

Introducing optimization with Excel Solver

In many situations, you want to find the best way to do something. More formally, you want to find the values of certain cells in a worksheet that *optimize* (maximize or minimize) a certain objective. Microsoft Excel Solver helps you answer optimization problems such as the following ones:

- How can a large drug company determine the monthly product mix that maximizes corporate profitability at its Indianapolis plant?
- If Microsoft produces Xbox consoles at three locations, how can it minimize the cost of meeting demand for them?
- What price for Xbox consoles and games will maximize profit from Xbox sales for Microsoft?
- Microsoft would like to undertake 20 strategic initiatives that will tie up money and skilled programmers for the next five years. It does not have enough resources for all 20 projects; which ones should it undertake?
- How do bookmakers find the best set of ratings to set accurate point spreads for NFL teams?
- How should I allocate my retirement portfolio among high-tech stocks, value stocks, bonds, cash, and gold?

An optimization model has three parts: the target cell, the changing cells, and the constraints. The *target cell* represents the objective or goal. You want to minimize or maximize the amount in the target cell. In the question about a drug company's product mix given earlier, the plant manager would presumably want to maximize the profitability of the plant during each month. The cell that measures profitability would be the target cell. The target cells for each situation described at the beginning of the chapter are listed in Table 28-1.

Keep in mind, however, that in some situations, you might have multiple target cells. For example, Microsoft might have a secondary goal to maximize Xbox market share.

TABLE 28-1 List of target cells

Model	Maximize or minimize	Target cell
Drug company product mix	Maximize	Monthly profit
Xbox shipping	Minimize	Distribution costs
Xbox pricing	Maximize	Profit from Xbox consoles and games
Microsoft project initiatives	Maximize	Net present value (NPV) contributed by selected projects
NFL ratings	Minimize	Difference between scores predicted by ratings and actual game scores
Retirement portfolio	Minimize	Risk factor of portfolio

Changing cells are the worksheet cells that you can change or adjust to optimize the target cell. In the drug company example, the plant manager can adjust the amount produced for each product during a month. The cells in which these amounts are recorded are the changing cells in this model. Table 28-2 lists the appropriate changing cell definitions for the models described at the beginning of the chapter, and Table 28-3 lists the problem constraints.

TABLE 28-2 List of changing cells

Model	Changing cells
Drug company product mix	Amount of each product produced during the month
Xbox shipping	Amount produced at each plant each month that is shipped to each customer
Xbox pricing	Console and game prices
Microsoft program initiatives	Which projects are selected
NFL ratings	Team ratings
Retirement portfolio	Fraction of money invested in each asset class

TABLE 28-3 List of problem constraints

Model	Constraints
Drug company product mix	Product mix uses no more resources than are available Do not produce more of a product than can be sold
Xbox shipping	Do not ship more units each month from a plant than plant capacity Make sure that each customer receives the number of Xbox consoles that it needs
Xbox pricing	Prices can't be too far out of line with competitors' prices
Microsoft project initiatives	Projects selected can't use more money or skilled programmers than are available
NFL ratings	None
Retirement portfolio	Invest all our money somewhere (cash is a possibility) Obtain an expected return of at least 10 percent on our investments

The best way to understand how to use Solver is by looking at some detailed examples. In later chapters, you learn how to use Solver to address each of the situations described in this chapter as well as several other important business problems.

To activate Solver, click the **File** tab, choose **Options**, and then click **Add-Ins**. In the **Manage** box at the bottom of the dialog box, select **Excel Add-Ins** and click **Go**. Select the **Solver Add-In** check box in the **Add-Ins** dialog box and click **OK**. After Solver is activated, you can run Solver by clicking **Solver** in the **Analysis** group on the **Data** tab. Figure 28-1 shows the **Solver Parameters** dialog box. In the next few chapters, you see how to use this dialog box to configure the target cell, changing cells, and constraints for a Solver model.

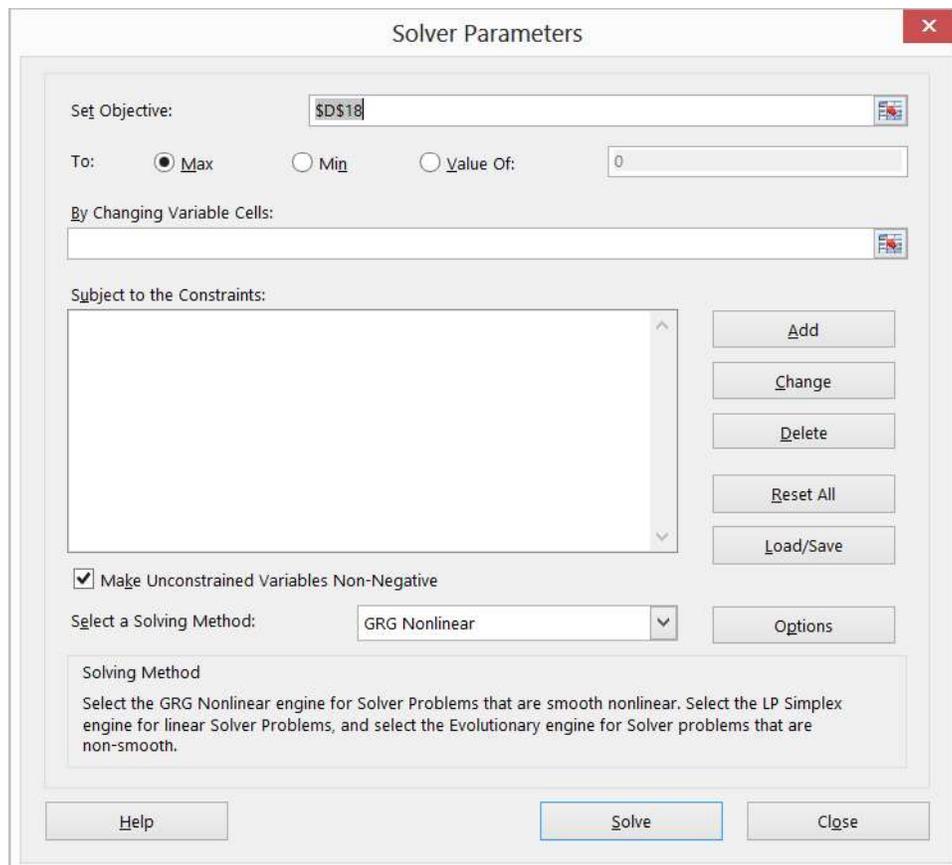


FIGURE 28-1 This figure shows the Solver Parameters dialog box.

Beginning with Excel 2010, Solver has been greatly revamped and improved. The primary change is the presence of the Select A Solving Method list. From this list, you must select the appropriate solution engine for your optimization problem:

- The Simplex LP engine solves linear optimization problems. As you will see in Chapters 29–32, a linear optimization problem is one in which the target cell and constraints are all created by adding terms of the (changing cell)*(constant) form.
- The GRG nonlinear engine solves optimization problems in which the target cell or some of the constraints are not linear and are computed by using common mathematical operations

such as multiplying or dividing changing cells, raising changing cells to a power, exponential or trigonometric functions involving changing cells, and so on. The GRG engine includes a powerful Multistart option that enables you to solve many problems that were solved incorrectly with previous versions of Excel. Chapters 33 through 35 discuss the GRG nonlinear engine.

- The Evolutionary Solver engine is used when your target cell or constraints contain non-smooth functions that reference changing cells. A nonsmooth function is one whose slope abruptly changes. For example, when $x = 0$, the slope of the absolute value of x abruptly changes from -1 to 1 . If your target cell or constraints contain *IF*, *SUMIF*, *COUNTIF*, *SUMIFS*, *COUNTIFS*, *AVERAGEIF*, *AVERAGEIFS*, *ABS*, *MAX*, or *MIN* functions that reference the changing cells, you are using nonsmooth functions, and the Evolutionary Solver engine probably has the best shot at finding a good solution to your optimization problem. The Evolutionary Solver engine is discussed in Chapter 36, “Penalties and the Evolutionary Solver,” and Chapter 37, “The traveling salesperson problem.”

After you have input the target cell, changing cells, and constraints, what does Solver do? To answer this question, you need some background in Solver terminology. Any specification of the changing cells that satisfies the model’s constraints is a *feasible solution*. For instance, any product mix that satisfies the following three conditions would be a feasible solution:

- Does not use more raw material or labor than is available.
- Does not produce more of each product than is demanded.
- Does not produce a negative amount of any product.

Essentially, Solver searches all feasible solutions and finds the one that has the best target cell value (the largest value for maximum optimization, the smallest for minimum optimization). Such a solution is called an *optimal solution*. As you’ll see in Chapter 29, “Using Solver to determine the optimal product mix,” some Solver models have no optimal solution, and some have a unique solution. Other Solver models have multiple (actually, an infinite number of) optimal solutions. In the next chapter, you’ll begin your study of Solver examples by examining the drug company product mix problem.

Problems

For each situation described, identify the target cell, changing cells, and constraints:

1. I am borrowing \$100,000 for a 15-year mortgage. The annual rate of interest is 8 percent. How can I determine my monthly mortgage payment?
2. How should an auto company allocate its advertising budget between different advertising formats?
3. How should cities transport students to more distant schools to obtain racial balance?
4. If a city has only one hospital, where should it be located?

5. How should a drug company allocate its sales force's efforts among its products?
6. A drug company has allocated \$2 billion to purchasing biotech companies. Which companies should it buy?
7. The tax rate charged to a drug company depends on the country in which a product is produced. How can a drug company determine where each drug should be made?

Using Solver to determine the optimal product mix

Questions answered in this chapter:

- How can I determine the monthly product mix that maximizes profitability?
- Does a Solver model always have a solution?
- What does it mean if a Solver model yields the Set Values Do Not Converge result?

Answers to this chapter's questions

This section provides the answers to the questions that are listed at the beginning of the chapter.

How can I determine the monthly product mix that maximizes profitability?

Companies often need to determine the quantity of each product to produce on a monthly basis. In its simplest form, the product mix problem involves how to determine the amount of each product that should be produced during a month to maximize profits. Product mix must usually adhere to the following constraints:

- Product mix can't use more resources than are available.
- Because demand for each product is limited, you don't want to produce more of it during a month than demand dictates because the excess production is wasted (for example, a perishable drug).

Now, solve the following example of the product mix problem. You can find a feasible solution to this problem in the Feasible Solution worksheet of the Prodmix.xlsx file, shown in Figure 29-1. Trial values for the amount produced of each drug have been entered in row 2.

	B	C	D	E	F	G	H	I
1								
2		Pounds made	150	160	170	180	190	200
3	Available	Product	1	2	3	4	5	6
4	4500	Labor	6	5	4	3	2.5	1.5
5	1600	Raw Material	3.2	2.6	1.5	0.8	0.7	0.3
6		Unit price	\$ 12.50	\$ 11.00	\$ 9.00	\$ 7.00	\$ 6.00	\$ 3.00
7		Variable cost	\$ 6.50	\$ 5.70	\$ 3.60	\$ 2.80	\$ 2.20	\$ 1.20
8		Demand	960	928	1041	977	1084	1055
9		Unit profit cont.	\$ 6.00	\$ 5.30	\$ 5.40	\$ 4.20	\$ 3.80	\$ 1.80
10								
11								
12		Profit	\$ 4,504.00					
13					Available			
14		Labor Used	3695 <=		4500			
15		Raw Material Used	1488 <=		1600			

FIGURE 29-1 This figure shows a feasible product mix.

You work for a drug company that produces six products at its plant. Production of each product requires labor and raw material. Row 4 in Figure 29-1 shows the hours of labor needed to produce a pound of each product, and row 5 shows the pounds of raw material needed to produce a pound of each product. For example, producing a pound of Product 1 requires six hours of labor and 3.2 pounds of raw material. For each drug, the price per pound is given in row 6, the unit cost per pound is given in row 7, and the profit contribution per pound is given in row 9. For example, Product 2 sells for \$11.00 per pound, incurs a unit cost of \$5.70 per pound, and contributes \$5.30 profit per pound. The month's demand for each drug is given in row 8. For example, demand for Product 3 is 1,041 pounds. This month, 4,500 hours of labor and 1,600 pounds of raw material are available. How can this company maximize its monthly profit?

If you knew nothing about Solver, you might attack this problem by constructing a worksheet to track profit and resource usage associated with the product mix. Then you could use trial and error to vary the product mix to optimize profit without using more labor or raw material than is available or producing any drug in excess of demand. You use Solver in this process only at the trial-and-error stage. Essentially, Solver is an optimization engine that flawlessly performs the trial-and-error search.

A key to solving the product mix problem is to compute the resource usage and profit associated with any given product mix efficiently. An important tool that you can use to make this computation is the *SUMPRODUCT* function. It multiplies corresponding values in cell ranges and returns the sum of those values. Each cell range used in a *SUMPRODUCT* evaluation must have the same dimensions, so you can use *SUMPRODUCT* with two rows or two columns but not with one column and one row.

As an example of how you can use the *SUMPRODUCT* function in the product mix example, compute resource usage. Labor usage is calculated by the following formula:

$$(\text{Labor used per pound of drug 1}) * (\text{Drug 1 pounds produced}) + (\text{Labor used per pound of drug 2}) * (\text{Drug 2 pounds produced}) + \dots + (\text{Labor used per pound of drug 6}) * (\text{Drug 6 pounds produced})$$

You could compute labor usage in a more tedious fashion as $D2*D4+E2*E4+F2*F4+G2*G4+H2*H4+I2*I4$. Similarly, raw material usage could be computed as $D2*D5+E2*E5+F2*F5+G2*G5+H2*H5+I2*I5$. However, entering these formulas in a worksheet for six products is time-consuming. Imagine how long it would take if you were working with a company that produced, for example, 50 products at its plant. A much easier way to compute labor and raw material usage is to copy the `SUMPRODUCT(D2:I2,D4:I4)` formula from D14 to D15. This formula computes $D2*D4+E2*E4+F2*F4+G2*G4+H2*H4+I2*I4$ (which is the labor usage) but is much easier to enter. Use the \$ sign with the D2:I2 range so that when you copy the formula, you still capture the product mix from row 2. The formula in cell D15 computes raw material usage.

In a similar fashion, profit is determined by the following formula:

$(\text{Drug 1 profit per pound}) * (\text{Drug 1 pounds produced}) + (\text{Drug 2 profit per pound}) * (\text{Drug 2 pounds produced}) + \dots + (\text{Drug 6 profit per pound}) * (\text{Drug 6 pounds produced})$

Profit is easily computed in cell D12 with the `SUMPRODUCT(D9:I9,D2:I2)` formula.

You now can identify the three components of the product mix Solver model:

- **Target cell** The goal is to maximize profit (computed in cell D12).
- **Changing cells** The number of pounds produced of each product is listed in the D2:I2 cell range.
- **Constraints** You have the following constraints:
 - Do not use more labor or raw material than is available. That is, the values in cells D14:D15 (the resources used) must be less than or equal to the values in cells F14:F15 (the available resources).
 - Do not produce more of a drug than is in demand. That is, the values in the D2:I2 cells (pounds produced of each drug) must be less than or equal to the demand for each drug (listed in cells D8:I8).
 - You can't produce a negative amount of any drug.

This chapter shows you how to enter the target cell, changing cells, and constraints into Solver. Then all you need to do is click the Solve button to find a profit-maximizing product mix.

To begin, click the **Data** tab and, in the **Analysis** group, click **Solver**.



Note As explained in Chapter 28, “Introducing optimization with Excel Solver,” Solver is activated by clicking the **File** tab and then choosing **Options** and **Add-Ins**. In the **Manage** list, click **Excel Add-ins** and click **Go**. Select the **Solver Add-In** check box and then click **OK**.

The **Solver Parameters** dialog box opens, as shown in Figure 29-2.

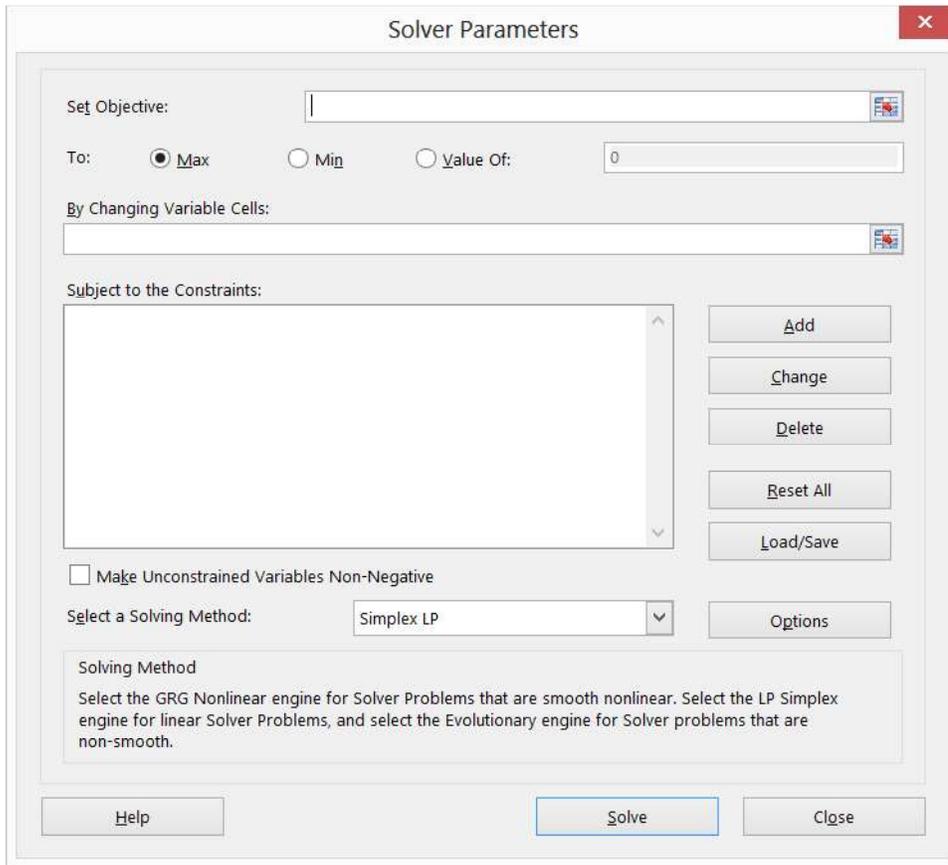


FIGURE 29-2 This figure shows the Solver Parameters dialog box.

Click the **Set Objective** box and then select the profit cell (cell D12). Click the **By Changing Variable Cells** box and then point to the D2:I2 range, which contains the pounds produced of each drug. The dialog box should now look like Figure 29-3.

