

Chapter 16

Optimization in Several Variables with Constraints¹

In a previous chapter, you explored the idea of slope (rate of change, also known as the derivative) and applied it to locating maxima and minima of a function of one variable (the process was referred to as optimization). However, we know that most functions that model real world data are composed of several variables, so we need slightly different techniques for this. If you recall the one-variable case, we only needed to set that derivative to zero to find the local maxima and minima. When there are n independent variables, there are n different partial derivatives. We can find the location of the maxima and minima by find the points at which all n of these derivatives are zero at the same time (simultaneously). This involves a great deal of algebra, and is not always possible to do without resorting to numerical methods that only find approximate locations.

To make matters worse, we also find that rarely are we optimizing a function by itself. Consider, for example, revenue for selling a certain number of products. The more you sell, the more you earn, so there is no maximum revenue; we can make as many as we want and still earn more revenue. But in the real world, we have to account for the cost of the objects we are selling, which includes raw materials, labor and equipment to produce them, marketing, distribution, and other costs. These extra conditions, known as constraints, make finding an optimum solution much more difficult. In this chapter, we will focus on defining such constraints and phrasing them mathematically. We will then see how to set up a spreadsheet to solve the optimization problem under these constraints.

- *As a result of this chapter, students will learn*
 - ✓ What constraint functions typically look like
 - ✓ About sensitivity analysis

- *As a result of this chapter, students will be able to*
 - ✓ Formulate constraints for optimizing a function
 - ✓ Formulate a constrained optimization problem using the "Solver" package in Excel

¹©2011 Kris H. Green and W. Allen Emerson

16.1 Constraints on Optimization

So far, we have analyzed data by building models of the data and then interpreted those models. We have worked with models as equations that take one or more variables as input and have even worked with nonlinear functions. But analyzing the data and building the model is only part of the process. It is important that our model be useful for answering questions about the underlying situation and that we be able to use our model to make decisions. One of the most common uses of a model is in optimization, where we seek to make some quantity (such as profit or cost) either as large as possible (for profit) or as small as possible (for cost). In an earlier chapter, we did this with functions of a single variable, making use of a concept from calculus: the derivative. We found that when the derivative of a function is zero, the function is at a critical point, and that critical points are the only candidates for being optimum values of the function.

But this process ignores two things. The first is that most functions or models have several independent variables. Consider, for example, the commuter rail system examples we have used before. In that case, we built a model with a total of four variables. Our one-variable optimization process won't work here. The second thing we have ignored is that we are seldom free to choose just any values of the independent variables in order to achieve our optimum results. We are often constrained by resources. These resource constraints could involve time, money, personnel or just about anything that could limit our ability to reach certain values or combinations of values for the independent variables.

To correctly deal with the first problem, multiple-variable functions, we need to use partial derivatives (one for each independent variable) and solve several equations simultaneously. The idea is similar to the one variable case, but we now need all of the partial derivatives to be zero at the same exact point (set of values of the independent variables). We will not be looking into this here, because most of the common multivariable functions, linear and multiplicative, do not have critical points, and so we find no optimum solutions. Instead, we'll focus on the second aspect of optimization, applying constraints.

To begin, we must learn how to formulate the constraints. Typically, these will take the form of inequalities, rather than equations. After all, if the most you have to spend on production is \$100,000 and you can achieve a slightly higher profit by using on \$95,000, why not do it? So, rather than forcing our constraints to be equations, where quantities are equal to each other, we will use inequalities, where some quantity is either less than, greater than, less than or equal to, or greater than or equal to, some other quantity. We'll also see that most optimization problems involve multiple constraint conditions. For example, one constraint may involve time, one might involve cost of raw materials, one might involve equipment, and one might involve distribution.

16.1.1 Definitions and Formulas

Constraint Anything (such as time, money or budget, personnel or other resources) that limits your options regarding possible values of the variables in your model

Inequality An expression similar to an equation, except the quantities are related by either be less than ($<$), greater than ($>$), less than or equal to or greater than or equal to

(<= or >=) each other, rather than requiring them to be exactly equal. For example, the inequality

$$10*(\text{Labor Hours}) + 2*(\text{Items made}) \leq \$1000$$

Could be used to express the situation where we pay \$10/hour for labor and each item we make uses \$2 of raw goods, electricity and equipment time. If we must keep daily costs below \$1000, then the total cost spent of labor (= \$10/hour * number of hours used) plus the total cost in materials (= \$2/item * number of items made) must be less than or equal to the budget, in this case \$1000.

Objective function This is the quantity (or quantities) that you are trying to optimize. It is sometimes referred to as a target or target cell in spreadsheets.

Optimization variables These are the variables you can change (sometimes called the changing variables) in order to achieve your optimum solution.

Maximize In some optimization problems, you seek to make the objective function as large as possible. Such problems are maximization problems.

Minimize In some optimization problems, you seek to make the objective function as small as possible. Such problems are minimization problems.

Explicit constraint Explicit constraints describe those items that clearly given to you as goals during your optimization process. For example, limitations on resources (materials, labor) and limitations on demand are often stated explicitly.

Implicit constraint Implicit constraints refer to those quantities that you must recognize are also constraints on your optimization process. For example, in optimizing company profits by producing different quantities of different goods, the number of units of each good to produce might need to be an integer (unless it's measured in pounds). The quantity produced must also be non-negative. You must learn to recognize these types of constraints as well.

16.1.2 Worked Examples

Example 16.1. Understanding the optimization problem

We'll start with a typical sort of optimization problem. Suppose you want to maximize the profit from selling a variety of products, each of which requires different amounts of different materials that have different costs. Each of these products has a different demand function and sells for a different amount of money. As is common, you are required to remain under budget for both materials and labor.

Suppose our factory makes three products: kitchen tables, kitchen chairs and juice carts. We would like to maximize our profits from selling these three items, given the following information about the price, cost and production time for making each item (each value in the table is given per item produced):

Item	Assembly time	Finishing time	Materials cost	Selling price
Chairs	1 hr	1 hr	\$5	\$18
Tables	1 hr.	4 hr	\$15	\$54
Carts	2 hr	2 hr	\$10	\$36

In addition, at our factory we have the following weekly constraints:

	Maximum hours available	Cost per hour
Assembly	250 hr	\$4
Finishing	350 hr	\$7

We now need to take this raw information and frame it into a solvable optimization problem.

First, what is it we wish to optimize? That would be the profit from selling these items. Thus, the objective function is profit. In this case, the profit is given by the following:

Objective Function (Maximize)

$$\begin{aligned}
 \text{Profit} &= \text{Revenue from selling items} - \text{Cost for producing items} \\
 &= (\text{Revenue from chairs} + \text{Revenue from tables} + \text{Revenue from carts}) \\
 &\quad - (\text{Cost of making chairs} + \text{Cost of making tables} + \text{Cost of making carts}) \\
 &= \text{Profit from chairs} + \text{Profit from tables} + \text{Profit from carts}
 \end{aligned}$$

Note that one of our assumptions will be that we sell all of the chairs, tables and carts we make, and that all of these items sell for the listed price.

What are the input variables? In this problem, the only things we can change are the number of each item produced. Thus, we have three input variables: the number of chairs produced, the number of tables produced, and the number of juice carts produced. Once we know these numbers, we can compute the total cost for production, the total time for production and the profits we will earn.

What constraints do face? From the given information, it seems that we must keep the total amount of time to produce our chairs, tables, and carts within the amounts listed in the second table. Thus, we have two explicit constraints: the total time for assembly must be less than 250 hrs, and the total time for finishing must be less than 350 hrs.

Constraint #1 (Assembly time \leq 200 hr)

Assembly time = chair assembly time + table assembly time + cart assembly time \leq 250

Constraint #2 (Finishing time \leq 200 hr)

Finishing time = time to finish chairs + time to finish tables + time to finish carts \leq 350

At the same time, we have several implicit constraints. These are not stated directly, but are relevant to making sure we have a problem with a sensible solution. Thus, we need to be sure that we produce either zero or a positive number of each product. This gives us the following constraints:

Constraint #3 (Number of chairs produced ≥ 0)

Constraint #4 (Number of tables produced ≥ 0)

Constraint #5 (Number of carts produced ≥ 0)

We have three more implicit constraints: the number of each product that we make and sell must be an integer. After all, we cannot produce and sell one-third of a cart to make one-third of the revenue.

Constraint #6 (Number of chairs produced must be an integer)

Constraint #7 (Number of tables produced must be an integer)

Constraint #8 (Number of carts produced must be an integer)

Thus, even this relatively simple problem leads to eight constraints on the three input variables to maximize a single objective function.

