

Chapter 9

Solver, Scenarios, and Goal Seek Tools



Contents

9.1	Introduction	311
9.2	Solver–Constrained Optimization	313
9.3	Example–York River Archaeology Budgeting	314
9.3.1	Formulation	316
9.3.2	Formulation of YRA Problem	318
9.3.3	Preparing a Solver Worksheet	318
9.3.4	Using Solver	322
9.3.5	Solver Reports	323
9.3.6	Some Questions for YRA	328
9.4	Scenarios	334
9.4.1	Example 1—Mortgage Interest Calculations	334
9.4.2	Example 2—An Income Statement Analysis	337
9.5	Goal Seek	338
9.5.1	Example 1—Goal Seek Applied to the PMT Cell	339
9.5.2	Example 2—Goal Seek Applied to the CUMIPMT Cell	341
9.6	Summary	342
	Key Terms	343
	Problems and Exercises	344

9.1 Introduction

Chapters 1 through 8 have introduced us to some very powerful tools for formulating, solving, and analyzing complex problems. From the seemingly endless array of cell functions, to the sorting, filtering, and querying, to PivotTable and PivotChart reports, to the Data Analysis tools, there seems to be a tool for almost every analytical need. Yet, there are a number of tools we have not exercised; in particular, the optimization tool Solver. Solver permits us to utilize a problem structure known as constrained optimization. Decision makers can search for a best solution for some desired objective, under constrained conditions. (It will also allow unconstrained problem structures.) We will learn more about this *prescriptive* analysis procedure later in this chapter.

Also, we will be introduced to the *Scenarios* and *Goal Seek* tools. The **Scenario** tool is used to implement what we have referred to as *what-if* analysis. Simply put, with Scenarios we have an efficient tool for automating the variation of inputs for a problem formulation, and then recording the resulting output. Its function is to organize and structure. Without the formal structure of Scenarios, it is very easy to lose important information in your analysis.

Goal Seek is also a member of the What-if Analysis sub-group in the Data Tools group. In situations where we know the outcome we desire from a problem, **Goal Seek**, as the name implies, is an approach that *seeks* to find the input that leads to our *desired* outcome. For example, if we are calculating the constant and periodic payments for a mortgage, we may ask the question—what interest rate will lead to a monthly payment of \$1000 for a loan, with a term of 240 monthly periods and principal value of \$100,000?

Before we become acquainted with these new tools, let us take stock of where we have been thus far on our analytical journey. We began with a careful classification of data from categorical to ratio, and we discussed the implications of the data type on the forms of analysis that could be applied. Our classification focused on *quantitative* and *qualitative (non-quantitative)* techniques for presenting, summarizing, and analyzing data. Through Chap. 6 we assumed that the data under consideration were available from the collection of primary¹ data or a secondary source. Additionally, our analysis was **descriptive** in nature, generally attempting to describe some characteristic of a population by analyzing a sample from that population; for example, comparing the mean of a sample to the mean of the population, and then determining if the sample mean could have come from the population, at some level of statistical significance.

In Chaps. 7 and 8, Modeling and Simulation, we created models in which we generated data that we could then analyze. In these chapters, we began our modeling effort by using the descriptive procedures discussed above to define our model, and then we generated data from the model to use for prescriptive purposes. For example, in Chap. 8 we used our understanding of the operating behavior of Autohaus to create a Monte Carlo simulation. The data generated by the simulation allowed us to prescribe to Inez the possible design selections for the Autohaus operational processes.

The models that we create with Solver have an even stronger prescriptive character than those encountered thus far. In using Solver, our goal is to determine the values of decision variables that will *minimize* or *maximize* some objective, while adhering to the technological constraints of the problem. Thus, the solution will *prescribe* very specific action about decision variables. As you can see from this discussion, our models have become ever more prescriptive in nature, as we have

¹Primary data is collected by someone in the role of collecting data for a specific purpose, and comes from sources that are generally not available as a result of other studies. For example, a survey study performed by a company to determine their customers' satisfaction with their service is primary data, while data on similar industry-wide service that can be purchased from a consulting firm is secondary data.

progressed through our early chapters. Before we explore Solver, we begin with a discussion of constrained optimization.

9.2 Solver–Constrained Optimization

Constrained optimization is a subfield of numerical optimization, with the critical distinction arising from the term *constrained*. Without a doubt, the most commonly used type of constrained optimization is **Linear Programming (LP)**. There is an abundance of software packages to provide solutions to LPs, including Solver. Additionally, there are numerous important applications of LP ranging from the distribution of electrical power, to scheduling, to menu planning. Some problems that can be formulated as LPs occur with such frequency, that they have received their own special designations; for example, the blending problem, the knapsack problem, the nurse or shift scheduling problem, and the capital budgeting problem. These are just a few of the generic formulations that have been developed over time as standard problem structures.

Linear programming, as a branch of constrained optimization, dates back to the 1940s. Over the years, as new innovations in LP techniques, computer programming, and digital computers have developed, the applications of LP have grown dramatically. Problems of such immense computational difficulty that a timely solution was once unimaginable can now be solved in fractions of a second by personal computers. Couple this with the availability of Solver, or other advanced Solvers that are add-ins to Excel, and now an ordinary personal computer user has a tool that was once the domain of firms that could afford costly LP software and the experts to conduct LP analysis.

Yet, the adage that “a little knowledge is a dangerous thing” certainly applies to LP. Before we cavalierly rush into formulating LPs, we need to spend some time considering the structure of an LP. We could also invest time in understanding the underlying algorithms that provide a computational platform for solving LPs, but that is a very sophisticated topic, and it will not be our focus. We will leave that for a text on the mathematical techniques of LP. Our focus will on *formulating* LPs; that is, structuring the important elements of a problem and then applying Solver to obtain a solution. Additionally, we will apply a special form of sensitivity analysis which is unique to LP solutions.

As we have discussed earlier, LPs consist of three major components:

1. **Decision Variables**—the continuous, non-negative variables that are selected to minimize or maximize an objective function, which is subject to constraints. As continuous variables, they can take on fractional values. The objective function and constraints are comprised of *linear*, algebraic combinations of the decision variables. Thus, no powers, other than 1, of the decision variables (no values like X^2 , $X^{0.5}$) are allowed, and only constants can be used as multipliers.

2. **Objective Function**—the *linear* combination of a decision maker’s decision variables that are to be either minimized or maximized. For example, we may want to maximize revenue or profit, and we will want to minimize cost or a deviation from a desired target level. We will refer to the objective function as *Z*.
3. **Constraints**—the *linear* combination of decision variables that represent how the decision variables *can* and *cannot* be used. Constraints are often referred to as the **technology of the linear program**.

As you can see, the term *linear* is important throughout the definitions above. If the conditions of linearity are not met, then you obviously do not have a LP. Yet, there are techniques for solving *non-linear* programs, and later we will briefly discuss the topic. It must be noted that it is far more difficult to solve non-linear programs, and it may also be very difficult to do so with Solver. There are software packages available for the specific purpose of solving non-linear programs. Now, let us consider a problem that can be formulated as a LP.

9.3 Example—York River Archaeology Budgeting

The principles of a thriving, but cash strapped business, York River Archaeology (YRA), are holding their quarterly budget meeting to decide upon the most critical budget that they have yet to produce. YRA does archaeological studies of sites that are to be developed for various uses—family housing, retail sales, public buildings, etc. Many cities and states require these studies. They are used to determine if important historical sites will be disturbed or destroyed due to development, and are quite similar to environmental impact studies. Thus, their services are in great demand in regions where significant building activity takes place.

In attendance at the meeting are the three principles, Nick, Matt, and Garrett, all accomplished Ph.D. archaeologists and historians. Also, their trusted business advisor, Elizabeth, is present. Nick spoke to the group first— “The importance of this budget cannot be overstated. In the past 10 years, we have not paid much attention to budgets, and it has gotten us into serious dilemmas with our cash flow. Back then, we could finance operations from our cash flow only, but now we have grown to the point where we need a substantial line of credit to keep operations running smoothly and finance our growth. Elizabeth has been telling us for years that it is important to create realistic budgets, and to stick to them, but we haven’t listened. Now, if we want our bank to provide us a reasonable line of credit, we have no choice but to do precisely as she has advised, and to show the bank that we can work *within* a budget.”

YRA is in the business of managing archeological projects. Just prior to the beginning of each quarter, the principals and Elizabeth meet to select the projects they will accept for the forthcoming quarter. From the projects selected and the resources consumed in the process, they are able to produce a detailed quarterly budget. YRA has been quite successful, so much so, that they have more requests for

Project Type	No. Projects Available	Rev.	Total Possible Rev	Res-A hrs. used	Res-B hrs. used	Res-C hrs. used	Res-D hrs. used
1	25	45,000	1,125,000	6	12	0	5
2	30	63,000	1,890,000	9	16	4	8
3	47	27,500	1,292,500	4	10	4	5
4	53	19,500	1,033,500	4	5	0	7
5	16	71,000	1,136,000	7	10	8	4
6	19	56,000	1,064,000	10	5	7	0
7	36	48,500	1,746,000	6	7	10	3
Total Available				Total Resource Hrs. Available in Quarter			
	226		9,287,000	800	900	700	375

Fig. 9.1 Projects for YRA budget

projects than they are capable of accepting. Thus, their resource limitation is the only factor that restricts them from accepting all the available project contracts.

Elizabeth has convinced them that they can model the selection process with LP. The LP will select the projects that maximize revenue while maintaining feasible usage of resources; that is, the LP will *not* select a set of projects that consumes more than the limited resources available, and it will do so while maximizing the objective function (in terms of revenue).

In preparation for creating an LP, Elizabeth sets a project pricing policy that estimates the revenue and outside resource usage associated with each project. Although there are some variations in her estimates, the revenues and outside resource usage are very nearly constant over short spans of time (3–6 months), certainly well within the quarterly planning horizon. The outside resources needed are stated in terms of hours, and they represent hours of outside consulting that must be obtained to perform a project. These include services like fieldwork by contract archeologists, report writing, and other types of services. YRA has decided it is wiser to obtain outside services than it is to hire permanent staff; it permits them to be very agile and responsive to varying demand. We will assume four resource types—*Res-A*, *Res-B*, etc.

Figure 9.1 provides a convenient format for assembling the relevant data for YRA’s project selection. For example, you can see that there are 25 projects of Type 1 available in the quarter, and that each project produces revenue of \$45,000. The resources used in Type 1 projects are 6 hours of *Res-A*, 12 of *Res-B*, 0 of *Res-C*, and 5 of *Res-D*. Note that there are seven types of projects available—Project 1, Project 2, etc.

An examination of the data suggests that there is no possibility to accept all projects, because doing so would result in a total of 4400 hours.² Since there are only 2775³ hours available, some projects cannot be accepted. The LP will choose the best combination of projects that meets the constraints imposed by the resources, while maximizing the revenue returned. It is not obvious which projects should be accepted. To simply assume that the highest revenue projects should be selected until resources run out is not advisable. There may be lower revenue projects that consume fewer resources that are more attractive from a cost/benefit perspective. Fortunately for us, the LP will make this determination and guarantee an “optimal solution”, or indicate that no solution is possible (**infeasibility**).

One of the features of an LP solution is the continuous nature of the resulting decision variables. Thus, it is likely that our solution will suggest fractional values of contracts; for example, it is possible that a solution could suggest 12.345 units of Project Type 1. Obviously, accepting fractional contracts is a problem; but if the numbers of projects that can be accepted are relatively large, rounding these fractional numbers up or down, while making sure that we do not exceed important constraints, should lead to a near optimal solution. Later, we will see that by imposing a constraint on a variable to be either binary (0 or 1) or integer (1, 2, 3, etc.) we convert our LP into a non-linear program (NLP). As we suggested earlier, these are problems that are far more complex to solve, and will require careful use of the Excel Solver to guarantee that an optimal solution has been achieved.

