

21

CHAPTER

The Art of Modeling with Spreadsheets

A key step in nearly any OR study is to formulate a mathematical model to represent the problem of interest. You have seen numerous examples of mathematical models throughout this book. These mathematical models generally have been formulated in an algebraic format.

However, the emergence of powerful spreadsheet technology in recent years now provides an alternative way of displaying a mathematical model for a problem that is small enough to fit comfortably into a spreadsheet. This often provides a convenient and intuitive way of representing the problem. The algebra of the model is still there, but it is hidden away in the formulas entered into certain cells of the spreadsheet. This can greatly aid communications between an OR team and a decision maker who may be uncomfortable with algebra. Spreadsheet software (such as the Excel add-in called Solver) includes basic OR algorithms, so various types of spreadsheet models can be solved as soon as they have been formulated. This also makes it easy to do basic sensitivity analysis by simply re-solving the model after changing some of its parameters that are entered into the corresponding cells of the spreadsheet.

Section 3.5 introduced spreadsheet modeling in the context of linear programming problems. Spreadsheet models also were formulated in several other chapters. However, those presentations focused mostly on the characteristics of spreadsheet models that fit the specific types of applications being considered in those chapters. We devote this chapter instead to the general art of formulating spreadsheet models to fit any application. (The discussion assumes that Microsoft Excel is being used, but the same principles also will apply when using other commercially available spreadsheet packages.)

Modeling in spreadsheets is more an art than a science. There is no systematic procedure that invariably will lead to a single correct spreadsheet model. For example, if two OR teams were to be given exactly the same problem to analyze with a spreadsheet, their spreadsheet models will likely look quite different. There is no one right way of modeling any given problem. However, some models will be better than others.

Although no completely systematic procedure is available for modeling in spreadsheets, there is a general process that should be followed. This process has four major steps: (1) *plan* the spreadsheet model, (2) *build* the model, (3) *test* the model, and (4) *analyze* the model and its results. (This process is a streamlined version of both the OR

modeling approach described in Chap. 2 and the outline of a major simulation study presented in Sec. 20.5.) After introducing a case study in Sec. 21.1, the next section will describe this plan-build-test-analyze process in some detail and illustrate the process in the context of the case study. Section 21.2 also will discuss some ways of overcoming common stumbling blocks in the modeling process.

Unfortunately, despite its helpful logical approach, there is no guarantee that the plan-build-test-analyze process will lead to a “good” spreadsheet model. Section 21.3 presents some guidelines for building such models. This section also uses the case study in Sec. 21.1 to illustrate the difference between appropriate formulations and poor formulations of a model.

Even with an appropriate formulation, the initial versions of large spreadsheet models commonly will include some small but troublesome errors, such as inaccurate references to cell addresses or typographical errors when entering equations into cells. These errors often can be difficult to track down. Section 21.4 presents some helpful ways to debug a spreadsheet model and to root out such errors.

The goal of this chapter is to provide a solid foundation for becoming a successful spreadsheet modeler.

21.1 A CASE STUDY: THE EVERGLADE GOLDEN YEARS COMPANY CASH FLOW PROBLEM

This case study involves a problem in cash flow management that the Everglade Golden Years Company faced in late 2009.

The Everglade Golden Years Company operates upscale retirement communities in certain parts of southern Florida. The company was founded in 1946 by Alfred Lee, who was in the right place at the right time to enjoy many successful years during the boom in the Florida economy when many wealthy retirees moved into the region. Today, the company continues to be run by the Lee family, with Alfred’s grandson, Sheldon Lee, as the CEO.

The past few years have been difficult ones for Everglade. The demand for retirement community housing has been light, and Everglade has been unable to maintain full occupancy. However, this market has picked up recently, and the future is looking brighter. Everglade has recently broken ground for the construction of a new retirement community and has more new construction planned over the next 10 years.

Julie Lee is the chief financial officer (CFO) at Everglade. She has spent the last week in front of her computer trying to come to grips with the company’s imminent cash flow problem. Julie has projected Everglade’s net cash flows over the next 10 years as shown in Table 21.1. With less money currently coming in than would be provided by full occupancy and with all the construction costs for the new retirement community, Everglade will have negative cash flow for the next few years. With only \$1 million in cash reserves, it appears that Everglade will need to take out loans in order to meet its financial obligations. Also, to protect against uncertainty, company policy dictates maintaining a balance of at least \$500,000 in cash reserves at all times.

The company’s bank has offered two types of loans to Everglade. The first is a 10-year loan with interest-only payments made annually and then the entire principal repaid in a single balloon payment after 10 years. The fixed interest rate on this long-term loan is a favorable 7 percent per year. The disadvantage is that the interest must be paid on the full loan throughout the 10 years even during those years when some or all of the loan money is not needed. The second option is a series of 1-year loans. These loans can be taken out each year as needed, but each must be repaid (with interest) the following year. Each new loan can be used to help repay the loan for the preceding year if needed. The interest rate for these short-term loans currently is projected to be 10 percent per year. Because of the

■ **TABLE 21.1** Projected net cash flows for the Everglade Golden Years Company over the next 10 years

Year	Projected Net Cash Flow (millions of dollars)
2010	-8
2011	-2
2012	-4
2013	3
2014	6
2015	3
2016	-4
2017	7
2018	-2
2019	10

uncertainty about how interest rates will evolve in the future, planning will be done on the basis of this projection of 10 percent per year. The third option is to use some combination of a 10-year loan and a series of 1-year loans.

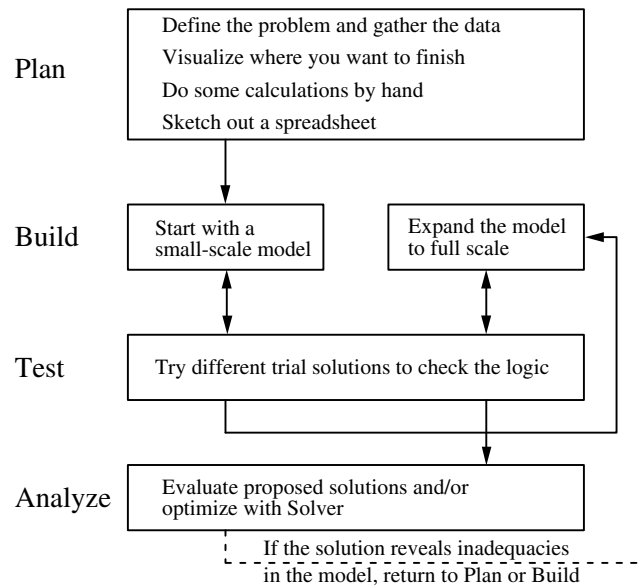
Armed with her cash flow projections and the loan options from the bank, Julie meets with the CEO, Sheldon Lee, to further define the problem. While discussing the three types of loan options, Julia asks two questions. What are the constraints on what can be done? When evaluating the various alternative plans, what should be the measure of performance for choosing the best plan? Sheldon indicates that any of the loan options would be acceptable as long as they observe the company policy of maintaining a balance of at least \$500,000 in cash reserves at all times. He also says that the objective should be to have as large a cash balance as possible at the end of the 10 years after paying off all the loans.

Given these guidelines, you'll see in the next two sections how Julie carefully develops her spreadsheet model for this cash flow problem.

■ 21.2 OVERVIEW OF THE PROCESS OF MODELING WITH SPREADSHEETS

When presented with a problem like Everglade's cash flow problem, the temptation is to jump right in, launch Excel, and start entering a model. Resist this urge. Developing a spreadsheet model without proper planning inevitably leads to a model that is poorly organized and difficult to interpret. To provide you with some structure as you begin learning the art of modeling with spreadsheets, we suggest that you follow the modeling process depicted in Fig. 21.1.

As suggested by this figure, the four major steps in this process are to (1) plan, (2) build, (3) test, and (4) analyze the spreadsheet model. The process mainly flows in this order. However, the two-headed arrows between Build and Test indicate a recursive process where testing frequently results in returning to the Build step to fix some problems discovered during the Test step. This back and forth movement between Build and Test may occur several times until the modeler is satisfied with the model. At the same time that this back and forth movement is occurring, the modeler may be involved with further building of the model. One strategy is to begin with a small version of the model to establish its basic logic and then, after testing verifies its accuracy, to expand to a full-scale model. Even after completing the testing and then analyzing the model, the process may return to the Build step or even the Plan step if the Analysis step reveals inadequacies in the model.



■ **FIGURE 21.1**
A flow diagram for the
general plan-build-test-
analyze process for modeling
with spreadsheets.

Each of these four major steps may also include some detailed steps. For example, Fig. 21.1 lists four detailed steps within the Plan step. Initially, when dealing with a fairly complicated problem, it is helpful to take some time to perform each of these detailed steps manually one at a time. However, as you become more experienced with modeling in spreadsheets, you may find yourself merging some of the detailed steps and quickly performing them mentally. An experienced modeler often is able to do some of these steps mentally, without working them out explicitly on paper. However, if you find yourself getting stuck, it is likely that you are missing a key element from one of the previous detailed steps. You then should go back a step or two and make sure that you have thoroughly completed those preceding steps.

We now describe the various components of the modeling process in the context of the Everglade cash flow problem. At the same time, we also point out some common stumbling blocks encountered while building a spreadsheet model and how these can be overcome.

Plan: Define the Problem and Gather the Data

Before sitting down to start planning how to organize the spreadsheet model, it is necessary to thoroughly understand what the problem is. Therefore, the first order of business is to develop a well-defined statement of the problem being considered. What are the decisions to be made? What are the constraints on these decisions? What is the overall measure of performance for these decisions? These are the kinds of questions that need to be addressed by the members of management who are responsible for making the decisions. This input enables an OR analyst (or team) to identify the “right” problem from management’s viewpoint. After defining this problem, the analyst can then undertake the sometimes lengthy process of gathering the relevant data for analyzing the problem. (See Sec. 2.1 for a more detailed discussion of this process of defining the problem and gathering the data.)