Example 16.2. Mathematizing the optimization problem

Now that we have examined the problem from example 1 (page 467) and determined the constraints, we can convert everything into mathematical notation.

Assign names to the variables

The first thing you must do is assign names to the variables. In this case, we will choose the obvious names:

C = number of chairs to be produced and sold

T = number of tables to be produced and sold

J = number of juice carts to be produced and sold

Converting the objective function

To write the objective function as a mathematical expression we need to understand what goes into it. The final form of the profit is probably easiest to work with:

Profit = Profit from chairs + Profit from tables + Profit from juice carts

To determine the profit from selling chairs, we look at four things: the revenue from selling C chairs, the materials cost for making C chairs, the assembly cost for making C chairs and the finishing cost for making C chairs.

$$\begin{aligned} \text{Revenue from chairs} &= (\text{selling price of each chair}) * (\text{number of chairs}) = 18C \\ \text{Materials cost for chairs} &= (\text{materials cost for one chair}) * (\text{number of chairs}) = 5C \\ \text{Assembly cost for chairs} &= (\text{hours per chair}) * (\text{cost per hour}) * (\text{number of chairs}) \\ &= (1\text{hours}/\text{chair}) * (\$4\text{per hour}) * (C\text{chairs}) = 4C \\ \text{Finishing cost for chairs} &= (\text{hours per chair}) * (\text{cost per hour}) * (\text{number of chairs}) \\ &= (1\text{hours}/\text{chair}) * (\$7\text{per hour}) * (C\text{chairs}) = 7C \end{aligned}$$

Thus, the profit for making and selling C chairs is

$$\begin{aligned} \text{Profit for chairs} &= \text{Chair revenue} - \text{Material cost} - \text{Assembly cost} - \text{Finishing cost} \\ &= 18C - 5C - 4C - 7C \\ &= 2C \end{aligned}$$

In exactly the same fashion, we compute the following

Item	Revenue	Material cost	Assembly cost	Finishing cost	Profit
Chairs, C	$18C$	$5C$	$4C$	$7C$	$2C$
Tables, T	$54T$	$15T$	$4T$	$28T$	$7T$
Juice carts, J	$36J$	$10J$	$8J$	$14J$	$4J$

Thus, the objective function is given by the expression:

$$\text{Profit} = 2C + 7T + 4J$$

Notice that this expression is easy to interpret. For each chair we sell, we earn \$2 in profit, for each table we sell, we earn \$7 and for each juice cart, \$4. In the absence of constraints, the only way to maximize this profit is to make as many of each product as possible. Based on the coefficients, it would seem that making tables is the best way to go. However, that approach would completely ignore the time constraints on making the different products. At 1 hour per table, we could assemble a total of 200 tables. But it would take $200 \cdot 4 = 800$ hours of finishing time to complete those tables. We must keep finishing time below 350, so we can make at most $350/4 = 87.5$ tables. But this leaves over half of the assembly hours unused. Perhaps there is some way to make just the right number of each product to use the available resources and increase profits.

Converting the constraints

The constraints on time can easily be converted over. For each chair, we use 1 hour of assembly time. Thus, to make C chairs, we use C hours of assembly time. To make T tables, using 1 hour of assembly time per table, we use up T hours of assembly time. To make J juice carts at 2 hours per cart, we use $2J$ hours of assembly time.

Constraint #1 (Assembly time ≤ 250): $C + T + 2J \leq 250$

Similarly, we find the constraint on finishing time:

Constraint #2 (Finishing time ≤ 350): $C + 4T + 2J \leq 350$

The other constraints are simple to write down, but harder to encode as an expression to be used in doing algebra.

Constraint #3 (number of chairs ≥ 0): $C \geq 0$

Constraint #4 (number of tables ≥ 0): $T \geq 0$

Constraint #5 (number of juice carts ≥ 0): $J \geq 0$

Constraint #6 (number of chairs is integer): C integer

Constraint #7 (number of tables is integer): T integer

Constraint #8 (number of carts is integer): J integer

We have now converted all the expressions into mathematical notation. We could now apply the simplex method or some other technique to solve the problem. Notice that one

of the drawbacks to solving the problem this way is that we have all the numbers "hard coded" into the problem. From the final expressions, it is difficult to see how changing some of the initial information, like the number of hours to assemble a cart or the total number of finishing hours available, will change the expressions without starting over from scratch. In the next section, we will formulate our optimization problems for solution in a spreadsheet. This has the advantage of automatically updating the formulas and expressions based on the new information.

Example 16.3. Minimizing Shipping Costs

CompuTek produces laptops in two cities, Spokane, WS and Atlanta, GA. It purchases screens for these from a manufacturer, Clear Viewing, that has two production facilities, one located in Topeka, KS and the other located in Rochester, NY. CompuTek needs these items shipped to its two facilities. The plant in Topeka can produce at most 2000 units/week, while the plant in Rochester can produce 1800 units per week. Given the schedule below for how much it costs to ship a unit of product from each plant to the different cities where CompuTek needs them, how many should be sent from each plant to each city, if CompuTek needs 1000 units in Spokane and 1200 units in Atlanta?

Shipping Costs From	To	
	Spokane	Atlanta
Topeka	\$3	\$2
Rochester	\$4	\$5

This problem is obviously a minimization problem: we want to keep the shipping costs (our objective function) down to the lowest possible amount. We seem to have four input variables: the amounts shipped from each plant to each final city. And we seem to have two explicit constraints. (1) We cannot ship more from a plant than the plant can produce. (2) We need to ship the right number of units to each city so that CompuTek's order is filled.

Let's introduce some variables and express our problem in terms of these variables. We'll use the following variable names for each of the input variables:

TS = the number of units shipped from Topeka to Spokane
 TA = the number of units shipped from Topeka to Atlanta
 RS = the number of units shipped from Rochester to Spokane
 RA = the number of units shipped from Rochester to Atlanta

Then, we have the objective function, which is the total shipping cost (TSC):

$$\begin{aligned}
 TSC = & \text{\#of units shipped Topeka to Spokane} * \text{Unit price from Topeka to Spokane} \\
 & + \text{\#of units shipped Topeka to Atlanta} * \text{Unit price from Topeka to Atlanta} \\
 & + \text{\#of units shipped Rochester to Spokane} * \text{Unit price from Rochester to Spokane} \\
 & + \text{\#of units shipped Rochester to Atlanta} * \text{Unit price from Rochester to Atlanta}
 \end{aligned}$$

Thus, we see that the total shipping cost is

$$TSC = 3TS + 2TA + 4RS + 5RA$$

We want to make this number as small as possible subject to the following constraints:

We cannot ship more from Topeka than 2000 units:	$TS + TA \leq 2000$
We cannot ship more from Rochester than 1800 units:	$RS + RA \leq 1800$
We need exactly 1000 units in Spokane:	$TS + RS = 1000$
We need exactly 1200 units in Atlanta:	$TA + RA = 1200$
All quantities need to be positive:	$TA, TS, RA, RS \geq 0$
All quantities need to be integers:	TA, TS, RA, RS integer

16.1.3 Exploration 16A: Setting up Optimization Problems

Break down each of the problems described below so that you know (1) what the objective function is, (2) whether it is to be maximized or minimized, (3) the input variables, (4) the explicit constraints, and (5) the implicit constraints. For most of these situations, the problem is very vaguely stated, and you will need to also describe (6) what further information you will need in order to completely set up the problem.

Situation A.

How should a fast food chain allocate its advertising budget among different advertising formats?

Objective function	(max or min?)
Input variables	
Explicit constraints	
Implicit constraints	
Further information needed	

Situation B.

Where should a small town locate its only school?

Objective function	(max or min?)
Input variables	
Explicit constraints	
Implicit constraints	
Further information needed	

Situation C.

A large drug company has \$5 billion available to acquire small, start-up biotechnology companies. Which companies should it acquire?

Objective function	(max or min?)
Input variables	
Explicit constraints	
Implicit constraints	
Further information needed	

Situation D.

Your company has a list of 15 different strategic initiatives that it would like to undertake. Each will tie up some of your skilled labor pool for the next 4 years, and you do not have the resources to take on all 15 projects. Which projects should you select?

Objective function	(max or min?)
Input variables	
Explicit constraints	
Implicit constraints	
Further information needed	

16.1.4 How To Guide

Setting up constraints in Excel

In order to use Excel's routines for solving constrained optimization problems, you must first set up the spreadsheet so that (a) all the information is present and easily understood by a person reading it and (b) it contains all the formulas needed to connect the optimization variables and constraints to the objective function.

