APPLIED SYSTEMS ANALYSIS

Engineering Planning and Technology Management

Richard de Neufville
Massachusetts Institute of Technology

McGraw-Hill Publishing Company
New York St. Louis San Francisco Auckland Bogotá Caracas Hamburg
Lisbon London Madrid Mexico Milan Montreal New Delhi Oklahoma City
Paris San Juan São Paulo Singapore Sydney Tokyo Toronto

5.16. *Beans and Corn*
A farmer is planning to produce beans and corn. The unit profit on a pound of beans is \$0.50 for the first 1000 lbs. After this the unit profit is \$0.35. Corn profits are divided into three categories. The unit profit for the first 800 lbs is \$0.75. The next 400 lbs yield \$0.40/lb. Any additional lbs of corn will yield \$0.25.

The farmer has a limited storage capacity and thus can produce no more than 5000 lbs of produce, in any combination of beans and corn.

The amount of land available is also a restriction on the output. The farmer has 10 acres of workable land for any combination of the two crops. He has determined that every 2000 lbs of beans requires 4 acres. Every 2000 lbs of corn requires 5 acres of land.

The farmer must use a cooperative machine to harvest these crops. He will have access to this machine for 20 hr. He can use it to harvest 200 lbs of beans/hr or 250 lbs of corn/hr.

Formulate an LP to determine the optimal quantity of each crop the farmer should plant, being sure to define your variables.

5.17. *Chemical Company*
A company makes two chemicals, 1 and 2. These chemicals can be manufactured by three different processes using two different raw materials and a fuel. Production data are given below. Formulate an LP model to estimate the time required to run each process in order to maximize the total amount of chemicals manufactured.

	Requirements per unit time			Output per unit time	
Process	**Raw material 1**	**Raw material 2**	**Fuel**	**Chemical 1**	**Chemical 2**
1	9	5	50	9	6
2	6	8	75	7	10
3	4	11	100	10	6
Amount Available	200	400	1850		

5.18. *Irrigation Problem*
A fertile but arid area in Southern California is to be cleared, leveled, fertilized, irrigated, and turned into prime farmland. Dealing specifically with the irrigation problem, a linear program is written to select, from a large number of possible canal and pipeline routes, the cheapest set of irrigation links that will provide adequate water to the entire area.

Given the characteristics of LP, explain how each of the following four considerations can be included in the LP, or why it cannot be included.

(*a*) It is futile to construct feeder line #2 with a greater capacity than link #1, which connects link #2 to the main supply line.
(*b*) The cost of a pipeline is a nonlinear function of its capacity.
(*c*) As bulldozer-hours/week increase, the cost of a bulldozer-hour increases.
(*d*) Water demand is a function of weather and therefore varies from year to year.

CHAPTER 6

SENSITIVITY ANALYSIS

6.1 CONCEPT

Sensitivity analysis is the process of investigating the dependence of an optimal solution to changes in the way a problem is formulated. Doing a sensitivity analysis is a key part of the design process, equal in importance to the optimization process itself.

The significance of sensitivity analysis stems from the fact that the mathematical problem we solve in any optimization is only an approximation of the real problem. The exact solution we obtain and use to represent reality is thus not an exact solution to the real problem of design. At best, the optimization process provides a good approximation to the best design of a real system.

None of our mathematical models will ever represent systems exactly. All these representations are approximations in some way. They each differ from reality in any or all of the following three ways:

- Structurally, because the overall nature of the equations does not correspond precisely to the actual situation.
- Parametrically, as we are not able to determine all coefficients precisely.
- Probabilistically, in that we typically assume that the situation is deterministic when it is generally variable.

Structural differences arise as a matter of course in the modeling process. The way we typically construct a mathematical model of a system is to imagine some form we believe is appropriate or useful, and then to match the real situation to this structure. This is the process by which most production functions are made: starting from a Cobb-Douglas or translog formula, we estimate coefficients (see Section 2.3). Similarly, when we wish to attack a problem with linear programming, we deliberately construct a set of linear approximations to the reality. These approximations are useful, informative, and efficient. The point is, however, that they are indeed approximations.

Parametric differences between a model and reality reflect our ultimate inability to measure every aspect of a system absolutely accurately. On the one hand we simply never have enough time or money to measure everything we might want to. On the other, we will frequently run into difficulties knowing how precisely to define a quantity so that we can measure it (see Chapter 14 for an extended presentation of that point).

Finally, most of our models assume that the real situation is deterministic, that it can be described by fixed quantities for any particular situation. Most real problems, however, are highly probabilistic; their performance or characteristics can be best described by a probability distribution of some sort. For example, the thermal efficiency of an electric power plant may be 40% on average but varies over a range depending on the variations in the fuels used, the level of maintenance, and the day to day performance of the operations. As Chapter 15 presents in detail, there is always uncertainty in our estimates and descriptions.

Because the real situation never corresponds exactly to the mathematical model, we cannot be satisfied when we have merely obtained the optimum answer for the mathematical problem. We must consider what the answer or design would be if the model were changed to represent what the system might actually be. We must do this both to understand how sensitive or vulnerable our design might be, and so that we can effectively respond when we do perceive changes. This is the concept behind sensitivity analysis.

This chapter presents sensitivity analysis principally in the context of linear programming. This is because the solutions to linear programming problems automatically include most of the sensitivity information a designer could wish for, and thus the linear programs constitute the predominant basis for sensitivity analysis. Additionally, the linearity of linear programming makes it easier to explain key concepts, which the student can then extend to other forms of optimization.

Most of this chapter is devoted to the two most important aspects of sensitivity analysis, the concept and use of

- shadow prices (Section 6.2)
- opportunity costs (Section 6.3)

Sections 6.4 and 6.5 illustrate their use and explain their relationship to each other. Additionally, the chapter describes

- break-even analysis (Section 6.6)

as a practical means of dealing with a special kind of problem.

6.2 SHADOW PRICES

A shadow price is the rate of change of the objective function with respect to a particular constraint. As indicated in Section 3.3, it is essentially equivalent to the Lagrangean multiplier.

> **Semantic caution:** The shadow price has no necessary connection with money, despite its name. Its units are those of the objective function divided by the constraint. For example, in producing electric power subject to environmental constraints, the shadow price might be in terms of kilowatts per gram of pollutant. The shadow price is expressed in dollars only when the objective function is also, say dollars of profit.

Use of shadow prices. Knowledge of the shadow prices is important because, by indicating the quantitative effects of changes in the constraints, they enable the designer to

- *identify* which constraints might most beneficially be changed, and to *initiate* these changes as a fundamental means to improve the design.
- *react* appropriately when external circumstances create opportunities or threats to change the constraints.

Designers working under constraints—such as a budget, an environmental standard, or the percent of a constituent in an alloy—should identify the shadow prices on these constraints to the extent possible. They then can determine which constraints have the greatest effect on their ability to optimize the design and, from this information, initiate proposals to change these standards to the extent desirable (see box on following page).

Knowledge of shadow prices also helps designers react to possible changes to constraints that may be generated externally. It often happens that customers or clients of a system propose that the specifications of the design be changed: Is this a good idea? Alternatively the legislative or administrative process may propose to change the rules: Will this degrade performance significantly? See box (page 109) for examples.

The use of shadow prices to examine the validity of the constraints is sometimes referred to as *specification sensitivity analysis*, because many of the constraints can be interpreted as specifications. This is a useful way to think of the process: rather than simply accept the rules as given, the designer should question these rules and get them changed when it seems best.

Sign of shadow prices. A key practical question with regard to shadow prices is: What is the sign of the shadow price? In which direction should one change a constraint to improve a design?

The right answer to this question is not obvious. The answer that immediately seems logical to many people is in fact wrong: A positive change in the constraint does not necessarily imply a positive shadow price. As it turns out,

Using Shadow Prices to Initiate Changes of Constraints

Consider again the problem of designing the water supply system for New York City. Suppose the problem had been formulated as

Minimize: Total cost of systems

Subject to (among others): Water pressure at curb $\geq$ 40 psi

Suppose further that we found that

Shadow price on pressure constraint = \$78 million/psi

This means that there is the potential for great improvements in design—savings in cost in this case—if we were to change the constraint. We could then ask:

- Is this constraint absolute, determined by some law of nature, or was it set for convenience at some time?
- If the latter, is the constraint still justified? Can it be changed?
- Is the shadow price sufficiently large to make it worthwhile to go to the trouble of getting the constraint changed?

For New York City in particular, the standard was set rather arbitrarily (at 40 as a round number rather than 38, say), years ago when buildings were different. There is thus a reasonable basis for changing the standard if the improvements in design justify it. In our analyses at the time we worked on the problem, it appeared that it would indeed have been worthwhile to alter the constraint at least a few psi, because of the huge savings in cost.

---

there is no evident connection between the change in constraint and the sign of the shadow price.

To understand the relationship between the sign of the shadow price and the constraints, it is useful to divide them into

- *upper bounds*, which place a maximum constraint on a feasible region, for example: $\sum_i a_{ji}X_i \leq b_j$.
- *lower bounds*, which place a minimum constraint on a feasible region, such as $\sum_i a_{ki}X_i \geq b_k$ (see Section 3.4).

