. A basic solution to the primal (respectively, dual) which is primal (respectively, dual) feasible, is a basic feasible solution to the primal (respectively, dual).

We now have a geometric interpretation of the dual simplex method: at every iteration, we have a basic feasible solution to the dual problem. The basic feasible solutions obtained at any two consecutive iterations have $m-1$ linearly independent active constraints in common (the reduced costs of the $m-1$ variables that are common to both bases are zero); thus, consecutive basic feasible solutions are either adjacent or they coincide.

Example 4.8 Consider the following standard form problem and its dual:

$$\begin{array}{ll} \text{minimize} & x_1 + x_2 \\ \text{subject to} & x_1 + 2x_2 - x_3 = 2 \\ & x_1 - x_4 = 1 \\ & x_1, x_2, x_3, x_4 \ge 0, \end{array} \qquad \begin{array}{ll} \text{maximize} & 2p_1 + p_2 \\ \text{subject to} & p_1 + p_2 \le 1 \\ & 2p_1 \le 1 \\ & p_1, p_2 \ge 0. \end{array}$$

The feasible set of the primal problem is 4-dimensional. If we eliminate the variables x_3 and x_4, we obtain the equivalent problem

$$\begin{array}{ll} \text{minimize} & x_1 + x_2 \\ \text{subject to} & x_1 + 2x_2 \ge 2 \\ & x_1 \ge 1 \\ & x_1, x_2 \ge 0. \end{array}$$

The feasible sets of the equivalent primal problem and of the dual are shown in Figures 4.5(a) and 4.5(b), respectively.

There is a total of five different bases in the standard form primal problem and five different basic solutions. These correspond to the points A, B, C, D, and E in Figure 4.5(a). The same five bases also lead to five basic solutions to the dual problem, which are points A, B, C, D, and E in Figure 4.5(b).

For example, if we choose the columns $\mathbf{A}_3$ and $\mathbf{A}_4$ to be the basic columns, we have the infeasible primal basic solution $\mathbf{x} = (0, 0, -2, -1)$ (point A). The corresponding dual basic solution is obtained by letting $\mathbf{p}'\mathbf{A}_3 = c_3 = 0$ and $\mathbf{p}'\mathbf{A}_4 = c_4 = 0$, which yields $\mathbf{p} = (0, 0)$. This is a basic feasible solution of the dual problem and can be used to start the dual simplex method. The associated initial tableau is

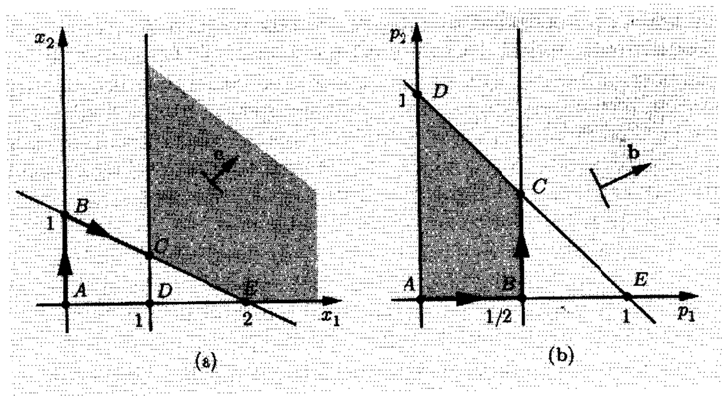

Figure 4.5: The feasible sets in Example 4.8.

		x_1	x_2	x_3	x_4
	0	1	1	0	0
$x_3 =$	−2	−1	−2*	1	0
$x_4 =$	−1	−1	0	0	1

We carry out two iterations of the dual simplex method to obtain the following two tableaux:

		x_1	x_2	x_3	x_4
	−1	1/2	0	1/2	0
$x_2 =$	1	1/2	1	−1/2	0
$x_4 =$	−1	−1*	0	0	1

		x_1	x_2	x_3	x_4
	−3/2	0	0	1/2	1/2
$x_2 =$	1/2	0	1	−1/2	1/2
$x_1 =$	1	1	0	0	−1

This sequence of tableaux corresponds to the path $A-B-C$ in either figure. In the primal space, the path traces a sequence of infeasible basic solutions until, at

optimality, it becomes feasible. In the dual space, the algorithm behaves exactly like the primal simplex method: it moves through a sequence of (dual) basic feasible solutions, while at each step improving the cost function.

Having observed that the dual simplex method moves from one basic feasible solution of the dual to an adjacent one, it may be tempting to say that the dual simplex method is simply the primal simplex method applied to the dual. This is a somewhat ambiguous statement, however, because the dual problem is not in standard form. If we were to convert it to standard form and then apply the primal simplex method, the resulting method is not necessarily identical to the dual simplex method (Exercise 4.25). A more accurate statement is to simply say that the dual simplex method is a variant of the simplex method tailored to problems defined exclusively in terms of linear inequality constraints.

Duality and degeneracy

Let us keep assuming that we are dealing with a standard form problem in which the rows of the matrix $\mathbf{A}$ are linearly independent. Any basis matrix $\mathbf{B}$ leads to an associated dual basic solution given by $\mathbf{p}' = \mathbf{c}'_B\mathbf{B}^{-1}$. At this basic solution, the dual constraint $\mathbf{p}'\mathbf{A}_j = c_j$ is active if and only if $\mathbf{c}'_B\mathbf{B}^{-1}\mathbf{A}_j = c_j$, that is, if and only if the reduced cost $\bar{c}_j$ is zero. Since $\mathbf{p}$ is m-dimensional, dual degeneracy amounts to having more than m reduced costs that are zero. Given that the reduced costs of the m basic variables must be zero, dual degeneracy is obtained whenever there exists a nonbasic variable whose reduced cost is zero.

The example that follows deals with the relation between basic solutions to the primal and the dual in the face of degeneracy.

Example 4.9 Consider the following standard from problem and its dual:

$$\begin{array}{llllll} \text{minimize} & 3x_1 + x_2 & & \qquad & \text{maximize} & 2p_1 \\ \text{subject to} & x_1 + x_2 - x_3 & = 2 & & \text{subject to} & p_1 + 2p_2 \le 3 \\ & 2x_1 - x_2 \quad - x_4 & = 0 & & & p_1 - p_2 \le 1 \\ & x_1, x_2, x_3, x_4 \ge 0, & & & & p_1, p_2 \ge 0. \end{array}$$

We eliminate x_3 and x_4 to obtain the equivalent primal problem

$$\begin{array}{ll} \text{minimize} & 3x_1 + x_2 \\ \text{subject to} & x_1 + x_2 \ge 2 \\ & 2x_1 - x_2 \ge 0 \\ & x_1, x_2 \ge 0. \end{array}$$

The feasible set of the equivalent primal and of the dual is shown in Figures 4.6(a) and 4.6(b), respectively.

There is a total of six different bases in the standard form primal problem, but only four different basic solutions [points A, B, C, D in Figure 4.6(a)]. In the dual problem, however, the six bases lead to six distinct basic solutions [points A, A', A'', B, C, D in Figure 4.6(b)].

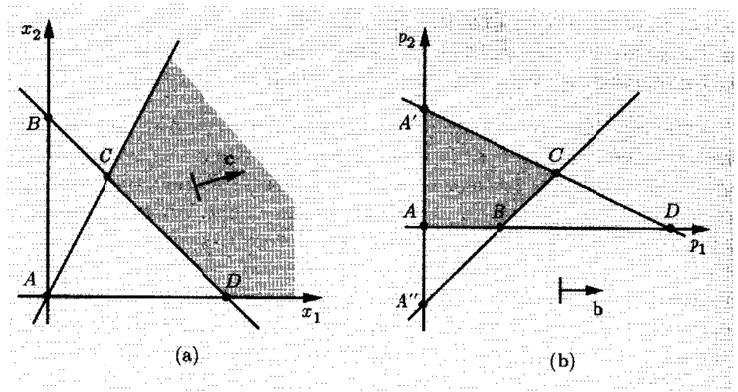

Figure 4.6: The feasible sets in Example 4.9.

For example, if we let columns $\mathbf{A}_3$ and $\mathbf{A}_4$ be basic, the primal basic solution has $x_1 = x_2 = 0$ and the corresponding dual basic solution is $(p_1, p_2) = (0, 0)$. Note that this is a basic feasible solution of the dual problem. If we let columns $\mathbf{A}_1$ and $\mathbf{A}_3$ be basic, the primal basic solution has again $x_1 = x_2 = 0$. For the dual problem, however, the equations $\mathbf{p}'\mathbf{A}_1 = c_1$ and $\mathbf{p}'\mathbf{A}_3 = c_3$ yield $(p_1, p_2) = (0, 3/2)$, which is a basic feasible solution of the dual, namely, point A' in Figure 4.6(b). Finally, if we let columns $\mathbf{A}_2$ and $\mathbf{A}_3$ be basic, we still have the same primal solution. For the dual problem, the equations $\mathbf{p}'\mathbf{A}_2 = c_1$ and $\mathbf{p}'\mathbf{A}_3 = c_3$ yield $(p_1, p_2) = (0, -1)$, which is an infeasible basic solution to the dual, namely, point A'' in Figure 4.6(b).

Example 4.9 has established that different bases may lead to the same basic solution for the primal problem, but to different basic solutions for the dual. Furthermore, out of the different basic solutions to the dual problem, it may be that some are feasible and some are infeasible.

We conclude with a summary of some properties of bases and basic solutions, for standard form problems, that were discussed in this section.

(a) Every basis determines a basic solution to the primal, but also a corresponding basic solution to the dual, namely, $\mathbf{p}' = \mathbf{c}'_B\mathbf{B}^{-1}$.

(b) This dual basic solution is feasible if and only if all of the reduced costs are nonnegative.

(c) Under this dual basic solution, the reduced costs that are equal to zero correspond to active constraints in the dual problem.

(d) This dual basic solution is degenerate if and only if some nonbasic variable has zero reduced cost.

4.6 Farkas' lemma and linear inequalities

Suppose that we wish to determine whether a given system of linear inequalities is infeasible. In this section, we approach this question using duality theory, and we show that infeasibility of a given system of linear inequalities is equivalent to the feasibility of another, related, system of linear inequalities. Intuitively, the latter system of linear inequalities can be interpreted as a search for a *certificate of infeasibility* for the former system.

To be more specific, consider a set of standard form constraints $\mathbf{Ax} = \mathbf{b}$ and $\mathbf{x} \geq \mathbf{0}$. Suppose that there exists some vector $\mathbf{p}$ such that $\mathbf{p}'\mathbf{A} \geq \mathbf{0}'$ and $\mathbf{p}'\mathbf{b} < 0$. Then, for any $\mathbf{x} \geq \mathbf{0}$, we have $\mathbf{p}'\mathbf{Ax} \geq \mathbf{0}$ and since $\mathbf{p}'\mathbf{b} < 0$, it follows that $\mathbf{p}'\mathbf{Ax} \neq \mathbf{p}'\mathbf{b}$. We conclude that $\mathbf{Ax} \neq \mathbf{b}$, for all $\mathbf{x} \geq \mathbf{0}$. This argument shows that if we can find a vector $\mathbf{p}$ satisfying $\mathbf{p}'\mathbf{A} \geq \mathbf{0}'$ and $\mathbf{p}'\mathbf{b} < 0$, the standard form constraints cannot have any feasible solution, and such a vector $\mathbf{p}$ is a certificate of infeasibility. Farkas' lemma below states that whenever a standard form problem is infeasible, such a certificate of infeasibility $\mathbf{p}$ is guaranteed to exist.

Theorem 4.6 (Farkas' lemma) *Let $\mathbf{A}$ be a matrix of dimensions $m \times n$ and let $\mathbf{b}$ be a vector in $\Re^m$. Then, exactly one of the following two alternatives holds:*

(a) *There exists some $\mathbf{x} \geq \mathbf{0}$ such that $\mathbf{Ax} = \mathbf{b}$.*

(b) *There exists some vector $\mathbf{p}$ such that $\mathbf{p}'\mathbf{A} \geq \mathbf{0}'$ and $\mathbf{p}'\mathbf{b} < 0$.*

Proof. One direction is easy. If there exists some $\mathbf{x} \geq \mathbf{0}$ satisfying $\mathbf{Ax} = \mathbf{b}$, and if $\mathbf{p}'\mathbf{A} \geq \mathbf{0}'$, then $\mathbf{p}'\mathbf{b} = \mathbf{p}'\mathbf{Ax} \geq 0$, which shows that the second alternative cannot hold.

Let us now assume that there exists no vector $\mathbf{x} \geq \mathbf{0}$ satisfying $\mathbf{Ax} = \mathbf{b}$. Consider the pair of problems

$$\begin{array}{ll} \text{maximize} & \mathbf{0}'\mathbf{x} \\ \text{subject to} & \mathbf{Ax} = \mathbf{b} \\ & \mathbf{x} \geq \mathbf{0}, \end{array} \qquad \begin{array}{ll} \text{minimize} & \mathbf{p}'\mathbf{b} \\ \text{subject to} & \mathbf{p}'\mathbf{A} \geq \mathbf{0}', \end{array}$$

and note that the first is the dual of the second. The maximization problem is infeasible, which implies that the minimization problem is either unbounded (the optimal cost is $-\infty$) or infeasible. Since $\mathbf{p} = \mathbf{0}$ is a feasible solution to the minimization problem, it follows that the minimization problem is unbounded. Therefore, there exists some $\mathbf{p}$ which is feasible, that is, $\mathbf{p}'\mathbf{A} \geq \mathbf{0}'$, and whose cost is negative, that is, $\mathbf{p}'\mathbf{b} < 0$. □

We now provide a geometric illustration of Farkas' lemma (see Figure 4.7). Let $\mathbf{A}_1, \ldots, \mathbf{A}_n$ be the columns of the matrix $\mathbf{A}$ and note that $\mathbf{Ax} = \sum_{i=1}^{n} \mathbf{A}_i x_i$. Therefore, the existence of a vector $\mathbf{x} \geq \mathbf{0}$ satisfying

$\mathbf{Ax} = \mathbf{b}$ is the same as requiring that $\mathbf{b}$ lies in the set of all nonnegative linear combinations of the vectors $\mathbf{A}_1, \ldots, \mathbf{A}_n$, which is the shaded region in Figure 4.7. If $\mathbf{b}$ does not belong to the shaded region (in which case the first alternative in Farkas' lemma does not hold), we expect intuitively that we can find a vector $\mathbf{p}$ and an associated hyperplane $\{\mathbf{z} \mid \mathbf{p}'\mathbf{z} = 0\}$ such that $\mathbf{b}$ lies on one side of the hyperplane while the shaded region lies on the other side. We then have $\mathbf{p}'\mathbf{b} < 0$ and $\mathbf{p}'\mathbf{A}_i \geq 0$ for all i, or, equivalently, $\mathbf{p}'\mathbf{A} \geq \mathbf{0}'$, and the second alternative holds.

