

# Chapter 1

---

## Introduction to Spreadsheet Models for Optimization

This is a book about optimization with an emphasis on building models and using spreadsheets. Each facet of this theme—models, spreadsheets, and optimization—has a role in defining the emphasis of our coverage.

A *model* is a simplified representation of a situation or problem. Models attempt to capture the essential features of a complicated situation so that it can be studied and understood more completely. In the worlds of business, engineering, and science, models aim to improve our understanding of practical situations. Models can be built with tangible materials, or words, or mathematical symbols and expressions. A *mathematical model* is a model that is constructed—and also analyzed—using mathematics. In this book, we focus on mathematical models. Moreover, we work with *decision models*, or models that contain representations of decisions. The term also refers to models that support decision-making activities.

A *spreadsheet* is a row-and-column layout of text, numerical data, and logical information. The spreadsheet version of a model contains the model's elements, linked together by specific logical information. Electronic spreadsheets, like those built using Microsoft Excel<sup>®</sup>, have become familiar tools in the business, engineering, and scientific worlds. Spreadsheets are relatively easy to understand, and people often rely on spreadsheets to communicate their analyses. In this book, we focus on the use of spreadsheets to represent and analyze mathematical models.

This text is written for an audience that already has some familiarity with Excel. Our coverage assumes a level of facility with Excel comparable to a beginner's level. Someone who has used other people's spreadsheets and built simple spreadsheets for some purpose—either personal or organizational—has probably developed this skill level. Box 1.1 describes the Excel skill level assumed. Readers without this level of background are encouraged to first work through some introductory materials, such as the books by McFedries (1) and Reding and Wermers (2).

*Optimization* is the process of finding the best values of the variables for a particular criterion or, in our context, the best decisions for a particular measure of performance. The elements of an optimization problem are a set of decisions, a criterion, and

**BOX 1.1** *Excel Skills Assumed as Background for this Book*

Navigating in workbooks, worksheets, and windows.  
Using the cursor to select cells, rows, columns, and noncontiguous cell ranges.  
Entering text and data; copying and pasting; filling down or across.  
Formatting cells (number display, alignment, font, border, and protection).  
Editing cells (using the formula bar and cell-edit capability [F2]).  
Entering formulas and using the function wizard.  
Using relative and absolute addresses.  
Using range names.  
Creating charts and graphs.

perhaps a set of required conditions, or *constraints*, that the decisions must satisfy. These elements lend themselves to description in a mathematical model. The term optimization sometimes refers specifically to a procedure that is implemented by software. However, in this book, we expand that perspective to include the model-building process as well as the process of finding the best decisions.

Not all mathematical models are optimization models. Some models merely describe the logical relationship between inputs and outputs. Optimization models are a special kind of model in which the purpose is to find the best value of a particular output measure and the choices that produce it. Optimization problems abound in the real world, and if we're at all ambitious or curious, we often find ourselves seeking solutions to those problems. Business firms are very interested in optimization because making good decisions helps a firm run efficiently, perform profitably, and compete effectively. In this book, we focus on optimization problems expressed in the form of spreadsheet models and solved using a spreadsheet-based approach.

## 1.1. ELEMENTS OF A MODEL

To restate our premise, we are interested in mathematical models. Specifically, we are interested in two forms—algebraic and spreadsheet models. In the former, we use algebraic notation to represent elements and relationships, and in the latter, we use spreadsheet entries and structure. For example, in an algebraic statement, we might use the variable  $x$  to represent a quantitative decision, and we might use some function  $f(x)$  to represent the measure of performance that results from choosing decision  $x$ . Then we might adopt the letter  $z$  to represent a criterion for decision making and construct the equation  $z = f(x)$  to guide the choice of a decision. Algebra is the basic language of analysis largely because it is precise and compact.

As an introductory modeling example, let's consider the price decision in the scenario of Example 1.1.

**EXAMPLE 1.1** *Price, Demand, and Profit*

Our firm's production department has carried out a cost accounting study and found that the unit cost for one of its main products is \$40. Meanwhile, the marketing department has estimated the relationship between price and sales volume (the so called *demand curve* for the product) as follows:

$$y = 800 - 5x \quad (1.1)$$

where  $y$  represents quarterly demand and  $x$  represents the selling price per unit. We wish to determine a selling price for this product, given the information available. ■

In Example 1.1, the decision is the unit price, and the consequence of that decision is the level of demand. The demand curve in Equation 1.1 expresses the relationship of demand and price in algebraic terms. Another equation expresses the calculation of profit contribution, by multiplying the demand  $y$  by the unit profit contribution ( $x - 40$ ) on each item

$$z = (x - 40)y \quad (1.2)$$

where  $z$  represents our product's quarterly profit contribution.

We can substitute Equation 1.1 into 1.2 if we want to write  $z$  algebraically as a function of  $x$  alone. As a result, we can express the profit contribution as

$$z = 1000x - 5x^2 - 32,000 \quad (1.3)$$

This step embodies the algebraic principle that simplification is always desirable. Here, simplification reduces the number of variables in the expression for profit contribution. Simplification, however, is not necessarily a virtue when we use a spreadsheet model.

Example 1.1, simple as it is, has some important features. First, our model contains three numerical inputs: 40 (the unit cost),  $-5$  (the marginal effect of price on demand) and 800 (the maximum demand). Numerical inputs such as these are called *parameters*. In some models, parameters correspond to raw data, but in many cases, parameters are summaries drawn from a more primitive data set. They may also be estimates made by a knowledgeable party, forecasts derived from statistical analyses, or predictions chosen to reflect a future scenario.