As a member of Everglade’s top management, Julie Lee was able to undertake a major part of this process of defining the company’s cash flow problem by herself. She identified the nature of the problem (projected cash deficits in some future years), the alternative

courses of action (the different types of loan options), and the decisions to be made (the size of the long-term 10-year loan and the sizes of the short-term 1-year loans in the respective years). She also gathered the relevant data for analyzing the problem. However, because the ultimate responsibility for making the decisions rests with Everglade's CEO, Sheldon Lee, Julie was careful to consult with Sheldon before proceeding further. Sheldon imposed a constraint on the decisions by reaffirming that the company would need to continue to observe the policy of maintaining a balance of at least \$500,000 in cash reserves at all times. Sheldon also identified the objective as maximizing the cash balance at the end of the 10 years after paying off all the loans.

Plan: Visualize Where You Want to Finish

Having defined the problem clearly and gathered the relevant data, you now are ready to begin the process of formulating the spreadsheet model. One common stumbling block in the modeling process occurs right at the very beginning. Given a complicated situation like the one facing Julie at Everglade, it sometimes can be difficult to decide how to even get started. At this point, it can be helpful to think about where you want to end up. For example, what information should Julie provide in her report to Sheldon? What should the "answer" look like when presenting the recommended approach to the problem? What kinds of numbers need to be included in the recommendation? The answers to these questions can quickly lead you to the heart of the problem and help get the modeling process started.

The question that Julie is addressing is which loan, or combination of loans, to use and in what amounts. The long-term loan is taken in a single lump sum. Therefore, the "answer" should include a single number indicating how much money to borrow now at the long-term rate. The short-term loan can be taken in any or all of the 10 years, so the "answer" should include 10 numbers indicating how much to borrow at the short-term rate in each given year. These will be the **changing cells** (the cells containing the values of the decision variables) in the spreadsheet model.

What other numbers should Julie include in her report to Sheldon? The key numbers would be the projected cash balance at the end of each year, the amount of the interest payments, and when loan payments are due. These will be **output cells** (the cells that show quantities that are calculated from the changing cells) in the spreadsheet model.

It is important to distinguish between the numbers that represent decisions (changing cells) and those that represent results (output cells). For instance, it may be tempting to include the cash balances as changing cells. These cells clearly change depending on the decisions made. However, the cash balances are a *result* of how much is borrowed, how much is paid, and all of the other cash flows. They cannot be chosen independently, but instead are a function of the other numbers in the spreadsheet. The distinguishing characteristic of changing cells (the loan amounts) is that they do not depend on anything else. They represent the independent decisions being made. They impact the other numbers, but not vice versa.

At this stage in the process, you should have a clear idea of what the answer will look like, including what and how many changing cells are needed, and what kind of results (output cells) should be obtained.

Plan: Do Some Calculations by Hand

When building a model, another common stumbling block can arise when trying to enter a formula in one of the output cells. For example, just how does Julie keep track of the cash balances in the Everglade cash flow problem? What formulas need to be entered? There are a lot of factors that enter into this calculation, so it is easy to get overwhelmed.

If you are getting stuck at this point, it can be a very useful exercise to do some calculations by hand. Just pick some numbers for the changing cells and determine with a calculator or pencil and paper what the results should be. For example, pick some loan amounts for Everglade, and then calculate the company's resulting cash balance at the end of the first couple years. Let's say Everglade takes a long-term loan of \$6 million, and then adds short-term loans of \$2 million in 2010 and \$5 million in 2011. How much cash would the company have left at the end of 2010 and at the end of 2011?

These two quantities can be calculated by hand as follows. In 2010, Everglade has some initial money in the bank (\$1 million), a negative cash flow from its business operations (-\$8 million), and a cash inflow from the long-term and short-term loans (\$6 million and \$2 million, respectively). Thus, the ending balance for 2010 would be:

Ending Balance (2010) = Starting Balance	\$1 million
+ Cash Flow (2010)	- \$8 million
+ LT Loan (2010)	+ \$6 million
+ ST Loan (2010)	+ <u>\$2 million</u>
	\$1 million

The calculations for the year 2011 are a little more complicated. In addition to the starting balance left over from 2010 (\$1 million), negative cash flow from business operations for 2011 (-\$2 million), and a new short-term loan for 2011 (\$5 million), the company will need to make interest payments on its 2010 loans as well as pay back the short-term loan from 2010. The ending balance for 2011 is therefore:

Ending Balance (2011) = Starting Balance (from 2010)	\$1 million
+ Cash Flow (2011)	- \$2 million
+ ST Loan (2011)	+ \$5 million
- LT Interest Payment	- (7%)(6 million)
- ST Interest Payment	- (10%)(2 million)
- ST Loan Payback (2010)	- <u>\$2 million</u>
	\$1.38 million

Doing calculations by hand can help in a couple of ways. First, it can help clarify what formula should be entered for an output cell. For instance, looking at the by-hand calculations above, it appears that the formula for the ending balance for a particular year should be

$$\text{Ending balance} = \text{starting balance} + \text{cash flow} + \text{loans} - \text{interest payments} - \text{loan paybacks.}$$

It now will be a simple exercise to enter the proper cell references in the formula for the ending balance in the spreadsheet model. Second, hand calculations can help to verify the spreadsheet model. By plugging in a long-term loan of \$6 million, along with short-term loans of \$2 million in 2010 and \$5 million in 2011, into a completed spreadsheet, the ending balances should be the same as calculated above. If they're not, this suggests an error in the spreadsheet model (assuming the hand calculations are correct).

Plan: Sketch Out a Spreadsheet

Any model typically has a large number of different elements that need to be included on the spreadsheet. For the Everglade problem, these would include some data cells (interest rates, starting balance, minimum balances, and cash flows), some changing cells (loan amounts), and a number of output cells (interest payments, loan paybacks, and ending balances). Therefore, a potential stumbling block can arise when trying to organize and lay

out the spreadsheet model. Where should all the pieces fit on the spreadsheet? How do you begin putting together the spreadsheet?

Before firing up Excel and blindly entering the various elements, it can be helpful to sketch a layout of the spreadsheet. Is there a logical way to arrange the elements? A little planning at this stage can go a long way toward building a spreadsheet that is well organized. Don't bother with numbers at this point. Simply sketch out blocks on a piece of paper for the various data cells, changing cells, and output cells, and label them. (The **data cells** are the cells that show the data for the problem.) Concentrate on the layout. Should a block of numbers be laid out in a row or a column, or as a two-dimensional table? Are there common row or column headings for different blocks of cells? If so, try to arrange the blocks in consistent rows or columns so they can utilize a single set of headings. Try to arrange the spreadsheet so that it starts with the data at the top and progresses logically toward the **target cell** (the output cell that contains the value of the objective function) at the bottom. This will be easier to understand and follow than if the data cells, changing cells, output cells, and target cell are all scattered throughout the spreadsheet.

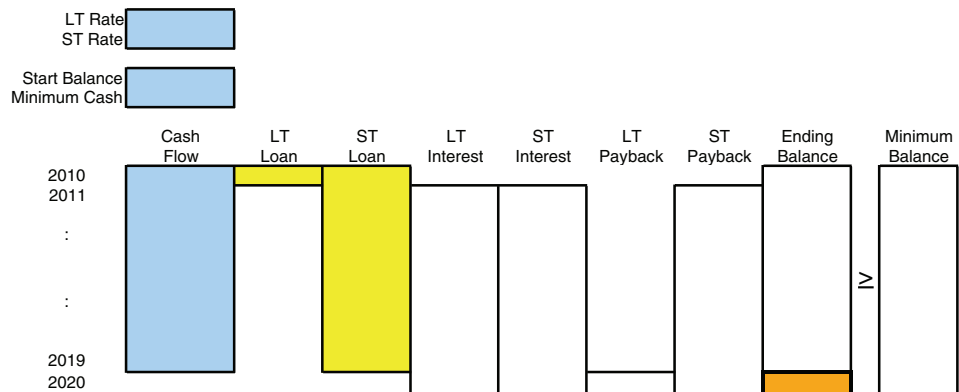
A sketch of a potential spreadsheet layout for the Everglade problem is shown in Fig. 21.2. The data cells for the interest rates, starting balance, and minimum cash balance are at the top of the spreadsheet. All of the remaining elements in the spreadsheet then follow the same structure. The rows represent the different years (from 2010 through 2020). All the various cash inflows and outflows are then broken out in the columns, starting with the projected cash flow from the business operations (with data for each of the 10 years), continuing with the loan inflows, interest payments, and loan paybacks, and culminating with the ending balance (calculated for each year). The long-term loan is a one-time loan (in 2010), so it is sketched as a single cell. The short-term loan can occur in any of the 10 years (2010 through 2019), so it is sketched as a block of cells. The interest payments start one year after the loans. The long-term loan is paid back 10 years later (2020).

Organizing the elements with a consistent structure, like in Fig. 21.2, not only saves having to retype the year labels for each element, but also makes the model easier to understand. Everything that happens in a given year is arranged together in a single row.

It is generally easiest to start sketching the layout with the data. The structure of the rest of the model should then follow the structure of the data cells. For example, once the projected cash flows data are sketched as a vertical column (with each year in a row), then it follows that the other cash flows should be structured the same way.

There is also a logical progression to the spreadsheet. The data for the problem are located at the top and left of the spreadsheet. Then, since the cash flow, loan amounts,

■ **FIGURE 21.2**
Sketch of the spreadsheet for Everglade's cash flow problem.



interest payments, and loan paybacks are all part of the calculation for the ending balance, the columns are arranged this way, with the ending balance directly to the right of all these other elements. Since Sheldon has indicated that the objective is to maximize the ending balance in 2020, this cell is designated to be the target cell.

