
What Is Optimization?

1.1 Introduction

Optimization, or constrained optimization, or mathematical programming, is a mathematical procedure for determining optimal allocation of scarce resources. Optimization, and its most popular special form, Linear Programming (LP), has found practical application in almost all facets of business, from advertising to production planning. Transportation and aggregate production planning problems are the most typical objects of LP analysis. The petroleum industry was an early intensive user of LP for solving fuel blending problems.

It is important for the reader to appreciate at the outset that the “programming” in Mathematical Programming is of a different flavor than the “programming” in Computer Programming. In the former case, it means to plan and organize (as in “Get with the program!”). In the latter case, it means to write instructions for performing calculations. Although aptitude in one suggests aptitude in the other, training in the one kind of programming has very little direct relevance to the other.

For most optimization problems, one can think of there being *two important classes of objects*. The first of these is *limited resources*, such as land, plant capacity, and sales force size. The second is *activities*, such as “produce low carbon steel,” “produce stainless steel,” and “produce high carbon steel.” *Each activity consumes* or possibly *contributes* additional amounts of the *resources*. The problem is to determine the best combination of activity levels that does not use more resources than are actually available. We can best gain the flavor of LP by using a simple example.

1.2 A Simple Product Mix Problem

The Enginola Television Company produces two types of TV sets, the “Astro” and the “Cosmo”. There are two production lines, one for each set. The Astro production line has a capacity of 60 sets per day, whereas the capacity for the Cosmo production line is only 50 sets per day. The labor requirements for the Astro set is 1 person-hour, whereas the Cosmo requires a full 2 person-hours of labor. Presently, there is a maximum of 120 man-hours of labor per day that can be assigned to production of the two types of sets. If the profit contributions are \$20 and \$30 for each Astro and Cosmo set, respectively, what should be the daily production?

2 Chapter 1 What is Optimization?

A structured, but verbal, description of what we want to do is:

Maximize Profit contribution
subject to Astro production less-than-or-equal-to Astro capacity,
Cosmo production less-than-or-equal-to Cosmo capacity,
Labor used less-than-or-equal-to labor availability.

Until there is a significant improvement in artificial intelligence/expert system software, we will need to be more precise if we wish to get some help in solving our problem. We can be more precise if we define:

A = units of Astros to be produced per day,
 C = units of Cosmos to be produced per day.

Further, we decide to measure:

Profit contribution in dollars,
Astro usage in units of Astros produced,
Cosmo usage in units of Cosmos produced, and
Labor in person-hours.

Then, a precise statement of our problem is:

Maximize $20A + 30C$ (Dollars)
subject to $A \leq 60$ (Astro capacity)
 $C \leq 50$ (Cosmo capacity)
 $A + 2C \leq 120$ (Labor in person-hours)

The first line, “Maximize $20A+30C$ ”, is known as the *objective function*. The remaining three lines are known as *constraints*. Most optimization programs, sometimes called “solvers”, assume all variables are constrained to be nonnegative, so stating the constraints $A \geq 0$ and $C \geq 0$ is unnecessary.


Using the terminology of resources and activities, there are three resources: Astro capacity, Cosmo capacity, and labor capacity. The activities are Astro and Cosmo production. It is generally true that, with each constraint in an optimization model, one can associate some resource. For each decision variable, there is frequently a corresponding physical activity.

1.2.1 Graphical Analysis

The Enginola problem is represented graphically in Figure 1.1. The feasible production combinations are the points in the lower left enclosed by the five solid lines. We want to find the point in the feasible region that gives the highest profit.

To gain some idea of where the maximum profit point lies, let’s consider some possibilities. The point $A = C = 0$ is feasible, but it does not help us out much with respect to profits. If we spoke with the manager of the Cosmo line, the response might be: “The Cosmo is our more profitable product. Therefore, we should make as many of it as possible, namely 50, and be satisfied with the profit contribution of $30 \times 50 = \$1500$.”