Farkas' lemma predates the development of linear programming, but duality theory leads to a simple proof. A different proof, based on the geometric argument we just gave, is provided in the next section. Finally, there is an equivalent statement of Farkas' lemma which is sometimes more convenient.

Corollary 4.3 *Let $\mathbf{A}_1, \ldots, \mathbf{A}_n$ and $\mathbf{b}$ be given vectors and suppose that any vector $\mathbf{p}$ that satisfies $\mathbf{p}'\mathbf{A}_i \geq 0$, $i = 1, \ldots, n$, must also satisfy $\mathbf{p}'\mathbf{b} \geq 0$. Then, $\mathbf{b}$ can be expressed as a nonnegative linear combination of the vectors $\mathbf{A}_1, \ldots, \mathbf{A}_n$.*

Our next result is of a similar character.

Theorem 4.7 *Suppose that the system of linear inequalities $\mathbf{Ax} \leq \mathbf{b}$ has at least one solution, and let d be some scalar. Then, the following are equivalent:*

(a) Every feasible solution to the system $\mathbf{Ax} \leq \mathbf{b}$ satisfies $\mathbf{c}'\mathbf{x} \leq d$.

(b) There exists some $\mathbf{p} \geq \mathbf{0}$ such that $\mathbf{p}'\mathbf{A} = \mathbf{c}'$ and $\mathbf{p}'\mathbf{b} \leq d$.

Proof. Consider the following pair of problems

$$\begin{array}{llcll} \text{maximize} & \mathbf{c}'\mathbf{x} & \qquad & \text{minimize} & \mathbf{p}'\mathbf{b} \\ \text{subject to} & \mathbf{Ax} \leq \mathbf{b}, & & \text{subject to} & \mathbf{p}'\mathbf{A} = \mathbf{c}' \\ & & & & \mathbf{p} \geq \mathbf{0}, \end{array}$$

and note that the first is the dual of the second. If the system $\mathbf{Ax} \leq \mathbf{b}$ has a feasible solution and if every feasible solution satisfies $\mathbf{c}'\mathbf{x} \leq d$, then the first problem has an optimal solution and the optimal cost is bounded above by d. By the strong duality theorem, the second problem also has an optimal solution $\mathbf{p}$ whose cost is bounded above by d. This optimal solution satisfies $\mathbf{p}'\mathbf{A} = \mathbf{c}'$, $\mathbf{p} \geq \mathbf{0}$, and $\mathbf{p}'\mathbf{b} \leq d$.

Conversely, if some $\mathbf{p}$ satisfies $\mathbf{p}'\mathbf{A} = \mathbf{c}'$, $\mathbf{p} \geq \mathbf{0}$, and $\mathbf{p}'\mathbf{b} \leq d$, then the weak duality theorem asserts that every feasible solution to the first problem must also satisfy $\mathbf{c}'\mathbf{x} \leq d$. □

Results such as Theorems 4.6 and 4.7 are often called *theorems of the*

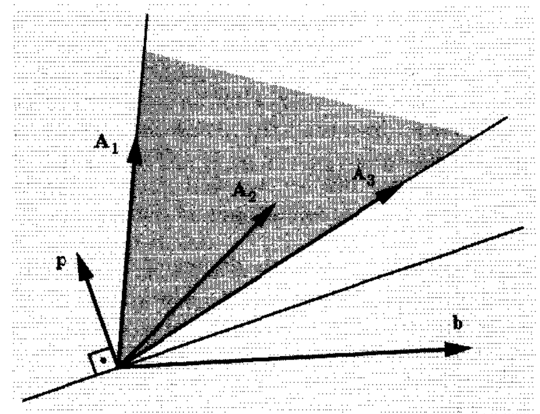

Figure 4.7: If the vector $\mathbf{b}$ does not belong to the set of all nonnegative linear combinations of $\mathbf{A}_1, \ldots, \mathbf{A}_n$, then we can find a hyperplane $\{\mathbf{z} \mid \mathbf{p}'\mathbf{z} = 0\}$ that separates it from that set.

alternative. There are several more results of this type; see, for example, Exercises 4.26, 4.27, and 4.28.

Applications of Farkas' lemma to asset pricing

Consider a market that operates for a single period, and in which n different assets are traded. Depending on the events during that single period, there are m possible states of nature at the end of the period. If we invest one dollar in some asset i and the state of nature turns out to be s, we receive a payoff of r_{si}. Thus, each asset i is described by a payoff vector $(r_{1i}, \ldots, r_{mi})$. The following $m \times n$ payoff matrix gives the payoffs of each of the n assets for each of the m states of nature:

$$\mathbf{R} = \begin{bmatrix} r_{11} & \ldots & r_{1n} \\ \vdots & \ddots & \vdots \\ r_{m1} & \ldots & r_{mn} \end{bmatrix}.$$

Let x_i be the amount held of asset i. A portfolio of assets is then a vector $\mathbf{x} = (x_1, \ldots, x_n)$. The components of a portfolio $\mathbf{x}$ can be either positive or negative. A positive value of x_i indicates that one has bought x_i units of asset i and is thus entitled to receive $r_{si}x_i$ if state s materializes. A negative value of x_i indicates a "short" position in asset i: this amounts to selling $|x_i|$ units of asset i at the beginning of the period, with a promise to buy them back at the end. Hence, one must pay out $r_{si}|x_i|$ if state s occurs, which is the same as receiving a payoff of $r_{si}x_i$.

The wealth in state s that results from a portfolio $\mathbf{x}$ is given by

$$w_s = \sum_{i=1}^{n} r_{si} x_i.$$

We introduce the vector $\mathbf{w} = (w_1, \ldots, w_m)$, and we obtain

$$\mathbf{w} = \mathbf{R}\mathbf{x}.$$

Let p_i be the price of asset i in the beginning of the period, and let $\mathbf{p} = (p_1, \ldots, p_n)$ be the vector of asset prices. Then, the cost of acquiring a portfolio $\mathbf{x}$ is given by $\mathbf{p}'\mathbf{x}$.

The central problem in asset pricing is to determine what the prices p_i should be. In order to address this question, we introduce the *absence of arbitrage* condition, which underlies much of finance theory: asset prices should always be such that no investor can get a guaranteed nonnegative payoff out of a negative investment. In other words, any portfolio that pays off nonnegative amounts in every state of nature, must be valuable to investors, so it must have nonnegative cost. Mathematically, the absence of arbitrage condition can be expressed as follows:

$$\text{if } \mathbf{R}\mathbf{x} \geq \mathbf{0}, \quad \text{then we must have } \mathbf{p}'\mathbf{x} \geq 0.$$

Given a particular set of assets, as described by the payoff matrix $\mathbf{R}$, only certain prices $\mathbf{p}$ are consistent with the absence of arbitrage. What characterizes such prices? What restrictions does the assumption of no arbitrage impose on asset prices? The answer is provided by Farkas' lemma.

Theorem 4.8 *The absence of arbitrage condition holds if and only if there exists a nonnegative vector* $\mathbf{q} = (q_1, \ldots, q_m)$, *such that the price of each asset* i *is given by*

$$p_i = \sum_{s=1}^{m} q_s r_{si}.$$

Proof. The absence of arbitrage condition states that there exists no vector $\mathbf{x}$ such that $\mathbf{x}'\mathbf{R}' \geq \mathbf{0}'$ and $\mathbf{x}'\mathbf{p} < 0$. This is of the same form as condition (b) in the statement of Farkas' lemma (Theorem 4.6). (Note that here $\mathbf{p}$ plays the role of $\mathbf{b}$, and $\mathbf{R}'$ plays the role of $\mathbf{A}$.) Therefore, by Farkas' lemma, the absence of arbitrage condition holds if and only if there exists some nonnegative vector $\mathbf{q}$ such that $\mathbf{R}'\mathbf{q} = \mathbf{p}$, which is the same as the condition in the theorem's statement. □

Theorem 4.8 asserts that whenever the market works efficiently enough to eliminate the possibility of arbitrage, there must exist "state prices" q_s

that can be used to value the existing assets. Intuitively, it establishes a nonnegative price q_s for an elementary asset that pays one dollar if the state of nature is s, and nothing otherwise. It then requires that every asset must be consistently priced, its total value being the sum of the values of the elementary assets from which it is composed. There is an alternative interpretation of the variables q_s as being (unnormalized) probabilities of the different states s, which, however, we will not pursue. In general, the state price vector $\mathbf{q}$ will not be unique, unless the number of assets equals or exceeds the number of states.

The no arbitrage condition is very simple, and yet very powerful. It is the key element behind many important results in financial economics, but these lie beyond the scope of this text. (See, however, Exercise 4.33 for an application in options pricing.)

4.7 From separating hyperplanes to duality*

Let us review the path followed in our development of duality theory. We started from the fact that the simplex method, in conjunction with an anticycling rule, is guaranteed to terminate. We then exploited the termination conditions of the simplex method to derive the strong duality theorem. We finally used the duality theorem to derive Farkas' lemma, which we interpreted in terms of a hyperplane that separates $\mathbf{b}$ from the columns of $\mathbf{A}$. In this section, we show that the reverse line of argument is also possible. We start from first principles and prove a general result on separating hyperplanes. We then establish Farkas' lemma, and conclude by showing that the duality theorem follows from Farkas' lemma. This line of argument is more elegant and fundamental because instead of relying on the rather complicated development of the simplex method, it only involves a small number of basic geometric concepts. Furthermore, it can be naturally generalized to nonlinear optimization problems.

Closed sets and Weierstrass' theorem

Before we proceed any further, we need to develop some background material. A set $S \subset \Re^n$ is called *closed* if it has the following property: if $\mathbf{x}^1, \mathbf{x}^2, \ldots$ is a sequence of elements of S that converges to some $\mathbf{x} \in \Re^n$, then $\mathbf{x} \in S$. In other words, S contains the limit of any sequence of elements of S. Intuitively, the set S contains its boundary.

Theorem 4.9 *Every polyhedron is closed.*

Proof. Consider the polyhedron $P = \{\mathbf{x} \in \Re^n \mid \mathbf{A}\mathbf{x} \geq \mathbf{b}\}$. Suppose that $\mathbf{x}^1, \mathbf{x}^2, \ldots$ is a sequence of elements of P that converges to some $\mathbf{x}^*$. We have to show that $\mathbf{x}^* \in P$. For each k, we have $\mathbf{x}^k \in P$ and, therefore, $\mathbf{A}\mathbf{x}^k \geq \mathbf{b}$. Taking the limit, we obtain $\mathbf{A}\mathbf{x}^* = \mathbf{A}\big(\lim_{k\to\infty} \mathbf{x}^k\big) = \lim_{k\to\infty} \big(\mathbf{A}\mathbf{x}^k\big) \geq \mathbf{b}$, and $\mathbf{x}^*$ belongs to P. □

The following is a fundamental result from real analysis that provides us with conditions for the existence of an optimal solution to an optimization problem. The proof lies beyond the scope of this book and is omitted.

Theorem 4.10 (Weierstrass' theorem) *If $f : \Re^n \mapsto \Re$ is a continuous function, and if S is a nonempty, closed, and bounded subset of $\Re^n$, then there exists some $\mathbf{x}^* \in S$ such that $f(\mathbf{x}^*) \leq f(\mathbf{x})$ for all $\mathbf{x} \in S$. Similarly, there exists some $\mathbf{y}^* \in S$ such that $f(\mathbf{y}^*) \geq f(\mathbf{x})$ for all $\mathbf{x} \in S$.*

Weierstrass' theorem is not valid if the set S is not closed. Consider, for example, the set $S = \{x \in \Re \mid x > 0\}$. This set is not closed because we can form a sequence of elements of S that converge to zero, but $x = 0$ does not belong to S. We then observe that the cost function $f(x) = x$ is not minimized at any point in S; for every $x > 0$, there exists another positive number with smaller cost, and no feasible x can be optimal. Ultimately, the reason that S is not closed is that the feasible set was defined by means of *strict* inequalities. The definition of polyhedra and linear programming problems does not allow for strict inequalities in order to avoid situations of this type.

The separating hyperplane theorem

The result that follows is "geometrically obvious" but nevertheless extremely important in the study of convex sets and functions. It states that if we are given a closed and nonempty convex set S and a point $\mathbf{x}^* \notin S$, then we can find a hyperplane, called a *separating hyperplane*, such that S and $\mathbf{x}^*$ lie in different halfspaces (Figure 4.8).

Theorem 4.11 (Separating hyperplane theorem) *Let S be a nonempty closed convex subset of $\Re^n$ and let $\mathbf{x}^* \in \Re^n$ be a vector that does not belong to S. Then, there exists some vector $\mathbf{c} \in \Re^n$ such that $\mathbf{c}'\mathbf{x}^* < \mathbf{c}'\mathbf{x}$ for all $\mathbf{x} \in S$.*

Proof. Let $\|\cdot\|$ be the Euclidean norm defined by $\|\mathbf{x}\| = (\mathbf{x}'\mathbf{x})^{1/2}$. Let us fix some element $\mathbf{w}$ of S, and let

$$B = \big\{\mathbf{x} \mid \|\mathbf{x} - \mathbf{x}^*\| \leq \|\mathbf{w} - \mathbf{x}^*\|\big\},$$

and $D = S \cap B$ [Figure 4.9(a)]. The set D is nonempty, because $\mathbf{w} \in D$.

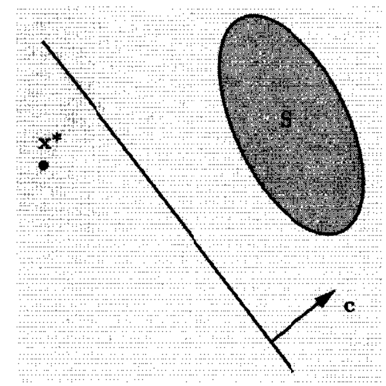

Figure 4.8: A hyperplane that separates the point $\mathbf{x}^*$ from the convex set S.