Each year, the balance must be greater than the minimum required balance (\$500,000). Since this will be a constraint in the model, it is logical to arrange the balance and minimum balance blocks of numbers adjacent to each other in the spreadsheet. You can put the \geq signs on the sketch to remind yourself that these will be constraints.

Build: Start with a Small Version of the Spreadsheet

Once you've thought about a logical layout for the spreadsheet, it is finally time to open a new worksheet in Excel and start building the model. If it is a complicated model, you may want to start by building a small, readily manageable version of the model. The idea is to first make sure that you've got the logic of the model worked out correctly for the small version before expanding the model to full scale.

For example, in the Everglade problem, we could get started by building a model for just the first two years (2010 and 2011), like the spreadsheet shown in Fig. 21.3. This spreadsheet is set up to follow the layout suggested in the sketch of Fig. 21.2. The loan amounts are in columns D and E. Since the interest payments are not due until the following year, the formulas in columns F and G refer to the loan amounts from the preceding year (LTLoan, or D11, for the long-term loan, and E11 for the short-term loan). The loan payments are calculated in columns H and I. Column H is blank because the long-term loan does not need to be repaid until 2020. The short-term loan is repaid one year later, so the

■ FIGURE 21.3

A small version (years 2010 and 2011 only) of the spreadsheet for the Everglade cash flow management problem.

	A	B	C	D	E	F	G	H	I	J	K	L
1	Everglade Cash Flow Management Problem (Years 2010 and 2011)											
2												
3		LT Rate	7%									
4		ST Rate	10%									
5												
6		Start Balance	1			(all cash figures in millions of dollars)						
7		MinimumCash	0.5									
8												
9			Cash	LT	ST	LT	ST	LT	ST	Ending		Minimum
10		Year	Flow	Loan	Loan	Interest	Interest	Payback	Payback	Balance		Balance
11		2010	-8	6	2					1.00	>=	0.50
12		2011	-2		5	-0.42	-0.20			-2.00	>=	0.50

	F	G	H	I	J	K	L
9	LT	ST	LT	ST	Ending		Minimum
10	Interest	Interest	Payback	Payback	Balance		Balance
11					=StartBalance+SUM(C11:I11)	>=	=MinimumCash
12	=LTRate*LTLoan	=STRate*E11		=E11	=J11+SUM(C12:I12)	>=	=MinimumCash

Range Name	Cell
LTLoan	D11
LTRate	C3
MinimumCash	C7
StartBalance	C6
STRate	C4

formula in cell I12 refers to the short-term loan taken the preceding year (cell E11). The ending balance in 2010 is the starting balance plus the sum of all the various cash flows that occur in 2010 (cells C11:I11). The ending balance in 2011 is the ending balance in 2003 (cell J11) plus the sum of all the various cash flows that occur in 2011 (cells C12:I12). All these formulas are summarized below the spreadsheet in Fig. 21.3.

The bottom of Fig. 21.3 shows the “range names” given to certain cells. A **range name** is a descriptive name given to a cell or a block of cells that immediately identifies what is there. As illustrated by certain formulas (especially the one in cell F12) below the spreadsheet, writing a formula in terms of range names instead of cell addresses makes the formula much easier to interpret. (We will discuss range names and their usefulness further in Sec. 21.3.)

Building a small version of the spreadsheet works very well for spreadsheets that have a time dimension. For example, instead of jumping right into a 10-year planning problem, you can start with the simpler problem of just looking at a couple of years. Once this smaller model is working correctly, you then can expand the model to 10 years.

Even if a spreadsheet model does not have a time dimension, the same concept of starting small can be applied. For example, if certain constraints considerably complicate a problem, start by working on a simpler problem without the difficult constraints. Get the simple model working, and then move on to tackle the difficult constraints. If a model has many sets of output cells, you can build up a model piece by piece by working on one set of output cells at a time, making sure each set works correctly before moving on to the next.

Test: Test the Small Version of the Model

If you do start with a small version of the model first, be sure to test this version thoroughly to make sure that all the logic is correct. It is far easier to fix a problem early, while the spreadsheet is still a manageable size, rather than later after an error has been propagated throughout a much larger spreadsheet.

To test the spreadsheet, try entering values in the changing cells for which you know what the values of the output cells should be, and then see if the spreadsheet gives the results that you expect. For example, in Fig. 21.3, if zeroes are entered for the loan amounts, then the interest payments and loan payback quantities should also be zero. If \$1 million is borrowed for both the long-term loan and the short-term loan, then the interest payments the following year should be \$70,000 and \$100,000, respectively. (Recall that the interest rates are 7 percent and 10 percent, respectively.) If Everglade takes out a \$6 million long-term loan and a \$2 million short-term loan in 2010, plus a \$5 million short-term loan in 2011, then the ending balances should be \$1 million for 2010 and \$1.38 million for 2011 (based on the calculations done earlier by hand). All these tests work correctly for the spreadsheet in Fig. 21.3, so we can be fairly certain that it is correct.

If the output cells are not giving the results that you expect, then carefully look through the formulas to see if you can determine and fix the problem. Section 21.4 will give further guidance on some ways to debug a spreadsheet model.

Build: Expand the Model to Full-Scale Size

Once a small version of the spreadsheet has been tested to make sure all the formulas are correct and everything is working properly, the model can be expanded to full-scale size. Excel’s fill commands often can be used to quickly copy the formulas into the remainder of the model. For Fig. 21.3, the formulas in columns F, G, I, J, and L can be copied using the Fill Down command in the Editing Group of the Home tab (for Excel 2007) or in the Edit menu (for earlier versions of Excel) to obtain all the formulas shown in Fig. 21.4. For example, selecting cells G12:G21 and choosing Fill Down will take the formula in cell G12 and copy it (after adjusting the cell address in Column E for the formula) into cells G13 through G21.

When using the fill commands, it is important to understand the difference between relative and absolute references. Consider the formula in cell G12 ($=-STRate*E11$).

21-10

CHAPTER 21 THE ART OF MODELING WITH SPREADSHEETS

	A	B	C	D	E	F	G	H	I	J	K	L
1	Everglade Cash Flow Management Problem											
2												
3		LT Rate	7%									
4		ST Rate	10%									
5												
6		Start Balance	1	(all cash figures in millions of dollars)								
7		Minimum Cash	0.5									
8												
9			Cash	LT	ST	LT	ST	LT	ST	Ending		Minimum
10		Year	Flow	Loan	Loan	Interest	Interest	Payback	Payback	Balance		Balance
11		2010	-8	6	2					1.00	>=	0.5
12		2011	-2		5	-0.42	-0.20			-2	>=	1.38
13		2012	-4		0	-0.42	-0.50			-5	>=	0.5
14		2013	3		0	-0.42	0			0	>=	0.5
15		2014	6		0	-0.42	0			0	>=	0.5
16		2015	3		0	-0.42	0			0	>=	0.5
17		2016	-4		0	-0.42	0			0	>=	0.5
18		2017	7		0	-0.42	0			0	>=	0.5
19		2018	-2		0	-0.42	0			0	>=	0.5
20		2019	10		0	-0.42	0			0	>=	0.5
21		2020				-0.42	0	-6	0	5.10	>=	0.5

	F	G	H	I	J	K	L
9	LT	ST	LT	ST	Ending		Minimum
10	Interest	Interest	Payback	Payback	Balance		Balance
11					=StartBalance+SUM(C11:I11)	>=	=MinimumCash
12	=-LTRate*LTLoan	=-STRate*E11		=-E11	=J11+SUM(C12:I12)	>=	=MinimumCash
13	=-LTRate*LTLoan	=-STRate*E12		=-E12	=J12+SUM(C13:I13)	>=	=MinimumCash
14	=-LTRate*LTLoan	=-STRate*E13		=-E13	=J13+SUM(C14:I14)	>=	=MinimumCash
15	=-LTRate*LTLoan	=-STRate*E14		=-E14	=J14+SUM(C15:I15)	>=	=MinimumCash
16	=-LTRate*LTLoan	=-STRate*E15		=-E15	=J15+SUM(C16:I16)	>=	=MinimumCash
17	=-LTRate*LTLoan	=-STRate*E16		=-E16	=J16+SUM(C17:I17)	>=	=MinimumCash
18	=-LTRate*LTLoan	=-STRate*E17		=-E17	=J17+SUM(C18:I18)	>=	=MinimumCash
19	=-LTRate*LTLoan	=-STRate*E18		=-E18	=J18+SUM(C19:I19)	>=	=MinimumCash
20	=-LTRate*LTLoan	=-STRate*E19		=-E19	=J19+SUM(C20:I20)	>=	=MinimumCash
21	=-LTRate*LTLoan	=-STRate*E20	=-LTLoan	=-E20	=J20+SUM(C21:I21)	>=	=MinimumCash

Range Name	Cells
CashFlow	C11:C20
EndBalance	J21
Ending Balance	J11:J21
LTLoan	D11
LTRate	C3
MinimumBalance	L11:L21
MinimumCash	C7
StartBalance	C6
STLoan	E11:E20
STRate	C4

■ **FIGURE 21.4**

A complete spreadsheet model for the Everglade cash flow management problem, including the equations entered into the target cell EndBalance (J21) and all the other output cells, before calling on the Excel Solver. The entries in the changing cells, LTLoan (D11) and STLoan (E11:E20), are only a trial solution at this stage.

References to cells or ranges within a formula (like E11) are usually based upon their position relative to the cell containing the formula. Thus, E11 is two cells to the left and one cell up. This is known as a **relative reference**. When this formula is copied to a new cell, the reference is automatically adjusted to refer to the new cell that is at the same relative location (two cells to the left and one cell up). For example, the formula copied to G13 refers to cell E12, the one in G14 refers to cell E13, and so on. This is exactly what we want, since we always want the interest payment to be based on the short-term loan that was taken one year ago (two cells to the left and one cell up).

