

3.3 – Objective Functions and Solving an LP

So far we’ve defined the general form of a linear programming model (LP), how to use it within the modeling process, and what it means to have a feasible solution. We now turn to identifying the best, or optimal, solution from among the feasible options.

From the initial reading in this block where we learned the form of an LP we know that an objective function is the mathematical expression that represents the goal of the model [1]. The general form for our objective function is:

$$\text{Minimize } \sum_{j=1}^n c_j x_j$$

where c_j represents the cost associated to the variable x_j for all j .

1 Formulating an Objective Function

Let’s explore how to formulate an objective function algebraically with an example. We’ll use the same context as the last reading, but with one key change; we are trying to *maximize* profit.

Problem 3.3.1: The Wyndor Problem [1]

The Wyndor Glass Co. produces high-quality glass products, including windows and glass doors. The company has three plants that simultaneously produce the components of its products. There are three plants that produce these items for the company. Doors require one hour of production time at Plant 1 and three hours of production time at Plant 3. Windows require two hour each at Plants 2 and 3. Plant 1 has 4 hours available each week, Plant 2 has 12 hours available each week, and Plant 3 has 18 hours available each week.

If Wyndor makes \$300 on each door they sell and \$500 on each window they sell, how many windows and doors should Wyndor Glass Co. produce each week to maximize profit?

Just as with constraints, we need to ensure we have an understanding of the decision that needs to be made. In Problem 1, the decision is still how many windows and doors to produce each week. The *objective* or goal of our problem is to maximize profit, but the total profit is not the decision. Profit, in this context, is how we evaluate and pick the best solution out of all of our feasible solutions. Therefore, the variables we defined in the last reading still hold here:

- x_1 = the number of doors produced in one week by Wyndor Glass Co.
- x_2 = the number of windows produced in one week by Wyndor Glass Co.

Next we need to define our cost coefficients. Remember that while we call them “cost” coefficients, the parameter

could take on a different meaning, like profit or risk. If our objective is to maximize profit, how do we calculate the profit Wyndor Glass Co. will make? We need to multiply the amount they make on each door and window by the number of doors and windows they make. This means our “cost” coefficients are the dollar value that Wyndor makes on each sale of a door or window. We can generally define our cost coefficient as

$$c_j = \text{the amount of money in \$ that Wyndor makes on the sale of one of product } j.$$

As before we can also specify the definition for each one if that make it easier:

- c_1 = the amount of money in \$ that Wyndor makes on the sale of one door.
- c_2 = the amount of money in \$ that Wyndor makes on the sale of one window.

The last piece required to define our objective function is the goal of our problem. While the general form of an LP model uses “Minimize,” we can just as easily use “Maximize” when our goal is to increase a quantity like profit. So our full algebraic formulation for our objective function to complete the solve step is:

$$\text{Maximize } c_1 x_1 + c_2 x_2$$

Our **full model** algebraic formulation for the Wyndor Problem that would be expected as the output of the transform step is:

$$\begin{array}{rcll} \text{Maximize} & c_1 & x_1 & + & c_2 & x_2 & & & \\ \text{subject to} & a_{11} & x_1 & + & a_{12} & x_2 & \leq & b_1 & \\ & a_{21} & x_1 & + & a_{22} & x_2 & \leq & b_2 & \\ & a_{31} & x_1 & + & a_{32} & x_2 & \leq & b_3 & \\ & & & & & & & & & x_1 & \geq & 0 & \\ & & & & & & & & & x_2 & \geq & 0 & \end{array}$$

Note: This model is considered complete as long as the the definitions we developed in these last two readings are also included.

