CHAPTER 1

Introduction

This book is mostly about a subject called Linear Programming. Before defining what we mean, in general, by a linear programming problem, let us describe a few practical real-world problems that serve to motivate and at least vaguely to define this subject.

1. Managing a Production Facility

Consider a production facility for a manufacturing company. The facility is capable of producing a variety of products that, for simplicity, we enumerate as $1, 2, \ldots, n$. These products are constructed/manufactured/produced out of certain raw materials. Suppose that there are m different raw materials, which again we simply enumerate as $1, 2, \ldots, m$. The decisions involved in managing/operating this facility are complicated and arise dynamically as market conditions evolve around it. However, to describe a simple, fairly realistic optimization problem, we consider a particular snapshot of the dynamic evolution. At this specific point in time, the facility has, for each raw material $i = 1, 2, \ldots, m$, a known amount, say b_i, on hand. Furthermore, each raw material has at this moment in time a known unit market value. We denote the unit value of the ith raw material by ρ_i.

In addition, each product is made from known amounts of the various raw materials. That is, producing one unit of product j requires a certain known amount, say a_{ij} units, of raw material i. Also, the jth final product can be sold at the known prevailing market price of σ_j dollars per unit.

Throughout this section we make an important assumption:

> The production facility is small relative to the market as a whole and therefore cannot through its actions alter the prevailing market value of its raw materials, nor can it affect the prevailing market price for its products.

We consider two optimization problems related to the efficient operation of this facility (later, in Chapter 5, we will see that these two problems are in fact closely related to each other).

1.1. Production Manager as Optimist. The first problem we wish to consider is the one faced by the company's production manager. It is the problem of how to use the raw materials on hand. Let us assume that she decides to produce x_j units of the jth product, $j = 1, 2, \ldots, n$. The revenue associated with the production of one unit of product j is σ_j. But there is also a cost of raw materials that must be

R.J. Vanderbei, *Linear Programming*, International Series in Operations Research & Management Science 196, DOI 10.1007/978-1-4614-7630-6_1,

considered. The cost of producing one unit of product j is $\sum_{i=1}^{m} \rho_i a_{ij}$. Therefore, the net revenue associated with the production of one unit is the difference between the revenue and the cost. Since the net revenue plays an important role in our model, we introduce notation for it by setting

$$c_j = \sigma_j - \sum_{i=1}^{m} \rho_i a_{ij}, \qquad j = 1, 2, \ldots, n.$$

Now, the net revenue corresponding to the production of x_j units of product j is simply $c_j x_j$, and the total net revenue is

$$\sum_{j=1}^{n} c_j x_j. \tag{1.1}$$

The production planner's goal is to maximize this quantity. However, there are constraints on the production levels that she can assign. For example, each production quantity x_j must be nonnegative, and so she has the constraints

$$x_j \geq 0, \qquad j = 1, 2, \ldots, n. \tag{1.2}$$

Secondly, she can't produce more product than she has raw material for. The amount of raw material i consumed by a given production schedule is $\sum_{j=1}^{n} a_{ij} x_j$, and so she must adhere to the following constraints:

$$\sum_{j=1}^{n} a_{ij} x_j \leq b_i \qquad i = 1, 2, \ldots, m. \tag{1.3}$$

To summarize, the production manager's job is to determine production values x_j, $j = 1, 2, \ldots, n$, so as to maximize (1.1) subject to the constraints given by (1.2) and (1.3). This optimization problem is an example of a linear programming problem. This particular example is often called the *resource allocation problem*.

1.2. Comptroller as Pessimist. In another office at the production facility sits an executive called the comptroller. The comptroller's problem (among others) is to assign a value to the raw materials on hand. These values are needed for accounting and planning purposes to determine the *cost of inventory*. There are rules about how these values can be set. The most important such rule (and the only one relevant to our discussion) is the following:

> The company must be willing to sell the raw materials should an outside firm offer to buy them at a price consistent with these values.