In contrast, the reference to STRate (C4) in the formula for cell G12 is called an **absolute reference**. These references do not change when they are filled into other cells. That is, wherever this formula is copied, the formula will still refer to the cell STRate (C4).

To make a relative reference, simply enter the cell address (e.g., E11). To make an absolute reference, either use a range name for the cell (e.g., STRate) or put \$ signs in front of the letter and number of the cell reference (e.g., \$E\$11). Similarly, you can make the column absolute and the row relative (or vice versa) by putting a \$ sign in front of only the letter (or number) of the cell reference. For example, if a reference to \$E11 in a formula is copied to a new location, the \$E will remain constant, but the row number will adjust. In the case of the formula for cell G12 in Fig. 21.4, \$E11 could have been used for the cell reference since column E will remain constant, but the \$ sign is not necessary (and so was not used) when copying down column G since the relative location of column E (two columns to the left) always remains the same.

After using the Fill Down command to copy the formulas in columns F, G, I, J, and L, and entering the LT loan payback into cell H21, the complete model appears as shown in Fig. 21.4.

Test: Test the Full-Scale Version of the Model

Just as it was important to test the small version of the model, it needs to be tested again after it is expanded to full-scale size. The procedure is the same one followed for testing the small version, including the ideas that will be presented in Sec. 21.4 for debugging a spreadsheet model.

Analyze: Analyze the Model

Before using the Solver dialogue box to prepare for applying Solver, the spreadsheet in Fig. 21.4 is merely an evaluative model for Everglade. It can be used to evaluate any proposed solution, including quickly determining what interest and loan payments will be required and what the resulting balances will be at the end of each year. For example, LTLoan (D11) and STLoan (E11:E20) in Fig. 21.4 show one possible plan, which turns out to be unacceptable because EndingBalance (J11:J21) indicates that a negative ending balance would result in four of the years.

To optimize the model, the Solver dialogue box is used as shown in Fig. 21.5 to specify the target cell, the changing cells, and the constraints. (Even when constraints already are displayed in the spreadsheet, as in columns J, K, and L of this figure, Excel allows these constraints to be violated unless they also are specified in the Solver dialogue box.) Everglade management wants to find a combination of loans that will keep the company solvent throughout the next 10 years (2010–2019) and then will leave as large a cash balance as possible in 2020 after paying off all the loans. Therefore, the target cell to be maximized is EndBalance (J21), and the changing cells are the loan amounts LTLoan (D11) and STLoan (E11:E20). To assure that Everglade maintains a minimum balance of at least \$500,000 at the end of each year, the constraints for the model are EndingBalance (J11:J21) \geq MinimumBalance (L11:L21).

	A	B	C	D	E	F	G	H	I	J	K	L
1	Everglade Cash Flow Management Problem											
2												
3		LT Rate	7%									
4		ST Rate	10%									
5												
6		Start Balance	1	(all cash figures in millions of dollars)								
7		Minimum Cash	0.5									
8												
9		Cash	LT	ST	LT	ST	LT	ST	Ending		Minimum	
10	Year	Flow	Loan	Loan	Interest	Interest	Payback	Payback	Balance			Balance
11	2010	-8	6.65	0.85					0.50	>=		0.50
12	2011	-2		3.40	-0.47	-0.09			-0.85			0.50
13	2012	-4		8.21	-0.47	-0.34			-3.40			0.50
14	2013	3		6.49	-0.47	-0.82			-8.21			0.50
15	2014	6		1.61	-0.47	-0.65			-6.49			0.50
16	2015	3		0	-0.47	-0.16			-1.61			1.27
17	2016	-4		3.70	-0.47	0			0			0.50
18	2017	7		0	-0.47	-0.37			-3.70			2.97
19	2018	-2		0	-0.47	0			0			0.50
20	2019	10		0	-0.47	0			0			10.03
21	2020				-0.47	0	-6.65	0	2.92	>=		0.50

	F	G	H	I	J	K	L
9	LT	ST	LT	ST	Ending		Minimum
10	Interest	Interest	Payback	Payback	Balance		Balance
11					=StartBalance+SUM(C11:I11)	>=	=MinimumCash
12	=-LTRate*LTLoan	=-STRate*E11		=E11	=J11+SUM(C12:I12)	>=	=MinimumCash
13	=-LTRate*LTLoan	=-STRate*E12		=E12	=J12+SUM(C13:I13)	>=	=MinimumCash
14	=-LTRate*LTLoan	=-STRate*E13		=E13	=J13+SUM(C14:I14)	>=	=MinimumCash
15	=-LTRate*LTLoan	=-STRate*E14		=E14	=J14+SUM(C15:I15)	>=	=MinimumCash
16	=-LTRate*LTLoan	=-STRate*E15		=E15	=J15+SUM(C16:I16)	>=	=MinimumCash
17	=-LTRate*LTLoan	=-STRate*E16		=E16	=J16+SUM(C17:I17)	>=	=MinimumCash
18	=-LTRate*LTLoan	=-STRate*E17		=E17	=J17+SUM(C18:I18)	>=	=MinimumCash
19	=-LTRate*LTLoan	=-STRate*E18		=E18	=J18+SUM(C19:I19)	>=	=MinimumCash
20	=-LTRate*LTLoan	=-STRate*E19		=E19	=J19+SUM(C20:I20)	>=	=MinimumCash
21	=-LTRate*LTLoan	=-STRate*E20	=-LTLoan	=E20	=J20+SUM(C21:I21)	>=	=MinimumCash

Solver Parameters

Set Target Cell:

Equal To: Max Min

By Changing Cells:

Subject to the Constraints:

Solver Options

Assume Linear Model

Assume Non-Negative

Range Name	Cells
CashFlow	C11:C20
EndBalance	J21
Ending Balance	J11:J21
LTLoan	D11
LTRate	C3
MinimumBalance	L11:L21
MinimumCash	C7
StartBalance	C6
STLoan	E11:E20
STRate	C4

■ FIGURE 21.5

A complete spreadsheet model for the Everglade cash flow management problem after calling on the Excel Solver to obtain the optimal solution shown in the changing cells LTLoan (D11) and STLoan (E11:E20). The target cell EndBalance (J21) indicates that the resulting cash balance in 2020 will be \$2.92 million if all the data cells prove to be accurate.

21.3 SOME GUIDELINES FOR BUILDING “GOOD” SPREADSHEET MODELS 21-13

After running Solver, the optimal solution is shown in Fig. 21.5. The changing cells, LTLoan (D11) and STLoan (E11:E20) give the loan amounts in the various years. The target cell EndBalance (J21) indicates that the ending balance in 2020 will be \$2.92 million.

Conclusion of the Case Study

The spreadsheet model developed by Everglade’s CFO, Julie Lee, is the one shown in Fig. 21.5. Her next step is to submit to her CEO, Sheldon Lee, a report that recommends the plan obtained by this model.

Soon thereafter, Sheldon and Julie meet to discuss her report. The one concern that Sheldon raises is that the cash flows in the coming years shown in column C of Fig. 21.5 are only estimates. When there is a shift in the economy, or when other unexpected developments occur that impact on the company, those cash flows can change substantially. Would the recommended plan still be a good one if those kinds of changes were to occur? Julie and Sheldon agree that some sensitivity analysis should be done to check on the effect of such changes. Fortunately, Julie had set up the spreadsheet properly (providing a data cell for the cash flow in each of the next 10 years) to enable performing sensitivity analysis immediately by simply trying different numbers in some of these data cells. After spending half an hour trying different numbers, Sheldon and Julie conclude that the plan in Fig. 21.5 will be a sound initial financial plan for the next 10 years even if future cash flows deviate somewhat from current forecasts. If deviations do occur, adjustments will of course need to be made in the short-term loan amounts. At any point, Julie also will have the option of returning to the company’s bank to try to arrange another long-term loan for the remainder of the 10 years at a lower interest rate than for short-term loans. If so, essentially the same spreadsheet model as in Fig. 21.5 can be used, along with the Excel Solver, to find the optimal adjusted financial plan for the remainder of the 10 years.

21.3 SOME GUIDELINES FOR BUILDING “GOOD” SPREADSHEET MODELS

There are many ways to set up a model on a spreadsheet. While one of the benefits of spreadsheets is the flexibility they offer, this flexibility also can be dangerous. Although Excel provides many features (such as range names, shading, borders, etc.) that allow you to create “good” spreadsheet models that are easy to understand, easy to debug, and easy to modify, it is also easy to create “bad” spreadsheet models that are difficult to understand, difficult to debug, and difficult to modify. The goal of this section is to provide some guidelines that will help you create “good” spreadsheet models.

Enter the Data First

Any spreadsheet model is driven by the data in the spreadsheet. The form of the entire model is built around the structure of the data. Therefore, it is always a good idea to enter and carefully lay out all the data before you begin to set up the rest of the model. The model structure then can conform to the layout of the data as closely as possible.

Often, it is easier to set up the rest of the model when the data are already on the spreadsheet. In the Everglade problem (see Fig. 21.5), the data for the cash flows have been laid out in the first columns of the spreadsheet (B and C), with the year labels in column B and the data in cells C11:C20. Once the data are in place, the layout for the rest of the model quickly falls into place around the structure of the data. It is only logical to lay out the changing cells and output cells using the same structure, with each of the various cash flows in columns that utilize the same row labels from column B.

Now reconsider the spreadsheet model developed in Sec. 3.5 for the Wyndor Glass Co. problem. This spreadsheet model is repeated here as Fig. 21.6. The data for the Hours Used Per Batch Produced have been laid out in the center of the spreadsheet in cells C7:D9.

