CHAPTER 12: A BASIC HARVEST SCHEDULING MODEL

Forest management problems are complex because of the diversity of the landscape, the unpredictability of the natural processes that occur in forests, the myriad and often unknown inter-relationships between the different components of the forest, and because of the many and diverse human values and objectives associated with natural resources. Mathematical programming techniques, including linear programming (LP), nonlinear programming, integer programming, and other, more specialized algorithms, can be useful tools for dealing with this complexity. Mathematical programming allows managers to model complex problems involving hundreds – even thousands – of management units, thousands of management options, and a wide range of concerns at multiple scales. These concerns potentially include such varied goals as sustaining yields of products, maintaining desired species compositions and habitats, minimizing management costs, even selecting biodiversity reserve locations.

This chapter discusses how to formulate a relatively simple forest harvest scheduling problem as an LP. This should help you gain an appreciation for how real-world problems can be expressed and solved using mathematical programming techniques. Understanding the modeling approach should help you understand the data required for a modeling exercise like this and also give you some appreciation of the limitations of LP. The problems discussed in this book will typically have between eighteen and fifty variables and nine to thirty constraints. While these problems may seem large to you, they are actually highly simplified. Real-world harvest scheduling problems sometimes have hundreds of thousands of variables and thousands of constraints. Obviously, these real-world problems are not formulated or solved by hand. Computer programs have been written to facilitate the formulation, solution, and interpretation of problems like these. When working with such large problems, it is easy to get lost in the details. Thus, it is extremely important that you have an intuitive understanding of the situation being modeled. Experience with small problems, such as those discussed here, should help you build some of this intuition.

1. Planning Periods and the Planning Horizon

We expect our forests to continue to be productive forever – or at least for a very long time. However, forest management scheduling models can only explicitly recognize some finite period of time. This period is called the *planning horizon*. While we will only explicitly model a finite period, there are ways of ensuring the sustainability of the forest beyond the planning horizon that will be discussed later. The planning horizon is divided up into *planning periods*. As with the forest regulation problems we solved, it is helpful to use planning periods that equal the width of our age classes. We will use ten-year age classes in our example forest. We will therefore also use ten-year planning periods. Ten-year periods are used here to keep the problem size relatively small. The period length could be 5 years, or any other reasonable number – for example, between 1 and 20. Figure 12.1 illustrates a 30-year planning horizon consisting of three

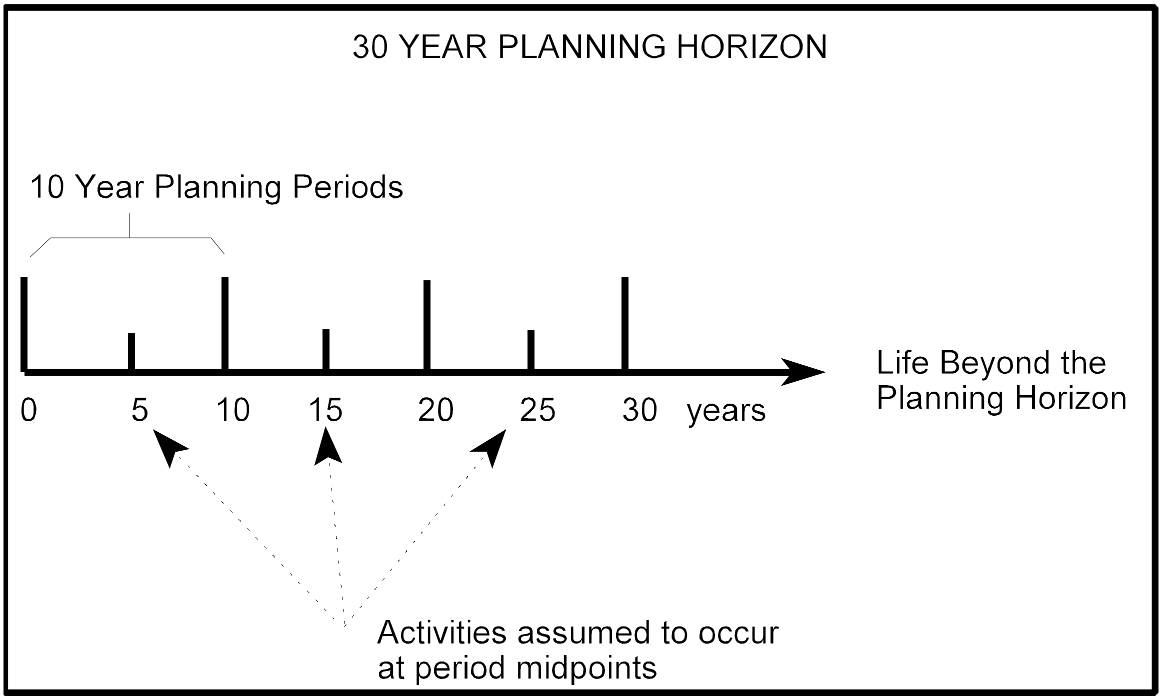

Figure 12.1. Illustration of a 30-year planning horizon divided into three 10-year planning periods.

10-year planning periods. In order to keep the example problems in this chapter small, we will use a 20-year planning horizon, divided into two 10-year planning periods.

All activities that take place during a given planning period will be assumed to take place at the same time. For example, in Figure 12.1 all activities are assumed to occur in year 5 if they are scheduled for period 1, year 15 if they are scheduled for period 2, and year 25 if they are scheduled for period 3. We will assume that activities (in this case, harvests) occur at the period midpoint because this is the average time when actual activities will take place during that planning period – if the plan were actually implemented on a real forest, the activities would generally be spread out over the planning period.

It might seem desirable to use shorter planning periods – for example, 5 years – as this would give more specific information about when management activities are supposed to occur. However, using 5-year planning periods instead of 10-year planning periods would require twice as many planning periods to cover the same time horizon. It would also require the forest age-class distribution to be subdivided into twice as many age classes. This would make the model at least four times as large with 5-year planning periods instead of 10-year planning periods. So, there is a trade-off between having shorter time periods and having a large model. It might also seem desirable to have as long a planning horizon as possible. However, longer planning horizons will also result in larger, more complex models. Thus, model size considerations will have to be balanced against the desire to use very long planning horizons and very short planning periods when selecting the length of the planning periods and the planning horizon.

So, how long should the planning periods and the planning horizon be? A good rule of thumb with planning periods is that they should be between one fifth and one tenth of a rotation. If planning periods are too long, it you will the precision of the yield estimates will suffer because areas that are quite different in age will be lumped together and there can be a substantial difference in the yield if an area is harvested at the beginning of the planning period rather than the end. On the other hand, our information on the age of a stand and our yield estimates may not be that precise and using planning periods that are too short may assume that our data and yield estimates are more precise than they really are. Using somewhat longer planning periods has the added advantage in that they provide some flexibility about when actions really have to occur on the ground, which may make implementation of the plan easier.