2 Feasible vs. Optimal Solutions

Just as before, once we have completed the generic formulation in the transform step, we can apply the actual values to the model in the solve step. The formulation for the LP in Problem 1 that we would use to solve the problem is:

$$\begin{array}{rcll} \text{Maximize} & 300 & x_1 & + & 500 & x_2 & & & \\ \text{subject to} & & & & & & & & & x_1 & \leq & 4 & \\ & & & & & & & & & & & & & 2 & x_2 & \leq & 12 & \\ & 3 & x_1 & + & 2 & x_2 & \leq & 18 & \\ & & & & & & & & & x_1 & \geq & 0 & \\ & & & & & & & & & & & & & & x_2 & \geq & 0 & \end{array}$$

In the previous reading we developed the set of all feasible solutions by graphing the feasible region. We said

3.3 – Objective Functions and Solving an LP

that these represented all of the solutions that simultaneously satisfied all of our constraints, but they were not necessarily the best solution. Now with our object function defined, we can begin evaluating feasible solutions to determine which one is optimal, or the one that results in the greatest profit.

upper right corner of the graph, we can see the colors shift from blue to green to yellow, representing increasing profit values.

Definition 3.3.1 (Optimal Solution [2])
 For a minimization problem, the optimal solution is the point, (x_1, x_2) , in the feasible region with the smallest objective function value.
 For a maximization problem, the optimal solution is the point, (x_1, x_2) , in the feasible region with the largest objective function value.

Let’s look at the two feasible solutions we identified in the previous reading $(x_1, x_2) = (3, 4)$ and $(x_1, x_2) = (2, 4)$. Substituting each feasible solution into the objective function we can get a value for the profit made by Wyndor at each solution.

$$P(x_1, x_2) = 300x_1 + 500x_2$$

$$\begin{aligned} P(3, 4) &= 300(3) + 500(4) \\ &= 900 + 2000 \\ &= \$2900 \end{aligned}$$

$$\begin{aligned} P(2, 4) &= 300(2) + 500(4) \\ &= 600 + 2000 \\ &= \$2600 \end{aligned}$$

Comparing these two solutions, we can see that $P(3, 4) = 2900 > 2600 = P(2, 4)$, meaning that the solution to produce three doors and four windows results in more profit than producing two doors and four windows. Does this make $(3, 4)$ the optimal solution or is there a better feasible solution in our feasible region?

3 Solving for an Optimal Solution

Before we try to see if there is a better feasible solution let’s look at what is happening graphically. Just as we graphed our feasible region, we can graph the objective function. Figure 1 depicts the objective function graphed on top of the feasible region. The objective function is a plane in three dimensional space that rises over the feasible region, with the objective function values being along the z-axis.

On the objective function plane in Figure 1, the purple and blue colors represent lower profit values, beginning at $P = 0$ when $x_1, x_2 = 0$. As we move toward the




Figure 1: The objective function for Problem 1 graphed in conjunction with the feasible region. The corners of the feasible region are projected using the dashed lines to the corresponding points on the objective function as a reference.

Using this information, let’s look at the solution $(2, 4)$ and see where that lies on our objective function plane. Figure 2 plots the solution $(2, 4)$ within the feasible region as well as on the objective function plane at $(2, 4, 2600)$. We can also see the red line through the objective function plane that highlights every combination of x_1 and x_2 points that has a profit value of \$2600.

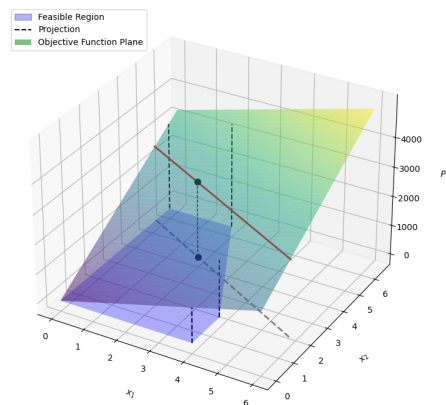


Figure 2: The solution $(2, 4)$ plotted within the feasible region and again on the objective function plane at $(2, 4, 2600)$. The red line along the objective function plane represents the line where every profit value is \$2600. The dashed gray line below the red line is the projection of objective function onto the two-dimensional plane represented by the x_1 - x_2 axes.