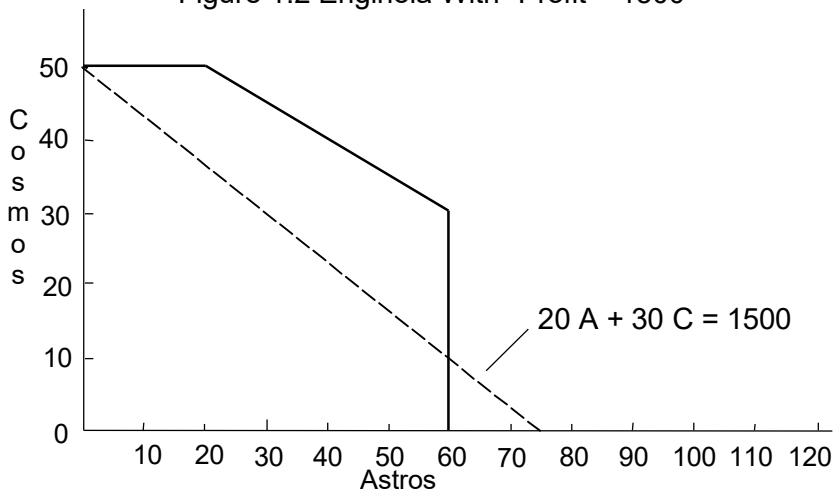
Figure 1.1 Feasible Region for Enginola



You, the thoughtful reader, might observe there are many combinations of A and C , other than just $A = 0$ and $C = 50$, that achieve \$1500 of profit. Indeed, if you plot the line $20A + 30C = 1500$ and add it to the graph, then you get Figure 1.2. Any point on the dotted line segment achieves a profit of \$1500. Any line of constant profit such as that is called an iso-profit line (or iso-cost in the case of a cost minimization problem).


If we next talk with the manager of the Astro line, the response might be: “If you produce 50 Cosmos, you still have enough labor to produce 20 Astros. This would give a profit of $30 \times 50 + 20 \times 20 = \1900 . That is certainly a respectable profit. Why don’t we call it a day and go home?”

Figure 1.2 Enginola With "Profit = 1500"

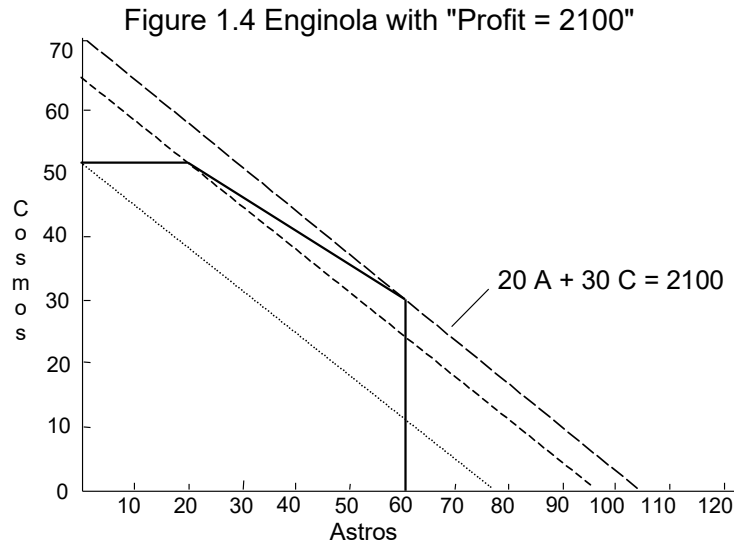


4 Chapter 1 What is Optimization?

Our ever-alert reader might again observe that there are many ways of making \$1900 of profit. If you plot the line $20A + 30C = 1900$ and add it to the graph, then you get Figure 1.3. Any point on the higher rightmost dotted line segment achieves a profit of \$1900.



Now, our ever-perceptive reader makes a leap of insight. As we increase our profit aspirations, the dotted line representing all points that achieve a given profit simply shifts in a parallel fashion. Why not shift it as far as possible for as long as the line contains a feasible point? This last and best feasible point is $A = 60$, $C = 30$. It lies on the line $20A + 30C = 2100$. This is illustrated in Figure 1.4. Notice, even though the profit contribution per unit is higher for Cosmo, we did not make as many (30) as we feasibly could have made (50). Intuitively, this is an optimal solution and, in fact, it is. The graphical analysis of this small problem helps understand what is going on when we analyze larger problems.



1.3 Linearity

We have now seen one example. We will return to it regularly. This is an example of a linear mathematical program, or LP for short. Solving linear programs tends to be substantially easier than solving more general mathematical programs. Therefore, it is worthwhile to dwell for a bit on the linearity feature.

Linear programming applies *directly* only to situations in which the effects of the different activities in which we can engage are linear. For practical purposes, we can think of the linearity requirement as consisting of three features:

1. *Proportionality*. The effects of a single variable or activity by itself are proportional (e.g., doubling the amount of steel purchased will double the dollar cost of steel purchased).
2. *Additivity*. The interactions among variables must be additive (e.g., the dollar amount of sales is the sum of the steel dollar sales, the aluminum dollar sales, etc.; whereas the amount of electricity used is the sum of that used to produce steel, aluminum, etc).
3. *Continuity*. The variables must be continuous (i.e., fractional values for the decision variables, such as 6.38, must be allowed). If both 2 and 3 are feasible values for a variable, then so is 2.51.