9.3.1 Formulation

Now, let us state our YRA problem in mathematical detail. First, the *decision variables* from which we can construct our objective function and constraints are:

X_1, X_2, \dots, X_7 = the number of each of the seven types of projects selected for the quarter

For example, X_4 is the number of projects of type 4 selected. The decision variables must all be non-negative values; selecting a negative number of projects does not make any practical sense.

Next, consider the *objective function*, Z , for the YRA in terms of the decision variables:

²If we multiply the “No. Projects Available” row by the individual “Res-A” through “Res-D” rows and sum the results, we find that a total of 4400 hours would be required to satisfy *all* the available projects. For example, for Res-A a total of 1338 hours would be required ($6*25 + 9*30 + 47*4 + 53*4 + 16*7 + 19*10 + 36*6 = 1150$). Similarly, 1625, 825, and 800 are needed for Res-B through Res-D, respectively. Thus, $1150 + 1625 + 825 + 800 = 4400$.

³Total resource hours available in the quarter are 2775 ($800 + 900 + 700 + 375 = 2775$) for Res-A through Res-D, respectively.

$$Z = 45,000X_1 + 63,000X_2 + 27,500X_3 + 19,500X_4 + 71,000X_5 + 56,000X_6 + 48,500X_7$$

The objective function sums the revenue contribution of the projects that are selected. If X_1 is selected to be 10, then the contribution of X_1 projects to the objective function, Z , is \$450,000 ($10 * 45,000 = 450,000$). We must also provide the type of optimality we are seeking—maximization or minimization. In the case YRA we are maximizing revenue.

Finally, we consider the *constraints* that are relevant to YRA. There are several constraint types that must be met, and the first relates to project availability:

$$X_1 \leq 25; X_2 \leq 30; X_3 \leq 47; X_4 \leq 53; X_5 \leq 16; X_6 \leq 19; X_7 \leq 36$$

Note that these seven constraints restrict the number of project types that are selected to *not* exceed the maximum available. For example, $X_1 \leq 25$ insures that the number of type 1 projects selected cannot exceed 25, while permitting values less than or equal to 25. Although we also want to restrict variables to non-negative values, this can be easily and universally handled with an option available in Solver—*Assume Non-Negative*. This condition is particularly important in minimization problems, since values for decision variables that are negative can contribute to the minimization of an objective function. Thus, in such a case, unless we set the non-negative condition, the LP will attempt to make the values of decision variables more and more negative to achieve a lower and lower Z value.

But, we are not done with the constraints, yet. We still have a set of constraints to consider that relate to the consumption of resource hours. For example, there is a maximum of 800 Res-A hours available in the quarter. Similarly, there are 900, 700, and 375 available hours of Res-B, Res-C and Res-D, respectively. The consumption of Res-A occurs when the various projects are selected. Thus, if we multiply each of the decision variables by the number of hours consumed, the resulting linear constraint relationships are:

$$\text{Res-A constraint. . . } 6X_1 + 9X_2 + 4X_3 + 4X_4 + 7X_5 + 10X_6 + 6X_7 \leq 800$$

$$\text{Res-B constraint. . . } 12X_1 + 16X_2 + 10X_3 + 5X_4 + 10X_5 + 5X_6 + 7X_7 \leq 900$$

$$\text{Res-C constraint. . . } 0X_1 + 4X_2 + 4X_3 + 0X_4 + 8X_5 + 7X_6 + 10X_7 \leq 700$$

$$\text{Res-D constraint. . . } 5X_1 + 8X_2 + 5X_3 + 7X_4 + 4X_5 + 0X_6 + 3X_7 \leq 375.$$

Let's take a close look at the first inequality, *Res-A constraint*, above. The **coefficients** of the decision variables represent the *technology* of how the maximum available hours, 800, are consumed. Each unit of project type 1, X_1 , that is selected results in the consumption of 6 hours of resource A; each unit of project type 2, X_2 , consumes 9 hours, etc. Also, recall that since LP's permit continuous variables, it is possible that a fraction of a project will be selected. As we mentioned above, later we

will deal with the issue of continuous variables by imposing integer restrictions on the decision variables.

Our LP for YRA is now complete. We have defined our *decision variables*, constructed our *objective function*, and identified our *constraints*. Below you can see the **LP formulation**, as it is called, of the problem. This format is often used to provide the complete structure of the LP. The problem can be read as follows: Given the decision variables which we have defined, maximize the objective function Z , subject to the constraint set. Next, we will see how to use the Solver to select the optimal values of the decision variables. As usual, dialogue boxes will prompt the user to input data and designate particular cells and ranges on a spreadsheet for data input and calculations.

By creating the LP formulation with the data input requirements of Solver, we permit the LP algorithm to perform its analysis. A consistent and predictable format for the formulation is important to reduce possible errors and facilitate entry of data.

9.3.2 Formulation of YRA Problem

Maximize:

$$Z = 45,000X_1 + 63,000X_2 + 27,500X_3 + 19,500X_4 + 71,000X_5 + 56,000X_6 + 48,500X_7$$

Subject to:

$$X_1 \leq 25; X_2 \leq 30; X_3 \leq 47; X_4 \leq 53; X_5 \leq 16; X_6 \leq 19; X_7 \leq 36$$

$$6X_1 + 9X_2 + 4X_3 + 4X_4 + 7X_5 + 10X_6 + 6X_7 \leq 800$$

$$12X_1 + 16X_2 + 10X_3 + 5X_4 + 10X_5 + 5X_6 + 7X_7 \leq 900$$

$$0X_1 + 4X_2 + 4X_3 + 0X_4 + 8X_5 + 7X_6 + 10X_7 \leq 700$$

$$5X_1 + 8X_2 + 5X_3 + 7X_4 + 4X_5 + 0X_6 + 3X_7 \leq 375$$

All X_i , where $i = 1$ to 7 , are non-negative.

9.3.3 Preparing a Solver Worksheet

Let us begin the process of using Solver to solve our LP. First, we create a worksheet that enables Solver to perform its analysis. Figure 9.2 appears to be quite similar to Fig. 9.1 in terms of the data contained, but there are some important differences. The

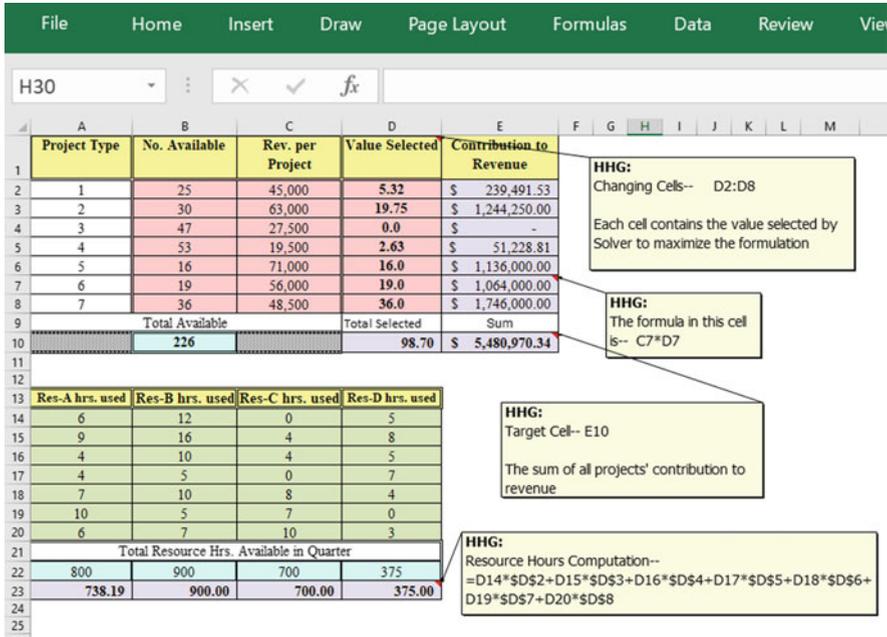


Fig. 9.2 Spreadsheet structure for solver

differences are essential to the calculations that will be performed, and the worksheet must be prepared in a manner that Solver understands.

There are two areas in the worksheet that must be set in order to execute Solver: a *target cell* and *changing cells*. The **target cell** is the cell in which we provide Solver with algebraic instructions on how to calculate the value of the objective function. This is the value of Z in our formulation, and as you recall, it is the sum of the number of projects of each type multiplied by the revenue that each project returns (coefficients). In Fig. 9.2 it is cell E10, and it is the summation of E2 through E8, the individual contributions by each project type. The values shown in the *target* and *changing* cells are calculated by Solver and will not contain these values until you have executed Solver. Once a solution is obtained, you can perform sensitivity analysis by changing the parameters currently in the problem. The values in the *changing cells* and the *target cell* will be recalculated after the next execution of Solver.

Although we have yet to perform the procedure, a solution is shown in Fig. 9.2, and we can see that Solver uses the **changing cells** as the cell locations for storing the values of the decision variables, in this case the optimal. Similarly, the optimal value of Z is stored in the *target cell*. This is quite convenient and makes reading the solution simple.

Note below that Solver has produced a solution that selected non-zero, positive values for project types 1, 2, 4, 5, 6, and 7, and a value of 0 for project type 3:

$$X_1 = 5.32; X_2 = 19.75; X_3 = 0; X_4 = 2.63; X_5 = 16.0; X_6 = 19.0; X_7 = 36.0$$

Value of the maximum Z	\$5,480,970.34
Total Projects selected	98.70.

The solution resulted in fractional values for project types 1, 2, and 4. How we deal with this situation depends on how strongly we feel that the constraints must be strictly maintained. For example, we could simply round every variable in the solution *up* to the next highest value: $X_1 = 6$; $X_2 = 20$; $X_4 = 3$. Rounding these fractional values *up* will require more resource hours. This will violate our constraints, but may or may not create a problem for YRA. To determine the effect on the use of resource hours, insert these values into the worksheet area containing the values of the decision variables, D2:D8.

It is easy to assume a *false-precision* when devising the constraints of a formulation. The constant values of the **right-hand side** of the constraints (which we call RHS) may simply be estimates. For example, the RHS of the Res-A hours constraint, 800, could be 812 or 789. So, there may be some latitude in our ability to deal with fractional results. Conversely, if the constraints are *hard and fast*, a solution could be implemented that rounds all, or some, of the fractional values *down* to insure strict adherence to the constraints. YRA should have an understanding of the precision of the data used in the formulation, and they will have to decide how to deal with this issue. Regardless, they have the formulation available to test the adjusted solutions that they develop. We will see later that one of the reports that is produced by Solver will provide information on the **slack** (unused resources) for each constraint. If there is no slack in a constraint, then the entire RHS (maximum amount of the resource) is consumed.

Let us now turn to the use of the Solver dialogue boxes to input our formulation. In Fig. 9.3 we can see the *target or objective cell* (E10), and the *changing cells* (D2:D8), where decision variable values are changed in search of an optimum, in the *Solver Parameters* dialogue box. Now, we introduce the constraints of the formulation. In row 23 (Fig. 9.2) we calculate the use of resource hours. The cell comment in D23 shows the calculation for the resource D hour usage. Generally, the formula sums the products of the number of projects selected times the hours consumed by each project type. These four cells (A23:D23) will be used by Solver to determine the use of the RHS (800, 900, 700, and 375,) and to ensure that the individual RHS's are not exceeded. We must also account for the maximum available projects for each type in B2:B8 (Fig. 9.2).

How do we enter this information into our formulation? We begin by selecting the Solver in the Analysis group in the Data ribbon. See Fig. 9.3. Solver should appear in the *Analyze* subgroup. But, if it does not, it is because you have not enabled the Solver add-in. To enable the add-in, select the File ribbon, then Options. One of the options is Excel Add-ins. It is at the bottom of the dialogue box. This is similar to how we enabled the Data Analysis tools.