Our model also contains a decision—an unknown quantity yet to be determined. In traditional algebraic formulations, unknowns are represented as variables. Quantitative representations of decisions are therefore called *decision variables*. The decision variable in our model is the unit price  $x$ .

Our model contains the equation that relates demand to price. We can think of this relationship as part of the model's logic. In that role, the demand curve prescribes a relationship between two variables—price and demand—that must always hold. Thus, in our model, the only admissible values of  $x$  and  $y$  are those that satisfy Equation 1.1.

Finally, our model contains a calculation of quarterly profit contribution, which is the performance measure of interest and a quantity that we wish to maximize. This output variable measures the consequence of selecting any particular price decision

in the model. In optimization models, we are concerned with maximizing or minimizing some measure of performance, expressed as a mathematical function, and we refer to it as the *objective function*, or simply the *objective*.

## 1.2. SPREADSHEET MODELS

Algebra is an established language that works well for describing problems, but not always for obtaining solutions. Algebraic solutions tend to occur in formulas, not numbers, but numbers most often represent decisions in the practical world. By contrast, spreadsheets represent a practical language—one that works very effectively with numbers. Like algebraic models, spreadsheets can be precise and compact, but there are also complications that are unique to spreadsheets. For example, there is a difference between form and content in a spreadsheet. Two spreadsheets may look the same in terms of layout and the numbers displayed on a computer screen, but the underlying formulas in corresponding cells could differ. Because the information behind the display can be different even when two spreadsheets have the same on-screen appearance, we can't always tell the logical content from the form of the display. Another complication is the lack of a single, well accepted way to build a spreadsheet representation of a given model. In an optimization model, we want to represent decision variables, an objective function, and constraints. However, that still leaves a lot of flexibility in choosing how the logic of a particular model is incorporated into a spreadsheet. Such flexibility would ordinarily be advantageous if the only use of a spreadsheet were to help individuals solve problems. However, spreadsheets are perhaps even more important as vehicles for communication. When we use spreadsheets in this role, flexibility can sometimes lead to confusion and disrupt the intended communication.

We will try to mitigate these complications with some design guidelines. For example, it is helpful to create separate modules in the spreadsheet for decision variables, objective function, and constraints. To the extent that we follow such guidelines, we may lose some flexibility in building a spreadsheet model. Moving the design process toward standardization will, however, make the content of a spreadsheet more understandable from its form, so differences between form and content become less problematic.

With optimization, a spreadsheet model contains the analysis that ultimately provides decision support. For this reason, the spreadsheet model should be intelligible to its users, not just to its developer. On some occasions, a spreadsheet might come into routine use in an organization, even when the developer moves on. New analysts may inherit the responsibilities associated with the model, so it is vital that they, too, understand how the spreadsheet works. For that matter, the decision maker may also move on. For the organization to retain the learning that has taken place, successive decision makers must also understand the spreadsheet. In yet another scenario, the analyst develops a model for one-time use but then discovers a need to reuse it several months later in a different context. In such a situation, it's important that the analyst understands the original model, lest the passage of time obscure its purpose and

logic. In all of these cases, the spreadsheet model fills a significant communications need. Thus, it is important to keep the role of communication in mind while developing a spreadsheet.

A spreadsheet version of our pricing model might look like the one in Figure 1.1. This spreadsheet contains a cell (C9) that holds the unit price, a cell (C12) that holds the level of demand, and a cell (C15) that holds the total profit contribution. Actually, cell C12 holds Equation 1.1 in the form of the Excel formula  $=C4+C5*C9$ . Similarly, cell C15 holds Equation 1.2 with the formula  $=(C9-C6)*C12$ . In cell C9, the unit price is initially set to \$80. For this choice, demand is 400. The quarterly profit contribution is \$16,000.

In a spreadsheet model, there is usually no premium on being concise, as there is when we use algebra. In fact, when conciseness begins to interfere with a model's transparency, it becomes undesirable. Thus, in Figure 1.1, the model retains the demand equation and displays the demand quantity explicitly; we have not tried to incorporate Equation 1.3. This form allows a user to see how price influences profit contribution through demand because all of these quantities are explicit. Furthermore, it is straightforward to trace the connection between the three input parameters and the calculation of profit contribution.

To summarize, our model consists of three parameters and a decision variable, together with some intermediate calculations, all leading to an objective function that we want to maximize. In algebraic terms, the model consists of Equations 1.1 and 1.2, with the prescription that we want to maximize Equation 1.2. In spreadsheet terms, the model consists of the spreadsheet in Figure 1.1, with the prescription that we want to maximize the value in cell C15.

	A	B	C
1	<b>Price, Demand, and Profit</b>		
2			
3	<b>inputs</b>		
4		<i>Max. demand</i>	800
5		<i>Slope</i>	-5
6		<i>Unit cost</i>	40
7			
8	<b>Decision</b>		
9		<i>Price</i>	\$ 80
10			
11	<b>Calculation</b>		
12		<i>Demand</i>	400
13			
14	<b>Outcome</b>		
15		<i>Profit</i>	\$ 16,000
16			

Figure 1.1. Spreadsheet model for determining price.

6 Chapter 1 Introduction to Spreadsheet Models for Optimization

The spreadsheet is organized into four modules: Inputs, Decision, Calculation, and Outcome, separating different kinds of information. In spreadsheet models, it is a good idea to separate input data from decisions and decisions from outcome measures. Intermediate calculations that do not lead directly to the outcome measure should also be kept separate.

