

Solver

We know Solver is a numerical methods approach (as opposed to analytical) to solving an optimization problem. We saw it in action as it minimized the time it takes to reach the victim in the Lifeguard Problem for running speeds of 5 m/sec and 10 m/sec.

This section is devoted to these two questions:

1. How does Solver actually work?
2. Can we always count on it giving us the correct answer?

The answer to the first question is explored in more detail below, but in a nutshell, Solver hunts and pecks really fast and converges to a solution. The answer to the second question is shorter: no. We will see that Solver can fail in two ways: *miserable* and *disastrous*. It can announce that it cannot find the answer (that is miserable) or, worse, it can claim to have an answer that is wrong. This is disastrous because you think it is right, but it is not. The bottom line is that you always need to be on high alert when using Solver—it is not a silver bullet.

How Does Solver Actually Work?

We begin our explanation of how Solver works with a simple profit maximization problem. A firm sells its product at \$4/unit and it has costs of production given by the square of the number of units it produces. So, it costs \$9 to make 3 units and \$25 to make 5 units. Since the price is \$4, the firm's revenue function is $4x$. If it sells 3 units, it makes \$12, while 5 units generates \$20 of revenue.

The firm seeks to maximize profits (denoted by convention with Greek letter pi, π) which are equal to revenues minus costs. The profit function below says that 3 units produces \$3 of profit (\$12 of revenue minus \$9 of costs) and 5 units leaves the firm with a loss of \$5 since $20 - 25 = -5$.

$$\max_x \pi = 4x - x^2$$

If you know calculus, you could solve this problem analytically by taking the derivative with respect to x , setting it equal to zero, and solving for x^* . But even without calculus, the solution is obvious if you tabulate (a single entry table that lists the values of a function at discrete points) and graph the profit function.

STEP Create data and make a chart of the profit function with the x -axis going from 0 to 5 in increments of 0.5. If you have trouble, see the appendix.

Both the data you created and the corresponding chart show that the answer is $x^* = 2$ leading to maximum profits of \$4. Can Solver also find the optimal solution? Yes, it can.

STEP Use cell A17 for x and enter the profit formula in cell A18. Next, call Solver and fill in the dialog box and click *Solve*. Again, if you have trouble, see the appendix.

As expected, Solver finds the optimal solution, but how does Solver do it? In a word, *iteration*, which means repetition. Solver runs the three steps below over and over until it cannot improve much:

1. Using the starting value in cell A17 (this is zero if left blank), it evaluates cell A18.
2. It then moves away from the starting value. The amount it moves is determined by the particular recipe—a popular one (and Solver’s default) is called Newton-Raphson steepest descent.
3. It compares the new value of profits to the original. If profits are higher, it continues in that direction. If lower, it goes in the opposite direction.

Solver continues iterating, changing cell A17 and checking how profits respond, until the improvement in profits is “small,” an amount determined by the convergence criterion. The last change it made to cell A17 is the answer.




Figure 2.10: Solver in action.

The *stylized graph* (which means it represents an idea without using actual data) in Figure 2.10 shows that Solver works by trying different values and seeing how much improvement occurs. The path of the choice variable (on the x -axis) is determined by Solver's internal optimization algorithm. See FrontlineSolvers at <https://www.solver.com/excel-solver-online-help> to learn more about Excel's Solver.

When Solver takes a step that improves the value of the objective function by very little, determined by the convergence criterion (adjustable via the *Options* button), it stops searching and announces success. In Figure 2.10, Solver is missing the optimal solution by a little bit because, if we zoomed in on the graph, the objective function would be almost flat at the top. When Solver computes minuscule additional improvement, it stops and announces it has found a solution.

This is all too abstract. Let's see it in action.

STEP Set cell A17 to 1.1 and widen column A to make sure it displays many decimal points. Call Solver, but this time, click the *Options* button and check *Show Iteration Results* (in the *All Methods* tab). Click Solve. Click Continue as many times as needed and watch the values on the spreadsheet.

Excel displays each trial solution. Solver might (this is not guaranteed) even hit 2 as a trial solution, but then it moves off it because it does not know this is the exactly correct answer.


It is quite likely that, starting from 1.1, Solver's answer is not exactly 2. You might have a result like this, 1.99999999467892. When we say Solver got the answer, we mean this in a practical sense. If Solver is off the exact answer by a tiny amount, that is success, for all practical purposes.

There is a danger to avoid when using Solver (or any numerical optimization method that relies on convergence): it is easy to conclude that Solver must give an exact answer because it displays so many decimal places. This is incorrect. Solver's solution is an example of *false precision*. It is not true that the many digits provide useful information. The exact answer is 2.

Usually, Solver does not find the exactly correct answer, yet it displays a number with many digits. This is *Solver noise* (a less technical way of saying false precision). You must learn to interpret Solver's results as inexact and not report all of the decimal places. You must use words like *roughly* or *approximately* when reporting answers produced by Solver.