Let w_i denote the assigned unit value of the ith raw material, $i = 1, 2, \ldots, m$. That is, these are the numbers that the comptroller must determine. The *lost opportunity cost* of having b_i units of raw material i on hand is $b_i w_i$, and so the total lost opportunity cost is

$$\sum_{i=1}^{m} b_i w_i. \tag{1.4}$$

The comptroller's goal is to minimize this lost opportunity cost (to make the financial statements look as good as possible). But again, there are constraints. First of all, each assigned unit value w_i must be no less than the prevailing unit market value ρ_i, since if it were less an outsider would buy the company's raw material at a price lower than ρ_i, contradicting the assumption that ρ_i is the prevailing market price. That is,

$$w_i \geq \rho_i, \qquad i = 1, 2, \ldots, m. \tag{1.5}$$

Similarly,

$$\sum_{i=1}^{m} w_i a_{ij} \geq \sigma_j, \qquad j = 1, 2, \ldots, n. \tag{1.6}$$

To see why, suppose that the opposite inequality holds for some specific product j. Then an outsider could buy raw materials from the company, produce product j, and sell it at a lower price than the prevailing market price. This contradicts the assumption that σ_j is the prevailing market price, which cannot be lowered by the actions of the company we are studying. Minimizing (1.4) subject to the constraints given by (1.5) and (1.6) is a linear programming problem. It takes on a slightly simpler form if we make a change of variables by letting

$$y_i = w_i - \rho_i, \qquad i = 1, 2, \ldots, m.$$

In words, y_i is the increase in the unit value of raw material i representing the "mark-up" the company would charge should it wish simply to act as a reseller and sell raw materials back to the market. In terms of these variables, the comptroller's problem is to minimize

$$\sum_{i=1}^{m} b_i y_i$$

subject to

$$\sum_{i=1}^{m} y_i a_{ij} \geq c_j, \qquad j = 1, 2, \ldots, n$$

and

$$y_i \geq 0, \qquad i = 1, 2, \ldots, m.$$

Note that we've dropped a term, $\sum_{i=1}^{m} b_i \rho_i$, from the objective. It is a constant (the market value of the raw materials), and so, while it affects the value of the function being minimized, it does not have any impact on the actual optimal values of the variables (whose determination is the comptroller's main interest).

2. The Linear Programming Problem

In the two examples given above, there have been variables whose values are to be decided in some optimal fashion. These variables are referred to as *decision variables*. They are usually denoted as

$$x_j, \qquad j = 1, 2, \ldots, n.$$

In linear programming, the objective is always to maximize or to minimize some linear function of these decision variables

$$\zeta = c_1x_1 + c_2x_2 + \cdots + c_nx_n.$$

This function is called the *objective function*. It often seems that real-world problems are most naturally formulated as minimizations (since real-world planners always seem to be pessimists), but when discussing mathematics it is usually nicer to work with maximization problems. Of course, converting from one to the other is trivial both from the modeler's viewpoint (either minimize cost or maximize profit) and from the analyst's viewpoint (either maximize ζ or minimize $-\zeta$). Since this book is primarily about the mathematics of linear programming, we usually take the optimist's view of maximizing the objective function.

In addition to the objective function, the examples also had constraints. Some of these constraints were really simple, such as the requirement that some decision variable be nonnegative. Others were more involved. But in all cases the constraints consisted of either an equality or an inequality associated with some linear combination of the decision variables:

$$a_1x_1 + a_2x_2 + \cdots + a_nx_n \left\{ \begin{array}{c} \leq \\ = \\ \geq \end{array} \right\} b.$$