Furthermore, D is the intersection of the closed set S with the closed set B and is also closed. Finally, D is a bounded set because B is bounded. Consider the quantity $\|\mathbf{x} - \mathbf{x}^*\|$, where $\mathbf{x}$ ranges over the set D. This is a continuous function of $\mathbf{x}$. Since D is nonempty, closed, and bounded, Weierstrass' theorem implies that there exists some $\mathbf{y} \in D$ such that

$$\|\mathbf{y} - \mathbf{x}^*\| \leq \|\mathbf{x} - \mathbf{x}^*\|, \qquad \forall\, \mathbf{x} \in D.$$

For any $\mathbf{x} \in S$ that does not belong to D, we have $\|\mathbf{x} - \mathbf{x}^*\| > \|\mathbf{w} - \mathbf{x}^*\| \geq \|\mathbf{y} - \mathbf{x}^*\|$. We conclude that $\mathbf{y}$ minimizes $\|\mathbf{x} - \mathbf{x}^*\|$ over all $\mathbf{x} \in S$.

We have so far established that there exists an element $\mathbf{y}$ of S which is closest to $\mathbf{x}^*$. We now show that the vector $\mathbf{c} = \mathbf{y} - \mathbf{x}^*$ has the desired property [see Figure 4.9(b)].

Let $\mathbf{x} \in S$. For any λ satisfying $0 < \lambda \leq 1$, we have $\mathbf{y} + \lambda(\mathbf{x} - \mathbf{y}) \in S$, because S is convex. Since $\mathbf{y}$ minimizes $\|\mathbf{x} - \mathbf{x}^*\|$ over all $\mathbf{x} \in S$, we obtain

$$\begin{aligned}\|\mathbf{y} - \mathbf{x}^*\|^2 &\leq \|\mathbf{y} + \lambda(\mathbf{x} - \mathbf{y}) - \mathbf{x}^*\|^2 \\ &= \|\mathbf{y} - \mathbf{x}^*\|^2 + 2\lambda(\mathbf{y} - \mathbf{x}^*)'(\mathbf{x} - \mathbf{y}) + \lambda^2\|\mathbf{x} - \mathbf{y}\|^2,\end{aligned}$$

which yields

$$2\lambda(\mathbf{y} - \mathbf{x}^*)'(\mathbf{x} - \mathbf{y}) + \lambda^2\|\mathbf{x} - \mathbf{y}\|^2 \geq 0.$$

We divide by λ and then take the limit as λ decreases to zero. We obtain

$$(\mathbf{y} - \mathbf{x}^*)'(\mathbf{x} - \mathbf{y}) \geq 0.$$

[This inequality states that the angle θ in Figure 4.9(b) is no larger than 90 degrees.] Thus,

$$(\mathbf{y} - \mathbf{x}^*)'\mathbf{x} \geq (\mathbf{y} - \mathbf{x}^*)'\mathbf{y}$$

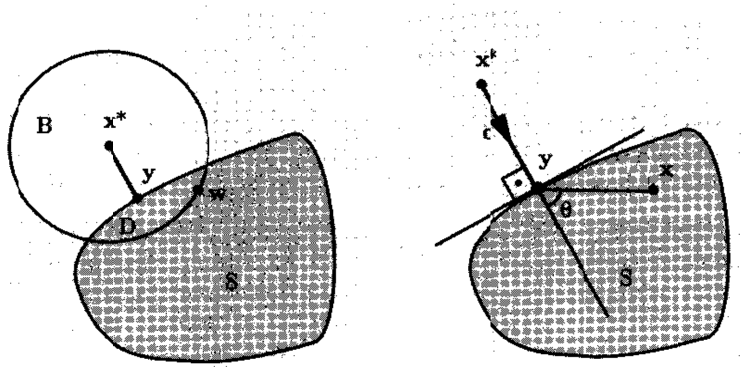

(a) (b)

Figure 4.9: Illustration of the proof of the separating hyperplane theorem.

$$= (\mathbf{y}-\mathbf{x}^*)'\mathbf{x}^* + (\mathbf{y}-\mathbf{x}^*)'(\mathbf{y}-\mathbf{x}^*)$$
$$> (\mathbf{y}-\mathbf{x}^*)'\mathbf{x}^*.$$

Setting $\mathbf{c} = \mathbf{y} - \mathbf{x}^*$ proves the theorem. □

Farkas' lemma revisited

We now show that Farkas' lemma is a consequence of the separating hyperplane theorem.

We will only be concerned with the difficult half of Farkas' lemma. In particular, we will prove that if the system $\mathbf{Ax} = \mathbf{b}$, $\mathbf{x} \geq \mathbf{0}$, does not have a solution, then there exists a vector $\mathbf{p}$ such that $\mathbf{p}'\mathbf{A} \geq \mathbf{0}'$ and $\mathbf{p}'\mathbf{b} < 0$.

Let

$$S = \{\mathbf{Ax} \mid \mathbf{x} \geq \mathbf{0}\}$$
$$= \Big\{\mathbf{y} \;\Big|\; \text{there exists } \mathbf{x} \text{ such that } \mathbf{y} = \mathbf{Ax},\ \mathbf{x} \geq \mathbf{0}\Big\},$$

and suppose that the vector $\mathbf{b}$ does not belong to S. The set S is clearly convex; it is also nonempty because $\mathbf{0} \in S$. Finally, the set S is closed; this may seem obvious, but is not easy to prove. For one possible proof, note that S is the projection of the polyhedron $\{(\mathbf{x},\mathbf{y}) \mid \mathbf{y} = \mathbf{Ax},\ \mathbf{x} \geq \mathbf{0}\}$ onto the $\mathbf{y}$ coordinates, is itself a polyhedron (see Section 2.8), and is therefore closed. An alternative proof is outlined in Exercise 4.37.

We now invoke the separating hyperplane theorem to separate $\mathbf{b}$ from S and conclude that there exists a vector $\mathbf{p}$ such that $\mathbf{p}'\mathbf{b} < \mathbf{p}'\mathbf{y}$ for every

$\mathbf{y} \in S$. Since $\mathbf{0} \in S$, we must have $\mathbf{p}'\mathbf{b} < 0$. Furthermore, for every column $\mathbf{A}_i$ of $\mathbf{A}$ and every $\lambda > 0$, we have $\lambda\mathbf{A}_i \in S$ and $\mathbf{p}'\mathbf{b} < \lambda\mathbf{p}'\mathbf{A}_i$. We divide both sides of the latter inequality by λ and then take the limit as λ tends to infinity, to conclude that $\mathbf{p}'\mathbf{A}_i \geq 0$. Since this is true for every i, we obtain $\mathbf{p}'\mathbf{A} \geq \mathbf{0}'$ and the proof is complete.

The duality theorem revisited

We will now derive the duality theorem as a corollary of Farkas' lemma. We only provide the proof for the case where the primal constraints are of the form $\mathbf{Ax} \geq \mathbf{b}$. The proof for the general case can be constructed along the same lines at the expense of more notation (Exercise 4.38). We also note that the proof given here is very similar to the line of argument used in the heuristic explanation of the duality theorem in Example 4.4.

We consider the following pair of primal and dual problems

$$\begin{array}{ll} \text{minimize} & \mathbf{c}'\mathbf{x} \\ \text{subject to} & \mathbf{Ax} \geq \mathbf{b}, \end{array} \qquad \begin{array}{ll} \text{maximize} & \mathbf{p}'\mathbf{b} \\ \text{subject to} & \mathbf{p}'\mathbf{A} = \mathbf{c}' \\ & \mathbf{p} \geq \mathbf{0}, \end{array}$$

and we assume that the primal has an optimal solution $\mathbf{x}^*$. We will show that the dual problem also has a feasible solution with the same cost. Once this is done, the strong duality theorem follows from weak duality (cf. Corollary 4.2).

Let $I = \{i \mid \mathbf{a}_i'\mathbf{x}^* = b_i\}$ be the set of indices of the constraints that are active at $\mathbf{x}^*$. We will first show that any vector $\mathbf{d}$ that satisfies $\mathbf{a}_i'\mathbf{d} \geq 0$ for every $i \in I$, must also satisfy $\mathbf{c}'\mathbf{d} \geq 0$. Consider such a vector $\mathbf{d}$ and let ϵ be a positive scalar. We then have $\mathbf{a}_i'(\mathbf{x}^* + \epsilon\mathbf{d}) \geq \mathbf{a}_i\mathbf{x}^* = b_i$ for all $i \in I$. In addition, if $i \notin I$ and if ϵ is sufficiently small, the inequality $\mathbf{a}_i'\mathbf{x}^* > b_i$ implies that $\mathbf{a}_i'(\mathbf{x}^* + \epsilon\mathbf{d}) > b_i$. We conclude that when ϵ is sufficiently small, $\mathbf{x}^* + \epsilon\mathbf{d}$ is a feasible solution. By the optimality of $\mathbf{x}^*$, we obtain $\mathbf{c}'\mathbf{d} \geq 0$, which establishes our claim. By Farkas' lemma (cf. Corollary 4.3), $\mathbf{c}$ can be expressed as a nonnegative linear combination of the vectors $\mathbf{a}_i$, $i \in I$, and there exist nonnegative scalars p_i, $i \in I$, such that

$$\mathbf{c} = \sum_{i \in I} p_i \mathbf{a}_i. \tag{4.3}$$

For $i \notin I$, we define $p_i = 0$. We then have $\mathbf{p} \geq \mathbf{0}$ and Eq. (4.3) shows that the vector $\mathbf{p}$ satisfies the dual constraint $\mathbf{p}'\mathbf{A} = \mathbf{c}'$. In addition,

$$\mathbf{p}'\mathbf{b} = \sum_{i \in I} p_i b_i = \sum_{i \in I} p_i \mathbf{a}_i'\mathbf{x}^* = \mathbf{c}'\mathbf{x}^*,$$

which shows that the cost of this dual feasible solution $\mathbf{p}$ is the same as the optimal primal cost. The duality theorem now follows from Corollary 4.2.

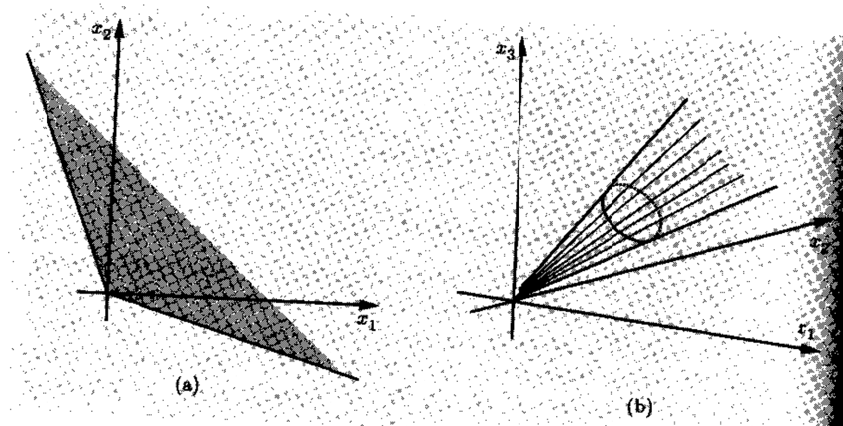

Figure 4.10: Examples of cones.

In conclusion, we have accomplished the goals that were set out in the beginning of this section. We proved the separating hyperplane theorem, which is a very intuitive and seemingly simple result, but with many important ramifications in optimization and other areas in mathematics. We used the separating hyperplane theorem to establish Farkas' lemma, and finally showed that the strong duality theorem is an easy consequence of Farkas' lemma.

4.8 Cones and extreme rays

We have seen in Chapter 2, that if the optimal cost in a linear programming problem is finite, then our search for an optimal solution can be restricted to finitely many points, namely, the basic feasible solutions, assuming one exists. In this section, we wish to develop a similar result for the case where the optimal cost is $-\infty$. In particular, we will show that the optimal cost is $-\infty$ if and only if there exists a cost reducing direction along which we can move without ever leaving the feasible set. Furthermore, our search for such a direction can be restricted to a finite set of suitably defined "extreme rays."

Cones

The first step in our development is to introduce the concept of a cone.

Definition 4.1 *A set $C \subset \Re^n$ is a* **cone** *if $\lambda \mathbf{x} \in C$ for all $\lambda \geq 0$ and all $\mathbf{x} \in C$.*

Notice that if C is a nonempty cone, then $\mathbf{0} \in C$. To this see, consider an arbitrary element $\mathbf{x}$ of C and set $\lambda = 0$ in the definition of a cone; see also Figure 4.10. A polyhedron of the form $P = \{\mathbf{x} \in \Re^n \mid \mathbf{Ax} \geq \mathbf{0}\}$ is easily seen to be a nonempty cone and is called a *polyhedral cone.*

Let $\mathbf{x}$ be a nonzero element of a polyhedral cone C. We then have $3\mathbf{x}/2 \in C$ and $\mathbf{x}/2 \in C$. Since $\mathbf{x}$ is the average of $3\mathbf{x}/2$ and $\mathbf{x}/2$, it is not an extreme point and, therefore, the only possible extreme point is the zero vector. If the zero vector is indeed an extreme point, we say that the cone is *pointed.* Whether this will be the case or not is determined by the criteria provided by our next result.

Theorem 4.12 *Let $C \subset \Re^n$ be the polyhedral cone defined by the constraints $\mathbf{a}_i'\mathbf{x} \geq 0$, $i = 1, \ldots, m$. Then, the following are equivalent:*

(a) *The zero vector is an extreme point of C.*

(b) *The cone C does not contain a line.*

(c) *There exist n vectors out of the family $\mathbf{a}_1, \ldots, \mathbf{a}_m$, which are linearly independent.*

Proof. This result is a special case of Theorem 2.6 in Section 2.5. ☐

Rays and recession cones

Consider a nonempty polyhedron

$$P = \{\mathbf{x} \in \Re^n \mid \mathbf{Ax} \geq \mathbf{b}\},$$

and let us fix some $\mathbf{y} \in P$. We define the *recession cone at* $\mathbf{y}$ as the set of all directions $\mathbf{d}$ along which we can move indefinitely away from $\mathbf{y}$, without leaving the set P. More formally, the recession cone is defined as the set

$$\{\mathbf{d} \in \Re^n \mid \mathbf{A}(\mathbf{y} + \lambda \mathbf{d}) \geq \mathbf{b}, \text{ for all } \lambda \geq 0\}.$$

It is easily seen that this set is the same as

$$\{\mathbf{d} \in \Re^n \mid \mathbf{Ad} \geq \mathbf{0}\},$$

and is a polyhedral cone. This shows that the recession cone is independent of the starting point $\mathbf{y}$; see Figure 4.11. The nonzero elements of the recession cone are called the *rays* of the polyhedron P.