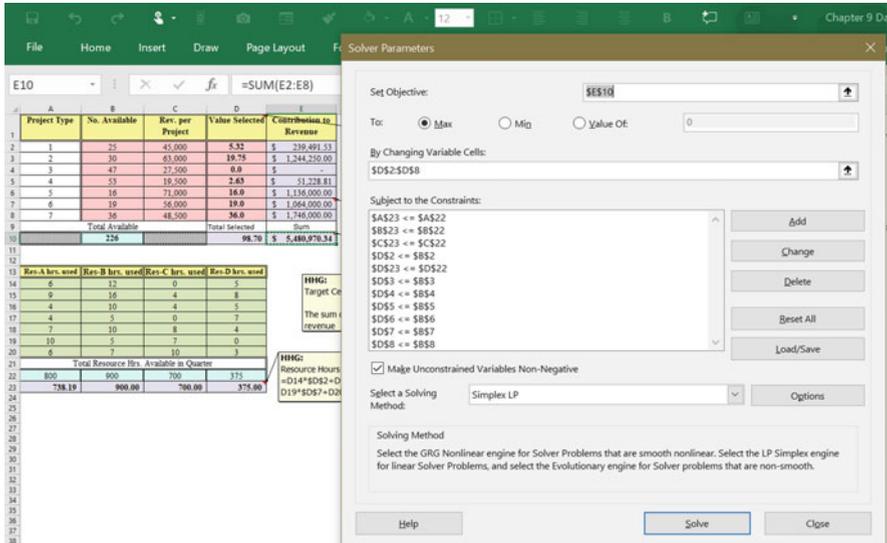


Fig. 9.3 Solver parameters dialogue box

The *Solver Parameters* dialogue box is relatively simple, but does require some forethought to use effectively. The entry of the *target cell* and the *changing cells* is straightforward; the entry of constraints is slightly more complex. You can see in the *Subject to the Constraints* area a number of constraints that have been entered manually. Imagine that we are beginning the process of entering the two types of constraints in our formulation. The first is the maximum number of project types that can be selected in a quarter. By selecting the *Add* button in the constraint area of the dialogue box, the *Add Constraint* dialogue entry box appears. Figure 9.4 shows the entry of $X_1 \leq 25$. The *Cell Reference* entry is the *changing cell* that contains the value for the project type 1 units selected, D2. The *Constraint* is the RHS of the constraint, B2 (25). Rather than enter 25 into the *Constraint*, I have use a cell location, B2. This permits a value change of cell B2 that results in the automatic update to the constraint. The new value is used in the next Solver execution.

The entry of the resource hour constraints is a more complex task. Above we noted that cells A23:D23 are the summation of the products of the decision variables and hours consumed by the use of each project type. Thus, each of these four sums represents the use of each resource hour type, due to the selected number of projects. We will use these cells, A23:D23, as the *Cell Reference* entries, and the corresponding RHS entries (*Constraint*) are found in A22:D22. In Fig. 9.3, you can see the first entry ($A23 \leq A22$) represents the constraint for resource A. The creation of the formula in cell A23 greatly simplifies data entry. Yet, it is possible to enter the constraints directly into the *Constraint* section of the *Add Constraint* dialogue box. We would do so by typing the same formula that appears in cell A23, and then entering the RHS, A22, in the *Cell Reference*. Note that you will have

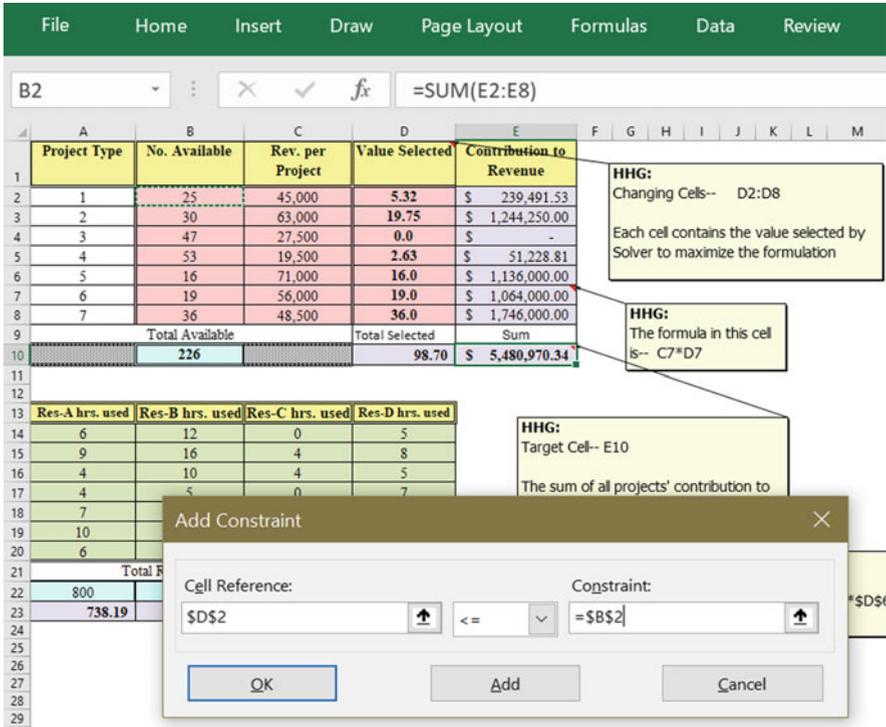


Fig. 9.4 Constraint entry for $X_1 \leq 25$

to change the sense of the inequality between these quantities to \geq , to preserve the proper inequality of the constraint.

9.3.4 Using Solver

Now, let us consider the application of Solver and examine the resulting solution. Assuming we have entered all the necessary formulation data, and we have correctly prepared the worksheet for Solver, we are now ready to execute the tool. Figure 9.5 shows the dialogue box called *Solver Options*. This dialogue box is available via a button in the *Solver Parameters* box. In previous versions of solver, the user had the option of making changes to these settings. I suggest that you not change these settings, unless you have a well-informed justification to do so. These settings control the technical aspects of the solution algorithms associated with Solver, and they are important to its precision and functioning.

In Fig. 9.3, we can see that the *Select a Solving Method* is set to *Simplex LP*, and the *Make Unconstrained Variables Non-Negative* box is checked. Both these conditions apply to our LP model and result in improved calculation performance. When

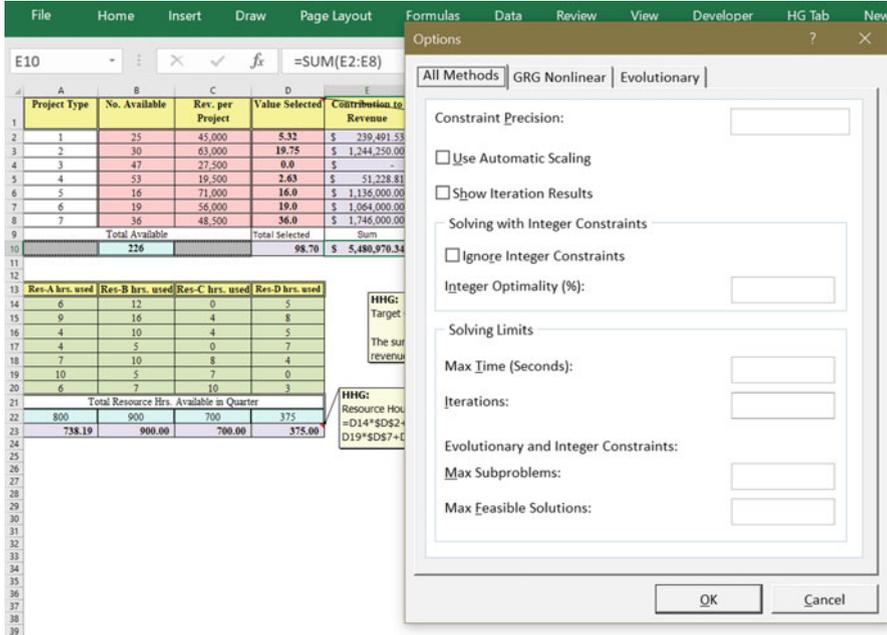


Fig. 9.5 Running the solver tool

Solver can assume LP characteristics, it simplifies and streamlines the solution procedure, and it also affects the solution output format. You can now return to the *Solver Parameters* box by selecting the *OK* button, and you are ready to launch Solver by selecting the *Solve* button.

Finally, make sure you have selected the *Max* button in *Solver Parameters* dialogue box, since we are maximizing our objective function. The moment of truth arrives when you receive verification that Solver has produced a feasible solution: a solution that does not violate the constraints and meets the optimality condition of the problem. This appears in the *Solver Results* dialogue box in Fig. 9.6. It is possible that your formulation may not have a solution, either because the constraints you have imposed simply do not permit a solution, or you have made a mistake in declaring your constraints. In both cases, no combination of feasible decision variable values can be found that satisfy your constraints.

9.3.5 Solver Reports

There are three reports that can be selected in the *Solver Results* box: *Answer*, *Sensitivity*, and *Limits*. We will focus on the *Answer* and *Sensitivity* reports. Figure 9.7 shows the *Answer Report*. It provides the basic information for the optimal solution in three output sections:

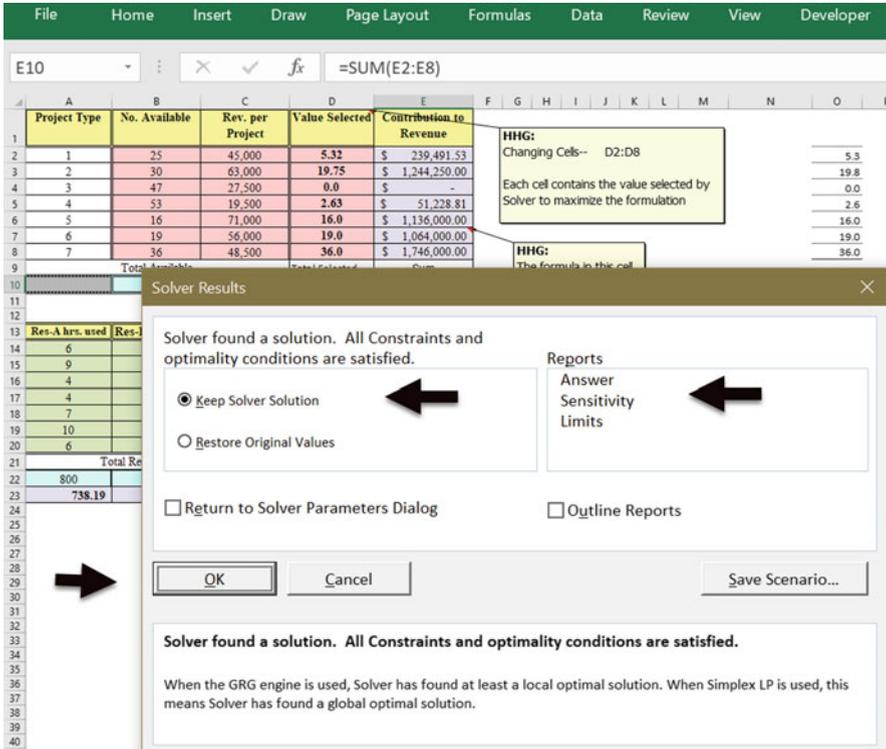


Fig. 9.6 Implementing solver solutions

1. *Objective Cell (Max)* section: the optimal value of the objective function (Z) determined by Solver—\$5,480,970.34
2. *Variable Cells* section: the values of the seven decision variables that lead to the optimal solution— $X_1 = 5.30$, $X_2 = 19.8$, etc.
3. *Constraints* section: information relating to the constraints and the amount of resources consumed by the optimal solution.