Let's return to examples 1 (page 467) and 2 (page 469), where we have translated an optimization problem into mathematical language as a preparation for solution. If you recall, in these problems we were trying to determine the optimum mix of three products (chairs, tables and juice carts) to manufacture in order to maximize our profits.

The screenshot below shows this problem set up (see C16 Furniture.xls). Notice that the first block of data, in cells B3:E8 is the given information about producing each of the three products (chairs, tables and juice carts). The only part of this that is not given is the second row (B4:E4) where we start with a guess regarding the number of each product to make in order to maximize our profit. In this case, we have simply guessed at 50 of each.

	A	B	C	D	E	F
1						
2						
3		Product Information	Chairs	Tables	Juice Carts	
4		# Produced	50	50	50	
5		Assembly Time (hours)	1	1	2	
6		Finishing Time (hours)	1	4	2	
7		Materials Cost	\$5	\$15	\$10	
8		Selling Price	\$18	\$54	\$36	
9						
10		Labor Information	Cost (\$/hr)	Labor Used		Available hours
11		Assembly	\$4	200 <=		250
12		Finishing	\$7	350 <=		350
13						
14		Profit Information				
15		Total Revenue	\$ 5,400			
16		Total Material Cost	\$ 1,500			
17		Assembly Cost	\$ 800			
18		Finishing Cost	\$ 2,450			
19		Net Profit	\$ 650			
20						
21						
22						

Figure 16.1: Setting up the constraints and functions for optimization with solver.

One of the reasons for setting up the spreadsheet as shown is that calculating the various costs and total times for assembly and finishing can be done efficiently with the SUMPRODUCT formula in Excel.

The next block of information contains the labor statistics (cells B10:F12). A few words about this are needed. First, the symbols in E11:E12 are not used by Excel in solving

this problem; these are given strictly to remind the user about the set up of this problem and some of the constraints present. Second, the "total labor used" in cells D10:D12 are calculations, based on the number of each item produced and the number of hours each step takes for that item. In cell D11, the following formula has been entered:

$$=\text{SUMPRODUCT}(\$C\$4:\$E\$4, C5:E5)$$

This calculation is an Excel shorthand formula for " $=C4*C5 + D4*D5 + E4*E5$ " which goes through the list of products and computes the assembly time needed for each product (number of items*hours per item) and then adds these times together. Using the sumproduct formula makes it easy to generalize this calculation to any number of products you might be considering, without having to type an unnecessarily long formula in. Using the absolute cell references for the product amounts (C4:E4) means that you can then directly copy the formula in D11 to cell D12 to compute the total hours of finishing labor used for this mix of products.

The third block of information (B14:C14) contains the calculations for the total revenue from selling the items, the total labors costs (broken down by assembly and finishing), the total materials costs and the net profit from sales. Each formula is a sumproduct, multiplying the row of information about quantities if each item produced ($\$C\$4:\$E\4) by the appropriate information (for revenue and materials costs) or it is a simple product (e.g., assembly cost in cell C17 is simply the product of total assembly hours used and cost per assembly hour = $D11*C11$). The net profit is then the revenue (C15) minus the total costs involved. Thus, profit is given by

$$=C15-(C16+C17+C18)$$

It is this quantity, profit, that we seek to maximize (objective function) by changing the amounts of each product we make (input variables) subject to various constraints.

16.2 Using Solver Table

There are many ways to go about solving an optimization problem with constraints, once you have the constraints written down mathematically as we did in the previous section. Some of these techniques involve a lot of algebra to solve equations, others use matrices and linear algebra to solve large systems of equations simultaneously, others are essentially based on guessing and checking, then changing your guess based on some calculus concepts.

Linear programming (LP) can be used to find solutions to optimization problems when the constraint functions and the objective function are all linear. Such problems typically occur in many operations research situations, such as studying the flow of materials through a network. Because the problem involves only linear functions, there are a variety of algebraic tools to solve them. One commonly used technique is called the simplex algorithm. But don't let the name fool you; this can be a very sophisticated technique, involving things like slack variables and shadow prices. Nonlinear and dynamic programming methods are modifications to LP that relax some of the conditions so that you can apply these techniques to nonlinear situations.

A more general method of optimization comes from multivariable calculus. Once you have defined all of your constraints, you have determined a feasible region for your solutions. You can then use calculus techniques to find any solutions inside this region, and can use the method of Lagrange multipliers to find possible solutions on the boundary of this region. The downside is that, since any kinds of constraints and objective functions can be used, there is no guarantee that you can solve the problem algebraically. You may need to resort to numerical tools to approximate the solution.

If your problem has only two independent variables, you can always graph the objective function and the constraints and used graphical techniques to approximate the solution. This method fails, though, when dealing with more than two independent variables because you cannot graph such functions easily.

The method we will discuss in this section is a combination of these techniques built into Excel called solver table. This lets us create an objective function, define our input variables, and develop a variety of constraints, both explicit and implicit. The solver table will then locate solutions based on the tools we tell it to use. It will also help us perform sensitivity analysis to find out how much our solutions will change if the inputs or constraints change. It is one thing to give your boss a report claiming that the company needs to do X under conditions Y, but it is much more useful to give him a report that says "under conditions Y we should do X, but if conditions Z change, we should modify our approach in the following ways..." This second version not only provides knowledge about the current situation, but also helps the decision maker respond to other situations and to anticipate the possible future directions he or she may need to explore.

16.2.1 Definitions and Formulas

Feasibility region This is the set of all possible values of the independent variables for which it even makes sense to talk about the optimization problem. For example, if you are optimizing profits from selling different quantities of different goods, you do not expect to find the optimum profit when you sell a negative number of an item. That

negative value is outside the feasibility region.

Feasible solution After solving the optimization problem, you need to check whether the solution is inside the feasibility region. If it is, you have a feasible solution.

Gradient function The gradient of a function of several variables represents a list, each item of which is one of the partial derivatives of the function. When solving an unconstrained optimization problem, you are trying to find the value(s) of each independent variable that make the partial derivatives in this list equal to zero simultaneously.

Lagrange multipliers This is a method of solving constrained optimization problems in n variables with m constraints. This method is similar to that of solving unconstrained optimization problems, but instead of trying to make n derivatives simultaneously zero we have to solve $n + m$ equations, a considerably harder task.

Linear programming This is a technique for solving optimization problems when you have a single linear objective function and your constraints are all linear functions of the different variables. It is guaranteed to find a solution, but may not find all the solutions or the best possible solution if there are multiple solutions possible. There are many software packages designed to solve linear programming problems. In fact, Excel's solver table may be manipulated to solve LP problems.

Dynamic programming This is a generalized version of linear programming that can be used to solve optimization problems with more complicated objective and constraint functions. By default, this is the way solver table in Excel works.

Sensitivity and sensitivity analysis When solving an optimization problem with constraints, we often need to find not only the solution to a particular problem but also what happens if we change some of the information slightly. For example, we can optimize profits for various anticipated budget and materials constraints, but need to know how this solution will change if the budget is a little less (or more) or the materials are a little harder to get. This additional analysis helps us determine how sensitive the optimum solution is to changes in the inputs and constraints.

Solver Table This is a program in Excel that enables you to solve various types of optimization problems

16.2.2 Worked Examples

Example 16.4. Solving an optimization problem with Solver Table

In this example, we will actually solve and interpret the results for the production example described in the last section. We'll be using file "C16 Furniture.xls". This relates to the production of tables, chairs and carts that we set up in examples 1 (page 467) and 2 (page 469). For information on how to set this up in Excel, see the How To Guide for section 16A.

Once the information is entered and the spreadsheet is prepared, you can complete the solution by invoking the solver routine. First, select the cell containing the profit (C19) and

then select "Solver" from the "Tools" menu. You should get a dialog box like the one below, already set up for you (because the file was saved with this solver scenario already set up). Notice that the first box contains the "Target Cell" which is our objective function, the profit (\$C\$19). The next row of information tells solver to maximize the target cell. Notice that you could select "Min" for minimization problems, or you could set this target cell to a particular value. This last is useful for scenarios of like the one below:

I need to make \$700 profit. How many of each product do I need to make?

Problems like this are not, strictly speaking, optimization problems. They are much more like the problems one might use Goal Seek to solve, except solver allows multiple input variables, while Goal Seek only allows one.