We can simplify our view to build greater understand-

3.3 – Objective Functions and Solving an LP

ing by focusing on the projected line representing all the points that produce the same objective function value in two dimensions as shown in Figure 3. The dashed gray line that bisects the feasible region shows every point with a profit, or objective function value, of \$2600. Any point where this line touches our feasible region will produce the same objective value. For example, say Wyndor decided to produce 1 door and 4.6 windows every week; their resulting profit is still \$2600. We can verify this with the graph, but also by solving the objective function.

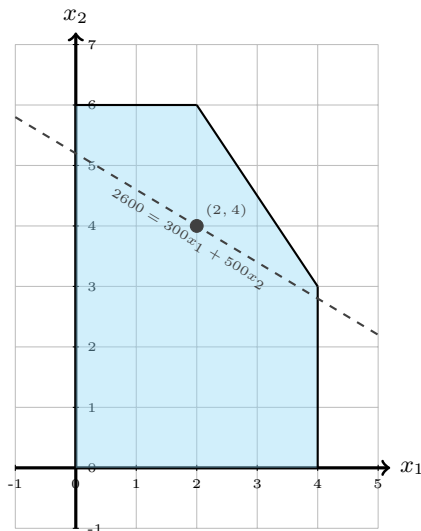


Figure 3: Projection of objective function $2600 = 300x_1 + 500x_2$ onto the 2D plane formed from the x_1 and x_2 axes for Problem 1. Note that the solution $(2, 4)$ falls onto this projection.

What about the solution $(3, 4)$ that resulted in a profit of \$2900? Is that the optimal solution? Hopefully, we can find one unique solution that is optimal. Let’s graph the objective function for $(3, 4)$ and see what happens.

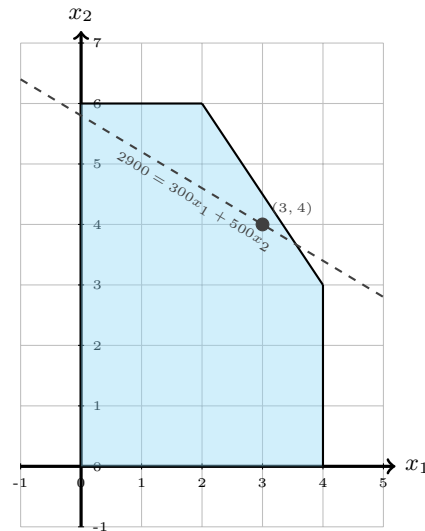


Figure 4: Projection of objective function $2900 = 300x_1 + 500x_2$ onto the 2D plane formed from the x_1 and x_2 axes for Problem 1. Note that the solution $(3, 4)$ falls onto this projection.

From Figure 4 we can see that at $(3, 4)$ the objective function still bisects the feasible region, meaning there are many solutions where the profit value is \$2900. We still don’t have a unique solution that results in the highest profit value. Let’s plot several objective function projections and see what is happening to build intuition of where we should look for the optimal solution.

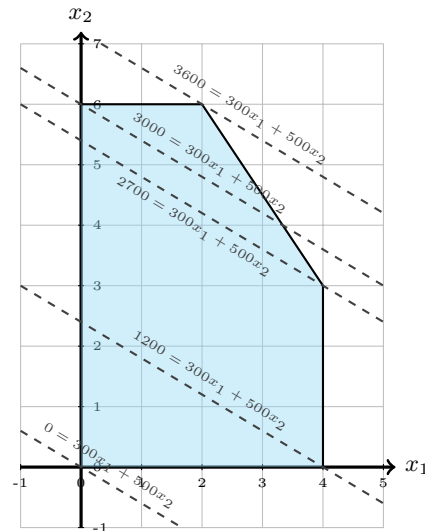


Figure 5: Projection of multiple objective functions with differing profit values onto the 2D plane formed from the x_1 and x_2 axes for Problem 1.

By examining multiple projected objective functions in Figure 5, we see how the level curves shift across the feasible region. Initially, we can see that when the profit value is \$0, the only solution is at $(0, 0)$, or to make 0 doors and 0 windows. As we increase the objective value, these lines shift parallel to each other through the feasible region and eventually intersect fewer and fewer points, until only one feasible point remains at $(2, 6)$. This represents the highest profit value we can get and still satisfy all the constraints at \$3600. Therefore, $(2, 6)$ is our optimal solution. This means that to maximize profit, Wyndor

should produce 2 doors and 6 windows every week.