The information relating to constraints comes in the form of a constraint designation found in the *Status* section, as either **Not Binding** or **Binding**. You can see that there are 11 items in the *Constraint* section, corresponding to the formulation’s 11 constraints. The first constraint is the resource A usage. This constraint is *Not Binding* and has a slack of 61.81. What does this mean? If we sum the *Slack* and the *Cell Value*, the result is 800 hours, which is the maximum available hours for resource A. Thus, we can interpret this result as follows: the optimal solution consumed 738.19 hours of resource A, and therefore Resource A had 61.81 hours of unused capacity, or *Slack*. The next constraint is for resource B hours, and the *Status* of the constraint is *Binding*. For this constraint all the available capacity, 900 hours, is consumed, and there is no unused capacity; so, *Slack* is 0. The same

	A	B	C	D	E	F	G	H
14	Objective Cell (Max)							
15	Cell		Name	Original Value	Final Value			
16	\$E\$10		Sum	--	\$ 5,480,970.34			
17								
18								
19	Variable Cells							
20	Cell		Name	Original Value	Final Value	Integer		
21	\$D\$2		Value Selected	0.00	5.32	Contin		
22	\$D\$3		Value Selected	0.00	19.75	Contin		
23	\$D\$4		Value Selected	0.00	0.0	Contin		
24	\$D\$5		Value Selected	0.00	2.63	Contin		
25	\$D\$6		Value Selected	0.00	16.0	Contin		
26	\$D\$7		Value Selected	0.00	19.0	Contin		
27	\$D\$8		Value Selected	0.00	36.0	Contin		
28								
29								
30	Constraints							
31	Cell	Name	Cell Value	Formula	Status	Slack		
32	\$A\$23	Res-A hrs. used	738.19	\$A\$23<=\$A\$22	Not Binding	61.809322		
33	\$B\$23	Res-B hrs. used	900.00	\$B\$23<=\$B\$22	Binding	0		
34	\$C\$23	Res-C hrs. used	700.00	\$C\$23<=\$C\$22	Binding	0		
35	\$D\$23	Res-D hrs. used	375.00	\$D\$23<=\$D\$22	Binding	0		
36	\$D\$2	Value Selected	5.32	\$D\$2<=\$B\$2	Not Binding	19.6779661		
37	\$D\$3	Value Selected	19.75	\$D\$3<=\$B\$3	Not Binding	10.25		
38	\$D\$4	Value Selected	0.0	\$D\$4<=\$B\$4	Not Binding	47		
39	\$D\$5	Value Selected	2.63	\$D\$5<=\$B\$5	Not Binding	50.3728814		
40	\$D\$6	Value Selected	16.0	\$D\$6<=\$B\$6	Binding	0		
41	\$D\$7	Value Selected	19.0	\$D\$7<=\$B\$7	Binding	0		
42	\$D\$8	Value Selected	36.0	\$D\$8<=\$B\$8	Binding	0		
43								

Fig. 9.7 Solver answer report

analysis can be applied to the decision variable maximum limits, e.g. $X_1 \leq 25$, $X_2 \leq 30$, etc. Project type 1, which can be found in cell D36, has a value of 5.32. This is less than the maximum value of 25, thus the *Slack* is approximately 19.68. We can state that 19.68 units of project type 1 were not utilized in the optimal solution.

How might we use this information? Let us consider Resource A. If more hours of resource A could be found and added to the RHS, would the objective function benefit by the addition? Why? Currently you have unused hours of A; the constraint has *slack*. The addition of a resource that is currently underutilized *cannot* be of any value to the objective function. The solution algorithm sees no value to the objective function by adding an additional unit of A. We are far wiser to acquire additional hours of resource B, C, and/or D, since their constraints are *binding* and have no *slack*. To demonstrate the point, I will change the formulation to increase the number of resource B hours by 1 hour. Although this is a very minor change, it does lead to different decision variable values and to a higher objective function value as seen in Fig. 9.8.

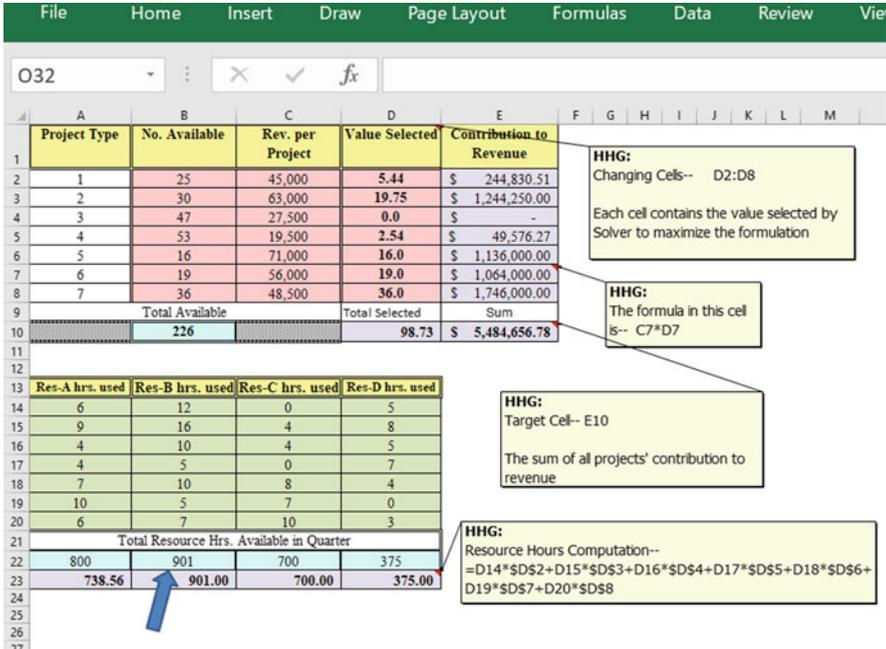


Fig. 9.8 Incremental change in resource B to 901

The new solution increases the number of project type 1, from 5.32 to 5.44, and reduces the number of project type 4 from 2.63 to 2.54. All other decision variables remain the same, including project type 2. Recall it was not at its maximum (30) in the previous solution (19.75), and the addition of a single unit of resource B has not caused it to change. The new value of the objective function is \$5,484,656.78, which is \$3686.44 greater than the previously optimal solution of \$5,480,970.34. The value of an additional hour of resource B is \$3686.44, and although the changes in decision variables are minor, the change has been beneficial to the objective function. We could perform this type of analysis with all our resource hours to determine the *marginal value* of an additional unit of resource (RHS) for binding constraints.

Although we do not perform the analysis here, if you increase resource A to 801 hours, I assure you that the solution will not change since it is non-binding and has slack of 61.44 hours.

These types of changes in the problem formulation are a form of sensitivity analysis. It would be convenient to have access to the results of these analyses without having to perform each individual change in the formulation. As you might guess, Excel does provide this type of analysis. It comes in the form of one of the three *Report* options in the *Solver Results* dialogue box (see Fig. 9.6)—the *Sensitivity Report*. Simply left-click the desired report to create a worksheet containing the report.

Variable Cells						
Cell	Name	Final Value	Reduced Cost	Objective Coefficient	Allowable Increase	Allowable Decrease
\$D\$2	Value Selected	5.440677966	0	45000	1800	6160.958904
\$D\$3	Value Selected	19.75	0	63000	6098.305085	2796.610169
\$D\$4	Value Selected	0	-12923.72881	27500	12923.72881	1E+30
\$D\$5	Value Selected	2.542372881	0	19500	10312.5	750
\$D\$6	Value Selected	16	27932.20339	71000	1E+30	27932.20339
\$D\$7	Value Selected	19	32673.72881	56000	1E+30	32673.72881
\$D\$8	Value Selected	36	15245.76271	48500	1E+30	15245.76271

Constraints						
Cell	Name	Final Value	Shadow Price	Constraint R.H. Side	Allowable Increase	Allowable Decrease
\$A\$23	Res-A hrs. used	738.5635593	0	800	1E+30	61.43644068
\$B\$23	Res-B hrs. used	901	3686.440678	901	30	45.85714286
\$C\$23	Res-C hrs. used	700	699.1525424	700	17.83333333	64.11111111
\$D\$23	Res-D hrs. used	375	152.5423729	375	64.2	12.5

Fig. 9.9 Sensitivity report for YRA

Figure 9.9 shows the Sensitivity Report for the YRA formulation. Let us focus on the *Constraints* section of the report, and particularly the column entitled *Shadow Price*. **Shadow prices** are associated with constraints, and result from the solution algorithm used to determine an optimal solution. Each of our four resource hour constraints can have a shadow price, and it can be interpreted as the benefit to the objective function if the RHS of a constraint is increased by one unit. Additionally, only binding constraints have shadow prices; thus, the resource A constraint will *not* have a shadow price since it is non-binding. Resource B, C, and D constraints do have a shadow price due to their binding condition, as can be seen in Fig. 9.9.

As we examine the shadow prices, we can see a very interesting similarity. The shadow price of an additional hour of resource B is 3686.44, which is precisely the change in objective function value we obtained earlier when we *manually* increased the RHS to 901. The shadow prices of resources C and D are less, 699.15 and 152.54, respectively. Thus, if I am allowed to increase any RHS of the four resource hour constraints, the most valuable hour is for resource B, followed by C, then D. Of course, it is important to consider the cost of each extra unit of resource. For example, if I can obtain an additional hour of resource D for \$155, it would not be wise to make the investment given that the return will only lead to \$152.54 in benefit for the objective function. The results would be a net loss of \$2.46.

In the last two columns of the *Constraints* area, additional valuable information is provided regarding the **Allowable Increase** and **Allowable Decrease** for the change in the RHS. As the titles suggest, these are the allowable changes in the RHS of each constraint for the shadow price to remain at the same value. So, we can state that if our increase of resource B had been for a 30-hour increase (the *Allowable Increase* is 31), our return would be 110,593.20 (30*3686.44). Beyond 31 units, there may still be a benefit to the objective function, but it will not have the same shadow price. In fact, the shadow price will be lower.

In the *Adjustable Cells* section of the *Sensitivity Report*, we also have information regarding the *estimated* opportunity cost for variables. These opportunity costs are found in the **Reduced Cost** column, and they apply to variables in our solution that are currently “0”; they relate to the *damage* done to the objective function if we force a variable to enter the solution. We have one such variable, X_3 . If we force the type 3 project variable to equal 1, that is, we require our solution to accept exactly one of these contracts, the reduction in the objective function value will be approximately \$12,923.90. Why do we expect a reduction in the objective function value Z ? The value of Z will be reduced because forcing the entry of a variable that was heretofore set to “0” by the optimization algorithm, will obviously damage, rather than help, our solution. The reduced costs for the X_5 , X_6 , and X_7 , although appearing to have a value, *cannot* be interpreted as the positive effect on the objective function if we increase the RHS by one unit, e.g. $X_5 = 17$.

One last, but important, use of the *Sensitivity Report* is related to the *Allowable increase* and *Allowable decrease* in the coefficients of the objective function variables (revenues for projects). These increases and decreases represent changes that can be made, one at a time, without changing the solution values for the decision variables. The objective function value will change, but not the selected decision variable values. As often is the case with complex problems, uncertainty will play a part in our analysis. The objective function coefficients of our decision variables are likely to be estimates that are dependent on many factors. This form of analysis will therefore be very useful in understanding the range of possible coefficient change, without affecting the current decision variable values.

9.3.6 Some Questions for YRA

Now, let us put our understanding of LP to the test. Elizabeth has received an urgent message from Nick. He wants her to meet him at a local coffee shop, the Daily Grind, to discuss some matters related to the quarterly budgeting process they have been working on. The following discussion occurs at the Daily Grind:

Elizabeth: Nick, your phone call sounded like something related to the budget needs urgent attention.

Nick: It does. I have been talking to a couple of our big clients and they are upset with the idea that we may not choose to accept their projects. They are threatening to take their business elsewhere.

Elizabeth: That is a problem. What project types are they talking about?

Nick: The problem is that they want us to take on 2 projects of type 3, and as you know, the solution that you determined suggested that we should not take on *any* type 3 projects. They are terrible projects from the stand- point of the low revenue they return and the high resource hours they consume.

Elizabeth: Yes, I know, they are a disaster for YRA, but if we want to keep these clients I guess we have to bend a bit and accept a couple of type 3 projects. It's the price of being customer focused.

Nick: There is also another issue we have to talk about. We need to reconsider the revenue we assumed for project type 2 and 4. I believe these are coefficients in the objective function of the formulation. Well, I think that we may have been inaccurate for both. In the case of project type 2, the revenue is more accurately between \$60,500 and \$68,000. For project type 4, the range is likely to be between \$19,000 and \$32,000.

Elizabeth: Nick, this could be a serious issue. I'll have to take a look at the sensitivity analysis for the solution. We may be fine with the current selection of projects, but with the requirement that we must select 2 projects of type 3, we will clearly have to make changes.