With the planning horizon, the most important consideration is to have a long enough planning horizon to make it feasible to achieve the target forest condition. This will typically take at least one rotation, and possibly longer. Thus, the planning horizon should generally be at least one rotation in length and possibly up to two rotations. If the planning horizon is long enough to allow the target forest condition to be met, is there any benefit to using a longer planning horizon? Generally not. A second test for whether the planning horizon is long enough is whether increasing the length of the horizon by one period significantly changes the activities that are planned for the first planning period. The first period is, after all, the only period for which the plan is likely to be implemented. It is likely that the basic information with which the model was developed will have changed before one period is over, and the planning process will have to be repeated before the second period is implemented. Thus, if extending the planning horizon further does not change the activities scheduled for the first period, it is probably long enough. On the other hand, if adding another period does change the activities scheduled for the first period, then the planning horizon probably is not long enough since the length of the planning horizon is a relatively arbitrary modeling decision that should not affect what is done on the ground. You may wonder why one should even bother modeling more than one period if we are only going to use the results from the first period. The answer is that, while we probably will only implement the first period, we want to know what the impacts of our decisions in the first period will be on the state of the forest in future periods an on the options that are available in future periods. This is a fundamental aspect of forest management that makes being a forester so interesting and rewarding – that we must consider the impact of the actions we take today on the resources and management options available in the future. The example problem in this chapter will violate these rules, as the purpose here is to keep the model as simple as possible. As may also be the case in some real-world planning exercises, model size considerations may rule the day in determining the length of the planning horizon.

2. The Example Forest

This chapter describes the formulation and solution to a harvest scheduling LP for the example forest described in Tables 12.1 to 12.3. In general, three types of information are needed: 1) information about the initial age-class distribution, 2) yield tables, and 3) economic information such as prices, costs and a discount rate.

Initial Age-Class Distribution

Table 12.1 shows the initial age-class distribution of the example forest. It includes a total of 37,000 acres. It has been divided into three 10-year age classes and two site classes. This results in six basic categories of forest land. It is common in the forest management literature to refer to these basic categories of land as *analysis areas.* As you will see, the number of analysis areas in the forest directly affects the size of the LP problem. For now, it is best to use an example with only a few analysis areas in order to keep the problem as simple as possible.

Table 12.1. Initial acreage by site and age class for the example forest.

Age Classes	Acres by site class	
	Site I	Site II
0 to 10	3,000	8,000
11 to 20	6,000	4,000
21 to 30	9,000	7,000
Total	18,000	19,000

Yield Data

Table 12.2 presents yield data for the example forest. Note that the yields are reported for ages that are even multiples of the period length. These areas also correspond to the upper bound on each age class. These are the ages at which we will assume that areas are cut. For example, consider acres that are initially in the 0 - 10 year age class. At time zero, the average age of these areas is 5 years. If they are scheduled for harvest in period 1, they are assumed to be harvested at the midpoint of the period – i.e., in year 5 (see Fig. 1). By this time, these acres will, on average, be 10 years old. Similarly, if an area that starts out in age class 3 (average initial age of 25) is harvested in period 2 (at year 15), the area

Table 12.2. Expected yield by site and age class.

Harvest Age	Cords per acre by site class	
	Site I	Site II
10	2	5
20	10	14
30	20	27
40	31	38
50	37	47
60	42	54
70	46	60

will be 40 years old at harvest. Note also that even though the oldest acres currently in the forest are only 30 years old, yield data for older stands will be needed in order to model the yields of stands that might be harvested in later periods.

We will assume that only one product is produced on the forest. Again, this is a simplifying assumption. More products can be recognized in an LP – in fact, this is one of the strengths of LP relative to techniques such as area or volume control. However, the idea here is to keep the initial problem as small and simple as possible. You can see from the yields that site class II is the better of the two site classes. Volumes are measured in cords and areas in acres. These units have been chosen because they are commonly used. Cubic feet and acres or cubic meters and hectares could have been used instead.

Economic Data

Table 12.3 gives the economic data for the example. These data are used for calculating the costs and returns for each management alternative. Wood price refers to the stumpage price per unit volume (cords in our example). We will assume that when an area is harvested it will be regenerated at a cost of $100 per acre, which includes site preparation, seedling, and planting costs. Timber sale preparation costs are assumed to depend on both the area and the volume sold. Although it is not precisely correct to do so, the per-acre cost will be called a fixed cost and the per cord cost a variable cost. The real interest rate (above the rate of inflation) is assumed to be 4%.

Table 12.3. Basic Economic Data for the Example Problem.

Item	Symbol	Amount
Wood Price	P	$25.00/cd
Planting Cost	E	$100.00/ac
Timber Sales Cost		
- per acre	s_f	$15.00/ac
- per cord	s_v	$0.20/cd
Interest Rate	r	4%

Land Expectation Value (LEV) Analysis

You have already learned to use several analytical tools that can give you a lot of information about how to manage this forest. One of the most important is the land expectation value (LEV). The LEV is the net present value of an infinite series of identical future rotations. The optimal economic rotation for a single stand, without recognition of forest-wide considerations such as sustained yield, is that rotation age that results in the greatest LEV. The LEVs for this example for each site class and each rotation age from 20 to 60 years are shown in Table 12.4. We can use the LEVs here to determine the optimal economic rotation for each site class. The rotations with the highest LEVs are 40 years for site class I and 30 years for site class II. You should verify some of the numbers in Table 12.4 for yourself.

Table 12.4. Land Expectation Values by Site Class for Rotation Ages 20 to 60.

Rotation Age	Land Expectation Value	
	Site I	Site II
20	$11.66	$94.94
30	$69.83	**$147.21**
40	**$72.01**	$117.68
50	$31.43	$72.04
60	-$2.66	$28.60
R^* (yr)	40	30

The general formula for the LEVs in Table 12.4 is:

$$LEV_{R,s} = \frac{(P - s_v)Y_{R,s} - s_f - E(1+r)^R}{(1+r)^R - 1}$$

where: $LEV_{R,s}$ = the LEV for site class s at rotation age R,
$Y_{R,s}$ = the yield for site class s at rotation age R,
P = the wood price,
s_v = the variable (per cord) timber sale cost,
s_f = the fixed (per acre) timber sale cost,
E = the stand establishment (regeneration) cost, and
r = the real interest rate.

Long-Term Sustained Yield (LTSY)

Another key management parameter for this forest is the long-term sustained yield, or LTSY. This, you may recall, is the annual harvest volume that the forest would produce if it was regulated. The formula for the LTSY of a forest is:

$$LTSY = \sum_{s=1}^{S} [MAI_{R_s^*} \times A_s]$$

where MAI_{R^*s} = the mean annual increment for site class s at the optimal economic rotation age, R^*, and
A_s = the area in site class s.

For the example forest, the LTSY is:

$$\text{LTSY} = 0.775 \ \tfrac{cords}{ac \cdot yr} \times 18{,}000 \ acres + 0.9 \ \tfrac{cords}{ac \cdot yr} \times 19{,}000 \ acres = 31{,}050 \ cords/yr$$

If the example forest was regulated, each year 450 acres (18,000/40) would be harvested from site class I and 633.3 acres (19,000/30) would be harvested from site class II, producing a total of 31,050 cords per year.

3. Formulating the Example Problem as a Cost Minimization Linear Program

Harvest scheduling models can be formulated with a variety of objective functions. Two common types of objective functions are cost minimization and profit maximization. We will begin here with the cost-minimization formulation because it is somewhat simpler. Chapter 13 discusses formulating the model as a profit maximization problem.

To formulate the problem as an LP, we will follow the three basic formulation steps:

1. define the variables,
2. formulate the objective function, and
3. formulate the constraints.

We will discuss four types of constraints:

1. area constraints,
2. harvest target constraints,
3. ending age constraints, and
4. non-negativity constraints.