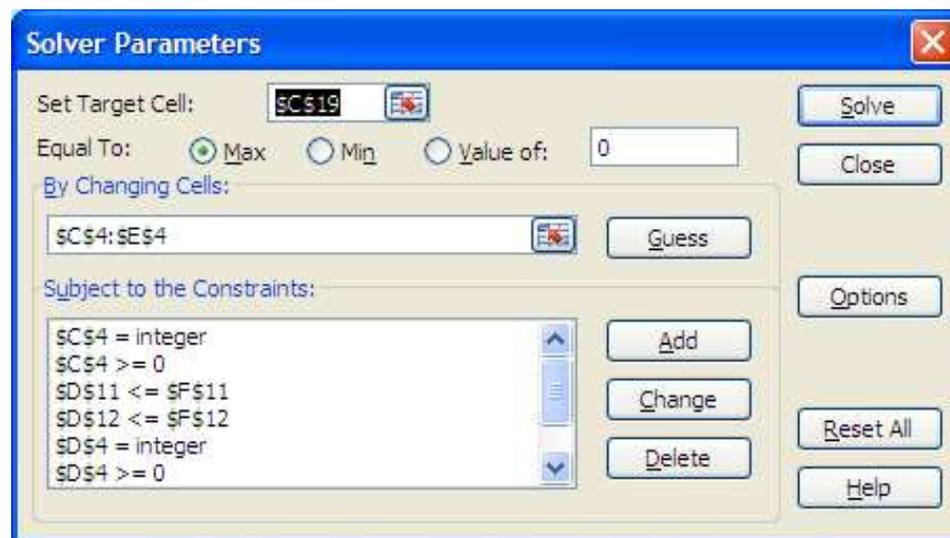


Figure 16.2: Solver dialog box filled in with constraints for solving the furniture problem.

The most difficult part about setting up the solver routine though is the final box, where the constraints are entered. Scrolling down the listed constraints should show you all eight of the constraints identified in example 1 (page 467). To add another constraint, you simply click the "Add" button and fill in the dialog box (see the How To Guide for details). Once all the constraints are entered, it can be as simple as hitting the "Solve" button to get a solution. In this case, solver returns the following screen.

The "Solver Results" box shows that the routine has found a solution. This solution is displayed in the spreadsheet, and it seems that if we make 62 chairs, 34 tables and 76 juice carts we can maximize our profit at \$666. Notice that this solution leaves us with only 2 unused hours of labor (in the assembly area). Clicking on the "Answer" under "Reports" and then clicking "OK" causes solver to create a new sheet (titled "Answer Report 1") summarizing the scenario (see the screen below for what this looks like). The report summarizes the initial information and compares this to the final solution. It also summarizes all of the constraints in the problem and whether the constraint was met exactly or whether some "slack" was allowed.

In this example, we know all of the constraints and the objective function are linear, so we could force Excel to use this information. Clicking on the "options" button in the solver routine, we can check off the "Assume Linear" box. Notice, though, that the solution determined in this way is very different from the solution determined otherwise, even though the maximum profit amount, \$666, is the same!

We could also have left out the constraints that force the quantity of each product to be positive. Instead, we could have selected the "Assume non-negative" option in the "Options" dialog box. The result is the same solution we had originally determined.

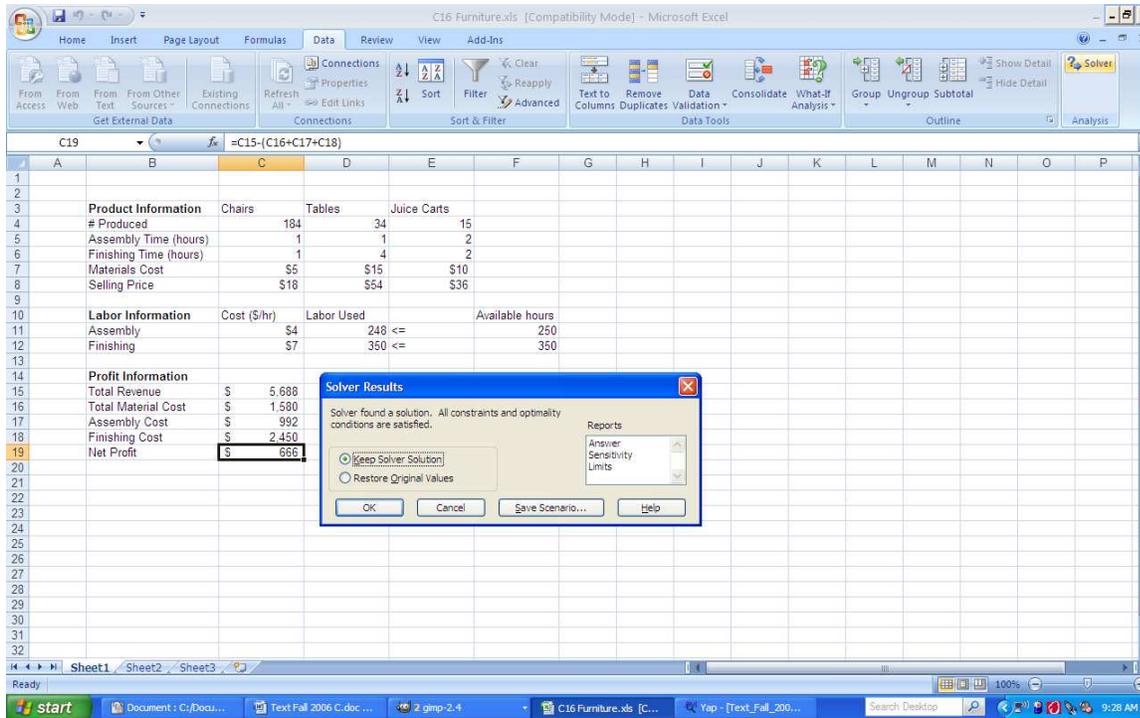


Figure 16.3: Results of solver routine on the furniture problem.

Example 16.5. Solving a minimization problem

Let's solve the problem related to shipping from example 3 (page 471). We need to organize our spreadsheet carefully so that we can easily change any of the given information in case the scenario changes. We also want to set up the spreadsheet so that the calculations are easy to copy. The screen shot below shows how we have chosen to organize the information.

Notice that the given information about the shipping costs from each plant to the final cities is listed at the top (cells B3:D5), and they are organized into a table much like the one that originally stated the information. This structure will be repeated throughout the setup of this problem. Below that, we have the temporary values for our input variables. These are also organized into a table in cells B8:D10. To the right of the variables are some calculations. Cell E9 contains the total number of units shipped from Topeka (=C9+D9), which is then copied to cell E10 to get the number of units shipped from Rochester. Cells F9:F10 contain symbols to remind us what the constraints are ($j=$) and G9:G10 contain the

Microsoft Excel 11.0 Answer Report
 Worksheet: [C16 Furniture.xls]Sheet1
 Report Created: 7/15/2006 7:42:33 AM

Target Cell (Max)
 NONE

Adjustable Cells

Cell	Name	Original Value	Final Value
\$C\$4	# Produced Chairs	50	62
\$D\$4	# Produced Tables	50	34
\$E\$4	# Produced Juice Carts	50	76

Constraints

Cell	Name	Cell Value	Formula	Status	Slack
\$D\$11	Assembly Labor Used	248	\$D\$11<=\$F\$11	Not Binding	2
\$D\$12	Finishing Labor Used	350	\$D\$12<=\$F\$12	Binding	0
\$C\$4	# Produced Chairs	62	\$C\$4>=0	Not Binding	62
\$D\$4	# Produced Tables	34	\$D\$4>=0	Binding	0
\$E\$4	# Produced Juice Carts	76	\$E\$4>=0	Not Binding	1

Figure 16.4: Answer report for solver solution of the furniture problem.

maximum number of units that can be shipped from the two plants. Below the variables, in cells B12:D13, we have information about the number of units needed in each city.

The next table, cells B15:D17 contains calculations for the total cost of shipping from each plant to each city. For example, in cell C16, the formula “=C9*C4” has been entered; this computes the cost for shipping a total of C9 units from Topeka to Spokane at a cost of C4 dollars per unit. This formula can then be directly copied to the other cells in the table to compute the corresponding amounts of shipping from each plant to each city.

The final piece of the setup is in cell C19: the total shipping cost. This is simply the formula

“=SUM(C16:D17)”. This will be the objective function that we minimize.