For the case of a nonempty polyhedron $P = \{\mathbf{x} \in \Re^n \mid \mathbf{Ax} = \mathbf{b},\ \mathbf{x} \geq \mathbf{0}\}$ in standard form, the recession cone is seen to be the set of all vectors $\mathbf{d}$ that satisfy

$$\mathbf{Ad} = \mathbf{0}, \qquad \mathbf{d} \geq \mathbf{0}.$$

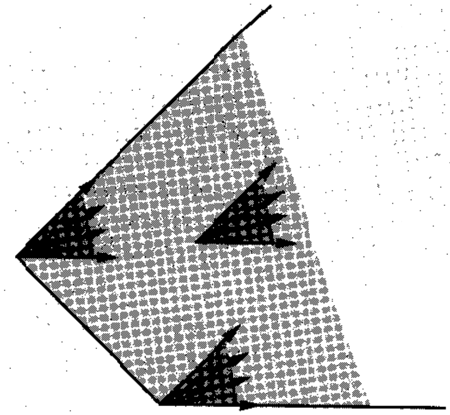

Figure 4.11: The recession cone at different elements of a polyhedron.

Extreme rays

We now define the extreme rays of a polyhedron. Intuitively, these are the directions associated with "edges" of the polyhedron that extend to infinity; see Figure 4.12 for an illustration.

> **Definition 4.2**
>
> **(a)** *A nonzero element* $\mathbf{x}$ *of a polyhedral cone* $C \subset \Re^n$ *is called an* **extreme ray** *if there are* $n-1$ *linearly independent constraints that are active at* $\mathbf{x}$.
>
> **(b)** *An extreme ray of the recession cone associated with a nonempty polyhedron* P *is also called an* **extreme ray** *of* P.

Note that a positive multiple of an extreme ray is also an extreme ray. We say that two extreme rays are *equivalent* if one is a positive multiple of the other. Note that for this to happen, they must correspond to the same $n-1$ linearly independent active constraints. Any $n-1$ linearly independent constraints define a line and can lead to at most two nonequivalent extreme rays (one being the negative of the other). Given that there is a finite number of ways that we can choose $n-1$ constraints to become active, and as long as we do not distinguish between equivalent extreme rays, we conclude that the number of extreme rays of a polyhedron is finite. A finite collection of extreme rays will be said to be a *complete set of extreme rays* if it contains exactly one representative from each equivalence class.

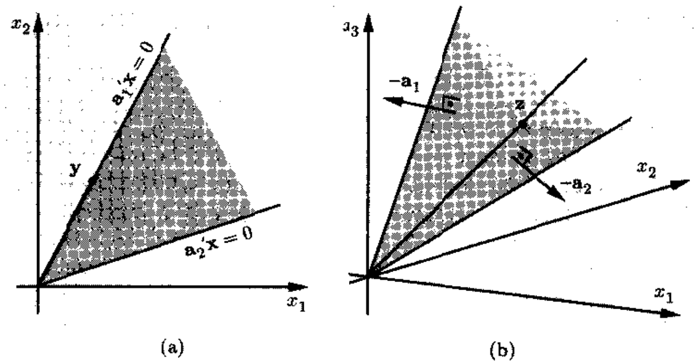

Figure 4.12: Extreme rays of polyhedral cones. (a) The vector $\mathbf{y}$ is an extreme ray because $n = 2$ and the constraint $\mathbf{a}_1'\mathbf{x} = 0$ is active at $\mathbf{y}$. (b) A polyhedral cone defined by three linearly independent constraints of the form $\mathbf{a}_i'\mathbf{x} \geq 0$. The vector $\mathbf{z}$ is an extreme ray because $n = 3$ and the two linearly independent constraints $\mathbf{a}_1'\mathbf{x} \geq 0$ and $\mathbf{a}_2'\mathbf{x} \geq 0$ are active at $\mathbf{z}$.

The definition of extreme rays mimics the definition of basic feasible solutions. An alternative and equivalent definition, resembling the definition of extreme points of polyhedra, is explored in Exercise 4.39.

Characterization of unbounded linear programming problems

We now derive conditions under which the optimal cost in a linear programming problem is equal to $-\infty$, first for the case where the feasible set is a cone, and then for the general case.

> **Theorem 4.13** *Consider the problem of minimizing* $\mathbf{c}'\mathbf{x}$ *over a pointed polyhedral cone* $C = \{\mathbf{x} \in \Re^n \mid \mathbf{a}_i'\mathbf{x} \geq 0,\ i = 1, \ldots, m\}$. *The optimal cost is equal to* $-\infty$ *if and only if some extreme ray* $\mathbf{d}$ *of* C *satisfies* $\mathbf{c}'\mathbf{d} < 0$.

Proof. One direction of the result is trivial because if some extreme ray has negative cost, then the cost becomes arbitrarily negative by moving along this ray.

For the converse, suppose that the optimal cost is $-\infty$. In particular, there exists some $\mathbf{x} \in C$ whose cost is negative and, by suitably scaling $\mathbf{x}$, we can assume that $\mathbf{c}'\mathbf{x} = -1$. In particular, the polyhedron

$$P = \{\mathbf{x} \in \Re^n \mid \mathbf{a}_1'\mathbf{x} \geq 0, \dots, \mathbf{a}_m'\mathbf{x} \geq 0, \mathbf{c}'\mathbf{x} = -1\}$$

is nonempty. Since C is pointed, the vectors $\mathbf{a}_1, \dots, \mathbf{a}_m$ span $\Re^n$ and this implies that P has at least one extreme point; let $\mathbf{d}$ be one of them. At $\mathbf{d}$, we have n linearly independent active constraints, which means that $n-1$ linearly independent constraints of the form $\mathbf{a}_i'\mathbf{x} \geq 0$ must be active. It follows that $\mathbf{d}$ is an extreme ray of C. □

By exploiting duality, Theorem 4.13 leads to a criterion for unboundedness in general linear programming problems. Interestingly enough, this criterion does not involve the right-hand side vector $\mathbf{b}$.

Theorem 4.14 *Consider the problem of minimizing $\mathbf{c}'\mathbf{x}$ subject to $\mathbf{A}\mathbf{x} \geq \mathbf{b}$, and assume that the feasible set has at least one extreme point. The optimal cost is equal to $-\infty$ if and only if some extreme ray $\mathbf{d}$ of the feasible set satisfies $\mathbf{c}'\mathbf{d} < 0$.*

Proof. One direction of the result is trivial because if an extreme ray has negative cost, then the cost becomes arbitrarily negative by starting at a feasible solution and moving along the direction of this ray.

For the proof of the reverse direction, we consider the dual problem:

$$\begin{aligned} \text{maximize} \quad & \mathbf{p}'\mathbf{b} \\ \text{subject to} \quad & \mathbf{p}'\mathbf{A} = \mathbf{c}' \\ & \mathbf{p} \geq \mathbf{0}. \end{aligned}$$

If the primal problem is unbounded, the dual problem is infeasible. Then, the related problem

$$\begin{aligned} \text{maximize} \quad & \mathbf{p}'\mathbf{0} \\ \text{subject to} \quad & \mathbf{p}'\mathbf{A} = \mathbf{c}' \\ & \mathbf{p} \geq \mathbf{0}, \end{aligned}$$

is also infeasible. This implies that the associated primal problem

$$\begin{aligned} \text{minimize} \quad & \mathbf{c}'\mathbf{x} \\ \text{subject to} \quad & \mathbf{A}\mathbf{x} \geq \mathbf{0}, \end{aligned}$$

is either unbounded or infeasible. Since $\mathbf{x} = \mathbf{0}$ is one feasible solution, it must be unbounded. Since the primal feasible set has at least one extreme point, the rows of A span $\Re^n$, where n is the dimension of $\mathbf{x}$. It follows that the recession cone $\{\mathbf{x} \mid \mathbf{A}\mathbf{x} \geq \mathbf{0}\}$ is pointed and, by Theorem 4.13, there exists an extreme ray $\mathbf{d}$ of the recession cone satisfying $\mathbf{c}'\mathbf{d} < 0$. By definition, this is an extreme ray of the feasible set. □

The unboundedness criterion in the simplex method

We end this section by pointing out that if we have a standard form problem in which the optimal cost is $-\infty$, the simplex method provides us at termination with an extreme ray.

Indeed, consider what happens when the simplex method terminates with an indication that the optimal cost is $-\infty$. At that point, we have a basis matrix $\mathbf{B}$, a nonbasic variable x_j with negative reduced cost, and the jth column $\mathbf{B}^{-1}\mathbf{A}_j$ of the tableau has no positive elements. Consider the jth basic direction $\mathbf{d}$, which is the vector that satisfies $\mathbf{d}_B = -\mathbf{B}^{-1}\mathbf{A}_j$, $d_j = 1$, and $d_i = 0$ for every nonbasic index i other than j. Then, the vector $\mathbf{d}$ satisfies $\mathbf{A}\mathbf{d} = \mathbf{0}$ and $\mathbf{d} \geq \mathbf{0}$, and belongs to the recession cone. It is also a direction of cost decrease, since the reduced cost $\bar{c}_j$ of the entering variable is negative.

Out of the constraints defining the recession cone, the jth basic direction $\mathbf{d}$ satisfies $n-1$ linearly independent such constraints with equality: these are the constraints $\mathbf{A}\mathbf{d} = \mathbf{0}$ (m of them) and the constraints $d_i = 0$ for i nonbasic and different than j ($n - m - 1$ of them). We conclude that $\mathbf{d}$ is an extreme ray.

4.9 Representation of polyhedra

In this section, we establish one of the fundamental results of linear programming theory. In particular, we show that any element of a polyhedron that has at least one extreme point can be represented as a convex combination of extreme points plus a nonnegative linear combination of extreme rays. A precise statement is given by our next result. A generalization to the case of general polyhedra is developed in Exercise 4.47.

Theorem 4.15 (Resolution theorem) *Let*

$$P = \{\mathbf{x} \in \Re^n \mid \mathbf{A}\mathbf{x} \geq \mathbf{b}\}$$

be a nonempty polyhedron with at least one extreme point. Let $\mathbf{x}^1, \dots, \mathbf{x}^k$ be the extreme points, and let $\mathbf{w}^1, \dots, \mathbf{w}^r$ be a complete set of extreme rays of P. Let

$$Q = \left\{ \sum_{i=1}^{k} \lambda_i \mathbf{x}^i + \sum_{j=1}^{r} \theta_j \mathbf{w}^j \;\middle|\; \lambda_i \geq 0,\ \theta_j \geq 0,\ \sum_{i=1}^{k} \lambda_i = 1 \right\}.$$

Then, $Q = P$.

Proof. We first prove that $Q \subset P$. Let

$$\mathbf{x} = \sum_{i=1}^{k} \lambda_i \mathbf{x}^i + \sum_{j=1}^{r} \theta_j \mathbf{w}^j$$

be an element of Q, where the coefficients λ_i and θ_j are nonnegative, and $\sum_{i=1}^{k} \lambda_i = 1$. The vector $\mathbf{y} = \sum_{i=1}^{k} \lambda_i \mathbf{x}^i$ is a convex combination of elements of P. It therefore belongs to P and satisfies $\mathbf{A}\mathbf{y} \geq \mathbf{b}$. We also have $\mathbf{A}\mathbf{w}^j \geq \mathbf{0}$ for every j, which implies that the vector $\mathbf{z} = \sum_{j=1}^{r} \theta_j \mathbf{w}^j$ satisfies $\mathbf{A}\mathbf{z} \geq \mathbf{0}$. It then follows that the vector $\mathbf{x} = \mathbf{y} + \mathbf{z}$ satisfies $\mathbf{A}\mathbf{x} \geq \mathbf{b}$ and belongs to P.

For the reverse inclusion, we assume that P is not a subset of Q and we will derive a contradiction. Let $\mathbf{z}$ be an element of P that does not belong to Q. Consider the linear programming problem

$$\begin{array}{lll} \text{maximize} & \displaystyle\sum_{i=1}^{k} 0\lambda_i + \sum_{j=1}^{r} 0\theta_j & \\ \text{subject to} & \displaystyle\sum_{i=1}^{k} \lambda_i \mathbf{x}^i + \sum_{j=1}^{r} \theta_j \mathbf{w}^j = \mathbf{z} & \\ & \displaystyle\sum_{i=1}^{k} \lambda_i = 1 & \\ & \lambda_i \geq 0, & i = 1, \ldots, k, \\ & \theta_j \geq 0, & j = 1, \ldots, r, \end{array} \qquad (4.4)$$

which is infeasible because $\mathbf{z} \notin Q$. This problem is the dual of the problem

$$\begin{array}{lll} \text{minimize} & \mathbf{p}'\mathbf{z} + q & \\ \text{subject to} & \mathbf{p}'\mathbf{x}^i + q \geq 0, & i = 1, \ldots, k, \\ & \mathbf{p}'\mathbf{w}^j \geq 0, & j = 1, \ldots, r. \end{array} \qquad (4.5)$$

Because the latter problem has a feasible solution, namely, $\mathbf{p} = \mathbf{0}$ and $q = 0$, the optimal cost is $-\infty$, and there exists a feasible solution $(\mathbf{p}, q)$ whose cost $\mathbf{p}'\mathbf{z} + q$ is negative. On the other hand, $\mathbf{p}'\mathbf{x}^i + q \geq 0$ for all i and this implies that $\mathbf{p}'\mathbf{z} < \mathbf{p}'\mathbf{x}^i$ for all i. We also have $\mathbf{p}'\mathbf{w}^j \geq 0$ for all j. [1]

Having fixed $\mathbf{p}$ as above, we now consider the linear programming problem

$$\begin{array}{ll} \text{minimize} & \mathbf{p}'\mathbf{x} \\ \text{subject to} & \mathbf{A}\mathbf{x} \geq \mathbf{b}. \end{array}$$

If the optimal cost is finite, there exists an extreme point $\mathbf{x}^i$ which is optimal. Since $\mathbf{z}$ is a feasible solution, we obtain $\mathbf{p}'\mathbf{x}^i \leq \mathbf{p}'\mathbf{z}$, which is a

[1] For an intuitive view of this proof, the purpose of this paragraph was to construct a hyperplane that separates $\mathbf{z}$ from Q.

contradiction. If the optimal cost is $-\infty$, Theorem 4.14 implies that there exists an extreme ray $\mathbf{w}^j$ such that $\mathbf{p}'\mathbf{w}^j < 0$, which is again a contradiction. □

Example 4.10 Consider the unbounded polyhedron defined by the constraints

$$\begin{aligned} x_1 - x_2 &\geq -2 \\ x_1 + x_2 &\geq 1 \\ x_1, x_2 &\geq 0 \end{aligned}$$

(see Figure 4.13). This polyhedron has three extreme points, namely, $\mathbf{x}^1 = (0,2)$, $\mathbf{x}^2 = (0,1)$, and $\mathbf{x}^3 = (1,0)$. The recession cone C is described by the inequalities $d_1 - d_2 \geq 0$, $d_1 + d_2 \geq 0$, and $d_1, d_2 \geq 0$. We conclude that $C = \{(d_1, d_2) \mid 0 \leq d_2 \leq d_1\}$. This cone has two extreme rays, namely, $\mathbf{w}^1 = (1,1)$ and $\mathbf{w}^2 = (1,0)$. The vector $\mathbf{y} = (2,2)$ is an element of the polyhedron and can be represented as

$$\mathbf{y} = \begin{bmatrix} 2 \\ 2 \end{bmatrix} = \begin{bmatrix} 0 \\ 1 \end{bmatrix} + \begin{bmatrix} 1 \\ 1 \end{bmatrix} + \begin{bmatrix} 1 \\ 0 \end{bmatrix} = \mathbf{x}^2 + \mathbf{w}^1 + \mathbf{w}^2.$$

However, this representation is not unique; for example, we also have

$$\mathbf{y} = \begin{bmatrix} 2 \\ 2 \end{bmatrix} = \frac{1}{2}\begin{bmatrix} 0 \\ 1 \end{bmatrix} + \frac{1}{2}\begin{bmatrix} 1 \\ 0 \end{bmatrix} + \frac{3}{2}\begin{bmatrix} 1 \\ 1 \end{bmatrix} = \frac{1}{2}\mathbf{x}^2 + \frac{1}{2}\mathbf{x}^3 + \frac{3}{2}\mathbf{w}^1.$$

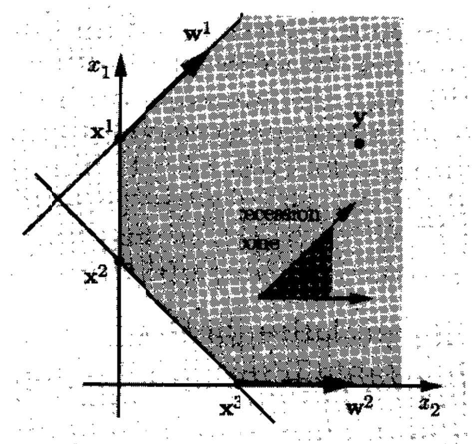

Figure 4.13: The polyhedron of Example 4.10.