Variable Definitions

The first step in formulating the example problem as an LP is to specify the variables. We have six analysis areas, and we will consider three possible prescriptions that could be applied on each analysis area. They are: 1) harvest in period one, 2) harvest in period two, and 3) do not harvest during the planning horizon of the problem. The problem can be viewed as determining the number of acres from each analysis area to assign to each of these prescriptions. Thus, the variables will be defined as follows:

X_{sap} = the number of acres cut from site class s (where s = 1 or 2 in this example) and initial age class a (where a = 1, 2, or 3) in period p, (where p = 0, 1, or 2 and where p = 0 means no harvest during the planning horizon).

For example, X_{231} is the number of acres from site class 2 (II), initial age class 3 assigned to be harvested in period 1. The total number of decision variables is equal to the number of site classes times the number of initial age classes times the number of possible prescriptions (the number of periods plus the do-not-cut option). For the current example, the LP formulation contains 18 decision variables (2 site classes × 3 initial age classes × 3 prescriptions). The complete list of variables is: X_{110}, X_{111}, X_{112}, X_{120}, X_{121}, X_{122}, X_{130}, X_{131}, X_{132}, X_{210}, X_{211}, X_{212}, X_{220}, X_{221}, X_{222}, X_{230}, X_{231}, and X_{232}.

The Cost-Minimization Objective Function

The objective in this example is to minimize the present value of the cost of meeting certain pre-determined harvest targets during the planning horizon. (The harvest targets will be discussed in a later section.) The general form of the objective function will be:

$$Min\ \ Z = \sum_{s=1}^{2} \sum_{a=1}^{3} \sum_{p=0}^{2} c_{sap} \cdot X_{sap}$$

where c_{sap} = the present value of the cost of assigning one acre to the variable X_{sap}.

The triple summation sign in this formula may be intimidating at first. The notation means that you start out with s=1, a=1 and p=0; you then increment p by one until p=2; then you reset p to 0 and increment a by 1; this continues on until s=2, a=3, and p=2. Using this logic, the objective function above can also be written as follows (both are equivalent):

$$\begin{aligned} Min\ Z = \ & c_{110} \cdot X_{110} + c_{111} \cdot X_{111} + c_{112} \cdot X_{112} + c_{120} \cdot X_{120} + c_{121} \cdot X_{121} + c_{122} \cdot X_{122} + \\ & c_{130} \cdot X_{130} + c_{131} \cdot X_{131} + c_{132} \cdot X_{132} + c_{210} \cdot X_{210} + c_{211} \cdot X_{211} + c_{212} \cdot X_{212} + \\ & c_{220} \cdot X_{220} + c_{221} \cdot X_{221} + c_{222} \cdot X_{222} + c_{230} \cdot X_{230} + c_{231} \cdot X_{231} + c_{232} \cdot X_{232} \end{aligned}$$

The problem now is to determine the values of the coefficients (c_{sap}). First, consider the units of the objective function value, Z. The objective is to minimize the discounted cost of meeting certain harvest targets. Thus, the units of Z are discounted dollars of cost. The units of the variables are acres. Therefore, the units of the coefficients must be discounted dollars of cost per acre.

Costs are incurred only when a stand is harvested. Therefore, the objective function coefficients corresponding to variables that do not involve a harvest are zero. These would include any points where p = 0, or all the coefficients of the form c_{sa0}. All harvested acres are assumed to be reforested by planting. Thus, when a harvest occurs, both timber sale and reforestation costs are incurred. The reforestation cost is a per-acre cost, and the timber sale cost has a per-acre component and a component that depends on the harvest volume (see Table 12.3). Also, since the cost coefficients are discounted costs, the costs have to be discounted by an appropriate number of years.

Consider the following example. The cost coefficient for acres from site class 2, initial age class 3, assigned to be harvested in period 1 (i.e., acres assigned to the variable X_{231}) will be:

$$c_{231} = \frac{E + s_f + s_v v_{231}}{(1+r)^5} = \frac{\$100/ac + \$15/ac + \$0.20\,/\,cd \times 27cd/ac}{(1.04)^5} = \$98.96/ac$$

where E = the stand establishment (regeneration) cost,
s_f = the per-acre timber sale cost,
s_v = the per-cord timber sale cost,
v_{231} = the harvest volume for each acre assigned to the variable X_{231}, and
r = the real interest rate.

The hardest part of calculating the cost coefficients for the objective function is determining the harvest volume, v_{sap}, or, in this case, v_{231}. The key to determining the harvest volume is determining the age at harvest. Acres assigned to the variable X_{231} start out at time zero in the third age class, ages 21 to 30. Thus, on average, these acres will be 25 years old at time zero. Furthermore, the acres assigned to X_{231} are scheduled to be harvested in the first period, and this harvest is assumed to happen at the midpoint of the period – year 5. Thus, these acres were 25 years old at time zero, and at time 5 when they are harvested they will be five years older, or 30 years old. Thus, these acres will be 30 years old at harvest. The harvest volume is also determined by the site class, which in this case happens to be site class II. Table 12.2 indicates that acres harvested from site class II stands that are 30 years old should yield 27 cords per acre.

Determining the amount harvested per acre for each variable can be tricky at first. Remember that the first step is to determine the age the acres will be when they are harvested. Ask yourself how old the acres were at the beginning of the time horizon and how many more years will pass before they are cut. The sum of these two values gives you the age at harvest. Once you know the age at harvest and the site class, you can look up the harvest volume in Table 12.2. Table 12.5 shows the values of the harvest volume coefficients corresponding to each variable in this problem. (Note: a simple rule for determining the age at harvest is to add the *a* subscript to the *p* subscript, subtract one, and multiply the result by the period length. Once you convince yourself that this works this is a quick way to get the age at harvest. But please don't substitute a simple rule for having an understanding of what the variables represent and how the acres assigned to a variable will be managed.)

Table 12.5. Harvest ages and volumes for each initial age class and harvest period.

Initial Age Class	Harvest Age		Volume — Site I		Volume — Site II	
	Harvest period					
	Period 1	Period 2	Period 1	Period 2	Period 1	Period 2
0 to 10	10	20	$v_{111} = 2$	$v_{112} = 10$	$v_{211} = 5$	$v_{212} = 14$
11 to 20	20	30	$v_{121} = 10$	$v_{122} = 20$	$v_{221} = 14$	$v_{222} = 27$
21 to 30	30	40	$v_{121} = 20$	$v_{132} = 31$	$v_{231} = 27$	$v_{232} = 38$

The general equation for the coefficients of the objective function is:

$$c_{sap} = \frac{E + s_f + s_v v_{sap}}{(1+r)^{10 \cdot p - 5}} \qquad for\ p > 0;\ for\ p = 0,\ c_{sap} = 0$$

where all of the symbols are as previously defined.

Note that the expression 10·p - 5 will equal 5 when the period is 1; it will equal 15 if the period is 2; it will equal 25 if the period is 3, and so on. This is just a general expression for the midpoint of the period p. The specific objective function for the example problem is:

$$\text{Min Z} = 94.85\ X_{111} + 64.97\ X_{112} + 96.17\ X_{121} + 66.08\ X_{122} + 97.81\ X_{131} + 67.30\ X_{132} + 95.34\ X_{211} + 65.41\ X_{212} + 96.82\ X_{221} + 66.85\ X_{222} + 98.96\ X_{231} + 68.08\ X_{232}$$

You should be able to follow the above example to calculate any of the coefficients in this equation.