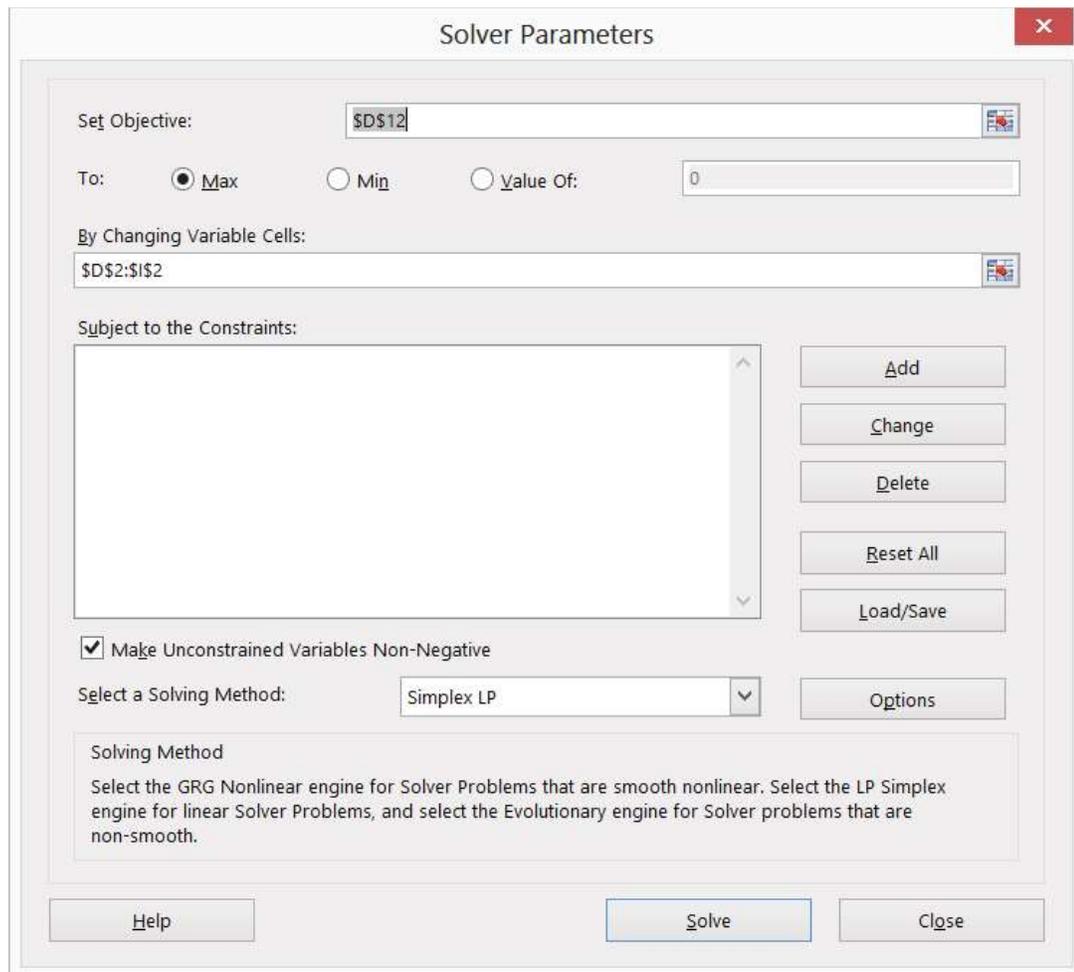


FIGURE 29-3 This figure shows the Solver Parameters dialog box with target cell and changing cells defined.

You're now ready to add constraints to the model. Click the **Add** button. The **Add Constraint** dialog box opens, as shown in Figure 29-4.

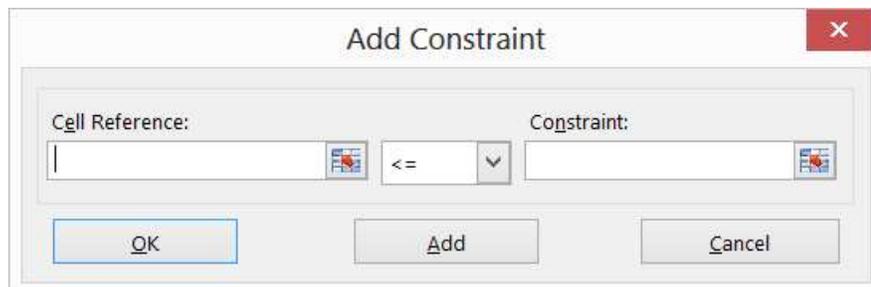


FIGURE 29-4 This figure shows the Add Constraint dialog box.

To add the resource usage constraints, click the **Cell Reference** box and then select the D14:D15 range. After confirming the selection of \leq from the middle list, click the **Constraint** box and then select the F14:F15 cell range. The **Add Constraint** dialog box should now look like Figure 29-5.

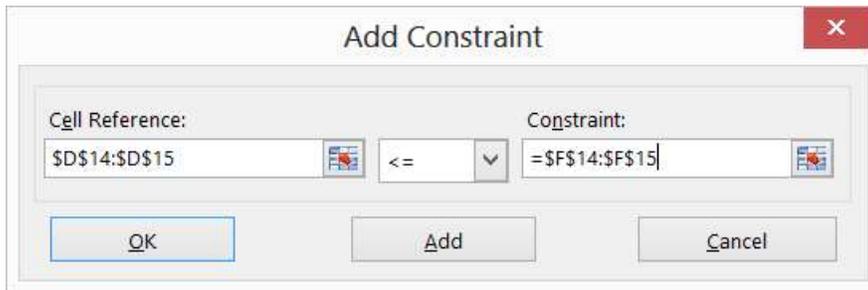


FIGURE 29-5 This figure shows the Add Constraint dialog box with the resource usage constraints entered.

You have now ensured that when Solver tries different values for the changing cells, only combinations that satisfy both $D14 \leq F14$ (labor used is less than or equal to labor available) and $D15 \leq F15$ (raw material used is less than or equal to raw material available) are considered. Click **Add** to enter the demand constraints, filling in the **Add Constraint** dialog box as shown in Figure 29-6.

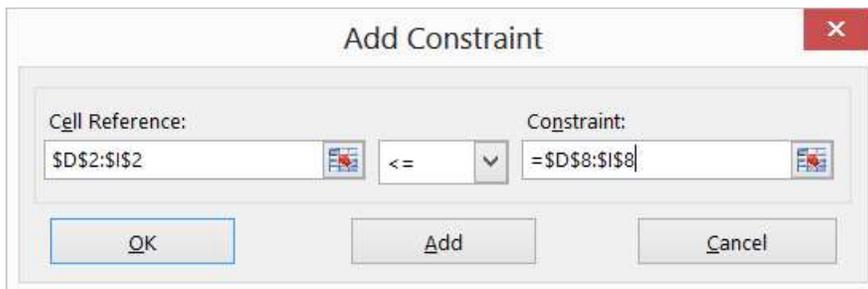


FIGURE 29-6 This figure shows the Add Constraint dialog box with the demand constraints entered.

Adding these constraints ensures that when Solver tries different combinations for the changing cell values, only combinations that satisfy the following parameters are considered:

- $D2 \leq D8$ (the amount produced of Drug 1 is less than or equal to the demand for Drug 1)
- $E2 \leq E8$ (the amount produced of Drug 2 is less than or equal to the demand for Drug 2)
- $F2 \leq F8$ (the amount produced of Drug 3 is less than or equal to the demand for Drug 3)
- $G2 \leq G8$ (the amount produced of Drug 4 is less than or equal to the demand for Drug 4)
- $H2 \leq H8$ (the amount produced of Drug 5 is less than or equal to the demand for Drug 5)
- $I2 \leq I8$ (the amount produced of Drug 6 is less than or equal to the demand for Drug 6)

Click **OK** in the **Add Constraint** dialog box. The **Solver Parameters** dialog box should look like Figure 29-7.

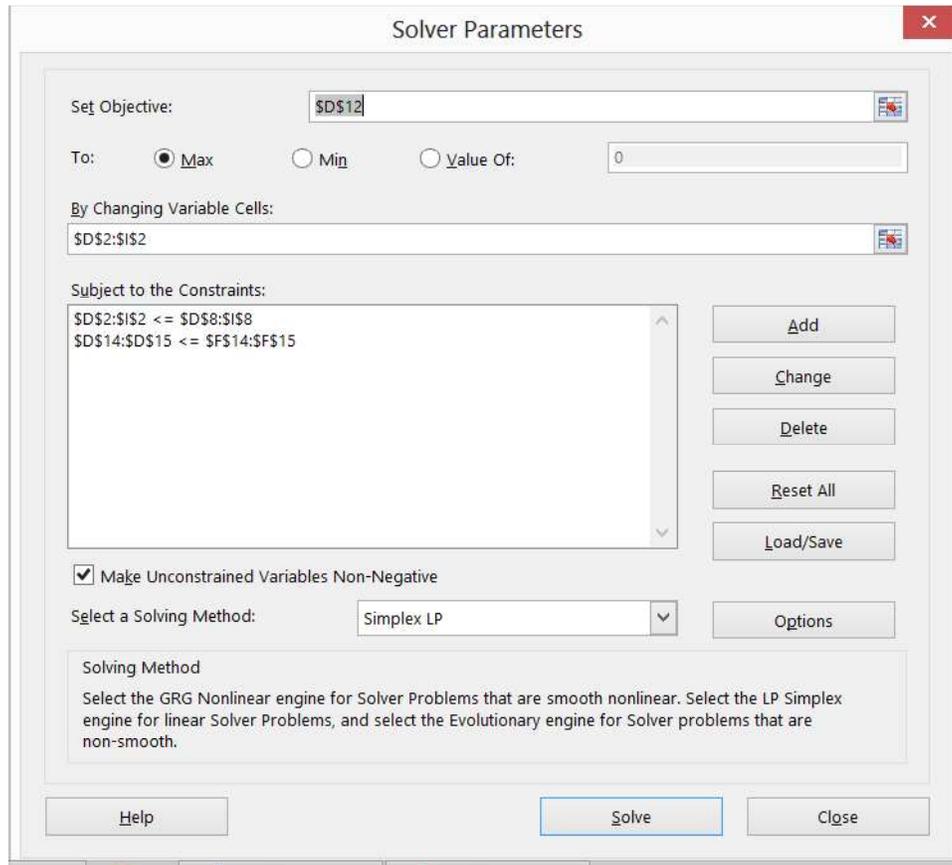


FIGURE 29-7 This is the final Solver Parameters dialog box for the product mix problem.

Selecting the **Make Unconstrained Variables Non-Negative** check box ensures that all the changing cells are forced to be greater than or equal to 0.

Next, choose Simplex LP from the Select A Solving Method list. You choose the Simplex LP engine because the product mix problem is a special type of Solver problem called a *linear* model. Essentially, a Solver model is linear under the following conditions:

- The target cell is computed by adding together the terms of the (changing cell)*(constant) form.
- Each constraint satisfies the *linear model requirement*. This means that each constraint is evaluated by adding together the terms of the (changing cell)*(constant) form and comparing the sums to a constant.

Why is this Solver problem linear? Your target cell (profit) is computed as follows:

$$(\text{Drug 1 profit per pound}) * (\text{Drug 1 pounds produced}) + (\text{Drug 2 profit per pound}) * (\text{Drug 2 pounds produced}) + \dots + (\text{Drug 6 profit per pound}) * (\text{Drug 6 pounds produced})$$

This computation follows a pattern in which the target cell's value is derived by adding together terms of the (changing cell)*(constant) form.

Your labor constraint is evaluated by comparing the value derived from (Labor used per pound of Drug 1)*(Drug 1 pounds produced)+(Labor used per pound of Drug 2)*(Drug 2 pounds produced)+ . . . (Labor used per pound of Drug 6)*(Drug 6 pounds produced) to the labor available.

Therefore, the labor constraint is evaluated by adding together the terms of the (changing cell)*(constant) form and comparing the sums to a constant. Both the labor constraint and the raw material constraint satisfy the linear model requirement.

The demand constraints take the following form:

```
(Drug 1 produced)<=(Drug 1 Demand)
(Drug 2 produced)<=(Drug 2 Demand)
...
(Drug 6 produced)<=(Drug 6 Demand)
```

Each demand constraint also satisfies the linear model requirement because each is evaluated by adding together the terms of the (changing cell)*(constant) form and comparing the sums to a constant.

Knowing that the product mix model is a linear model, why should you care?

- If a Solver model is linear and you select Simplex LP, Solver is guaranteed to find the optimal solution to the Solver model. If a Solver model is not linear, Solver might or might not find the optimal solution.
- If a Solver model is linear and you select Simplex LP, Solver uses a very efficient algorithm (the simplex method) to find the model's optimal solution. If a Solver model is linear and you do not select Simplex LP, Solver uses a very inefficient algorithm (the GRG2 method) and might have difficulty finding the model's optimal solution.

After you click **Solve**, Solver calculates an optimal solution (if one exists) for the product mix model. As stated in Chapter 28, an optimal solution to the product mix model would be a set of changing cell values (pounds produced of each drug) that maximizes profit over the set of all feasible solutions. Again, a feasible solution is a set of changing cell values satisfying all constraints. The changing cell values shown in Figure 29-8 are a feasible solution because all production levels are nonnegative, production levels do not exceed demand, and resource usage does not exceed available resources.

	C	D	E	F	G	H	I
1							
2	Pounds made	150	160	170	180	190	200
3	Product	1	2	3	4	5	6
4	Labor	6	5	4	3	2.5	1.5
5	Raw Material	3.2	2.6	1.5	0.8	0.7	0.3
6	Unit price	\$ 12.50	\$ 11.00	\$ 9.00	\$ 7.00	\$ 6.00	\$ 3.00
7	Variable cost	\$ 6.50	\$ 5.70	\$ 3.60	\$ 2.80	\$ 2.20	\$ 1.20
8	Demand	960	928	1041	977	1084	1055
9	Unit profit cont.	\$ 6.00	\$ 5.30	\$ 5.40	\$ 4.20	\$ 3.80	\$ 1.80
10							
11							
12	Profit	\$ 4,504.00					
13				Available			
14	Labor Used	3695	<=	4500			
15	Raw Material Used	1488	<=	1600			

FIGURE 29-8 A feasible solution to the product mix problem fits within constraints.

The changing cell values shown in Figure 29-9 represent an *infeasible solution* for the following reasons:

- You produce more of Drug 5 than the demand for it.
- You use more labor than is available.
- You use more raw material than is available.

	B	C	D	E	F	G	H	I
1								
2		Pounds made	300	0	0	0	1085	1000
3	Available	Product	1	2	3	4	5	6
4	4500	Labor	6	5	4	3	2.5	1.5
5	1600	Raw Material	3.2	2.6	1.5	0.8	0.7	0.3
6		Unit price	\$ 12.50	\$ 11.00	\$ 9.00	\$ 7.00	\$ 6.00	\$ 3.00
7		Variable cost	\$ 6.50	\$ 5.70	\$ 3.60	\$ 2.80	\$ 2.20	\$ 1.20
8		Demand	960	928	1041	977	1084	1055
9		Unit profit cont.	\$ 6.00	\$ 5.30	\$ 5.40	\$ 4.20	\$ 3.80	\$ 1.80
10								
11								
12		Profit	\$ 7,723.00					
13					Available			
14		Labor Used	6012.5	<=	4500			
15		Raw Material Used	2019.5	<=	1600			

FIGURE 29-9 An infeasible solution to the product mix problem doesn't fit within the defined constraints.

	B	C	D	E	F	G	H	I
1								
2		Pounds made	0	0	0	596.6667	1084	0
3	Available	Product	1	2	3	4	5	6
4	4500	Labor	6	5	4	3	2.5	1.5
5	1600	Raw Material	3.2	2.6	1.5	0.8	0.7	0.3
6		Unit price	\$ 12.50	\$ 11.00	\$ 9.00	\$ 7.00	\$ 6.00	\$ 3.00
7		Variable cost	\$ 6.50	\$ 5.70	\$ 3.60	\$ 2.80	\$ 2.20	\$ 1.20
8		Demand	960	928	1041	977	1084	1055
9		Unit profit cont.	\$ 6.00	\$ 5.30	\$ 5.40	\$ 4.20	\$ 3.80	\$ 1.80
10								
11								
12		Profit	\$ 6,625.20					
13					Available			
14		Labor Used	4500	<=	4500			
15		Raw Material Used	1236.13333	<=	1600			

FIGURE 29-10 This is the optimal solution to the product mix problem.

Your drug company can maximize its monthly profit at a level of \$6,625.20 by producing 596.67 pounds of Drug 4; 1,084 pounds of Drug 5; and none of the other drugs. You can't determine whether you can achieve the maximum profit of \$6,625.20 in other ways. All you can be sure of is that with your limited resources and demand, you cannot make more than \$6,625.20 this month.

Does a Solver model always have a solution?

Suppose that demand for each product *must* be met. (See the No Feasible Solution worksheet in the Prodmix.xlsx file.) You then have to change your demand constraints from D2:I2<=D8:I8 to D2:I2>=D8:I8. To do this, open **Solver**, select the D2:I2<=D8:I8 constraint, and then click **Change**. The **Change Constraint** dialog box, shown in Figure 29-11, opens.

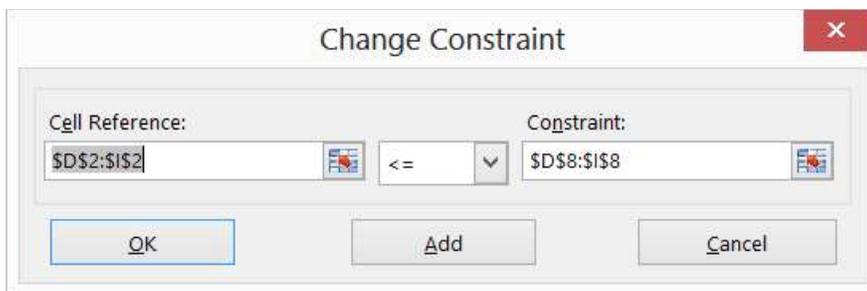


FIGURE 29-11 This is the Change Constraint dialog box.

As shown in Figure 29-12, select **>=** and then click **OK**. You've now ensured that Solver will consider changing only cell values that meet all demands. When you click **Solve**, you'll see the message, "Solver could not find a feasible solution." This message does not mean that you made a mistake in your model but, rather, that with limited resources, you can't meet demand for all products. Solver is

simply telling you that if you want to meet demand for each product, you need to add more labor, more raw material, or both.

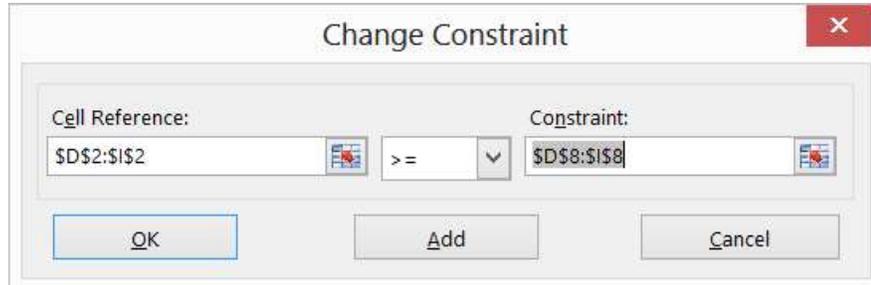


FIGURE 29-12 Demand constraint now ensures that demand is met.

What does it mean if a Solver model yields the Set Values Do Not Converge result?

What happens if you allow unlimited demand for each product and negative quantities to be produced of each drug? (You can see this Solver problem on the Set Values Do Not Converge worksheet in the Prodmix.xlsx file.) To find the optimal solution for this situation, open **Solver** and clear the **Make Unconstrained Variables Non-Negative** check box. In the **Solver Parameters** dialog box, select the D2:I2<=D8:I8 demand constraint and then click **Delete** to remove the constraint. When you click **Solve**, Solver returns the Set Cell Values Do Not Converge message. This message means that if the target cell is to be maximized (as in this example), there are feasible solutions with arbitrarily large target cell values. For example, if you wanted to make a billion dollars in profit, Solver could find a feasible solution to make that level of profit (see Problem 7). (If the target cell is to be minimized, the Set Cell Values Do Not Converge message means there are feasible solutions with arbitrarily small target cell values.) In this situation, by allowing negative production of a drug, you create resources that can be used to produce arbitrarily large amounts of other drugs. Given your unlimited demand, you can make unlimited profit. In a real situation, you can't make an infinite amount of money. In short, if you see Set Values Do Not Converge, your model does have an error.

Problems

1. Suppose your drug company can purchase up to 500 hours of labor at \$1 more per hour than current labor costs. How can it maximize profit?
2. At a chip manufacturing plant, four technicians (A, B, C, and D) produce three products (Products 1, 2, and 3). This month, the chip manufacturer can sell 80 units of Product 1, 50 units of Product 2, and, at most, 50 units of Product 3. Technician A can make only Products 1 and 3. Technician B can make only Products 1 and 2. Technician C can make only Product 3. Technician D can make only Product 2. For each unit produced, the products contribute the

following profit: Product 1, \$6; Product 2, \$7; and Product 3, \$10. The time (in hours) each technician needs to manufacture a product is as follows:

Product	Technician A	Technician B	Technician C	Technician D
1	2	2.5	Cannot do	Cannot do
2	Cannot do	3	Cannot do	3.5
3	3	Cannot do	4	Cannot do

Each technician can work up to 120 hours per month. How can the chip manufacturer maximize its monthly profit? Assume a fractional number of units can be produced.

3. A computer manufacturing plant produces mice, keyboards, and video game joysticks. The per-unit profit, per-unit labor usage, monthly demand, and per-unit machine-time usage are given in the following table:

	Mice	Keyboards	Joysticks
Profit/unit	\$8	\$11	\$9
Labor usage/unit	.2 hour	.3 hour	.24 hour
Machine time/unit	.04 hour	.055 hour	.04 hour
Monthly demand	15,000	29,000	11,000

Each month, a total of 13,000 labor hours and 3,000 hours of machine time are available. How can the manufacturer maximize its monthly profit contribution from the plant?

