1 Introduction

In the previous reading we discussed what happens if our goal or our objective changes. We saw that this could potentially change our solution. We also asked the question, "is this solution sensitive to changes in the model?" To answer this question, we now explore what happens when the inputs to our existing model change, without altering the model's structure. This type of analysis is called *what-if analysis*. In what-if analysis, we aren't changing the objective or our variables and we aren't adding or subtracting constraints; we are instead focused on the model we have and how changes in the parameters affect our optimal solution.

2 What-if Analysis and Sensitivity

What-if analysis addresses questions about what would happen to the optimal solution if we made different assumptions [1], and specifically if we change assumptions regarding the parameters within our model. Typically, most parameters in an LP are estimates of quantities that we either don't have exact values for or are from some predictive model [1]. The goal of what-if analysis is to identify parameter sensitivity, or the range of values that a parameter can hold and still maintain the current optimal solution. What-if analysis is conducted as a part of the Interpret step of our modeling process. It provides more in-depth analysis of the solution and helps identify when the recommendation provided will no longer be useful [1].

2.1 Changes to the Constraints

When looking at conducting what-if analysis as it relates to the constraint parameters, first identify which constraints are binding (those that limit the current solution) versus non-binding (those that do not). Binding constraints for the problem will be those constraints that intersect to form the vertex where your optimal solution is. Non-binding constraints are everything else.

Definition 3.4.1 (Binding vs. Non-binding Constraints [2])
A constraint is **binding** if the left-hand side and the right-hand side of the constraint are equal when the optimal values of the variables are substituted into the constraint.

A constraint is **non-binding** if the left-hand side and the right-hand side of the constraint are unequal when the optimal values of the decision variables are substituted into the constraint.

In the context of the Wyndor problem we have been using over the last several readings, our binding constraints are:

$$2x_2 \leq 12 \quad \text{and} \quad 3x_1 + 2x_2 \leq 18$$

This is because at our optimal solution $(2, 6)$, our left-hand side equals the right-hand side of our constraints. This means that the binding constraints are the ones that physically form the edge, or boundary, of the feasible region at the optimal point. If you were to remove or relax one of these constraints, the feasible region would expand, and the optimal solution could shift. In contrast, non-binding constraints do not touch the optimal point and have no immediate effect on the solution if changed slightly. We can also see this graphically looking at Figure 1 with the optimal solution plotted at the intersection of the constraint $2x_2 \leq 12$ and $3x_1 + 2x_2 \leq 18$. Our non-binding constraints are $x_1 \leq 4$ and the nonnegativity constraints.

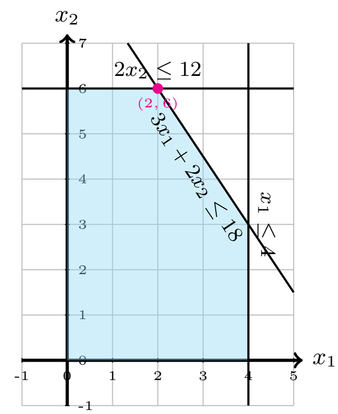

Figure 1: Identifying binding and non-binding constraints by graphing our optimal solution on the feasible region and identifying the intersecting constraints that make up our optimal solution.

2.1.1 Non-Binding Constraints

Let's start with assessing the sensitivity of the non-binding constraint $x_1 \leq 4$. Let's remember the different parameter values associated with this constraint. First we have the technical coefficients on the left-hand side of the equation representing the hours it takes to produce a product at Plant 1, 1 and 0, for x_1 and x_2 respectively. We also have the right-hand side value of 4 which is the total time in hours available at Plant 1.

We'll start examining the sensitivity of the the right-hand side value. When we conduct what-if analysis, we are going to analyze the sensitivity of the parameter as if it is the only parameter that will change. So, when we ask, "what if Plant 1 increases or decreases its total available hours," we are going to hold everything else in place. Now refer back to Figure 1. If we increase the number of hours available at Plant 1, will that change our solution? Increasing the number of available hours shifts the vertical line of my feasible region right, which does not affect my optimal solution. What about if we decrease the hours available at Plant 1? What if, for example, there are only 2 hours available at Plant 1?

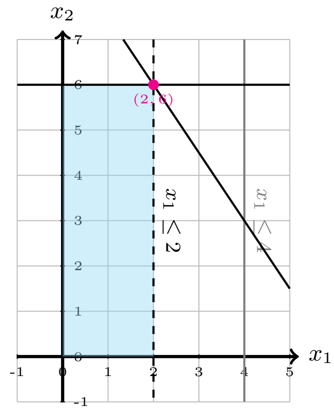

Figure 2: Feasible region created if Plant 1 time constraint is reduced to 2 hours.