We note that the set Q in Theorem 4.15 is the image of the polyhedron

$$H = \left\{ (\lambda_1, \ldots, \lambda_k, \theta_1, \ldots, \theta_r) \;\middle|\; \sum_{i=1}^{k} \lambda_i = 1,\ \lambda_i \geq 0,\ \theta_j \geq 0 \right\},$$

under the linear mapping

$$(\lambda_1,\ldots,\lambda_k,\theta_1,\ldots,\theta_r) \mapsto \sum_{i=1}^{k}\lambda_i \mathbf{x}^i + \sum_{j=1}^{r}\theta_j \mathbf{w}^j.$$

Thus, one corollary of the resolution theorem is that every polyhedron is the image, under a linear mapping, of a polyhedron H with this particular structure.

We now specialize Theorem 4.15 to the case of bounded polyhedra, to recover a result that was also proved in Section 2.7, using a different line of argument.

Corollary 4.4 *A nonempty bounded polyhedron is the convex hull of its extreme points.*

Proof. Let $P = \{\mathbf{x} \mid \mathbf{A}\mathbf{x} \geq \mathbf{b}\}$ be a nonempty bounded polyhedron. If $\mathbf{d}$ is a nonzero element of the cone $C = \{\mathbf{x} \mid \mathbf{A}\mathbf{x} \geq \mathbf{0}\}$ and $\mathbf{x}$ is an element of P, we have $\mathbf{x} + \lambda\mathbf{d} \in P$ for all $\lambda \geq 0$, contradicting the boundedness of P. We conclude that C consists of only the zero vector and does not have any extreme rays. The result then follows from Theorem 4.15. □

There is another corollary of Theorem 4.15 that deals with cones, and which is proved by noting that a cone can have no extreme points other than the zero vector.

Corollary 4.5 *Assume that the cone $C = \{\mathbf{x} \mid \mathbf{A}\mathbf{x} \geq \mathbf{0}\}$ is pointed. Then, every element of C can be expressed as a nonnegative linear combination of the extreme rays of C.*

Converse to the resolution theorem

Let us say that a set Q is *finitely generated* if it is specified in the form

$$Q = \left\{ \sum_{i=1}^{k}\lambda_i \mathbf{x}^i + \sum_{j=1}^{r}\theta_j \mathbf{w}^j \;\middle|\; \lambda_i \geq 0,\ \theta_j \geq 0,\ \sum_{i=1}^{k}\lambda_i = 1 \right\}, \qquad (4.6)$$

where $\mathbf{x}^1,\ldots,\mathbf{x}^k$ and $\mathbf{w}^1,\ldots,\mathbf{w}^r$ are some given elements of $\Re^n$. The resolution theorem states that a polyhedron with at least one extreme point is a finitely generated set (this is also true for general polyhedra; see Exercise 4.47). We now discuss a converse result, which states that every finitely generated set is a polyhedron.

As observed earlier, a finitely generated set Q can be viewed as the image of the polyhedron

$$H = \left\{ (\lambda_1,\ldots,\lambda_k,\theta_1,\ldots,\theta_r) \;\middle|\; \sum_{i=1}^{k}\lambda_i = 1,\ \lambda_i \geq 0,\ \theta_j \geq 0 \right\}$$

under a certain linear mapping. Thus, the results of Section 2.8 apply and establish that a finitely generated set is indeed a polyhedron. We record this result and also present a proof based on duality.

Theorem 4.16 *A finitely generated set is a polyhedron. In particular, the convex hull of finitely many vectors is a (bounded) polyhedron.*

Proof. Consider the linear programming problem (4.4) that was used in the proof of Theorem 4.15. A given vector $\mathbf{z}$ belongs to a finitely generated set Q of the form (4.6) if and only if the problem (4.4) has a feasible solution. Using duality, this is the case if and only if problem (4.5) has finite optimal cost. We convert problem (4.5) to standard form by introducing nonnegative variables $\mathbf{p}^+, \mathbf{p}^-, q^+, q^-$, such that $\mathbf{p} = \mathbf{p}^+ - \mathbf{p}^-$, and $q = q^+ - q^-$, as well as surplus variables. Since standard form polyhedra contain no lines, Theorem 4.13 shows that the optimal cost in the standard form problem is finite if and only if

$$(\mathbf{p}^+)'\mathbf{z} - (\mathbf{p}^-)'\mathbf{z} + q^+ - q^- \geq 0,$$

for each one of its finitely many extreme rays. Hence, $\mathbf{z} \in Q$ if and only if $\mathbf{z}$ satisfies a finite collection of linear inequalities. This shows that Q is a polyhedron. □

In conclusion, we have two ways of representing a polyhedron:

(a) in terms of a finite set of linear constraints;

(b) as a finitely generated set, in terms of its extreme points and extreme rays.

These two descriptions are mathematically equivalent, but can be quite different from a practical viewpoint. For example, we may be able to describe a polyhedron in terms of a small number of linear constraints. If on the other hand, this polyhedron has many extreme points, a description as a finitely generated set can be much more complicated. Furthermore, passing from one type of description to the other is, in general, a complicated computational task.

4.10 General linear programming duality*

In the definition of the dual problem (Section 4.2), we associated a dual variable p_i with each constraint of the form $\mathbf{a}_i'\mathbf{x} = b_i$, $\mathbf{a}_i'\mathbf{x} \geq b_i$, or $\mathbf{a}_i'\mathbf{x} \leq b_i$.

However, no dual variables were associated with constraints of the form $x_i \geq 0$ or $x_i \leq 0$. In the same spirit, and in a more general approach to linear programming duality, we can choose arbitrarily which constraints will be associated with price variables and which ones will not. In this section, we develop a general duality theorem that covers such a situation.

Consider the primal problem

$$\begin{aligned} \text{minimize} \quad & \mathbf{c}'\mathbf{x} \\ \text{subject to} \quad & \mathbf{A}\mathbf{x} \geq \mathbf{b} \\ & \mathbf{x} \in P, \end{aligned}$$

where P is the polyhedron

$$P = \{\mathbf{x} \mid \mathbf{D}\mathbf{x} \geq \mathbf{d}\}.$$

We associate a dual vector $\mathbf{p}$ with the constraint $\mathbf{A}\mathbf{x} \geq \mathbf{b}$. The constraint $\mathbf{x} \in P$ is a generalization of constraints of the form $x_i \geq 0$ or $x_i \leq 0$ and dual variables are not associated with it.

As in Section 4.1, we define the *dual objective* $g(\mathbf{p})$ by

$$g(\mathbf{p}) = \min_{\mathbf{x} \in P} \left[\mathbf{c}'\mathbf{x} + \mathbf{p}'(\mathbf{b} - \mathbf{A}\mathbf{x})\right]. \tag{4.7}$$

The *dual problem* is then defined as

$$\begin{aligned} \text{maximize} \quad & g(\mathbf{p}) \\ \text{subject to} \quad & \mathbf{p} \geq \mathbf{0}. \end{aligned}$$

We first provide a generalization of the weak duality theorem.

Theorem 4.17 (Weak duality) *If $\mathbf{x}$ is primal feasible ($\mathbf{A}\mathbf{x} \geq \mathbf{b}$ and $\mathbf{x} \in P$), and $\mathbf{p}$ is dual feasible ($\mathbf{p} \geq \mathbf{0}$), then $g(\mathbf{p}) \leq \mathbf{c}'\mathbf{x}$.*

Proof. If $\mathbf{x}$ and $\mathbf{p}$ are primal and dual feasible, respectively, then $\mathbf{p}'(\mathbf{b} - \mathbf{A}\mathbf{x}) \leq 0$, which implies that

$$\begin{aligned} g(\mathbf{p}) &= \min_{\mathbf{y} \in P} \left[\mathbf{c}'\mathbf{y} + \mathbf{p}'(\mathbf{b} - \mathbf{A}\mathbf{y})\right] \\ &\leq \mathbf{c}'\mathbf{x} + \mathbf{p}'(\mathbf{b} - \mathbf{A}\mathbf{x}) \\ &\leq \mathbf{c}'\mathbf{x}. \end{aligned}$$

□

We also have the following generalization of the strong duality theorem.

Theorem 4.18 (Strong duality) *If the primal problem has an optimal solution, so does the dual, and the respective optimal costs are equal.*

Proof. Since $P = \{\mathbf{x} \mid \mathbf{D}\mathbf{x} \geq \mathbf{d}\}$, the primal problem is of the form

$$\begin{aligned} \text{minimize} \quad & \mathbf{c}'\mathbf{x} \\ \text{subject to} \quad & \mathbf{A}\mathbf{x} \geq \mathbf{b} \\ & \mathbf{D}\mathbf{x} \geq \mathbf{d}, \end{aligned}$$

and we assume that it has an optimal solution. Its dual, which is

$$\begin{aligned} \text{maximize} \quad & \mathbf{p}'\mathbf{b} + \mathbf{q}'\mathbf{d} \\ \text{subject to} \quad & \mathbf{p}'\mathbf{A} + \mathbf{q}'\mathbf{D} = \mathbf{c}' \\ & \mathbf{p} \geq \mathbf{0} \\ & \mathbf{q} \geq \mathbf{0}, \end{aligned} \tag{4.8}$$

must then have the same optimal cost. For any fixed $\mathbf{p}$, the vector $\mathbf{q}$ should be chosen optimally in the problem (4.8). Thus, the dual problem (4.8) can also be written as

$$\begin{aligned} \text{maximize} \quad & \mathbf{p}'\mathbf{b} + f(\mathbf{p}) \\ \text{subject to} \quad & \mathbf{p} \geq \mathbf{0}, \end{aligned}$$

where $f(\mathbf{p})$ is the optimal cost in the problem

$$\begin{aligned} \text{maximize} \quad & \mathbf{q}'\mathbf{d} \\ \text{subject to} \quad & \mathbf{q}'\mathbf{D} = \mathbf{c}' - \mathbf{p}'\mathbf{A} \\ & \mathbf{q} \geq \mathbf{0}. \end{aligned} \tag{4.9}$$

[If the latter problem is infeasible, we set $f(\mathbf{p}) = -\infty$.] Using the strong duality theorem for problem (4.9), we obtain

$$f(\mathbf{p}) = \min_{\mathbf{D}\mathbf{x} \geq \mathbf{d}} (\mathbf{c}'\mathbf{x} - \mathbf{p}'\mathbf{A}\mathbf{x}).$$

We conclude that the dual problem (4.8) has the same optimal cost as the problem

$$\begin{aligned} \text{maximize} \quad & \mathbf{p}'\mathbf{b} + \min_{\mathbf{D}\mathbf{x} \geq \mathbf{d}} (\mathbf{c}'\mathbf{x} - \mathbf{p}'\mathbf{A}\mathbf{x}) \\ \text{subject to} \quad & \mathbf{p} \geq \mathbf{0}. \end{aligned}$$

By comparing with Eq. (4.7), we see that this is the same as maximizing $g(\mathbf{p})$ over all $\mathbf{p} \geq \mathbf{0}$. □

The idea of selectively assigning dual variables to some of the constraints is often used in order to treat "simpler" constraints differently than more "complex" ones, and has numerous applications in large scale optimization. (Applications to integer programming are discussed in Section 11.4.) Finally, let us point out that the approach in this section extends to certain nonlinear optimization problems. For example, if we replace the linear cost function $\mathbf{c}'\mathbf{x}$ by a general convex function $c(\mathbf{x})$, and the polyhedron P by a general convex set, we can again define the dual objective according to the formula

$$g(\mathbf{p}) = \min_{\mathbf{x}\in P} \Big[c(\mathbf{x}) + \mathbf{p}'(\mathbf{b} - \mathbf{A}\mathbf{x})\Big].$$

It turns out that the strong duality theorem remains valid for such nonlinear problems, under suitable technical conditions, but this lies beyond the scope of this book.

4.11 Summary

We summarize here the main ideas that have been developed in this chapter.

Given a (primal) linear programming problem, we can associate with it another (dual) linear programming problem, by following a set of mechanical rules. The definition of the dual problem is consistent, in the sense that the duals of equivalent primal problems are themselves equivalent.

Each dual variable is associated with a particular primal constraint and can be viewed as a penalty for violating that constraint. By replacing the primal constraints with penalty terms, we increase the set of available options, and this allows us to construct primal solutions whose cost is less than the optimal cost. In particular, every dual feasible vector leads to a lower bound on the optimal cost of the primal problem (this is the essence of the weak duality theorem). The maximization in the dual problem is then a search for the tightest such lower bound. The strong duality theorem asserts that the tightest such lower bound is equal to the optimal primal cost.