4. Resolve your drug example, assuming that a minimum demand of 200 units for each drug must be met.
5. Jason makes diamond bracelets, necklaces, and earrings. He wants to work a maximum of 160 hours per month. He has 800 ounces of diamonds. The profit, labor time, and ounces of diamonds required to produce each product are as follows. If demand for each product is unlimited, how can Jason maximize his profit?

Product	Unit profit	Labor hours per unit	Ounces of diamonds per unit
Bracelet	\$300	.35	1.2
Necklace	\$200	.15	.75
Earrings	\$100	.05	.5

In your product mix example, suppose that whenever you sell more than 400 pounds of any product, you must give a \$1 per pound discount on each pound above 400 sold. How does this change the answer to the problem?

6. If demand is unlimited and negative production is allowed, find a feasible solution that earns a billion dollars.
7. Assuming unlimited demand and unlimited resources, find a feasible solution that earns \$1 billion in profit.

Using Solver to schedule your workforce

Question answered in this chapter:

- How can I schedule my workforce efficiently to meet labor demands?

Many organizations (such as banks, restaurants, and postal service companies) know what their labor requirements are at different times of the day, and they need a method to schedule their workforce efficiently. You can use Solver to solve workforce scheduling problems easily.

Answer to this chapter’s question

This section provides the answer to the question that is listed at the beginning of the chapter.

How can I schedule my workforce efficiently to meet labor demands?

Bank 24 processes checks seven days a week. The number of workers needed each day to process checks is shown in row 14 of the Bank24.xlsx file, which appears in Figure 30-1. For example, 13 workers are needed on Tuesday, 15 workers are needed on Wednesday, and so on. All bank employees work five consecutive days. What is the minimum number of employees Bank 24 can have and still meet its labor requirements?

	A	B	C	D	E	F	G	H	I	J
2	Total				There are multiple solutions all of which use 20 workers.					
3	20		Working?							
4	Number starting	Day worker starts	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday	
5	1	Monday		1	1	1	1	0	0	
6	3	Tuesday		0	1	1	1	1	0	
7	0	Wednesday		0	0	1	1	1	1	
8	4	Thursday		1	0	0	1	1	1	
9	1	Friday		1	1	0	0	1	1	
10	2	Saturday		1	1	1	0	0	1	
11	9	Sunday		1	1	1	1	0	0	1
12		Number working		17	16	15	17	9	10	16
13			>=	>=	>=	>=	>=	>=	>=	
14		Number needed		17	13	15	17	9	9	12

FIGURE 30-1 This figure shows the data you’ll use to work through the bank workforce scheduling problem.

- **Target cell** Minimize total number of employees.
- **Changing cells** Number of employees who start work (the first of five consecutive days) each day of the week. Each changing cell must be a nonnegative integer.
- **Constraints** For each day of the week, the number of employees who are working must be greater than or equal to the number of employees required—(Number of employees working) \geq (Needed employees).

To set up the model for this problem, you need to track the number of employees working each day. Begin by entering trial values in the A5:A11 cell range for the number of employees who start their five-day shift each day. For example, in A5, 1 is entered, indicating that 1 employee begins work on Monday and works Monday through Friday. Each day's required workers are entered in the C14:I14 range.

To track the number of employees working each day, a 1 or a 0 is entered in each cell in the C5:I11 range. The value 1 indicates that the employees who started working on the day designated in the cell's row are working on the day associated with the cell's column. For example, the 1 in cell G5 indicates that employees who started working on Monday are working on Friday; the 0 in cell H5 indicates that the employees who started working on Monday are not working on Saturday.

By copying the `=SUMPRODUCT(A5:A11,C5:C11)` formula from C12 to D12:I12, you can compute the number of employees working each day. For example, in cell C12, this formula evaluates to `=A5+A8+A9+A10+A11`, which equals (Number starting on Monday) + (Number starting on Thursday) + (Number starting on Friday) + (Number starting on Saturday) + (Number starting on Sunday). This total is indeed the number of people working on Monday.

After computing the total number of employees in cell A3 with the `=SUM(A5:A11)` formula, you can enter a model in Solver as shown in Figure 30-2.

In the target cell (A3), you want to minimize the number of total employees. The `C12:I12>=C14:I14` constraint ensures that the number of employees working each day is at least as large as the number needed each day. The `A5:A11=integer` constraint ensures that the number of employees beginning work each day is an integer. To add this constraint, click **Add** in the **Solver Parameters** dialog box and fill in the **Add Constraint** dialog box as shown in Figure 30-3.

Note that this model is linear because the target cell is created by adding together changing cells, and the constraint is created by comparing the result obtained by adding together the product of each changing cell times a constant (either 1 or 0) to the required number of workers. Therefore, you select the Simplex LP engine. Because you cannot start a negative number of workers on a day, select **Make Unconstrained Variables Non-Negative**. After clicking **Solve**, you find the optimal solution that was shown earlier, in Figure 30-1.

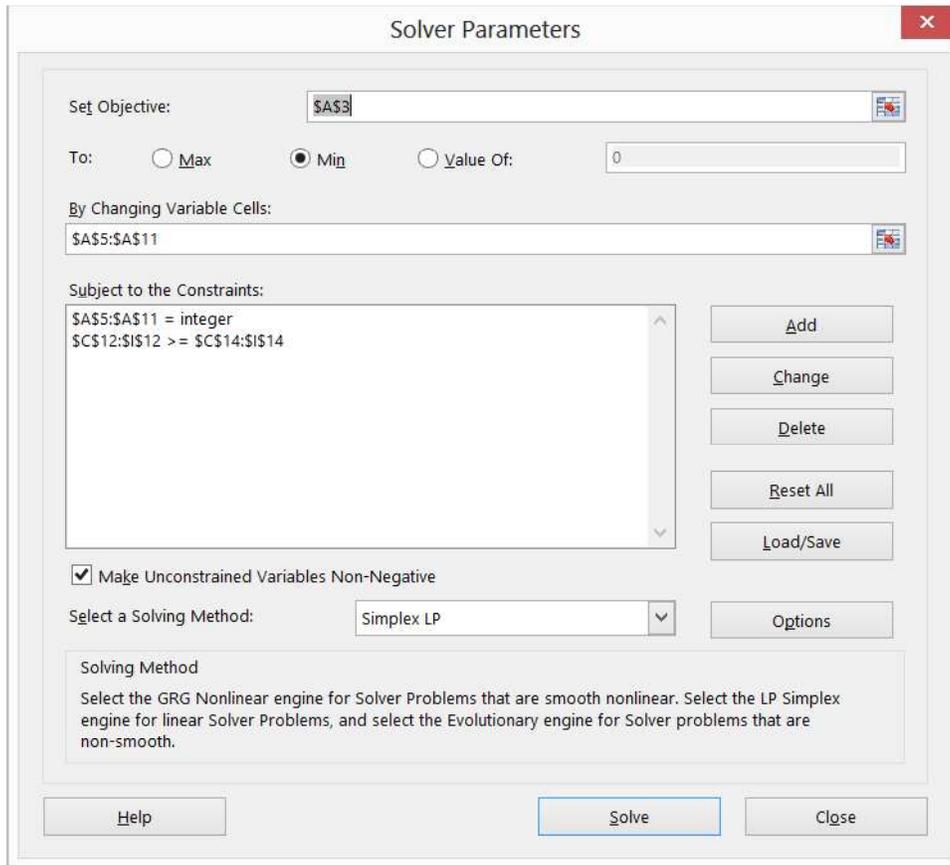


FIGURE 30-2 The Solver Parameters dialog box is filled in to solve the workforce problem.

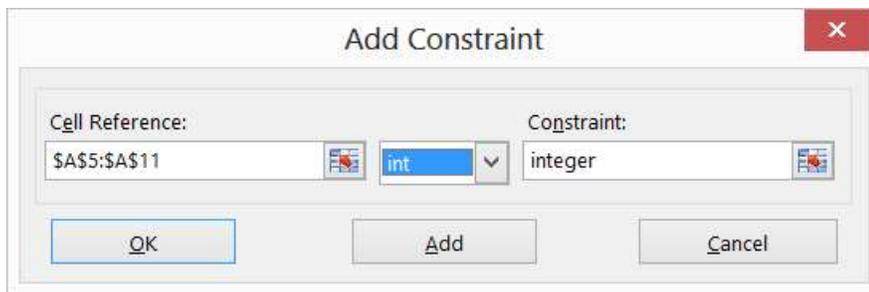


FIGURE 30-3 This constraint defines the number of workers who start each day as an integer.

A total of 20 employees is needed. One employee starts on Monday, three start on Tuesday, four start on Thursday, one starts on Friday, two start on Saturday, and nine start on Sunday. Note that this model actually has multiple optimal solutions that use 20 workers. If you run Solver again, you might very well find one of these alternative optimal solutions.

Problems

1. Suppose Bank 24 had 22 employees and that the goal was to schedule employees so that they would have the maximum number of weekend days off. How should the workers be scheduled?
2. Suppose Bank 24 employees are paid \$150 per day the first five days they work and can work a day of overtime at a cost of \$350. How should the bank schedule its employees?
3. The number of telephone reservation operators needed by an airline during each time of day is as follows:

Time	Operators needed
Midnight–4 A.M.	12
4 A.M.–8 A.M.	16
8 A.M.–noon	22
Noon–4 P.M.	30
4 P.M.–8 P.M.	31
8 P.M.–midnight	22

Each operator works one of the following six-hour shifts: midnight to 6:00 A.M., 6:00 A.M. to noon, noon to 6:00 P.M., 6:00 P.M. to midnight. What is the minimum number of operators needed?

4. Shown in Figure 30-4 are the number of people in different demographic groups who watch various TV shows and the cost (in thousands of dollars) of placing a 30-second ad on each show. For example, it costs \$160,000 to place a 30-second ad on *Friends*. The show is watched by 6 million males between the ages of 18 and 35, 3 million males between 36 and 55, 1 million males over 55, 9 million females between 18 and 35, 4 million females between 36 and 55, and 2 million females over 55. The data also includes the number of people in each group (in millions) that you want to see the ad. For example, the advertiser wants at least 60 million 18 to 35-year-old males to see its ads. What is the cheapest way to meet these goals?

	D	E	F	G	H	I	J
4	needed	60	60	28	60	60	28
5	Show	M 18-35	M 36-55	M >55	W 18-35	W 36-55	W >55
6	Friends	6	3	1	9	4	2
7	MNF	6	5	3	1	1	1
8	Malcolm	5	2	0	4	2	0
9	Sports Center	0.5	0.5	0.3	0.1	0.1	0
10	MTV	0.7	0.2	0	0.9	0.1	0
11	Lifetime	0.1	0.1	0	0.6	1.3	0.4
12	CNN	0.1	0.2	0.3	0.1	0.2	0.3
13	Jag	1	2	4	1	3	4

FIGURE 30-4 This figure shows the data for problem 4.

5. The Pine Valley Credit Union is trying to schedule bank tellers. The credit union is open from 8 A.M. to 6 P.M. and needs the following number of tellers each hour:

Needed	Time
4	8–9
8	9–10
6	10–11
4	11–12
9	12–1
8	1–2
5	2–3
4	3–4
4	4–5
5	5–6

Full-time tellers can work from 8 A.M. to 5 P.M. (with a 12 to 1 P.M. lunch hour) or from 9 A.M. to 6 P.M. (with a 1 P.M. to 2 P.M. lunch hour).

Part-time tellers work from 10 A.M. to 2 P.M. Full-time tellers receive \$300 per day, and part-time tellers \$60 per day. At most, four part-time tellers can be hired. How can the credit union minimize its daily teller salary cost?

Using Solver to solve transportation or distribution problems

Question answered in this chapter:

- How can a drug company determine at which location it should produce drugs and from which location it should ship drugs to customers?

Many companies manufacture products at different locations (often called *supply points*) and ship their products to customers (often called *demand points*). A natural question is, “What is the least expensive way to produce and ship products to customers and still meet demand?” This type of problem is called a *transportation problem*. A transportation problem can be set up as a linear Solver model with the following specifications:

- **Target cell** Minimize total production and shipping cost.
- **Changing cells** This is the amount produced at each supply point that is shipped to each demand point.
- **Constraints** The amount shipped from each supply point can’t exceed plant capacity. Each demand point must receive its required demand. In addition, each changing cell must be nonnegative.

Answer to this chapter’s question

This section provides the answer to the question that is listed at the beginning of the chapter.

How can a drug company determine at which location it should produce drugs and from which location it should ship drugs to customers?

You can follow along with this problem by looking at the *Transport.xlsx* file. Suppose a company produces a certain drug at its Los Angeles, Atlanta, and New York facilities. Each month, the Los Angeles plant can produce up to 10,000 pounds of the drug. Atlanta can produce up to 12,000 pounds, and New York can produce up to 14,000 pounds. The company must ship the number of pounds listed in cells B2:E2 each month to the four regions of the United States—East, Midwest, South, and West—as

shown in Figure 31-1. For example, the West region must receive at least 13,000 pounds of the drug each month. The cost per pound of producing a drug at each plant and shipping the drug to each region of the country is given in cells B4:E6. For example, it costs \$3.50 to produce one pound of the drug in Los Angeles and ship it to the Midwest region. What is the cheapest way to deliver the quantity of the drug each region needs?

	A	B	C	D	E	F	G	H
1								
2	DEMAND	9000	6000	6000	13000			
3		EAST	MIDWEST	SOUTH	WEST	CAPACITY		
4	LA	\$ 5.00	\$ 3.50	\$ 4.20	\$ 2.20	10000		
5	ATLANTA	\$ 3.20	\$ 2.60	\$ 1.80	\$ 4.80	12000		
6	NEW YORK CITY	\$ 2.50	\$ 3.10	\$ 3.30	\$ 5.40	14000		
7								
8	Shipments							
9		EAST	MIDWEST	SOUTH	WEST	Sent		Capacity
10	LA	0	0	0	10000	10000	<=	10000
11	ATLANTA	0	3000	6000	3000	12000	<=	12000
12	NEW YORK CITY	9000	3000	0	0	12000	<=	14000
13	Received	9000	6000	6000	13000			
14		>=	>=	>=	>=			
15	Demand	9000	6000	6000	13000			
16								
17								
18	Total Cost	\$ 86,800.00						

FIGURE 31-1 This is data for a transportation problem.

To express the target cell, you need to track total shipping cost. After entering trial values in the B10:E12 cell range for shipments from each supply point to each region, you can compute total shipping cost as follows:

$$\begin{aligned}
 & (\text{Amount sent from LA to East}) * (\text{Cost per pound of sending drug from LA to East}) \\
 & + (\text{Amount sent from LA to Midwest}) * (\text{Cost per pound of sending drug from LA to Midwest}) \\
 & + (\text{Amount sent from LA to South}) * (\text{Cost per pound of sending drug from LA to South}) \\
 & + (\text{Amount sent from LA to West}) * (\text{Cost per pound of sending drug from LA to West}) \\
 & + \dots (\text{Amount sent from New York City to West}) \\
 & * (\text{Cost per pound of sending drug from New York City to West})
 \end{aligned}$$

The *SUMPRODUCT* function can multiply corresponding elements in two rectangles (as long as the rectangles are the same size) and add the products together. The B4:E6 cell range has been named Costs and the changing cells range (B10:E12) Shipped. Therefore, total shipping and production cost is computed in cell B18 with the *SUMPRODUCT*(costs,shipped) formula.

To express the problem's constraints, first compute the total shipped from each supply point. By entering the *SUM*(B10:E10) formula in cell F10, you can compute the total number of pounds shipped from Los Angeles as (LA shipped to East) + (LA shipped to Midwest) + (LA shipped to South) + (LA shipped to West). Copying this formula to F11:F12 computes the total shipped from Atlanta and

New York City. Later, you'll add a constraint (called a *supply constraint*) that ensures that the amount shipped from each location does not exceed the plant's capacity.

Next, compute the total received by each demand point. Begin by entering the SUM(B10:B12) formula in cell B13. This formula computes the total number of pounds received in the East as (Pounds shipped from LA to East) + (Pounds shipped from Atlanta to East) + (Pounds shipped from New York City to East). By copying this formula from B13 to C13:E13, you compute the pounds of the drug received by the Midwest, South, and West regions. Later, you'll add a constraint (called a *demand constraint*) that ensures that each region receives the amount of the drug it requires.

Open the **Solver Parameters** dialog box (click **Solver** in the **Analysis** group on the **Data** tab) and fill it in as shown in Figure 31-2.

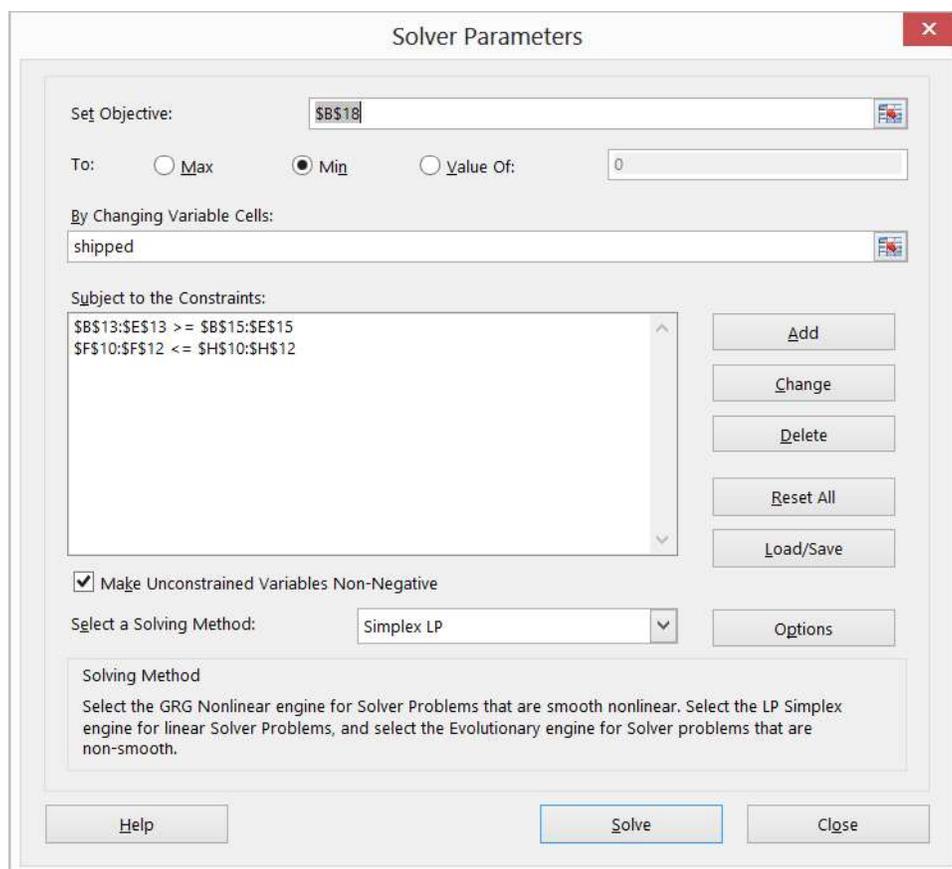


FIGURE 31-2 Solver is set up to solve your transportation problem.

You want to minimize total shipping cost (computed in cell B18). The changing cells are the number of pounds shipped from each plant to each region of the country. (These amounts are listed in the range named Shipped, consisting of cells B10:E12.) The F10:F12<=H10:H12 constraint (the supply constraint) ensures that the amount sent from each plant does not exceed its capacity. The

B13:E13>=B15:E15 constraint (the demand constraint) ensures that each region receives at least the amount of the drug it needs.

This model is a linear Solver model because the target cell is created by adding together the terms of the form (changing cell)*(constant), and both your supply and demand constraints are created by comparing the sum of changing cells to a constant. Because the model is linear, choose the Simplex LP engine. Clearly, shipments must be nonnegative, so select the **Make Unconstrained Variables Non-Negative** check box.