Area Constraints

The first set of constraints we will discuss consists of the area constraints. It seems fairly obvious that you can't manage more acres than you have. But even the obvious has to be explicitly included in the formulation. Area constraints – one for each analysis area – specify this restriction for the LP. One constraint is needed for each analysis area stating that the sum of the areas allocated from the analysis area to each potential prescription must be no more than the total area available from the analysis area. For example, there are 3,000 acres in site class I, initial age class 1. The three potential prescriptions for acres in each analysis area are: cut in period 1, cut in period 2, and do not cut the acres during the planning horizon. An area constraint must require that the total acres from site class I, initial age class 1 that are assigned to these three prescriptions must be less than 3,000. This constraint can be written as:

$$X_{110} + X_{111} + X_{112} \le 3{,}000$$

One of these constraints has to be written for each analysis area. The rest of the area constraints for this problem are:

$$X_{120} + X_{121} + X_{122} \le 6{,}000$$
$$X_{130} + X_{131} + X_{132} \le 9{,}000$$
$$X_{210} + X_{211} + X_{212} \le 8{,}000$$
$$X_{220} + X_{221} + X_{222} \le 4{,}000$$
$$X_{230} + X_{231} + X_{232} \le 7{,}000$$

Note that these equations could be expressed as equalities, rather than inequalities. However, for technical reasons, it works better if they are expressed as inequalities. It is generally best to avoid using equality constraints in formulating LP problems to give the solution algorithm more flexibility. If any of these constraints are not binding, any slack should be assigned in the interpretation stageto the no-cut prescription. The general formula for these constraints is:

$$\sum_{p=0}^{2} X_{sap} \le A_{sa} \quad s = 1,2 \quad a = 1,2,3$$

where A_{sa} = the total number of acres in site class s, initial age class a.

Harvest Target Constraints

The harvest target constraints require the production of some minimum output of timber in each period. Without these constraints, the cost-minimizing solution would be to harvest nothing. An obvious question that arises when formulating these constraints is: "what target should I use?" The volume control formulas presented in Chapter 10 can provide a good starting point. For example, Hundeshagen's formula could be used to obtain a harvest target for the first decade. If the example forest was regulated using the optimal economic rotation for each site class, its inventory would be 419,583 cords. The inventory of the current forest is 375,500. (You should be able to verify these numbers.) Earlier, we calculated the LTSY for this forest, and found that it was 31,050 cds/yr. Thus, by Hundeshagen's formula, the harvest in the first decade should be:

$$H_1 = \frac{375{,}500\,ac}{419{,}583\,ac} \times 31{,}050\,cd\ /\ yr = 27{,}788\,cd\ /\ yr$$

Thus, as a rough estimate, it should be possible to harvest about 28,000 cds/yr, or 280,000 cds in the first period. For the second period, we will set our target approximately at the average of this harvest target and the LTSY – 29,500 cd/yr, or a total of 295,000 cds in the second period. This is not the only way to go about setting harvest targets, but it provides a good starting point.

A harvest target constraint must be formulated for each period. Let's start with the harvest target constraint for period 1. Clearly, unless a variable involves cutting in period 1, it will not appear in this constraint. Thus, the variables with $p = 0$ or 2 will not appear in the harvest constraint for period 1 (in other words, their coefficients in this constraint will be zero); only those variables with $p = 1$ will have non-zero coefficients in this constraint. The right-hand-side of this constraint will be the harvest target, in cords. The units of the variables are acres. Thus, the units of the coefficients in this constraint should be cords per acre. Thus, the coefficients in this constraint should indicate the volume of wood that will be produced in period 1 for each acre assigned to that variable. You have already seen these values – they are the v_{sap} parameters that were calculated earlier (see Table 12.5). Thus, the harvest target constraint for period 1 can be written, generally, as:

$$\sum_{s=1}^{2}\sum_{a=1}^{3} v_{sa1} \cdot X_{sa1} \ge H_1$$

where H_1 = the harvest target for decade 1 (in cords), and, again,
v_{sa1} = the harvest volume for each acre assigned to the variable X_{sa1}.

You may want to review how the values in Table 12.5 were determined. Recall that the first step is to determine the age of the stand at harvest. This is determined by identifying the age of the stand at the beginning of the planning horizon and adding the number of years before the stand is to be harvested. For harvests that occur in period 1, add five years to the initial age of the stand. Once the age at harvest is determined and the site class noted, the yield per acre can be identified from the yield table (Table 12.2).

The harvest target constraint for period 1 with the specific values of the coefficients is:

$$2\ X_{111} + 10\ X_{121} + 20\ X_{131} + 5\ X_{211} + 14\ X_{221} + 27\ X_{231} >= 280{,}000$$

The specific harvest target constraint for period 2 is:

$$10\ X_{112} + 20\ X_{122} + 31\ X_{132} + 14\ X_{212} + 27\ X_{222} + 38\ X_{232} >= 295{,}000$$

The general form of the harvest target constraints is:

$$\sum_{s=1}^{2}\sum_{a=1}^{3} v_{sap} \cdot X_{sap} \geq H_p \quad p = 1, 2$$

where H_p = the harvest target for decade p.

Ending Age Constraints

LP models can produce some very undesirable effects if they are not formulated carefully. In a sense, in the case of profit maximization models, they are like greedy monsters, and in the case of cost minimization models, they are very lazy. Profit maximization models are, in fact, designed to be greedy because they are designed to earn as much profit as they can. As a result, they will tend to do so without consideration of other objectives unless they are forced to consider those objectives through either modification of the objective function or by formulating constraints that address those objectives. If we do not restrain these greedy beasts appropriately (or train them differently by modifying their objective functions), they may find very undesirable solutions. The greedy nature of LP harvest scheduling models becomes apparent near the end of the planning horizon. On the other hand, cost minimization models are lazy. Because almost any kind of activity involves cost, they won't do anything they are not told to. In either case, the LP monster does not "see" what happens after the end of the planning horizon, so it does not care at all about what happens then. The cost-minimizing model will not achieve any goal that it is not forced to achieve and, if no constraints are included in the model formulation to prevent it, at the end of the planning horizon the profit-maximizing model may schedule for harvest any part of the forest where any immediate profit can be made. Obviously, this is will not lead to a desirable management plan. The problem, therefore, is to find ways to ensure that the models create they type of forest that we desire to have at the end of the planning horizon.

One approach to addressing this problem would be to require the LP to leave a specific age-class distribution at the end of the planning horizon. This could, in fact, be done, but it would be difficult to say in advance what that age-class distribution should be. Even if you have a particular age-class distribution in mind, this approach tends to be too restrictive, not allowing the LP much latitude to achieve any other goals. Alternatively, the model can be required to leave a forest with a minimum amount inventory volume at the end of the planning horizon, or it can be required to ensure that the average age of the forest should be at least a certain target age at the end of the planning horizon. Both of these are workable approaches. We will use the latter approach here only because it is simpler.

Before going into the specifics of formulating this type of constraint, consider how you would calculate the average age of a forest with a simple age-class distribution. For example, consider the age-class distribution in Table 12.6. Assume that the average age of the acres in the 0 to 10 yr age class is 5 years old, the average age of acres in the 11 to 20 yr age class is 15 years old, and the average age of acres in the 21 to 30 yr age class is 25 years old. What is the average age of this forest?

Table 12.6. A simple age-class distribution.