A model that includes the two decision variables “price per unit sold” and “quantity of units sold” is probably not linear. The proportionality requirement is satisfied. However, the interaction between the two decision variables is multiplicative rather than additive (i.e., $dollar\ sales = price \times quantity$, not $price + quantity$).

If a supplier gives you quantity discounts on your purchases, then the cost of purchases will not satisfy the proportionality requirement (e.g., the total cost of the stainless steel purchased may be less than proportional to the amount purchased).


A model that includes the decision variable “number of floors to build” might satisfy the proportionality and additivity requirements, but violate the continuity conditions. The recommendation to build 6.38 floors might be difficult to implement unless one had a designer who was ingenious with split level designs. Nevertheless, the solution of an LP might recommend such fractional answers.

The possible formulations to which LP is applicable are substantially more general than that suggested by the example. The objective function may be minimized rather than maximized; the direction of the constraints may be \geq rather than \leq , or even $=$; and any or all of the parameters (e.g., the 20, 30, 60, 50, 120, 2, or 1) may be negative instead of positive. The principal restriction on the class of problems that can be analyzed results from the linearity restriction.

Fortunately, as we will see later in the chapters on integer programming and quadratic programming, there are other ways of accommodating these violations of linearity.

Figure 1.5 illustrates some nonlinear functions. For example, the expression $X \times Y$ satisfies the proportionality requirement, but the effects of X and Y are not additive. In the expression $X^2 + Y^2$, the effects of X and Y are additive, but the effects of each individual variable are not proportional.

Figure 1.5: Nonlinear Relations




1.4 Analysis of LP Solutions

When you direct the computer to solve a math program, the possible outcomes are indicated in Figure 1.6.

For a properly formulated LP, the leftmost path will be taken. The solution procedure will first attempt to find a feasible solution (i.e., a solution that simultaneously satisfies all constraints, but does not necessarily maximize the objective function). The rightmost, “No Feasible Solution”, path will be taken if the formulator has been too demanding. That is, two or more constraints are specified that cannot be simultaneously satisfied. A simple example is the pair of constraints $x \leq 2$ and $x \geq 3$. The nonexistence of a feasible solution does not depend upon the objective function. It depends solely upon the constraints. In practice, the “No Feasible Solution” outcome might occur in a large complicated problem in which an upper limit was specified on the number of productive hours available and an unrealistically high demand was placed on the number of units to be produced. An alternative message to “No Feasible Solution” is “You Can’t Have Your Cake and Eat It Too”.

Figure 1.6 Solution Outcomes



If a feasible solution has been found, then the procedure attempts to find an optimal solution. If the “Unbounded Solution” termination occurs, it implies the formulation admits the unrealistic result that an infinite amount of profit can be made. A more realistic conclusion is that an important constraint has been omitted or the formulation contains a critical typographical error.

We can solve the Enginola problem in LINGO by typing the following:

```

MODEL:
  MAX = 20*A + 30*C;
  A    <= 60;
  C    <= 50;
  A + 2*C <= 120;

END
  
```

We can solve the problem in the Windows version of LINGO by clicking on the red “bullseye” icon. We can get the following solution report by clicking on the “X=” icon”:

```

Objective value: 2100.000

Variable      Value      Reduced Cost
A             60.00000      0.00000
C             30.00000      0.00000

Row   Slack or Surplus   Dual Price
1     2100.00000          1.00000
2     0.00000             5.00000
3     20.00000            0.00000
4     0.00000            15.00000
  
```

The output has three sections, an informative section, a “variables” section, and a “rows” section. The second two sections are straightforward. The maximum profit solution is to produce 60 Astros and 30 Cosmos for a profit contribution of \$2,100. This solution will leave zero slack in row 2 (the constraint $A \leq 60$), a slack of 20 in row 3 (the constraint $C \leq 50$), and no slack in row 4 (the constraint $A + 2C \leq 120$). Note $60 + 2 \times 30 = 120$.

The third column contains a number of opportunity or marginal cost figures. These are useful by-products of the computations. The interpretation of these “reduced costs” and “dual prices” is discussed in the next section. The reduced cost/dual price section is optional and can be turned on or off by clicking on LINGO | Options | General Solver | Dual Computations | Prices.

1.5 Sensitivity Analysis, Reduced Costs, and Dual Prices

Realistic LPs require large amounts of data. Accurate data are expensive to collect, so we will generally be forced to use data in which we have less than complete confidence. A time-honored adage in data processing circles is “garbage in, garbage out”. A user of a model should be concerned with how the recommendations of the model are altered by changes in the input data. Sensitivity analysis is the term applied to the process of answering this question. Fortunately, an LP solution report provides supplemental information that is useful in sensitivity analysis. This information falls under two headings, reduced costs and dual prices.