In the spreadsheet model, cell borders and shading draw attention to the decision (cell C9) and the objective (cell C15) as the two most important elements of the optimization model. No matter how complicated a spreadsheet model may become, we want the decisions and the objective to be located easily by someone who looks at the display.

In the spreadsheet of Figure 1.1, the input parameters appear explicitly. It would not be difficult to skip the Inputs section entirely and express the demand function in cell C12 with the formula  $=800 - 5 * C9$ , or to express the profit contribution in cell C15 with the formula  $=(C9 - 40) * C12$ . This approach, however, places the numerical parameters in formulas, so a user would not see them at all when looking at the spreadsheet. Good practice calls for displaying parameters explicitly in the spreadsheet, as we have done in Figure 1.1, rather than burying them in formulas.

The basic version of our model, shown in Figure 1.1, is ready for optimization. But let's look at an alternative, shown in Figure 1.2. This version contains the four modules, and the numerical inputs are explicit but placed differently than in Figure 1.1. The main difference is that demand is treated as a decision variable, and the demand curve is expressed as an explicit constraint. Specifically, this form of the model treats both price and demand as variables in cells C9:C10, as if the two choices could be made arbitrarily. However, the Constraints module describes a relationship between the two variables in the form of Equation 1.1, which can

	A	B	C	D	E
1	<b>Price, Demand, and Profit</b>				
2					
3	<b>Inputs</b>				
4		Max. demand	800		
5		Slope	-5		
6		Cost	40		
7					
8	<b>Decisions</b>				
9		Price	\$ 80		
10		Demand	250		
11					
12	<b>Constraints</b>				
13		Demand	650	=	800
14					
15	<b>Outcomes</b>				
16		Profit	\$ 10,000		
17					

Figure 1.2. Alternative spreadsheet model for determining price.

equivalently be expressed as

$$y + 5x = 800 \quad (1.4)$$

We can meet this constraint by forcing cell C13 to equal cell E13, a condition that does not yet hold in Figure 1.2. Cell C13 contains the formula on the left-hand side of Equation 1.4, and cell E13 contains a reference to the parameter 800. The equals sign between them, in cell D13, signifies the nature of the constraint relationship to someone who is looking at the spreadsheet and trying to understand its logic. Equation 1.4 collects all the terms involving decision variables on the left-hand side (in cell C13) and places the constant term on the right-hand side (in cell E13). This is a standard form for expressing a constraint in a spreadsheet model. The spreadsheet itself displays, but does not actually enforce, this constraint. The enforcement task is left to the optimization software. Once the constraint is met, the corresponding decisions are called *feasible*.

This is a good place to include a reminder about the software that accompanies this book. The software contains important files and programs. In terms of files, the book's website<sup>1</sup> contains all of the spreadsheets shown in the figures. Figures 1.1 and 1.2, for example, can be found in the file that contains the spreadsheets for Chapter 1. Those files should be loaded, or else built from scratch, before continuing with the text. As we proceed through the chapters, the reader is welcome to load each file that appears in a figure, for hands-on examination.

### 1.3. A HIERARCHY FOR ANALYSIS

Before we proceed, some background on the development of models in organizations may be useful. Think about the person who builds a model as an *analyst*, someone who provides support to a decision maker or *client*. (In some cases, the analyst and the client are the same.) The development, testing, and application of a model constitute support for the decision maker—a service to the client. The application phase of this process includes some standard stages of model use.

When a model is built as an aid to decision making, the first stage often involves building a prototype, or a series of prototypes, leading to a model that the analyst and the client accept as a usable decision-support tool. That model provides quantitative analysis of a base-case scenario. In Example 1.1, suppose we set a tentative price of \$80. This price might be called a *base case*, in the sense that it represents a tentative decision. As we have seen, this price leads to demand of 400 and profit contribution of \$16,000.

After establishing a base case, it is usually appropriate to investigate the answers to a number of “what-if” questions. We ask, what if we change a numerical input or a decision in the model—what impact would that change have? Suppose, for example, that the marginal effect of price on demand (the slope of the demand curve) were  $-4$  instead of  $-5$ . What difference would this make? Retracing our algebraic steps, or

<sup>1</sup>The URL for the book's website is <http://mba.tuck.dartmouth.edu/opt/>

revising the spreadsheet in Figure 1.1, we can determine that the profit contribution would be \$19,200.

Systematic investigations of this kind are called *sensitivity analyses*. They explore how sensitive the results and conclusions are to changes in assumptions. Typically, we start by varying one assumption at a time and tracing the impact. Then we might try varying two or more assumptions, but such probing can quickly become difficult to follow. Therefore, most sensitivity analyses are performed one assumption at a time. Sometimes, it is useful to explore the what-if question in reverse. That is, we might ask, for the result to attain a given outcome level, what would the numerical input have to be? For example, starting with the base-case model, we might ask, what price would generate a profit contribution of \$17,000? We can answer this question algebraically, by setting  $z = 17,000$  in Equation 1.3 and solving for  $x$ , or, with the spreadsheet model, we can invoke Excel's Goal Seek tool to discover that the price would have to be about \$86.

Sensitivity analyses are helpful in determining the robustness of the results and any risks that might be present. They can also reveal how to achieve improvement from better choices in decision making. However, locating improvements this way is something of a trial-and-error process, and trial-and-error probing is inefficient. Faster and more reliable ways of locating improvements are available. Moreover, with trial-and-error approaches, we seldom know how far improvements can potentially reach, so a best outcome could exist that we never detect.