	A	B	C	D	E	F	G
1	Wyndor Glass Co. Product-Mix Problem						
2							
3			Doors	Windows			
4		Profit Per Batch	\$3,000	\$5,000			
5					Hours		Hours
6			Hours Used Per Batch Produced		Used		Available
7		Plant 1	1	0	2	<=	4
8		Plant 2	0	2	12	<=	12
9		Plant 3	3	2	18	<=	18
10							
11			Doors	Windows			Total Profit
12		Batches Produced	2	6			\$36,000

Solver Parameters

Set Target Cell: TotalProfit

Equal To: Max Min

By Changing Cells: BatchesProduced

Subject to the Constraints: HoursUsed <= HoursAvailable

	E
5	Hours
6	Used
7	=SUMPRODUCT(C7:D7,BatchesProduced)
8	=SUMPRODUCT(C8:D8,BatchesProduced)
9	=SUMPRODUCT(C9:D9,BatchesProduced)

	G
11	Total Profit
12	=SUMPRODUCT(ProfitPerBatch,BatchesProduced)

Solver Options

Assume Linear Model

Assume Non-Negative

Range Name	Cells
BatchesProduced	C12:D12
HoursAvailable	G7:G9
HoursUsed	E7:E9
HoursUsedPerBatchProduced	C7:D9
ProfitPerBatch	C4:D4
TotalProfit	G12

■ **FIGURE 21.6**

The spreadsheet model for the Wyndor Glass Co. product-mix problem introduced in Sec. 3.1.

The output cells, HoursUsed (E7:E9), then have been placed immediately to the right of these data and to the left of the data on HoursAvailable (G7:G9), where the row labels for these output cells are the same as for all these data. This makes it easy to interpret the three constraints being laid out in rows 7–9 of the spreadsheet model. Next, the changing cells and target cell have been placed together in row 12 below the data, where the column labels for the changing cells are the same as for the columns of data above.

The locations of the data occasionally will need to be shifted somewhat to better accommodate the overall model. However, with this caveat, the model structure generally should conform to the data as closely as possible.

Organize and Clearly Identify the Data

Related data should be grouped together in a convenient format and entered into the spreadsheet with labels that clearly identify the data. For data laid out in tabular form, the table should have a heading that provides a general description of the data, and then each row and column should have a label that will identify each entry in the table. The units of the data also should be identified. Different types of data should be well separated in the spreadsheet. However, if two tables need to use the same labels for either their rows or columns, then be consistent in making them either rows in both tables or columns in both tables.

In the Wyndor Glass Co. problem (Fig. 21.6), the three sets of data have been grouped into tables and clearly labeled Profit Per Batch, Hours Used Per Batch Produced, and Hours Available. The units of the data are identified (dollar signs are included in the unit profit data, and hours are indicated in the labels of the time data). Finally, all three data tables make consistent use of rows and columns. Since the Profit Per Batch data have their product labels (Doors and Windows) in columns C and D, the Hours Used Per Batch Produced data use this same structure. This structure also is carried through to the changing cells (Batches Produced). Similarly, the data for each plant are in the rows (row 7–9) for both the Hours Used Per Batch Produced data *and* the Hours Available data. Keeping the data oriented the same way is not only less confusing, but also makes it possible to use the SUMPRODUCT function. The SUMPRODUCT function introduced in Sec. 3.6 assumes that the two ranges are exactly the same shape (i.e., the same number of rows *and* columns). If the Profit Per Batch data and the Batches Produced data had not been oriented the same way (e.g., one in a column and the other in a row), it would not have been possible to use the SUMPRODUCT function to sum the product of each of the individual terms in the two ranges of cells in the Total Profit calculation.

Similarly, for the Everglade problem (Fig. 21.5), the five sets of data have been grouped into cells and tables and clearly labeled ST Rate, LT Rate, Start Balance, Cash Flow, and Minimum Cash. The units of the data are identified (cells F6:H6 specify that all cash figures are in millions of dollars), and all the tables make consistent use of rows and columns (years in the rows).

Enter Each Piece of Data into One Cell Only

If a piece of data is needed in more than one formula, then refer to the original data cell rather than repeating the data in additional places. This makes the model much easier to modify. If the value of that piece of data changes, it only needs to be changed in one place. You do not need to search through the entire model to find all the places where the data value appears.

For example, in the Everglade problem (Fig. 21.5), there is a company policy of maintaining a cash balance of at least \$500,000 at all times. This translates into a constraint for the minimum balance of \$500,000 at the end of each year. Rather than entering the minimum cash position of 0.5 (in millions of dollars) into all the cells in column L, it is entered once in MinimumCash (C7) and then referred to by the cells in MinimumBalance (L11:L21). Then, if this policy were to change to, say, a minimum of \$200,000 cash, the number would need to be changed in only one place.

Separate Data from Formulas

Avoid using numbers directly in formulas. Instead, enter any needed numbers into data cells, and then refer to the data cells as needed. For example, in the Everglade problem (Fig. 21.5), all the data (the interest rates, starting balance, minimum cash, and projected cash flows) are entered into separate data cells on the spreadsheet. When these numbers are needed to calculate the interest charges (in columns F and G), loan payments

(in column H and I), ending balances (column J), and minimum balances (column L), the data cells are referred to rather than entering these numbers directly in the formulas.

Separating the data from the formulas has a couple advantages. First, all the data are visible on the spreadsheet rather than buried in formulas. Seeing all the data makes the model easier to interpret. Second, the model is easier to modify since changing data only requires modifying the corresponding data cells. You don't need to modify any formulas. This proves to be very important when it comes time to perform sensitivity analysis to see what the effect would be if some of the estimates in the data cells were to take on other plausible values.

Keep it Simple

Avoid the use of powerful Excel functions when simpler functions are available that are easier to interpret. As much as possible, stick to SUMPRODUCT or SUM functions. This makes the model easier to understand and also helps to ensure that the model will be linear. (Linear models are considerably easier to solve than others.) Try to keep formulas short and simple. If a complicated formula is required, break it out into intermediate calculations with subtotals. For example, in the Everglade spreadsheet, each element of the loan payments is broken out explicitly: LT Interest, ST Interest, LT Payback, and ST Payback. Some of these columns could have been combined (e.g., into two columns with LT Payments and ST Payments, or even into one column for all Loan Payments). However, this makes the formulas more complicated, and also makes the model harder to test and debug. As laid out, the individual formulas for the loan payments are so simple that their values can be predicted easily without even looking at the formula. This simplifies the testing and debugging of the model.

Use Range Names

One way to refer to a block of related cells (or even a single cell) in a spreadsheet formula is to use its cell address (e.g., L11:L21 or C3). However, when reading the formula, this requires looking at that part of the spreadsheet to see what kind of information is given there. As mentioned in Sec. 21.2, a better alternative often is to assign a descriptive *range name* to the block of cells that immediately identifies what is there. (This is done by selecting the block of cells, clicking on the name box on the left of the formula bar above the spreadsheet, and then typing a name.) This is especially helpful when writing a formula for an output cell. Writing the formula in terms of range names instead of cell addresses makes the formula much easier to interpret. Range names also make the description of the model in the Solver dialogue box much easier to understand.

Figure 21.5 illustrates the use of range names for the Everglade spreadsheet model, where these range names are listed in the upper right-hand corner of the figure. (Spaces are not allowed in range names, so we have used capital letters to distinguish the start of each new word in a range name.) For example, consider the formula for long-term interest in cell F12. Since the long-term rate is given in cell C3 and the long-term loan amount is in cell D11, the formula for the long-term interest could have been written as $=-C3*D11$. However, by using the range name *LTRate* for cell C3 and the range name *LTLoan* for cell D11, the formula instead becomes $=-LTRate*LTLoan$, which is much easier to interpret at a glance.

On the other hand, be aware that it is easy to get carried away with defining range names. Defining too many range names can be more trouble than it is worth. For example, when related data are grouped together in a table, we recommend giving a range name only for the entire table rather than for the individual rows and columns. In general, we suggest defining range names only for each group of data cells, the changing cells, the target cell, and both sides of each group of constraints (the left-hand side and the right-hand side).

Care also should be taken to assure that it is easy to quickly identify which cells are referred to by a particular range name. Use a name that corresponds exactly to the label

21.3 SOME GUIDELINES FOR BUILDING “GOOD” SPREADSHEET MODELS 21-17

on the spreadsheet. For example, in Fig. 21.5, columns J and L are labeled Ending Balance and Minimum Balance on the spreadsheet, so we use the range names EndingBalance and MinimumBalance. Using exactly the same name as the label on the spreadsheet makes it quick and easy to find the cells that are referred to by a range name.

When desired, a list of all the range names and their corresponding cell addresses can be pasted directly into the spreadsheet by choosing Paste from the Use in Formula menu on the Formulas tab (for Excel 2007) or Name\Paste from the Insert menu (for earlier versions of Excel), and then clicking Paste List. Such a list (after reformatting) is included below many of the spreadsheets displayed in this chapter.

When modifying an existing model that utilizes range names, care should be taken to assure that the range names continue to refer to the correct range of cells. When inserting a row or column into a spreadsheet model, it is helpful to insert the row or column into the middle of a range rather than at the end. For example, to add another product to a product-mix model with four products, add a column between products 2 and 3 rather than after product 4. This will automatically extend the relevant range names to span across all five columns since these range names will continue to refer to everything between product 1 and product 4, including the newly inserted column for the fifth product. Similarly, deleting a row or column from the middle of a range will contract the span of the relevant range names appropriately. You can double-check the cells that are referred to by a range name by choosing that range name from the name box (on the left of the formula bar above the spreadsheet). This will highlight the cells that are referred to by the chosen range name.

Use Relative and Absolute Referencing to Simplify Copying Formulas

Whenever multiple related formulas will be needed, try to enter the formula just once and then use Excel's fill commands to replicate the formula. Not only is this quicker than retyping the formula, but it is also less prone to error.