Sensitivity analysis can reveal which pieces of information should be estimated most carefully. For example, if it is blatantly obvious that a certain product is unprofitable, then little effort need be expended in accurately estimating its costs. The first law of modeling is "do not waste time accurately estimating a parameter if a modest error in the parameter has little effect on the recommended decision".

1.5.1 Reduced Costs

Associated with each variable in any solution is a quantity known as the *reduced cost*. If the units of the objective function are dollars and the units of the variable are gallons, then the units of the reduced cost are dollars per gallon. The reduced cost of a variable is the amount by which the profit contribution of the variable must be improved (e.g., by reducing its cost) before the variable in question would have a positive value in an optimal solution. Obviously, a variable that already appears in the optimal solution will have a zero reduced cost.

It follows that a second, correct interpretation of the reduced cost is that it is the rate at which the objective function value will deteriorate if a variable, currently at zero, is arbitrarily forced to increase a small amount. Suppose the reduced cost of x is \$2/gallon. This means, if the profitability of x were increased by \$2/gallon, then 1 unit of x (if 1 unit is a “small change”) could be brought into the solution without affecting the total profit. Clearly, the total profit would be reduced by \$2 if x were increased by 1.0 without altering its original profit contribution.

1.5.2 Dual Prices

Associated with each constraint is a quantity known as the *dual price*. If the units of the objective function are cruzeiros and the units of the constraint in question are kilograms, then the units of the dual price are cruzeiros per kilogram. The dual price of a constraint is the rate at which the objective function value will improve as the right-hand side or constant term of the constraint is increased a small amount.

Different optimization programs may use different sign conventions with regard to the dual prices. The LINGO computer program uses the convention that a positive dual price means increasing the right-hand side in question will improve the objective function value. On the other hand, a negative dual price means an increase in the right-hand side will cause the objective function value to deteriorate. A zero dual price means changing the right-hand side a small amount will have no effect on the solution value.

It follows that, under this convention, \leq constraints will have nonnegative dual prices, \geq constraints will have nonpositive dual prices, and $=$ constraints can have dual prices of any sign. Why?

Understanding Dual Prices. It is instructive to analyze the dual prices in the solution to the Enginola problem. The dual price on the constraint $A \leq 60$ is \$5/unit. At first, one might suspect this quantity should be \$20/unit because, if one more Astro is produced, the simple profit contribution of this unit is \$20. An additional Astro unit will require sacrifices elsewhere, however. Since all of the labor supply is being used, producing more Astros would require the production of Cosmos to be reduced in order to free up labor. The labor tradeoff rate for Astros and Cosmos is $\frac{1}{2}$. That is, producing one more Astro

implies reducing Cosmo production by $\frac{1}{2}$ of a unit. The net increase in profits is $\$20 - (1/2)*\$30 = \$5$, because Cosmos have a profit contribution of $\$30$ per unit.

Now, consider the dual price of $\$15/\text{hour}$ on the labor constraint. If we have 1 more hour of labor, it will be used solely to produce more Cosmos. One Cosmo has a profit contribution of $\$30/\text{unit}$. Since 1 hour of labor is only sufficient for one half of a Cosmo, the value of the additional hour of labor is $\$15$.

1.6 Unbounded Formulations

If we forget to include the labor constraint and the constraint on the production of Cosmos, then an unlimited amount of profit is possible by producing a large number of Cosmos. This is illustrated here:


$$\begin{aligned} \text{MAX} &= 20 * A + 30 * C; \\ A &\leq 60; \end{aligned}$$

This generates an error window with the message:

UNBOUNDED SOLUTION

There is nothing to prevent C from being infinitely large. The feasible region is illustrated in Figure 1.7. In larger problems, there are typically several unbounded variables and it is not as easy to identify the manner in which the unboundedness arises.