From this perspective, optimization can be viewed as a sophisticated form of sensitivity analysis that seeks the best values for the decisions and the best value for the performance measure. Optimization takes us beyond mere improvement; we look for the very best outcome in our model: the maximum possible benefit or the minimum possible cost. If we have constraints in our model, then optimization also tells us which of those conditions ultimately limit what we want to accomplish. Optimization can also reveal what we might gain if we can find a way to overcome those constraints and proceed beyond the limitations they impose.

## 1.4. OPTIMIZATION SOFTWARE

Optimization procedures find the best values of the decision variables in a given model. In the case of Excel, the optimization software is known as *Solver*, which is a standard tool available on the Data ribbon. (The generic term *solver* often refers to optimization software, whether or not it is implemented in a spreadsheet.) Optimization tools have been available on computers for several decades, prior to the widespread use of electronic spreadsheets. Before spreadsheets became popular, optimization was available as stand-alone software; it relied on an algebraic approach, but it was often accessible only by technical experts. Decision makers and even their analysts had to rely on those experts to build and solve optimization models. Spreadsheets, if they were used at all, were limited to small examples. Now, however, the spreadsheet allows decision makers to develop their own models, without having to learn specialized software, and to find optimal solutions for those models using Solver. Two trends account for the popularity of spreadsheet optimization. First,

familiarity with spreadsheets has become almost ubiquitous, at least in the business world. The spreadsheet has come to represent a common language for analysis. Second, the software packages available for spreadsheet-based optimization now include some of the most powerful tools available. The spreadsheet platform need not be an impediment to solving practical optimization problems.

Spreadsheet-based optimization has several advantages. The spreadsheet allows model inputs to be documented clearly and systematically. Moreover, if it is necessary to convert raw data into other forms for the purposes of setting up a model, the required calculations can be performed and documented conveniently in the same spreadsheet, or at least on another sheet in the same workbook. This allows integration between raw data and model data. Without this integration, errors or omissions are more likely, and maintenance becomes more difficult. Another advantage is algorithmic flexibility: The spreadsheet has the ability to call on several different optimization procedures, but the process of preparing the model is mostly the same no matter which procedure is applied. Finally, spreadsheet models have a certain amount of intrinsic credibility because spreadsheets are now so widely used for other purposes. Although spreadsheets can contain errors (and often do), there is at least some comfort in knowing that logic and discipline must be applied in the building of a spreadsheet.

Table 1.1 summarizes and compares the advantages of spreadsheet and algebraic (3,4) software approaches to optimization problems. The main advantage of algebraic approaches is the efficiency with which models can be specified. With spreadsheets, the elements of a model are represented explicitly. Thus, if the model requires a thousand variables, then the model builder must designate a thousand cells to hold their respective values. Algebraic codes use a different method. If a model contains a thousand variables, the code might refer to  $x(k)$ , with a specification that  $k$  may take on values from 1 to 1000, but  $x(k)$  need not be represented explicitly for each of the thousand values.

A second advantage of algebraic approaches is the fact that they can sometimes be tailored to a particular application. For example, the very large crew-scheduling applications used by airlines exhibit a special structure. To exploit this structure in the solution procedure, algebraic codes are sometimes enhanced with specialized subroutines that add solution efficiencies when solving a crew-scheduling problem.

A disadvantage of using spreadsheets is that they are not always transparent. As noted earlier, the analyst has a lot of flexibility in the layout and organization of a spreadsheet, but this flexibility, taken too far, may detract from effective communication. In this book, we try to promote better communication by suggesting

**Table 1.1.** Advantages of Spreadsheet and Algebraic Solution Approaches

Spreadsheet approaches	Algebraic approaches
Several algorithms available in one place	Large problem sizes accommodated
Integration of raw data and model data	Concise model specification
Flexibility in layout and design	Standardized model description
Ease of communication with nonspecialists	Enhancements possible for special cases
Intrinsic credibility	

standard forms for particular types of models. By using some standardization, we make it easier to understand and debug someone else's model. Algebraic codes usually have very detailed specifications for model format, so once we're familiar with the specifications, we should be able to read and understand anyone else's model.

In brief, commercially available algebraic solvers represent an alternative to spreadsheet-based optimization. In this book, our focus on a spreadsheet approach allows the novice to learn basic concepts of mathematical programming, practice building optimization models, obtain solutions readily, and interpret and apply the results of the analysis. All these skills can be developed in the accessible world of spreadsheets. Moreover, these skills provide a solid foundation for using algebraic solvers at some later date, when and if the situation demands it.

## 1.5. USING SOLVER

Purchasers of this book may download a powerful software package called *Risk Solver Platform (RSP)* that was developed by the same team that created Excel's Solver and that accommodates all Excel Solver models. (Before continuing with the text, the reader should install the software by following the guidelines and instructions in Appendix 1.) RSP is an integrated software package that includes more than just optimization capabilities, but this book focuses on optimization. Hence, the installation instructions recommend setting this software to operate in *Premium Solver Platform mode*, which exposes all of the optimization features, but hides other features such as Monte Carlo simulation and decision trees. Once the software is installed, a new Risk Solver Platform tab appears in Excel, with its own ribbon of commands. Under Premium Solver Platform mode, a Premium Solver Platform tab appears instead. In addition, the Add-Ins tab contains a Premium Solver choice which displays a Solver Parameters dialog that closely resembles the standard Excel Solver but uses the more powerful optimization capabilities of RSP. For our purposes, these tabs contain the equivalent optimization capabilities, and we may refer to either one.

In the remainder of this book, we assume the use of RSP, but we refer to it simply as Solver. The book covers its four main optimization procedures:

- The nonlinear solver
- The linear solver
- The integer solver
- The evolutionary solver

As in all matters involving software, the user should be aware of the copyright privileges and restrictions that apply to the use of RSP.