We saw a good example of this when discussing the expansion of the model to full-scale size in the preceding section. Starting with the 2-year spreadsheet in Fig. 21.3, fill commands were used to copy the formulas in columns F, G, I, J, and L for the remaining years to create the full-scale, 10-year spreadsheet in Fig. 21.4.

Use Borders, Shading, and Colors to Distinguish between Cell Types

It is important to be able to easily distinguish between the data cells, changing cells, output cells, and target cell in a spreadsheet. One way to do this is to use different borders and cell shading for each of these different types of cells. For example, data cells could appear lightly shaded with a light border, changing cells darkly shaded with a heavy border, output cells with no shading, and the target cell darkly shaded with a double border.

Another option would be to use different colors for the different types of cells. For example, data cells could appear blue, changing cells yellow, output cells white, and the target cell green.

Obviously, you may use any scheme that you like. The important thing is to be consistent, so that you can quickly recognize the types of cells. Then, when you want to examine the cells of a certain type, the shading or color will immediately guide you there.

Show the Entire Model on the Spreadsheet

The Solver uses a combination of the spreadsheet and the Solver dialogue box to specify the model to be solved. Therefore, it is possible to include certain elements of the model (such as the \leq , $=$, or \geq signs and/or the right-hand sides of the constraints) in the Solver dialogue box without displaying them in the spreadsheet. However, we strongly recommend that *every* element of the model be displayed *on the spreadsheet*. Every person using or adapting the model, or referring back to it later, needs to be able to interpret the

model. This is much easier to do by viewing the model on the spreadsheet than by trying to decipher it from the Solver dialogue box. Furthermore, a printout of the spreadsheet does not include information from the Solver dialogue box.

In particular, all the elements of a constraint should be displayed on the spreadsheet, even though the constraint will be enforced only after it is listed in the Solver dialogue box. For each constraint, three adjacent cells should be used for the total of the left-hand side, the \leq , $=$, or \geq sign in the middle, and the right-hand side. (Note in Fig. 21.5 that this was done in columns J, K, and L of the spreadsheet for the Everglade problem.) As mentioned earlier, the changing cells and target cell should be highlighted in some manner (e.g., with borders and/or cell shading and coloring). A good test is that you should not need to go to the Solver dialogue box to determine any element of the model. You should be able to identify the changing cells, the target cell, and all the constraints in the model just by looking at the spreadsheet.

A Poor Spreadsheet Model

It is certainly possible to set up a linear programming spreadsheet model without utilizing any of these ideas. Figure 21.7 shows an alternative spreadsheet formulation for the Everglade problem that violates nearly every one of these guidelines. This formulation can still be solved using Solver, which in fact yields the same optimal solution as in Fig. 21.5. However, the formulation has many problems. It is not clear which cells yield the solution (borders, shading,

FIGURE 21.7
A poor formulation of the spreadsheet model for the Everglade cash flow management problem.

	A	B	C	D	E	F
1	A Poor Formulation of the Everglade Cash Flow Problem					
2						
3			LT	ST	Ending	
4		Year	Loan	Loan	Balance	
5		2010	6.65	0.85	0.50	
6		2011		3.40	0.50	
7		2012		8.21	0.50	
8		2013		6.49	0.50	
9		2014		1.61	0.50	
10		2015		0	1.27	
11		2016		3.70	0.50	
12		2017		0	2.97	
13		2018		0	0.50	
14		2019		0	10.03	
15		2020			2.92	

Solver Parameters

Set Target Cell:

Equal To: Max Min

By Changing Cells:

Subject to the Constraints:

Solver Options

Assume Linear Model

Assume Non-Negative

	E
3	Ending
4	Balance
5	=1-8+C5+D5
6	=E5-2+D6-\$C\$5*(0.07)-D5*(1.1)
7	=E6-4+D7-\$C\$5*(0.07)-D6*(1.1)
8	=E7+3+D8-\$C\$5*(0.07)-D7*(1.1)
9	=E8+6+D9-\$C\$5*(0.07)-D8*(1.1)
10	=E9+3+D10-\$C\$5*(0.07)-D9*(1.1)
11	=E10-4+D11-\$C\$5*(0.07)-D10*(1.1)
12	=E11+7+D12-\$C\$5*(0.07)-D11*(1.1)
13	=E12-2+D13-\$C\$5*(0.07)-D12*(1.1)
14	=E13+10+D14-\$C\$5*(0.07)-D13*(1.1)
15	=E14+D15-\$C\$5*(1.07)-D14*(1.1)

or coloring are not used to highlight the changing cells and target cell). Without going to the Solver dialogue box, the constraints in the model cannot be identified (the spreadsheet does not show the entire model). The spreadsheet also does not show most of the data. For example, to determine the data used for the projected cash flows, the interest rates, or the starting balance, you need to dig into the formulas in column E (the data are not separate from the formulas). If any of these data change, the actual formulas need to be modified rather than simply changing a number on the spreadsheet. Furthermore, the formulas and the model in the Solver dialogue box are difficult to interpret (range names are not utilized).

Compare Figs. 21.5 and 21.7. Applying the guidelines for good spreadsheet models (as is done for Fig. 21.5) results in a model that is easier to understand, easier to debug, and easier to modify. This is especially important for models that will have a long life span. If this model is going to be reused months later, the “good” model of Fig. 21.5 immediately can be understood, modified, and reapplied as needed, whereas deciphering the spreadsheet model of Fig. 21.7 again would be a great challenge.

21.4 DEBUGGING A SPREADSHEET MODEL

No matter how carefully it is planned and built, even a moderately complicated model usually will not be error-free the first time it is run. Often the mistakes are immediately obvious and quickly corrected. However, sometimes an error is harder to root out. Following the guidelines in Sec. 21.3 for developing a good spreadsheet model can make the model *much* easier to debug. Even so, much like debugging a computer program, debugging a spreadsheet model can be a difficult task. This section presents some tips and a variety of Excel features that can make debugging easier.

As a first step in debugging a spreadsheet model, test the model using the principles discussed in the first subsection on testing in Sec. 21.2. In particular, try different values for the changing cells for which you can predict the correct result in the output cells and see if they calculate as expected. Values of 0 are good ones to try initially because usually it is then obvious what should be in the output cells. Try other simple values, such as all 1s, where the correct results in the output cells are reasonably obvious. For more complicated values, break out a calculator and do some manual calculations to check the various output cells. Include some very large values for the changing cells to ensure that the calculations are behaving reasonably for these extreme cases.

If you have defined range names, be sure that they still refer to the correct cells. Sometimes they can become disjointed when you add rows or columns to the spreadsheet. To test the range names, you can either select the various range names in the name box, which will highlight the selected range in the spreadsheet, or paste the entire list of range names and their references into the spreadsheet.

Carefully study each formula to be sure it is entered correctly. A very useful feature in Excel for checking formulas is the **toggle** to switch back and forth between viewing the formulas in the worksheet and viewing the resulting values in the output cells. By default, Excel shows the values that are calculated by the various output cells in the model. Typing control-~ on a PC (or command-~ on a Mac) switches the current worksheet to instead display the formulas in the output cells, as shown in Fig. 21.8. Typing control-~ again switches back to the standard view of displaying the values in the output cells (like Fig. 21.5).

Another useful set of features built into Excel are the **auditing tools**. The auditing tools are available in the Formula Auditing group of the Formulas Tab (for Excel 2007) or under the Tools\Auditing menu (for earlier versions of Excel).

The auditing tools can be used to graphically display which cells make direct links to a given cell. For example, selecting LTLoan (D11) in Fig. 21.5 and then Trace Dependents generates the arrows on the spreadsheet shown in Fig. 21.9.

21-20 CHAPTER 21 THE ART OF MODELING WITH SPREADSHEETS

	A	B	C	D	E	F	G	H	I	J	K
1	Cash Flow Management Problem										
2											
3		LT Rate	0.07								
4		ST Rate	0.1								
5											
6		Start Balance	1			(all cash figures in millions of dollars)					
7		Minimum Cash	0.5								
8											
9			Cash	LT	ST	LT	ST	LT	ST	Ending	
10			Flow	Loan	Loan	Interest	Interest	Payback	Payback	Balance	
11	2010	-8	6.64945	0.850540						=StartBalance+SUM(C11:I11)	>= =Mir
12	2011	-2		3.401057	=-LTRate*LTLoan	=-STRate*E11		=-E11		=J11+SUM(C12:I12)	>= =Mir
13	2012	-4		8.206625	=-LTRate*LTLoan	=-STRate*E12		=-E12		=J12+SUM(C13:I13)	>= =Mir
14	2013	3		6.492749	=-LTRate*LTLoan	=-STRate*E13		=-E13		=J13+SUM(C14:I14)	>= =Mir
15	2014	6		1.607486	=-LTRate*LTLoan	=-STRate*E14		=-E14		=J14+SUM(C15:I15)	>= =Mir
16	2015	3		0	=-LTRate*LTLoan	=-STRate*E15		=-E15		=J15+SUM(C16:I16)	>= =Mir
17	2016	-4		3.699159	=-LTRate*LTLoan	=-STRate*E16		=-E16		=J16+SUM(C17:I17)	>= =Mir
18	2017	7		0	=-LTRate*LTLoan	=-STRate*E17		=-E17		=J17+SUM(C18:I18)	>= =Mir
19	2018	-2		0	=-LTRate*LTLoan	=-STRate*E18		=-E18		=J18+SUM(C19:I19)	>= =Mir
20	2019	10		0	=-LTRate*LTLoan	=-STRate*E19		=-E19		=J19+SUM(C20:I20)	>= =Mir
21	2020				=-LTRate*LTLoan	=-STRate*E20	=-LTLoan	=-E20		=J20+SUM(C21:I21)	>= =Mir

■ FIGURE 21.8

The spreadsheet obtained by toggling the spreadsheet in Fig. 21.5 once to replace the values in the output cells by the formulas entered into these cells. Using the toggle feature in Excel once more will restore the view of the spreadsheet shown in Fig. 21.5.