Figure 1.7 Graph of Unbounded Formulation



1.7 Infeasible Formulations

An example of an infeasible formulation is obtained if the right-hand side of the labor constraint is made 190 and its direction is inadvertently reversed. In this case, the most labor that can be used is to produce 60 Astros and 50 Cosmos for a total labor consumption of $60 + 2 \times 50 = 160$ hours. The formulation and attempted solution are:

```
MAX = (20 * A) + (30 * C);
A <= 60;
C <= 50;
A + 2 * C >= 190;
```

A window with the error message:

```
NO FEASIBLE SOLUTION.
```


will print. The reports window will generate the following:

Variable	Value	Reduced Cost
A	60.00000	0.0000000
C	50.00000	0.0000000
Row	Slack or Surplus	Dual Price
1	2700.000	0.0000000
2	0.0000000	1.000000
3	0.0000000	2.000000
4	-30.00000	-1.000000

This “solution” is infeasible for the labor constraint by the amount of 30 person-hours ($190 - (1 \times 60 + 2 \times 50)$). The dual prices in this case give information helpful in determining how the infeasibility arose. For example, the +1 associated with row 2 indicates that increasing its right-hand side by one will decrease the infeasibility by 1. The +2 with row 3 means, if we allowed 1 more unit of Cosmo production, the infeasibility would decrease by 2 units because each Cosmo uses 2 hours of labor. The -1 associated with row 4 means that decreasing the right-hand side of the labor constraint by 1 would reduce the infeasibility by 1.

Figure 1.8 illustrates the constraints for this formulation.

Figure 1.8 Graph of Infeasible Formulation




1.8 Multiple Optimal Solutions and Degeneracy

For a given formulation that has a bounded optimal solution, there will be a unique optimum objective function value. However, there may be several different combinations of decision variable values (and associated dual prices) that produce this unique optimal value. Such solutions are said to be degenerate in some sense. In the Enginola problem, for example, suppose the profit contribution of *A* happened to be \$15 rather than \$20. The problem and a solution are:

```

MAX = 15 * A + 30 * C;
A <= 60;
C <= 50;
A + 2 * C <= 120;
Optimal solution found at step:          1
Objective value:                       1800.000
Variable      Value      Reduced Cost
  A           20.00000    0.0000000
  C           50.00000    0.0000000
Row   Slack or Surplus   Dual Price
  1           1800.000     1.000000
  2           40.00000     0.000000
  3           0.0000000     0.000000
  4           0.0000000     15.00000
    
```

Figure 1.9 Model with Alternative Optima



The feasible region, as well as a “profit = 1500” line, are shown in Figure 1.9. Notice the lines $A + 2C = 120$ and $15A + 30C = 1500$ are parallel. It should be apparent that any feasible point on the line $A + 2C = 120$ is optimal.

The particularly observant may have noted in the solution report that the constraint, $C \leq 50$ (i.e., row 3), has both zero slack and a zero dual price. This suggests the production of Cosmos could be decreased a small amount without any effect on total profits. Of course, there would have to be a compensatory increase in the production of Astros. We conclude that there must be an alternate optimum solution that produces more Astros, but fewer Cosmos. We can discover this solution by increasing the profitability of Astros ever so slightly. Observe:

```

MAX = 15.0001 * A + 30 * C;
A <= 60;
C <= 50;
A + 2 * C <= 120;

Optimal solution found at step:          1
Objective value:                      1800.006

Variable          Value          Reduced Cost
   A             60.00000          0.0000000
   C             30.00000          0.0000000

Row  Slack or Surplus          Dual Price
  1             1800.006          1.00000
  2             0.0000000          0.1000000E-03
  3             20.00000          0.0000000
  4             0.0000000          15.00000

```

As predicted, the profit is still about \$1800. However, the production of Cosmos has been decreased to 30 from 50, whereas there has been an increase in the production of Astros to 60 from 20.

1.8.1 The “Snake Eyes” Condition

Alternate optima may exist only if some row in the solution report has zeroes in both the second and third columns of the report, a configuration that some applied statisticians call “snake eyes”. That is, alternate optima may exist only if some variable has both zero value and zero reduced cost, or some constraint has both zero slack and zero dual price. Mathematicians, with no intent of moral judgment, refer to such solutions as degenerate.

If there are alternate optima, you may find your computer gives a different solution from that in the text. However, you should always get the same objective function value.

There are, in fact, two ways in which multiple optimal solutions can occur. For the example in Figure 1.9, the two optimal solution reports differ only in the values of the so-called primal variables (i.e., our original decision variables A , C) and the slack variables in the constraint. There can also be situations where there are multiple optimal solutions in which only the dual variables differ. Consider this variation of the Enginola problem in which the capacity of the Cosmo line has been reduced to 30.