Selecting C19 and invoking the solver routine, we make sure to select “minimize” and then we enter our four input variable, found in cells \$C\$9:\$D\$10. There are three constraints:

Constraint	Notation in Solver
Total Shipped <= Maximum Output	\$E\$9:\$E\$10 <= \$G\$9:\$G\$10
Total Received = Total Needed	\$C\$11:\$D\$11 = \$C\$13:\$D\$13
All amounts are integers	\$C\$9:\$D\$10 integer

Then, we set the options for the solver routine so that “Assume non-negative” is checked,

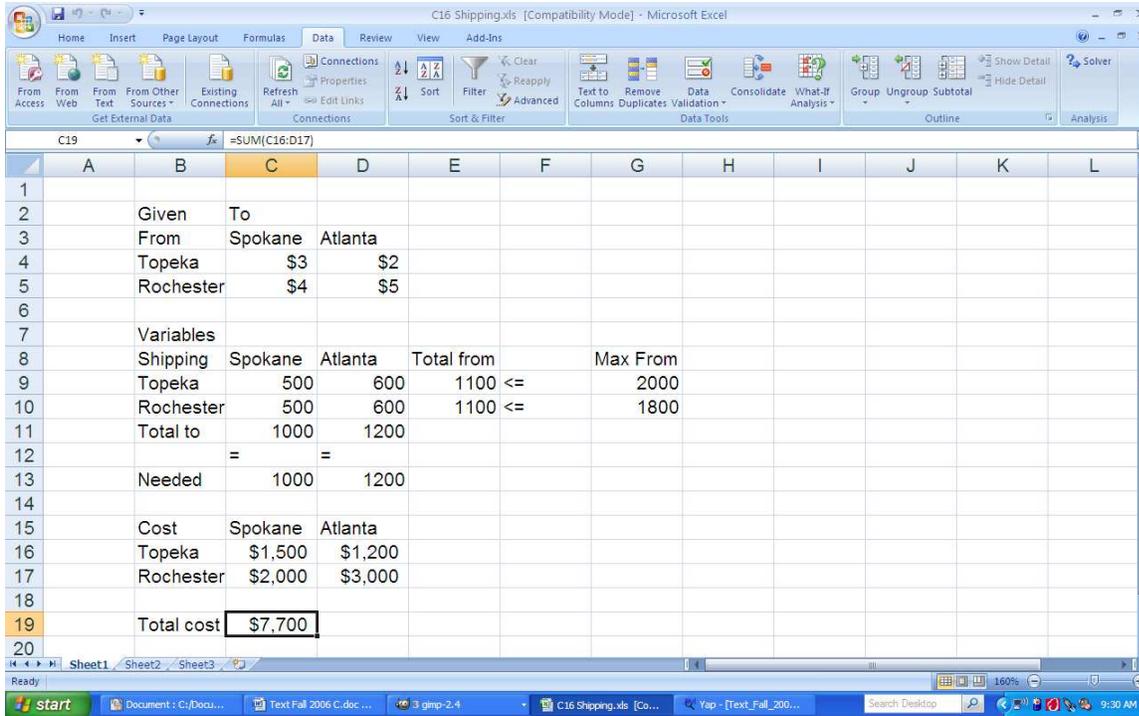


Figure 16.5: Setup for minimizing shipping costs in example 5.

and solve the problem. The answer report from the solver routine is shown below. It seems we should $TS = 800$, $TA = 1200$, $RS = 200$, and $RA = 0$.

Example 16.6. Special case - two input variables and one constraint

In some situations, optimization problems for multivariable functions are much easier to solve. One of these situations, one that occurs frequently, is when your objective function has two input variables and only one constraint function. In principle, any problem of this kind can be converting into a single variable optimization problem like those solved in chapter 14. Let’s see how this works.

Consider a typical model of production for an economy using a Cobb-Douglas (multiplicative) model like the one below. The two input variables are K , the units of capital invested in the economy, and L , the units of labor invested in the economy.

$$\text{Production} = 300K^{0.6}L^{0.4}$$

In most situations, we would like to maximize the productivity of the economy by setting the labor and capital investments appropriately. Thus, the productivity becomes our objective function. But we cannot choose just any values for K and L . Suppose that each unit of labor costs 85 and each unit of capital costs 130. If costs must be maintained below 100,000, then we want to maximize the productivity above, subject to the constraint that

$$\text{Total Cost} = 130K + 85L \leq 100,000$$

Microsoft Excel - C16 Shipping 2.xls

Microsoft Excel 11.0 Answer Report
 Worksheet: [C16 Shipping 2.xls]Sheet1
 Report Created: 7/15/2006 2:12:09 PM

Target Cell (Min)
 NONE

Adjustable Cells

Cell	Name	Original Value	Final Value
\$C\$9	Topeka Spokane	500	800
\$D\$9	Topeka Atlanta	600	1200
\$C\$10	Rochester Spokane	500	200
\$D\$10	Rochester Atlanta	600	0

Constraints

Cell	Name	Cell Value	Formula	Status	Slack
\$C\$11	Total to Spokane	1000	\$C\$11=\$C\$13	Binding	0
\$D\$11	Total to Atlanta	1200	\$D\$11=\$D\$13	Binding	0
\$E\$9	Topeka Total from	2000	\$E\$9<=\$G\$9	Binding	0
\$E\$10	Rochester Total from	200	\$E\$10<=\$G\$10	Not Binding	1600

Figure 16.6: Answer report for solver solution of the shipping problem.

Notice that for the type of function we have we can increase the objective function by making the input variables as large as possible. This, and other obvious logic, lead us to conclude that the only place we will get a solution to this problem is if the total cost is exactly equal to its maximum value, 100,000. (If we went for less, then another unit of capital or labor would push the cost closer to the maximum possible cost and would increase the productivity of the economy.) This means that we will take the constraint to hold at equality. That gives us a linear equation, which we can solve for one of the two variables:

$$130K + 85L = 100,000 \rightarrow K = (100000 - 85L)/130$$

If we plug this value of K into the objective function, we get

$$P = 300 \left(\frac{1000000 - 85L}{130} \right)^{0.6} L^{0.4}$$

This is a function of one variable! We can graph it easily and estimate the value of L for which the productivity, P , is maximum. We can then plug this value of L into the equation for K to find the number of units of capital we need to reach maximum productivity. The graph is shown below, indicating that about 500 units of labor are needed. Using the

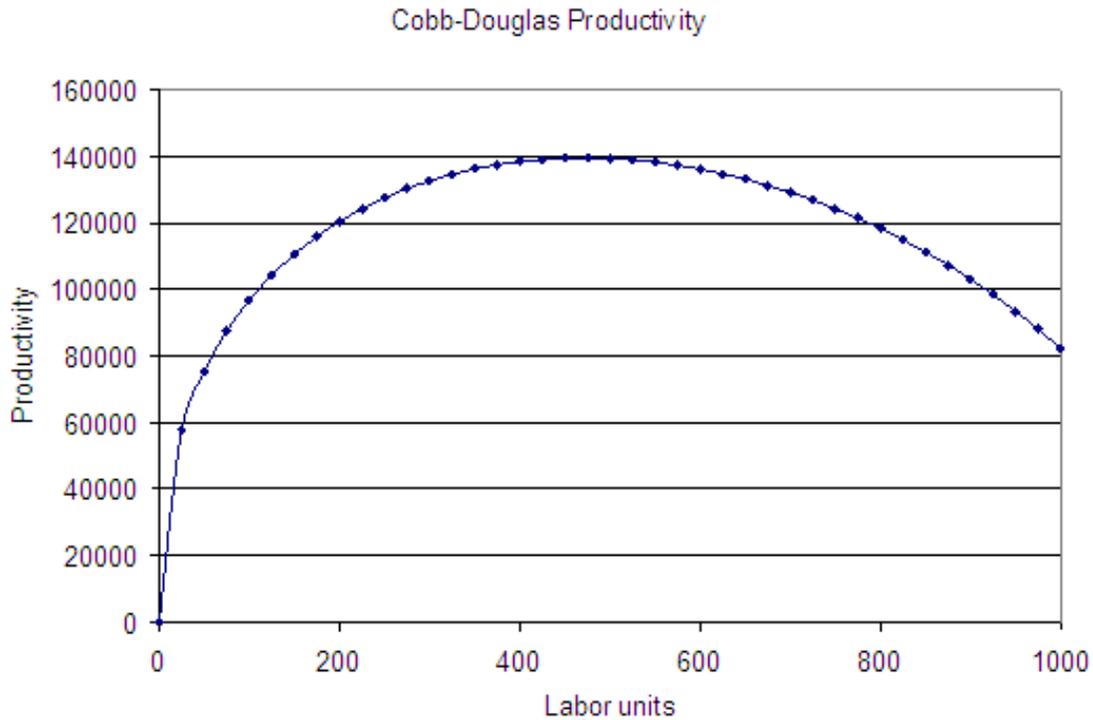


Figure 16.7: Cobb-Douglas productivity function with a constraint that reduces it to a one-variable model.

methods from chapter 14, we find that a better estimate is 470.59 units of labor. This results in 461.53 units of capital needed to achieve maximum productivity.