Nick: Also, while you are looking, can you also consider an increase in resource C of 12 hours. You know that many of the numbers we used in the formulation were point estimates of some range of values, and I have a feeling I was incorrect about that one.

Elizabeth: Nick, I want to use the current *Sensitivity Report* wherever I can, but I may have to reformulate and solve the problem again. Don't worry; these things happen all the time. Solver can answer many of the questions.

Nick: Thanks for stopping by; we really need to answer these questions quickly.

Elizabeth: You are right. I should have answers for you by tomorrow.

Elizabeth considered the tasks ahead of her. She decided to start with the issue of accepting 2 of the unprofitable type 3 projects. Depending on the results with this analysis, she then will consider the issues related to the revenues (coefficients of the objective function), and, finally, the change in the resource hours.

She begins by returning to Fig. 9.9 where she finds that the *Reduced Cost* of the *Project Type 3* is approximately—12,924 (see cell E11). As you recall, this number represents the penalty associated with forcing a unit of a decision variable into the solution that is currently set to 0. Therefore, an estimate of the reduction of the objective function by including 2 units of type 3 projects is approximately \$25,848. Figure 9.10 verifies that this is approximately the change to the objective function if we re-solve the LP with a change in the constraint for resource type 3.

The new solution, appearing in the *Final Value* column, is shown next to the old solution values, in the *Original Value* column. The inclusion of an *Original Value* and a *Final Value* is a convenient mechanism for comparison. The change from the previous objective function value, \$5,480,970.34, is to a value \$25,847.46 lower, \$5,455,122.88. The reduction is very close to \$25,848, as predicted by the reduced cost. The new decision variables selected look familiar, but with a few significant changes:

$$X_1 = 6.24; X_2 = 17.75; X_3 = 2.0; X_4 = 2.83; X_5 = 16.0; X_6 = 19.0; X_7 = 36.0$$

	A	B	C	D	E	F	G
12							
13							
14		Objective Cell (Max)					
15		Cell	Name	Original Value	Final Value		
16		\$E\$10	Sum	\$5,458,809.32	\$5,455,122.88		
17							
18							
19		Variable Cells					
20		Cell	Name	Original Value	Final Value	Integer	
21		\$D\$2	Value Selected	6.36	6.24	Contin	
22		\$D\$3	Value Selected	17.75	17.75	Contin	
23		\$D\$4	Value Selected	2.0	2.0	Contin	
24		\$D\$5	Value Selected	2.75	2.83	Contin	
25		\$D\$6	Value Selected	16.0	16.0	Contin	
26		\$D\$7	Value Selected	19.0	19.0	Contin	
27		\$D\$8	Value Selected	36.0	36.0	Contin	
28							
29							
30		Constraints					
31		Cell	Name	Cell Value	Formula	Status	Slack
32		\$D\$23	Res-D hrs. used	375.00	\$D\$23<=\$D\$22	Binding	0
33		\$B\$23	Res-B hrs. used	900.00	\$B\$23<=\$B\$22	Binding	0
34		\$C\$23	Res-A hrs. used	700.00	\$C\$23<=\$C\$22	Binding	0
35		\$A\$23	Res-C hrs. used	734.50	\$A\$23<=\$A\$22	Not Binding	65.5042373
36		\$D\$4	Value Selected	2.0	\$D\$4<=\$B\$4	Binding	0
37		\$D\$3	Value Selected	17.75	\$D\$3<=\$B\$3	Not Binding	12.25
38		\$D\$4	Value Selected	2.0	\$D\$4=2	Binding	0
39		\$D\$5	Value Selected	2.83	\$D\$5<=\$B\$5	Not Binding	50.1694915
40		\$D\$8	Value Selected	36.0	\$D\$8<=\$B\$8	Binding	0
41		\$D\$7	Value Selected	19.0	\$D\$7<=\$B\$7	Binding	0
42		\$D\$6	Value Selected	16.0	\$D\$6<=\$B\$6	Binding	0
43		\$D\$2	Value Selected	6.24	\$D\$2<=\$B\$2	Not Binding	18.7627119
44							

Fig. 9.10 Answer report for YRA with 2 units project type 3

Elizabeth’s reply to Nick’s first issue is that complying with the customer’s request for selecting 2 units of project type 3 will lead to a \$25,848 loss, and the number of the other projects selected will change slightly. All in all, this is not a terribly disruptive outcome, given the necessity to satisfy important clients.

Now for the remaining questions, consider the action Elizabeth has taken to re-solve the problem under the new constraint condition. If she is to answer the question related to the change in revenue, she must be careful with the project type 3 changes imposed on the solution, above. In the case of the small change Nick suggested (2 projects), we can see from comparing the *Sensitivity Reports* in Figs. 9.9 and 9.11, *pre* and *post* inclusion of project 2, that the *Adjustable Cells* area is unchanged (the Constraints area is changed). The reduced costs for both reports are identical, but if the change to the project type 3 was greater, for example 29 units, the two reports would not be the same. Thus, she can answer the questions without regard to the change (2 required) in type 3 projects. Since it is difficult to know how large a change is necessary to change the *Reduced Costs* and related *Allowable Increase* and *Decrease*, it is wise to re-solve the problem, and not take the

Variable Cells		Final Value	Reduced Cost	Objective Coefficient	Allowable Increase	Allowable Decrease
Cell	Name					
\$D\$2	Value Selected	6.237288136	0	45000	1800	6160.958904
\$D\$3	Value Selected	17.75	0	63000	6098.305085	2796.610169
\$D\$4	Value Selected	2	-12923.72881	27500	12923.72881	1E+30
\$D\$5	Value Selected	2.830508475	0	19500	10312.5	750
\$D\$6	Value Selected	16	27932.20339	71000	1E+30	27932.20339
\$D\$7	Value Selected	19	32673.72881	56000	1E+30	32673.72881
\$D\$8	Value Selected	36	15245.76271	48500	1E+30	15245.76271

Constraints		Final Value	Shadow Price	Constraint R.H. Side	Allowable Increase	Allowable Decrease
Cell	Name					
\$A\$23	Res-A hrs. used	734.4957627	0	800	1E+30	65.50423729
\$B\$23	Res-B hrs. used	900	3686.440678	900	33.4	52.57142857
\$C\$23	Res-C hrs. used	700	699.1525424	700	20.44444444	61.5
\$D\$23	Res-D hrs. used	375	152.5423729	375	73.6	13.91666667

Fig. 9.11 Sensitivity report for new solution

risk of missing the change. Now, the answer to Nick’s question about changes in revenue:

1. The allowable change in the revenue of type 2 projects is from \$60,203.39 to \$69,098.31. This range includes the \$60,500.00 to \$68,000.00 that Nick has proposed; thus, the current solution will *not* be changed by Nick’s change in type 2 revenue.
2. The allowable change for type 4 projects is \$18,750.00 to \$29,812.50, which includes Nick’s lower boundary of \$19,000.00, but the upper boundary, \$32,000.00, exceeds the allowable upper boundary. Therefore, if Nick believes that the revenue can indeed be greater than \$29,812.50, then the decision variables in the current optimal solution will be changed. The current revenue of \$19,500.00 appears to be quite low given Nick’s new range, and he should probably revisit his point estimate, and attempt to improve his estimate.

Finally, Nick has asked Elizabeth to determine the effect of increasing the RHS of the resource C by 12 hours from 700 to 712. We can see from Fig. 9.11 that the *Shadow Price* for the resource is 699.15, and we are within the allowable *Constraint R.H. Side*, 638.5 (700–61.5) and 720.44 (700 + 20.44). The results will be an increase to the objective function of \$8389.80 (12*699.15).

The value of LP is not just the determination of a set of decision variables that optimize an objective function. As important as the optional solution, are the uses of the sensitivity analysis associated with the solution. We have seen the power of the *shadow price*, *reduced cost*, and *allowable changes of coefficients* and *RHS’s* in our YRA example. Since LP is a deterministic procedure (all parameters are point estimates), sensitivity analysis permits the consideration of ranges of parameter values, even though we have represented the ranges with point estimates. Thus, the reality of the uncertainty associated with parameters can be recognized and investigated.

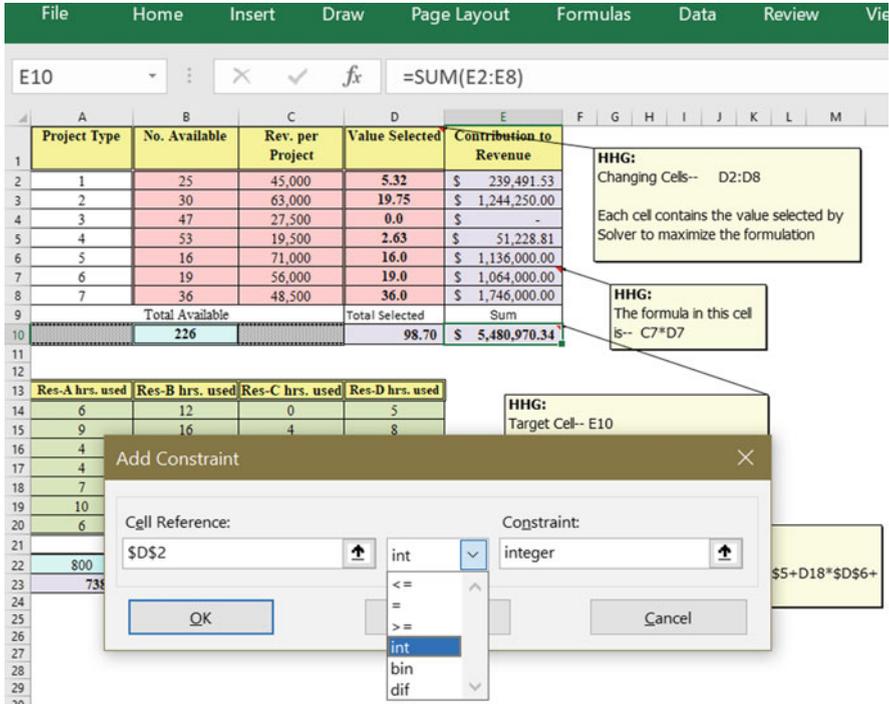


Fig. 9.12 Imposition of integer decision variables

Solver is not restricted to the solution of linear programs. A related class of problems is known as **Non-Linear Programs** (NLP) that, broadly defined, contain non-linear relationships in the objective function and/or in constraints. These can be very difficult problems to solve. Additionally, there are **Integer Programs** (IP) where the decision variables are restricted to integer values, **Mixed Integer Programs** (MIP) where decision variables can be both continuous (fractional value) and integer values, and **0–1 Integer Programs** where variables are binary (having two states). Again, these conditions require substantially more complicated solution algorithms than LP.