The formulation is:

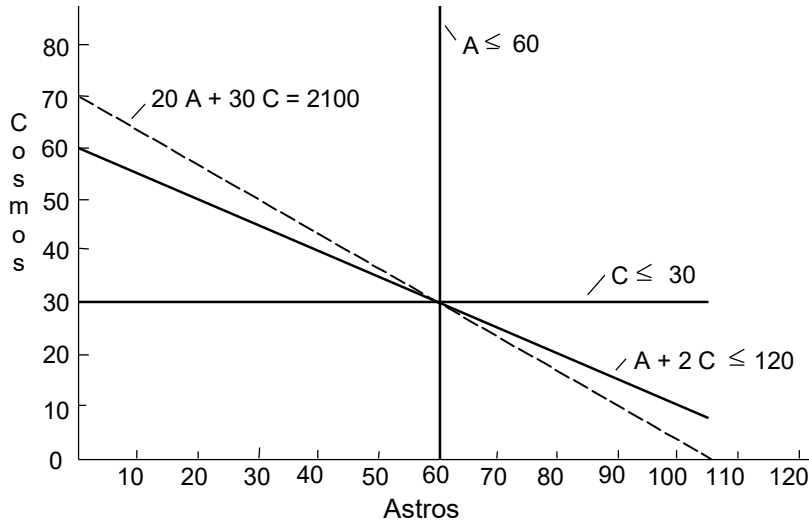
```
MAX = 20 * A + 30 * C;
A < 60;
!note that < and <= are equivalent;
!in LINGO;
C < 30;
A + 2 * C < 120;
```

The corresponding graph of this problem appears in Figure 1.10. An optimal solution is:

Optimal solution found at step:		0
Objective value:		2100.000
Variable	Value	Reduced Cost
A	60.00000	0.000000
C	30.00000	0.000000
Row	Slack or Surplus	Dual Price
1	2100.000	1.000000
2	0.000000	20.00000
3	0.000000	30.00000
4	0.000000	0.000000

Again, notice the “snake eyes” in the solution (i.e., the pair of zeroes in a row of the solution report). This suggests the capacity of the Cosmo line (the RHS of row 3) could be changed without changing the objective value. Figure 1.10 illustrates the situation. Three constraints pass through the point $A = 60$, $C = 30$. Any two of the constraints determine the point. In fact, the constraint $A + 2C \leq 120$ is mathematically redundant (i.e., it could be dropped without changing the feasible region).

Figure 1.10 Alternate Solutions in Dual Variables



If you decrease the RHS of row 3 very slightly, you will get essentially the following solution:

Optimal solution found at step:		0
Objective value:		2100.000
Variable	Value	Reduced Cost
A	60.00000	0.000000
C	30.00000	0.000000
Row	Slack or Surplus	Dual Price
1	2100.000	1.000000
2	0.000000	5.000000
3	0.000000	0.000000
4	0.000000	15.00000

Notice this solution differs from the previous one only in the dual values.

We can now state the following rule: If a solution report has the “snake eyes” feature (i.e., a pair of zeroes in any row of the report), then there may be an alternate optimal solution that differs either in the primal variables, the dual variables, or in both.

If a solution report exhibits the “snake eyes” configuration, a natural question to ask is: can we determine from the solution report alone whether the alternate optima are in the primal variables or the dual variables? The answer is “no”, as the following two related problems illustrate.

Problem D	Problem P
MAX = X + Y;	MAX = X + Y;
X + Y + Z <= 1;	X + Y + Z <= 1;
X + 2 * Y <= 1;	X + 2 * Z <= 1;

Both problems possess multiple optimal solutions. The ones that can be identified by the standard simplex solution methods are:

Solution 1

Problem D			Problem P		
OBJECTIVE VALUE			OBJECTIVE VALUE		
1)	1.00000000		1)	1.00000000	
Variable	Value	Reduced Cost	Variable	Value	Reduced Cost
X	1.0000000	0.0000000	X	1.0000000	0.0000000
Y	0.0000000	0.0000000	Y	0.0000000	0.0000000
Z	0.0000000	1.0000000	Z	0.0000000	1.0000000
Row	Slack or Surplus	Dual Prices	Row	Slack or Surplus	Dual Prices
2)	0.0000000	1.0000000	2)	0.0000000	1.0000000
3)	0.0000000	0.0000000	3)	0.0000000	0.0000000

Solution 2

Problem D			Problem P		
OBJECTIVE VALUE			OBJECTIVE VALUE		
1)	1.00000000		1)	1.00000000	
Variable	Value	Reduced Cost	Variable	Value	Reduced Cost
X	1.0000000	0.0000000	X	0.0000000	0.0000000
Y	0.0000000	1.0000000	Y	1.0000000	0.0000000
Z	0.0000000	0.0000000	Z	0.0000000	1.0000000
Row	Slack or Surplus	Dual Prices	Row	Slack or Surplus	Dual Prices
2)	0.0000000	0.0000000	2)	0.0000000	1.0000000
3)	0.0000000	1.0000000	3)	1.0000000	0.0000000

Notice that:

- *Solution 1* is exactly the same for both problems;
- *Problem D* has multiple optimal solutions in the dual variables (only); while
- *Problem P* has multiple optimal solutions in the primal variables (only).