In order to illustrate the use of Solver, we return to Example 1.1. The optimization problem is to find a price that maximizes quarterly profit contribution. An algebraic statement of the problem is

$$\begin{array}{ll} \text{Maximize} & z = (x - 40)y \quad (\text{objective}) \\ \text{subject to} & y + 5x = 800 \quad (\text{constraint}) \end{array}$$

This form of the model corresponds to Figure 1.2, which contains two decision variables ( $x$  and  $y$ , or price and demand) and one constraint on the decision variables. The spreadsheet model in Figure 1.2 is ready for optimization.

To start, we select the Risk Solver Platform tab and click on the Model icon (on the left side of its ribbon). This step opens the *task pane* on the right-hand side of the Excel window. The task pane contains four tabs: Model, Platform, Engine, and Output. Initially, the Model tab displays a window listing several components of the software, including Optimization. In Figure 1.3, we have expanded the Optimization entry on the Model tab. As we specify the elements of our model, they are recorded in the folder icons of this window. At the top of the model tab five icons appear:

- Green “plus” sign, to *Add* model specifications
- Red “delete” sign, to *Remove* specifications
- Orange paired sheets with small blue arrows, to *Refresh* the display after changes
- Green checked sheet, to *Analyze* the model
- Green triangle, to *Solve* the specified optimization problem.

To specify the model we first select the decision cells (C9:C10) and then on the drop-down menu of the *Add* icon, select *Add Variable*. The range \$C\$9 : \$C\$10 immediately appears in the Model window, in the folder for Normal Variables. (Another way

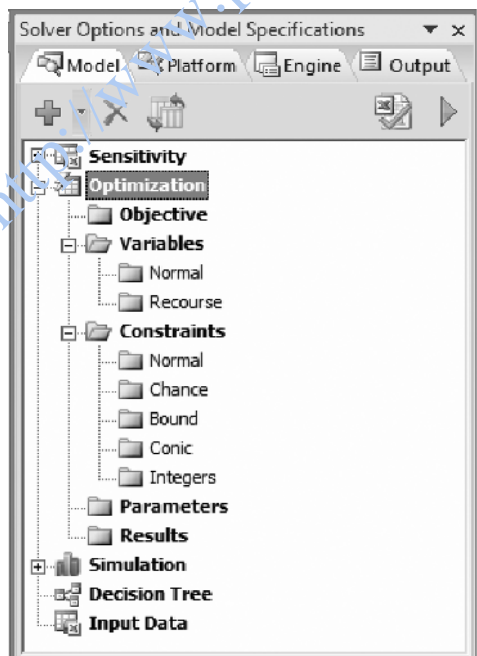


Figure 1.3. Model tab on the initial task pane.

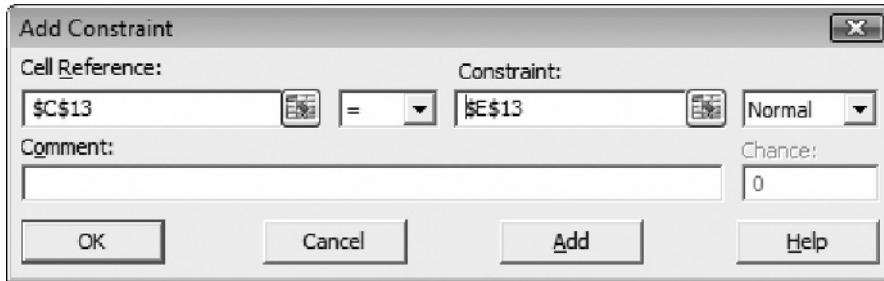


Figure 1.4. Add Constraint window.

to accomplish this step without using the drop-down menu is to highlight the Normal Variables folder icon and simply click the *Add* icon.)

Next, we select the objective cell (C16) and on the drop-down menu of the *Add* icon, select Add Objective. The cell address \$C\$16 immediately appears in the Model window, in the folder for Objective. By default, the specification assumes that the objective is to maximize this value. (We can implement this step by highlighting the Objective folder and simply clicking the *Add* icon.)

Next, we select the left-hand side of the constraint (C13) and on the drop-down menu of the *Add* icon, select Add Constraint. (Alternatively, we can highlight the Normal Constraints folder icon and click the *Add* icon.) The Add Constraint window appears, with the cell address \$C\$13 in the Cell Reference box, as shown in Figure 1.4. On the drop-down menu to its right, we select “=” and enter E13 in the Constraint box (or, with the cursor in the box, select cell E13).

When specifying constraints, one of our design guidelines for Solver models is to reference a cell containing a *formula* in the Cell Reference box and to reference a cell containing a *number* in the Constraint box. The use of cell references keeps the key parameters visible on the spreadsheet, rather than in the less accessible windows of Solver’s interface. The principle at work here is to communicate as much as possible about the model using the spreadsheet itself. Ideally, another person would not have to examine the task pane to understand the model. (Although Solver permits us to enter numerical values directly into the Constraint box, this form is less effective for communication and complicates sensitivity analysis. It would be reasonable only in special cases where the model structure is obvious from the spreadsheet and where we expect to perform no sensitivity analyses for the corresponding parameter.)