Earlier, we ignored the integer nature of the decision variables by assuming that we could simply round the variables, and not worry about the potential violation of constraints, as long as they were not severe. What if the integer condition is important? We can impose integer values on the decision variables by adding constraints to the problem. Figure 9.12 demonstrates how to use the *Add Constraint* dialogue box to declare the number of project type 1 to be restricted to integer. By selecting the *int* designation in the pull-down menu, between the *Cell Reference* and *Constraint*, we can convert the decision variables to integers. The *Add Constraint* dialogue box shows the integer constraint for cell D2, project type 1. Figure 9.13 shows the new solution (after imposing integer conditions on all decision variables),

The screenshot shows the Solver Parameters dialog box with the following data:

Objective Cell (Max)				
Cell	Name	Original Value	Final Value	
\$E\$10	Sum	\$ 5,480,970.34	\$ 5,477,500.00	

Variable Cells					
Cell	Name	Original Value	Final Value	Integer	
\$D\$2	Value Selected	5.32	7.00	Integer	
\$D\$3	Value Selected	19.75	19.00	Integer	
\$D\$4	Value Selected	0.0	0.0	Integer	
\$D\$5	Value Selected	2.63	1.00	Integer	
\$D\$6	Value Selected	16.0	16.0	Integer	
\$D\$7	Value Selected	19.0	19.0	Integer	
\$D\$8	Value Selected	36.0	36.0	Integer	

Constraints						
Cell	Name	Cell Value	Formula	Status	Slack	
\$D\$23	Res-D hrs. used	366.00	\$D\$23<=\$D\$22	Not Binding	9	
\$B\$23	Res-B hrs. used	900.00	\$B\$23<=\$B\$22	Binding	0	
\$C\$23	Res-C hrs. used	697.00	\$C\$23<=\$C\$22	Not Binding	3	
\$A\$23	Total Resource Hrs. Available in Quarter	735.00	\$A\$23<=\$A\$22	Not Binding	65	
\$D\$3	Value Selected	19.00	\$D\$3<=\$B\$3	Binding	0	
\$D\$4	Value Selected	0.0	\$D\$4<=\$B\$4	Binding	0	
\$D\$5	Value Selected	1.00	\$D\$5<=\$B\$5	Not Binding	1	
\$D\$8	Value Selected	36.0	\$D\$8<=\$B\$8	Binding	0	
\$D\$7	Value Selected	19.0	\$D\$7<=\$B\$7	Binding	0	
\$D\$6	Value Selected	16.0	\$D\$6<=\$B\$6	Binding	0	
\$D\$2	Value Selected	7.00	\$D\$2<=\$B\$2	Not Binding	18	
\$D\$2	Value Selected	7.00	\$D\$2<=\$B\$2	Not Binding	18	
\$D\$3=integer						
\$D\$4=integer						
\$D\$5=integer						
\$D\$6=integer						
\$D\$7=integer						
\$D\$2=integer						
\$D\$8=integer						

Fig. 9.13 Answer report for integer variables

where we can see that all variables result in integer values, and the value of the associated objective function is \$5,477,500.00, not a great deal below the continuous variable solution. The projects selected are:

$$X_1 = 7.0; X_2 = 19.0; X_3 = 0.0; X_4 = 1.0; X_5 = 16.0; X_6 = 19.0; X_7 = 36.0$$

Note though, that the *Sensitivity Report* for integer and binary programs are no longer valid due to the algorithm used to solve these problems. This of course is unfortunate, and a weakness of the approach, since we will no longer have the sensitivity analysis available to answer questions. Now, let us move on to another applications tool—*Scenarios*.

9.4 Scenarios

The Scenarios tool is one of the *what-if analysis* tools in Data ribbon of Excel. It also has been incorporated into Solver as a button on the *Solver Results* dialogue box, *Save Scenario*. The basic function of Scenarios is to simplify the process of management, record keeping, and entry of data for the repeated calculation of a spreadsheet. It is often the case that we are interested in asking repeated *what-if* questions of a spreadsheet model. The questions are generally of the form—what if we change the inputs to our model to *this*, then to *this*, then to *this*, etc.

You will recall that we dealt with this question when we introduced Data Tables. Data Tables display the value of a particular calculation, as one or two inputs are varied. Although this is a powerful tool, what if we have many more than two inputs to vary? We may need to construct many Data Tables, but the comparison between tables will be difficult at best. Scenarios permit you to determine the changes in a calculated value while varying as many as 32 inputs, and each different set of input values will represent a scenario.

9.4.1 Example 1—Mortgage Interest Calculations

After many years of hard work, Padcha Chakravarty has experienced great success in her import-export business. So much so, that she is considering the purchase of a yacht that she can claim as a second home. It meets the United States Internal Revenue Service criteria for a second home, by being capable of providing “sleeping, cooking, and toilet facilities”, and it is a very convenient way to reduce her tax burden in the coming years. The mortgage interest deduction is one of the few remaining personal income tax deductions available in the US tax code.

Padcha has decided that a short-term mortgage of 4–6 years (these are the shortest terms she can find) is in her best interest, since she may sell the yacht soon (2–3 years) after the capture of the initial tax advantages. Knowing that mortgage payments consist overwhelmingly of interest in early years, she is interested in finding a loan structure that will lead to a beneficial interest tax deduction while satisfying other criteria.

Padcha decides to construct a spreadsheet that calculates the cumulative interest paid over 2 years for numerous scenarios of principal, term, and interest rate. She has discussed the problem with a yacht broker in Jakarta, Indonesia, and he has provided six yacht options for her to consider. He is willing finance the purchase, and has forwarded the following scenarios to Padcha. See Table 9.1:

A spreadsheet for the calculation of the scenarios is shown in Fig. 9.14. In Fig. 9.14, we introduce a new cell formula (see C18 and C19) that is part of the financial cell formulas contained in Excel—**CUMIPMT** (rate, nper, pv, start_period, end_period, type). It calculates the cumulative interest paid over a specified number of time periods, and contains the same arguments as the PMT

Table 9.1 Scenarios for Yacht purchase

Yacht	A	B	C	D	E	F
Interest (%)	7	6.75	6.5	6.25	6	5.75
No. of periods	72	72	60	60	48	48
Principal	160,000	150,000	140,000	180,000	330,000	360,000

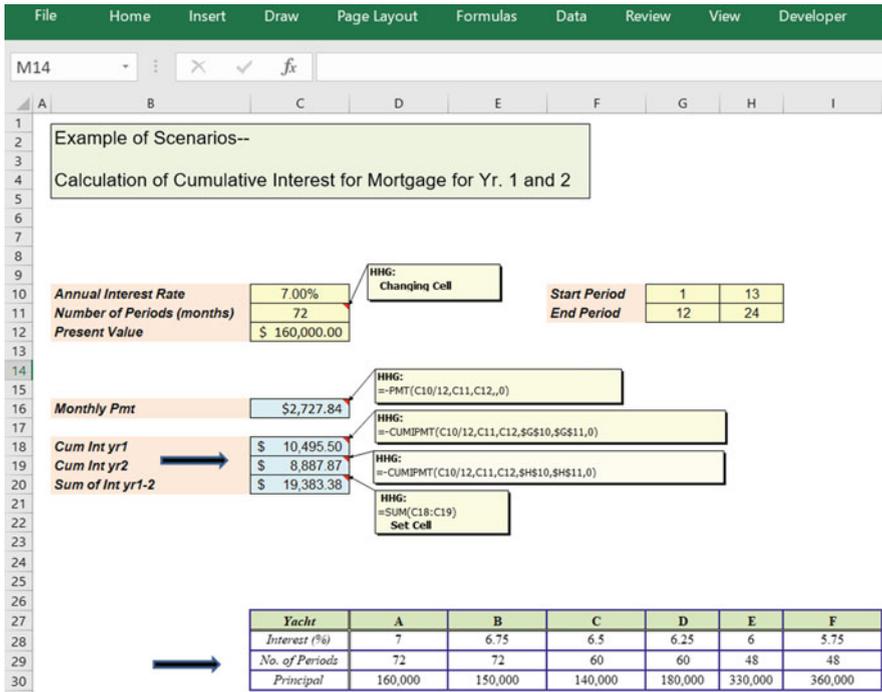


Fig. 9.14 Scenarios example for mortgage problem

cell formula. There are also two additional inputs, start_period and end_period; they identify the period over which to accumulate interest payments.

For Padcha’s mortgage problem, the periods of interest are the first year (months 1–12) and the second year (months 13–24). Thus, the first payment will begin in January and the last will be in December. Since income taxes are paid annually, it makes sense to accumulate over a yearly time horizon. If payments do not begin in January, we must select the end_period to reflect the true number of payments of interest in the initial year. For example, if we begin payment in the month of September, the start_period is 1 and the end_period is 4, indicating that we accumulated interest payments for 4 months, September through December. At the bottom of Fig. 9.14 are the values of the six scenarios, A through F, for Padcha’s model.

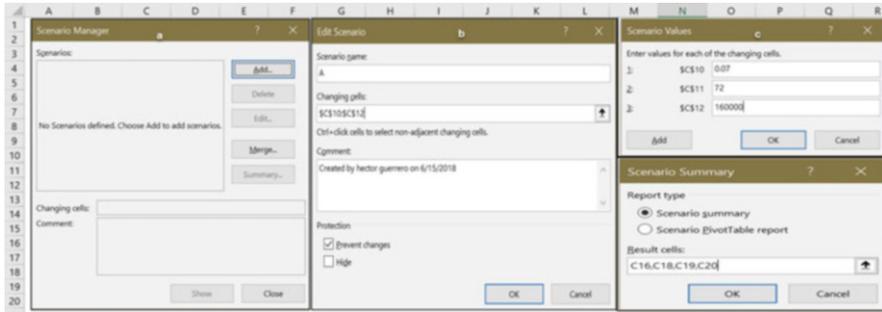


Fig. 9.15 Creating scenarios

So how do we create a scenario? The process of creating scenarios is shown in Fig. 9.15a, b, and c and is described as follows:

1. We begin by engaging the What-If Analysis tools in the Data Ribbon and Data Tools Group. Section *a* of Fig. 9.15, *Scenario Manager*, shows the first dialogue box encountered. As you can see, no scenarios are currently defined.
2. In section *b* of Fig. 9.15 we depress the *Add* button, and the *Edit Scenario* dialogue box becomes available. Here we name scenarios and identify *Changing Cells*: the cells that contain the data inputs for the calculations of interest.
3. Next, the *Scenarios Values* dialogue box permits entry of the individual values for the cells, as shown in section *c* of Fig. 9.15. Note that Excel recognizes the cell name for C10:C12—C10 as *IntRate*, C11 as *NumPeriods*, and C12 as *Principal*. The cell ranges were *named* in the spreadsheet; thus, providing easy identification for *Changing Cells* and *Result Cells*.
4. The process is repeated for each scenario by selecting the *Add* button on the *Scenarios Values* dialogue box.
5. When you return to the *Scenario Manager* dialogue box by selecting *OK*, the named scenarios will appear in the window.
6. Finally, we are able to select the *Summary* button to generate a report, as either a *Scenario summary* or a *Scenario PivotTable report*, as shown in the *Scenario Summary* dialogue box in section *c* of Fig. 9.15.

The resulting *Scenario summary* report is shown in Fig. 9.16. In this report, we can see the data of interest: (1) *MntlyPmt* is monthly payment for the mortgage, (2) *CumIntyr1* and *CumIntyr2* are cumulative interest payments in years 1 and 2, respectively, and (3) *SumIntyr1_2* is the sum of year 1 and 2 cumulative interest. The report provides a convenient format for presenting comparative results. If Padcha believes she would like to generate the highest interest deduction possible, she may consider either scenarios E or F. If more modest interest deductions are more appealing, then scenarios B and C are possible. Regardless, she has the entire array of possibilities to choose from, and she may be able to generate others based on the results she has observed, for example the Current Values shown in column

	A	B	C	D	E	F	G	H	I	J
1	Scenario Summary									
2	Current Values:									
3		Created by Author on 6/15/2018	Created by Hector guerrero on 6/15/2018							
4	Changing Cells:									
5		7.00%	7.00%	6.75%	6.50%	6.25%	6.00%	5.75%		
6	IntRate	7.00%	7.00%	6.75%	6.50%	6.25%	6.00%	5.75%		
7	NumPeriods	72	72	72	60	60	48	48		
8	Principal	\$ 160,000.00	\$ 160,000.00	\$ 150,000.00	\$ 140,000.00	\$ 180,000.00	\$ 330,000.00	\$ 360,000.00		
9	Result Cells:									
10	MnlyPymt	\$2,727.84	\$2,727.84	\$2,539.38	\$2,739.26	\$3,500.87	\$7,750.06	\$8,413.41		
11	CumIntYr1	\$ 10,495.50	\$ 10,495.50	\$ 9,483.54	\$ 8,378.87	\$ 10,353.36	\$ 17,753.05	\$ 18,550.64		
12	CumIntYr2	\$ 8,887.87	\$ 8,887.87	\$ 8,022.12	\$ 6,738.58	\$ 8,317.13	\$ 13,111.94	\$ 13,685.15		
13	SumInt Yr1-2	\$ 19,383.38	\$ 19,383.38	\$ 17,505.66	\$ 15,117.46	\$ 18,670.49	\$ 30,864.99	\$ 32,235.79		
14	Notes: Current Values column represents values of changing cells at									
15	time Scenario Summary Report was created. Changing cells for each									
16	scenario are highlighted in gray.									