A helpful mental image here is that the feasible region is the space inside the constraints. The upper and lower bounds then constitute the roof and the floor of this space.

Secondly, the concept of relaxation needs to be understood. *A constraint is relaxed* if it is changed so as to increase the size of the feasible region, that is,

Using Shadow Prices to React to Changes in Constraints

Suppose that, as the analyst for a refinery, you have determined that your shadow price on the percent of sulfur in the fuel oil is

Shadow Price = \$3.81 a barrel/% sulfur

If a prospective customer who does not care about the sulfur content offers to take a lower grade of fuel oil with 1% more sulfur if you will give a \$5 a barrel discount in price, do you accept? No: the change in the constraint would only save you \$3.81 and you would lose money (3.81 – 5.00) if you accepted the offer. You might respond by offering a \$3 a barrel discount, which would leave you ahead by \$0.81.

Alternatively, if the local authorities proposed to force you to take out more sulfur, say 2% more, you could respond by indicating this requirement would raise your price by \$7.62 (= 2 × 3.81) a barrel and—let us say—make it impossible to produce profitably. Knowledge of the shadow price thus gives you specific grounds for reacting to proposed changes in specifications.

---

if an upper bound is increased or a lower bound is decreased. The idea here is that the constraints place pressures on the designer and that moving them away, making the feasible region larger, relaxes these pressures.

The relationship between the nature of the shadow prices and the changes in constraints is that

- Relaxing the constraints leads to improvements in the optimum, either increasing a maximum or decreasing a minimum.

Visualizing the feasible region as a space, this rule can be stated colloquially as

- Changes in constraints that "raise the roof" or "lower the floor" will tend to improve the optimum.

The logic behind this rule is simple. When one relaxes a constraint, the feasible region expands—it is the previous region plus an increment. Thus, one can do at least as well after the relaxation as before—never worse. On the other hand, there can be better solutions in the expanded region. The optimum solution after the relaxation of a constraint must necessarily be equal or better.

Conversely, if a constraint is tightened, thus reducing the feasible region, this may exclude the previous optimum solution. Therefore, the optimum solution after a constraint is tightened is necessarily worse if not equal.

It is important to notice that there is absolutely no simple relationship between the sign of the change in constraint and the sign of the shadow price. This is because an increase in the constraint can either relax or tighten a constraint,

Lack of Evident Relationship Between the Signs of Changes in Constraints and the Objective

Consider the problem:

Maximize: X_2

Subject to:

$$X_1 - 2X_2 \geq 0$$
$$X_1 \leq 2$$
$$X_1, \quad X_2 \geq 0$$

The feasible region for this problem appears in Figure 6.1, interior to the shadowed lines, that is, the triangle between the origin and point A. The optimum, the highest feasible value of X_2, is obviously at $A(X_1 = 2, X_2 = 1)$.

Suppose one now relaxes the first constraint by 1. Since it is a lower bound, this means we decrease the constraint and have

$$X_1 - 2X_2 \geq -1$$

This expands the feasible region, allowing the optimum to raise to point B $(X_1 = 2, X_2 = 1.5)$. We have a positive change in the objective function associated with a negative change in the constraint.

Relaxing the second constraint by 1 means increasing this upper bound to

$$X_2 \leq 3$$

Again the feasible region is expanded, and the optimum is at $C(X_1 = 3, X_2 = 1.5)$. Here a positive change in the objective is associated with a positive change in constraint.

---

depending on whether it is an upper or lower bound and if one is maximizing or minimizing. The preceding box illustrates the phenomenon graphically.

Because there is no simple sign convention relating to shadow prices, they are often reported in absolute values. (See box on page 112.) This places the responsibility on the analyst to determine what kind of change improves the optimum. The minimum value of the shadow price is thus typically taken to be zero.

> **Semantic caution:** There may be a difference between the formal definition of an upper or lower bound, and what you may visualize. Consider Figure 6.1; the constraint $X_1 - 2X_2 \geq 0$ is a lower bound, but it effectively looks like a bound on the maximum or uppermost value of the objective function.

Range of shadow prices. The concept of the range of a shadow price applies only, in practice, to linear programming. In general, the shadow price is the instantaneous change in the objective function with respect to a specific constraint, $\partial Y / \partial b_j$, as defined in Section 3.2. This rate can vary with the decision variables, and normally will when the constraints are nonlinear, as the application in Section 3.4 illustrates.

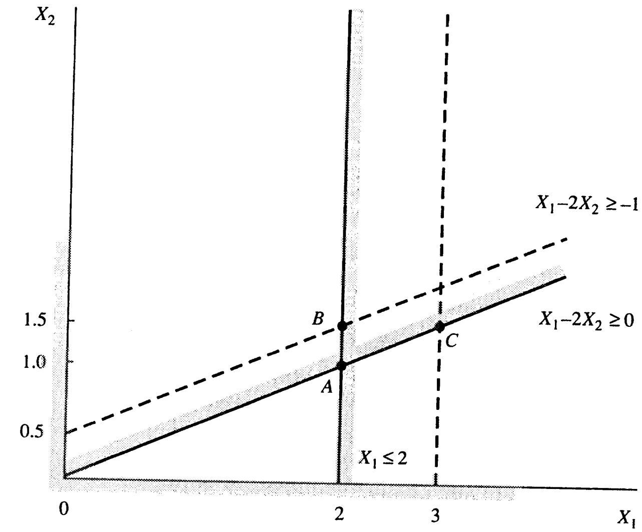

FIGURE 6.1
Illustration for demonstration of lack of evident relationship between signs of changes in constraints and objective.

The peculiarity of linear programming in this regard is that the shadow prices are constant over a range, rather than varying continuously. The output of a linear program correspondingly consists not only of the value of the shadow price but also defines the range of the change in the constraint over which the given value holds.

The existence of a range of constancy for the shadow price is, generally, an artifact of the way we have formulated the problem as a linear program. In reality, one can assume that the linear equations are approximations and that, if we really described the problem accurately with the appropriate nonlinear equations, the shadow prices could usually vary instantaneously. Even though the range of constancy of the shadow prices is thus an artificial result, the concept is most useful in practice.

The range of constancy of shadow prices is useful because it further indicates how sensitive the optimum solution is to the constraint. Indeed, if the range is narrow, we know that even small changes in the constraint could lead to quite different shadow prices, thus that the shadow prices may change rapidly. If, on the contrary, the range is large, we can infer that the shadow prices vary slowly.

The constancy of any shadow price over a range results from the nature of the optimum solution in linear programming. The optimum is at a corner point of the feasible region, defined by the intersection of linear equations of constraints; as one of these equations is changed by a constant amount, so is the intersection and so is the value of the objective function since it is also linear. This constancy goes on indefinitely until something fundamental changes.

Range of Shadow Prices

Consider the same problem as before, augmented by a third constraint as shown in Figure 6.2:

Maximize: X_2

Subject to:

$$\begin{aligned} X_1 - 2X_2 &\geq 0 \\ X_1 &\leq 2 \\ X_1 + X_2 &\leq 4 \\ X_1, \quad X_2 &\geq 0 \end{aligned}$$

As shown in the previous box, the shadow price on the first constraint is $\frac{1}{2}$; when the feasible region is expanded by lowering this lower bound by 1 (to $X_1 - 2X_2 \geq -1$), the optimum increases by $\frac{1}{2}$.

This shadow price will prevail until the first constraint has reached the intersection of constraints at D(at $X_1 - 2X_2 \geq -2$). Further relaxation of the first constraint then still leads to increases in the optimum, but these will be at a different rate. The optimum corner point will no longer be defined along the second constraint ($X_1 \leq 2$) but along the third ($X_1 + X_2 \leq 4$). The rate of change of the objective function will then be different.

In general, there are upper and lower limits to the range. In this case, the limit in the range in the other direction is point E(at $X_1 - 2X_2 \geq 2$). At that point the feasible region disappears. Further increases in the constraint are then meaningless since there is no solution, let alone an optimum.

The range on the shadow price on the first constraint is therefore from −2 to +2.

The range of the shadow price is defined by the intersections of constraints adjacent to the one that defines the optimum of the linear program. Ultimately, as a constraint varies, it intersects with other constraints. The equations defining the optimum corner then change, leading to a change in the shadow prices. This occurs at other intersections. (See box above for an example.)

In general there can be a limit to the range of a shadow price for both increases and decreases. It will often be, however, that the range in either direction is unbounded. This occurs typically with constraints that are not binding at the optimum (slack $\neq 0$, so shadow price $= 0$). The constraint then has a zero shadow price indefinitely in the direction moving ever further away from the optimum; in the other direction, the shadow price is zero only until the constraint is binding (slack $= 0$).

Change in the shadow prices. For linear programming, the shadow prices change in a distinctive manner beyond the range of their constancy. The rule is

- the shadow price reduces beyond its range if the constraint is being relaxed.
- the shadow price increases beyond its range if the constraint is being tightened.