After clicking **Solve** in the **Solver Parameters** dialog box, you are presented with the optimal solution shown earlier, in Figure 31-1. The minimum cost of meeting customer demand is \$86,800. This minimum cost can be achieved if the company uses the following production and shipping schedule:

- Ship 10,000 pounds from Los Angeles to the West region.
- Ship 3,000 pounds from Atlanta to the West region and the same amount from Atlanta to the Midwest region. Ship 6,000 pounds from Atlanta to the South region.
- Ship 9,000 pounds from New York City to the East region and 3,000 pounds from New York City to the Midwest region.

Problems

1. The following table gives the distances between Boston, Chicago, Dallas, Los Angeles, and Miami. Each city requires 40,000 kilowatt hours (kWh) of power, and Chicago, Dallas, and Miami are capable of producing 70,000 kWh. Assume that shipping 1,000 kWh over 100 miles costs \$4. From where should power be sent to minimize the cost of meeting each city's demand?

	Boston	Chicago	Dallas	Los Angeles	Miami
Chicago	983	0	1205	2112	1390
Dallas	1815	1205	0	801	1332
Miami	1539	1390	1332	2757	0

2. Resolve this chapter's example assuming that demand in the West region increases to 13,000.
3. Your company produces and sells drugs at several locations. The decision of where to produce goods for each sales location can have a huge impact on profitability. The model here is similar to the model used in this chapter to determine where drugs should be produced. Use the following assumptions:
 - You produce drugs at six locations and sell to customers in six areas.

- The tax rate and variable production cost depend on where the drug is produced. For example, any units produced at Location 3 cost \$6 per unit to produce; profits from these goods are taxed at 20 percent.
- The sales price of each drug depends on where the drug is sold. For example, each product sold in Location 2 is sold for \$40:

Production location	1	2	3	4	5	6
Sales price	\$45	\$40	\$38	\$36	\$39	\$34
Tax rate	31%	40%	20%	40%	35%	18%
Variable production cost	\$8	\$7	\$6	\$9	\$7	\$7

- Each of the six plants can produce up to 6 million units per year.
- The annual demand (in millions) for your product in each location is as follows:

Sales location	1	2	3	4	5	6
Demand	1	2	3	4	5	6

- The unit shipping cost depends on the plant where the product is produced and where the product is sold:

	Sold 1	Sold 2	Sold 3	Sold 4	Sold 5	Sold 6
Plant 1	\$3	\$4	\$5	\$6	\$7	\$8
Plant 2	\$5	\$2	\$6	\$9	\$10	\$11
Plant 3	\$4	\$3	\$1	\$6	\$8	\$6
Plant 4	\$5	\$5	\$7	\$2	\$5	\$5
Plant 5	\$6	\$9	\$6	\$5	\$3	\$7
Plant 6	\$7	\$7	\$8	\$9	\$10	\$4

For example, if you produce a unit at Plant 1 and sell it in Location 3, it costs \$5 to ship it.

How can you maximize after-tax profit with your limited production capacity?

4. Suppose that each day, northern, central, and southern California each uses 100 billion gallons of water. Also, assume that northern California and central California have 120 billion gallons of water available, whereas southern California has 40 billion gallons of water available. The cost of shipping 1 billion gallons of water between the three regions is as follows:

	Northern	Central	Southern
Northern	\$5,000	\$7,000	\$10,000
Central	\$7,000	\$5,000	\$6,000
Southern	\$10,000	\$6,000	\$5,000

You will not be able to meet all demand for water, so assume that each billion gallons of unmet demand incurs the following shortage costs:

	Northern	Central	Southern
Shortage cost/billion gallons short	\$6,000	\$5,500	\$9,000

How should California's water be distributed to minimize the sum of shipping and shortage costs?

Using Solver for capital budgeting

Question answered in this chapter:

- How can a company use Solver to determine which projects it should undertake?

Each year, a company such as Eli Lilly needs to determine which drugs to develop; a company such as Microsoft, which software programs to develop; a company such as Proctor & Gamble, which new consumer products to develop. Solver can help a company make these decisions.

Answer to this chapter's question

This section provides the answer to the question that is listed at the beginning of the chapter.

How can a company use Solver to determine which projects it should undertake?

Most corporations want to undertake projects that contribute the greatest net present value (NPV), subject to limited resources (usually capital and labor). Say that a software development company is trying to determine which of 20 software projects it should undertake. The NPV (in millions of dollars) contributed by each project as well as the capital (in millions of dollars) and the number of programmers needed during each of the next three years is given on the Basic Model worksheet in the Capbudget.xlsx file, shown in Figure 32-1. For example, Project 2 yields \$908 million. It requires \$151 million during Year 1, \$269 million during Year 2, and \$248 million during Year 3. Project 2 requires 139 programmers during Year 1, 86 programmers during Year 2, and 83 programmers during Year 3. Cells E4:G4 show the capital (in millions of dollars) available during each of the three years, and cells H4:J4 indicate how many programmers are available. For example, during Year 1, up to \$2.5 billion in capital and 900 programmers are available.

	A	B	C	D	E	F	G	H	I	J
1		Total NPV								
2		9293		Used	2460	2684	2742	876	895	702
3					<=	<=	<=	<=	<=	<=
4				Available	2500	2800	2900	900	900	900
5	Do IT?	NPV			Cost Year 1	Cost Year 2	Cost Year 3	Labor Year 1	Labor Year 2	Labor Year 3
6	0	Project 1	928		398	180	368	111	108	123
7	1	Project 2	908		151	269	248	139	86	83
8	1	Project 3	801		129	189	308	56	61	23
9	0	Project 4	543		275	218	220	54	70	59
10	0	Project 5	944		291	252	228	123	141	70
11	1	Project 6	848		80	283	285	119	84	37
12	1	Project 7	545		203	220	77	54	44	42
13	1	Project 8	808		150	113	143	67	101	43
14	1	Project 9	638		282	141	160	37	55	64
15	1	Project 10	841		214	254	355	130	72	62
16	0	Project 11	664		224	271	130	51	79	58
17	0	Project 12	546		225	150	33	35	107	63
18	0	Project 13	699		101	218	272	43	90	71
19	1	Project 14	599		255	202	70	3	75	83
20	1	Project 15	903		228	351	240	60	93	80
21	1	Project 16	859		303	173	431	60	90	41
22	0	Project 17	748		133	427	220	59	40	39
23	0	Project 18	668		197	98	214	95	96	74
24	1	Project 19	888		313	278	291	66	75	74
25	1	Project 20	655		152	211	134	85	59	70

FIGURE 32-1 This is data you will use with Solver to determine which projects to undertake.

The company must decide whether it should undertake each project. Assume that the company can't undertake a fraction of a software project; if 0.5 of the needed resources is allocated, for example, the company would have a nonworking program that would bring in \$0 revenue!

The trick in modeling situations in which you either do or don't do something is to use *binary changing cells*. A binary changing cell always equals 0 or 1. When a binary changing cell that corresponds to a project equals 1, you do the project. If a binary changing cell that corresponds to a project equals 0, you don't do the project. You set up Solver to use a range of binary changing cells by adding a constraint; select the changing cells you want to use and then choose **Bin** from the list in the **Add Constraint** dialog box.

With this background, you're ready to solve the software project selection problem. As with any Solver model, you should begin by identifying the target cell, the changing cells, and the constraints:

- **Target cell** Maximize the NPV generated by selected projects.
- **Changing cells** Look for a 0 or 1 binary changing cell for each project. Locate these cells in the A6:A25 range (named the *doit* range). For example, a 1 in cell A6 indicates that you undertake Project 1; a 0 in cell A6 indicates that you don't undertake Project 1.
- **Constraints** You need to ensure that for each Year t ($t = 1, 2, 3$), Year t capital used is less than or equal to Year t capital available, and Year t labor used is less than or equal to Year t labor available.

As you can see, the worksheet must compute the NPV, the capital used annually, and the programmers used each year for any selection of projects. In cell B2, use the `SUMPRODUCT(doit, NPV)` formula to compute the total NPV generated by selected projects. (The NPV range name refers to the C6:C25

range.) For every project with a 1 in column A, this formula picks up the NPV of the project, and for every project with a 0 in column A, this formula does not pick up the NPV of the project. Therefore, you can compute the NPV of all projects, and the target cell is linear because it is computed by summing terms that follow the (changing cell)*(constant) form. In a similar fashion, compute the capital used each year and the labor used each year by copying the *SUMPRODUCT(doit,E6:E25)* formula from E2 to F2:J2.

Now fill in the **Solver Parameters** dialog box as shown in Figure 32-2.

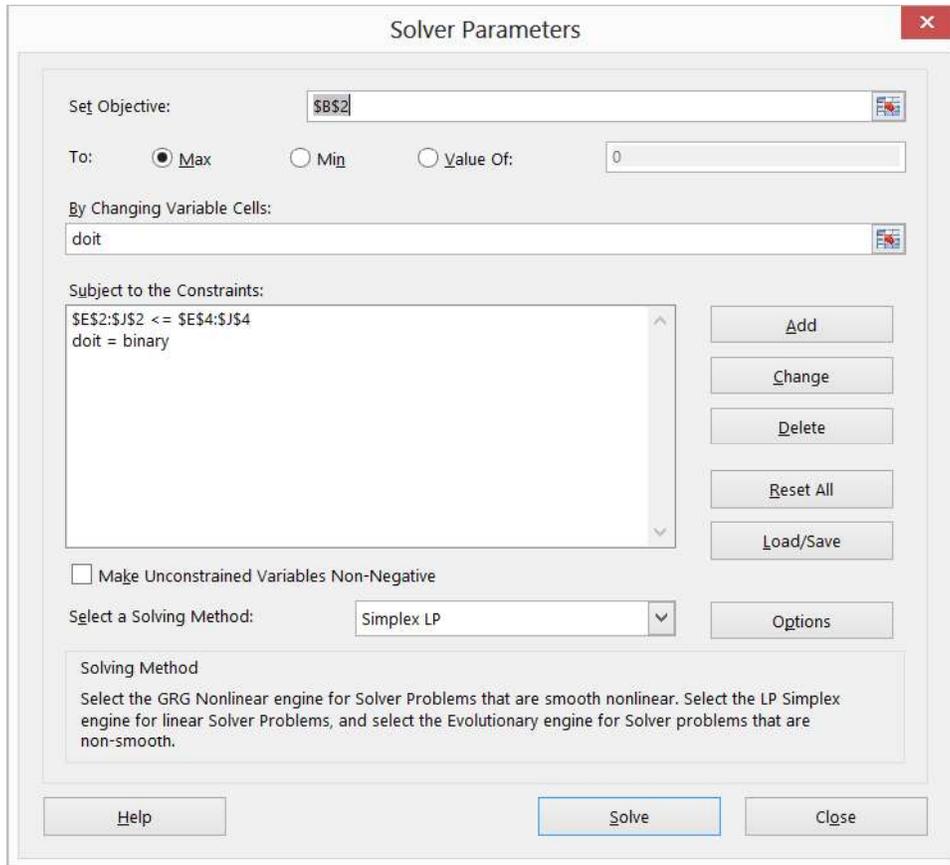


FIGURE 32-2 The Solver Parameters dialog box is set up for the project selection model.

The goal is to maximize the NPV of selected projects (cell B2). The changing cells (the range named *doit*) are the binary changing cells for each project. The $E2:J2 \leq E4:J4$ constraint ensures that during each year, the capital and labor used are less than or equal to the capital and labor available. To add the constraint that makes the changing cells binary, click **Add** in the **Solver Parameters** dialog box and then select **Bin** from the list in the middle of the **Add Constraint** dialog box. The **Add Constraint** dialog box should appear as shown in Figure 32-3.



FIGURE 32-3 Use Bin in the Add Constraint dialog box to set up binary changing cells—cells that display either a 0 or a 1.

This model is linear because the target cell is computed as the sum of terms that have the (changing cell)*(constant) form and because the resource usage constraints are computed by comparing the sum of (changing cells)*(constants) to a constant. Therefore, select the Simplex LP engine.

With the **Solver Parameters** dialog box filled in, click **Solve** and get the results shown earlier, in Figure 32-1. The company can obtain a maximum NPV of \$9,293 billion (\$9.293 billion) by choosing Projects 2, 3, 6–10, 14–16, 19, and 20.

Handling other constraints

Sometimes project selection models have other constraints. For example, suppose that if you select Project 3, you must also select Project 4. Because the current optimal solution selects Project 3 but not Project 4, this tells you that the current solution can't remain optimal. To solve this problem, add the constraint that the binary changing cell for Project 3 is less than or equal to the binary changing cell for Project 4.

You can find this example on the If 3 then 4 worksheet in the Capbudget.xlsx file, which is shown in Figure 32-4. Cell L9 refers to the binary value related to Project 3, and cell L12 to the binary value related to Project 4. By adding the $L9 \leq L12$ constraint, if you choose Project 3, L9 equals 1, and this constraint forces L12 (the Project 4 binary) to equal 1. This constraint must also leave the binary value in the changing cell of Project 4 unrestricted if you do not select Project 3. If you do not select Project 3, L9 equals 0, and the constraint allows the Project 4 binary to equal 0 or 1, which is what you want. The new optimal solution is shown in Figure 32-4.

Now suppose that you can do only four projects from among Projects 1 through 10. (See the At Most 4 Of P1–P10 worksheet, shown in Figure 32-5.) In cell L8, you compute the sum of the binary values associated with Projects 1 through 10 with the $SUM(A6:A15)$ formula. Then add the $L8 \leq L10$ constraint, which ensures that at most, 4 of the first 10 projects are selected. The new optimal solution is shown in Figure 32-5. The NPV has dropped to \$9.014 billion.

	A	B	C	D	E	F	G	H	I	J	K	L
1		Total NPV										
2		9157	Used		2444	2760	2837	866	895	659		
3				<=	<=	<=	<=	<=	<=			
4			Available		2500	2800	2900	900	900	900		
5	Do IT?	NPV		Cost Year 1	Cost Year 2	Cost Year 3	Labor Year 1	Labor Year 2	Labor Year 3			
6	0	Project 1	928	398	180	368	111	108	123			
7	1	Project 2	908	151	269	248	139	86	83			
8	1	Project 3	801	129	189	308	56	61	23	Proj 3		
9	1	Project 4	543	275	218	220	54	70	59		1	
10	0	Project 5	944	291	252	228	123	141	70	<=		
11	1	Project 6	848	80	283	285	119	84	37	Proj 4		
12	1	Project 7	545	203	220	77	54	44	42		1	
13	1	Project 8	808	150	113	143	67	101	43			
14	1	Project 9	638	282	141	160	37	55	64			
15	0	Project 10	841	214	254	355	130	72	62			
16	0	Project 11	664	224	271	130	51	79	58			
17	0	Project 12	546	225	150	33	35	107	63			
18	0	Project 13	699	101	218	272	43	90	71			
19	0	Project 14	599	255	202	70	3	75	83			
20	1	Project 15	903	228	351	240	60	93	80			
21	1	Project 16	859	303	173	431	60	90	41			
22	1	Project 17	748	133	427	220	59	40	39			
23	1	Project 18	668	197	98	214	95	96	74			
24	1	Project 19	888	313	278	291	66	75	74			
25	0	Project 20	655	152	211	134	85	59	70			

FIGURE 32-4 This is the new optimal solution for If Project 3's selection requires we do Project 4.

	A	B	C	D	E	F	G	H	I	J	K	L
1		Total NPV										
2		9014	Used		2378	2734	2755	778	896	702		
3				<=	<=	<=	<=	<=	<=			
4			Available		2500	2800	2900	900	900	900		
5	Do IT?	NPV		Cost Year 1	Cost Year 2	Cost Year 3	Labor Year 1	Labor Year 2	Labor Year 3			
6	0	Project 1	928	398	180	368	111	108	123			
7	0	Project 2	908	151	269	248	139	86	83			
8	1	Project 3	801	129	189	308	56	61	23			
9	0	Project 4	543	275	218	220	54	70	59			
10	0	Project 5	944	291	252	228	123	141	70			
11	0	Project 6	848	80	283	285	119	84	37			
12	1	Project 7	545	203	220	77	54	44	42			
13	1	Project 8	808	150	113	143	67	101	43			
14	0	Project 9	638	282	141	160	37	55	64			
15	1	Project 10	841	214	254	355	130	72	62			
21	1	Project 16	859	303	173	431	60	90	41			
22	1	Project 17	748	133	427	220	59	40	39			
23	1	Project 18	668	197	98	214	95	96	74			
24	1	Project 19	888	313	278	291	66	75	74			
25	1	Project 20	655	152	211	134	85	59	70			

FIGURE 32-5 Optimal solution when only 4 of Projects 1–10 can be selected.

Solving binary and integer programming problems

Linear Solver models in which some or all changing cells are required to be binary or an integer are usually harder to solve than linear models in which all changing cells are allowed to be fractions. For this reason, analysts are often satisfied with a near-optimal solution to a binary or integer programming problem. If your Solver model runs for a long time, you might want to consider adjusting the Integer Optimality (formerly called Tolerance) setting in the Solver Options dialog box. (See Figure 32-6.) For example, a Tolerance setting of .5 means that Solver will stop the first time it finds a feasible solution that is within 0.5 percent of the theoretical optimal target cell value. (The theoretical optimal target cell value is the optimal target value found when the binary and integer constraints are

omitted.) Often you'll be faced with a choice between finding an answer within 10 percent of optimal in 10 minutes or finding an optimal solution in two weeks of computer time! The default Integer Optimality value is 5%, which means that Solver stops when it finds a target cell value within 5 percent of the theoretical optimal target cell value. When the software development example was first solved, the Integer Optimality was set to 5% and found an optimal target cell value of 9269. When the Integer Optimality value was changed to 0.50%, a better target cell value (9293) was obtained.

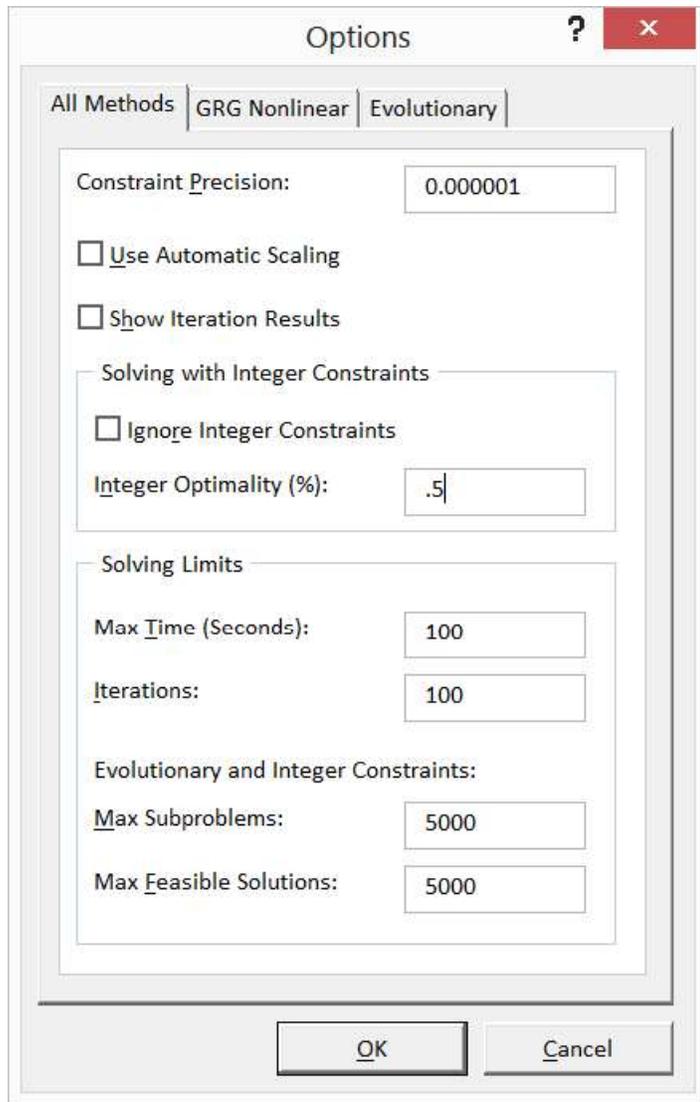


FIGURE 32-6 This figure shows adjustment of the Integer Optimality option.