Fig. 9.16 Scenario summary for mortgage problem

D. This ability to manage multiple scenarios is a very attractive feature in spreadsheet analysis.

9.4.2 Example 2—An Income Statement Analysis

We now consider a slightly more complex model for scenario analysis. In this example, we focus on a standard income statement and a related set of scenarios that are provided by a decision maker. The decision maker would like to determine the *bottom-line* (net profit) that results from various combinations of input values. In Fig. 9.17 we can see that we have 7 input variables, and each variable has two possible values. This is not a particularly complex problem, but with a greater number of possible input values, this problem could easily become quite cumbersome. The 7 input values represent standard inputs that are often estimated in *proforma* Income Statement analysis:

- Sales Revenue = (Volume)*(Price)
- COGS = (percentage⁴)*(Sales Revenue)
- Variable Operating Expense = (percentage)*(Sales Revenue)
- Fixed Operating Expenses
- Depreciation Expense
- Interest Expense

Obviously, we cannot use a two variable Data Table for this type of analysis; there are too many variables to consider simultaneously. This example is an

⁴The estimation of Cost of Goods Sold (COGS) and Variable Operating Expense as a percentage (%) of Sales Revenue is common approach to estimation of Income Statements, but not an approach without its detractors.

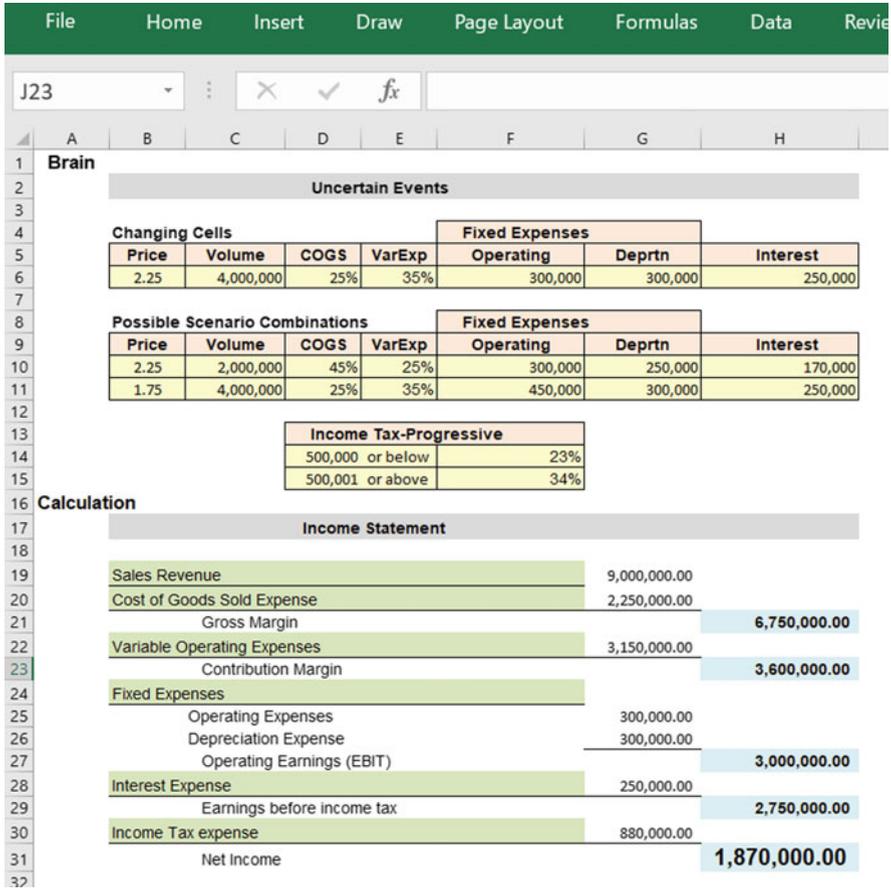


Fig. 9.17 Income statement analysis example

excellent use of the Scenarios tools. Figure 9.18 shows the results of the 4 scenarios. They range from a loss of \$300,000 to a gain of \$2,417,800.

9.5 Goal Seek

There is another tool in Excel’s *What-If Analysis* sub-group, Goal Seek. It is similar to Solver, except that it functions in reverse: it determines the value of an input that will *result* in a specified output. While Solver can manipulate numerous variables, and has a generalized goal to maximize or minimize the objective function, Goal Seek knows *a priori* the goal, and must find a single variable value, among several, to arrive at the goal. For example, assume that you want to have a payment of exactly

	A	B	C	D	E	F	G	H
1								
2		Scenario Summary						
3		Current Values:		A	B	C	D	
4		Changing Cells:						
5		Price	2.25	2.25	1.75	1.75	2.25	
6		Volume	4,000,000	4,000,000	2,000,000	2,000,000	4,000,000	
7		COGS	25%	25%	25%	45%	25%	
8		VarExp	35%	35%	25%	35%	25%	
9		FixOpExp	300,000	300,000	300,000	450,000	450,000	
10		Deptn	300,000	300,000	300,000	300,000	300,000	
11		Interest	250,000	250,000	170,000	250,000	170,000	
12								
13		Result Cells:						
14		NetIncome	1,870,000.00	1,870,000.00	701,800.00	(300,000.00)	2,417,800.00	
15		Notes: Current Values column represents values of changing cells at time Scenario Summary Report was created. Changing cells for each scenario are highlighted in gray.						
16								
17								

Fig. 9.18 Income statement scenarios

\$1000 for a loan. There are three inputs in the PMT function—interest rate, number of periods, and present value of the loan principal. Goal Seek will allow the user to select one of the three inputs to result in a payment of \$1000 per period. It is a limited tool, in that it will permit only a single variable to be changed to arrive at the goal. Thus, it is not possible to vary interest rate and number of periods and present value, simultaneously.

In the next section, we will examine two examples that demonstrate the power and the pitfalls of Goal Seek. The first example is relatively simple and relates to the calculation of Padcha’s loan, in particular the PMT function. The second example is a more complex application related to Padcha’s problem of accumulating interest in years 1 and 2, and it utilizes the CUMIPMT cell function. Although the PMT function is similar to the CUMIPMT function, the application of Goal Seek to the latter cell function is somewhat problematic.

9.5.1 Example 1—Goal Seek Applied to the PMT Cell

Consider the mortgage example we introduced in the Scenarios section. Imagine that Padcha has determined the yacht that she will purchase, the Queen of Malacca, along with its price, \$240,000. The broker for the yacht has agreed to finance at an interest rate of 7%; he is anxious to sell the Queen of Malacca due to some rather unfortunate history of the yacht’s previous owners—pirates and gun runners. He is not concerned with the *term* of the loan, as long as he gets an agreement to purchase. Padcha sees an opportunity to *set* a loan payment, and determine the term that will be implied given the broker’s interest rate and the principal of the loan. She decides that \$5000 per month is a very manageable loan sum for her. Figure 9.19 shows the Goal Seek dialogue box for Padcha’s problem. There are three entries:

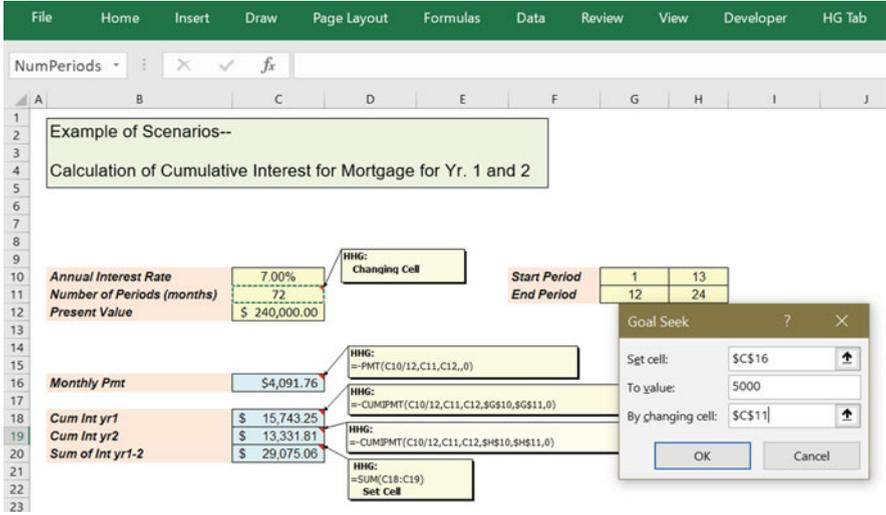


Fig. 9.19 Goal seek for term of PMT function

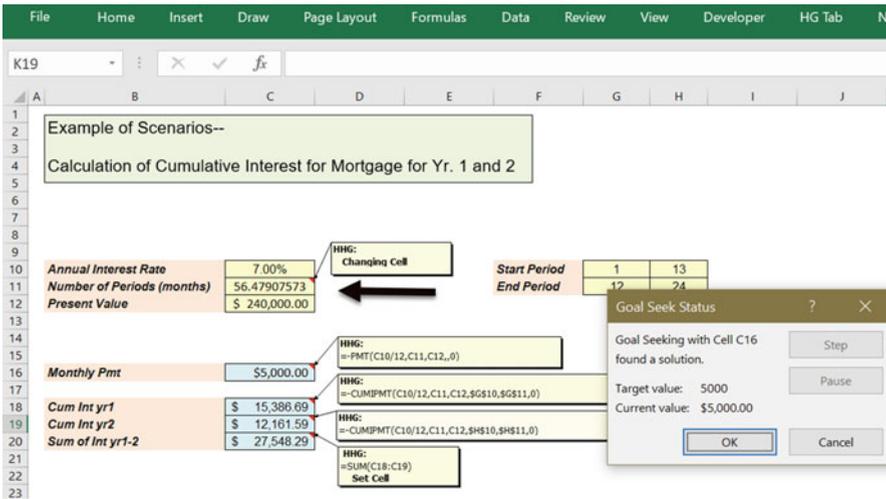


Fig. 9.20 Goal seek solution to PMT of \$5000

1. Set cell entry is the cell that she will set as a goal—*Monthly Pmt*, C16.
2. To value is the value she selects for the Set cell—\$5000.
3. By changing cell is the cell where changes will be permitted—*Number of periods (months)*, C11.

Figure 9.20 shows the results of the Goal Seek. The term that will lead to a loan payment of \$5000 per month is 56.47907573 months, or approximately 56. The

solution is found in a fraction of a second; thus, you could perform many *what-if* scenarios with little effort and in a minimal amount of time. Now, let us move to the next example to see how we might run into problems with Goal Seek in more complex situations.

9.5.2 Example 2—Goal Seek Applied to the CUMIPMT Cell

Suppose that Padcha, after some consideration, has decided that she would like the sum of 2 years of cumulative interest to be exactly \$25,000: this is her new goal. As before, she has decided on the level of investment she would like to make, \$240,000, and the interest rate that the yacht broker will offer on financing the purchase is 7%. Thus, the variable that is available to achieve her goal is the term of the loan. This appears to be an application of the Goal Seek tool quite similar to Example 1. As before, the tool seeks to obtain a goal for a calculated value, by manipulating a single input. Note that the calculated value is much more complex than before (CUMIPMT), but why should that make a difference? In fact, this more complex calculation may make a very significant difference in the application of Goal Seek.