We could also solve this by setting the derivative of the production function to zero with Goal Seek. Computing the derivative of the objective function (after substituting in the constraint) will require the product rule and the chain rule.

16.2.3 Exploration 16B: Sensitivity Analysis

In this exploration, we will learn a little about "Sensitivity analysis". We'll use the example of producing chairs, tables, and juice carts (Examples 1 (page 467), 2 (page 469) and 4 (page 478)) and see how the optimal solution changes as we adjust the information about costs and labor hours. Use the file "C16 Furniture2.xls". Each of the following scenarios is a slight modification to the existing data; simply change the numbers from the original values and run the solver routine to see what happens. Summarize your results in the tables provided.

Scenario A.

It is probably not realistic to assume that each and every chair is assembled in 1 hour. Suppose your assembly crews are a little faster (0.9 assembly hours per chair) or a little slower (1.1 assembly hours per chair). How does the optimal solution change?

Value being adjusted	Assembly time for chairs			
Change being made	Value	Optimal # chairs	Optimal # tables	Optimal # carts
Lower	0.9			
Original value	1.0	62	34	76
Higher	1.1			

Scenario B.

The market probably changes a little, resulting in slight differences in actual selling prices for the products. What happens if the selling price of juice carts fluctuates between \$35 and \$37?

Value being adjusted	Selling price for juice carts			
Change being made	Value	Optimal # chairs	Optimal # tables	Optimal # carts
Lower	\$35			
Original value	\$36	62	34	76
Higher	\$37			

Scenario C.

Just as your selling prices fluctuate, you should expect your material costs to change slightly. Explore what happens if the material cost for the tables is between \$14 and \$16.

Value being adjusted	Material cost for tables			
Change being made	Value	Optimal # chairs	Optimal # tables	Optimal # carts
Lower	\$14			
Original value	\$15	62	34	76
Higher	\$16			

Scenario D.

Suppose the union demands that the workers in the assembly portion get higher wages. Explore what happens if the wages are raised from \$4 to \$5 or \$6 per hour of assembly work.

Value being adjusted	Labor cost for assembly			
Change being made	Value	Optimal # chairs	Optimal # tables	Optimal # carts
Original value	\$4	62	34	76
Increase \$1	\$5			
Increase \$2	\$6			

16.2.4 How To Guide

Installing the SOLVER add-in

The solver add-in is not always installed when you load Excel. To make it available, click on the "Tools" menu and select "Add-ins". Regardless of what other add-ins you have installed, you should see "Solver add-in" in the list of available add-ins. Check the box next to it, and then hit "OK" and from now on, solver should load in Excel and be available for solving optimization problems.

Adding constraints in Solver

In the solver dialog box, clicking on the "Add" button under "constraints" brings up a dialog box like the one in figure 16.8.

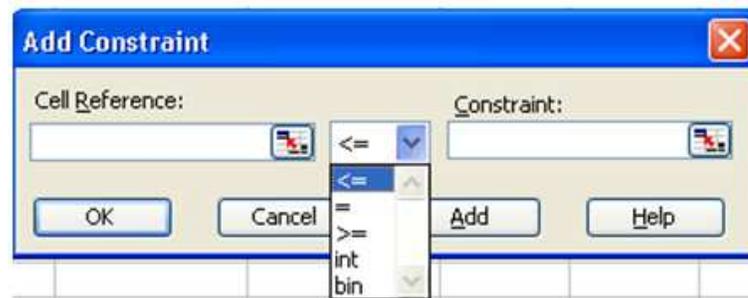


Figure 16.8: Adding a constraint in the solver routine.

Notice that there are three pieces of information needed for each constraint: the "Cell reference" for the constraint (which must somehow be connected to the input variables so that it changes as different input variable values are tested), the "constraint" (the value that you want to meet) and how the cell is constrained with respect to the value you provide. You have five possible options for this:

- <= Forces the value in the cell reference to be less than or equal to the constraint
- = Forces the value in the cell reference to be equal to the constraint
- >= Forces the value in the cell reference to be greater than or equal to the constraint
- Int Forces the cell value to be an integer
- Bin Forces the cell value to be a binary (either 0 or 1)

Clicking "Add" will enter your constraint and allow you to enter another constraint. Clicking "OK" enters the current constraint and returns you to the solver dialog box.

There is also a clever way to enter several constraints at once, if each is essentially the same type of constraint and the spreadsheet is properly organized. Consider the constraints shown in figure 16.9. (These are taken from a modified form of the solver scenario in C16 Furniture.xls.) Rather than entering one constraint for each of the product quantities to be integer, we have entered that "\$C\$4:\$E\$4 = integer" which forces each of the cells in the cell reference (C4, D4, and E4 in this case) to be an integer; this is much more efficient than

entering each one separately, especially when you could have hundreds of products in the scenario.

Likewise, instead of entering the labor constraints separately, one for assembly hours and one for finishing hours, we have entered a single constraint " $\$D\$11:\$D\$12 \leq \$F\$11:\$F\12 ". Solver goes through the list of cells on the left of the constraint and pairs them up with the corresponding cell on the right. Very efficient!

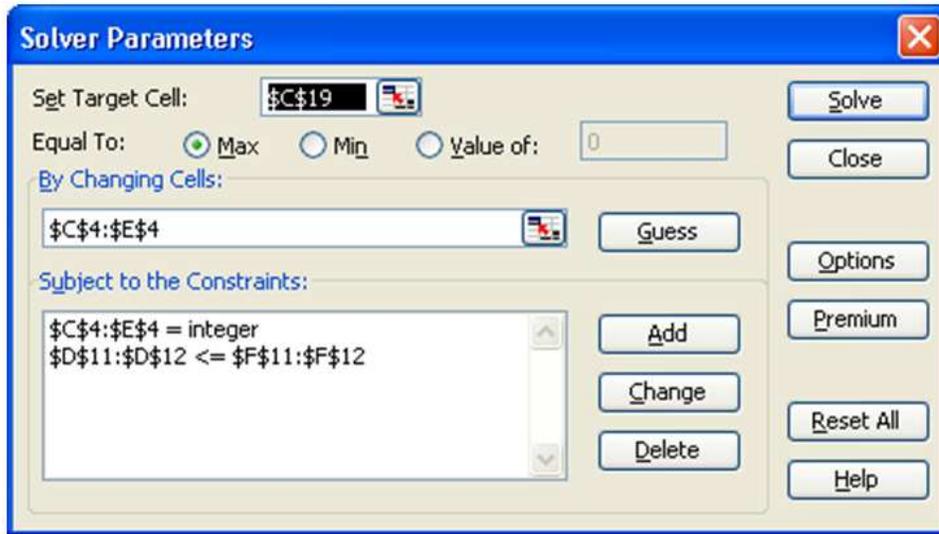


Figure 16.9: Setting up the solver parameters and constraints efficiently.

Changing constraints in solver

You can easily change an existing constraint in solver. Select the constraint in the solver dialog box, click "Change" and you can set all three of the options available when adding a new constraint.

Deleting constraints

Deleting a constraint is easy. Simply click on the constraint so that it is highlighted in the solver dialog box, and then click the "Delete" button.

Options in solver

Most of the time, you will not want to change many of the options in the solver options dialog box (see figure 16.10.) The first block of options all deals with the routine's limitations. The "Max time" is simply the maximum amount of time you want solver to search for a solution. The iterations relates to the number of times it loops through its procedures - there is almost no need to change these two options. Precision, tolerance and convergence all relate to the fact that, as numerical calculations, there is some rounding being done. Since all answers

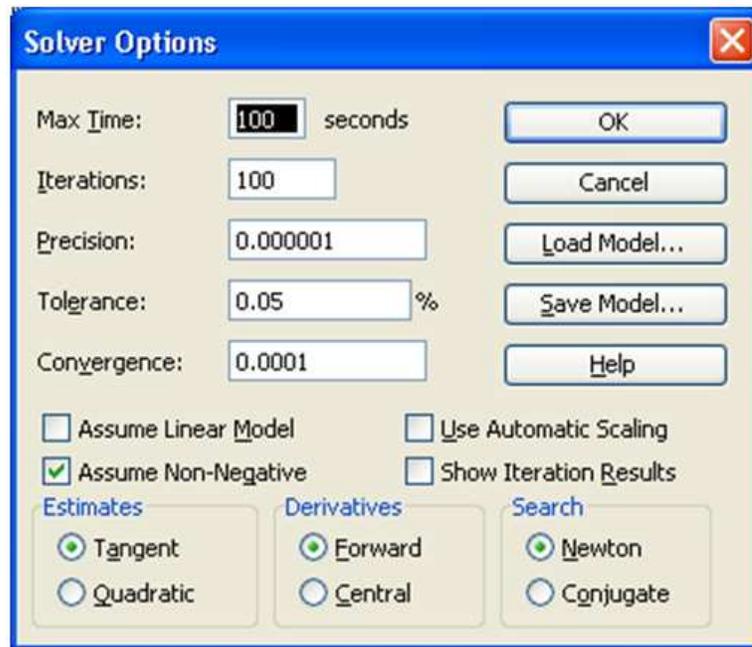


Figure 16.10: Clicking on the "Options" button in solver brings up this dialog box.

are approximate, these numbers let you specify how close solver must get to the specified values before it considers its work to be complete.