Figure 2 plots this new constraint and updates the feasible region. We see that my optimal solution remains the same if $x_1 \le 2$. However, at our optimal solution, $x_1 \le 2$ becomes a binding constraint. Now if we continue to shrink the number of hours available at Plant 1, my optimal solution is going to have to change. So the range for my right-hand side parameter is $b_1 \in [2, \infty)$. What then can we say about the sensitivity of b_1, or the hours available at Plant 1?

Interpretation: The available hours at Plant 1 are sensitive to decreasing. If the hours available at Plant 1 decrease below 2 hours per week, then our recommendation will change.

Before we move on, remember what the technological coefficient matrix, $\boldsymbol{A}$, is for the current problem:

$$\boldsymbol{A} = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \\ a_{31} & a_{32} \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 2 \\ 3 & 2 \end{bmatrix}$$

Now let's examine the technological coefficients in the non-binding constraint $x_1 \le 4$, which correspond to the first row of the matrix $\boldsymbol{A}$. We'll start with $a_{11} = 1$, or that it takes one hour at Plant 1 to produce a door. Keeping the total time available at Plant 1 fixed at 4 hours, what happens to the feasible region if this parameter changes?

Figure 3 shows how the feasible region changes as the time required to produce a door at Plant 1 increases or decreases. If the production time increases to 2 hours per door, the optimal solution $(2, 6)$ still holds. However, once the production time exceeds 2 hours per door, $(2, 6)$ is no longer optimal. The feasible region shifts, and a new corner point becomes optimal. On the other hand, if the production time decreases below 1 hour per door, the optimal solution remains unchanged. Thus, the allowable range for this parameter is $a_{11} \in [0, 2]$.

Interpretation: The optimal solution is sensitive to increases in door production at Plant 1. If production time for doors at Plant 1 increases above two hours per door, the recommendation to produce $(2, 6)$ is no longer valid.

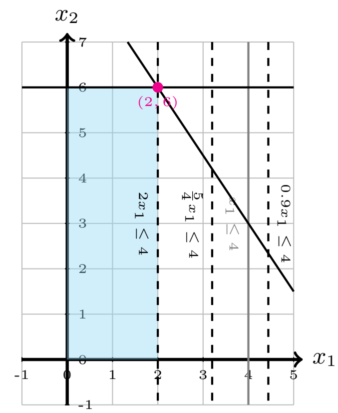

Figure 3: Feasible region created if production time to create one door at Plant 1 is increased to 2 hours, maintaining the current optimal solution. Other dark dashed lines show the constraint at different production rates.

Next, consider the other coefficient in the constraint: $a_{12} = 0$, or that manufacturing a window requires no time needed at Plant 1. But what if that changes? Suppose that windows begin requiring production time at Plant 1. How much production time can be added before our optimal solution changes?

Figure 4 shows the resulting shifts in the feasible region. With all other parameters held constant, we see that if window production time at Plant 1 increases beyond one-third of an hour per window, the optimal solution shifts away from $(2, 6)$. Therefore, the acceptable range is $a_{12} \in [0, \frac{1}{3}]$.

Interpretation: The solution $(2, 6)$ remains optimal only if windows require no more than one-third of an hour per unit at Plant 1. Any value beyond that changes the recommended production plan.

2.1.2 Binding Constraints

Now let's return to one of the original binding constraints in our model: $3x_1 + 2x_2 \le 18$. We'll conduct what-if analysis on this constraint, just as we did with the non-binding constraints earlier by adjusting one parameter at a time. But because this constraint is binding at the optimal solution, what do you think happens when we slightly change a single parameter? *The optimal solution shifts!*

This is a key insight: for our specific objective function and feasible region, any small change to a parameter in a binding constraint will move the location of the optimal point. So how do we think about sensitivity analysis when *any* change leads to a different solution?

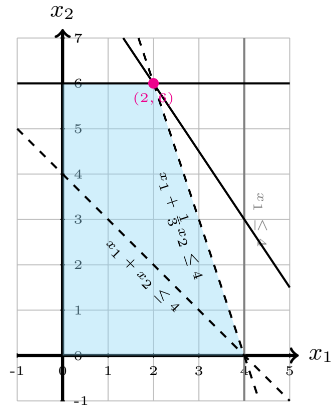

Figure 4: Feasible region created if windows require production time at Plant 1, and the current optimal solution is kept. The other dashed line shows what would happen if the required production time for one window at Plant 1 is increased to one hour.