Thus, one cannot determine from the solution report alone the kind of alternate optima that might exist. You can generate *Solution 1* by setting the RHS of row 3 and the coefficient of *X* in the objective to slightly larger than 1 (e.g., 1.001). Likewise, *Solution 2* is generated by setting the RHS of row 3 and the coefficient of *X* in the objective to slightly less than 1 (e.g., 0.9999).

Some authors refer to a problem that has multiple solutions to the primal variables as *dual degenerate* and a problem with multiple solutions in the dual variables as *primal degenerate*. Other authors say a problem has multiple optima only if there are multiple optimal solutions for the primal variables.


1.8.2 Degeneracy and Redundant Constraints

In small examples, degeneracy usually means there are redundant constraints. In general, however, especially in large problems, degeneracy does not imply there are redundant constraints. The constraint set below and the corresponding Figure 1.11 illustrate:

$$2x - y \leq 1$$

$$\begin{aligned}
 2x - z &\leq 1 \\
 2y - x &\leq 1 \\
 2y - z &\leq 1 \\
 2z - x &\leq 1 \\
 2z - y &\leq 1
 \end{aligned}$$

Figure 1.11 Degeneracy but No Redundancy



These constraints define a cone with apex or point at $x=y=z=1$, having six sides. The point $x=y=z=1$ is degenerate because it has more than three constraints passing through it. Nevertheless, none of the constraints are redundant. Notice the point $x=0.6, y=0, z=0.5$ violates the first constraint, but satisfies all the others. Therefore, the first constraint is nonredundant. By trying all six permutations of 0.6, 0, 0.5, you can verify each of the six constraints are nonredundant.


1.9 Nonlinear Models and Global Optimization

Throughout this text the emphasis is on formulating linear programs. Historically nonlinear models were to be avoided, if possible, for two reasons: a) they take much longer to solve, and b) once “solved” traditional solvers could only guarantee that you had a locally optimal solution. A solution is a local optimum if there is no better solution nearby, although there might be a much better solution some distance away. Traditional nonlinear solvers are like myopic mountain climbers, they can get you to the top of the nearest peak, but they may not see and get you to the highest peak in the mountain range. Versions of LINGO from LINGO 8 onward have a global solver option. If you check the global

solver option, then you are guaranteed to get a global optimum, if you let the solver run long enough. To illustrate, suppose our problem is:

$$\begin{aligned} \text{Min} &= \sin(x) + .5*\text{abs}(x-9.5); \\ x &\leq 12; \end{aligned}$$

The graph of the function appears in Figure 1.12.



If you apply a traditional nonlinear solver to this model you might get one of three solutions: either $x = 0$, or $x = 5.235987$, or $x = 10.47197$. If you check the Global solver option in LINGO, it will report the solution $x = 10.47197$ and label it as a global optimum. Be forewarned that the global solver does not eliminate drawback (a), namely, nonlinear models may take a long time to solve to guaranteed optimality. Nevertheless, the global solver may give a very good, even optimal, solution very quickly but then take a long time to prove that there is no other better solution.

1.10 Problems

1. Your firm produces two products, Thyristors (T) and Lozenges (L), that compete for the scarce resources of your distribution system. For the next planning period, your distribution system has available 6,000 person-hours. Proper distribution of each T requires 3 hours and each L requires 2 hours. The profit contributions per unit are 40 and 30 for T and L , respectively. Product line considerations dictate that at least 1 T must be sold for each 2 L 's.
 - (a) Draw the feasible region and draw the profit line that passes through the optimum point.
 - (b) By simple common sense arguments, what is the optimal solution?

2. Graph the following LP problem:

$$\begin{aligned} &\text{Minimize } 4X + 6Y \\ &\text{subject to } 5X + 2Y \geq 12 \\ &\quad 3X + 7Y \geq 13 \\ &\quad X \geq 0, Y \geq 0. \end{aligned}$$

In addition, plot the line $4X + 6Y = 18$ and indicate the optimum point.