Finally, we press OK and observe that the task pane displays the model’s specification, as shown in Figure 1.5. In summary, our model specification is the following:

Objective: C16 (maximize)  
 Variables: C9:C10  
 Constraint: C13 = E13

This model is simple enough that we need not address the information on the Platform tab. (However, it is generally a good idea to set the Nonsmooth Model

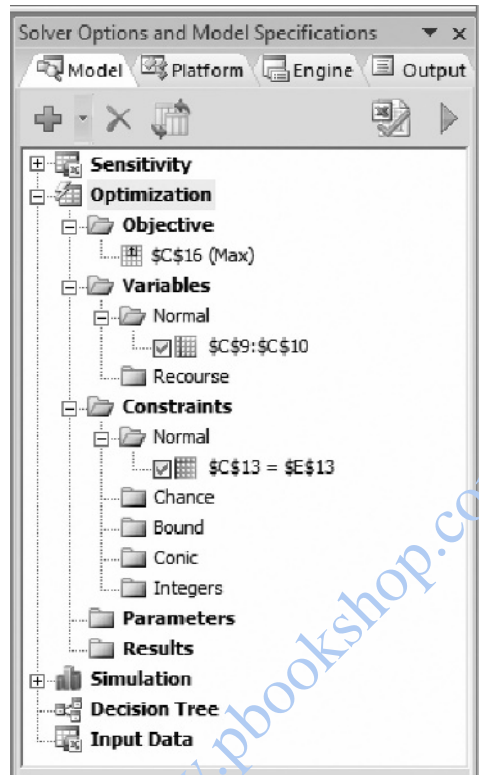


Figure 1.5. Model specification.

Transformation option to *Never*.) At the top of the Engine tab, we observe the default selection of the Standard GRG Nonlinear Engine, which we refer to as the *nonlinear solver*. (To ensure this selection, we uncheck the box for Automatically Select Engine.) This solution algorithm is appropriate for our optimization problem, and we do not need to address most of the other information on the tab. However, one of the options is important.

Although we may guess that the optimal price is a positive quantity, the model as specified permits the price decision to be negative. Such an outcome would not make sense in this problem, so it may be a good idea to limit the model to nonnegative prices. In fact, virtually all of the models in this book involve decision variables that make practical sense only when they are nonnegative, so we will impose this restriction routinely. On the Engine tab of the task pane, we find the *Assume Non-Negative* option in the General group and change it to True, using the drop-down menu on the right-hand side, as shown in Figure 1.6.

Finally, we proceed to the Output tab (or return to the Model tab) and click the *Solve* icon. Solver searches for the optimal price and ultimately places it in the price cell. In this case, the optimal price is \$100, and the corresponding quarterly profit contribution is \$18,000 as shown in Figure 1.7.

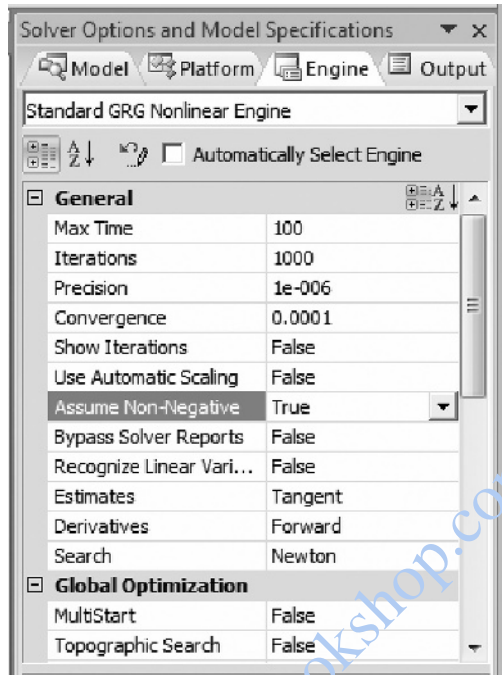


Figure 1.6. Setting the Assume Non-Negative option.

	A	B	C	D	E
1	<b>Price, Demand, and Profit</b>				
2					
3	<b>Inputs</b>				
4		<i>Max. demand</i>	800		
5		<i>Slope</i>	-5		
6		<i>Cost</i>	40		
7					
8	<b>Decisions</b>				
9		<i>Price</i>	\$ 100		
10		<i>Demand</i>	300		
11					
12	<b>Constraints</b>				
13		<i>Demand</i>	800	=	800
14					
15	<b>Outcomes</b>				
16		<i>Profit</i>	\$ 18,000		
17					

Figure 1.7. Optimal solution for Example 1.1.

Meanwhile, the Output tab's window displays the solution log for the optimization run. (The detail in this log is controlled by the Log Level option on the Platform tab, but the default setting of *Normal* is usually adequate.) The most important part of the log is the Solver Results message, which in this case states:

Solver found a solution. All constraints and optimality conditions are satisfied.

This *optimality message*, which is repeated at the very bottom of the task pane, tells us that no problems arose during the optimization and Solver was able to find an optimal solution. The profit-maximizing price is \$100, yielding an optimal profit of \$18,000. No other price can achieve more than this level. Thus, if we are confident that the demand curve continues to hold, the profit-maximizing decision would be to set price at \$100.

We have used Example 1.1 to introduce Solver and its interface. The task pane contains many user-selected options that are not a concern in this problem. In later chapters, we cover many of these settings and discuss when they become relevant. We also discuss the variations that can occur in optimization runs. For example, depending on the initial values of the decision variables, the nonlinear solver may generate the following result message in the solution log.

Solver has converged to the current solution. All constraints are satisfied.

This *convergence message* indicates that Solver has not been able to confirm optimality. Usually, this condition occurs because of numerical issues in the solution algorithm, and the resolution is to rerun Solver from the point where convergence occurred. Normally, one or two iterations are sufficient to produce the optimality message. We discuss Solver's result messages in more detail later.