Age Class	Avg. Age	Acres
0 to10 yr	5 yr	300
11 to20 yr	15 yr	100
21 to30 yr	25 yr	250

You should be able to tell by looking at the age-class distribution that the average age of the forest will be a little less than 15 years old – there are somewhat more acres in the 0 to 10 year age class than there are in the 21 to 30 year age class. But how does one calculate the exact average age of the forest? The answer is with a weighted average, using the acres in each age class as the "weights." The weighted average age of the forest in Table 12.6 is:

$$\overline{Age} = \frac{300\times 5+100\times 15+250\times 25}{300+100+250} = \frac{300}{650}5+\frac{100}{650}15+\frac{250}{650}25 = 14.23\,years$$

The general formula for the weighted average age of a forest is:

$$\overline{Age} = \sum_{i=1}^{n} \frac{Area_i}{\sum_{j=1}^{n} Area_j} Age_i$$

where $\overline{Age}$ = the average age of the forest,
$Area_i$ = the area in the i^{th} unit of the forest, and
Age_i = the age of the i^{th} unit of the forest.

How can this formula be used to formulate a constraint for the average age of the forest at the end of the planning horizon? Well, the variables in the problem formulation – the X_{sap}'s – each represent the acres that will be in different blocks of the forest at the end of the planning horizon. Thus, the term $Area_i$ in the above formula can be replaced with the variables X_{sap}. Furthermore, the term in the denominator is just the total area of the forest.

Now, all we need is the age that acres assigned to each variable will be at the end of the planning horizon. Let:

Age_{sap}^{20} = the age in year 20 of acres in site class *s*, initial age class *a*, that are scheduled to be harvested in period *p* (where *p*=0 implies no harvest during the planning horizon) – in other words, acres assigned to the variable X_{sap}.

The following formula gives the average age of the forest at the end of the planning horizon:

$$\overline{Age}^{20} = \frac{\sum_{s=1}^{2}\sum_{a=1}^{3}\sum_{p=0}^{2} Age_{sap}^{20} \times X_{sap}}{\sum_{s=1}^{2}\sum_{a=1}^{3}\sum_{p=0}^{2} X_{sap}} = \frac{\sum_{s=1}^{2}\sum_{a=1}^{3}\sum_{p=0}^{2} Age_{sap}^{20} \times X_{sap}}{TotalArea}$$

where $\overline{Age}^{20}$ = the average age of the forest in year 20,
$TotalArea$ = the total area of the forest, and
all of the other symbols are as previously defined.

This formula can be rearranged by multiplying both sides by the total area of the forest. We also need to convert it to a greater-than-or-equal constraint. This gives us the following general form of the ending average age constraint for our problem formulation:

$$\sum_{s=1}^{2}\sum_{a=1}^{3}\sum_{p=0}^{2} Age_{sap}^{20} \times X_{sap} \geq \overline{Age}^{20} \times TotalArea$$

where $\overline{Age}^{20}$ = is the target minimum average age of the forest in year 20,
and all of the other symbols are as previously defined.

How the values of the Age_{sap}^{20} parameters are determined? These parameters should equal the age in year 20 (at the end of the planning horizon) of acres assigned to the corresponding variable X_{sap}. First, observe that the age of the acres will not depend on the site class. Next, consider the age of acres assigned to do-not-cut prescriptions ($p = 0$). If an area is not cut, in year 20 it will be 20 years older than it was at the beginning of the planning horizon. Thus, if $p = 0$ and $a = 1$, the acres were 5 years old (on average) at the beginning of the planning horizon, and they will be 25 years old (on average) at the end of the planning horizon (in year 20). If $p = 0$ and $a = 2$, the acres were 15 years old at the beginning of the planning horizon, and they will be 35 years old at the end. Finally, if $p = 0$ and $a = 3$, the acres were 25 years old at the beginning of the planning horizon, and they will be 35 years old at the end.

Next, consider acres assigned to prescriptions involving a harvest in period 1 ($p = 1$). Regardless of their initial age, these acres will be 15 years old in year 20. Since they were harvested in period 1, they were 0 years old (i.e., "born") in year 5. By year 20, they will be 15 years old. Similarly, acres scheduled to be harvested in period 2 ($p = 2$) will be 5 years old in year 20 regardless of their initial age. These acres will be harvested and regenerated in year 15, so they will be 5 years old in year 20. Table 12.7 summarizes the values of the Age_{sap}^{20} parameters.

Table 12.7. Ending ages for each initial age class and harvest period.

Initial Age Class	Site I			Site II		
	Harvest Period					
	Not Cut	Period 1	Period 2	Not Cut	Period 1	Period 2
0 to 10	$Age_{110}^{20} = 25$	$Age_{111}^{20} = 15$	$Age_{112}^{20} = 5$	$Age_{210}^{20} = 25$	$Age_{211}^{20} = 15$	$Age_{212}^{20} = 5$
11 to 20	$Age_{120}^{20} = 35$	$Age_{112}^{20} = 15$	$Age_{122}^{20} = 5$	$Age_{220}^{20} = 35$	$Age_{212}^{20} = 15$	$Age_{222}^{20} = 5$
21 to 30	$Age_{130}^{20} = 45$	$Age_{113}^{20} = 15$	$Age_{132}^{20} = 5$	$Age_{230}^{20} = 45$	$Age_{213}^{20} = 15$	$Age_{232}^{20} = 5$

The final question regarding this constraint that needs to be addressed is how one should identify what the minimum average age of the forest should be at the end of the planning horizon. Again, it is useful here to consider what the forest would be like if it was regulated. What would the average age of the forest be if it was regulated using the optimal economic rotation? Since the optimal rotations of the two site classes are different, we will have to consider each separately.

First, consider site class I. The optimal rotation for site I acres is 40 years. The average age of a forest regulated on a 40 year rotation will be approximately 20 years old – half a rotation. Actually, it will be 20½, which is (40 + 1)/2 or (R + 1)/2. To see this, consider a regulated forest with just 3 acres: one of the acres is 1 year old, another acre is 2 years old, and the third acre is 3 years old. This would be a 3-acre forest, regulated on a rotation age of 3 years. The average age of this forest is obviously 2 years, which is consistent with the formula: (R + 1)/2. If you still aren't convinced, calculate the average age of a regulated forest with 4 acres, regulated on a rotation age of 4.

Now, consider site class II. The optimal rotation age for this site class is 30 years. The average age of the acres in site class II, therefore, would be 15½ years if they were regulated on a 30-year rotation. Thus, if our forest was regulated, it would have 18,000 acres with an average age of 20½ and 19,000 acres with an average age of 15½. Thus, the average age of the regulated forest will be 17.93 years [(20½·18,000 + 15½·19,000)/37,000]. This is a useful guideline, but we may not want to require that the forest fully reach this average age in only 20 years. Thus, for the example, we will require that the forest have an average age of at least 17 years at the end of the planning horizon. Now, we can write out the specific form of the ending age constraint for the example problem:

$$25\ X_{110} + 15\ X_{111} + 5\ X_{112} + 35\ X_{120} + 15\ X_{121} + 5\ X_{122} + 45\ X_{130} + 15\ X_{131} + 5\ X_{132}$$
$$25\ X_{210} + 15\ X_{211} + 5\ X_{212} + 35\ X_{220} + 15\ X_{221} + 5\ X_{222} + 45\ X_{230} + 15\ X_{231} + 5\ X_{232} \geq 629{,}000$$

Note that the right-hand-side coefficient on this constraint is 17×37,000 – the minimum average ending age times the total area of the forest.