You will frequently encounter the need to set the "Assume linear model" and "Assume non-negative" options. Most of the rest of the options relate to the specific solution techniques used by solver. Details of these are fairly technical, and you rarely need to change them, but if you encounter a stubborn problem that defies solution, you may want to try solving it with a variety of options on these, to see if a solution can be found using alternative methods.

Errors in Solver

In example 4 (page 478), suppose we had union requirements that led to a minimum number of labor hours at our company. If we change the constraints from "Assembly hours ≤ 250 " and "Finishing hours ≤ 350 " to be minimum hours, swapping all the " \leq " for " \geq " what happens?

Solving this model (with the labor hours as minimums, rather than maximums) produces the following screen. Excel tells you that "the set cells do not converge" which means that you need to produce an infinite amount of each item in order to maximize profit. This is not a mistake or error in Excel; it is a problem with the scenario we set up. After all, if we have no maximum number of labor hours, then we can make as many of each item as we want; making more always results in more profit, so we should make as much as possible, which is infinite in this case. We can return this to a problem with a feasible solution by adding a constraint. For example, we may have a maximum amount of money available for materials.

Another possible error that could occur is "Solver could not find a feasible solution".

This doesn't mean the computer is broken; it simply means that the constraints you have specified over-determine the problem and cannot all be met simultaneously. You will have to carefully consider what constraints can be relaxed, and re-run the solver.

There are other possible errors that Solver may encounter. For more information, consult the help features in Excel under "Troubleshoot Solver".

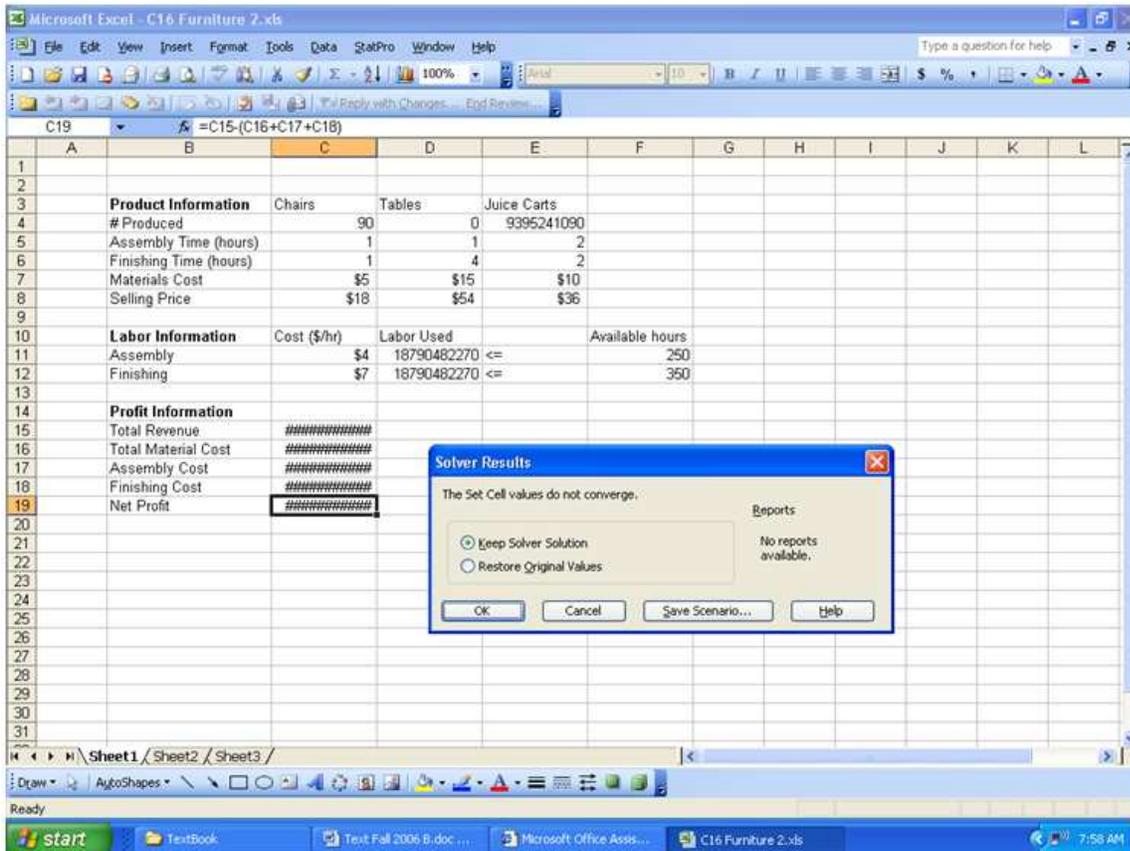


Figure 16.11: A typical solver error when there is no solution to the optimization problem.

More tricks solver will play on you

One final note about solver. Since its routine is numerical in nature, essentially a sophisticated way of guessing an answer, checking whether it works, and then calculating how to adjust the guess, the results are highly dependent on your initial guess for the input variables. Changing them could drastically change your optimal solution, especially if your situation is highly sensitive, or if there are many possible solutions to the problem. For example, in file C16 Furniture.xls, if you change the initial values from 50, 50, and 50 for the production of chairs, tables, and juice carts, the solver routine will give very different results, summarized below.

Initial Guess (C, T, J)	Optimal Solution (C, T, J)
(50, 50, 50)	(62, 34, 76)
(20, 20, 20)	(52, 34, 81)
(20, 50, 45)	(42, 34, 86)
(0, 0, 0)	(44, 34, 85)

Notice, though, that regardless of the initial values of the three variables, the number of tables in the optimal solution seems to be 34. Thus, we are simply trading off juice carts and chairs in the various solutions.

16.3 Homework

16.3.1 Mechanics and Techniques Problems

In the situations described below in problems 1 and 2, identify the objective function, whether it is a maximization or minimization problem. Also identify the free variables, the explicit constraints and the implicit constraints.

16.1. A company manufactures two products, widgets and greebles, on two machines, I and II. It has been determined that the company will realize a profit of \$6 on each widget and a profit of \$8 on each greeble. To manufacture a widget requires 9 minutes on machine I and 7 minutes on machine II. To manufacture a greeble requires 13 minutes on machine I and 5 minutes on machine II. There are 5 hours of machine time available on machine I and 3 hours of machine time available on machine II in each work shift. How many units of each product should be produced?

16.2. Harbor Tours operates a fleet of ocean vessels for touring the coast around the Alaska. The fleet has two types of vessels: Cruisers have 60 deluxe cabins and 160 standard cabins, while Corvettes have 80 deluxe and 120 standard cabins. Under a chartered agreement with Glacier Travel Agency, Harbor Tours is to provide Glacier Travel with a minimum of 410 deluxe and 720 standard cabins for their 15-day cruise in July. It costs \$65,000 to operate a cruiser and \$82,000 to operate a corvette for that period of time. How many of each type vessel should be used?

16.3. A financier plans to invest up to \$5 million in three projects. She estimates that project A will yield a return of 12% on the investment, project B will yield a return of 16% on the investment and project C will yield a return of 23% on her investment. Because of the risks associated with the investments, she decides to put not more than 20% of the total investment into project C. She also decided that her investments in projects B and C should not exceed 60% of the total investment. Finally she decided that her investment in project A should be at least 60% of her investment in projects B and C.