An optimal dual variable can also be interpreted as a marginal cost, that is, as the rate of change of the optimal primal cost when we perform a small perturbation of the right-hand side vector $\mathbf{b}$, assuming nondegeneracy.

A useful relation between optimal primal and dual solutions is provided by the complementary slackness conditions. Intuitively, these conditions require that any constraint that is inactive at an optimal solution carries a zero price, which is compatible with the interpretation of prices as marginal costs.

We saw that every basis matrix in a standard form problem determines not only a primal basic solution, but also a basic dual solution. This observation is at the heart of the dual simplex method. This method is similar to the primal simplex method in that it generates a sequence of primal basic solutions, together with an associated sequence of dual basic solutions. It is different, however, in that the dual basic solutions are dual feasible, with ever improving costs, while the primal basic solutions are infeasible (except for the last one). We developed the dual simplex method by simply describing its mechanics and by providing an algebraic justification.

Nevertheless, the dual simplex method also has a geometric interpretation. It keeps moving from one dual basic feasible solution to an adjacent one and, in this respect, it is similar to the primal simplex method applied to the dual problem.

All of duality theory can be developed by exploiting the termination conditions of the simplex method, and this was our initial approach to the subject. We also pursued an alternative line of development that proceeded from first principles and used geometric arguments. This is a more direct and more general approach, but requires more abstract reasoning.

Duality theory provided us with some powerful tools based on which we were able to enhance our geometric understanding of polyhedra. We derived a few theorems of the alternative (like Farkas' lemma), which are surprisingly powerful and have applications in a wide variety of contexts. In fact, Farkas' lemma can be viewed as the core of linear programming duality theory. Another major result that we derived is the resolution theorem, which allows us to express any element of a nonempty polyhedron with at least one extreme point as a convex combination of its extreme points plus a nonnegative linear combination of its extreme rays; in other words, every polyhedron is "finitely generated." The converse is also true, and every finitely generated set is a polyhedron (can be represented in terms of linear inequality constraints). Results of this type play a key role in confirming our intuitive geometric understanding of polyhedra and linear programming. They allow us to develop alternative views of certain situations and lead to deeper understanding. Many such results have an "obvious" geometric content and are often taken for granted. Nevertheless, as we have seen, rigorous proofs can be quite elaborate.

4.12 Exercises

Exercise 4.1 Consider the linear programming problem:

$$\begin{array}{lrcrcrcrcl}
\text{minimize} & x_1 & - & x_2 & & & & & & \\
\text{subject to} & 2x_1 & + & 3x_2 & - & x_3 & + & x_4 & \leq & 0 \\
& 3x_1 & + & x_2 & + & 4x_3 & - & 2x_4 & \geq & 3 \\
& -x_1 & - & x_2 & + & 2x_3 & + & x_4 & = & 6 \\
& x_1 \leq 0 & & & & & & & & \\
& x_2, x_3 \geq 0. & & & & & & & &
\end{array}$$

Write down the corresponding dual problem.

Exercise 4.2 Consider the primal problem

$$\begin{array}{ll}
\text{minimize} & \mathbf{c}'\mathbf{x} \\
\text{subject to} & \mathbf{A}\mathbf{x} \geq \mathbf{b} \\
& \mathbf{x} \geq \mathbf{0}.
\end{array}$$

Form the dual problem and convert it into an equivalent minimization problem. Derive a set of conditions on the matrix $\mathbf{A}$ and the vectors $\mathbf{b}$, $\mathbf{c}$, under which the dual is identical to the primal, and construct an example in which these conditions are satisfied.

Exercise 4.3 The purpose of this exercise is to show that solving linear programming problems is no harder than solving systems of linear inequalities.

Suppose that we are given a subroutine which, given a system of linear inequality constraints, either produces a solution or decides that no solution exists. Construct a simple algorithm that uses a single call to this subroutine and which finds an optimal solution to any linear programming problem that has an optimal solution.

Exercise 4.4 Let $\mathbf{A}$ be a symmetric square matrix. Consider the linear programming problem

$$\begin{array}{ll} \text{minimize} & \mathbf{c}'\mathbf{x} \\ \text{subject to} & \mathbf{A}\mathbf{x} \geq \mathbf{c} \\ & \mathbf{x} \geq \mathbf{0}. \end{array}$$

Prove that if $\mathbf{x}^*$ satisfies $\mathbf{A}\mathbf{x}^* = \mathbf{c}$ and $\mathbf{x}^* \geq \mathbf{0}$, then $\mathbf{x}^*$ is an optimal solution.

Exercise 4.5 Consider a linear programming problem in standard form and assume that the rows of $\mathbf{A}$ are linearly independent. For each one of the following statements, provide either a proof or a counterexample.

(a) Let $\mathbf{x}^*$ be a basic feasible solution. Suppose that for every basis corresponding to $\mathbf{x}^*$, the associated basic solution to the dual is infeasible. Then, the optimal cost must be strictly less that $\mathbf{c}'\mathbf{x}^*$.

(b) The dual of the auxiliary primal problem considered in Phase I of the simplex method is always feasible.

(c) Let p_i be the dual variable associated with the ith equality constraint in the primal. Eliminating the ith primal equality constraint is equivalent to introducing the additional constraint $p_i = 0$ in the dual problem.

(d) If the unboundedness criterion in the primal simplex algorithm is satisfied, then the dual problem is infeasible.

Exercise 4.6* (Duality in Chebychev approximation) Let $\mathbf{A}$ be an $m \times n$ matrix and let $\mathbf{b}$ be a vector in $\Re^m$. We consider the problem of minimizing $\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_\infty$ over all $\mathbf{x} \in \Re^n$. Here $\|\cdot\|_\infty$ is the vector norm defined by $\|\mathbf{y}\|_\infty = \max_i |y_i|$. Let v be the value of the optimal cost.

(a) Let $\mathbf{p}$ be any vector in $\Re^m$ that satisfies $\sum_{i=1}^m |p_i| = 1$ and $\mathbf{p}'\mathbf{A} = \mathbf{0}'$. Show that $\mathbf{p}'\mathbf{b} \leq v$.

(b) In order to obtain the best possible lower bound of the form considered in part (a), we form the linear programming problem

$$\begin{array}{ll} \text{maximize} & \mathbf{p}'\mathbf{b} \\ \text{subject to} & \mathbf{p}'\mathbf{A} = \mathbf{0}' \\ & \displaystyle\sum_{i=1}^m |p_i| \leq 1. \end{array}$$

Show that the optimal cost in this problem is equal to v.

Exercise 4.7 (Duality in piecewise linear convex optimization) Consider the problem of minimizing $\max_{i=1,\ldots,m}(\mathbf{a}_i'\mathbf{x} - b_i)$ over all $\mathbf{x} \in \Re^n$. Let v be the value of the optimal cost, assumed finite. Let $\mathbf{A}$ be the matrix with rows $\mathbf{a}_1, \ldots, \mathbf{a}_m$, and let $\mathbf{b}$ be the vector with components $b_1, \ldots, b_m$.

(a) Consider any vector $\mathbf{p} \in \Re^m$ that satisfies $\mathbf{p}'\mathbf{A} = \mathbf{0}'$, $\mathbf{p} \geq \mathbf{0}$, and $\sum_{i=1}^m p_i = 1$. Show that $-\mathbf{p}'\mathbf{b} \leq v$.

(b) In order to obtain the best possible lower bound of the form considered in part (a), we form the linear programming problem

$$\begin{array}{ll} \text{maximize} & -\mathbf{p}'\mathbf{b} \\ \text{subject to} & \mathbf{p}'\mathbf{A} = \mathbf{0}' \\ & \mathbf{p}'\mathbf{e} = 1 \\ & \mathbf{p} \geq \mathbf{0}, \end{array}$$

where $\mathbf{e}$ is the vector with all components equal to 1. Show that the optimal cost in this problem is equal to v.

Exercise 4.8 Consider the linear programming problem of minimizing $\mathbf{c}'\mathbf{x}$ subject to $\mathbf{A}\mathbf{x} = \mathbf{b}$, $\mathbf{x} \geq \mathbf{0}$. Let $\mathbf{x}^*$ be an optimal solution, assumed to exist, and let $\mathbf{p}^*$ be an optimal solution to the dual.

(a) Let $\tilde{\mathbf{x}}$ be an optimal solution to the primal, when $\mathbf{c}$ is replaced by some $\tilde{\mathbf{c}}$. Show that $(\tilde{\mathbf{c}} - \mathbf{c})'(\tilde{\mathbf{x}} - \mathbf{x}^*) \leq 0$.

(b) Let the cost vector be fixed at $\mathbf{c}$, but suppose that we now change $\mathbf{b}$ to $\tilde{\mathbf{b}}$, and let $\tilde{\mathbf{x}}$ be a corresponding optimal solution to the primal. Prove that $(\mathbf{p}^*)'(\tilde{\mathbf{b}} - \mathbf{b}) \leq \mathbf{c}'(\tilde{\mathbf{x}} - \mathbf{x}^*)$.

Exercise 4.9 (Back-propagation of dual variables in a multiperiod problem) A company makes a product that can be either sold or stored to meet future demand. Let $t = 1, \ldots, T$ denote the periods of the planning horizon. Let b_t be the production volume during period t, which is assumed to be known in advance. During each period t, a quantity x_t of the product is sold, at a unit price of d_t. Furthermore, a quantity y_t can be sent to long-term storage, at a unit transportation cost of c. Alternatively, a quantity w_t can be retrieved from storage, at zero cost. We assume that when the product is prepared for long-term storage, it is partly damaged, and only a fraction f of the total survives. Demand is assumed to be unlimited. The main question is whether it is profitable to store some of the production, in anticipation of higher prices in the future. This leads us to the following problem, where z_t stands for the amount kept in long-term storage, at the end of period t:

$$\begin{array}{lll} \text{maximize} & \displaystyle\sum_{t=1}^T \alpha^{t-1}(d_t x_t - c y_t) + \alpha^T d_{T+1} z_T & \\ \text{subject to} & x_t + y_t - w_t = b_t, & t = 1, \ldots, T, \\ & z_t + w_t - z_{t-1} - f y_t = 0, & t = 1, \ldots, T, \\ & z_0 = 0, & \\ & x_t, y_t, w_t, z_t \geq 0. & \end{array}$$

Here, d_{T+1} is the salvage prive for whatever inventory is left at the end of period T. Furthermore, α is a discount factor, with $0 < \alpha < 1$, reflecting the fact that future revenues are valued less than current ones.

(a) Let p_t and q_t be dual variables associated with the first and second equality constraint, respectively. Write down the dual problem.

(b) Assume that $0 < f < 1$, $b_t \geq 0$, and $c \geq 0$. Show that the following formulae provide an optimal solution to the dual problem:

$$\begin{aligned} q_T &= \alpha^T d_{T+1}, \\ p_T &= \max\left\{\alpha^{T-1} d_T,\ f q_T - \alpha^{T-1} c\right\}, \\ q_t &= \max\left\{q_{t+1}, \alpha^{t-1} d_t\right\}, \qquad & t = 1, \ldots, T-1, \\ p_t &= \max\left\{\alpha^{t-1} d_t,\ f q_t - \alpha^{t-1} c\right\}, \qquad & t = 1, \ldots, T-1. \end{aligned}$$

(c) Explain how the result in part (b) can be used to compute an optimal solution to the original problem. Primal and dual nondegeneracy can be assumed.

Exercise 4.10 (Saddle points of the Lagrangean) Consider the standard form problem of minimizing $\mathbf{c}'\mathbf{x}$ subject to $\mathbf{A}\mathbf{x} = \mathbf{b}$ and $\mathbf{x} \geq \mathbf{0}$. We define the *Lagrangean* by

$$L(\mathbf{x}, \mathbf{p}) = \mathbf{c}'\mathbf{x} + \mathbf{p}'(\mathbf{b} - \mathbf{A}\mathbf{x}).$$

Consider the following "game": player 1 chooses some $\mathbf{x} \geq \mathbf{0}$, and player 2 chooses some $\mathbf{p}$; then, player 1 pays to player 2 the amount $L(\mathbf{x}, \mathbf{p})$. Player 1 would like to minimize $L(\mathbf{x}, \mathbf{p})$, while player 2 would like to maximize it.

A pair $(\mathbf{x}^*, \mathbf{p}^*)$, with $\mathbf{x}^* \geq \mathbf{0}$, is called an *equilibrium* point (or a *saddle point*, or a *Nash equilibrium*) if

$$L(\mathbf{x}^*, \mathbf{p}) \leq L(\mathbf{x}^*, \mathbf{p}^*) \leq L(\mathbf{x}, \mathbf{p}^*), \qquad \forall\ \mathbf{x} \geq \mathbf{0},\ \forall\ \mathbf{p}.$$

(Thus, we have an equilibrium if no player is able to improve her performance by unilaterally modifying her choice.)

Show that a pair $(\mathbf{x}^*, \mathbf{p}^*)$ is an equilibrium if and only if $\mathbf{x}^*$ and $\mathbf{p}^*$ are optimal solutions to the standard form problem under consideration and its dual, respectively.

Exercise 4.11 Consider a linear programming problem in standard form which is infeasible, but which becomes feasible and has finite optimal cost when the last equality constraint is omitted. Show that the dual of the original (infeasible) problem is feasible and the optimal cost is infinite.

Exercise 4.12* (Degeneracy and uniqueness – I) Consider a general linear programming problem and suppose that we have a nondegenerate basic feasible solution to the primal. Show that the complementary slackness conditions lead to a system of equations for the dual vector that has a unique solution.

Exercise 4.13* (Degeneracy and uniqueness – II) Consider the following pair of problems that are duals of each other:

$$\begin{array}{llcl} \text{minimize} & \mathbf{c}'\mathbf{x} & & \\ \text{subject to} & \mathbf{A}\mathbf{x} & = & \mathbf{b} \\ & \mathbf{x} & \geq & \mathbf{0}, \end{array} \qquad\qquad \begin{array}{llcl} \text{maximize} & \mathbf{p}'\mathbf{b} & & \\ \text{subject to} & \mathbf{p}'\mathbf{A} & \leq & \mathbf{c}'. \end{array}$$

(a) Prove that if one problem has a nondegenerate and unique optimal solution, so does the other.

(b) Suppose that we have a nondegenerate optimal basis for the primal and that the reduced cost for one of the basic variables is zero. What does the result of part (a) imply? Is it true that there must exist another optimal basis?

Exercise 4.14 (Degeneracy and uniqueness – III) Give an example in which the primal problem has a degenerate optimal basic feasible solution, but the dual has a unique optimal solution. (The example need not be in standard form.)