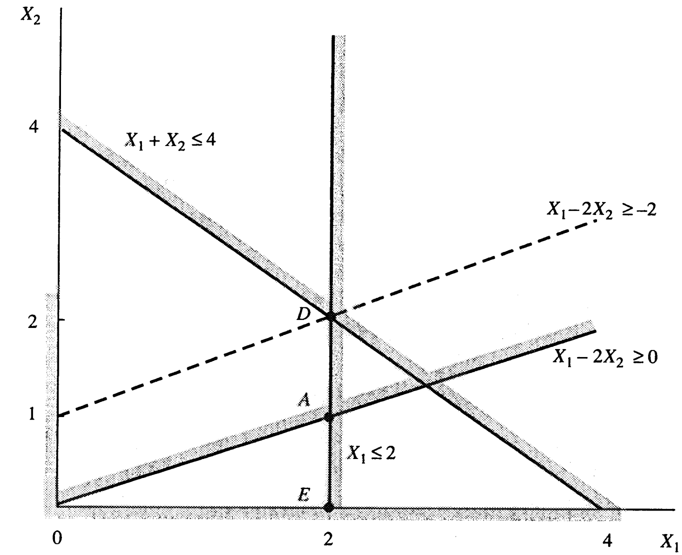

FIGURE 6.2
Illustration for demonstration of range of shadow prices in linear programming.

Because of the way shadow prices are typically reported, in absolute values, shadow prices can only reduce to zero. At that point slack $\neq 0$ and changes in the constraint have no effect. Shadow prices may, on the other hand, increase without limit.

The reason for this rule lies in the convexity of the feasible region and the linearity of the objective function. As we meet up with other constraints, which define a jog in the shape of the feasible region and thus a change in shadow price, the nature of the change in shadow price must be as defined (see box for an example).

Changes in Shadow Prices

Figure 6.3 shows the optimization of Figure 6.2 with the constraint at the end of its range for the shadow price and passing through D. The shadow price from the optimum (at A) to D was equal to $\frac{1}{2}$. What is it beyond the range?

If the constraint is further relaxed by one to $X_1 - 2X_2 \geq -3$, the new optimum will be at F, the intersection of this constraint with $X_1 + X_2 \leq 4$. The coordinates of F are $(\frac{5}{3}, \frac{7}{3})$. The change in X_2, the objective function, associated with the unit change in the constraint and the move from D to F is thus $\frac{1}{3}$. As expected, the shadow price steadily reduces, beyond its range of constancy, in the direction of relaxation.

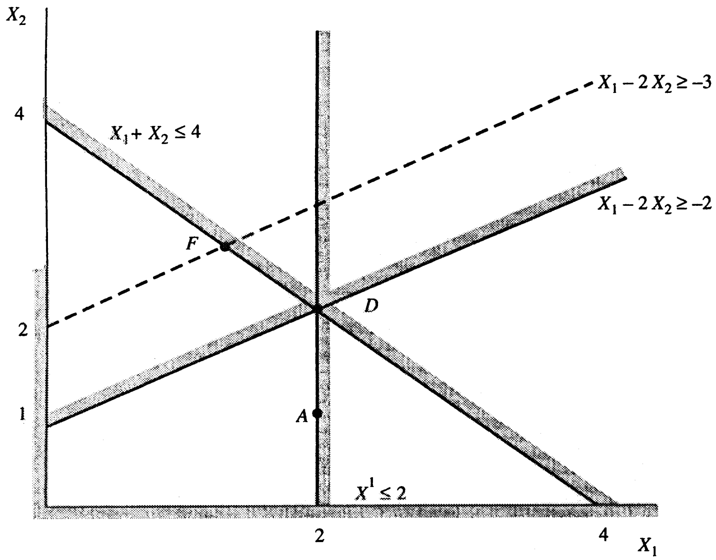

FIGURE 6.3
Illustration for demonstration of nature of change of shadow prices beyond range.

6.3 OPPORTUNITY COSTS

Opportunity costs, in the context of sensitivity analysis, relate to the coefficients of the decision variables in the objective function. In general terms they define the "cost"—in terms of degradation of the optimum, whether the objective function refers to cost, product of some other quantity—of using decision variables that are not part of the optimal design. They have a specialized meaning, which must be carefully understood. This meaning must also be distinguished from a more general one commonly used in economics.

> **Semantic caution:** In economics and management one often refers to "opportunity cost" in the sense of meaning the highest value one could obtain for some resource. Thus the term "opportunity cost of capital" refers to the highest rate of return one could obtain per extra unit available (see Section 12.3). For linear programming and sensitivity analysis, however, where one would be formally maximizing return or profit subject to constraints on the budget, this notion is embodied in the concept of the shadow price of the constraint on available capital. Thus the other usage is particularly confusing because it refers to a totally different sensitivity than the one accepted in optimization.

Definition of opportunity costs. To define opportunity costs formally, we first need to examine the nature of the optimum solution to an optimization. In general, the set of optimal decision variables, $\mathbf{X}^*$, can be divided into two categories. One

category consists of all the decision variables that are not used in the optimum design, that is, whose optimal value is zero, $X_i^* = 0$. Generally this category is quite large in any practical problem as there are indeed many possible decisions that we do not take. For example, if we intend to supply New York with fuel oil, we may well use plenty from Venezuela, the Texas Gulf, and elsewhere, but none from Alaska because of poor connections and other opportunities (see box on page 117). The two categories for the optimum set of the decision variables, $\mathbf{X}^*$, are thus the

- *optimal variables*, those with nonzero values at the optimum ($X_i^* \neq 0$). These are said to be "in the solution."
- *nonoptimal variables*, those equal to zero at the optimum ($X_i^* = 0$). These are said to be "not in the solution."

With this distinction in mind, we can now formally define opportunity costs in sensitivity analysis: The *opportunity cost* is the rate of degradation of the optimum per unit use of a nonoptimal variable in the design. The notion of degradation here is important: it refers to the worsening of an optimum. This may either be a decrease—if we are trying to maximize, or an increase—if we are trying to minimize.

It is important to recognize that the use of a nonoptimal variable in the design must make matters worse. If we should ever use a unit of a nonoptimal variable, we would necessarily—due to the constraints—displace some quantities of the optimal variables, thus leading to a worsening. Using again the hypothetical example of the shipment of fuel oil, if we insist on supplying New York from Alaska, we will need less from Venezuela or some other closer place and thus increase costs. Conversely, in the example shown in Figure 6.4, if we used some Activity 3 instead of some combination of Activities 1 and 2, we would reduce production.

Use of opportunity costs. The way the designer should use opportunity costs differs considerably from that of shadow prices. This difference results from the quite distinct functions of the constraints (**B**) and the coefficients of the objective function (**C**) in the optimization process.

The coefficients **C** represent some kind of technical reality of the design which the designer cannot, in general, influence directly. For example, if we are maximizing production, the objective represents the production function and the coefficients **C** each represent the best productivity of each activity or decision variable. Alternatively, if we are minimizing cost, the coefficients **C** are the prices of the items that could be purchased, and these should already be the best available. The constraints, on the other hand, do not represent any best cost or productivity, they are often fairly arbitrary and may be changed.

The result is that the analyst is unlikely to be able to or want to manipulate the coefficients **C**. Thus in general the variation of the optimum with every C_i is not of interest. The only useful variation is in fact that of the optimum with the coefficients of the nonoptimal decision variables, that is, the opportunity costs.

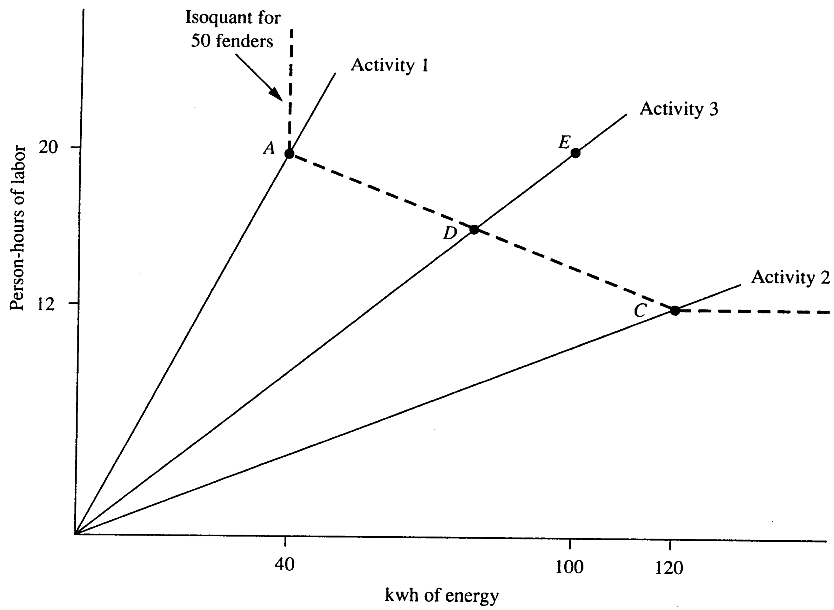

FIGURE 6.4
Illustration for discussion of opportunity cost of Activity 3.

The interest of the opportunity costs lies in the special information we can derive from them in a linear program. The opportunity cost associated with any nonoptimal variable in effect defines the value of the coefficient for that variable when its use would be optimal. This we call the *trigger price* for that variable (using price by analogy to opportunity cost, although the optimum may or may not be expressed in terms of money).