Problems

1. A company has nine projects under consideration. The NPV added by each project and the capital required for each project during the next two years is shown in the following table. (All numbers are in millions.) For example, Project 1 will add \$14 million in NPV and require expenditures of \$12 million during Year 1 and \$3 million during Year 2. During Year 1, \$50 million in capital is available for projects, and \$20 million is available during Year 2:

	NPV	Year 1 expenditure	Year 2 expenditure
Project 1	14	12	3
Project 2	17	54	7
Project 3	17	6	6
Project 4	15	6	2
Project 5	40	32	35
Project 6	12	6	6
Project 7	14	48	4
Project 8	10	36	3
Project 9	12	18	3

- If you can't undertake a fraction of a project but must undertake either all or none of it, how can you maximize NPV?
 - Suppose that if Project 4 is undertaken, Project 5 must be undertaken. How can you maximize NPV?
2. A publishing company is trying to determine which of 36 books it should publish this year. The `Pressdata.xlsx` file gives the following information about each book:
- Projected revenue and development costs (in thousands of dollars)
 - Pages in each book
 - Whether the book is geared toward an audience of software developers (indicated by a 1 in column E)
- The company can publish books with a total of up to 8,500 pages this year and must publish at least four books geared toward software developers. How can the company maximize its profit?
3. In the SEND + MORE = MONEY equation, each letter represents a different digit from 0–9. Which digit is associated with each letter?

4. Jill is trying to determine her class schedule for the next semester. A semester consists of two seven-week half semesters. Jill must take four courses during each half semester. There are five time slots during each semester. Of course, Jill cannot take the same course twice. Jill has associated a value with each course and time slot. This data is in the `Classdata.xlsx` file. For example, course 1 during time slot 5 in semester 1 has a value of 5. Which courses should Jill take during each semester to maximize her total value from the semester's courses?
5. Use conditional formatting to highlight in yellow fill each row corresponding to a selected project.
6. Use Solver to determine the minimum number of coins needed to make change for 92 cents.

Using Solver for financial planning

Questions answered in this chapter:

- Can I use Solver to verify the accuracy of the Excel *PMT* function or to determine mortgage payments for a variable interest rate?
- Can I use Solver to determine how much money I need to save for retirement?

Solver can be a powerful tool for analyzing financial planning problems. In many of these types of problems, a quantity such as the unpaid balance on a loan or the amount of money needed for retirement changes over time. For example, consider a situation in which you borrow money. Because only the noninterest portion of each monthly payment reduces the unpaid loan balance, you know that the following equation (Equation 1) is true:

$$\begin{aligned} (\text{Unpaid loan balance at end of period } t) &= (\text{Unpaid loan balance at beginning of period } t) \\ &- [(\text{Month } t \text{ payment}) - (\text{Month } t \text{ interest paid})] \end{aligned}$$

Now suppose that you are saving for retirement. Until you retire, you deposit at the beginning of each period (say *periods* equal *years*) an amount of money in your retirement account, and during the year, your retirement fund is invested and receives a return of some percentage. During retirement, you withdraw money at the beginning of each year, and your retirement fund still receives an investment return. You know that the following equation (Equation 2) describes the relationship among contributions, withdrawals, and return:

$$\begin{aligned} (\text{Retirement savings at end of Year } t+1) &= (\text{Retirement savings at end of Year } t + \text{retirement} \\ &\text{contribution at beginning of Year } t+1 - \text{Year } t+1 \text{ retirement withdrawal}) \\ &*(\text{Investment return earned during Year } t+1) \end{aligned}$$

Combining basic relationships such as these with Solver enables you to answer a myriad of interesting financial planning problems.

Answers to this chapter's questions

This section provides the answers to the questions that are listed at the beginning of the chapter.

Can I use Solver to verify the accuracy of the Excel *PMT* function or to determine mortgage payments for a variable interest rate?

Recall that in Chapter 9, “More Excel financial functions,” you found the monthly payment (assuming payments occur at the end of a month) on a 10-month loan for \$8,000 at an annual interest rate of 10 percent to be \$1,037.03. Could you have used Solver to determine your monthly payment? You’ll find the answer in the PMT By Solver worksheet in the Finmathsolver.xlsx file, which is shown in Figure 33-1.

	A	B	C	D	E
1			rate	0.00667	
2					
3		From PMT function	\$1,037.03		
4	Month	Beginning Balance	Payment	Interest Owed	Ending Balance
5	1	\$ 10,000.00	\$ 1,037.03	\$ 66.67	\$ 9,029.63
6	2	\$ 9,029.63	\$ 1,037.03	\$ 60.20	\$ 8,052.80
7	3	\$ 8,052.80	\$ 1,037.03	\$ 53.69	\$ 7,069.45
8	4	\$ 7,069.45	\$ 1,037.03	\$ 47.13	\$ 6,079.55
9	5	\$ 6,079.55	\$ 1,037.03	\$ 40.53	\$ 5,083.05
10	6	\$ 5,083.05	\$ 1,037.03	\$ 33.89	\$ 4,079.90
11	7	\$ 4,079.90	\$ 1,037.03	\$ 27.20	\$ 3,070.07
12	8	\$ 3,070.07	\$ 1,037.03	\$ 20.47	\$ 2,053.51
13	9	\$ 2,053.51	\$ 1,037.03	\$ 13.69	\$ 1,030.16
14	10	\$ 1,030.16	\$ 1,037.03	\$ 6.87	\$ 0.00

FIGURE 33-1 This is the Solver model for calculating the monthly payment for a loan.

The key to this model is to use Equation 1 (shown earlier in the chapter) to track the monthly beginning balance. The Solver target cell is to minimize the monthly payment. The changing cell is the monthly payment. The only constraint is that the ending balance in Month 10 equals 0.

Enter the beginning balance in cell B5. You can enter a trial monthly payment in cell C5 and then copy the monthly payment to the C6:C14 range. Because you assume that the payments occur at the end of each month, interest is incurred on the balance at the beginning of the month. The monthly interest rate (cell C1 is named Rate) is computed in D1 by dividing the annual rate of 0.08 by 12. The interest paid each month is computed by copying the rate*B5 formula from cell D5 to D6:D14. Each month, this formula computes the interest as .006666*(month’s beginning balance). By copying the (B5–(Payment–D5)) formula from cell E5 to E6:E14, you use Equation 1 to compute each month’s ending balance. Because (Month $t+1$ beginning balance) = (Month t ending balance), each month’s beginning balance is computed by copying the =E5 formula from cell B6 to B7:B14.

You are now ready to use Solver to determine the monthly payment. To see how you’ve set up the Solver Parameters dialog box, take a look at Figure 33-2.

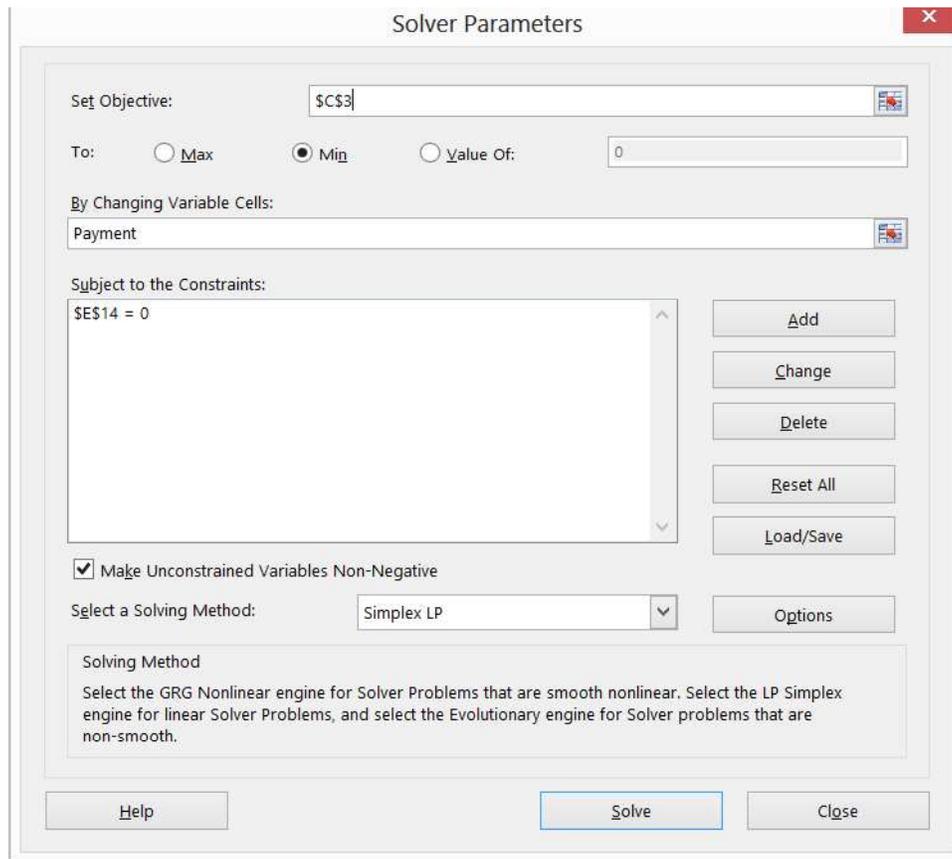


FIGURE 33-2 The Solver Parameters dialog box is set up to determine mortgage payments.

The goal is to minimize the monthly payment (cell C5). Note that the changing cell is the same as the target cell. The only constraint is that the ending balance for Month 10 must equal 0. Adding this constraint ensures that the loan is paid off. After you choose the Simplex LP engine and select the nonnegative variables option, Solver calculates a payment of \$1,037.03, which matches the amount calculated by the Excel *PMT* function.

This model is linear because the target cell equals the changing cell, and the constraint is created by adding multiples of changing cells.

It should be mentioned that when Solver models involve very large or very small numbers, Solver sometimes thinks models that *are* linear are *not* linear. To avoid this problem, it is good practice to select **Use Automatic Scaling** in the **Options** dialog box. This should ensure that Solver properly recognizes linear models as being linear.

Can I use Solver to determine how much money I need to save for retirement?

By using Equation 2 (shown earlier in the chapter), you can easily determine how much money a person needs to save for retirement. Here's an example.

You are planning for your retirement, and at the beginning of this year and each of the next 39 years, you will contribute some money to your retirement fund. Each year, you plan to increase your retirement contribution by \$500. When you retire in 40 years, you plan to withdraw (at the beginning of each year) \$100,000 per year for 20 years. You make the following assumptions about the yields for your retirement investment portfolio:

- During the first 20 years of your investing, the investments will earn 10 percent per year.
- During all other years, your investments will earn 5 percent per year.

You assume that all contributions and withdrawals occur at the beginning of the year. Given these assumptions, what is the least amount of money you can contribute this year and still have enough to make your retirement withdrawals?

You can find the solution to this question on the Retire worksheet in the Finmathsolver.xlsx file, shown in Figure 33-3. Note that many rows are hidden in the model.

This worksheet simply tracks your retirement balance during each of the next 60 years. Each year, you earn the indicated interest rate on the retirement balance. Begin by entering a trial value for the Year 1 payment in cell C6. Copying the C6+500 formula from cell C7 to C8:C45 ensures that the retirement contribution increases by \$500 per year during Years 2 through 40. Enter in column D the assumed return on your investments for each of the next 60 years. In cells E46:E65, enter the annual \$100,000 withdrawal for Years 41 through 60. Copying the (B6+C6-E6)*(1+D6) formula from F6 to F7:F65 uses Equation 2 to compute each year's ending retirement account balance. Copying the =F6 formula from cell B7 to B8:B65 computes the beginning balance for Years 2 through 60. Of course, the Year 1 initial balance is 0. Note that the 6.8704E-07 value in cell F65 is approximately 0, with the difference the result of a rounding error.

	A	B	C	D	E	F
5	Year	Initial balance	Contribution	Return	Withdrawal	Ending Balance
6	1	\$0.00	\$ 1,387.87	10%	\$0.00	\$1,526.65
7	2	\$1,526.65	\$ 1,887.87	10%	\$0.00	\$3,755.98
8	3	\$3,755.98	\$ 2,387.87	10%	\$0.00	\$6,758.23
44	39	\$1,146,596.10	\$ 20,387.87	5%	\$0.00	\$1,225,333.17
45	40	\$1,225,333.17	\$ 20,887.87	5%	\$0.00	\$1,308,532.09
46	41	\$1,308,532.09		5%	\$100,000.00	\$1,268,958.69
47	42	\$1,268,958.69		5%	\$100,000.00	\$1,227,406.62
62	57	\$372,324.80		5%	\$100,000.00	\$285,941.04
63	58	\$285,941.04		5%	\$100,000.00	\$195,238.10
64	59	\$195,238.10		5%	\$100,000.00	\$100,000.00
65	60	\$100,000.00		5%	\$100,000.00	\$0.00

FIGURE 33-3 Here is retirement planning data that can be set up for analysis with Solver.

The **Solver Parameters** dialog box for this model is shown in Figure 33-4. You want to minimize your Year 1 contribution (cell C6). The changing cell is also your Year 1 contribution (cell C6). Ensure that you never run out of money during retirement by adding the F46:F65 \geq 0 constraint so that the ending balance for Years 41 through 60 is nonnegative.

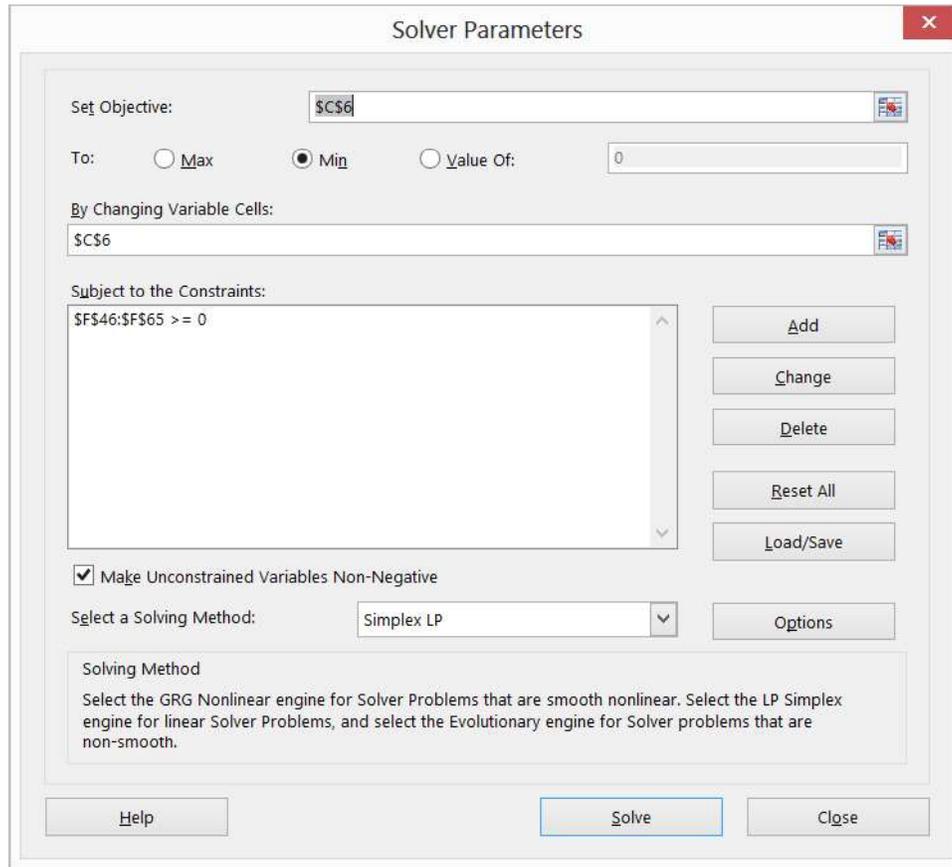


FIGURE 33-4 The Solver Parameters dialog box is set up for the retirement problem.

After choosing the Simplex LP engine and selecting **Make Unconstrained Variables Non-Negative** in the **Solver Parameters** dialog box, click **Solve** in the **Solver Parameters** dialog box; you find that the first year's contribution should equal \$1,387.87.

This model is linear because the target cell equals the changing cell, and the constraint is created by adding multiples of changing cells. Note that because the return on the investments is not the same each year, there is no easy way to use Excel financial functions to solve this problem. Solver provides a general framework that can be used to analyze financial planning problems when mortgage rates or investment returns are not constant.

Problems

1. I am borrowing \$15,000 to buy a new car. I will make 60 end-of-month payments. The annual interest rate on the loan is 10 percent. The car dealer is a friend of mine, and he will allow me to make the monthly payment for Months 1–30 equal to one-half the payment for Months 31 through 60. What is the payment during each month?
2. Solve the retirement planning problem, assuming that withdrawals occur at the end of each year and contributions occur at the beginning of each year.
3. Solve the mortgage example, assuming that payments are made at the beginning of each month.
4. In the retirement-planning example, suppose that during Year 1, your salary is \$40,000 and your salary increases 5 percent per year until retirement. You want to save the same percentage of your salary each year you work. What percentage of your salary should you save?
5. In the mortgage example, suppose that you want your monthly payment to increase by \$50 each month. What should each month's payment be?
6. Assume you want to take out a \$300,000 loan on a 20-year mortgage with end-of-month payments. The annual rate of interest is 6 percent. Twenty years from now, you need to make an ending balloon payment of \$40,000. Because you expect your income to increase, you want to structure the loan so that at the beginning of each year, your monthly payments increase by 2 percent. Determine the amount of each year's monthly payment.
7. Blair's mother is saving for Blair's college education. The following payments must be made at the indicated times:

4 years from now	5 years from now	6 years from now	7 years from now
\$24,000	\$26,000	\$28,000	\$30,000

The following investments are available:

- Today, one year from now, two years from now, three years from now, and four years from now, she can invest money for one year and receive a 6 percent return.
- Today, two years from now, and four years from now, she can invest money for two years and receive a 14 percent return.
- Three years from now, she can invest money for three years and receive an 18 percent return.
- Today, she can invest money for seven years and receive a 65 percent return.

What is the minimum amount that Blair's mother needs to commit today to Blair's college education that ensures that she can pay her college bills?

8. I owe \$10,000 on one credit card that charges 18 percent annual interest and \$5,000 on another credit card that charges 12 percent annual interest. Interest for the month is based on the month's beginning balance. I can afford to make total payments of \$2,000 per month, and the minimum monthly payment on each card is 10 percent of the card's unpaid balance at the beginning of the month. My goal is to pay off both cards in two years. What is the minimum amount of interest I need to pay?

Using Solver to rate sports teams

Question answered in this chapter:

- Can I use Excel to set NFL point spreads?

Many of us follow basketball, football, hockey, or baseball. Oddsmakers set point spreads on games in all these sports and others. For example, the bookmakers' best guess was that the Indianapolis Colts would win the 2010 Super Bowl by 7 points. Instead, the New Orleans Saints won the game. In this chapter, you see that Solver predicted that the Saints were the better team and should have been favored. Now see how Solver can estimate the relative ability of NFL teams accurately.

Using a simple Solver model, you can generate reasonable point spreads for games based on the scores of the 2009 season. The work is in file `Nfl2009april2010.xlsx`, shown in Figure 34-1. You use the score of each game of the 2009 NFL season as input data. The changing cell for the Solver model is a rating for each team and the size of the home field advantage. For example, if the Indianapolis Colts have a rating of +5 and the New York Jets have a rating of +7, the Jets are considered two points better than the Colts.

With regard to the home-field edge, in most years, college and professional football teams, as well as professional basketball teams, tend to win by an average of three points (whereas home college basketball teams tend to win by an average of five points). In your model, however, you will define the home edge as a changing cell and have Solver estimate the home edge. You can define the outcome of an NFL game to be the number of points by which the home team outscores the visitors and predict the outcome of each game by using the following equation (Equation 1):

(Predicted points by which home team outscores visitors) = (Home edge) + (Home team rating) – (Away Team rating)

For example, if the home-field edge equals three points, when the Colts host the Jets, the Colts will be a one-point favorite ($3 + 5 - 7$). If the Jets host the Colts, the Jets will be a five-point favorite ($3 + 7 - 5$). (The Tampa Bay–New England game was played in London, so there is no home edge for this game.)