You may have noticed that in Figure 5 we focused our attention on the corners, or vertices, of the feasible region. This is not a coincidence. In fact, it is a fundamental property of linear programming: when the objective function and all constraints are linear, and the feasible region is bounded, the optimal solution will always occur at one of the corner points of the feasible region. This makes our solution process much more efficient. Rather than testing random points throughout the feasible re-

3.3 – Objective Functions and Solving an LP

gion, we only need to evaluate the objective function at the coordinates of each corner. These corners represent the points where constraint boundary lines intersect, and they are the only locations where the objective function can achieve its maximum or minimum value. In our Wyndor example, this meant checking only a few candidate points and identifying (2, 6) as the optimal solution.

Graphically, this corresponds to “sliding” the objective function line across the feasible region. The last point it touches, before rising above or falling below the region, lies on one of the corners. This geometric insight is why solving a linear program often reduces to finding and evaluating just the corner points.

Note: In some problems, the optimal solution may occur not at a single point but along an entire edge of the feasible region. This happens when the objective function is parallel to one of the constraint boundaries. Also, if the feasible region is unbounded and the objective function improves indefinitely in some direction, there may be no optimal solution at all.

4 Solving an LP Using Excel

Graphical methods work well to solve LPs with only two variables, but what if we have a problem that has more than two variables? While there are methods to solving these by hand, for the purposes of this course, we are going to focus on using software to help us solve more complex LPs. While there are many different software applications that can solve LP’s, we will use Microsoft Excel Solver, a free add-on to Excel.

Excel spreadsheets can be built in many different ways, and it’s up to you to determine which way makes the most sense to you. This reading will present the Excel spreadsheet using matrices and vectors. Let’s think about the vectors and matrices we have: \mathbf{A} , \vec{x} , \vec{c} , and \vec{b} . We’ve already defined the components of these previously, but let’s review. Matrix \mathbf{A} is the matrix of technological coefficients:

$$\mathbf{A} = \begin{bmatrix} a_{11} & a_{12} & \dots & a_{1n} \\ a_{21} & a_{22} & \dots & a_{2n} \\ \vdots & \vdots & & \vdots \\ a_{m1} & a_{m2} & \dots & a_{mn} \end{bmatrix}$$

Vector \vec{x} is the column vector of variables:

$$\vec{x} = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix}$$

Vector \vec{c} is the row vector of cost coefficients:

$$\vec{c} = [c_1 \quad \dots \quad c_n]$$

And vector \vec{b} is the right-hand side column vector repre-

senting the total resource available:

$$\vec{b} = \begin{bmatrix} b_1 \\ \vdots \\ b_m \end{bmatrix}$$

We can apply the values from our problem to this and produce the following set of vectors and matrices:

$$\mathbf{A} = \begin{bmatrix} 1 & 0 \\ 0 & 2 \\ 3 & 2 \end{bmatrix}$$

$$\vec{x} = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}, \quad \vec{c} = [300 \quad 500], \quad \vec{b} = \begin{bmatrix} 4 \\ 12 \\ 18 \end{bmatrix}$$

Therefore the model used to solve, written using vectors and matrices looks like:

$$\begin{aligned} &\text{Maximize} && \vec{c} \vec{x} \\ &\text{subject to} && \mathbf{A} \vec{x} \leq \vec{b} \\ &&& \vec{x} \geq \vec{0} \end{aligned}$$

Or, if we substitute the matrices into the general notation:

$$\begin{aligned} &\text{Maximize} && [300 \quad 500] \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \\ &\text{subject to} && \begin{bmatrix} 1 & 0 \\ 0 & 2 \\ 3 & 2 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \leq \begin{bmatrix} 4 \\ 12 \\ 18 \end{bmatrix} \\ &&& \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \geq \begin{bmatrix} 0 \\ 0 \end{bmatrix} \end{aligned}$$

We are going to use this to build our spreadsheet as shown in Figure 6. While changing the color of the cells in the spreadsheet is not mandatory, it can be a helpful tool to visualize the different components within the spreadsheet. In Figure 6 all of the blue cells represent our parameters for the problem, \mathbf{A} , \vec{c} , and \vec{b} . The yellow cells represent our variable vector, \vec{x} , and, while it currently has a 0 placed for each cell, we are going to let Excel change these values when we use the software package.