The relationship between an opportunity cost and the trigger price for a variable is direct. Suppose X_k^* is a nonoptimal decision variable for a problem. This means that its contribution to the objective function per unit of X_k, that is C_k, is unattractive. Either C_k is too large if we are minimizing, or it is too small if we are maximizing. In either case, using any unit of X_k in the design worsens the optimum by the opportunity cost, OC_k. It follows that if the coefficient C_k somehow improved (became smaller for a minimization, larger for a maximization), the OC_k would become smaller. As relationships are linear in a linear program, it then follows that if C_k were improved by OC_k, there would no longer be any opportunity price for X_k and it could be used in the optimal design. Therefore:

Trigger price	$= C_k - OC_k$	minimization
	$= C_k + OC_k$	maximization

Nonoptimal Decision Variables

Consider the production process described in Section 5.6, Table 5.1, which could use either of two activities. We now add an Activity 3, which requires 100 kwh and 20 person-hours to produce 50 fenders. This is shown on Figure 6.4, which is simply an extension of Figure 5.5.

If one were trying to maximize production subject only to constraints on the resources, one would never use Activity 3. To produce 50 units it requires the resources shown by point E. In fact, that proportion of resources can be used much more efficiently by a combination of Activities 1 and 2, which combine to produce 50 units at point D or 20% more (62.5 units) at E (see previous discussion). Activity 3 is thus a nonoptimal variable in this situation.

Further, if C_k were improved by more than OC_k we would definitely want to use X_k in the optimum design (see box below).

The use of the opportunity costs is thus to define the coefficients of the decision variables which would lead to a change in design. The designer, having defined the optimum design, then continuously monitors the situation to determine when it has changed enough so that a new design ought to be used.

In practice, opportunity costs and trigger prices are most useful in problems of minimizing costs. This is because the prices for goods fluctuate fairly rapidly, as can be verified by following the quotes for commodities in the financial newspapers. Managers of large systems and of plants thus can continually monitor the fluctuations in prices and, by comparing the quotes to their trigger prices, identify possibilities for changing design to maintain optimality for the new circumstances. Alternatively they can define for their purchasing agents and suppliers the price at which they can afford to buy specific goods.

Opportunity Costs

Same problem as above. Activity 3 is nonoptimal. One unit of Activity 3, at E in Figure 6.4, produces 50, whereas the same resources, used in a combination of Activities 1 and 2, can produce 62.5.

These facts can be interpreted as follows:

Productivity per unit X_3: $C_3 = 50$

Opportunity cost of X_3: $OC_3 = 62.5 - 50 = 12.5$

Trigger "Price" (here in units of product) for $X_3 = 50 + OC_3 = 62.5$

6.4 ILLUSTRATION

A linear programming procedure normally provides all the sensitivity information previously discussed. The complete output you should expect to get then consists of two parts. First, the optimal answer:

- Optimum objective function, OF^*
- Optimal set of decision variables, $\mathbf{X}^*$
- Slack variables at optimum, $\mathbf{SV}^*$

Secondly, the sensitivity information:

- The set of shadow prices, $\mathbf{SP}^*$
- The ranges of the shadow prices
- The set of opportunity costs, $\mathbf{OC}^*$

These are now all illustrated by an example.

The example is deliberately simple so that the relationships between the different elements of the solution are easy to understand. The problem is sufficiently easy mathematically that it can be solved by inspection, thus avoiding any questions concerning the algorithms for solving linear programming problems. Yet it is complex enough to demonstrate all the interesting conceptual issues.

The problem is:

$$\begin{array}{llrcrcrl} \text{Minimize:} & \text{Cost} = & 4X_1 & + & 12X_2 & + & 10X_3 & \\ \text{Subject to:} & & 2X_1 & + & X_2 & & & \geq 2 = b_1 \\ & & X_1 & + & 2X_2 & + & X_3 & \geq 5 = b_2 \\ & & & & & & X_3 & \geq 1 = b_3 \\ & & X_1, & & X_2, & & X_3 & \geq 0 \end{array}$$

Physically, this might represent the cost minimization of a product subject to technical constraints. Thus it could be the manufacture of reinforced concrete using cement mix (X_1), special additives (X_2), and reinforcing steel rods (X_3). The constraints would then be interpreted as follows:

- A minimum amount of cement mix and additives is needed to bind everything together.
- The separate contributions of each input add up to at least the desired total.
- A minimum amount of steel is needed for strength.
- The inputs cannot be negative.

The optimal solution to this problem can be determined by inspection. We need to have at least $X_3 = 1$ to meet the third constraint, and then we want to use as much X_1 as possible to meet the other constraints, because it is cheapest to use. (In the first constraint, the effect of using one unit of X_1, which costs 4, is equivalent to using two units of X_2, which cost 24. Similarly, in the second constraint, X_1 is cheapest in meeting a unit's worth of constraint.) The best design is

$$\mathbf{X}^* = (4, 0, 1)$$

for a minimum objective function

$$OF^* = 26$$

The slack variables, calculated by substituting the optimal decision variables in the constraints, are

$$\mathbf{SV}^* = (6, 0, 0)$$

Shadow prices. The shadow price on the first constraint is obvious: since there is slack (=6), the shadow price on b_1 is zero. This follows both intuitively, from the observation that since we meet the constraint by a wide margin the optimum design is not affected by a unit change in b_1, and formally by the complementary slackness rule ($S\lambda = 0$).

The shadow price on the second constraint is not zero. Its value can be seen fairly directly: a unit increase in b_2 requires a unit change in the decision variables, the cheapest of which is X_1. Thus the shadow price on b_2 is 4.

The shadow price on the third constraint is also not zero. In this case, however, its value results from a fairly complex interplay between the decision variables—as is normally the case, in fact. In general, the values of the shadow prices are not at all obvious and cannot be estimated from any individual equations.

For the third constraint, a unit increase in b_3 implies a unit increase in X_3, which increases cost by 10. The increase in X_3, however, makes it possible to decrease other decision variables: the unit increase in X_3 contributes to meeting the second constraint and thus makes it possible to reduce X_1 by 1, which reduces cost by 4. The net change in the objective function per unit increase in b_3 is, then,

$$SP_3 = +10 - 4 = 6$$

In summary we have:

$$\mathbf{SP}^* = (0, 4, 6)$$

Note that the zeros of the optimal slack variables correspond to nonzero shadow prices, and vice-versa.

Range of shadow prices. The range of validity for the shadow price of a constraint that is not binding is easy to determine: the shadow price remains equal to zero as long as there is slack. In our problem, the first constraint is not binding when $b_1 = 2$, a point where the slack for the optimum solution is 6. There will continue to be slack for any value of b_1 up to 8 (=2 + slack). Conversely, there would always be slack if b_1 were decreased, so there is no lower bound on the range of the shadow price. The range of validity for $SP_1^* = 0$ is thus: $-\infty \leq b_1 \leq 8$.

The range of validity for a binding constraint is more difficult to determine. One needs to determine when another constraint begins to have an effect. For the second constraint we can lower b_2 (which implies lowering X_1, since X_3 is fixed by the third constraint) until the level of X_1 is restrained by the first constraint, that is until $X_1 = 1$ and $b_2 = 2$. If b_2 were less than 2, the first and third constraints would be binding, the second would not be and its shadow price would then be zero. Therefore the lower bound on the b_2 for $SP_2^* = 4$ is 2. When it comes to raising b_2, however, there is no other constraint that might lead to a change in SP_2^*, so that the range is unbounded. The range of validity for $SP_2^* = 4$ is thus:

$$2 \leq b_2 \leq +\infty$$

Finally, the range of validity for $SP_3^* = 6$ is:

$$0 \leq b_3 \leq 4$$

The shadow price becomes zero if $b_3 = 0$ and X_3 is no longer required. It also changes when b_3 is large enough to create slack in the second constraint and make the first constraint binding. This occurs at $b_3 = 4$: each additional unit beyond that requires an extra unit of X_3 (at a cost of 10) which can no longer be compensated by a decrease in X_1, since the first constraint forces X_1 to be at least 1. So for $b_3 > 4$, $SP_3 = 10$.

Change in the shadow prices. The third constraint provides a good illustration of the way shadow prices vary with constraints. As b_3 is increased, this lower bound constraint is tightened (the "floor is raised") as the shadow price increases. Conversely, as b_3 is lowered (as from $b_3 = 5$ to $b_3 = 3$), the constraint is relaxed (the "floor is lowered") and the shadow price decreases. Figure 6.5 shows the behavior.

Opportunity costs. Opportunity costs exist only for nonoptimal decision variables; they are zero for optimal decision variables by definition. Therefore, since X_1 and X_3 are both part of the optimal design, $OC_1 = OC_3 = 0$.

The opportunity cost of a decision variable that is not in the optimal solution, such as X_2, is determined by *both* the coefficient of that variable in the objective function *and* the nature of constraints. Introducing the nonoptimal variable in the design generally permits one to use less of others.