What target cell will yield reliable ratings? The goal is to find the set of values for team ratings and home-field advantage that best predicts the outcome of all games. In short, you want the prediction for each game to be as close as possible to the outcome of each game. This suggests that you want to minimize the sum over all games of *(Actual outcome) – (Predicted outcome)*. However, the problem with using this target is that positive and negative prediction errors cancel each other out. For example, if you over-predict the home-team margin by 50 points in one game and under-predict the

home-team margin by 50 points in another game, the target cell would yield a value of 0, indicating perfect accuracy, when in fact you were off by 50 points a game. You can remedy this problem by minimizing the sum over all games by using the $[(Actual\ Outcome) - (Predicted\ Outcome)]^2$ formula. Now positive and negative errors will not cancel each other out.

Answer to this chapter's question

This section provides the answer to the question that is listed at the beginning of the chapter.

Can I use Excel to set NFL point spreads?

See now how to determine accurate ratings for NFL teams by using the scores from the 2009 regular season. You can find the data for this problem in the `Nfl2009april2010.xlsx` file, which is shown in Figure 34-1.

	B	C	D	E	F	G	H	I	J	K
2										SSE
3										43108
4			Week Home		Away	Home Points	Away Points	Home margin	Prediction	Squared Error
5			1	Pittsburgh Steelers	Tennessee Titans	13	10	3	6.722344	13.856
6			1	Atlanta Falcons	Miami Dolphins	19	7	12	5.637121	40.486
7	home		1	Seattle Seahawks	St. Louis Rams	28	0	28	10.3964	309.89
8		2.260510793	1	New York Giants	Washington Redskins	23	17	6	6.910514	0.829
9			1	Baltimore Ravens	Kansas City Chiefs	38	24	14	18.15126	17.233
10	mean	-6E-10	1	Houston Texans	New York Jets	7	24	-17	-4.35591	159.87
11	Team	rating	1	New Orleans Saints	Detroit Lions	45	27	18	27.40592	88.471
12	Pittsburgh Steelers	1.682	1	Tampa Bay Buccaneers	Dallas Cowboys	21	34	-13	-10.4081	6.7179
13	Atlanta Falcons	5.046	1	Arizona Cardinals	San Francisco 49ers	16	20	-4	1.925441	35.111
14	Seattle Seahawks	-9.301	1	Carolina Panthers	Philadelphia Eagles	10	38	-28	0.172907	793.71
15	New York Giants	0.103	1	Green Bay Packers	Chicago Bears	21	15	6	13.51688	56.503
16	Baltimore Ravens	7.465	1	Cincinnati Bengals	Denver Broncos	7	12	-5	2.606708	57.862
17	Houston Texans	1.946	1	Cleveland Browns	Minnesota Vikings	20	34	-14	-13.2951	0.4969
18	New Orleans Saints	10.77	1	Indianapolis Colts	Jacksonville Jaguars	14	12	2	14.6823	160.84
19	Tampa Bay Buccaneers	-5.5	1	Oakland Raiders	San Diego Chargers	20	24	-4	-14.6496	113.41
20	Arizona Cardinals	-0.281	1	New England Patriots	Buffalo Bills	25	24	1	15.16782	200.73
21	Carolina Panthers	3.92	2	Tennessee Titans	Houston Texans	31	34	-3	-2.46553	0.2857
22	Green Bay Packers	7.371	2	Kansas City Chiefs	Oakland Raiders	10	13	-3	4.101831	50.436

FIGURE 34-1 This is the data rating NFL teams that you'll use with Solver.

To begin, place a trial home-field advantage value in cell B8.

Starting in row 5, columns E and F contain the home and away teams for each game. For example, the first game (listed in row 5) is Tennessee playing at Pittsburgh. Column G contains the home team's score, and column H contains the visiting team's score. As you can see, the Steelers beat the Titans 13–10. You can now compute the outcome of each game (the number of points by which the home team beats the visiting team) by entering the `=G5–H5` formula in cell I5. By pointing to the lower-right portion of this cell and double-clicking, you can copy this formula down to the last regular season game, which appears in row 260. (By the way, an easy way to select all the data is to press `Ctrl+Shift+down arrow`. This key combination takes you to the last row filled with data—row 260 in this case.)

In column J, use Equation 1 to generate the prediction for each game. The prediction for the first game is computed in cell J5 as follows:

$$=B\$8+VLOOKUP(E5,B\$12:C\$43,2,FALSE)-VLOOKUP(F5,B\$12:C\$43,2,FALSE)$$

This formula creates a prediction for the first game by adding the home edge to the home-team rating and then subtracting the visiting-team rating. (Note that in row 103, the term $B\$8$ is deleted from the formula because there was no home edge in the New England–Tampa Bay game.) The term $VLOOKUP(E5,B\$12:C\$43,2,FALSE)$ locates the home-team rating, and $VLOOKUP(F5,B\$12:C\$43,2,FALSE)$ looks up the visiting team's rating. (For more information about using lookup functions, see Chapter 2, "Lookup functions.") In column K, compute the squared error $(actual\ score - predicted\ score)^2$ for each game. Your squared error for the first game is computed in cell K5 with the $=(I5 - J5)^2$ formula. After selecting the I5:K5 cell range, you can double-click and copy the formulas down to row 260.

Next, compute the target cell in cell K3 by summing all the squared errors with the $SUM(J5:J260)$ formula.



Tip You can enter a formula for a large column of numbers such as this by typing $=SUM($ and then selecting the first cell in the range you want to add together. Press $Ctrl+Shift+down$ arrow to enter the range from the cell you selected to the bottom row in the column and then add the closing parenthesis.

It is convenient to make the average team rating equal to 0. A team with a positive rating is better than average and a team with a negative rating is worse than average. Compute the average team rating in cell C10 with the $AVERAGE(C12:C43)$ formula.

You can now fill in the **Solver Parameters** dialog box as shown in Figure 34-2.

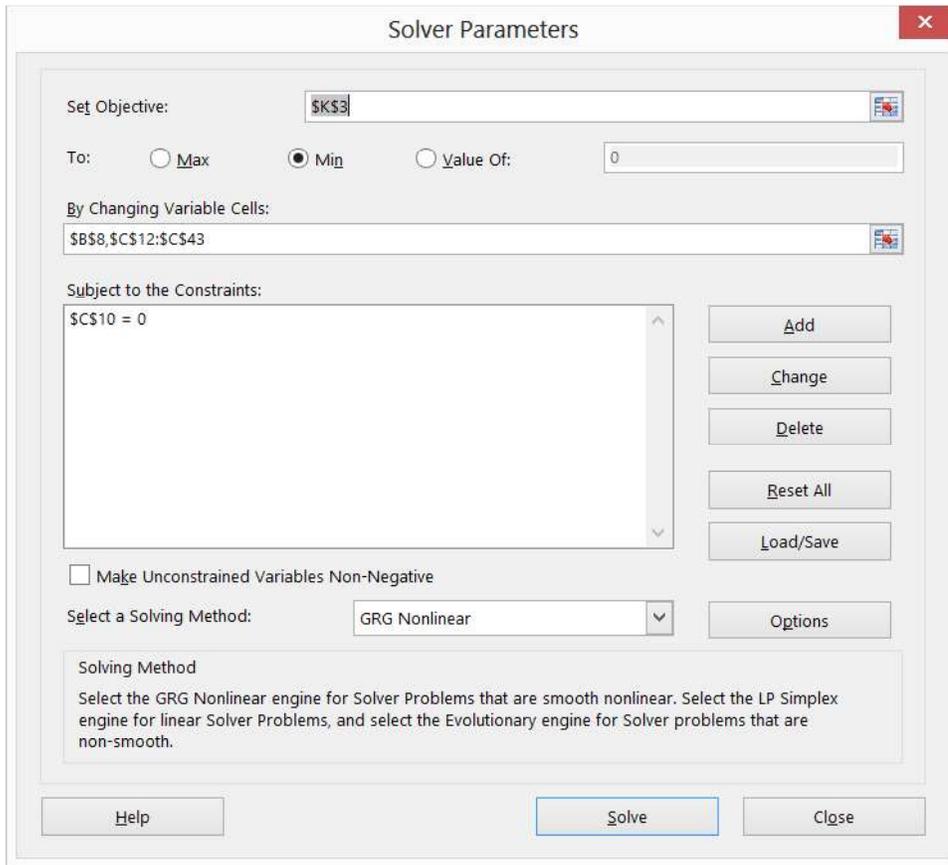


FIGURE 34-2 The Solver Parameters dialog box is set up for NFL ratings.

Minimize the sum of the squared prediction errors for all games (computed in cell K3) by changing each team's rating (listed in cells C12:C43) and the home advantage (cell B8). The C10=0 constraint ensures that the average team rating is 0. From Figure 34-1, you can see that the home team has an advantage of 2.26 points over the visiting team. The 15 highest-rated teams are shown in Figure 34-3. Remember that the ratings listed in cell range E3:E34 are computed by Solver. In the template file, you can start with any numbers in these cells, and Solver will still find the best ratings.

	L	M	N
12	Rank	Team	rating
13		1 New England Patriots	11.0695
14		2 New Orleans Saints	10.7697
15		3 New York Jets	8.56256
16		4 Baltimore Ravens	7.4648
17		5 Green Bay Packers	7.37132
18		6 Minnesota Vikings	7.16877
19		7 Dallas Cowboys	7.16853
20		8 San Diego Chargers	6.6428
21		9 Philadelphia Eagles	6.00733
22		10 Indianapolis Colts	5.91202
23		11 Atlanta Falcons	5.04647
24		12 Carolina Panthers	3.91973
25		13 Houston Texans	1.94614
26		14 Pittsburgh Steelers	1.68193
27		15 Miami Dolphins	1.66985

FIGURE 34-3 These are the top 15 teams for the NFL 2009 season.

These ratings have the Saints around 5 points better than the Colts, so this model would have predicted (before the playoffs) that the Saints would beat the Colts by 5 points.

Why is your model not a linear Solver model?

This model is not linear because the target cell adds together terms of the (Home Team Rating + Home-Field Edge – Visiting Team Rating)² form. Recall that for a Solver model to be linear, the target cell must be created by adding together terms with the (changing cell)*(constant) form. This relationship doesn't exist in this case, so the model is not linear. Solver does obtain the correct answer, however, for any sports-rating model in which the target cell minimizes the sum of squared errors. Note that the GRG nonlinear engine was chosen because this model is not linear and did not involve nonmathematical functions such as IF statements. I did not select Make Unconstrained Variables Non-Negative because to have the team ratings average 0, you must allow some of the team ratings to be negative.



Note Recently, I found that the GRG Solver engine works poorly when automatic scaling is checked. I recommend opening the **Options** dialog box and clearing **Use Automatic Scaling**.

Problems

- 1-4.** The Nfl0x.xlsx ($x = 1, 2, 3, 4$) files contain scores for every regular season game during the 200x NFL season. Rate the teams for each season. During each season, which teams would you forecast to have made the Super Bowl?
- 5.** For the 2004 season, devise a method to predict the actual score of each game. Hint: Give each team an offensive rating and a defensive rating. Who had the best offense? Who had the best defense?
- 6.** True or false? An NFL team could lose every game and be an above-average team.
- 7.** The Nba01_02.xlsx file contains scores for every game during the 2001–2002 NBA season. Rate the teams.
- 8.** The Nba02_03.xlsx file contains scores for every regular season game during the 2002–2003 NBA season. Rate the teams.
- 9.** The Worldball.xlsx file contains all scores from the 2006 World Basketball Championships. Rate the teams. Who were the best three teams?
- 10.** This method of rating teams works fine for football and basketball. What problems arise if you apply these methods to hockey or baseball?
- 11.** The NFL2012data.xlsx file contains scores of all NFL 2012 regular season games. Rate the teams. Even though the Colts were 10–6, your ratings have the Colts as a well-below-average team. Can you explain this anomaly?

Warehouse location and the GRG Multistart and Evolutionary Solver engines

Questions answered in this chapter:

- Where in the United States should an Internet shipping company locate a single warehouse to minimize the total distance that packages are shipped?
- Where in the United States should an Internet shipping company locate two warehouses to minimize the total distance that packages are shipped?

In Microsoft Excel 2013, Solver has been blessed with many new exciting capabilities. This chapter (and Chapter 36, “Penalties and the Evolutionary Solver,” and Chapter 37, “The traveling salesperson problem”) explains how these algorithms can help you solve many important optimization problems.

Understanding the GRG Multistart and Evolutionary Solver engines

As was pointed out in Chapter 28, “Introducing optimization with Excel Solver,” the Excel 2013 Solver uses three engines to solve optimization problems: Simplex LP, GRG Nonlinear, and Evolutionary. The following sections provide more details about how the latter two of these engines are used to solve optimization problems.

How does Solver solve linear Solver problems?

As was pointed out in Chapters 28 through 33, a Solver model is linear if all references to changing cells in the target cells and constraints are created by adding together terms of the (changing cells)*(constants) form. For linear models, you should always select the Simplex LP engine, which is designed to find solutions to linear Solver models efficiently. The Excel 2013 Solver can handle problems with up to 200 changing cells and 100 constraints. Versions of Solver that can handle larger problems are available from the Solver.com website.

How does the GRG Nonlinear engine solve nonlinear optimization models?

If your target cell, any of your constraints, or both contain references to changing cells that are not of the (changing cell)*(constant) form, you have a nonlinear model. If x and y are changing cells, references such as the following in the target cell, any constraints, or both make your model nonlinear:

- x^2
- xy
- $\sin x$
- e^x
- xe^{2y}

If your nonlinear formulas involve ordinary math operators such as the previous examples, proper use of the GRG Nonlinear engine should quickly find the optimal solution to your Solver model. To illustrate how the GRG Nonlinear engine works, suppose you want to maximize $-x^2 + 4x + 2$. This function is graphed in Figure 35-1.

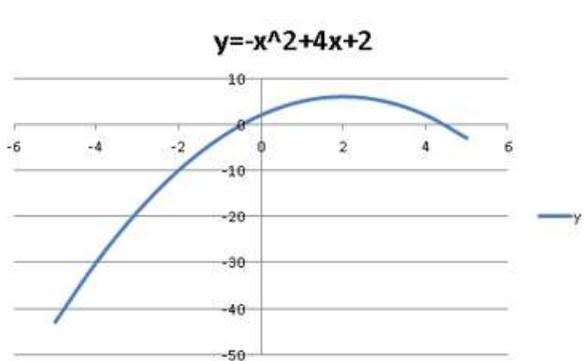


FIGURE 35-1 This is how the GRG Nonlinear engine maximizes a function.

You can see that this function is maximized for $x = 2$. Notice also that for $x = 2$, the function has a slope of 0. The GRG Nonlinear engine solves this problem by trying to find a point at which the slope of the function is 0. Similarly, if you want to minimize $y = x^2$, the GRG Nonlinear engine solves this problem by determining that the slope of this function is 0 for $x = 0$. See Figure 35-2.

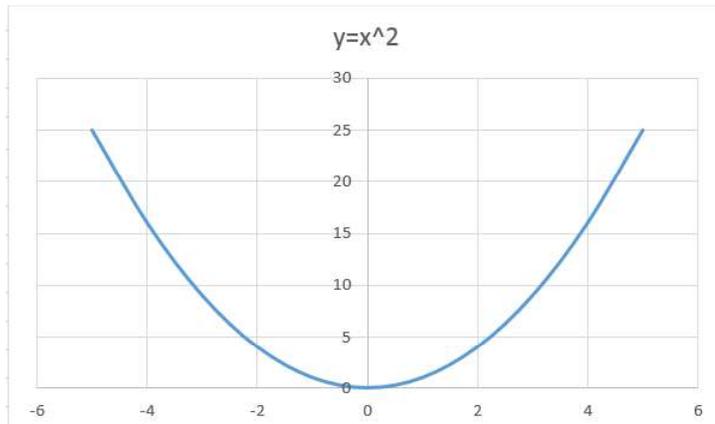


FIGURE 35-2 This is how the GRG Nonlinear engine minimizes a function.

Unfortunately, many functions cannot be maximized simply by locating a point where the function's slope equals 0. For example, suppose you want to maximize the function shown in Figure 35-3, when x ranges between -5 and $+10$.

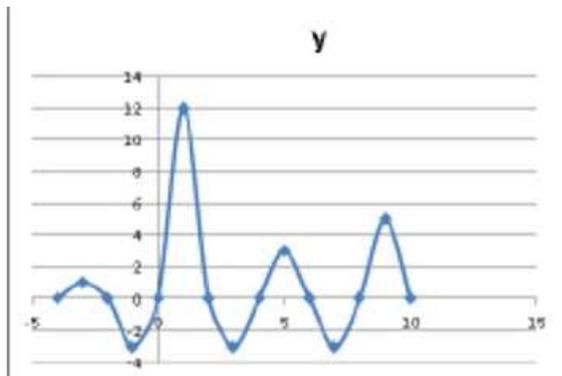


FIGURE 35-3 This is maximizing a function with multiple peaks.

You can see that this function has more than one peak. If you start with a value of x near 1, you will find the right solution to the ($x = 1$) problem. If you start near another peak—say near $x = 5$ —you will find a solution of $x = 5$, which is incorrect. Because in most problems (especially those with more than one changing cell) you do not know a good starting point, it appears you have a major hurdle to clear. Fortunately, Excel 2013 has a Multistart option. You can select **Multistart** after choosing **Options** and then clicking the **GRG Nonlinear** tab. When the Multistart option is selected, Excel chooses many starting solutions and finds the best answer after beginning with these starting points. This approach usually resolves the multiple peak and valley problem.

By the way, pressing Esc stops Solver. Also, keep in mind that GRG Multistart works best when you place reasonable upper and lower bounds on your changing cells. (For example, you do not specify changing cell ≤ 100 million.)

The GRG engine also runs into trouble if the target cell, constraints, or both use nonsmooth functions such as *MAX*, *MIN*, *ABS*, *IF*, *SUMIF*, *COUNTIF*, *SUMIFS*, *COUNTIFS*, and others that involve changing cells. These functions create points where there is no uniquely defined slope because the slope changes abruptly. For example, suppose an optimization problem requires you to model the value of a European call option with a \$40 exercise price. This call option enables you to buy the stock for \$40. If the stock price is s at expiration of the option, then the value of the call option may be computed with the $\max(0, s-40)$ or $\text{IF}(s>40, s-40, 0)$ formula relationship graphed in Figure 35-4. It is clear that when $s = 40$, the option value has no slope, so the GRG engine would break down.

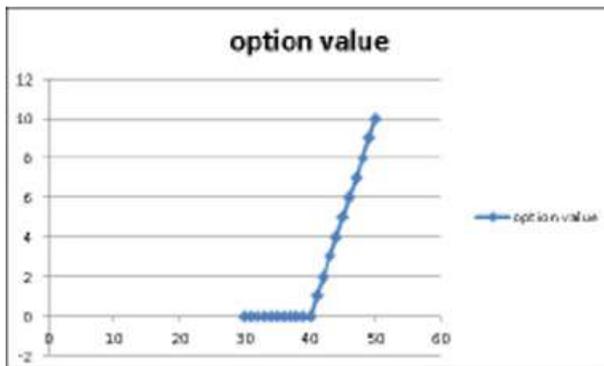


FIGURE 35-4 The option value has no slope for a \$40 stock price.

As Figure 35-5 shows, Solver models that include the absolute value function (recall that the absolute value of a number is just the distance of the number from 0) will have no slope for $x = 0$. In Excel, the *ABS(x)* function returns the absolute value of a number x .

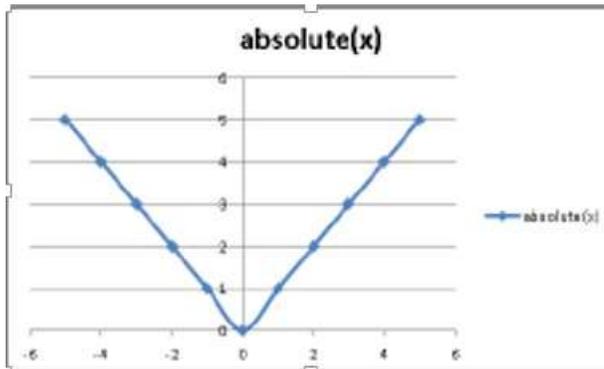


FIGURE 35-5 The absolute value function has no slope for $x = 0$.