In general, there are two reasons for really tiny disagreement between Solver and the exact answer.

1. Excel cannot display a number to an infinite number of decimal places. Most modern Excel software has 15 significant digits of precision. If the solution is a repeating decimal or irrational number, there is no exact decimal representation. Even if the number can be expressed as a decimal—for example, one-half is 0.5—precision error may occur during the computation of the final answer.

STEP Click in any cell and enter the formula =pi(). Now select the cell displaying π and add decimal places by repeatedly clicking the *Increase Decimal* button, , in the *Home* tab, *Number* group in the Ribbon. When Excel displays “#####,” it means the number does not fit in the column, so you need to widen the column. Keep adding decimal places until you start seeing zeroes. You have reached Excel's maximum precision.

2. Even if the answer is an integer, like 2, Excel's Solver often misses the exactly correct answer by small amounts. Solver's convergence criterion (that you can set via the *Options* button in the *Solver Parameters* dialog box) determines when it stops hunting for a better answer. Excel treats all of numbers the same—it has no way of saying, “Oh, that's really close to 2 so the answer must be 2.”

False precision or Solver noise is a serious issue. Be aware that Solver's answer is likely not the exactly right answer. There are other ways in which Solver can fail us that are more serious than misunderstanding precision.

Solver Behaving Badly

A *miserable result* (an actual, technical term in the numerical methods literature) occurs when an algorithm reports that it cannot find the answer or displays an obviously erroneous solution. We can easily create an example.

STEP Copy your existing sheet (right-click the sheet tab in the bottom left-hand corner, select *Move or Copy...* and check *Create a copy*). Select cell B2 and delete the $-A2^2$ part of the formula so you have only $=4*A2$. Fill down.

The chart updates and you have a straight line. You are graphing total revenue now instead of profits.

STEP Select cell A18 and change the formula to $=4*A17$ (just like you did to the chart). Call Solver and click Solve.

You are looking at a miserable result. Solver cannot find a solution so it announces (loudly, with a red exclamation mark) that it cannot find an answer.

In this case, that is actually reasonable. After all, this function does not have a well-defined maximum, unless you want to argue that the answer is positive infinity.

Solver will also give a message like this if it cannot perform a computation such as dividing by zero or taking the square root of a negative number. The algorithm fails, which is bad news, but at least it tells you that it cannot solve the problem.

When Solver fails like this, there are three basic strategies to get Solver to find a solution (but nothing is guaranteed):

1. Try different initial values in the changing cells. If you know roughly where the solution lies, start near it. Always avoid starting from zero or a blank cell.
2. Add more structure to the problem. Include non-negativity constraints on the endogenous variables, if appropriate.
3. Completely reorganize the problem. Instead of directly optimizing, you can put Solver to work on conditions that must be met.

None of these approaches will work on this problem because it does not have a solution.

There is a second way in which Solver can behave badly and it is much worse than a miserable result. Solver can give an answer that is wrong, yet it thinks it is right and reports it as the solution. This is called a *disastrous result*. Again, we can create an example to demonstrate this.

Our example requires use of a more complicated cost function. Instead of the simple square we used earlier, we will use a cubic cost function, $ax^3 + bx^2 + cx$. By strategically choosing the coefficients, $a = 1$, $b = -8$, and $c = 19$, we can get a cost curve that will give us a profit function with a trough and a peak, like a sine wave:

$$\max_x \pi = 4x - (x^3 - 8x^2 + 19x) = 4x - x^3 + 8x^2 - 19x$$

STEP Copy the original sheet again. Change cell B2's formula so it looks like this: $=4*A2-A2^3+8*A2^2-19*A2$. Fill it down.

The graph updates and now has a clear maximum near $x = 4$, but also has a valley as we increase output starting from zero. Can Solver find the optimal solution with this more complicated cost function?

STEP Edit cell A18 so that it uses the new profit function. The formula should look like this: $=4*A17-A17^3+8*A17^2-19*A17$. Call Solver and click Solve. How did it do?

If you started from the original sheet's optimal solution ($x = 2$) or near the maximum, then Solver will work fine. Notice that Solver's answer is a little higher than 4 and remember that Solver's answer is not the exact optimal solution, but it is really close (off by 0.000001 or so).

STEP Change cell A17 to 1 and run Solver. What happens now?

You just saw a disastrous result. Solver said it got an answer, but it is wrong. Your spreadsheet is showing cell A17 is zero, but we know the correct solution is a little over 4. This is really bad and shows that we must always be skeptical of Solver. It is not a silver bullet and mindless reliance on its output is a recipe for disaster.

Can you figure out why Solver failed?