The use of a unit of nonoptimal X_2 thus leads to an immediate cost of $C_2 = 12$. This is compensated, however, by the possibility of lowering X_1 (not X_3 since its quantity is limited by the third constraint). Specifically, using a unit of X_2 allows us to satisfy the second constraint with two less units of X_1, reducing cost on that account by 8. The net worsening of the objective function due to the use of a unit of the nonoptimal X_2 is thus:

$$OC_2^* = 12 - 8 = 4$$

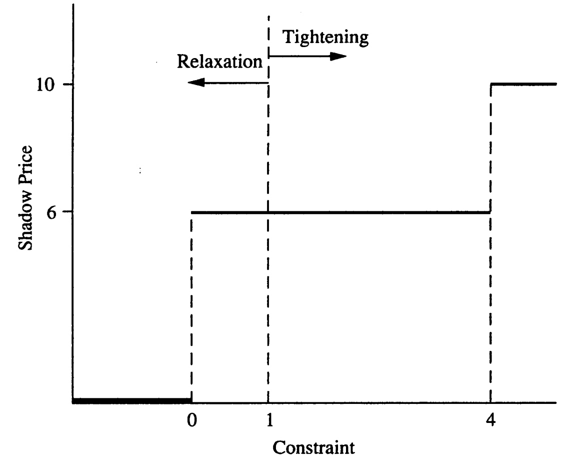

FIGURE 6.5
Shadow prices increase or stay constant as constraints are tightened in a linear program, and vice-versa.

In summary, we have:

$$\mathbf{OC}^* = (0, 4, 0)$$

Note that the zeros of the optimal opportunity costs correspond to the nonzero optimal decision variables, and vice-versa. This effect corresponds to the complementary slackness associated with shadow prices. The reason for this has to do with duality, as explained next.

6.5 PRIMAL-DUAL RELATIONSHIPS

For any given linear programming problem there is a dual—or "mirror image"—problem, which contains all the information of the original or primal problem (see Section 5.10). Particularly important aspects of this relationship are that

- the optimum value of the objective function for the primal problem equals that of the dual.
- the optimal solution to the dual contains all the sensitivity information in that of the primal, coded in different ways.

Specifically, it is true that

- the optimal values of the decision variables in the dual problem are the shadow prices of the primal (and vice-versa).
- the optimal value of the slack variables in the dual problem are the opportunity costs of the primal (and vice-versa).

For this reason, operations research jargon often refers to shadow prices as *dual variables*, and to opportunity costs as *dual slacks*.

For the user, really only interested in the results of the linear program, these relationships between the primal and dual problems are not especially useful. However, since much of the literature on linear programs uses these concepts to describe the results of an analysis, it is important to appreciate this language.

The equivalences stated above can be seen by looking at the example in the previous section. Since the dual to the minimization is

$$\text{Maximize:} \quad \mathbf{B}^T\mathbf{W}$$

$$\text{Subject to:} \quad \mathbf{A}^T\mathbf{W} \leq C^T \qquad \mathbf{W} \geq 0$$

we can write

$$\text{Maximize:} \quad 2W_1 + 5W_2 + W_3$$

$$\begin{aligned} \text{Subject to:} \quad 2W_1 + W_2 &\leq 4 \\ W_1 + 2W_2 &\leq 12 \\ W_2 + W_3 &\leq 10 \\ W_1, \quad W_2, \quad W_3 &\geq 0 \end{aligned}$$

By inspection, as before, we can see that we would like to use as much W_2 as possible, and then W_1 and W_3 as permitted by the constraints. This leads to

$$\mathbf{W}^* = (0, 4, 6)$$

which is equivalent to $\mathbf{SP}^*$ of the primal. We can also see that

$$\mathbf{OF}^* = 26$$

again, as for the primal. Further,

$$\mathbf{SV}^* = (0, 4, 0)$$

which is the $\mathbf{OC}^*$ of the primal.

A more detailed analysis would show that for this dual we would also have

$$\mathbf{SP}^* = (4, 0, 1)$$

the vector of $\mathbf{X}^*$ in the primal, and

$$\mathbf{OC}^* = (6, 0, 0)$$

the slack variables in the primal.

6.6 BREAK-EVEN ANALYSIS

Break-even analysis is an important form of sensitivity analysis. It is much used in practice, particularly in the comparison of alternative technologies for a system.

This method has the advantage of being able to deal with nonconvex feasible regions. It is thus not connected with linear programming, as are the previous methods. In fact, it is not associated with any form of optimization. Break-even analysis is conceptually and mathematically simple.

The essential idea derives from the fact that the creation of any product requires both an initial investment into plant and equipment and then, in production, the cost of labor, power, and supplies. The initial investment is known as the *capital cost*, C_k. The costs of production constitute the *variable cost*, C_v. (Section 14.4 presents these concepts in detail). The variable costs may be nonlinear, particularly as diseconomies of scale arise (see Section 4.5). In practice the variable costs are generally assumed to be linear, as a reasonable approximation when costs are difficult to define precisely. The total cost of production is then

$$\text{Total Cost} = C_k + C_v X$$

The relative size of the fixed and variable costs is a distinctive feature of different technologies. These costs provide the economic basis for choosing between alternative technologies. The typical situation is that one technology costs less to implement than the second, but has higher variable costs:

$$C_{k1} < C_{k2} \quad \text{but} \quad C_{v1} > C_{v2}$$

The question is: Which technology should be chosen? Assuming that they perform the same task, this resolves into an economic issue: Which costs less? The answer depends on the level of production: the first technology is cheaper when volume is low, the second will be when volume is high enough. The crucial issue thus concerns the *break-even point*, the level of production where the costs of one equal that of the other, where the economic advantage shifts from the first technology to the second. Break-even analysis defines this the point, which provides the economic basis for technology choice.

Figure 6.6 illustrates the situation. Notice that the feasible region of costs, above the hatched line, is nonconvex. This is typical for questions of technology choice. Switching to a technology that is cheaper at higher volumes of production is a prime reason for economies of scale (see discussion in Section 4.5).

The analysis itself merely consists of solving for the break-even point. As can be seen, this requires finding the volume of production, X, at which both technologies imply the same total cost. For linear costs this is:

$$C_{k1} + C_{v1}X = C_{k2} + C_{v2}X$$

The technologies with lower capital costs and higher variable costs are often low technologies. For example, a high-technology, automated factory costs more to install than an ordinary factory, but is designed to save extensively on the variable costs of labor. Yet important exceptions to this pattern exist: sometimes the higher technology has lower capital costs and higher variable costs. Thus it is for the cost of transport in undeveloped areas. Air transport, which only requires airfields at each destination, often costs less initially than road transport that involves the construction of a highway network. Once roads are built, however, trucks carry goods much more cheaply than aircraft. As part of our work for the U.S. AID and the Government of Colombia, referred to in the applications of Sections 2.6 and 3.4, we did break-even analyses for exactly this situation.

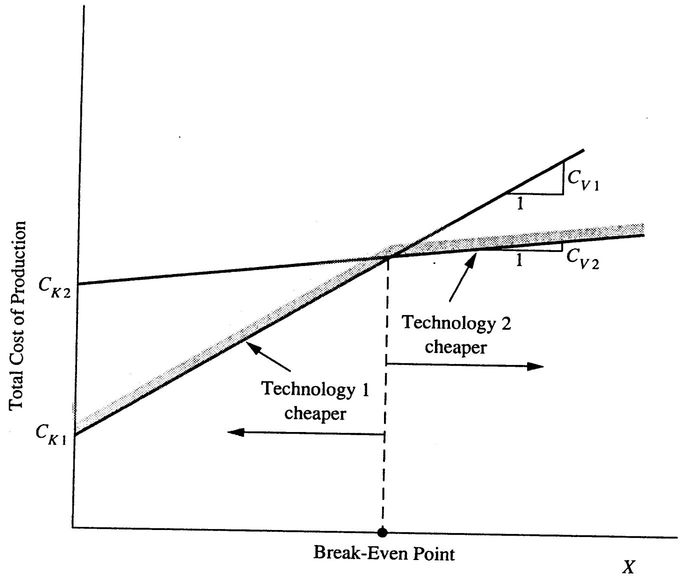

FIGURE 6.6
The break-even point is the volume of production where two technologies lead to the same total cost of production.

6.7 APPLICATION

The collection and disposal of municipal solid waste—garbage—illustrates well the role of sensitivity analyses. This section focuses on how these results were extracted from the optimization and then used in the design of the system. Garbage disposal is an expensive operation, especially in big cities. It requires a lot of transportation and a lot of labor. The amount of trucking required is surprisingly large, especially as existing disposal facilities reach capacity and the expanding cities force new sites to be further and further away. Distances of 10 to 20 miles are common, and some cities have been forced to use dumps 50 to 100 miles away. Garbage collection typically requires many workers to pick up barrels and tip them into the garbage trucks. An efficient system must use its labor and trucks effectively.