Non-negativity Constraints

As always, it is important to include the non-negativity constraints in the problem formulation. These can be written as follows:

$$X_{sap} \geq 0 \qquad s = 1, 2 \quad a = 1, 2, 3 \quad p = 0, 1, 2$$

The Complete Cost-Minimization Problem Formulation

The formulation process is now complete. The complete problem formulation can be written as follows. First, the general form of the cost-minimization problem is:

$$Min\ Z = \sum_{s=1}^{2} \sum_{a=1}^{3} \sum_{p=0}^{2} c_{sap} \cdot X_{sap} \qquad \text{(Objective function)}$$

Subject to:

$$\sum_{p=0}^{2} X_{sap} \leq A_{sa} \quad s = 1,2 \quad a = 1,2,3 \qquad \text{(Area constraints)}$$

$$\sum_{s=1}^{2} \sum_{a=1}^{3} v_{sap} \cdot X_{sap} \geq H_p \quad p = 1,2 \qquad \text{(Harvest constraints)}$$

$$\sum_{s=1}^{2} \sum_{a=1}^{3} \sum_{p=0}^{2} Age_{sap}^{20} \times X_{sap} \geq \overline{Age}^{20} \times TotalArea \qquad \text{(Ending age constraints)}$$

$$X_{sap} \geq 0 \quad s = 1, 2 \quad a = 1, 2, 3 \quad p = 0, 1, 2 \qquad \text{(Non-negativity constraints)}$$

where X_{sap} = the number of acres cut from site class s (where s = 1 or 2) and initial age class a (where a = 1, 2, or 3) in period p, (where p = 0, 1, or 2 and where p = 0 means no harvest during the planning horizon);

c_{sap} = the present value of the cost of assigning one acre to the variable X_{sap};

A_{sa} = the total number of acres in site class s, initial age class a;

v_{sap} = the harvest volume for each acre assigned to the variable X_{sap};

H_p = the harvest target for decade p (in cords);
Age^{20}_{sap} = the age in year 20 of acres assigned to the variable X_{sap};
$\overline{Age}^{20}$ = the target (minimum) average age of the forest in year 20; and
$TotalArea$ = the total area of the forest.

The specific formulation of the LP for the example forest is:

$$\text{Min } Z = 94.85\,X_{111} + 64.97\,X_{112} + 96.17\,X_{121} + 66.08\,X_{122} + 97.81\,X_{131} + 67.30\,X_{132} + 95.34\,X_{211} + 65.41\,X_{212} + 96.82\,X_{221} + 66.85\,X_{222} + 98.96\,X_{231} + 68.08\,X_{232}$$

Subject to:

$$X_{110} + X_{111} + X_{112} \le 3{,}000$$
$$X_{120} + X_{121} + X_{122} \le 6{,}000$$
$$X_{130} + X_{131} + X_{132} \le 9{,}000$$
$$X_{210} + X_{211} + X_{212} \le 8{,}000$$
$$X_{220} + X_{221} + X_{222} \le 4{,}000$$
$$X_{230} + X_{231} + X_{232} \le 7{,}000$$
$$2\,X_{111} + 10\,X_{121} + 20\,X_{131} + 5\,X_{211} + 14\,X_{221} + 27\,X_{231} >= 280{,}000$$
$$10\,X_{112} + 20\,X_{122} + 31\,X_{132} + 14\,X_{212} + 27\,X_{222} + 38\,X_{232} >= 295{,}000$$
$$25\,X_{110} + 15\,X_{111} + 5\,X_{112} + 35\,X_{120} + 15\,X_{121} + 5\,X_{122} + 45\,X_{130} + 15\,X_{131} + 5\,X_{132}$$
$$25\,X_{210} + 15\,X_{211} + 5\,X_{212} + 35\,X_{220} + 15\,X_{221} + 5\,X_{222} + 45\,X_{230} + 15\,X_{231} + 5\,X_{232} \ge 629{,}000$$
$$X_{110} \ge 0;\ X_{111} \ge 0;\ X_{112} \ge 0;\ X_{120} \ge 0;\ X_{121} \ge 0;\ X_{122} \ge 0;\ X_{130} \ge 0;\ X_{131} \ge 0;\ X_{132} \ge 0;$$
$$X_{210} \ge 0;\ X_{211} \ge 0;\ X_{212} \ge 0;\ X_{220} \ge 0;\ X_{221} \ge 0;\ X_{222} \ge 0;\ X_{230} \ge 0;\ X_{231} \ge 0;\ X_{232} \ge 0$$

At this point, the problem has been formulated. The next step is to get the formulation into a format that can be read by a program that will solve the problem. After the solution has been obtained, the next step is to interpret the results.

4. Interpreting the Solution to the Example Problem

LINDO[1] was used to solve the example problem. The output reported by LINDO is shown in Figure 12.2. There are five types of information presented in the LINDO output that you should be able to interpret: 1) the optimal objective function value, 2) the optimal values of the variables, 3) the reduced cost coefficients, 4) the slack or surplus values, and 5) the dual prices.[2]

---

[1] LINDO is a computer application specifically designed for solving LP problems.

[2] Similar information is provided by other programs that can be used to solve LP problems – for example, Excel.

The Optimal Objective Function Value

The objective function value reported by LINDO is $1,906,447. This is the minimum discounted cost of meeting the harvest targets of 28,000 cords per year in the first decade and 29,500 cords per year in the second decade – given the initial forest, the available management options, and the additional constraint requiring the average age of the forest to be at least 17 years at the end of the 20-year planning horizon.

Interpreting the Optimal Variable Values

The optimal variable values are listed in the second column of the first block of the output. These values indicate how many acres from each initial age/site class combination should be assigned to each prescription. The value of 0 assigned to X_{111}, for example, indicates that none of the acres in site class I, currently in age class 1 (ages from 1 to 10) should be harvested in the first period. Similarly, the value of 1,000 assigned to X_{122} indicates that 1,000 acres in site class I, currently in age class 2 (ages from 11 to 20) should be harvested in period 2. The remaining 5,000 acres in this analysis area (site class I, initial age class 2) should be left unharvested for the next 20 years ($X_{120} = 5{,}000$).

Table 12.8 summarizes the harvest schedule by analysis area, indicating the number of acres from each analysis area assigned to each prescription. Table 12.8 is organized from the perspective of the LP variables, but it is not necessarily the most intuitive way to present the results from the perspective of a forest manager — who typically wants to know how many acres of what type will be harvested at a given time. Table 12.9 presents the harvest schedule in a different way — by period and age at harvest. This organizes the harvest information in a more intuitive way, but it requires some additional interpretation on your part. From either table, you can tell that approximately 1,150 acres will be harvested each year. From Table 12.9, it is clearer that in period 1, all acres harvested will be 30 years old at harvest. In period 2, the average age at harvest goes up on site I lands to close to 40 years old. On the other hand, the average age at harvest in period 2 for site II lands is between 20 and 30 years.