In this problem, we are obviously trying to maximize the financier's return on her investments (the objective function) by altering the free variables (the total amount of money invested in each of the three projects). We have four explicit constraints: (1) Investment in B should be less than or equal to 20% of the total investment. (2) Investment in B and C should be less than or equal to 60% of the total investment. (3) Investment in A should be greater than or equal to 60% of the total investment in B and C. (4) The total of all the investments must not exceed \$5 million. We also have some implicit constraints: (1) All three investments must be positive or zero. (2) All three investments must be a monetary value with at most two decimal places, when measured in dollars.

1. Write the objective function to be optimized and the constraints (both explicit and implicit) algebraically in as simple a form as possible.
2. Set up a spreadsheet to help solve this problem (but do not solve it).

16.4. Ivory Keys manufactures upright and console pianos in two plants, Boise and Canton. The output of Boise is at most 300/month, whereas the output of Canton is at most 250/month. These pianos are shipped to three warehouses that serve as distribution centers for the company. To fill current and projected future orders, the warehouse in Seattle requires a minimum of 200 pianos/month, the warehouse in Dallas requires at least 150 pianos/month, and the warehouse in Pittsburgh requires at least 200 pianos/month. The shipping costs from Boise to Seattle, Dallas and Pittsburgh are \$60, \$60, and \$80, respectively. The shipping costs from Canton to Seattle, Dallas and Pittsburgh are \$80, \$70, and \$50, respectively.

Clearly, in this problem we want to minimize the shipping costs (objective function) by altering the free variables (the number of each type of piano shipped from Boise to Seattle, Dallas and Pittsburgh and the number of each type of piano shipped from Canton to Seattle, Dallas and Pittsburgh; that's six variables!). We have five explicit constraints: (1) The number of pianos shipped from Boise cannot exceed 300. (2) The number of pianos shipped from Canton cannot exceed 250/month. (3) Seattle needs at least 200 pianos/month. (4) Dallas needs at least 150 pianos/month. (5) Pittsburgh needs at least 200 pianos/month. Finally, the only implicit constraints are that each shipping value must be positive and an integer.

1. Write the objective function to be optimized and the constraints (both explicit and implicit) algebraically.
2. Set up a spreadsheet to help solve this problem (but do not solve it).

16.5. Return to the furniture production example explored in this chapter. Modify the spreadsheet to include the additional constraint that you must produce at least four times as many chairs as tables. Solve this problem with the additional constraint.

16.3.2 Application and Reasoning Problems

16.6. Consider the situation described in #1 above. Under the stated conditions, the maximum profit of \$192 per shift comes from making 16 Widgets and 12 Greebles. The following questions relate to the way the optimal solution changes under different conditions in the environment. Open file "C16 Problem6.xls". The problem has been set up for you, and the constraints and settings for Solver are already configured.

1. Suppose one of your work crews finds a way to cut the time for making the Greebles on machine I. Instead of needing 13 minutes, they need only 12 minutes on machine I. See the screen shot below for what this might look like in Excel. Notice that you now have enough time left to make either one more Widget or one more Greeble. Which would be better? Why?

	A	B	C	D	E	F
1						
2		Product	Machine I	Machine II	Profit	Number
3		Widgets	9	7	\$6.00	16
4		Greebles	12	5	\$8.00	12
5						
6		Minutes				
7		Machine I	288	\leq	300	
8		Machine II	172	\leq	180	
9						
10		Profit	\$192.00			

16.7. Open file "C16 Problem6.xls". The problem has been set up for you, and the constraints and settings for Solver are already configured. Use Solver to help you explore the following change to the situation: Suppose the materials costs for Widgets are expected to increase. This will reduce your profit on each Widget. At what point should we drastically cut back on Widget production? At what point should we cut Widget production altogether?

16.8. Consider the situation described in #2 above. Under the given conditions, the optimal solution is shown in figure 16.13. Data file "C16 Problem7.xls" has this problem set up for you, including the solver table constraints.

	A	B	C	D	E	F
1						
2			Deluxe	Standard	Number	Cost
3		Cruiser	60	160	3	\$65,000
4		Corvette	80	120	3	\$82,000
5						
6		Deluxe	420	\geq	410	
7		Standard	840	\geq	720	
8						
9		Total Cost				
10		\$441,000.00				

1. Given the current optimal solution, how many additional cabins of each type could be rented for the cruise?
2. Suppose the cost for operating the Corvettes is expected to increase. At what point should Harbor Tours supply Glacier Travel Agency with only Cruisers?

16.9. The file "C16 Problem8.xls" shows the set up for the optimization problem below. C-Vite Company has decided to introduce three fruit juices made from blending two or more

concentrates. These juices will be packaged in 2-qt (64 fluid oz) cartons. To make one carton of pineapple-orange juice requires 8 oz each of pineapple and orange juice concentrates. To make one carton of orange-guava juice requires 4 oz of guava concentrate and 12 oz of orange concentrate. Finally, to make one carton of pineapple-orange-guava juice requires 4 oz of pineapple juice concentrate, 4 oz of orange juice concentrate and 8 oz of guava juice concentrate. The company has decided to allot 22,000 oz of pineapple juice concentrate, 28,000 oz of orange juice concentrate and 12,000 oz of guava juice concentrate for the initial product run. The company also stipulated that the production of the pineapple-orange-guava juice should not exceed 900 cartons. Its profit on one carton of pineapple-orange juice is \$1.00; its profits on one carton of orange-guava juice is \$0.90; and its profit on one carton of pineapple-orange-guava is \$0.95.

1. Explain how the spreadsheet sets up this problem for solution. Point out any unnecessary information provided in the problem statement (such unnecessary items are often called "red herrings".)
2. Find the optimal solution using Solver.
3. C-Vite is worried about a bad winter which could damage much of the Florida orange crop. If C-Vite only has 20,000 oz of orange juice concentrate to allot for the blends, how does this change their optimal production plan?
4. After a taste test, C-Vite decides to change the blend in the Pineapple-Orange-Guava juice to 4 oz pineapple, 6 oz orange and 6 oz guava. What optimal product mix should they pursue?

16.10. Oregon Lumber has decided to enter the pre-fabricated housing market. For its initial venture, it plans to offer three models of homes: traditional, deluxe and luxury. Each house is prefabricated and partially assembled at the factory. The final completion of the home takes place on site. The table below shows the costs, in material and labor, and the profit from each type of home.

	Traditional	Deluxe	Luxury
Material	\$6,028	\$8,062	\$10,135
Factory Labor (hr)	245	222	200
On-Site Labor (hr)	178	211	300
Profit	\$3,400	\$4,000	\$5,000

During the first year, Oregon Lumber has \$9 million to spend on materials. They cannot exceed 230,000 hours of labor in the factory, and they cannot exceed 245,000 hours of labor on site. Assuming that the market can sell as many houses as Oregon Lumber makes, how many of each type should be made if they want to maximize their profit?

16.3.3 Memo Problem

To: Analysis Staff
 From: Advertising Director
 Date: May 30, 2008
 Re: Advertising for Oracular

Oracular Consulting wants to increase its exposure to the public in order to increase business. We are planning to split our advertising budget among three venues:

- Full-page ads in the Albatross Airlines magazine that every executive will see when flying
- Web-based ads in Business Breaks Online which most executives subscribe to
- Print ads in the Sunday business section of a large national chain of newspapers

We have a maximum of \$50,000 to spend on advertising per month. Ads in the airline magazine cost \$25,000 for a one page ad each month. The magazine is seen by about 600,000 people each month. The newspaper ads cost \$210 per square inch for an ad that runs in all four Sunday issues in one month. The smallest ad available is four square inches; the largest is a half-page ad (115 square inches). The newspaper has a circulation of 900,000. Market research has shown that the number of people who pay enough attention to the ad to count as having seen it is dependent on the number of square inches of the ad according to the formula (where "units" stands for the number of units, in square inches, of the ad):

$$\text{Exposure} = -68.05 \cdot \text{Units}^2 + 14504.6 \cdot \text{Units}$$

The web ads are randomly generated whenever someone visits the web site. The price we pay is based on the priority our ad is assigned. For every 0.1% chance of our ad being shown, we pay \$60. The website will not allow us to purchase more than 25% of the ad space. The web site gets an average of 800,000 hits per month. Obviously, we want Oracular to be seen by as many people as possible. Keep in mind, though, that many of the circulation numbers above are approximate, so give us some idea of what range of results we can expect if these numbers change. Also, you should be aware that Albatross Airlines has been having some difficulties lately, and the number of passengers is expected to drop; they will probably raise prices for ads to compensate a little. We need to know at what point we should stop advertising in their magazine. Finally, the monthly advertising budget fluctuates as much as \$10,000 per month; how will that effect our decisions about what ads to place?

Attachments: None