Optimization problems in which the target cell, any of the constraints, or both have no slope for any changing cell values are called *nonsmooth* optimization problems. Even GRG Multistart has difficulty with these types of problems. In these situations, you should apply the Solver Evolutionary engine. For nonlinear Solver models, Solver is limited to 100 changing cells and 100 constraints.

How does the Evolutionary Solver engine tackle nonsmooth optimization problems?

Evolutionary Solver in Excel 2013 is based on genetic algorithms, a concept discovered by John Holland, a computer science professor at the University of Michigan. To use the Evolutionary Solver, begin by taking 50 to 100 points in the problem's feasible region (that is, the set of points that meet the constraints). This set of points is called the *population*. Then, the target cell is evaluated for each point. Using the idea of survival of the fittest from the theory of evolution, you change the points in the population in a way that increases the likelihood that future population members are located near previous population members that have a good target cell value. Because this approach is based on target cell values and not on slopes, multiple peaks and valleys pose no problem. Also, functions that do not have slopes (the so-called nonsmooth functions) also become a less important issue. The Evolutionary Solver engine (like GRG Multistart) also works best when reasonable upper and lower bounds are placed on your changing cells. After you select the Evolutionary Solver engine, it's best to choose **Options**, click the **Evolutionary** tab, and change the mutation rate to .5. Also, select the **Required Bounds On Variables** check box and increase the maximum time without improvement to 3,600 seconds. Increasing the mutation rate decreases the likelihood that the Solver gets stuck near a poor solution. Increasing the maximum time without improvement to 3,600 seconds enables the Solver to run until it fails to improve the target cell for 3,600 seconds. That way, the Solver keeps running if you leave your computer.

Now, use Solver to solve two interesting facility location problems.

Answers to this chapter's questions

This section provides the answers to the questions that are listed at the beginning of the chapter.

Where in the United States should an Internet shipping company locate a single warehouse to minimize the total distance that packages are shipped?

The number of shipments (in thousands) made each year to various cities is shown in Figure 35-6. (See the One warehouse worksheet in the Warehouseloc.xlsx file.)

	A	B	C	D	E	F	G	H	I
1									
2									
3						Lat	Long		Mean
4					1	36.813439	92.48191		1125.827
5							Total Shipped*	252185.2	
6		City	Lat	Long	Shipments	Distance	Dist		
7		New York	40.7	73.9	15	1309.8969	19648.45		
8		Boston	42.3	71	8	1529.8326	12238.66		
9		Philadelphia	40	75.1	10	1219.3395	12193.39		
10		Charlotte	35.2	80.8	6	813.70342	4882.221		
11		Atlanta	33.8	84.4	11	595.15484	6546.703		
12		New Orleans	30	89.9	8	502.75019	4022.002		
13		Miami	25.8	80.2	13	1138.2724	14797.54		
14		Dallas	32.8	96.8	10	406.77005	4067.701		
15		Houston	29.8	95.4	12	524.14383	6289.726		
16		Chicago	41.8	87.7	14	476.71185	6673.966		
17		Detroit	42.4	83.1	11	753.42785	8287.706		
18		Cleveland	41.5	81.7	8	811.19304	6489.544		
19		Indy	39.8	86.1	7	486.18479	3403.294		
20		Denver	39.8	105	8	881.28021	7050.242		
21		Minneapolis	45	93.3	9	567.68613	5109.175		
22		Phoenix	33.5	112	11	1372.8197	15101.02		
23		Salt Lake City	40.8	112	10	1367.7932	13677.93		
24		LA	34.1	118	18	1798.1222	32366.2		
25		SF	37.8	123	12	2079.2628	24951.15		
26		SD	32.8	117	10	1721.0736	17210.74		
27		Seattle	41.6	122	13	2090.6012	27177.82		

FIGURE 35-6 Here is data for the single warehouse problem.

A key to this model is the following formula, which gives the approximate distance between two US cities having a latitude and longitude given by (Lat1, Long1) and (Lat2, Long2).

$$Distance = 69 * \sqrt{(Lat1 - Lat2)^2 + (Long1 - Long2)^2}$$

To begin, enter trial values in cells F4:G4 for the latitude and longitude of the warehouse. Next, by copying the $69 * \text{SQRT}((C7 - \text{F}\$4)^2 + (D7 - \text{G}\$4)^2)$ formula from F7 to F8:F27, you compute the approximate distance of each city from the warehouse. Next, copying the $E7 * F7$ formula from G7 to G8:G27 computes the distance traveled by the shipments to each city. In cell H5, the $\text{SUM}(G7:G27)$ formula computes the total distance traveled by all shipments. Your target cell is to minimize H5 by changing F4:G4. After you select the GRG Nonlinear engine, the **Solver Parameters** dialog box appears, as in Figure 35-7.

After you click **Solve**, you'll find that the warehouse should be located at 36.81 degrees latitude and 92.48 degrees longitude, which is near Springfield, Missouri. (See Figure 35-6.)

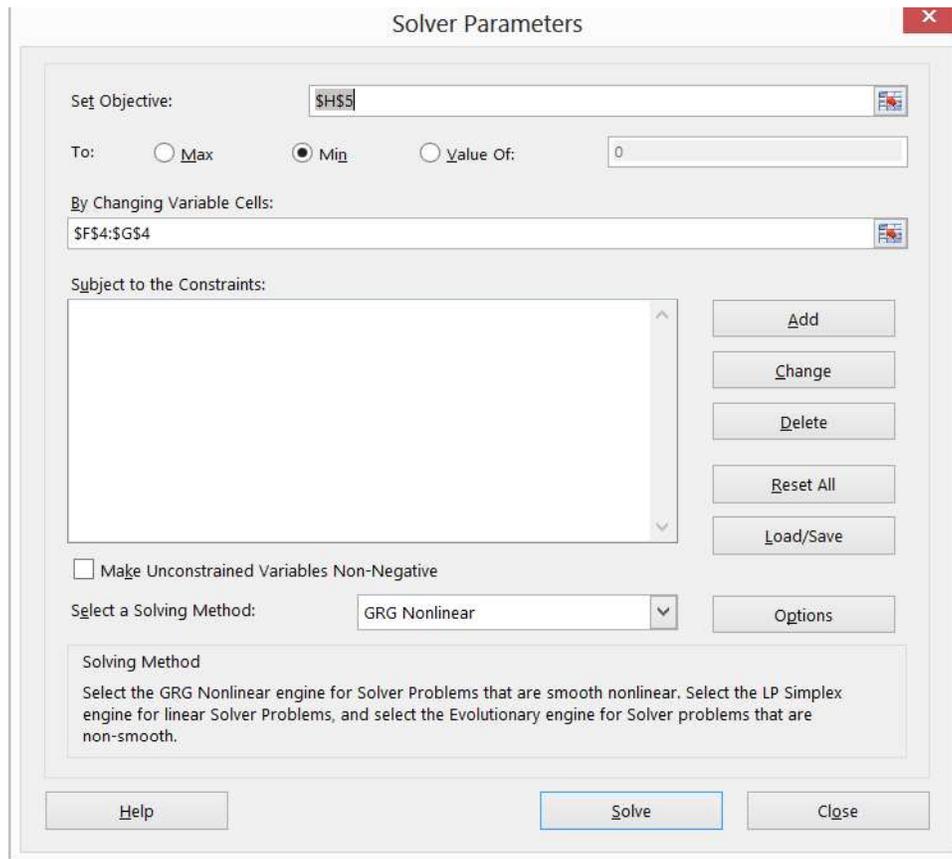


FIGURE 35-7 This is the Solver Parameters dialog box for the One warehouse problem.

Where in the United States should an Internet shipping company locate two warehouses to minimize the total distance that packages are shipped?

The work for this problem is in the Two warehouses worksheet in the Warehouseloc.xlsx file, shown in Figure 35-8.

To begin, enter trial latitudes and longitudes for the warehouses in F4:G5. Next, copy the $69 \cdot \text{SQRT}((C7 - \$F\$4)^2 + (D7 - \$G\$4)^2)$ formula from F7 to F8:F27 to compute the distance of each city from Warehouse 1. By copying the $69 \cdot \text{SQRT}((C7 - \$F\$5)^2 + (D7 - \$G\$5)^2)$ formula from G7 to G8:G27, you compute the distance from each city to Warehouse 2. Because the shipments from each city will be sent from the *closer* warehouse, you now compute the distance of each city to the closer warehouse by copying the $\text{MIN}(F7, G7)$ formula from H7 to H8:H27. In I7:I27, you compute the distance traveled by each city's shipments by copying the $H7 \cdot E7$ formula from I7 to I8:I27. In cell I5, you compute the total distance traveled by shipments with the $\text{SUM}(I7:I27)$ formula.

	A	B	C	D	E	F	G	H	I
1									
2									Mean distance
3						Lat	Long		501.811289
4					1	38.16405	84.02896		Total
5					2	34.93189	117.7916		119676.4655
6		City	Lat	Long	Shipments	Distance to 1	Distance to 2	Min Distance	Dist*Shipped
7		New York	40.7	73.9	15	720.47	3054.56	720.47	10807.05051
8		Boston	42.3	71	8	943.2073	3268.403	943.2073	7545.658519
9		Philadelphia	40	75.1	10	628.9874	2966.405	628.9874	6289.873795
10		Charlotte	35.2	80.8	6	302.4357	2552.487	302.4357	1814.614483
11		Atlanta	33.8	84.4	11	302.206	2305.344	302.206	3324.265857
12		New Orleans	30	89.9	8	693.8561	1954.375	693.8561	5550.848907
13		Miami	25.8	80.2	13	893.0923	2669.257	893.0923	11610.19974
14		Dallas	32.8	96.8	10	955.7745	1455.871	955.7745	9557.744684
15		Houston	29.8	95.4	12	973.9953	1585.079	973.9953	11687.94305
16		Chicago	41.8	87.7	14	356.5146	2129.715	356.5146	4991.204988
17		Detroit	42.4	83.1	11	299.2264	2448.557	299.2264	3291.490274
18		Cleveland	41.5	81.7	8	280.7258	2531.222	280.7258	2245.806491
19		Indy	39.8	86.1	7	182.1067	2212.369	182.1067	1274.747093
20		Denver	39.8	104.9	8	1444.519	950.8286	950.8286	7606.628829
21		Minneapolis	45	93.3	9	794.7959	1827.139	794.7959	7153.163389
22		Phoenix	33.5	112.1	11	1963.455	404.9579	404.9579	4454.537037
23		Salt Lake Cit	40.8	111.9	10	1931.683	573.7616	573.7616	5737.616195
24		LA	34.1	118.4	18	2388.123	71.11337	71.11337	1280.04075
25		SF	37.8	122.6	12	2661.52	386.3183	386.3183	4635.819785
26		SD	32.8	117.1	10	2311.723	154.6474	154.6474	1546.474352
27		Seattle	41.6	122.4	13	2658.195	559.2874	559.2874	7270.736808

FIGURE 35-8 This figure shows a model for locating two warehouses.

You're now ready to use Solver to determine the optimal warehouse locations. The setup for the Solver Parameters dialog box is shown in Figure 35-9.

Begin by selecting the GRG Nonlinear engine and then use the poor solution, which places each warehouse at 0 latitude and longitude. This solution is poor for two reasons: It locates the warehouses in Africa and it puts two warehouses in the same place. After running Solver, you find that Solver recommends locating both warehouses in the same place. Of course, this is a suboptimal solution. The problem is twofold: The *MIN* function creates situations with no slopes, and perhaps your target cell, as a function of the four changing cells, has multiple peaks and valleys. If your target cell has multiple peaks and valleys (in four dimensions), perhaps your poor starting solution is not near the lowest valley, which is the true optimal solution. When you suspect multiple peaks and valleys exist, it is a good idea to use GRG Multistart, which tries multiple starting points and finds the best answer from each. Most of the time, the best of the best Multistart finds will be the optimal solution to the problem. To use Multistart, place upper and lower bounds on the changing cells. For the bounds on latitude changing cells, you could select 0 and 90 degrees. This ensures that the warehouse is north of the equator. For the bounds on longitude changing cells, choose 0 and 150, which ensures that your location is west of Greenwich, England, and east of Anchorage, Alaska.

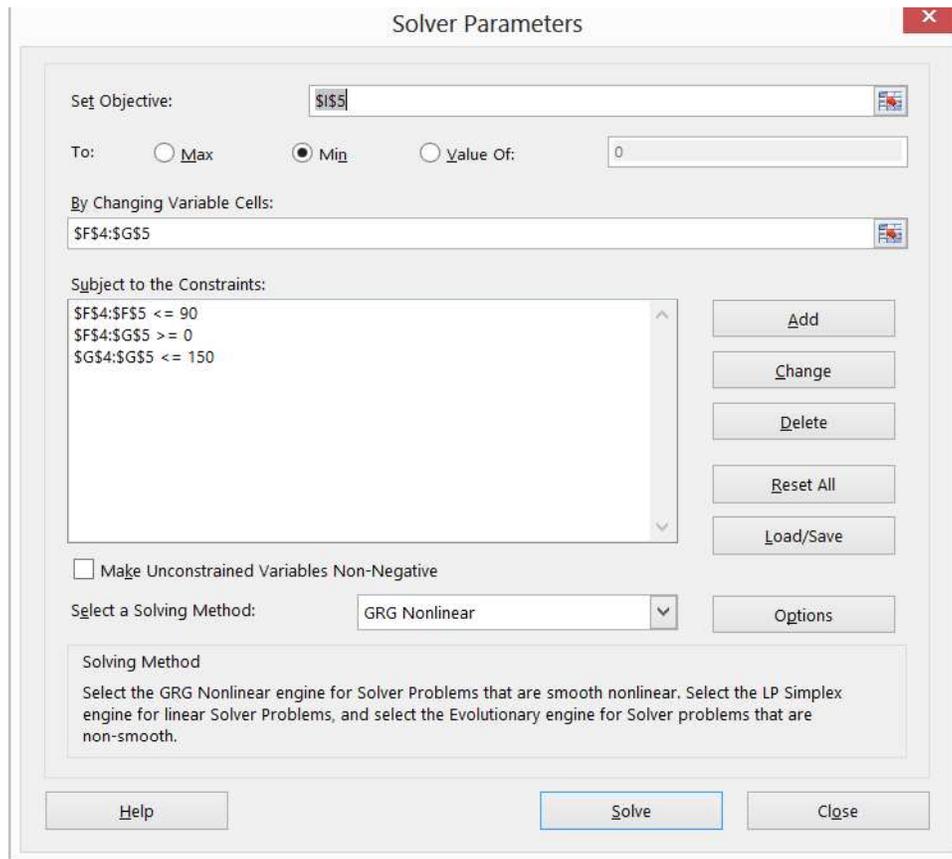


FIGURE 35-9 Solver is set up for locating two warehouses.

After running the GRG Multistart engine, you find the total distance traveled was 119,676 miles, and the average distance traveled per shipment is 502 miles. The locations of the warehouses are shown in cells F4:G5 of Figure 35-8. Warehouse 1 is located near Lexington, Kentucky; Warehouse 2 is located near Lancaster, California.

To confirm that Solver found the optimal solution, run the Evolutionary Solver engine. You find no improvement in the optimal solution.

Suppose you had set an upper bound for longitude of 110 degrees. After running Solver, you would have found that Solver recommends a longitude near 110 degrees. If you place bounds on a changing cell and the Solver forces the changing cell to assume a value near a bound, you should relax the bound.

Problems

1. Find the optimal solution to the warehouse problem if three warehouses are allowed.
2. Suppose you want to locate a single restroom so that company employees have to travel the smallest possible distance per day when going to the bathroom. Employees work in four locations within the plant as described in the following table:

X	Y	Number of employees
5	20	6
50	50	12
25	75	23
80	30	15

Assume that employees always walk in a north–south or east–west direction when going to and from the restroom. Where should the restroom be located? Solve Problem 2 if the company wants to locate two restrooms.

Penalties and the Evolutionary Solver

Questions answered in this chapter:

- What are the keys to using the Evolutionary Solver successfully?
- How can I use the Evolutionary Solver to assign 80 workers in Microsoft Finance to a job in one of four workgroups?

Answers to this chapter's questions

This section provides the answers to the questions that are listed at the beginning of the chapter.

What are the keys to using the Evolutionary Solver successfully?

Previously, this book stated that the Evolutionary Solver should be used to find solutions to optimization problems in which the target cell, changing cells, or both involve nonsmooth functions such as *IF*, *ABS*, *MAX*, *MIN*, *COUNTIF*, *COUNTIFS*, *SUMIF*, *SUMIFS*, *AVERAGEIF*, and *AVERAGEIFS*. Before solving a problem with the Evolutionary Solver, you should do the following in the **Solver Parameters** dialog box:

- Click **Options**, select the **Evolutionary** tab, and increase **Mutation Rate** to .50. Increasing the mutation rate enables the Solver to jump around in the set of possible solutions and avoid being stuck in a portion of the set of possible solutions that does not contain a good solution to the Solver model.
- Change **Maximum Time Without Improvement** to 3,600 seconds. Increasing Maximum Time Without Improvement ensures that if you leave your PC, Microsoft Excel will keep looking for a solution for 3,600 seconds. You can stop the Solver at any time by pressing the Esc key.
- Place reasonable lower and upper bounds on your changing cells. Placing bounds on the changing cells reduces the size of the region in which Solver searches for an optimal solution. This can greatly speed Solver's progress toward an optimal solution. The user can change many other Evolutionary Solver settings, but Mutation Rate and Maximum Time Without Improvement have been found to be the only settings that have a significant impact on the solution process.

Everything in life has an upside and a downside, and the Evolutionary Solver is no exception. The upside of the Evolutionary Solver is that it handles nonsmooth functions well. The downside is that constraints that are not linear functions of the changing cells are not handled very well. To handle most constraints with the Evolutionary Solver, you should penalize the target cell to make violation of a constraint a bad thing. Then the survival of the fittest will do away with any constraint violation. The chapter's next question shows how to use penalties with the Evolutionary Solver.

How can I use the Evolutionary Solver to assign 80 workers in Microsoft Finance to a job in one of four workgroups?

You need to assign 80 employees to four workgroups. The head of each workgroup has rated each employee's competence on a 0 to 10 scale (10 equals most competent). Each employee has rated his satisfaction with each job assignment (again on a 0 to 10 scale). For example, Worker 1 has been given a 9 rating from the leader of Workgroup 1, and Worker 1 gives Workgroup 4 a rating of 7.

The work for this question is in the Assign.xlsx file. (See Figure 36-1.) You want to assign between 18 and 22 people to each workgroup. You consider job competence to be twice as important as employee satisfaction. How can you assign employees to workgroups to maximize total satisfaction and ensure that each division has the required number of employees?

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O
1			Qual				Sat				576	499			
2	Division	Worker	1	2	3	4	1	2	3	4	Quality	Satisfaction			
3	4	1	9	8	6	8	1	2	6	7	8	7			
4	1	2	10	0	5	6	9	6	7	4	10	9			
5	3	3	5	8	10	5	1	7	7	3	10	7	Group #assigned	penalty	
6	1	4	4	0	5	2	9	1	0	3	4	9	1	19	0
7	2	5	9	10	4	5	9	8	8	3	10	8	2	18	0
8	3	6	5	2	7	3	2	8	1	5	7	1	3	22	0
9	1	7	8	3	1	2	1	8	2	2	8	1	4	21	0
10	3	8	2	2	9	2	8	3	1	6	9	1	Total pen		0
11	1	9	8	7	6	3	4	3	4	1	8	4			
12	4	10	7	0	1	8	4	1	5	4	8	4	Total		1651
13	3	11	8	1	6	6	2	0	9	3	6	9			
14	2	12	0	7	1	2	5	2	1	1	7	2			
15	1	13	9	0	5	4	3	0	7	8	9	3			
16	4	14	9	2	2	7	1	1	2	10	7	10			
17	3	15	1	3	8	4	9	8	6	8	8	6			
18	1	16	9	6	4	5	5	7	8	8	9	5			
19	1	17	8	0	5	0	5	7	2	4	8	5			
20	2	18	6	7	6	3	2	4	1	6	7	4			
21	3	19	3	4	5	4	8	7	6	6	5	6			

FIGURE 36-1 This is the data for the job assignment problem.