Both of the harvest target constraints are binding. The volume harvested in each period is therefore equal to the right-hand side of the harvest target constraint for that period. The gross revenue for each period will equal the volume harvested times the wood price:

Gross revenue for period 1 = 280,000 cords × $25 /cd = $7,000,000
Gross revenue for period 2 = 295,000 cords × $25 /cd = $7,375,000

The costs for each period are from replanting the harvested acres and conducting timber sales. The cost is a function of both the area planted and the volume harvested — $100 per acre for reforestation, $15 per acre for timber sale administration costs, and $.2 per cord for timber sale administration. The costs in each period will be:

```
LP OPTIMUM FOUND AT STEP     14

       OBJECTIVE FUNCTION VALUE

       1)      1906447.

 VARIABLE        VALUE          REDUCED COST
     X111          .000000        203.202000
     X112          .000000        127.710800
     X121          .000000         63.224650
     X122      1000.000000          .000000
     X131      4550.000000          .000000
     X132      4450.000000          .000000
     X211          .000000         78.873120
     X212      2075.000000          .000000
     X221          .000000        120.946800
     X222      4000.000000          .000000
     X231      7000.000000          .000000
     X232          .000000         66.606670
     X110      3000.000000          .000000
     X120      5000.000000          .000000
     X130          .000000        159.638500
     X210      5925.000000          .000000
     X220          .000000        223.493900
     X230          .000000        449.739100


      ROW   SLACK OR SURPLUS     DUAL PRICES
       2)          .000000        478.915600
       3)          .000000        670.481800
       4)          .000000       1021.687000
       5)          .000000        478.915600
       6)          .000000        893.975700
       7)          .000000       1311.787000
       8)          .000000        -41.607330
       9)          .000000        -32.038760
      10)          .000000        -19.156620

NO. ITERATIONS=      14
```

Figure 12.2. LINDO solution to the example cost-minimization harvest scheduling problem.

Costs for period 1 = 11,550 ac × \$115/ac + 280,000 cd × \$.2 /cd = \$1,384,250
Costs for period 2 = 11,525 ac × \$115/ac + 295,000 cd × \$.2 /cd = \$1,384,375

The net revenue for each period is just the gross revenue minus the cost:

Net revenue for period 1 = \$7,000,000 - \$1,384,250 = \$5,615,750
Net revenue for period 2 = \$7,375,000 - \$1,384,375 = \$5,990,625

Table 12.8. Acres assigned to each prescription, by site class and initial age class.

Site C lass	Initial Age Class	Harvest in Pd 1	Harvest in Pd 2	No Harvest
I	0-10 yr	0	0	3,000
	11-20 yr	0	1,000	5,000
	21-30 yr	4,550	4,450	0
II	0-10 yr	0	2,075	5,925
	11-20 yr	0	4,000	0
	21-30 yr	7,000	0	0
Total		11,550	11,525	13,925

Table 12.9. Acres harvested by period, by site class, and by age at harvest.

Planning Period	Age at Harvest	Site I	Site II	Total
1	30	4,550	7,000	11,550
2	20	0	2,075	2,075
	30	1,000	4,000	5,000
	40	4,450	0	4,450
	Total	5,450	6,075	11,525
Total acres harvested		10,000	13,075	23,075
Acres not harvested		8,000	5,925	13,925

Each of these cost and revenue figures is for a 10-year period. They should be divided by 10 to get average annual costs and revenues. Table 12.10 shows the annual costs and revenues for the example problem. You can use the cost data in Table 12.10 to check your calculations. You should verify that the discounted costs are equal to $1,906,447.

Before we move on to the interpretation of the other information given in the LP solution, it is also useful to consider what the age-class distribution of the forest will look like at the end of each planning period. Tables 10 and 11 show the age-class distribution for the forest at the end of periods 1 and 2, respectively. (You should be able to reproduce tables like these yourself.) Note how, within only 20 years, the age-class distribution of the forest is projected to be converted to one with a relatively even balance of age classes between 0 and 40 years for site I and between 0 and 30 years for site II.

Table 12.10. Acres and volume harvested, and revenues and costs by period for the example forest.

Quantity	Period 1	Period 2
Acres harvested	11,550	11,525
Volume harvested (cords)	280,000	295,000
Gross Revenues	$7,000,000	$7,375,000
Costs	$1,384,250	$1,384,375
Net Revenues	$5,615,750	$5,990,625

Table 12.11. Age-class distribution of the example forest after period 1.

Age Classes	Acres by site class	
	Site I	Site II
0 to 10	4,550	7,000
11 to 20	3,000	8,000
21 to 30	6,000	4,000
31 to 40	4,450	0
Total	18,000	19,000

Table 12.12. Age-class distribution of the example forest after period 2.

Age Classes	Acres by site class	
	Site I	Site II
0 to 10	5,450	6,075
11 to 20	4,550	7,000
21 to 30	3,000	5,925
31 to 40	5,000	0
Total	18,000	19,000

Interpreting the Reduced Cost Coefficients

Recall that the reduced cost coefficients indicate how much the objective function coefficient corresponding to each variable whose optimal value is zero would have to be improved in order for the value of that variable to become positive in the optimal solution. In this context, the reduced cost coefficients indicate how much the discounted cost per acre for the corresponding prescription would have to be reduced before that prescription would be optimal to apply on any acres from the corresponding analysis area.

For example, the reduced cost coefficient corresponding to the variable X_{111} is $203.20. This means that the discounted cost of harvesting an acre from site class 1, initial age class 1 in period 1 would have to be reduced by $203.20 before any acres from that analysis area would be harvested in period 1. Since the coefficient on X_{111} is only $94.85, this means that this activity would have to have a negative cost of more than $100 before it would be optimal. Obviously, this prescription is unlikely to be optimal under any circumstances.

The lowest reduced cost coefficient in the example – $63.22 – corresponds to the variable X_{121}. The discounted cost of assigning one acre to this variable – i.e., the discounted cost of harvesting an acre from site class 1, initial age class 2 in period 1 – is $96.17. If this cost could be reduced to $32.94, then it would be optimal to assign some acres from that analysis area to be harvested in period 5.

Of course, whenever the optimal value of a variable is positive, the corresponding reduced cost value is zero.

Interpreting the Slack/Surplus Coefficients and the-Dual Prices

All of the slack or surplus coefficients in this example are zero. This means that all of the constraints are binding. Since each constraint is binding, each has a non-zero dual price. Note that the dual prices for the less-than-or-equal constraints are positive and the dual prices for the greater-than-or-equal constraints are negative. The dual prices indicate how much the objective function would be improved if the right-hand side of the corresponding constraint was increased by 1. Increasing the right-hand side of a less-than-or-equal constraint increases the available amount of a limited resource. Since this will tend to improve the objective function, the dual prices corresponding to these constraints are positive. Increasing the right-hand side of a greater-than-or-equal constraint raises some minimum target that must be achieved. Since this makes the constraint harder to achieve, the objective function value will be worsened — hence these constraints generally have negative dual prices.

Consider the dual prices associated with the area constraints first. These dual prices indicate how much the objective function value — the minimum cost of achieving the harvest targets — will be improved if there was one more acre in the corresponding analysis area. These dual prices, therefore, give the marginal value of an acre in the corresponding analysis area. If the harvest scheduling model accurately represents the concerns of the organization, then these dual prices indicate what the organization should be willing to pay for one more acre in each analysis area. In this example, another acre in site class I between the ages of 0 and 10 would be worth $478.92. Similarly, another acre in site class II between the ages of 21 and 30 would be worth $1,311.79. As you might expect, acres tend to be worth more if they are older and if they are on a better site class. Interestingly, an acre between 0 and 10 years old is worth the same, regardless of its site class. This more likely represents a formulation problem than a real relationship. (*Can you think of what formulation problem might be causing this?*)

Now, consider the dual prices associated with the harvest target constraints. These values indicate how much the objective function value would be "improved" if one more cord had to be produced in that period. Therefore, these values give the discounted marginal cost of producing another cord of wood in each period. Note that if we want the future marginal cost, these values should be compounded forward to the midpoint of the respective period. For example, the marginal discounted cost per cord in period 1 is $41.61. The marginal cost in year 5 (the period midpoint) is $50.62/cd [$41.61\times(1.04)^5$]. Similarly, the discounted marginal cost of producing a cord in period 2 is $32.04. The future value of this cost, in year 15, is $57.70 [$32.04\times(1.04)^{15}$]. Therefore, even though the discounted cost for period 2 is lower than the discounted cost for period 1, the marginal cost of producing a cord of wood is higher in period 2 than in period 1. Note also that the marginal cost of producing wood in both periods is higher than the wood cost given in Table 12.3 ($25/cd). If wood can be purchased on the open market for $25 per cord, then it appears that the harvest targets would be cheaper to achieve by buying at least some of the required wood on the open market and reducing the harvest from the example forest. (The appropriate value to use in this comparison is the future marginal cost because the price of $25/cd is the real future price.)