We will repeat the Goal Seek for a new set of inputs, and now we will change the *Set cell* entry to C20, the sum of 2 years of accumulated interest, and the *To value* entry to \$25,000. The *Changing cell* entry will remain C11. In Fig. 9.21, we see the new Goal Seek entry data, and in Fig. 9.22 the results of the Goal Seek analysis. The results are a bit troubling in that the dialogue box indicates that the tool “*may not have found a solution.*” How is this possible? The algorithm used to find solutions is a search technique that does not guarantee a solution in all cases. Additionally, these types of algorithms are often very sensitive to where the search starts, i.e. they use

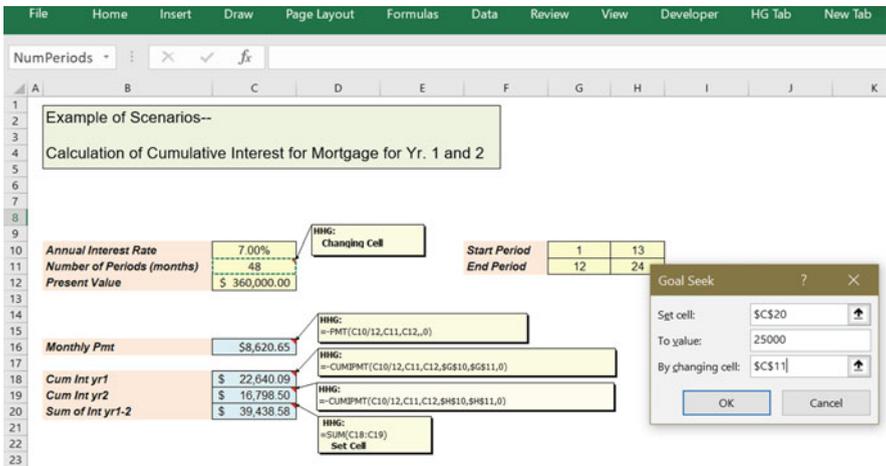


Fig. 9.21 Goal seek for cumulative interest payments

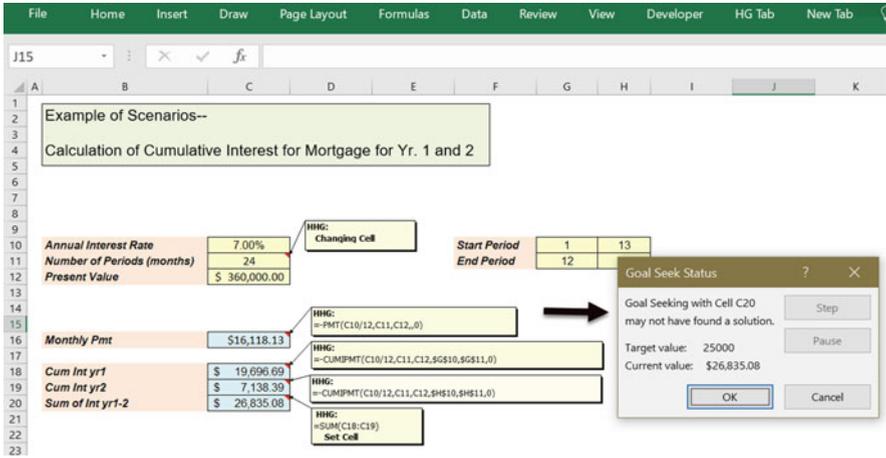


Fig. 9.22 Uncertain goal seek status of cumulative interest

the value that is currently in the cell to begin the search for the goal. In the case of Fig. 9.21, the changing cell contained 48 periods, so this is where the search began. The search terminated at 24 periods and a cumulative sum of \$26,835.08, but the tool was unsure of the solution. The problem we face is that it is impossible to achieve a \$25,000 in a term of greater than or equal to 24 months, and the problem required that 24 months be used in the calculation period. But, some experimentation shows that the end period in cell H11 can be changed to 18 and 19 months to achieve a value very near \$25,000, \$24,890.93 and \$25,443.72, respectively. Obviously, this is a complex condition and may take considerable experience before it is easily identified by an Excel analyst.

9.6 Summary

Solver, Scenarios, and Goal Seek are extremely powerful tools for quantitative analysis. Yet, we must be careful to use these tools with caution. In the words of US President Reagan—*Trust, but verify*. We have seen how a complex goal seek function can lead to problems, if some forethought is not applied to our analysis. The nature of the search algorithms that are used in the Solver and Goal Seek tools and the possible non-linear nature of the problem structure can baffle the search algorithm; it can lead to uncertainty in the veracity of the answer, or it can also lead to wrong answers. Although we did not spend much time discussing non-linear programs in our Solver section, other than to say they were very difficult to solve, it is not wise to assume that an optimal solution is always optimal. If the objective function and/or the constraints of a problem are non-linear, you might experience a solution that is a **local optimum**. A local optimum occurs when the search

algorithm assumes that it need not search any further for a better solution, but in doing so, it has actually ignored other regions of the function where better solutions are possible.

How are we to know when we have a local optimum, or that a solution that has been identified as optimal is possibly not optimal? A little common sense is invaluable in making this determination. Here are a few tips that might help you avoid accepting the claim of an optimal solution when it is not, or help you verify whether an uncertain solution is in fact optimal:

1. If you have a non-linear target cell or objective function for a formulation in a *single* variable, attempt to plot the function by using successive values of inputs to see if the function might be a candidate for a local optimum. You can do this by copying the function to a long column of cells and placing consecutive values of input in an adjacent column. Then plot the results and note the shape of the curve. Of course, this is only possible for a single variable, and in most problems we have far more than one input variable.
2. In the case of multi-variable problems, you may want to resort to simulation of inputs, and to see if you can find some combination that outperforms the so-called optimal solution.
3. If a solution is uncertain, but appears to be correct, investigate by examining values near the solution that is proposed. Be careful to consider a local optimum condition.
4. Be careful to note any odd solutions—negative values where none are possible and values that are either too large or too small to accept as possible.
5. Verify that the constraints that are imposed on a formulation are satisfied.
6. Remember that in spite of your best efforts you may still, on rare occasions, have problems dealing with these issues.

There is nothing more embarrassing than presenting a solution that contains a clear inconsistency in a solution which you have overlooked. Verification of an analysis is much like editing—it is not a pleasant task, but it is foolhardy to avoid it.

Key Terms

Solver	Changing Cell
Prescriptive analysis	Right-hand side (RHS)
Scenario	Slack
Goal seek	Not binding
Descriptive analysis	Binding
Constrained optimization	Shadow Price
Linear programming	Allowable increase
Decision variables	Allowable decrease
Objective function	Reduced cost

(continued)

Solver	Changing Cell
Constraints	Non-linear programs (NLP)
Technology of LP	Integer programs (IP)
Infeasibility	Mixed integer programs (MIP)
Coefficients	0–1 integer programs
LP formulation	CUMIMTP
Target cell	Local optimum

Problems and Exercises

- Name 2 types of Prescriptive Analysis and 2 types of Descriptive Analysis.
- Simulation is to Linear Programming as Descriptive is to?
- Constrained optimization optimizes an objective function without regard to factors that constrain the selection of decision variables—T or F?
- Decision variables in Linear Programming are always integer valued—T or F?
- Identify the following relationship as either linear or non-linear:
 - $2X + 3Y = 24$
 - $4/X + 3Y^2 = 45$
 - $3XY - 8Y = 0$
 - $4X = 6Y$
- For the following linear programs, what is the solution? Do not use Solver; use strict observation:
 - Maximize: $Z = 4X$; Subject to: $X \leq 6$
 - Maximize: $Z = 2X + 5Y$; Subject to: $X = 5$ and $X + Y \leq 12$
 - Minimize: $Z = 12X + 2Y$; Subject to: $X \geq 3$ and $Y \geq 4$
 - Minimize: $Z = X - Y$; Subject to: $X \geq 0$ and $Y \leq 26$
- Knapsack Problem*—Consider a number of possible investments in contracts for various projects that can fit into a budget (knapsack). Each investment in a contract has a cost and a return in terms of millions of dollars, and the contracts can be purchased in multiple quantities. This problem is known as the Knapsack Problem due to its general structure—selecting a number of items (projects in this case) for a limited budget (knapsack).

Available contract types	# of contract		
	investments available	Cost per contract	Value of contract
Project 1	4	1	2
Project 2	3	3	8
Project 3	2	4	11
Project 4	2	7	20

Budget ≤ 20 ; $X_j = \#$ of contracts of project j with possible values 0, 1, 2,

- (a) Formulate the LP for this problem.
 - (b) Solve the problem with Solver.
 - (c) What is the marginal effect on the objective function of adding one unit of each project to *# of contract investments available*? Recall this sensitivity analysis is done by looking at each decision variable, one at a time.
 - (d) How does the solution change if the budget constraint is reduced by 5 units to 15?
8. *Nutrition Problem*—A farmer raises hogs for profit. They are organic pigs of the highest caliber. The farmer is a former nutritionist and has devised the following table of possible nutrients and minimum daily requirements for 3 nutritional categories. For example, a Kilogram of corn provides 90 units of carbohydrates, 30 of protein, and 10 of vitamins. Also, there are 200 units of carbohydrates required daily, and corn costs \$35 per kilogram. The farmer is interested in knowing the kilograms of Corn, Tankage, and Alfalfa that will minimize his cost.

Daily nutritional requirements				
	Kg. of Corn	Kg. of Tankage	Kg. of Alfalfa	Min. daily requirements
Carbohydrates	90	20	40	200
Protein	30	80	60	180
Vitamins	10	20	60	150
Cost \$	\$35	\$30	\$25	

- (a) Formulate the LP for this problem.
 - (b) Solve the problem with Solver.
 - (c) What is the marginal effect on the objective function of adding one unit of each ingredient (Min. daily requirements)? Recall this sensitivity analysis is done by looking at each decision variable, one at a time.
 - (d) How does the solution change if an additional 15 units of RHS are added for each of the nutrition constraints? Add them one at a time; for example, change 200–215, then 180–195, etc.
9. You have joined a CD club and you are required to purchase two types of music CD's in a year: Country Music (X) and Easy Listening (Y). Your contract requires you to purchase a minimum of 20 Country CDs. Your contract also requires you to purchase at least 20 Easy Listening CDs. Additionally, you must purchase a minimum of 50 CDs (both types—Country Music and Easy Listening) yearly. If the Country CDs cost \$7 per CD and the Easy listening cost \$10 per CD, what is the solution that minimizes your yearly investment in CDs?
- (a) Solve this LP with Solver.
 - (b) Which constraints are binding?
 - (c) Will the optimal solution (number of each CD type) change if the cost of Country CD's increases to \$9?

10. You are interested in obtaining a mortgage to finance a home. You borrow a principal of \$150,000 for 30 years. If you would like to have a monthly payment of \$700, what is the interest rate that will permit this payment?
11. The CUMPRINC() function is similar to the CUMIPMT() function, except rather than calculate the cumulative interest paid, the function calculates the cumulative principal paid in a period of time. You have a mortgage loan at 6% over 30 years for \$150,000.
- Use Goal Seek to find the approximate period in which \$5000 in principal payments have been accumulated.
 - Use Goal Seek to find the approximate period in which \$75,000 in principal payments have been accumulated.
 - If the use of Goal Seek is problematic for problem 11b, what do you suggest as an alternative to finding the approximate period that still uses the CUMPRINC() function?
12. Advanced Problem—*Shift Scheduling Problem*—I run a call center for customers using my *personal sensitivity training* tapes. I guarantee that an *understanding and caring voice* (UCV) will be available to customers 24 hours a day. In order to satisfy this promise, I must schedule UCV's based on historical demand data shown in the table below. I must determine how many UCV's must report to work at the beginning of each *Period*. Once a worker reports for duty, they will work an 8-hour shift. There is no restriction in which of the periods a worker starts a shift. For example, if a worker begins at 3 pm, then they will work until 11 pm.
- Formulate and solve the LP that will minimize the total workers needed to cover the overall UCV historical demand. Note that the assignment of a partial worker is not a realistic situation. Therefore, you may have to consider a constraint that guarantees integer UCV's.

Period	Time	UCV historical demand
1	3–7 am	3
2	7–11 am	12
3	11 am–3 pm	16
4	3 pm–7 pm	9
5	7–11 pm	11
6	11 pm–3 am	4

- What happens to the solution if the demand for UCV's in period 6 changes to 8? What is the new solution?
- How would you handle varying cost for the time periods? For example, what if the cost of the period 5 (7 pm–11 pm) and period 6 (11 pm–3 am) time period is twice as high as other time periods. How does the objective function change if you want to cover the UCV demand at the minimum cost under the new condition that costs of UCV's are not equal?