You now can immediately see that LTLoan (D11) is used in the calculation of LT Interest for every year in column F, in the calculation of LTPayback (H21), and in the calculation of the balance in 2010 (J11). This can be very illuminating. Think about what output cells LTLoan should impact directly. There should be an arrow to each of these cells. If, for example, LTLoan is missing from any of the formulas in column F, the error will be immediately revealed by the missing arrow. Similarly, if LTLoan is mistakenly entered in any of the short-term loan output cells, this will show up as extra arrows.

You also can trace backward to see which cells provide the data for any given cell. These can be displayed graphically by choosing Trace Precedents. For example, choosing Trace Precedents for the ST Interest cell for 2011 (G12) displays the arrows shown in Fig. 21.10. These arrows indicate that the ST Interest cell for 2011 (G12) refers to the ST Loan in 2010 (E11) and to STRate (C4).

When you are done, choose Remove Arrows.

	A	B	C	D	E	F	G	H	I	J	K	L
1	Everglade Cash Flow Management Problem											
2												
3		LT Rate	7%									
4		ST Rate	10%									
5												
6		Start Balance	1	(all cash figures in millions of dollars)								
7		Minimum Cash	0.5									
8												
9			Cash	LT	ST	LT	ST	LT	ST	Ending	Minimum	
10		Year	Flow	Loan	Loan	Interest	Interest	Payback	Payback	Balance	Balance	
11		2010	-8	6.65	0.85					0.50	>=	0.5
12		2011	-2		3.40	-0.47	-0.09		-0.85	0.50	>=	0.5
13		2012	-4		8.21	-0.47	-0.34		-3.40	0.50	>=	0.5
14		2013	3		6.49	-0.47	-0.82		-8.21	0.50	>=	0.5
15		2014	6		1.61	-0.47	-0.65		-6.49	0.50	>=	0.5
16		2015	3		0	-0.47	-0.16		-1.61	1.27	>=	0.5
17		2016	-4		3.70	-0.47	0		0	0.50	>=	0.5
18		2017	7		0	-0.47	-0.37		-3.70	2.97	>=	0.5
19		2018	-2		0	-0.47	0		0	0.50	>=	0.5
20		2019	10		0	-0.47	0		0	10.03	>=	0.5
21		2020				-0.47	0	-6.65	0	2.92	>=	0.5

■ **FIGURE 21.9**

The spreadsheet obtained by using the Excel auditing tools to trace the dependents of the LT Loan value in cell D11 of the spreadsheet in Fig. 21.5.

	A	B	C	D	E	F	G	H	I	J	K	L
1	Everglade Cash Flow Management Problem											
2												
3		LT Rate	7%									
4		ST Rate	10%									
5												
6		Start Balance	1	(all cash figures in millions of dollars)								
7		Minimum Cash	0.5									
8												
9			Cash	LT	ST	LT	ST	LT	ST	Ending	Minimum	
10		Year	Flow	Loan	Loan	Interest	Interest	Payback	Payback	Balance	Balance	
11		2010	-8	6.65	0.85					0.50	>=	0.5
12		2011	-2		3.40	-0.47	-0.09		-0.85	0.50	>=	0.5
13		2012	-4		8.21	-0.47	-0.34		-3.40	0.50	>=	0.5
14		2013	3		6.49	-0.47	-0.82		-8.21	0.50	>=	0.5
15		2014	6		1.61	-0.47	-0.65		-6.49	0.50	>=	0.5
16		2015	3		0	-0.47	-0.16		-1.61	1.27	>=	0.5
17		2016	-4		3.70	-0.47	0		0	0.50	>=	0.5
18		2017	7		0	-0.47	-0.37		-3.70	2.97	>=	0.5
19		2018	-2		0	-0.47	0		0	0.50	>=	0.5
20		2019	10		0	-0.47	0		0	10.03	>=	0.5
21		2020				-0.47	0	-6.65	0	2.92	>=	0.5

■ **FIGURE 21.10**

The spreadsheet obtained by using the Excel auditing tools to trace the precedents of the ST Interest (2004) calculation in cell G12 of the spreadsheet in Fig. 21.5.

21.5 CONCLUSIONS

There is considerable art to modeling well with spreadsheets. This chapter focuses on providing a foundation for learning this art.

The general process of modeling in spreadsheets has four major steps: (1) plan the spreadsheet model, (2) build the model, (3) test the model, and (4) analyze the model and

its results. During the planning step, after defining the problem clearly and gathering the relevant data, it is helpful to begin by visualizing where you want to finish and then doing some calculations by hand to clarify the needed computations before starting to sketch out a logical layout for the spreadsheet. Then, when you are ready to undertake the building step, it is a good idea to start by building a small, readily manageable version of the model before expanding the model to full-scale size. This enables you to test the small version first to get all the logic straightened out correctly before expanding to a full-scale model and undertaking a final test. After completing all of this, you are ready for the analysis step, which involves applying the model to evaluate proposed solutions and perhaps using Solver to optimize the model.

Using this plan-build-test-analyze process should yield a spreadsheet model, but it doesn't guarantee that you will obtain a good one. Section 21.3 describes in detail the following guidelines for building "good" spreadsheet models.

- Enter the data first.
- Organize and clearly identify the data.
- Enter each piece of data into one cell only.
- Separate data from formulas.
- Keep it simple.
- Use range names.
- Use relative and absolute references to simplify copying formulas.
- Use borders, shading, and colors to distinguish between cell types.
- Show the entire model on the spreadsheet.

Even if all these guidelines are followed, a thorough debugging process may be needed to eliminate the errors that lurk within the initial version of the model. It is important to check whether the output cells are giving correct results for various values of the changing cells. Other items to check include whether range names refer to the appropriate cells and whether formulas have been entered into output cells correctly. Excel provides a number of useful features to aid in the debugging process. One is the ability to toggle the worksheet between viewing the results in the output cells and the formulas entered into those output cells. Several other helpful features are available from Excel's auditing tools.

■ SELECTED REFERENCES

1. Albright, S. C., and W. L. Winston: *Spreadsheet Modeling and Applications: Essentials of Practical Management Science*, South-Western College Publishing, Mason, OH, 2005.
2. Denardo, E. Y.: *The Science of Decision Making: A Problem-Based Approach Using Excel*, Wiley, New York, 2002.
3. Hillier, F. S., and M. S. Hillier, *Introduction to Management Science: A Modeling and Case Studies Approach with Spreadsheets*, 3rd ed., McGraw-Hill/Irwin, Burr Ridge, IL, 2008.
4. Powell, S. G., and K. R. Baker: *Management Science: The Art of Modeling with Spreadsheets*, 2nd ed., Wiley, New York, 2008.
5. Ragsdale, C. T., *Spreadsheet Modeling and Decision Analysis*, 5th ed., South-Western College Publishing, Mason, OH, 2007.
6. Winston, W. L., and S. C. Albright, *Practical Management Science*, 3rd ed., South-Western College Publishing, Mason, OH, 2007.

■ LEARNING AIDS FOR THIS CHAPTER ON THIS WEBSITE

Chapter 21 Excel Files:

Everglade Case Study

Wyndor Example
 Everglade Problem 21-9
 Everglade Problem 21-10

An Excel Add-in:

Premium Solver for Education

■ **PROBLEMS**

We have inserted the symbol E* (for Excel) to the left of each problem or part where Excel should be used.

E* **21-1.** Consider the Everglade cash flow problem discussed in this chapter. Suppose that extra cash is kept in an interest-bearing savings account. Assume that any cash left at the end of a year earns 3 percent interest the following year. Make any necessary modifications to the spreadsheet and re-solve.

21-2. The Pine Furniture Company makes fine country furniture. The company's current product lines consist of end tables, coffee tables, and dining room tables. The production of each of these tables requires 8, 15, and 80 pounds of pine wood, respectively. The tables are handmade, and require one hour, two hours, and four hours, respectively. Each table sold generates \$50, \$100, and \$220 profit, respectively. The company has 3,000 pounds of pine wood and 200 hours of labor available for the coming week's production. The chief operating officer (COO) has asked you to do some spreadsheet modeling with these data to analyze what the product mix should be for the coming week and make a recommendation.

- (a) Visualize where you want to finish. What numbers will the COO need? What are the decisions that need to be made? What should the objective be?
- (b) Suppose that Pine Furniture were to produce three end tables and three dining room tables. Calculate by hand the amount of pine wood and labor that would be required, as well as the profit generated from sales.
- (c) Make a rough sketch of a spreadsheet model, with blocks laid out for the data cells, changing cells, output cells, and target cell.
- E* (d) Build a spreadsheet model and then solve it.

21-3. Reboot, Inc. is a manufacturer of hiking boots. Demand for boots is highly seasonal. In particular, the demand in the next year is expected to be 3,000, 4,000, 8,000, and 7,000 pairs of boots in quarters 1, 2, 3, and 4, respectively. With its current production facility, the company can produce at most 6,000 pairs of boots in any quarter. Reboot would like to meet all the expected demand, so it will need to carry inventory to meet demand in the later quarters. Each pair of boots sold generates a profit of \$20 per pair. Each pair of boots in inventory at the end of a quarter incurs \$8 in storage and capital recovery costs. Reboot has 1,000 pairs of boots in inventory at the start of quarter 1. Reboot's top management has given you the assignment of doing some spreadsheet modeling to analyze what the production schedule should be for the next four quarters and make a recommendation.

- (a) Visualize where you want to finish. What numbers will top management need? What are the decisions that need to be made? What should the objective be?
- (b) Suppose that Reboot were to produce 5,000 pairs of boots in each of the first two quarters. Calculate by hand the ending inventory, profit from sales, and inventory costs for quarters 1 and 2.
- (c) Make a rough sketch of a spreadsheet model, with blocks laid out for the data cells, changing cells, output cells, and target cell.
- E* (d) Build a spreadsheet model for quarters 1 and 2, and then thoroughly test the model.
- E* (e) Expand the model to full scale and then solve it.