Two devices are used in the United States to improve the efficiency of the system. The first is a compactor truck, a vehicle with a mechanism that squeezes the garbage tightly as it is picked up; this permits much more to be loaded into the collection vehicles. The second is the use of transfer stations, at which garbage is unloaded from the compactor trucks for reloading on very large trucks for carting to the ultimate disposal site. The operation of transfer stations is expensive, but they offer two major possibilities for reducing costs:

they permit the use of semi-trailers, which could not be introduced directly in city streets; and they avoid wasting the time of the expensive compactor truck and its crew of garbage collectors on a long haul to a distant dump.

The optimization of a garbage collection system requires careful design of the size, number and location of the transfer points, as well as the corresponding routes for the trucks. The number of possible combinations is very, very large; the constraints on routes, working hours, vehicle capacities, and so on are also enormous. Fortunately it can be formulated and solved by linear programming.

The following illustrations of sensitivity analysis come from an extensive study of Baltimore. It is a good example of the general case of garbage collection, but naturally the specific parameters differ from city to city.

Specification sensitivity analysis played an essential part in the design. The procedure was to optimize the system for one set of specifications, then to vary these constraints, and finally to plot the results to determine the specifications that gave the least expensive results.

The capacity of the transfer point was one specification that was investigated in detail. In any particular analysis this specification was entered as an upper bound constraint on the capacity. Varying this constraint therefore defined the shadow price on the capacity. Figure 6.7 illustrates the results. At the lowest capacities, the size of the transfer point is a binding constraint and its shadow price is nonzero. As the capacity is enlarged, this upper bound constraint is

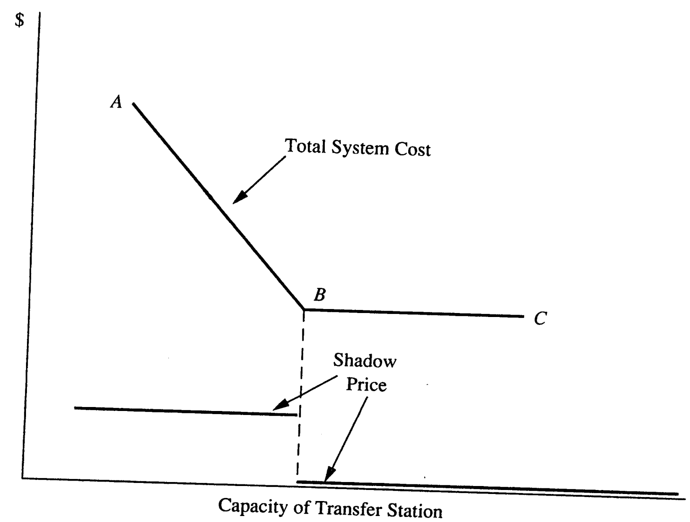

FIGURE 6.7
Changing the upper bound constraint on capacity reduces costs at a constant rate until this constraint is no longer binding.

relaxed, the total cost decreases at the constraint rate defined by the shadow price (as the line *AB* shows). Ultimately, the constraint is no longer binding, there is slack, the shadow price equals zero, and further increases in capacity no longer reduce cost (the line *BC*). (As a second-order effect, the cost actually rises from *B* to *C*, due to the capital cost of a larger transfer station, a factor not included in the minimization of the operating costs by the linear program.)

The number of garbage collections per week was a different kind of specification that was investigated. This was not directly entered into the formulation of the linear program as a constraint; it altered the amount of garbage that had to be picked up in any operation. The range of validity on the shadow price of that constraint did not go as far as the 50% increase implied by going from three collections a week to two, so that shadow prices were useless in this instance. Consequently, the optimization was simply performed twice and the effect noted: collecting three times a week only cost 4% more (assuming people do not create more garbage).

The lack of sensitivity of cost to the number of collections illustrates a point about linear programming and optimization in general: one's intuition about systems is frequently quite wrong. Most people, when asked about the difference in cost between two or three collections a week, immediately think that the better service—with half again as many pickups—should cost about 50% more. The reason this is not so has to do with the way constraints define the solution. In this case, the principal constraint is the volume of the truck, which defines the number of unproductive trips to the dump; if the amount of trash to be collected

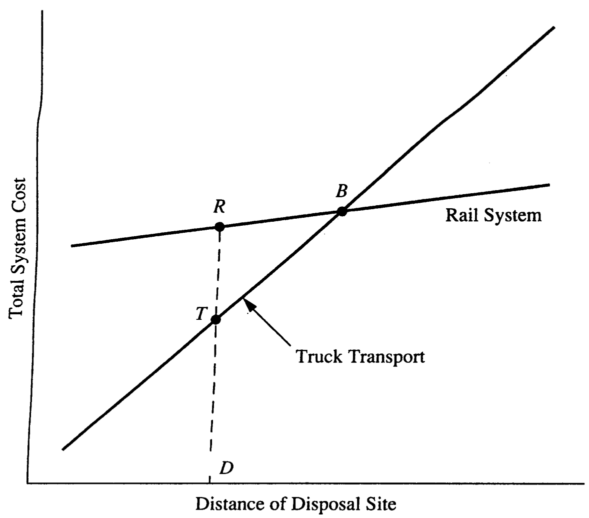

FIGURE 6.8
The relative costs of truck and rail transport define the opportunity costs of the more expensive system and the break-even point.

does not change, the number of unproductive trips stays about the same and so does the total cost. One of the advantages of linear programming is that it helps bring out these counter-intuitive aspects of a system.

Opportunity costs were used in the design to explore the viability of rail as an alternative way to carry the garbage from the transfer station to the disposal site. Figure 6.8 illustrates the approach. The analysis indicated that for the distance (*D*) involved for Baltimore, truck transport was cheaper (at *T*) than rail (at *R*). The opportunity cost of the uneconomic solution was thus the difference in total cost, $(R - T)$, divided by the tonnage involved. This gave the city the guideline it needed in negotiating with the railroad about better prices. If the reduction in the cost of rail did not exceed its opportunity costs, that alternative should not be part of the solution.

A break-even analysis can be derived from the same data. The intersection of the cost functions for rail and truck transport, at *B*, indicate the distance to the dump that would make rail attractive. This information is useful to planners as they visualize that future disposal sites are indeed going to be further and further from the city.

Overall, the optimization of the garbage collection system was particularly useful to Baltimore because it provided extensive sensitivity analysis. This information enabled the planners and city managers to design the system, to specify its constraints and the mode of operation that would be most economical. The original optimization was, in this case, principally useful as a starting point for the sensitivity analysis.

REFERENCES

de Neufville, R., and Marks, D. H. (1974). *Systems Planning and Design: Case Studies in Modeling, Optimization and Evaluation*, Prentice-Hall, Englewood Cliffs, NJ.

Eppen, G. D., and Gould, F. J. (1985). "Linear Programs: Computer Analysis and Interpreting Sensitivity Output," Ch. 6, *Quantitive Concepts for Management: Decision-Making Without Algorithms*, 2nd ed., Prentice-Hall, Englewood Cliffs, NJ.

Marks, D. H. and Liebman, J. C. (1971), "Location Models: A Solid Waste Collection Example," *ASCE Journal of the Urban Planning and Development Division*, 99, pp. 15–30.

White, J. A., Agee, M. H., Case, K. E. (1984). "Break-Even, Sensitivity and Risk Analysis," Chapter 8, *Principles of Engineering Economic Analysis*, 2nd ed., Wiley, New York.

PROBLEMS

6.1. *Sensitivity Exercise I*
For the linear program

Maximize: $Z = X_1 - X_2$

Subject to:

$$\begin{aligned} X_2 &\geq 5 \\ 3X_1 + 5X_2 &\leq 90 \\ -3X_1 + X_2 &\leq 0 \\ X_1 - X_2 &\leq 6 \\ -X_1 - 3X_2 &\leq 6 = b_5 \\ X_1,\ X_2 &\geq 0 \end{aligned}$$

(*a*) Sketch the feasible region.
(*b*) Determine the optimum solution.
(*c*) Find the shadow price on the fifth constraint (SP_5) at the optimum solution.
(*d*) Determine the range of b_5 for which the SP_5 is constant.
(*e*) Find SP_5 for b_5 above the range on the current value of SP_5; for b_5 below the range.
(*f*) Plot SP_5 as a function of b_5 for $-\infty < b_5 < \infty$.

6.2. *Sensitivity Exercise II*
Given the linear program

$$\begin{array}{ll} \text{Maximize:} & Z = X_1 + 2X_2 + 5X_3 \\ \text{Subject to:} & X_1 + 3X_2 \leq 6 \\ & 2X_2 + X_3 \leq 3 \\ & X_1 + X_3 \leq 15 \\ & X_1, \quad X_2 \geq 0 \end{array}$$

Determine by inspection and/or experimentation

(*a*) the optimal value of the objective function.
(*b*) the optimal values of the decision variables.
(*c*) the opportunity costs for any nonoptimal decision variables.
(*d*) the shadow prices on the constraints.
(*e*) the value of the slack variable associated with each constraint at the optimum.

6.3. *Sensitivity Exercise III*
Given the linear program

$$\begin{array}{ll} \text{Maximize:} & Z = 6X_1 + 3X_2 + 15X_3 \\ \text{Subject to:} & X_1 + X_3 \geq 1 \\ & 3X_1 + 2X_2 \geq 2 \\ & X_2 + X_3 \geq 5 \end{array}$$

answer the same questions as for Problem 6.2.