3. The Volkswagen Company produces two products, the Bug and the SuperBug, which share production facilities. Raw materials costs are \$600 per car for the Bug and \$750 per car for the SuperBug. The Bug requires 4 hours in the foundry/forge area per car; whereas, the SuperBug, because it uses newer more advanced dies, requires only 2 hours in the foundry/forge. The Bug requires 2 hours per car in the assembly plant; whereas, the SuperBug, because it is a more complicated car, requires 3 hours per car in the assembly plant. The available daily capacities in the two areas are 160 hours in the foundry/forge and 180 hours in the assembly plant. Note, if there are multiple machines, the total hours available per day may be greater than 24. The selling price of the Bug at the factory door is \$4800. It is \$5250 for the SuperBug. It is safe to assume whatever number of cars are produced by this factory can be sold.
 - (a) Write the linear program formulation of this problem.
 - (b) The above description implies the capacities of the two departments (foundry/forge and assembly) are sunk costs. Reformulate the LP under the conditions that each hour of foundry/forge time cost \$90; whereas, each hour of assembly time cost \$60. The capacities remain as before. Unused capacity has no charge.
4. The Keyesport Quarry has two different pits from which it obtains rock. The rock is run through a crusher to produce two products: concrete grade stone and road surface chat. Each ton of rock from the South pit converts into 0.75 tons of stone and 0.25 tons of chat when crushed. Rock from the North pit is of different quality. When it is crushed, it produces a "50-50" split of stone and chat. The Quarry has contracts for 60 tons of stone and 40 tons of chat this planning period. The cost per ton of extracting and crushing rock from the South pit is 1.6 times as costly as from the North pit.
 - (a) What are the decision variables in the problem?
 - (b) There are two constraints for this problem. State them in words.
 - (c) Graph the feasible region for this problem.
 - (d) Draw an appropriate objective function line on the graph and indicate graphically and numerically the optimal solution.
 - (e) Suppose all the information given in the problem description is accurate. What additional information might you wish to know before having confidence in this model?

5. A problem faced by railroads is of assembling engine sets for particular trains. There are three important characteristics associated with each engine type, namely, operating cost per hour, horsepower, and tractive power. Associated with each train (e.g., the Super Chief run from Chicago to Los Angeles) is a required horsepower and a required tractive power. The horsepower required depends largely upon the speed required by the run; whereas, the tractive power required depends largely upon the weight of the train and the steepness of the grades encountered on the run. For a particular train, the problem is to find that combination of engines that satisfies the horsepower and tractive power requirements at lowest cost.

In particular, consider the Cimarron Special, the train that runs from Omaha to Santa Fe. This train requires 12,000 horsepower and 50,000 tractive power units. Two engine types, the GM-I and the GM-II, are available for pulling this train. The GM-I has 2,000 horsepower, 10,000 tractive power units, and its variable operating costs are \$150 per hour. The GM-II has 3,000 horsepower, 10,000 tractive power units, and its variable operating costs are \$180 per hour. The engine set may be mixed (e.g., use two GM-I's and three GM-II's).

Write the linear program formulation of this problem.

6. Graph the constraint lines and the objective function line passing through the optimum point and indicate the feasible region for the Enginola problem when:
- (a) All parameters are as given except labor supply is 70 rather than 120.
 - (b) All parameters are as given originally except the variable profit contribution of a Cosmo is \$40 instead of \$30.
7. Consider the problem:

$$\begin{aligned}
 &\text{Minimize} && 4x_1 + 3x_2 \\
 &\text{Subject to} && 2x_1 + x_2 \geq 10 \\
 &&& -3x_1 + 2x_2 \leq 6 \\
 &&& x_1 + x_2 \geq 6 \\
 &&& x_1 \geq 0, x_2 \geq 0
 \end{aligned}$$

Solve the problem graphically.

8. The surgical unit of a small hospital is becoming more concerned about finances. The hospital cannot control or set many of the important factors that determine its financial health. For example, the length of stay in the hospital for a given type of surgery is determined in large part by government regulation. The amount that can be charged for a given type of surgical procedure is controlled largely by the combination of the market and government regulation. Most of the hospital's surgical procedures are elective, so the hospital has considerable control over which patients and associated procedures are attracted and admitted to the hospital. The surgical unit has effectively two scarce resources, the hospital beds available to it (70 in a typical week), and the surgical suite hours available (165 hours in a typical week). Patients admitted to this surgical unit can be classified into the following three categories:

Patient Type	Days of Stay	Surgical Suite Hours Needed	Financial Contribution
A	3	2	\$240
B	5	1.5	\$225
C	6	3	\$425

For example, each type B patient admitted will use (i) 5 days of the $7 \times 70 = 490$ bed-days available each week, and (ii) 1.5 hours of the 165 surgical suite hours available each week. One doctor has argued that the surgical unit should try to admit more type A patients. Her argument is that, “in terms of \$/days of stay, type A is clearly the best, while in terms of \$/(surgical suite hour), it is not much worse than B and C .”

Suppose the surgical unit can in fact control the number of each type of patient admitted each week (i.e., they are decision variables). How many of each type should be admitted each week?

Can you formulate it as an LP?