In cells A3:A82, enter trial assignments of workers to workgroups. Assigning each worker to Workgroup 1, for example, is an acceptable starting solution. Copying the HLOOKUP(A3,Qual,B3+1) formula from K3 to K3:K82 enables you to look up each employee's qualifications for her assigned job. Note that Qual refers to the C2:F82 range. Next, copying the HLOOKUP(A3,Satis,B3+1) formula from L3 to L3:L82 enables you to look up the employee's satisfaction with her assigned job. Satis is the range name for G2:J82.

To deal with the fact that each division needs between 18 and 22 employees, you need to count how many employees have been assigned to each workgroup. You can do this in cells N6:N9 by copying the COUNTIF(\$A\$3:\$A\$82,M6) formula from N6 to N7:N9. Next, in cells O6:O9, determine whether a workgroup has the incorrect number of employees by copying the IF(OR(N6<18,N6>22),1,0) formula from O6 to O7:O9.

Now you'll see how to work on computing the target cell. In K1:L1, you compute total competence and total job quality by copying the SUM(K3:K82) formula from K1 to K1:L1. To ensure that each workgroup will have between 18 and 22 workers, you can penalize the target cell. Choose a penalty of 1,000 for each workgroup that has fewer than 18 or more than 22 workers. There is no hard and fast rule to help you determine an appropriate penalty. In this situation, the average rating is 5. This yields a target cell of $2 \times 400 + 400 = 1,200$. Therefore, it seems likely that putting the wrong number of people in any division would not benefit the target cell by more than 1,000, so survival of the fittest will kill off any solution for which a workgroup has too many or too few workers. The appropriate penalty should not be too large (100,000) because it sometimes makes the Solver ignore the real problem. If the penalty is too small, the Solver will not achieve the goal you've set.

In cell O10, compute the total number of divisions that do not have the correct number of workers with the SUM(O6:O9) formula. Now you are finally ready to compute the target cell in cell O12 by adding twice the total competence to the total job satisfaction and subtracting a penalty of 1,000 for each group that does not have the correct number of employees. Your final target cell is computed with the $2 \times K1 + L1 - 1000 \times O10$ formula.

You now can create the Solver model for this problem. You need to use the Evolutionary Solver because the COUNTIF and IF functions are nonsmooth functions of the changing cells. The model is shown in Figure 36-2.

Maximize the weighted sum of workgroup and employee satisfaction less the penalty for an incorrect number of workers in a workgroup (cell O12) and then constrain each worker's assignment to be 1, 2, 3, or 4. The solution is shown in Figure 36-1. Each group has the right number of workers; average employee competence is 7.2, and average employee satisfaction is 6.3. Over all 80 workers, the average of their competence ratings is 4.4, and the overall average for satisfaction ratings is 5, so conditions have improved quite a lot over a random assignment.

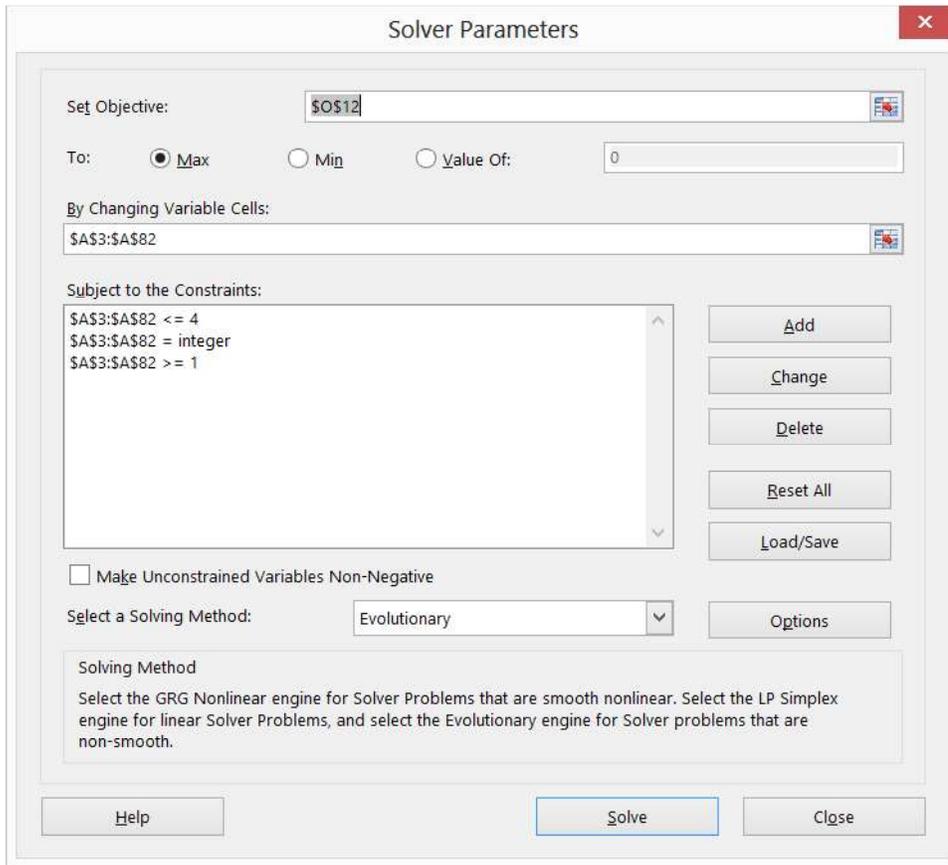


FIGURE 36-2 This is the Solver model for the worker assignment problem.

If you had tried the GRG Nonlinear engine (even with Multistart), the Solver would not have found the optimal solution because the model includes nonsmooth functions. Another tip about using the Evolutionary Solver is to use as few changing cells as possible, and you will usually be rewarded by Solver taking less time to find an optimal solution.

Using conditional formatting to highlight each employee's ratings

You can use the conditional formatting feature to highlight in yellow each employee's actual competence and satisfaction (based on his assignment). At cell C3, select the C3:J82 cell range. Select **New Rule** from **Conditional Formatting** on the **Home** tab, select **Use A Formula To Determine Which Cells To Format**, and fill in the dialog box as shown in Figure 36-3.

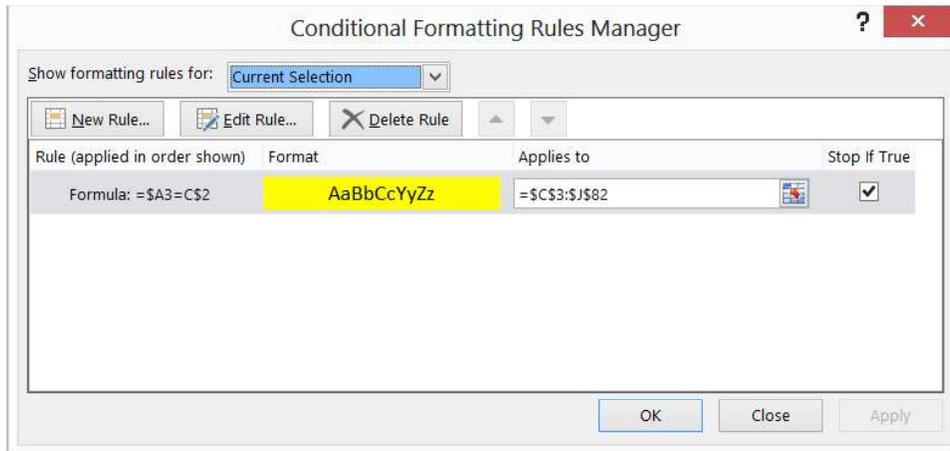


FIGURE 36-3 Use conditional formatting to highlight worker quality and satisfaction.

This formula enters a yellow format in cell C3 if and only if the first worker is assigned to Workgroup 1. Excel copies this formula across and down so that each worker’s quality and satisfaction ratings are highlighted only for the workgroup to which each employee is assigned.

Problems

1. Use the Evolutionary Solver to solve Problem 4 in Chapter 32, “Using Solver for capital budgeting.”
2. Solve the two-warehouse problems in Chapter 35, “Warehouse location and the GRG Multistart and Evolutionary Solver engines,” assuming that each warehouse can ship, at most, 120,000 units.
3. In the fictional state of Politicians Care about U.S., there are eight congressional districts. Each of 15 cities must be assigned to a congressional district, and each district must be assigned between 150,000 and 250,000 voters. The makeup of each district is given in the following table. Assign the cities to districts to maximize the number of districts won by the Democrats:

City	Rep	Dem
1	80	34
2	43	61
3	40	44
4	20	24
5	40	114
6	40	64
7	70	34
8	50	44

9	70	54
10	70	64
11	80	45
12	40	50
13	50	60
14	60	65
15	50	70

4. Solve the assignment of workers example, assuming that worker satisfaction is twice as important as the bosses' ratings.
5. Cook County General is attempting to develop the work schedule for its 20 nurses. Each nurse will work four consecutive days and is assigned to one of the following schedules:

Schedule	Days worked
1	Monday–Thursday
2	Tuesday–Friday
3	Wednesday–Saturday
4	Thursday–Sunday
5	Friday–Monday
6	Saturday–Tuesday
7	Sunday–Wednesday

Each nurse will be assigned for the week to either the ICU or a patient ward. Each nurse's satisfaction with her assignment is given in the Nursejackiedata.xlsx file. For example, if Nurse 5 is assigned to the ICU, Nurse 5 gives a perfect 10 rating to work schedule 3.

Each day, the ICU needs six nurses, and the patient wards need five nurses. Schedule the nurses to maximize their satisfaction and meet hospital needs.

You will need to know which schedules satisfy nurse demand for different days. For example, nurses starting on Monday, Friday, Saturday, and Sunday will work on Monday.

The traveling salesperson problem

Questions answered in this chapter:

- How can I use Excel to solve sequencing problems?
- How can I use Excel to solve a traveling salesperson problem (TSP)?

Answers to this chapter's questions

This section provides the answers to the questions that are listed at the beginning of the chapter.

How can I use Excel to solve sequencing problems?

Many business problems involve the choice of an optimal sequence. Here are two examples:

- In what order should a print shop work on 10 jobs to minimize the total time by which jobs fail to meet their due dates? Problems of this type are called job shop scheduling problems.
- A salesperson lives in Boston and wants to visit 10 other cities before returning home. In which order should he visit the cities to minimize the total distance he travels? This is an example of the classic TSP.

Here are two other examples of a TSP:

- A delivery driver needs to make 20 stops today. In which order should she deliver packages to minimize her time on the road?
- A robot must drill 10 holes to produce a single printed circuit board. Which order of drilling the holes minimizes the total time needed to produce a circuit board?

The Microsoft Excel 2013 Solver makes tackling sequencing problems very easy. Simply choose Evolutionary Solver, select your changing cells, and define All Different constraints. Configuring constraints with All Different ensures that if you have 10 changing cells, Excel will assign the values of 1, 2, . . . 10 to the changing cells, with each value occurring exactly once. In general, if you select a range of n changing cells to be different, Excel ensures that the changing cells assume the values of 1, 2, . . . , n , with each possible value occurring exactly once. See how to use Dif to solve a traveling salesperson problem easily.

How can I use Excel to solve a traveling salesperson problem (TSP)?

Solve the following problem.

Willie Lowman is a salesman who lives in Boston. He needs to visit each of the cities listed in Figure 37-1 and then return to Boston. In what order should Willie visit the cities to minimize the total distance he travels? Your work is in file Tsp.xlsx.

	E	F	G	H	I	J	K	L	M	N	O	P	Q
3			Boston	Chicago	Dallas	Denver	LA	Miami	NY	Phoenix	Pittsburgh	SF	Seattle
4		1 Boston	0	983	1815	1991	3036	1539	213	2664	792	2385	2612
5		2 Chicago	983	0	1205	1050	2112	1390	840	1729	457	2212	2052
6		3 Dallas	1815	1205	0	801	1425	1332	1604	1027	1237	1765	2404
7		4 Denver	1991	1050	801	0	1174	1332	1780	836	1411	1765	1373
8		5 LA	3036	2112	1425	1174	0	2757	2825	398	2456	403	1909
9		6 Miami	1539	1390	1332	1332	2757	0	1258	2359	1250	3097	3389
10		7 NY	213	840	1604	1780	2825	1258	0	2442	386	3036	2900
11		8 Phoenix	2664	1729	1027	836	398	2359	2442	0	2073	800	1482
12		9 Pittsburgh	792	457	1237	1411	2456	1250	386	2073	0	2653	2517
13		10 SF	2385	2212	1765	1765	403	3097	3036	800	2653	0	817
14		11 Seattle	2612	2052	2404	1373	1909	3389	2900	1482	2517	817	0
15		Order	Distance	City									
16		1	8	398	Phoenix								
17		2	3	1027	Dallas								
18		3	6	1332	Miami								
19		4	1	1539	Boston								
20		5	7	213	NY								
21		6	9	386	Pittsburgh								
22		7	2	457	Chicago								
23		8	4	1050	Denver								
24		9	11	1373	Seattle								
25		10	10	817	SF								
26		11	5	403	LA								
27		Total		8995									

FIGURE 37-1 Here is data for the TSP.

To model this problem in a spreadsheet, you should note that any ordering or permutation of the numbers 1 through 11 represents an order for visiting cities. For example, the order of 2-4-6-8-10-1-3-5-7-9-11 can be viewed as traveling from Boston (City 1) to Dallas (City 3), to LA (City 5), and finally to SF (City 10) before returning to Boston. Because the order is viewed from the location of City 1, there are $10! = 10 \times 9 \times 8 \times 7 \times 6 \dots \times 2 \times 1 = 3,628,800$ possible orderings for Willie to consider.

To begin, you need to determine the total distance traveled for any given order for visiting the cities. The *INDEX* function is perfect for this situation. Recall from Chapter 3, "The *INDEX* function," that the syntax of the *INDEX* function is *INDEX*(Range,row#,column#). Excel looks in the range of cells named Range and picks out the entry specified in row# and column#. In this case, you can use the *INDEX* function to find the total distance traveled in visiting all cities.

Begin by entering an order of the integers 1 through 11 in the F16:F26 range. Next, name the G4:Q14 range **distances** and enter the *INDEX*(distances,F26,F16) formula in cell G16. This formula determines the distance between the last city listed (in F26) and the first city listed (in F16). Enter the

INDEX(Distances,F16,F17) formula in cell G17 and copy it to the G18:G26 range. In G17, the formula computes the distance between the first and second city listed, the second and third city, and so on. Now you can compute the target cell (total distance traveled) in cell G27 with the SUM(G16:G26) formula.

At this point, you're ready to invoke the Evolutionary Solver. Minimize cell G27, click **Add Constraint**, and select the F16:F26 range. Select **Dif** for **All Different**. This ensures that the Solver always keeps the changing cells in the selected range, assuming the values 1, 2, up to 11. Each value will occur exactly once. The Solver Parameters dialog box is shown in Figure 37-2. Before running Solver, increase the mutation rate to .5.

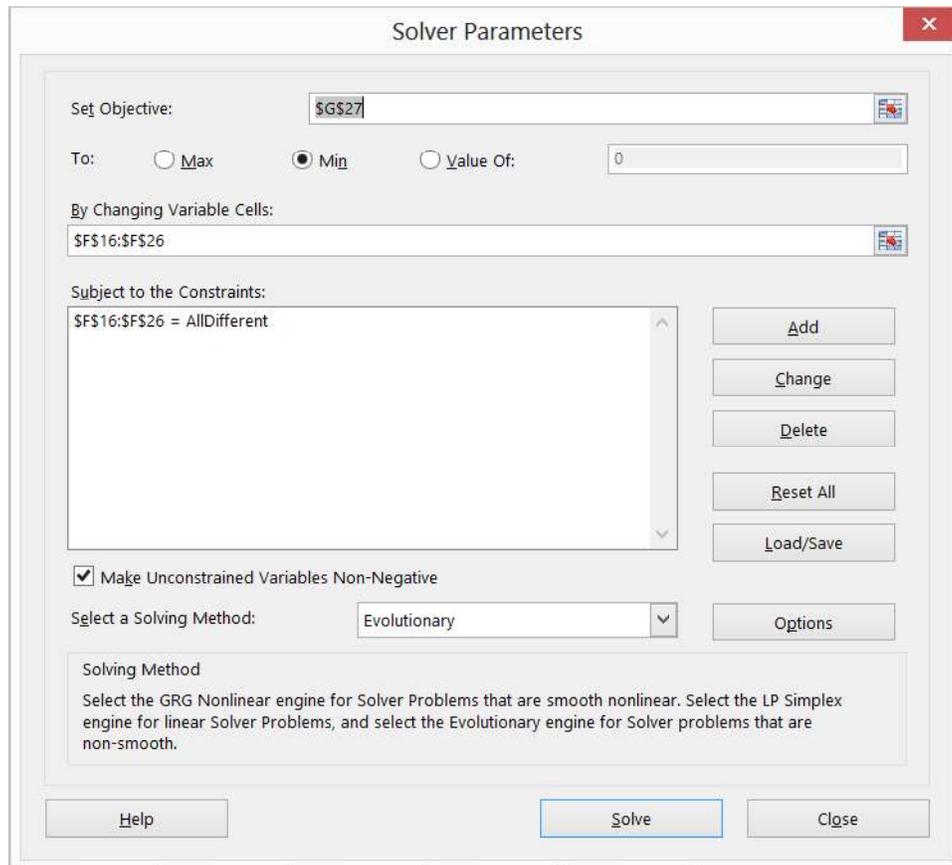


FIGURE 37-2 The Solver is set up for a traveling salesperson problem.

The minimum possible distance to travel is 8,995 miles. To see the order in which the cities are visited, begin in the row with a 1 (corresponding to Willie's home, Boston) and follow the cities in the listed sequence. The cities are visited in the following order: Boston–NY–Pittsburgh–Chicago–Denver–Seattle–San Francisco–Los Angeles–Phoenix–Dallas–Miami–Boston. Many other sequences for visiting the cities also yield the minimum total travel distance of 8,995 miles.

Problems

1. A small job shop needs to schedule six jobs. The due date and days needed to complete each job are given in the following table. In what order should the jobs be scheduled to minimize the total days the jobs are late?

Job	Processing time	Due date (measured from today)
1	9	32
2	7	29
3	8	22
4	18	21
5	9	37
6	6	28

2. The `Nbamiles.xlsx` file contains the distance between all NBA arenas. Suppose you live in New York and want to visit each arena once and return to New York. In what order should you visit the cities to minimize total distance traveled?
3. Suppose now that you live in Atlanta and are driving 29 general managers on this trip to NBA arenas. Each general manager wants to return to his home. Each time you visit an arena, you drop off a general manager at his home arena. In what order should you drop off the general managers to minimize the total distance traveled by the general managers?
4. In the Willy Lowman problem, suppose you must visit New York immediately after Denver. What is the solution to the problem?

Importing data from a text file or document

Question answered in this chapter:

- How can I import data from a text file into Excel so that I can analyze it?

Jeff Sagarin, the creator of the USA Today basketball and football ratings, and I have developed a system to rate NBA players that several NBA teams have used, including the Dallas Mavericks and New York Knicks. Every day during the season, Jeff's FORTRAN program produces a multitude of information, including ratings for each Dallas Maverick lineup during each game. Jeff's program produces this information in the form of a text file. In this chapter, you see how you can import a text file into Microsoft Excel to use it for data analysis.

Answer to this chapter's question

This section provides the answer to the question that is listed at the beginning of the chapter.

How can I import data from a text file into Excel so that I can analyze it?

You will likely often receive data in a Microsoft Word document or in a text (.txt) file that you need to import into Excel for analysis. To import a Word document into Excel, you should first save it as a text file. You can then use the Text Import Wizard to import the file. With the Text Import Wizard, you can break data in a text file into columns by using one of the following approaches:

- If you choose the fixed-width option, Excel guesses where the data should be broken into columns. You can easily modify the Excel assumptions.
- If you choose the delimited option, you pick delimiter characters (common choices are a comma, a space, or a plus sign), and Excel breaks the data into columns wherever it encounters the character(s) you choose.

As an example, the Lineupsch38.docx file (a sample of the data is shown in the following block) contains the length of time each lineup played for Dallas in several games during the 2002–2003 season. The file also contains the rating of the lineup. For example, the first two lines tell you that against Sacramento, the lineup of Bell, Finley, LaFrentz, Nash, and Nowitzki were on the court together for