6.4. *Graphic Analysis*
(*a*) Graph the following primal linear program and locate its optimal solution.

$$\begin{array}{ll} \text{Maximize:} & Z = 3X_1 + 2X_2 \\ \text{Subject to:} & 3X_1 + 4X_2 \leq 12 \\ & -2X_1 + X_2 \leq 4 \\ & X_1, \quad X_2 \geq 0 \end{array}$$

(*b*) How much could the objective function be changed without affecting the optimal combination of inputs?
(*c*) Find the shadow prices on the constraints by evaluating w_i in the vicinity of the primal optimum. What are the opportunity costs of using each of the inputs?

6.5. *Primal-Dual*
For the linear program

$$\begin{array}{ll} \text{Minimize:} & Z = 6X_1 + 12X_2 + 6X_3 \\ \text{Subject to:} & X_1 + 2X_3 \geq 6 \\ & 8X_2 + 3X_3 \geq 2 \\ & X_1 + 3X_2 \geq 3 \\ & X_i \geq 0 \end{array}$$

determine by inspection and/or experimentation

(*a*) the optimal value of the objective function.
(*b*) the optimal value of the decision variables.
(*c*) the opportunity cost for any nonoptimal decision variables.
(*d*) the shadow prices of the constraints.
(*e*) the value of the slack variable associated with each constraint at the optimum.
(*f*) the dual of the original problem.

For the dual, determine

(*g*) the optimal value of the objective function.
(*h*) the optimal value of the decision variables.
(*i*) the opportunity costs for any nonoptimal decision variables.
(*j*) the shadow prices on the constraints.
(*k*) the value of the slack variable associated with each constraint at the optimum.
(*l*) each of the variables (or set of variables) in parts (*a*) through (*e*) should correspond to one of the variables (or sets) in parts (*g*) through (*k*). Match them up.

6.6. *Iron Alloys*
As manager of an iron foundry, you need to determine the minimum cost charge to a cupola melting furnace producing a cast iron with the following constraints:

Impurities	**Other constraints**
3.2 ≤ Carbon ≤ 3.5%	Steel scrap #2 = 30%
1.4 ≤ Manganese ≤ 1.6%	
2.7 ≤ Silicon ≤ 3.0%	
0.30 ≤ Chromium ≤ 0.45%	

(All compositions are specified in percent by weight)

Assume you have already formulated a linear programming problem and the computer has determined the following:

Raw material	Value (% in mix)	Opportunity cost ($/ton)	Cost ($/ton)
Pig iron	72.1		58
Silvery pig	3.0		120
Ferrosilicon			
#1	0	12	128
#2	2.9		120
Master alloy			
#1	1.0		200
#2	0	38	260
#3	0	17	238
Silicon carbide	1.0		160
Steel scrap			
#1	0	6.7	39
#2	20.0		30
#3	0	4.2	33

	Lower bound (min. constraint)		Upper bound (max. constraint)		
Material	Value (% in mix)	Shadow price ($/ton)	Value (% in mix)	Shadow price ($/ton)	Range (% in mix)
Carbon	3.2	0	3.5	2.4	3.3–3.8
Manganese	1.4	0	1.6	5.0	1.5–1.7
Silica	2.7	0	3.0	3.0	2.9–3.3
Chromium	0.3	0	0.45	6.1	0.35–0.65
Steel scrap #2	—	—	30.0	7.0	22–31

(*a*) Assuming the costs of the other materials remain constant, at what price should you consider using master alloy #3 in your charge?

(*b*) The manager of a small factory offers to accept a ton of alloy with 0.2 percent greater carbon content for each 0.1 ton of the steel scrap #1 you are willing to purchase at $38/ton. Will you accept? How much will you profit (or lose) on the deal?

(*c*) The same factory is using another alloy with less desirable properties than your alloy for $1/ton less. Assuming that the addition of 0.2% greater impurities into your alloy will not greatly affect the mechanical properties, list two ways you might reduce your costs by $1/ton.

6.7. *Gravel Pit*

A company extracts gravel from several sources to satisfy the demand from various clients. A linear program has been run to optimize the distribution of gravel, yielding the following results:

Source	Client	Cost/ton($)	Tons Shipped	Opportunity Cost ($)
1	1	1.80	0	0.25
1	2	1.40	367	0
1	3	3.04	0	1.21
1	4	2.10	535	0
1	5	1.53	0	0
1	6	1.82	0	0.07
2	1	1.60	798	0
2	2	1.25	214	0
2	3	1.90	0	0.10
2	4	2.50	0	0.50
2	5	1.70	0	0.20
2	6	1.40	0	0.17
3	1	2.20	0	0.72
3	2	1.60	0	0.23
3	3	1.97	723	0
3	4	2.30	0	0.18

Source	Tons available	Slack variable	Shadow price ($/ton)	Range
1	1270	368	0	$b > 802$
2	1012	0	0.15	$1379 > b > 798$
3	1570	473	0	$b > 1097$
4	1120	0	0.21	$1309 > b > 1092$
5	2117	516	0	$b > 1601$

Client	Tons contracted	Slack variable	Shadow price ($/ton)	Range
1	798	0	1.75	$1012 > b > 431$
2	881	0	1.40	$1683 > b > 567$
3	1382	0	1.97	$2117 > b > 961$
4	535	0	2.10	$735 > b > 0$
5	2147	0	1.20	$2384 > b > 1602$
6	1346	0	1.31	$1521 > b > 0$

(*a*) Why are the slack variables for all the client (demand) constraints equal to zero?
(*b*) Client #6 would like 100 tons of his gravel to be the peculiarly colored stone from source 2. How much extra should you charge him for it, or how large a refund should you give him?
(*c*) Client #5 finds that she will need 700 tons less than she expected. How large a refund can you afford to give her for reducing the size of her order?
(*d*) How much would you lose if source #2 had 300 tons less gravel than you expected?

6.8. *Food Distribution*
An island is spanned by a network of railroads. Vegetables, grown in the western hills (nodes 1,2,3,4), must be shipped to the rest of the country. A linear program was written to minimize the transportation cost, subject to

- limited production at each supply node.
- limited capacity of each railroad link.
- demand that must be satisfied at each of the demand nodes.
- continuity that must be maintained at each node.

The computer output produces the following information:

	Link flows at optimum			Link capacity constraints		
Link	Cost/unit (pesos)	Units carried	Opportunity costs	Capacity (maxflow)	Shadow price	Range
1 → 2	42	0	50	30	0	
1 → 5	68	63	—	63	12	59 < maxflow < 100
1 → 6	76	32	—	70	*X*	32 < maxflow
2 → 1	42	30	—	30	6	
2 → 3	107	0	123	21	0	
2 → 5	143	0	21	27	0	
2 → 6	113	40	—	60	0	40 < maxflow
3 → 2	107	0	91	21	0	
3 → 4	231	0	211	15	0	

Demand node constraints

Node	Demand	SP	Range
5	50	143	41 < demand 5 < 57
6	40	113	38 < demand 6 < 45
7	100	201	93 < demand 7 < 117

Note that each link has an actual flow (the variable) and a maximum flow (the constraint).
(*a*) What is the value of *X*, the shadow price on the link 1 → 6 constraint?
(*b*) Improvements on link 2 → 5 will reduce its cost to 120 pesos/unit. Should 2 → 5 be used after the improvements are completed?
(*c*) How much is it worth paying to increase the capacity of link 1 → 5 by 5 units?
(*d*) It is possible to satisfy some of the existing demands at node 6 by growing vegetables at that node: one unit can be produced for 100 pesos, *or* four units can be grown for 350 pesos. What does the computer output tell you about these options?

6.9. *Gelderland*
A consultant is using LP to maximize the benefits, in Gelders, of a new port facility for Gelderland. The constraints are:

1. Actual Cost ≤ U.S. $20 M.
2. Probability of an oil spill ≤ 1%.
3. Construction time ≤ 5 years.
4. Local workers used in construction ≥ 500.
5. Port capacity ≥ 20 ships at any time.

The output includes the following information:

Constraint	Shadow price	Range
1	15 G/$	$15 \leq b_1 \leq 22$M
2	600,000 G/%	$0.8 \leq b_2 \leq 7$%
3	0 G/year	$3.5 \leq b_3$
4	2000 G/worker	$420 \leq b_4 \leq 700$
5	1M G/ship	$15 \leq b_5 \leq 23$

Also, the following possible features were not part of the optimal design:

Feature	Opportunity cost
deep water tanker berth	15M G
hydrofoil marina	3M G
seaplane facility	2M G

(*a*) Properly impressing some bank officials will probably increase the size of the loan by $1M. How many Gelders can the nation afford to spend on a "sales effort" directed at these officials?
(*b*) An oil spill would cost the tourist industry about 10M Gelders. Would it pay to relax the environmental standards? (Hint: calculate the expected cost of relaxed standards.)

(*c*) Skilled foreign labor will speed construction of the port, at somewhat greater cost. Is this idea worth pursuing?
(*d*) It has been suggested that the port be built to accommodate less than 20 ships and that excess ships could be served by barges costing 500,000 G. Should the consultant design for 20 ships? 18 ships? Get more information (what kind)?