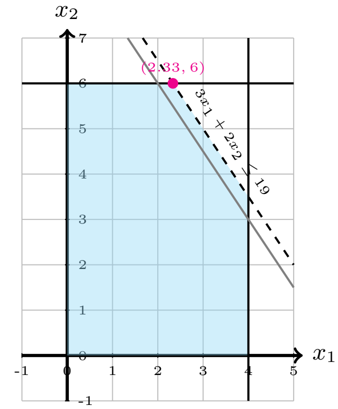

Figure 5: Identifying the new optimal solution for an increase of one hour in the total available hours at Plant 3.

When we conduct what-if analysis on a binding constraint, we are interested in determining the range of parameter values, whether on the left-hand or right-hand side, for which the current optimal solution remains both **feasible** and **binding**.

Let's start as we did before looking at the right-hand side $b_3 = 18$. When we change b_3 either up or down we are changing the total number of hours available at Plant 3 to produce windows and doors. If our production rates stay the same and we increase the number of hours, intuitively we would expect that the number of windows and doors we could produce will increase. Similarly, if the number of total hours available goes down, we would expect the number of windows and doors to also decrease. We can quantify this idea into something called a *shadow price.*

> **Definition 3.4.2** (**Shadow Price** [2])
> A shadow price is how much the objective function value is increased if the right-hand side value of a binding constraint is increased by 1, as long as it remains a binding constraint.

In Figure 5 we can see what happens to our feasible region when we increase the total number of hours available at Plant 3 by one hour. The two binding constraints remain binding, and our optimal solution shifts to $(\frac{7}{3}, 6)$. This shift increases our objective function value to \$3700. This makes our shadow price is \$100. Meaning that for every increase of one hour at Plant 3, we will increase profit made by \$100 and for every decrease of one hour at Plant 3 we should expect profit to decrease by \$100. This is true as long as $3x_1 + 2x_2 \leq 18$ and $2x_2 \leq 12$ remain our binding constraints.

So the next question is, for what values of b_3 would change our binding constraints? To answer this, we can examine where this constraint intersects the feasible region.

From the graphs, we observe that the constraint $3x_1 + 2x_2 \leq b_3$ fails to be binding when it shifts above the point $(4, 6)$. At that point, the constraints $2x_2 \leq 12$ and $x_1 \leq 4$ become the new binding constraints. On the other end, $2x_2 \leq 12$ ceases to be binding when $3x_1 + 2x_2 \leq b_3$ shifts below the point $(0, 6)$.

Figure 6 shows these boundary points projected onto the feasible region. To find the corresponding values of b_3, we substitute each point into the left-hand side of the constraint:

$$\text{At } (0,6): \quad 3(0) + 2(6) = 12 \quad \text{and,}$$

$$\text{At } (4,6): \quad 3(4) + 2(6) = 24$$

Therefore, for $b_3 \in [12, 24]$, the constraint $3x_1 + 2x_2 \leq b_3$ remains binding, along with $2x_2 \leq 12$, and the current optimal solution remains feasible.

In Figure 6, we also see that if b_3 decreases below 18, our current optimal solution at $(2, 6)$ becomes infeasible, and a different decision would be required. A clear way for a modeler to summarize the sensitivity of b_3 in an analysis or recommendation might be:

> **Modeler's Interpretation of Binding Constraint**
> Profit increases by \$100 for every additional hour available at Plant 3, up to a maximum of 24 hours, because we are able to make and sell 0.3 more doors. If hours available falls below 18 hours, the current production plan becomes infeasible and the model must be re-solved. Decision-makers can confidently plan for gains within the 18-24 hour range, but should reassess the model if time drops below 18 hours.

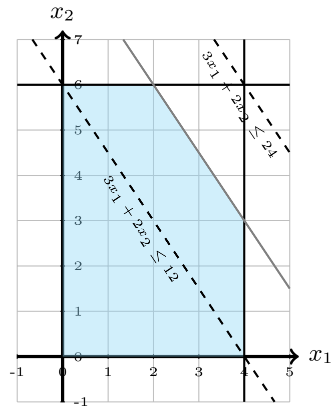

Figure 6: Identifying the range of values that b_3 can take without changing the binding constraints.

When we analyze the technological coefficients of a binding constraint, we are assessing under what conditions our solution remains feasible. If a shift in a technological coefficient causes our solution to become infeasible, our decision may not be very robust. We will talk more about robustness in the next reading. For now, let's focus on what changes in the technological coefficients will cause our solution to become infeasible.

What happens if the amount of time required to produce one door in Plant 3 is decreased to two hours from three? How does the feasible region shift? Figure 7 provides the updated feasible region with the current optimal solution plotted. We can see that the solution $(2, 6)$ remains in the feasible region, so it stays feasible.