STEP To help you, look at the chart and remember that we started from $x = 1$. Force your eyes to follow a straight line down from 1 (on the horizontal axis). You can see that you are to the left of the minimum of the profit function.

The fact that we are on a downward sloping part of the profit function explains why Solver failed. It started from there and the algorithm led it away from the true max (x becoming smaller) because profits rise (become less negative) as x decreases.

We can see that starting from near the maximum generates a good result because the function is a hill around the maximum. Moving away leads you to lower profits so the algorithm can find the correct solution.

STEP Start from 7 (set cell A17 to 7) and run Solver to see that starting values around the top of the hill enable Solver to succeed.

Will any starting value far away from zero work?

STEP Start from 100, then try 10.

This is surprising. If you start too far from the maximum, Solver takes a big first step that leads it to the wrong answer. Notice, however, that Solver does not give any indication that something may be wrong. To be clear, with a disastrous result, the computer's answer is incorrect, but you have no way of knowing this. That means you always have to be skeptical and vigilant when using Solver.

Unlike a miserable result where we know we have a problem, a disastrous result tricks us into thinking all is well. Even when Solver says it has an answer, you should question it. Try starting from a different place to see if you get the same answer. Always ask yourself if the answer makes sense. Be careful out there.

Takeaways

Solver is an example of a search algorithm. It works by exploring the objective function, like an ant crawling on a surface. If movement in a certain direction improves things, then it keeps going that way. Moves that generate worse results lead to reversing course.

Usually, this is an effective method. Many problems have solutions that can be found by numerical methods, using a computer to plug and chug through the problem.

Even when Solver finds the solution, always remember that its answer is not likely to be exactly correct. Solver usually suffers from what is formally known as false precision (and we call it, intuitively, Solver noise). This means that the last decimal places are not reliable and do not signal exactitude.

You want to be a sophisticated user and understand that, while powerful, Solver is not perfect. It can fail in two ways:

- 1) A miserable result is when Solver surrenders and announces that it cannot find an answer. This is disappointing.
- 2) A disastrous result is much worse. Solver says it has an answer, but it is wrong. This happens when you have a difficult problem, perhaps like asking an ant to find the highest point on a crumpled piece of paper. It is likely to find a local maximum, but not a true, global max.

Solver is a tool, like a chain saw. You can use it to cut down and trim trees really fast. You can also cut your leg off with it. Be careful.

Appendix

Charting the Profit Function

To make a chart of the profit function, begin with x -axis data. Create a label in cell A1 by entering the letter x . Below it, enter a 0, then a 0.5 below that. Select both cells, then fill down to cell A12, which should display 5.

For the y -axis, enter the label, profits, in cell B1. In the cell below it, enter the formula $=4*A2 - A2^2$. Fill it down. Format the profit values as \$.

Select the data and make a Scatter chart. Remember to add labels for the axes and an appropriate title. Your chart should look like Figure 2.11.

Using Solver to Find the Optimal Solution

Enter the labels in cells B17 and B18 shown in Figure 2.11. Enter the formula $=4*A17 - A17^2$ in cell A18. Call Solver (click the *Data* tab in the Ribbon and click *Solver*), then click on cell A18 for the objective function and cell A17 for the changing cells (as shown in Figure 2.11). Click *Solve* and click *OK* when Solver displays its results dialog box.

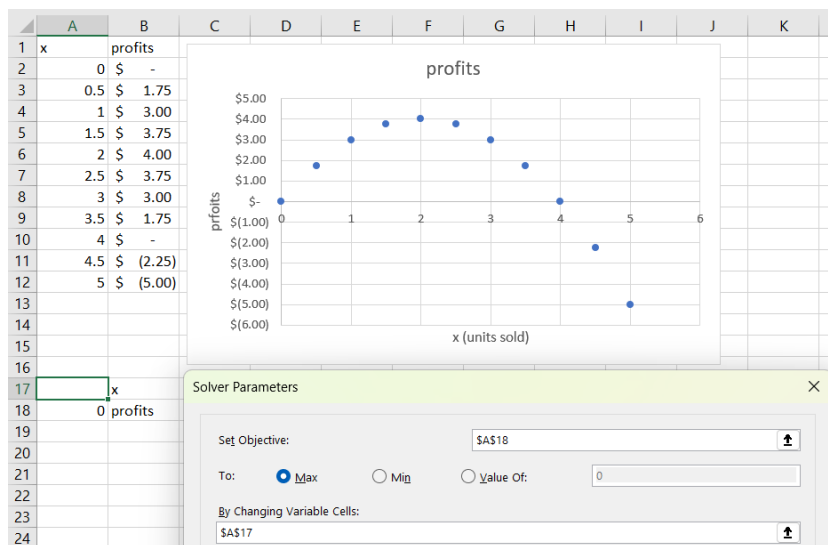


Figure 2.11: Tabulating and charting $4x - x^2$.