With Solver, we can minimize an objective function instead of maximizing it. We return to the specification in the window of the Model tab of the task pane and double-click on the entry in the Objective folder. The Change Objective window appears, as shown in Figure 1.8, and we can select the button for Min rather than Max. (A third option allows us to specify a target value and find a set of variables




Figure 1.8. Selecting minimization of an objective.

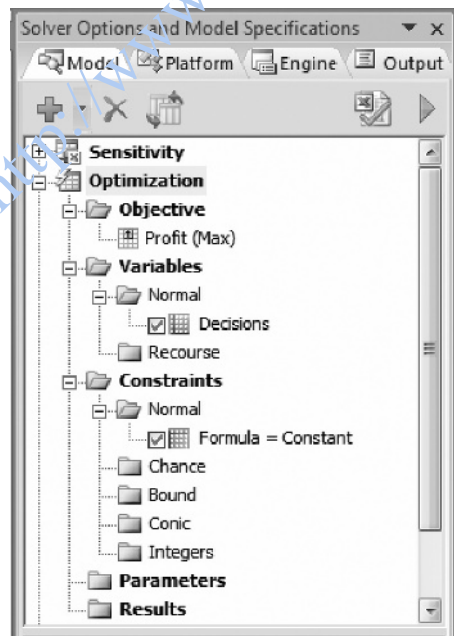
**BOX 1.2** *Excel Mini-Lesson: Using Range Names with Solver*

Excel offers the opportunity to refer to a cell range using a custom name. The range name can be entered in the Name Box (located just above the heading for column A) after selecting the desired range of cells. A one-cell range can be named in the same manner.

To illustrate the effect of using named ranges, suppose we return to the model of Figure 1.2 and name the following cells:

<i>Cells</i>	<i>Name</i>
C9:C10	Decisions
C13	Formula
E13	Constant
C16	Profit

Then the task pane window describes the model with range names instead of cell references, as shown in Figure 1.9. When a new user examines the model, this form is likely to be more meaningful than the use of literal cell references because the range names provide both description and documentation. Thus, range names are valuable for situations in which communicating the model to other audiences is an important consideration. When Solver is applied in an organizational setting, the use of range names is normally desirable. In the remainder of this book, however, we will continue to rely on cell references because they relate the information in the task pane directly to the contents of the spreadsheet display.



**Figure 1.9.** Model specification with range names.

that achieves the target value. This is not an optimization tool, and we will not pursue this particular capability.)

When an optimization model contains several decision variables, we can enter them one at a time, creating a list of Normal Variables in the task pane, each with its own checked box. More conveniently, we can arrange the spreadsheet so that all the variables appear in adjacent cells, as in Figure 1.2, and reference their cell range with just one entry in the Normal Variables folder. Because most optimization problems have several decision variables, we save time by placing them in adjacent cells. This layout also makes the information in the task pane easier to interpret when someone else is trying to audit our work, or if we are reviewing it after not having seen it for a long time. However, exceptions to this design guideline sometimes occur. Certain applications lead us to use nonadjacent locations for convenience in laying out the decision variable cells.

## SUMMARY

Many types of applications invite the use of Excel's Solver. In one sense, that is what this book is about—the problem types that Solver can handle and the use of Solver to obtain solutions. Thus, the book builds skill and confidence with spreadsheet applications because Solver is a spreadsheet tool. Actually, as mentioned earlier, Solver is a collection of procedures. Therefore, this book describes a variety of applications that can be addressed with spreadsheet capabilities.

In another sense, this book is about the problem types that Solver can handle, but the information on how to run Solver is incidental. The transcendent theme is the building of optimization models. If Solver wasn't around to produce solutions, then some other software would perform the computational task. The more basic skill is creating the model in the first place and recognizing its potential role in decision support.

Thus far, we have introduced six design guidelines for spreadsheet optimization models.

- Separate inputs from decisions and decisions from outputs.
- Create distinct modules for decision variables, objective function, and constraints.
- Display parameters explicitly on the spreadsheet, rather than in formulas.
- Enter parameters in the spreadsheet, rather than in the Add Constraints window.
- Place decision variables in adjacent cells.
- Highlight important cells, such as the decision variables and the objective.

Subsequent chapters introduce additional features of good spreadsheet design. This is not a claim that each example spreadsheet is the only possible way of designing a model, or even that it's the best way. A model should be easy to recognize, debug, use routinely, and pass on to others. A key feature of a good spreadsheet model is its ability to communicate clearly.

Chapters 2–5 deal with the linear solver, introducing many features of optimization analysis in the process. Chapters 6 and 7 deal with models that can be solved with the integer solver, and Chapter 8 deals with the nonlinear solver. The evolutionary solver, which is introduced in Chapter 9, is not properly an optimization procedure in the same sense as the others, but it applies in situations where the other solvers might fail. Each chapter is filled with illustrative examples and followed by a set of practice exercises. If readers work through the examples and the exercises they will develop a firm grasp on how to solve practical optimization problems using spreadsheets.

## EXERCISES

**1.1. Determining an Optimal Price** A firm's Marketing Department has estimated the demand curve of a product as  $y = 1100 - 7x$ , where  $y$  represents demand and  $x$  represents the unit selling price (in dollars) for the relevant decision period. The unit cost is known to be \$24. What price maximizes net income from sales of the product?

**1.2. Pricing in Two Markets** Global Products, Inc. has been making an electronic appliance for the domestic market. Demand for the appliance is price sensitive, and the demand curve is known to follow the linear function  $D = 4000 - 5P$ , where  $D$  represents annual demand and  $P$  represents selling price in the home currency, which is the Frank (F). The cost of manufacturing the appliance is 100F.