It is easy to convert constraints from one form to another. For example, an inequality constraint

$$a_1x_1 + a_2x_2 + \cdots + a_nx_n \leq b$$

can be converted to an equality constraint by adding a nonnegative variable, w, which we call a *slack variable*:

$$a_1x_1 + a_2x_2 + \cdots + a_nx_n + w = b, \qquad w \geq 0.$$

On the other hand, an equality constraint

$$a_1x_1 + a_2x_2 + \cdots + a_nx_n = b$$

can be converted to inequality form by introducing two inequality constraints:

$$\begin{aligned} a_1x_1 + a_2x_2 + \cdots + a_nx_n &\leq b \\ a_1x_1 + a_2x_2 + \cdots + a_nx_n &\geq b. \end{aligned}$$

Hence, in some sense, there is no a priori preference for how one poses the constraints (as long as they are linear, of course). However, we shall also see that, from a mathematical point of view, there is a preferred presentation. It is to pose the inequalities as less-thans and to stipulate that all the decision variables be nonnegative. Hence, the linear programming problem, as we study it, can be formulated as follows:

$$\begin{array}{lrcrcrcrcl} \text{maximize} & c_1x_1 & + & c_2x_2 & + \cdots + & c_nx_n & & \\ \text{subject to} & a_{11}x_1 & + & a_{12}x_2 & + \cdots + & a_{1n}x_n & \leq & b_1 \\ & a_{21}x_1 & + & a_{22}x_2 & + \cdots + & a_{2n}x_n & \leq & b_2 \\ & & & & \vdots & & & \\ & a_{m1}x_1 & + & a_{m2}x_2 & + \cdots + & a_{mn}x_n & \leq & b_m \\ & & & & x_1, \; x_2, \; \ldots & x_n & \geq & 0. \end{array}$$

We refer to linear programs formulated this way as linear programs in *standard form*. In our aim for consistency, we shall always use m to denote the number of constraints, and n to denote the number of decision variables.

A proposal of specific values for the decision variables is called a *solution*. A solution $(x_1, x_2, \ldots, x_n)$ is called *feasible* if it satisfies all of the constraints. It is called *optimal* if in addition it attains the desired maximum. Some problems are just simply infeasible, as the following example illustrates:

$$
\begin{array}{lrcr}
\text{maximize} & 5x_1 + 4x_2 & & \\
\text{subject to} & x_1 + x_2 & \leq & 2 \\
& -2x_1 - 2x_2 & \leq & -9 \\
& x_1,\ x_2 & \geq & 0.
\end{array}
$$

Indeed, the second constraint implies that $x_1 + x_2 \geq 4.5$, which contradicts the first constraint. If a problem has no feasible solution, then the problem itself is called *infeasible*.

At the other extreme from infeasible problems, one finds unbounded problems. A problem is *unbounded* if it has feasible solutions with arbitrarily large objective values. For example, consider

$$
\begin{array}{lrcr}
\text{maximize} & x_1 - 4x_2 & & \\
\text{subject to} & -2x_1 + x_2 & \leq & -1 \\
& -x_1 - 2x_2 & \leq & -2 \\
& x_1,\ x_2 & \geq & 0.
\end{array}
$$

Here, we could set x_2 to zero and let x_1 be arbitrarily large. As long as x_1 is greater than 2 the solution will be feasible, and as it gets large the objective function does too. Hence, the problem is unbounded. In addition to finding optimal solutions to linear programming problems, we shall also be interested in detecting when a problem is infeasible or unbounded.

Exercises

1.1 A steel company must decide how to allocate next week's time on a rolling mill, which is a machine that takes unfinished slabs of steel as input and can produce either of two semi-finished products: bands and coils. The mill's two products come off the rolling line at different rates:

Bands	200 tons/h
Coils	140 tons/h.

They also produce different profits:

Bands	\$25/ton
Coils	\$30/ton.

Based on currently booked orders, the following upper bounds are placed on the amount of each product to produce:

Bands	6,000 tons
Coils	4,000 tons.

Given that there are 40 h of production time available this week, the problem is to decide how many tons of bands and how many tons of coils should be produced to yield the greatest profit. Formulate this problem as a linear programming problem. Can you solve this problem by inspection?