Figure 6: The base Excel spreadsheet to solve Problem 1 using vectors and matrices.

The orange and non-colored cells are where we are going to do some calculations. In the orange cell, we use =MMULT(C3:D3,E7:E8) to multiply \vec{c} by \vec{x} . The function

3.3 – Objective Functions and Solving an LP

MMULT conducts the matrix multiplication for us automatically, and since we have a 1×2 row vector multiplied by a 2×1 column vector, the result is a 1×1 scalar value. Similarly, in the non-colored cells in column G, we multiply A by \vec{x} using =MMULT(C7:D9,E7:E8) in cell G7. When using MMULT, we must make sure that our dimensions match, or we will get a value error.

To solve this model in Excel, we are going to use Solver, which, after installing, can be found in the Data Tab on the top ribbon. The blank dialog window that you see should resemble what is found in Figure 7.

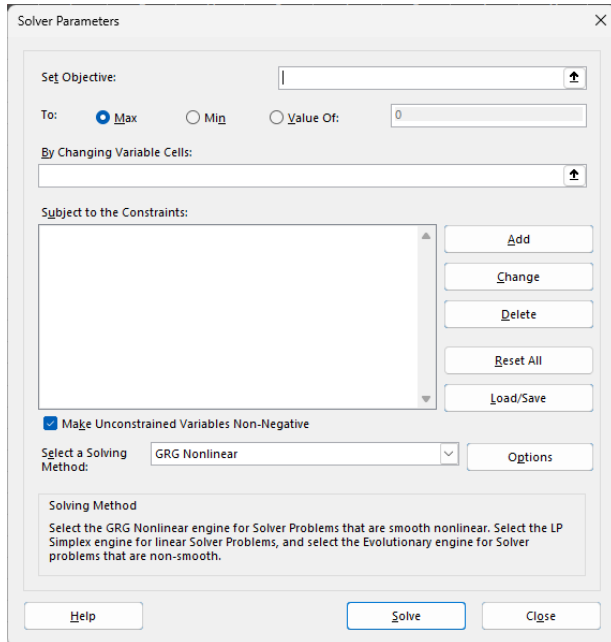


Figure 7: The blank Solver dialog window that appears when you open Excel Solver.

To fill in the dialog box to solve the LP, you start by clicking on the empty space next to “Set Objective”, then click on the corresponding cell in the spreadsheet. In this case, we click on the orange cell (G3) which contains the objective function value. The default objective in the Solver window is set to “Max”. For this problem, our goal is to maximize profit, so we can leave this as is. However, if we were trying to minimize, we would set the objective to “Min”. Then we select the variables. Click on the cell underneath “By Changing Variable Cells:”, then click and drag over the cells E7:E8, to select the yellow variable cells. Then we need to add constraints to Solver. Do this by clicking the “Add” button next to the constraint window. You will see a box as shown in Figure 8.

In the “Add Constraint” window, click on the space under “Cell Reference”, then click and drag across the cells that calculate $A\vec{x}$ (G7:G9). These cells are what Solver is going to compare against the “Constraint” cells to ensure the constraints are not violated. Next we need to ensure that we have the correct sign associated to our




Figure 8: The blank Add Constraint window for Excel Solver.

constraints. The default is \leq , representing \leq . In our case all of our constraints fall into this category. If we had different signs on the constraints, we would need to add them individually to account for them. Next, select the cell under the words “Constraint”, then click and drag over the cells of \vec{b} (I7:I9). Finally, click “OK”. You should be redirected back to the main Solver dialog window. We are almost done, but we can’t forget about our nonnegativity constraints. Luckily, Solver’s default is to have the box next to “Make Unconstrained Variables Non-Negative”. Having this box checked provides the nonnegativity constraints. If it is NOT checked, then our variables could be negative. Finally, select “Simplex LP” as the Solving Method. The full filled out Solver dialog box is shown in Figure 9.