For the coming year, Global is planning to sell the same product in a foreign market, where the currency is the Marc (M). From surveys, the demand curve in the foreign country is estimated to follow a different linear function,  $D = 2000 - 2P$ , where the price is denominated in Marcs.

All production will be carried out at Global's domestic plant, with the expectation that the unit cost will remain unchanged. The exchange rate is 1.5 M/F, and Global plans to offer an equivalent price in both markets.

- If Global were to operate exclusively in its domestic market, what would be its profit-maximizing price and its annual profit?
- When Global sells in both markets at one equivalent price, what is its profit-maximizing price and its annual profit?

**1.3. Locating a Distribution Center** Northeast Parts Supply is a wholesale distributor of components for printers, fax machines, scanners, and related equipment. Northeast stocks expensive spare parts, which dealers prefer not to hold, and offers same-day delivery on any order. The firm now serves eight dealers in the New England area and wishes to locate its distribution facility at a central point. In particular, its dealers have each been assigned a location on an  $x$ - $y$  grid, and Northeast would like to find the best location for the distribution facility.

The eight dealers and their grid locations are shown in the following table:

Dealer	1	2	3	4	5	6	7	8
$x$ -location	25	82	10	27	93	14	68	147
$y$ -location	32	36	71	58	68	163	149	192

- Determine the location that minimizes the sum of the distances from the distribution facility to the dealers.
  - Determine the location that minimizes the maximum distance from the distribution facility to any of the dealers.
- 1.4. Collecting Credit Card Debt** A bank offers a credit card that can be used in various locations. The bank's analysts believe that the percentage  $P$  of accounts receivable collected by  $t$  months after credit is issued increases at a decreasing rate. Historical data suggest the following function:

$$P = 0.9[1 - \exp(-0.6t)]$$

The average credit issued in any one month is \$125 million, and historical experience suggests that for new credit issued in any month, collection efforts cost \$1 million per month.

- (a) Determine the number of months that collection efforts should be continued if the objective is to maximize the net collections (dollars collected minus collection costs). Allow for fractional months.
- (b) Under the optimal policy in (a), what percentage of accounts receivable should be collected?

**1.5. Allocating Plant Output** A firm owns five manufacturing plants that are responsible for the quarterly production of an industrial solvent. The production process exhibits diseconomies of scale. At plant  $p$ , the cost of making  $x$  thousand pounds of the solvent is approximated by the quadratic function  $f(x) = (1/c_p)x^2$ . The parameters  $c_p$  are plant dependent, as shown in the table.

$p$	1	2	3	4	5
$c_p$	3	6	4	8	5

The quarterly volume requirement is 50,000 pounds.

How should production be allocated among the five plants in order to minimize the total cost of meeting the volume requirement?

**1.6. Determining Production Lot Sizes** Four products are routed through a machining center that is notorious for its delays. Each product has had stable demand for some time, so that average weekly demand is predictable over a 3–6 month time frame. However, in the short run, demand fluctuates a great deal, and the load at the machining center varies considerably. The production control system dictates the lot size for each of the products. These quantities are shown, along with other relevant information, in the following table.

Product no.	Demand (weekly)	Setup (hours)	Run time (hours/1000)	Lot size
1	100	3	30	100
2	500	15	45	500
3	50	6	75	100
4	250	24	150	1500

With the current lot sizes, the machining center is running at a utilization of about 76%, but long lead times, sometimes over 2 weeks, have discouraged production planners from increasing its load. (A week contains 120 productive hours.) In the past, lead times spiraled out of control when utilization grew to around 80%.

A lead time model for this problem has been constructed on a spreadsheet.<sup>2</sup> The model permits the user to select lot sizes and thereby influence the average lead time through the bottleneck work center. The lead time prediction is based on advanced modeling techniques, but the details of the model are not of primary importance.

What is the shortest possible lead time, and what lot sizes achieve this value?

<sup>2</sup>The lead time model is available in the DataSets workbook, at the book's website (<http://mba.tuck.dartmouth.edu/opt/>).

**1.7. Resolving a Construction Dilemma** A library building is about to undergo some renovations that will improve its structural integrity. As part of the process, a number of steel beams will be carried through the existing bookcases from a broad, open area around the entry point. The central aisle between the bookcases is 10 feet wide, while the side aisles (which run perpendicular to the central aisle) are 6 feet wide. The renovation will require that steel beams be carried through the stacks, down the main aisle and turning into the smaller aisles.

What is the longest steel beam that can be carried horizontally through this space to a construction point along the outer walls?

**1.8. Selecting the Number of Warehouses** The customers of a particular company are located throughout an area comprised of  $S$  square miles, and they are serviced from  $k$  warehouses. On average, the distance in miles between a warehouse and a customer is given by the formula  $(S/k)^{0.5}$ . The annual capital cost of building a warehouse is \$40,000 and the annual operating cost of running a warehouse is \$60,000. Annual shipping costs average \$1 per mile per customer.

Suppose that the current market size is 250,000 customers, spread out over an area of 500 square miles. What is the optimal number of warehouses for the firm to operate?

## REFERENCES

1. McFEDRIES, P. *Excel 2010 Simplified*. John Wiley and Sons, 2010.
2. REDING, E. and L. WERMERS. *Microsoft Office Excel 2010: Illustrated Introductory*. Cengage Learning, 2011.
3. SCHRAGE, L. *Optimization Modeling with LINGO*. Lindo Systems Inc., 2008.
4. FOURER, R., D.M. GAY, and B.W. KERNIGHAN. *AMPL: A Modeling Language for Mathematical Programming (Second Edition)*. Cengage Learning, 2003.