6.10. *SMC Factory II*
Having used LP to solve for the optimal design of SMC in Problem 5.14, you have obtained the following sensitivity information:

Raw material	% in mix (volume %)	Opportunity cost ($/lb)	Cost ($/lb)
Polyester	51	—	0.55
Vinylester	0	0.44	1.05
E-glass Fiber	24.5	—	1.00
Carbon Fiber	24.5	—	25.00
Kevlar Fiber	0	8	21.00

Constraint	Shadow price	Range: Max	Range: Min
Tensile < 200 ksi	0.019	296	103
Tensile > 100 ksi	0	—	—
Modulus > 10 msi	1.012	20	4.1
Specific gravity < 1.9	0	—	—
Fiber volume % < 78.5	0	—	—

(*a*) By how much must the price of Kevlar fall before it would be incorporated in the SMC formulation?
(*b*) Your client now needs SMC with a tensile strength not greater than 150 ksi. By how much and in what direction will the cost of this new formulation change (if at all)?
(*c*) The client is willing to relax the constraint on specific gravity by 10%. How much are you willing to pay for this change?
(*d*) If the modulus of the mixture must be at least 21 msi, will your production costs increase or fall and by how much?

6.11. *DEBON AIR*
DEBON AIR is looking into expanding into cargo transport from an already secure position in the crop dusting area. With a block of ready cash, DEBON plans to buy a new fleet. A linear program was commissioned to minimize the cost of this purchase, subject to various constraints. A portion of the results follow:

Aircraft type	Number purchased	Number available	*SP*($)	*OC*($)	Range
A	2	10	—	—	—
B	0	10	—	5×10^5	—
C	6	6	10^6	—	0–20

Capacity needs	Needed	Purchased	*SP*($)	Range
Short haul capacity	5×10^5 lbs	8×10^5 lbs	—	—
Long haul capacity	2×10^5 lbs	2×10^5 lbs	70	10^5–10^6

Fleet fuel efficiency requirement (fleet avg.)	Actual	*SP*($)	Range
10 mpg	10 mpg	10^6	1–15

Use the above information to determine quantitative answers. If a quantitative solution is not possible, state so and explain why.
(*a*) If one more plane of type C were available, how much less would the new fleet cost?
(*b*) If DEBON's long haul capacity needs had been 10^5 lbs higher, would the new fleet cost more or less than the current configuration? By how much?
(*c*) An aircraft manufacturer proposes an aircraft whose technical specifications and price are identical to those of plane A, with one exception—its fuel economy is 10% better. Should DEBON AIR buy this aircraft? Why or why not?
(*d*) The government circulates proposed regulations that would require future airline fleet expansions to be such that the mean fuel efficiency of the new craft is 12 mpg. In the DEBON protest, the cost of such a regulation to DEBON is detailed. What is its value?

6.12. *Computer Run*
This problem is to be solved using a PC. Formulate and solve the LP to answer the questions.

A small steel company makes specialty alloys to customer specification. One customer requires an alloy involving four metals, and has specified the following composition:

Metal	Requirements
A	No more than 15%
B	At least 25%
C	Between 40% and 50%
D	No more than 8%
Impurities	None

This steel company can acquire its ores from several sources. These ores contain both the required metals and impurities, which should be separated before the production of the alloy. The cost of this separation is the same for all ores and is quite small. The percent composition of these ores is

Ore	Metal A	Metal B	Metal C	Metal D	Impurities
1	20	20	40	0	20
2	15	0	20	5	60
3	0	40	40	0	20
4	15	15	20	0	50
5	20	20	20	20	20
6	5	8	17	10	60

Finally, there are limits to the amount of each ore available to the steelmaker:

Ore	Tons available	$/ton
1	2200	30.00
2	1800	10.00
3	1750	32.00
4	700	20.00
5	900	22.00
6	1700	12.00

(*a*) The customer wants 4000 tons of alloy. What is the minimum cost per ton of alloy? Which ores will be purchased? What percentage of the total metal requirement will ore 3 represent? What percentage of the total ore requirements?

(*b*) The customer calls to inform the steelmaker of an error in the specification of the alloy: there must now be no more than 14.5% of metal A in the alloy. What is the minimum penalty that should be charged to keep profits/ton fixed?

(*c*) The manager of the mine that supplies ore 4 is desperate for business. How much of a kickback must the steelmaker demand to assure the owner that 1% of the metal composing the alloy shipment will come from mine 4?

(*d*) The manager of the mine that supplies ore 3 proposes to announce a 15% increase in production. How would this change the cost/ton of the special alloy?

(*e*) The steelmaker feels that it may be possible to cheat on the 0% impurity level. How much could the profit/ton be improved by introducing a 1% impurity? (Assume that the cost of the impurity is zero.)

(*f*) In the discussion, the customer asks for an estimate on another order of special alloy. This alloy will be subject to the original (not the revised!!) specifications, except that the percentage of metal *D* must be no higher than 6%. Can you estimate the cost/ton of a 4000 ton order of this alloy?

(*g*) Suppose that the mine supplying ore 6 has just been struck. The manager of the mine can only supply 1400 tons of ore. How much of a penalty should the steelmaker exact from the mine manager to preserve the profit?

6.13. *Stomp II*

STOMP, Inc., having taken our advice from Problem 5.6, has prospered. Still manufacturing stereo nameplates and bird cage hinges, it has expanded its operations substantially. It now purchases varying grades of metal from several suppliers and has managed to expand the market for nameplates, which is now able to absorb 80 cases of plates a week. Luckily, the demand for bird cage hinges has remained strong and the market easily absorbs all that STOMP can produce. STOMP's only trouble is its union, which still will not allow its members to operate production machinery for more than 40 hours per week.

A new linear program, maximizing weekly profit, has been commissioned. A portion of the output follows.

Metal supplier	Price ($/ton)	Amount bought (ton/wk)	Supply constraint (ton/wk)	Shadow price	Range	Opportunity cost
I	20	1000	1000	1.5	$600 \le X \le 1300$	—
II	25	0	2000	—		2.5
III	22	400	1000	—		—
.	.	.	.	.		.
.	.	.	.	.		.
.	.	.	.	.		.

	Price	Amount bought	Supply constraint	Shadow price	Range
Machine Time	$25/hr	40 hr/wk	40 hr/wk	150	$35 \le X \le 46$
.	.	.	.	.	.
.	.	.	.	.	.
.	.	.	.	.	.

	Price	Amount sold	Demand constraint	Shadow price	Range
Cases of Nameplates	$200/case	80 cases/wk	80 cases/wk	150	$65 \le X \le 90$
.	.	.	.	.	.
.	.	.	.	.	.
.	.	.	.	.	.

(*a*) Supplier I wanted to cut his weekly deliveries to STOMP by 100 tons/wk. How much can STOMP afford to offer him to get him to maintain his current weekly deliveries?

(*b*) Supplier II wants to get some business from STOMP. What price should she charge in order to attract the attention of STOMP's metal buyers?

(*c*) STOMP is offered an overseas franchise in nameplates, which would allow STOMP to sell 15 more cases of nameplates each week. The price of the franchise is $2250/wk. Should STOMP buy the franchise?

6.14. *Gravel Company*
A gravel company has contracts to supply gravel to four construction projects in Massachusetts:

1. A highway in Berkshire county (1500 tons).
2. A shopping mall in Salem (1000 tons).
3. A runway extension at Logan Airport (1700 tons).
4. An urban development project in Worcester (500 tons).

The company owns three gravel pits in New Hampshire, in New York, and in Connecticut. The marginal cost of producing and shipping a ton of gravel from pit to user is about:

	Pit		
Client	**NH**	**NY**	**CT**
1	$2.20	$1.20	$1.70
2	1.60	2.20	1.90
3	2.50	3.10	2.70
4	2.10	2.00	1.80

No more than 40% of any project's gravel may come from the Connecticut pit, which contains a fair amount of silt. Only 800 tons of gravel are currently available at the New York pit. The objective is to minimize cost.
(*a*) What are the constraint equations?
(*b*) Indicate what sensitivity output or other data from your LP will answer each of the following questions:
(i) What is the most we should pay for offshore dredging of gravel for the Logan project?
(ii) How much more would it cost to require that 50 tons of the New York gravel be used on the Salem project?
(iii) What is the maximum price at which Connecticut gravel would be used for the Berkshire project?
(iv) How much of the 800 tons of New York gravel will be unused?

CHAPTER 7

DYNAMIC PROGRAMMING

7.1 CONCEPT

Dynamic programming is a powerful method of optimization which overcomes the three major drawbacks of linear programming. It can be applied to problems that

- are nonlinear.
- have nonconvex feasible regions.
- have discontinuous variables.

These substantial advantages are obtained at a significant price, however. Dynamic programming is limited to dealing with problems with

- relatively few constraints.

Dynamic programming complements linear programming. It works well on the nonlinear, nonconvex problems that are impossible for linear programming, but cannot handle the large number of constraints that linear programming deals with routinely. In practical situations, which may actually involve both nonconvex feasible regions and many constraints, analysts will have to make approximations