Exercise 4.15 (Degeneracy and uniqueness – IV) Consider the problem

$$\begin{array}{ll} \text{minimize} & x_2 \\ \text{subject to} & x_2 = 1 \\ & x_1 \geq 0 \\ & x_2 \geq 0. \end{array}$$

Write down its dual. For both the primal and the dual problem determine whether they have unique optimal solutions and whether they have nondegenerate optimal solutions. Is this example in agreement with the statement that nondegeneracy of an optimal basic feasible solution in one problem implies uniqueness of optimal solutions for the other? Explain.

Exercise 4.16 Give an example of a pair (primal and dual) of linear programming problems, both of which have multiple optimal solutions.

Exercise 4.17 This exercise is meant to demonstrate that knowledge of a primal optimal solution does not necessarily contain information that can be exploited to determine a dual optimal solution. In particular, determining an optimal solution to the dual is as hard as solving a system of linear inequalities, even if an optimal solution to the primal is available.

Consider the problem of minimizing $\mathbf{c}'\mathbf{x}$ subject to $\mathbf{A}\mathbf{x} \geq \mathbf{0}$, and suppose that we are told that the zero vector is optimal. Let the dimensions of $\mathbf{A}$ be $m \times n$, and suppose that we have an algorithm that determines a dual optimal solution and whose running time $O\big((m+n)^k\big)$, for some constant k. (Note that if $\mathbf{x} = \mathbf{0}$ is not an optimal primal solution, the dual has no feasible solution, and we assume that in this case our algorithm exits with an error message.) Assuming the availability of the above algorithm, construct a new algorithm that takes as input a system of m linear inequalities in n variables, runs for $O\big((m+n)^k\big)$ time, and either finds a feasible solution or determines that no feasible solution exists.

Exercise 4.18 Consider a problem in standard form. Suppose that the matrix $\mathbf{A}$ has dimensions $m \times n$ and its rows are linearly independent. Suppose that all basic solutions to the primal and to the dual are nondegenerate. Let $\mathbf{x}$ be a feasible solution to the primal and let $\mathbf{p}$ be a dual vector (not necessarily feasible), such that the pair $(\mathbf{x}, \mathbf{p})$ satisfies complementary slackness.

(a) Show that there exist m columns of $\mathbf{A}$ that are linearly independent and such that the corresponding components of $\mathbf{x}$ are all positive.

(b) Show that $\mathbf{x}$ and $\mathbf{p}$ are basic solutions to the primal and the dual, respectively.

(c) Show that the result of part (a) is false if the nondegeneracy assumption is removed.

Exercise 4.19 Let $P = \{\mathbf{x} \in \Re^n \mid \mathbf{Ax} = \mathbf{b},\ \mathbf{x} \geq \mathbf{0}\}$ be a nonempty polyhedron, and let m be the dimension of the vector $\mathbf{b}$. We call x_j a *null variable* if $x_j = 0$ whenever $\mathbf{x} \in P$.

(a) Suppose that there exists some $\mathbf{p} \in \Re^m$ for which $\mathbf{p}'\mathbf{A} \geq \mathbf{0}'$, $\mathbf{p}'\mathbf{b} = 0$, and such that the jth component of $\mathbf{p}'\mathbf{A}$ is positive. Prove that x_j is a null variable.

(b) Prove the converse of (a): if x_j is a null variable, then there exists some $\mathbf{p} \in \Re^m$ with the properties stated in part (a).

(c) If x_j is not a null variable, then by definition, there exists some $\mathbf{y} \in P$ for which $y_j > 0$. Use the results in parts (a) and (b) to prove that there exist $\mathbf{x} \in P$ and $\mathbf{p} \in \Re^m$ such that:

$$\mathbf{p}'\mathbf{A} \geq \mathbf{0}', \qquad \mathbf{p}'\mathbf{b} = 0, \qquad \mathbf{x} + \mathbf{A}'\mathbf{p} > \mathbf{0}.$$

Exercise 4.20 * (Strict complementary slackness)

(a) Consider the following linear programming problem and its dual

$$\begin{array}{llll} \text{minimize} & \mathbf{c}'\mathbf{x} & \qquad \text{maximize} & \mathbf{p}'\mathbf{b} \\ \text{subject to} & \mathbf{Ax} = \mathbf{b} & \qquad \text{subject to} & \mathbf{p}'\mathbf{A} \leq \mathbf{c}', \\ & \mathbf{x} \geq \mathbf{0}, & & \end{array}$$

and assume that both problems have an optimal solution. Fix some j. Suppose that every optimal solution to the primal satisfies $x_j = 0$. Show that there exists an optimal solution $\mathbf{p}$ to the dual such that $\mathbf{p}'\mathbf{A}_j < c_j$. (Here, $\mathbf{A}_j$ is the jth column of $\mathbf{A}$.) *Hint:* Let d be the optimal cost. Consider the problem of minimizing $-x_j$ subject to $\mathbf{Ax} = \mathbf{b}$, $\mathbf{x} \geq \mathbf{0}$, and $-\mathbf{c}'\mathbf{x} \geq -d$, and form its dual.

(b) Show that there exist optimal solutions $\mathbf{x}$ and $\mathbf{p}$ to the primal and to the dual, respectively, such that for every j we have either $x_j > 0$ or $\mathbf{p}'\mathbf{A}_j < c_j$. *Hint:* Use part (a) for each j, and then take the average of the vectors obtained.

(c) Consider now the following linear programming problem and its dual:

$$\begin{array}{llll} \text{minimize} & \mathbf{c}'\mathbf{x} & \qquad \text{maximize} & \mathbf{p}'\mathbf{b} \\ \text{subject to} & \mathbf{Ax} \geq \mathbf{b} & \qquad \text{subject to} & \mathbf{p}'\mathbf{A} \leq \mathbf{c}' \\ & \mathbf{x} \geq \mathbf{0}, & & \mathbf{p} \geq \mathbf{0}. \end{array}$$

Assume that both problems have an optimal solution. Show that there exist optimal solutions to the primal and to the dual, respectively, that satisfy *strict complementary slackness*, that is:

(i) For every j we have either $x_j > 0$ or $\mathbf{p}'\mathbf{A}_j < c_j$.

(ii) For every i, we have either $\mathbf{a}_i'\mathbf{x} > b_i$ or $p_i > 0$. (Here, $\mathbf{a}_i'$ is the ith row of $\mathbf{A}$.) *Hint:* Convert the primal to standard form and apply part (b).

(d) Consider the linear programming problem

$$\begin{array}{ll} \text{minimize} & 5x_1 + 5x_2 \\ \text{subject to} & x_1 + x_2 \geq 2 \\ & 2x_1 - x_2 \geq 0 \\ & x_1, x_2 \geq 0. \end{array}$$

Does the optimal primal solution $(2/3, 4/3)$, together with the corresponding dual optimal solution, satisfy strict complementary slackness? Determine all primal and dual optimal solutions and identify the set of *all* strictly complementary pairs.

Exercise 4.21 * (Clark's theorem) Consider the following pair of linear programming problems:

$$\begin{array}{llll} \text{minimize} & \mathbf{c}'\mathbf{x} & \qquad \text{maximize} & \mathbf{p}'\mathbf{b} \\ \text{subject to} & \mathbf{Ax} \geq \mathbf{b} & \qquad \text{subject to} & \mathbf{p}'\mathbf{A} \leq \mathbf{c}' \\ & \mathbf{x} \geq \mathbf{0}, & & \mathbf{p} \geq \mathbf{0}. \end{array}$$

Suppose that at least one of these two problems has a feasible solution. Prove that the set of feasible solutions to at least one of the two problems is unbounded. *Hint:* Interpret boundedness of a set in terms of the finiteness of the optimal cost of some linear programming problem.

Exercise 4.22 Consider the dual simplex method applied to a standard form problem with linearly independent rows. Suppose that we have a basis which is primal infeasible, but dual feasible, and let i be such that $x_{B(i)} < 0$. Suppose that all entries in the ith row in the tableau (other than $x_{B(i)}$) are nonnegative. Show that the optimal dual cost is $+\infty$.

Exercise 4.23 Describe in detail the mechanics of a revised dual simplex method that works in terms of the inverse basis matrix $\mathbf{B}^{-1}$ instead of the full simplex tableau.

Exercise 4.24 Consider the lexicographic pivoting rule for the dual simplex method and suppose that the algorithm is initialized with each column of the tableau being lexicographically positive. Prove that the dual simplex method does not cycle.

Exercise 4.25 This exercise shows that if we bring the dual problem into standard form and then apply the primal simplex method, the resulting algorithm is not identical to the dual simplex method.

Consider the following standard form problem and its dual.

$$\begin{array}{llll} \text{minimize} & x_1 + x_2 & \qquad \text{maximize} & p_1 + p_2 \\ \text{subject to} & x_1 = 1 & \qquad \text{subject to} & p_1 \leq 1 \\ & x_2 = 1 & & p_2 \leq 1. \\ & x_1, x_2 \geq 0 & & \end{array}$$

Here, there is only one possible basis and the dual simplex method must terminate immediately. Show that if the dual problem is converted into standard form and the primal simplex method is applied to it, one or more changes of basis may be required.

Exercise 4.26 Let $\mathbf{A}$ be a given matrix. Show that exactly one of the following alternatives must hold.

(a) There exists some $\mathbf{x} \neq \mathbf{0}$ such that $\mathbf{A}\mathbf{x} = \mathbf{0}$, $\mathbf{x} \geq \mathbf{0}$.

(b) There exists some $\mathbf{p}$ such that $\mathbf{p}'\mathbf{A} > \mathbf{0}'$.

Exercise 4.27 Let $\mathbf{A}$ be a given matrix. Show that the following two statements are equivalent.

(a) Every vector such that $\mathbf{A}\mathbf{x} \geq \mathbf{0}$ and $\mathbf{x} \geq \mathbf{0}$ must satisfy $x_1 = 0$.

(b) There exists some $\mathbf{p}$ such that $\mathbf{p}'\mathbf{A} \leq \mathbf{0}$, $\mathbf{p} \geq \mathbf{0}$, and $\mathbf{p}'\mathbf{A}_1 < 0$, where $\mathbf{A}_1$ is the first column of $\mathbf{A}$.

Exercise 4.28 Let $\mathbf{a}$ and $\mathbf{a}_1, \ldots, \mathbf{a}_m$ be given vectors in $\Re^n$. Prove that the following two statements are equivalent:

(a) For all $\mathbf{x} \geq \mathbf{0}$, we have $\mathbf{a}'\mathbf{x} \leq \max_i \mathbf{a}_i'\mathbf{x}$.

(b) There exist nonnegative coefficients λ_i that sum to 1 and such that $\mathbf{a} \leq \sum_{i=1}^m \lambda_i \mathbf{a}_i$.

Exercise 4.29 (Inconsistent systems of linear inequalities) Let $\mathbf{a}_1, \ldots, \mathbf{a}_m$ be some vectors in $\Re^n$, with $m > n + 1$. Suppose that the system of inequalities $\mathbf{a}_i'\mathbf{x} \geq b_i$, $i = 1, \ldots, m$, does not have any solutions. Show that we can choose $n + 1$ of these inequalities, so that the resulting system of inequalities has no solutions.

Exercise 4.30 (Helly's theorem)

(a) Let $\mathcal{F}$ be a finite family of polyhedra in $\Re^n$ such that every $n+1$ polyhedra in $\mathcal{F}$ have a point in common. Prove that all polyhedra in $\mathcal{F}$ have a point in common. *Hint:* Use the result in Exercise 4.29.

(b) For $n = 2$, part (a) asserts that the polyhedra $P_1, P_2, \ldots, P_K$ $(K \geq 3)$ in the plane have a point in common if and only if every three of them have a point in common. Is the result still true with "three" replaced by "two"?

Exercise 4.31 (Unit eigenvectors of stochastic matrices) We say that an $n \times n$ matrix $\mathbf{P}$, with entries p_{ij}, is *stochastic* if all of its entries are nonnegative and

$$\sum_{j=1}^n p_{ij} = 1, \qquad \forall\, i,$$

that is, the sum of the entries of each row is equal to 1.

Use duality to show that if $\mathbf{P}$ is a stochastic matrix, then the system of equations

$$\mathbf{p}'\mathbf{P} = \mathbf{p}', \qquad \mathbf{p} \geq \mathbf{0},$$

has a nonzero solution. (Note that the vector $\mathbf{p}$ can be normalized so that its components sum to one. Then, the result in this exercise establishes that every finite state Markov chain has an invariant probability distribution.)

Exercise 4.32 * (Leontief systems and Samuelson's substitution theorem) A *Leontief matrix* is an $m \times n$ matrix $\mathbf{A}$ in which every column has at most one positive element. For an interpretation, each column $\mathbf{A}_j$ corresponds to a production process. If a_{ij} is negative, $|a_{ij}|$ represents the amount of goods of type i consumed by the process. If a_{ij} is positive, it represents the amount of goods of type i produced by the process. If x_j is the intensity with which process j is used, then $\mathbf{A}\mathbf{x}$ represents the net output of the different goods. The matrix $\mathbf{A}$ is called *productive* if there exists some $\mathbf{x} \geq \mathbf{0}$ such that $\mathbf{A}\mathbf{x} > \mathbf{0}$.

(a) Let $\mathbf{A}$ be a square productive Leontief matrix ($m = n$). Show that every vector $\mathbf{z}$ that satisfies $\mathbf{A}\mathbf{z} \geq \mathbf{0}$ must be nonnegative. *Hint:* If $\mathbf{z}$ satisfies $\mathbf{A}\mathbf{z} \geq \mathbf{0}$ but has a negative component, consider the smallest nonnegative θ such that some component of $\mathbf{x} + \theta\mathbf{z}$ becomes zero, and derive a contradiction.

(b) Show that every square productive Leontief matrix is invertible and that all entries of the inverse matrix are nonnegative. *Hint:* Use the result in part (a).

(c) We now consider the general case where $n \geq m$, and we introduce a constraint of the form $\mathbf{e}'\mathbf{x} \leq 1$, where $\mathbf{e} = (1, \ldots, 1)$. (Such a constraint could capture, for example, a bottleneck due to the finiteness of the labor force.) An "output" vector $\mathbf{y} \in \Re^m$ is said to be *achievable* if $\mathbf{y} \geq \mathbf{0}$ and there exists some $\mathbf{x} \geq \mathbf{0}$ such that $\mathbf{A}\mathbf{x} = \mathbf{y}$ and $\mathbf{e}'\mathbf{y} \leq 1$. An achievable vector $\mathbf{y}$ is said to be *efficient* if there exists no achievable vector $\mathbf{z}$ such that $\mathbf{z} \geq \mathbf{y}$ and $\mathbf{z} \neq \mathbf{y}$. (Intuitively, an output vector $\mathbf{y}$ which is not efficient can be improved upon and is therefore uninteresting.) Suppose that $\mathbf{A}$ is productive. Show that there exists a positive efficient vector $\mathbf{y}$. *Hint:* Given a positive achievable vector $\mathbf{y}^*$, consider maximizing $\sum_i y_i$ over all achievable vectors $\mathbf{y}$ that are larger than $\mathbf{y}^*$.