E* **21-4.** The Fairwinds Development Corporation is considering taking part in one or more of three different development projects—A, B, and C—that are about to be launched. Each project requires a significant investment over the next few years, and then would be sold upon completion. The projected cash flows (in millions of dollars) associated with each project are shown in the table below.

Year	Project A	Project B	Project C
1	-4	-8	-10
2	-6	-8	-7
3	-6	-4	-7
4	24	-4	-5
5	0	30	-3
6	0	0	44

Fairwinds has \$10 million available now and expects to receive \$6 million from other projects by the end of each year (1 through 6) that would be available for the ongoing investments the following year in projects A, B, and C. By acting now, the company may participate in each project either fully, fractionally (with other development partners), or not at all. If Fairwinds participates at less than 100 percent, then all the cash flows associated with that project are reduced proportionally. Company policy requires ending each year with a cash balance of at least \$1 million. Your assignment is to formulate a spreadsheet model to analyze the problem.

- (a) Visualize where you want to finish. What numbers are needed? What are the decisions that need to be made? What should the objective be?
- (b) Suppose that Fairwinds were to participate in Project A fully and in Project C at 50 percent. Calculate by hand what the ending cash position would be after year 1 and year 2.

21-24

CHAPTER 21 THE ART OF MODELING WITH SPREADSHEETS

- (c) Make a rough sketch of a spreadsheet model, with blocks laid out for the data cells, changing cells, output cells, and target cell.
- E* (d) Build a spreadsheet model for years 1 and 2, and then thoroughly test the model
- E* (e) Expand the model to full scale, and then solve it.

21-5. Refer to the scenario described in Prob. 3.4-10 (Chap. 3), but ignore the instructions given there. Focus instead on using spreadsheet modeling to address Web Mercantile’s problem by doing the following.

- (a) Visualize where you want to finish. What numbers will Web Mercantile require? What are the decisions that need to be made? What should the objective be?
- (b) Suppose that Web Mercantile were to lease 30,000 square feet for all five months and then 20,000 additional square feet for the last three months. Calculate the total costs by hand.
- (c) Make a rough sketch of a spreadsheet model, with blocks laid out for the data cells, changing cells, output cells, and target cell.
- E* (d) Build a spreadsheet model for months 1 and 2, and then thoroughly test the model.
- E* (e) Expand the model to full scale, and then solve it.

21-6. Refer to the scenario described in Prob. 3.4-11 (Chap. 3), but ignore the instructions given there. Focus instead on using spreadsheet modeling to address Larry Edison’s problem by doing the following.

- (a) Visualize where you want to finish. What numbers will Larry require? What are the decisions that need to be made? What should the objective be?
- (b) Suppose that Larry were to hire three full-time workers for the morning shift, two for the afternoon shift, and four for the

evening shift, as well as hire three part-time workers for each of the four shifts. Calculate by hand how many workers would be working at each time of the day and what the total cost would be for the entire day.

- (c) Make a rough sketch of a spreadsheet model, with blocks laid out for the data cells, changing cells, output cells, and target cell.
- E* (d) Build a spreadsheet model and then solve it.

21-7. Refer to the scenario described in Prob. 3.4-13 (Chap.3), but ignore the instructions given there. Focus instead on using spreadsheet modeling to address Al Ferris’ problem by doing the following.

- (a) Visualize where you want to finish. What numbers will Al require? What are the decisions that need to be made? What should the objective be?
- (b) Suppose that Al were to invest \$20,000 each in investment A (year 1), investment B (year 2), and investment C (year 2). Calculate by hand what the ending cash position would be after each year.
- (c) Make a rough sketch of a spreadsheet model, with blocks laid out for the data cells, changing cells, output cells, and target cell.
- E* (d) Build a spreadsheet model for years 1 through 3, and then thoroughly test the model.
- E* (e) Expand the model to full scale, and then solve it.

21-8. In contrast to the spreadsheet model for the Wyndor Glass Co. product-mix problem shown in Fig. 21.6, the spreadsheet given next is an example of a poorly formulated spreadsheet model for this same problem. Identify each of the guidelines in Sec. 21.3 that is violated by this poor model. In each case, explain how it violates the guideline and why the model in Fig. 21.6 does a much better job of following the guideline.

	A	B	C	D
1	Wyndor Glass Co. (Poor Formulation)			
2				
3		Batches of Doors Produced	2	
4		Batches of Windows Produced	6	
5		Hours Used (Plant 1)	2	
6		Hours Used (Plant 2)	12	
7		Hours Used (Plant 3)	18	
8		Total Profit	\$36,000	

Solver Parameters

Set Target Cell:

Equal To: Max Min

By Changing Variable Cells:

Subject to the Constraints:

-
-
-

	B	C
5	Hours Used (Plant 1)	=1*C3+0*C4
6	Hours Used (Plant 2)	=0*C3+2*C4
7	Hours Used (Plant 3)	=3*C3+2*C4
8	Total Profit	=3000*C3+5000*C4

Solver Options

- Assume Linear Model
- Assume Non-Negative

E* 21-9. Refer to the spreadsheet file named “Everglade Problem 21-9” contained in the Excel files for this chapter on the book’s website. This file contains a formulation of the Everglade problem considered in this chapter. However, there are three errors in this formulation. Use the ideas presented in Sec. 21.4 for debugging a spreadsheet model to find the errors. In particular, try different trial values for which you can predict the correct results, use the toggle to examine all the formulas, and use the auditing toolbar to check precedence and dependence relationships among the various changing cells, data cells, and output cells. Describe the errors found and how you found them.

E* 21-10. Refer to the spreadsheet file named “Everglade Problem 21-10” contained in the Excel files for this chapter on the book’s website. This file contains a formulation of the Everglade problem considered in this chapter. However, there are three errors in this formulation. Use the ideas presented in Sec. 21.4 for debugging a spreadsheet model to find the errors. In particular, try different trial values for which you can predict the correct results, use the toggle to examine all the formulas, and use the auditing toolbar to check precedence and dependence relationships among the various changing cells, data cells, and output cells. Describe the errors found and how you found them.

■ **CASES**

CASE 21.1 Prudent Provisions for Pensions

Among its many financial products, the Prudent Financial Services Corporation (normally referred to as PFS) manages a well-regarded pension fund that is used by a number of companies to provide pensions for their employees. PFS’s management takes pride in the rigorous professional standards used in operating the fund. Since the Enron collapse in late 2001 and the subsequent tightening of federal and state regulations for operating pension funds, PFS has redoubled its efforts to provide prudent management of the fund. It is now December 2009. The total pension payments that will need to be made by the fund over the next 10 years are shown in the table below.

Year	Pension Payments (\$ millions)
2010	8
2011	12
2012	13
2013	14
2014	16
2015	17
2016	20
2017	21
2018	22
2019	24

By using interest as well, PFS currently has enough liquid assets to meet all these pension payments. Therefore, to safeguard the pension fund, PFS would like to make a number of investments whose payouts would match the pension payments over the next 10 years. The only investments that PFS trusts for the pension fund are a money market fund and bonds. The money market fund pays an annual interest rate of 5 percent. The characteristics of each unit of the four bonds under consideration are shown in the next table.

	Current Price	Coupon Rate	Maturity Date	Face Value
Bond 1	\$980	4%	Jan. 1, 2011	\$1,000
Bond 2	920	2	Jan. 1, 2013	1,000
Bond 3	750	0	Jan. 1, 2015	1,000
Bond 4	800	3	Jan. 1, 2018	1,000

All of these bonds will be available for purchase on January 1, 2010, in as many units as desired. The coupon rate is the percentage of the face value that will be paid in interest on January 1 of each year, starting one year after purchase and continuing until (and including) the maturity date. Thus, these interest payments on January 1 of each year are in time to be used toward the pension payments for that year. Any excess interest payments will be deposited into the money market fund. To be conservative in its financial planning, PFS assumes that all the pension payments for the year occur at the beginning of the year immediately after these interest payments (including a year’s interest from the money market fund) are received. The entire face value of a bond also will be received on its maturity date. Since the current price of each bond is less than its face value, the actual yield of the bond exceeds its coupon rate. Bond 3 is a zero-coupon bond, so it pays no interest but instead pays a face value on the maturity date that greatly exceeds the purchase price.

PFS would like to make the smallest possible investment (including any deposit into the money market fund) on January 1, 2010, to cover all its required pension payments through 2019. Some spreadsheet modeling needs to be done to see how to do this.

- (a) Visualize where you want to finish. What numbers are needed by PFS management? What are the decisions that need to be made? What should the objective be?
- (b) Suppose that PFS were to invest \$28 million in the money market fund and purchase 10,000 units each of bond 1 and bond 2



on January 1, 2010. Calculate by hand the payments received from bonds 1 and 2 on January 1 of 2011 and 2012. Also calculate the resulting balance in the money market fund on January 1 of 2010, 2011, and 2012 after receiving these payments, making the pension payments for the year, and depositing any excess into the money market fund.

- (c) Make a rough sketch of a spreadsheet model, with blocks laid out for the data cells, changing cells, output cells, and target cell.
- (d) Build a spreadsheet model for years 2010 through 2012, and then thoroughly test the model.
- (e) Expand the model to consider all years through 2019, and then solve it.

■ ACKNOWLEDGEMENT

This chapter (with slight differences) also appears as Chapter 4 in the 3rd edition of *Introduction to Management Science: A Modeling and Case Studies Approach with Spreadsheets* by Frederick S. Hillier and Mark S. Hillier, McGraw-Hill/Irwin, 2008. We gratefully acknowledge the major role that Mark S. Hillier played in developing this chapter.