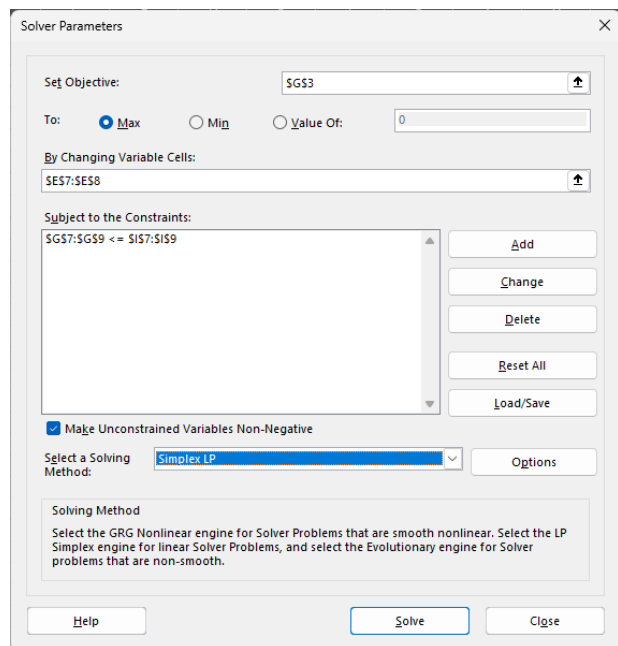


Figure 9: The fully filled out Solver dialog window for the spreadsheet show in Figure 6.

Now you are ready to hit “Solve” in the dialog window. Within a few seconds (for a small problem), Solver will indicate the results [1]. Typically it will indicate that it has found an optimal solution, as specified in the Solver Results dialog box shown in Figure 10. If the model has

3.3 – Objective Functions and Solving an LP

no feasible solutions or no optimal solution, the dialog box will indicate such [1]. The message “Solver could not find a feasible solution” means that there are no solutions that satisfy all the constraints [1]. The message “The Objective Cell values do not converge” means that Solver could not find a best solution because better solutions are always available [1].

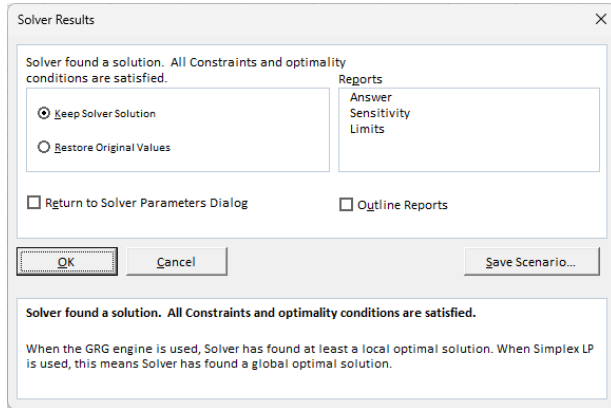


Figure 10: The Solver Results dialog box that indicates whether an optimal solution has been found.

After solving the model and clicking OK in the Solver Results dialog box, Solver replaces the original variable numbers with the optimal numbers, as shown in Figure 11 [1]. Therefore, the optimal solution is to produce two doors and six windows per week, just as we found graphically in the previous section. We can also see in the orange cell representing the total objective function value that the total profit Wyndor would make is \$3600.

	A	B	C	D	E	F	G	H	I	J
1	Wyndor Glass Co. Problem									
2							$\vec{c} \vec{x}$			
3		$\vec{c} =$	300	500			3600			
4										
5										
6			A		\vec{x}	$=$	$A \vec{x}$		\vec{b}	
7			1	0	2	$=$	2	\leq	4	
8			0	2	6	$=$	12	\leq	12	
9			3	2		$=$	18	\leq	18	
10										

Figure 11: The resulting Excel spreadsheet with the optimal variable values and objective function from Solver.