(d) Suppose that $\mathbf{A}$ is productive. Show that there exists a set of m production processes that are capable of generating all possible efficient output vectors $\mathbf{y}$. That is, there exist indices $B(1), \ldots, B(m)$, such that every efficient output vector $\mathbf{y}$ can be expressed in the form $\mathbf{y} = \sum_{i=1}^m \mathbf{A}_{B(i)} x_{B(i)}$, for some nonnegative coefficients $x_{B(i)}$ whose sum is bounded by 1. *Hint:* Consider the problem of minimizing $\mathbf{e}'\mathbf{x}$ subject to $\mathbf{A}\mathbf{x} = \mathbf{y}$, $\mathbf{x} \geq \mathbf{0}$, and show that we can use the same optimal basis for all efficient vectors $\mathbf{y}$.

Exercise 4.33 (Options pricing) Consider a market that operates for a single period, and which involves three assets: a stock, a bond, and an option. Let S be the price of the stock, in the beginning of the period. Its price $\overline{S}$ at the end of the period is random and is assumed to be equal to either Su, with probability β, or Sd, with probability $1 - \beta$. Here u and d are scalars that satisfy $d < 1 < u$. Bonds are assumed riskless. Investing one dollar in a bond results in a payoff of r, at the end of the period. (Here, r is a scalar greater than 1.) Finally, the option gives us the right to purchase, at the end of the period, one stock at a fixed price of K. If the realized price $\overline{S}$ of the stock is greater than K, we exercise the option and then immediately sell the stock in the stock market, for a payoff of $\overline{S} - K$. If on the other hand we have $\overline{S} < K$, there is no advantage in exercising the option, and we receive zero payoff. Thus, the value of the option at the end of the period is equal to $\max\{0, \overline{S} - K\}$. Since the option is itself an asset, it should have a value in the beginning of the time period. Show that under the absence of arbitrage condition, the value of the option must be equal to

$$\gamma \max\{0, Su - K\} + \delta \max\{0, Sd - K\},$$

where γ and δ are a solution to the following system of linear equations:

$$\begin{aligned} u\gamma \;+\; d\delta &= 1 \\ \gamma \;+\; \delta &= \frac{1}{r}. \end{aligned}$$

Hint: Write down the payoff matrix $\mathbf{R}$ and use Theorem 4.8.

Exercise 4.34 (Finding separating hyperplanes) Consider a polyhedron P that has at least one extreme point.

(a) Suppose that we are given the extreme points $\mathbf{x}^i$ and a complete set of extreme rays $\mathbf{w}^j$ of P. Create a linear programming problem whose solution provides us with a separating hyperplane that separates P from the origin, or allows us to conclude that none exists.

(b) Suppose now that P is given to us in the form $P = \{\mathbf{x} \mid \mathbf{a}_i'\mathbf{x} \geq b_i,\ i = 1, \ldots, m\}$. Suppose that $\mathbf{0} \notin P$. Explain how a separating hyperplane can be found.

Exercise 4.35 (Separation of disjoint polyhedra) Consider two nonempty polyhedra $P = \{\mathbf{x} \in \Re^n \mid \mathbf{A}\mathbf{x} \leq \mathbf{b}\}$ and $Q = \{\mathbf{x} \in \Re^n \mid \mathbf{D}\mathbf{x} \leq \mathbf{d}\}$. We are interested in finding out whether the two polyhedra have a point in common.

(a) Devise a linear programming problem such that: if $P \cap Q$ is nonempty, it returns a point in $P \cap Q$; if $P \cap Q$ is empty, the linear programming problem is infeasible.

(b) Suppose that $P \cap Q$ is empty. Use the dual of the problem you have constructed in part (a) to show that there exists a vector $\mathbf{c}$ such that $\mathbf{c}'\mathbf{x} < \mathbf{c}'\mathbf{y}$ for all $\mathbf{x} \in P$ and $\mathbf{y} \in Q$.

Exercise 4.36 (Containment of polyhedra)

(a) Let P and Q be two polyhedra in $\Re^n$ described in terms of linear inequality constraints. Devise an algorithm that decides whether P is a subset of Q.

(b) Repeat part (a) if the polyhedra are described in terms of their extreme points and extreme rays.

Exercise 4.37 (Closedness of finitely generated cones) Let $\mathbf{A}_1, \ldots, \mathbf{A}_n$ be given vectors in $\Re^m$. Consider the cone $C = \left\{\sum_{i=1}^n \mathbf{A}_i x_i \mid x_i \geq 0\right\}$ and let $\mathbf{y}^k$, $k = 1, 2, \ldots$, be a sequence of elements of C that converges to some $\mathbf{y}$. Show that $\mathbf{y} \in C$ (and hence C is closed), using the following argument. With $\mathbf{y}$ fixed as above, consider the problem of minimizing $\|\mathbf{y} - \sum_{i=1}^n \mathbf{A}_i x_i\|_\infty$, subject to the constraints $x_1, \ldots, x_n \geq 0$. Here $\|\cdot\|_\infty$ stands for the maximum norm, defined by $\|\mathbf{x}\|_\infty = \max_i |x_i|$. Explain why the above minimization problem has an optimal solution, find the value of the optimal cost, and prove that $\mathbf{y} \in C$.

Exercise 4.38 (From Farkas' lemma to duality) Use Farkas' lemma to prove the duality theorem for a linear programming problem involving constraints of the form $\mathbf{a}_i'\mathbf{x} = b_i$, $\mathbf{a}_i'\mathbf{x} \geq b_i$, and nonnegativity constraints for some of the variables x_j. *Hint:* Start by deriving the form of the set of feasible directions at an optimal solution.

Exercise 4.39 (Extreme rays of cones) Let us define a nonzero element $\mathbf{d}$ of a pointed polyhedral cone C to be an *extreme ray* if it has the following property: if there exist vectors $\mathbf{f} \in C$ and $\mathbf{g} \in C$ and some $\lambda \in (0, 1)$ satisfying $\mathbf{d} = \mathbf{f} + \mathbf{g}$, then both $\mathbf{f}$ and $\mathbf{g}$ are scalar multiples of $\mathbf{d}$. Prove that this definition of extreme rays is equivalent to Definition 4.2.

Exercise 4.40 (Extreme rays of a cone are extreme points of its sections) Consider the cone $C = \{\mathbf{x} \in \Re^n \mid \mathbf{a}_i'\mathbf{x} \geq 0,\ i = 1, \ldots, m\}$ and assume that the first n constraint vectors $\mathbf{a}_1, \ldots, \mathbf{a}_n$ are linearly independent. For any nonnegative scalar r, we define the polyhedron P_r by

$$P_r = \left\{\mathbf{x} \in C \;\middle|\; \sum_{i=1}^n \mathbf{a}_i'\mathbf{x} = r\right\}.$$

(a) Show that the polyhedron P_r is bounded for every $r \geq 0$.

(b) Let $r > 0$. Show that a vector $\mathbf{x} \in P_r$ is an extreme point of P_r if and only if $\mathbf{x}$ is an extreme ray of the cone C.

Exercise 4.41 (Carathéodory's theorem) Show that every element $\mathbf{x}$ of a bounded polyhedron $P \subset \Re^n$ can be expressed as a convex combination of at most $n + 1$ extreme points of P. *Hint:* Consider an extreme point of the set of all possible representations of $\mathbf{x}$.

Exercise 4.42 (Problems with side constraints) Consider the linear programming problem of minimizing $\mathbf{c}'\mathbf{x}$ over a bounded polyhedron $P \subset \Re^n$ and subject to additional constraints $\mathbf{a}_i'\mathbf{x} = b_i$, $i = 1, \ldots, L$. Assume that the problem has a feasible solution. Show that there exists an optimal solution which is a convex combination of $L + 1$ extreme points of P. *Hint:* Use the resolution theorem to represent P.

Exercise 4.43

(a) Consider the minimization of $c_1x_1 + c_2x_2$ subject to the constraints

$$x_2 - 3 \leq x_1 \leq 2x_2 + 2, \qquad x_1, x_2 \geq 0.$$

Find necessary and sufficient conditions on (c_1, c_2) for the optimal cost to be finite.

(b) For a general feasible linear programming problem, consider the set of all cost vectors for which the optimal cost is finite. Is it a polyhedron? Prove your answer.

Exercise 4.44

(a) Let $P = \{(x_1, x_2) \mid x_1 - x_2 = 0,\ x_1 + x_2 = 0\}$. What are the extreme points and the extreme rays of P?

(b) Let $P = \{(x_1, x_2) \mid 4x_1 + 2x_2 \geq 8,\ 2x_1 + x_2 \leq 8\}$. What are the extreme points and the extreme rays of P?

(c) For the polyhedron of part (b), is it possible to express each one of its elements as a convex combination of its extreme points plus a nonnegative linear combination of its extreme rays? Is this compatible with the resolution theorem?

Exercise 4.45 Let P be a polyhedron with at least one extreme point. Is it possible to express an arbitrary element of P as a convex combination of its extreme points plus a nonnegative multiple of a single extreme ray?

Exercise 4.46 (Resolution theorem for polyhedral cones) Let C be a nonempty polyhedral cone.

(a) Show that C can be expressed as the union of a finite number $C_1, \ldots, C_k$ of pointed polyhedral cones. *Hint:* Intersect with orthants.

(b) Show that an extreme ray of C must be an extreme ray of one of the cones $C_1, \ldots, C_k$.

(c) Show that there exists a finite number of elements $\mathbf{w}^1, \ldots, \mathbf{w}^r$ of C such that

$$C = \left\{ \sum_{i=1}^{r} \theta_i \mathbf{w}^i \;\middle|\; \theta_1, \ldots, \theta_r \geq 0 \right\}.$$

Exercise 4.47 (Resolution theorem for general polyhedra) Let P be a polyhedron. Show that there exist vectors $\mathbf{x}^1, \ldots, \mathbf{x}^k$ and $\mathbf{w}^1, \ldots, \mathbf{w}^r$ such that

$$P = \left\{ \sum_{i=1}^{k} \lambda_i \mathbf{x}^i + \sum_{j=1}^{r} \theta_j \mathbf{w}^j \;\middle|\; \lambda_i \geq 0,\ \theta_j \geq 0,\ \sum_{i=1}^{k} \lambda_i = 1 \right\}.$$

Hint: Generalize the steps in the preceding exercise.

Exercise 4.48 * (Polar, finitely generated, and polyhedral cones) For any cone C, we define its *polar* $C^\perp$ by

$$C^\perp = \{\mathbf{p} \mid \mathbf{p}'\mathbf{x} \leq 0, \text{ for all } \mathbf{x} \in C\}.$$

(a) Let F be a finitely generated cone, of the form

$$F = \left\{ \sum_{i=1}^{r} \theta_i \mathbf{w}^i \;\middle|\; \theta_1, \ldots, \theta_r \geq 0 \right\}.$$

Show that $F^\perp = \{\mathbf{p} \mid \mathbf{p}'\mathbf{w}^i \leq 0,\ i = 1, \ldots, r\}$, which is a polyhedral cone.

(b) Show that the polar of $F^\perp$ is F and conclude that the polar of a polyhedral cone is finitely generated. *Hint:* Use Farkas' lemma.

(c) Show that a finitely generated pointed cone F is a polyhedron. *Hint:* Consider the polar of the polar.

(d) **(Polar cone theorem)** Let C be a closed, nonempty, and convex cone. Show that $(C^\perp)^\perp = C$. *Hint:* Mimic the derivation of Farkas' lemma using the separating hyperplane theorem (Section 4.7).

(e) Is the polar cone theorem true when C is the empty set?

Exercise 4.49 Consider a polyhedron, and let $\mathbf{x}$, $\mathbf{y}$ be two basic feasible solutions. If we are only allowed to make moves from any basic feasible solution to an adjacent one, show that we can go from $\mathbf{x}$ to $\mathbf{y}$ in a finite number of steps. *Hint:* Generalize the simplex method to nonstandard form problems: starting from a nonoptimal basic feasible solution, move along an extreme ray of the cone of feasible directions.

Exercise 4.50 We are interested in the problem of deciding whether a polyhedron

$$Q = \{\mathbf{x} \in \Re^n \mid \mathbf{A}\mathbf{x} \leq \mathbf{b},\ \mathbf{D}\mathbf{x} \geq \mathbf{d},\ \mathbf{x} \geq \mathbf{0}\}$$

is nonempty. We assume that the polyhedron $P = \{\mathbf{x} \in \Re^n \mid \mathbf{A}\mathbf{x} \leq \mathbf{b},\ \mathbf{x} \geq \mathbf{0}\}$ is nonempty and bounded. For any vector $\mathbf{p}$, of the same dimension as $\mathbf{d}$, we define

$$g(\mathbf{p}) = -\mathbf{p}'\mathbf{d} + \max_{\mathbf{x} \in P} \mathbf{p}'\mathbf{D}\mathbf{x}.$$

(a) Show that if Q is nonempty, then $g(\mathbf{p}) \geq 0$ for all $\mathbf{p} \geq \mathbf{0}$.

(b) Show that if Q is empty, then there exists some $\mathbf{p} \geq \mathbf{0}$, such that $g(\mathbf{p}) < 0$.

(c) If Q is empty, what is the minimum of $g(\mathbf{p})$ over all $\mathbf{p} \geq \mathbf{0}$?

4.13 Notes and sources

4.3. The duality theorem is due to von Neumann (1947), and Gale, Kuhn, and Tucker (1951).

4.6. Farkas' lemma is due to Farkas (1894) and Minkowski (1896). See Schrijver (1986) for a comprehensive presentation of related results. The connection between duality theory and arbitrage was developed by Ross (1976, 1978).

4.7. Weierstrass' Theorem and its proof can be found in most texts on real analysis; see, for example, Rudin (1976). While the simplex method is only relevant to linear programming problems with a finite number of variables, the approach based on the separating hyperplane theorem leads to a generalization of duality theory that covers more general convex optimization problems, as well as infinite-dimensional linear programming problems, that is, linear programming problems with infinitely many variables and constraints; see, e.g., Luenberger (1969) and Rockafellar (1970).

4.9. The resolution theorem and its converse are usually attributed to Farkas, Minkowski, and Weyl.