1.2 A small airline, Ivy Air, flies between three cities: Ithaca, Newark, and Boston. They offer several flights but, for this problem, let us focus on the Friday afternoon flight that departs from Ithaca, stops in Newark, and continues to Boston. There are three types of passengers:

(a) Those traveling from Ithaca to Newark.
(b) Those traveling from Newark to Boston.
(c) Those traveling from Ithaca to Boston.

The aircraft is a small commuter plane that seats 30 passengers. The airline offers three fare classes:

(a) Y class: full coach.
(b) B class: nonrefundable.
(c) M class: nonrefundable, 3-week advanced purchase.

Ticket prices, which are largely determined by external influences (i.e., competitors), have been set and advertised as follows:

	Ithaca–Newark	Newark–Boston	Ithaca–Boston
Y	300	160	360
B	220	130	280
M	100	80	140

Based on past experience, demand forecasters at Ivy Air have determined the following upper bounds on the number of potential customers in each of the nine possible origin-destination/fare-class combinations:

	Ithaca–Newark	Newark–Boston	Ithaca–Boston
Y	4	8	3
B	8	13	10
M	22	20	18

The goal is to decide how many tickets from each of the nine origin/destination/fare-class combinations to sell. The constraints are that the plane cannot be overbooked on either of the two legs of the flight and that the number of tickets made available cannot exceed the forecasted maximum demand. The objective is to maximize the revenue. Formulate this problem as a linear programming problem.

1.3 Suppose that Y is a random variable taking on one of n known values:

$$a_1, a_2, \dots, a_n.$$

Suppose we know that Y either has distribution p given by

$$\mathbb{P}(Y = a_j) = p_j$$

or it has distribution q given by

$$\mathbb{P}(Y = a_j) = q_j.$$

Of course, the numbers p_j, $j = 1, 2, \ldots, n$ are nonnegative and sum to one. The same is true for the q_j's. Based on a single observation of Y, we wish to guess whether it has distribution p or distribution q. That is, for each possible outcome a_j, we will assert with probability x_j that the distribution is p and with probability $1-x_j$ that the distribution is q. We wish to determine the probabilities x_j, $j = 1, 2, \ldots, n$, such that the probability of saying the distribution is p when in fact it is q has probability no larger than β, where β is some small positive value (such as 0.05). Furthermore, given this constraint, we wish to maximize the probability that we say the distribution is p when in fact it is p. Formulate this maximization problem as a linear programming problem.

Notes

The subject of linear programming has its roots in the study of linear inequalities, which can be traced as far back as 1826 to the work of Fourier. Since then, many mathematicians have proved special cases of the most important result in the subject—the *duality theorem*. The applied side of the subject got its start in 1939 when L.V. Kantorovich noted the practical importance of a certain class of linear programming problems and gave an algorithm for their solution—see Kantorovich (1960). Unfortunately, for several years, Kantorovich's work was unknown in the West and unnoticed in the East. The subject really took off in 1947 when G.B. Dantzig invented the *simplex method* for solving the linear programming problems that arose in U.S. Air Force planning problems. The earliest published accounts of Dantzig's work appeared in 1951 (Dantzig 1951a,b). His monograph (Dantzig 1963) remains an important reference. In the same year that Dantzig invented the simplex method, T.C. Koopmans showed that linear programming provided the appropriate model for the analysis of classical economic theories. In 1975, the Royal Swedish Academy of Sciences awarded the Nobel Prize in economic science to L.V. Kantorovich and T.C. Koopmans "for their contributions to the theory of optimum allocation of resources." Apparently the academy regarded Dantzig's work as too mathematical for the prize in economics (and there is no Nobel Prize in mathematics).

The textbooks by Bradley et al. (1977), Bazaraa et al. (1977), and Hillier and Lieberman (1977) are known for their extensive collections of interesting practical applications.