Finally, consider the dual price associated with the ending average age constraint. This dual price is difficult to interpret because the units of the right-hand-side coefficient of this constraint are years×acres. The units of the dual price are therefore discounted dollars per acre·year. The dual price can be interpreted as present value of the marginal cost of requiring one acre to be one year older at the end of the planning horizon.

5. <u>Study Questions for the Cost-Minimization Harvest Scheduling Model</u>

1. What are the advantages and disadvantages of longer planning horizons in harvest scheduling models?

2. What two basic guidelines should you use for determining what the planning horizon of a harvest scheduling model should be?

3. What is an *analysis area* in a harvest scheduling model?

4. Explain how the variables were defined in the harvest scheduling model presented in this chapter. What does the variable X_{231} represent?

5. How many decision variables would a similar problem have if there were three site classes, four initial age classes, and six possible prescriptions for each analysis area?

6. In the objective function of the example problem, the coefficient on the variable X_{221} is 96.82. What does this coefficient represent? Show how it was calculated.

7. What is the purpose of the area constraints? In general, how many area constraints will there be in a harvest scheduling model like the one presented in this chapter?

8. What is the purpose of the harvest target constraints? In general, how many harvest target constraints will there be in a harvest scheduling model like the one presented in this chapter?

9. In the example problem, the coefficient on the variable X_{232} in the harvest target constraint for period 2 is 38. What does this coefficient represent? Explain why the value of the coefficient is 38.

10. What procedure would you use to set the harvest targets for a forest?

11. What is the purpose of the average ending age constraint? In the example problem, the coefficient on the variable X_{121} in the average ending age constraint 15. What does this coefficient represent? Explain why the value of the coefficient is 15.

12. What is the key difference between Tables 7 and 8? Which would be a better table to give to field foresters charged with managing the forest? Why?

13. In general, how should the reduced cost coefficients be interpreted? What is their interpretation in the harvest scheduling model presented in this chapter?

14. In general, how should the dual price coefficients be interpreted? What is their more specific interpretation for each type of constraint in the harvest scheduling model presented in this chapter? (i.e., for the area constraints; for the harvest target constraints; for the average ending age constraint?)

15. Why are the dual prices negative for greater-than-or-equal constraints and positive for less-than-or-equal constraints?

6. Exercises

1. It is your job to develop a management plan for a 5,800 acre forest. The age class distribution by site class is given in Table 12.13. Table 12.14 gives the expected yield per acre by age and site class.

Table 12.13. Initial forest acreage by site and age class.

Age Classes	Acres by site class	
	Site I	Site II
0-10	700	800
11-20	800	1,300
21-30	1,000	1,200

Table 12.14. Expected yield by site and age.

Age	Cords/acre by site class	
	Site I	Site II
10	0	5
20	9	15
30	16	24
40	21	32
50	25	39

Formulate (Do not solve!) this management problem as a LP. Be sure to clearly define all of your variables. Use the following assumptions:

i) You want to minimize the discounted cost of producing at least 49,000 cords for the first decade and 45,000 cords in the second decade of a 20 year planning horizon using an interest rate of 4%.

ii) It costs $10 per acre plus $1.20 per cord to prepare a timber sale. It costs $50 per acre to regenerate stands that are harvested.

iii) You want the average age of your ending inventory to be at least 13 years old.

2. The following data are for a 5,800 acre forest. The age class distribution by site class is given in Table 12.15. Table 12.16 gives the expected yield per acre by age and site class.

Table 12.15. Initial forest acreage by site and age class.

Age Classes	Acres by site class	
	Site I	Site II
0-10	480	950
11-20	910	1,080
21-30	1,060	1,320

Table 12.16. Expected yield by site and age.

Age	Cords/acre by site class	
	Site I	Site II
10	0	2
20	14	18
30	22	31
40	34	43
50	40	53

a. Formulate, and solve using Excel, a harvest scheduling model for this forest using the following assumptions:

i) The objective is to maximize the discounted net revenue from the forest over a 20-year planning horizon using an interest rate of 4%.

ii) Real stumpage prices are $30 per cord. It costs $15 per acre plus $0.30 per cord to prepare a timber sale. It costs $150 per acre to re-plant stands that are harvested.

iii) Harvest levels in any decade should not be more than 10% larger than or less than 10% smaller than the harvest level in any other decade.

iv) The average age of the forest at the end of the 20-year planning horizon should be at least 17 years old.

b. Use this information from the Excel Answer Report to complete Tables 12.17 through 12.19 on the next pages.

Table 12.17. Acres assigned to each prescription, by site class and initial age class.

Site Class	Initial Age Class	Harvest in Pd 1	Harvest in Pd 2	No Harvest
I	0-10 yr			
	11-20 yr			
	21-30 yr			
II	0-10 yr			
	11-20 yr			
	21-30 yr			
Total				

Table 12.18. Acres harvested by period, by site class, and by age at harvest.

Planning Period	Age at Harvest	Site I	Site II	Total
1	20			
	30			
	Total			
2	20			
	30			
	40			
	Total			
Total acres harvested				
Acres not harvested				

Table 12.19. Acres and volume harvested, and revenues and costs by period for the example forest.

Quantity	Period 1	Period 2
Acres harvested		
Volume harvested (cords)		
Gross Revenues		
Costs		
Net Revenues		

3. It is your job to develop a management plan for a 1.2 million acre forest. The age class distribution by site class is given in Table 12.20. Table 12.21 gives the expected yield per acre by age and site class.

   a. Formulate this management problem as an LP. Be sure to clearly define all of your variables. Use the following assumptions:

      i) You want to minimize the discounted cost of producing at least 1 million cords for each decade in a 30 year planning horizon using an interest rate of 4%.

      ii) It costs \$50 per acre plus \$0.10 per cord to prepare a timber sale. It costs \$100 per acre to re-plant stands that are harvested.

      iii) You want the average age of your ending inventory to be at least 20 years old.

Table 12.20. Initial forest acreage by site and age class.

Age Classes	Acres by site class		
	Site I	Site II	Site III
0-10	120,000	70,000	90,000
11-20	80,000	130,000	130,000
21-30	90,000	170,000	100,000
31-40	140,000	40,000	40,000

Table 12.21. Expected yield by site and age class.

Age	Cords per acre by site class		
	Site I	Site II	Site III
10	2	5	7
20	10	14	17
30	17	22	26
40	21	26	30
50	22	28	33
60	23	29	36
70	24	30	38

b. Solve the problem you formulated in part a using Excel or LINDO. Use the solution to the problem to create a harvest schedule table and a summary table similar to Tables 12.18 and 12.19.