5 Interpreting the Solution

Although we’ve already been interpreting solutions informally, this section makes that process explicit. Interpreting the solution is the final step of the modeling process and is just as important as solving the math. In the previous two sections we have solved the LP, but when we report the result, we need to make sure we interpret the result. In the Wyndor example, providing $(x_1, x_2) = (2, 6)$ is not a sufficient answer because it doesn’t make sense without having the full variable definitions and a base

understanding of the solving technique used. To communicate the solution clearly, we must translate it into plain language with proper context:

To maximize profit, Wyndor should produce 2 doors and 6 windows per week.

We can go one step further and include the profit that is expected and the resources used. So a complete real-world recommendation may read:

To maximize profit, Wyndor should produce 2 doors and 6 windows per week. This will result in a weekly profit of \$3600 and uses all the time available in Plants 2 and 3.

This allows the decision maker to understand the solution and how the available resources are used in only two sentences.

Beyond interpretation, we must also critically evaluate whether the solution is appropriate and realistic. As modelers, we should consider such questions as:

- Does this solution make sense in context?
- Are any variables fractions when they must be whole? Does rounding provide an approximate estimate, or does it change the feasibility of the solution?
- Does it use all resources? Would we expect all resources to be used?
- Is this solution sensitive to changes in the model?

We will look at more detailed analysis of the last question in the next reading, but thinking through these questions can help you as a modeler avoid unrealistic recommendations and build confidence in our analysis. A good model is not only technically correct; it is useful, responsible, and transparent. Interpretation ensures our model meets these criteria.

6 Role of the Objective Function in the Solution

What if our objective function changes? What if instead of trying to maximize profit, Wyndor Glass Co. wants to maximize labor or maximize the number of doors they produce? Would that same solution still be optimal? The next reading will examine what happens if we change the parameters within our defined problem using the same objective function. It is important to understand that the solution we found in our example problem is only valid for the objective to maximize profit. Let’s look at these other two options and see what happens to our objective function.

If Wyndor is trying to maximize labor use instead of profit, we need to redefine the coefficients in our objective function. For example, suppose each door requires

3.3 – Objective Functions and Solving an LP

25 hours of labor to produce, while each window requires 10 hours. Then the objective function would be:

$$\text{Maximize } 25x_1 + 10x_2$$

Or if Wyndor is trying to maximize the number of doors they produce, then the objective function looks like:

$$\text{Maximize } x_1$$

If we plot these alternative objective functions at the solution we previously found to maximize profit, as shown in Figure 12, we can see how the optimal solution changes. At the point (2, 6), the red line representing the objective to maximize available labor cuts through the feasible region, indicating that (2, 6) is no longer optimal under this new goal. A different point, where more labor hours are consumed, would yield a higher value for this objective.

We also see that the objective function to maximize the number of doors is a horizontal line across the top of the feasible region. In this case, there are infinitely many solutions along the top edge that achieve the same maximum number of doors. The point (2, 6) still satisfies this objective, but it is not unique, any point along that edge is equally optimal if our goal is to maximize door production.

maximizing profit might shift its objective after securing a contract, prompting a reevaluation of the model and its recommendations.

References

- [1] Frederick Hillier et al. *MA103 Mathematical Modeling: Introduction to Management, Science, & Business Analytics with Connect*. McGraw-Hill, 2024.
- [2] Wayne Winston and Munirpallam Venkataramanan. *Introduction to Mathematical Programming*. Brooks/Cole, 2003.




Figure 12: Different objective functions plotted at the optimal solution for maximized profit. The gray dashed line is the objective function representing maximized profit. The red dashed line is the objective function to maximize labor use, and the blue dashed line is the objective function to maximize the number of doors produced.

What is important to note here is that it’s essential to clearly define what you are trying to optimize. Solving an LP is about finding the best solution for a specific goal, and that goal may change depending on updated priorities. For example, a company originally focused on