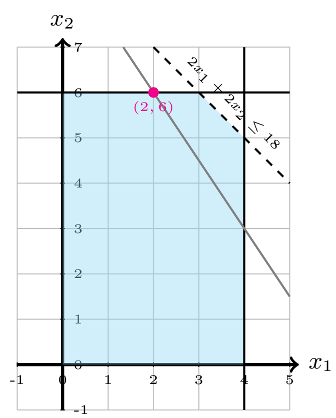

Figure 7: Feasible region produced if the time required to produce one door in Plant 3 is decreased from three to two hours.

However, if the amount of time required to produce one door in Plant 3 increases from three to four hours, as shown in Figure 8, the current optimal solution of $(2, 6)$ lies outside the feasible region, and is now infeasible.

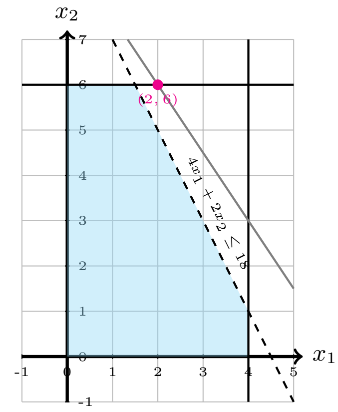

Figure 8: Feasible region produced if the time required to produce one door in Plant 3 is increased from three to four hours.

A similar trend is noticed if we increase or decrease the amount of time it takes to produce a window at Plant 3. If we increase the time it takes to produce a window, the current optimal solution becomes infeasible, and if we decrease the time the current optimal solution remains feasible.

Interpretation: Therefore, we can conclude that the solution we solved for remains valid as long as the time it takes to produce a door or window at Plant 3 remains the same or decreases. If the time is increased, the solution is no longer usable.

2.2 Changes to the Objective Function

We've explored possible changes to the parameters in our constraints, but what-if the constraints are the same and the cost coefficient of our objective function change? We saw that when we change the parameter of our constraints, the feasible region shifts. When we hold the constraint parameters constant, the feasible region will stay the same. However, as we saw in the previous reading the optimal solution may change when we change the parameters for our objective function. While that reading focused on full changes to the overall goal of the objective, we can see similarities to this idea when we conduct what-if analysis on the cost coefficients.

What-if analysis on the parameters in the objective function help inform the decision maker on the tolerance of the weight associated for the cost coefficients. For example what-if analysis can help answer the question, what if the profit margins on windows decrease? Is the decision to produce two doors and six windows still the best decision?

Let's answer this question. Remember the objective function and associated profit value at the optimal solution is:

$$300x_1 + 500x_2 = 3600$$

Figure 9 shows the projected objective function as the profit margin for windows decreases. As we decrease the profit margins for windows, the projected objective function line gets steeper. When $c_2 = \$200$, we can see that it is on top of the boundary constraint $3x_1 + 2x_2 \leq 18$. As c_2 drops below $200 we can see that it now bisects our feasible region, meaning that $(2, 6)$ is no longer an optimal solution. It remains a feasible solution, but it is no longer the best feasible solution.

Interpretation: For the current solution to remain optimal the the profit margins for windows cannot decrease below $200, holding the profit for doors constant.

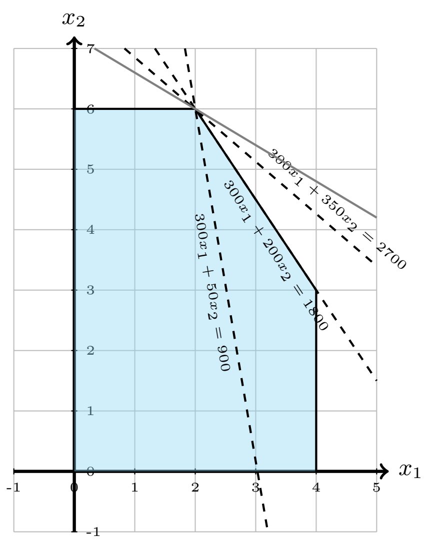

Figure 9: The objective function value projection lines plotted against the feasible region. Each dashed line represents a lower profit margin made on the sale of one window at the current optimal solution. The solid gray line represents the current objective function projection line associated with the optimal solution $300x_1+500x_2 = 3600$

We can also ask what happens if the profit margins on windows increase? We do the same process we just did, but now examining where the optimal solution would change as the profit margin increases, as shown in Figure 10. We can see as we increase the profit margin, the slope of the projected objective function line becomes less steep, getting closer and closer to a slope of 0. How can we determine how much we can increase the profit margin for c_2 before our optimal solution is no longer optimal?

We know that $(2, 6)$ is still an optimal solution if the slope of the objective function line is 0 where the projected objective function line overlays the constraint $2x_2 = 12$. We can rewrite the objective function line in y-intercept form as:

$$x_2 = -\frac{c_1}{c_2}x_1 + \frac{P}{c_2} \tag{1}$$

where P is the objective function value representing profit.

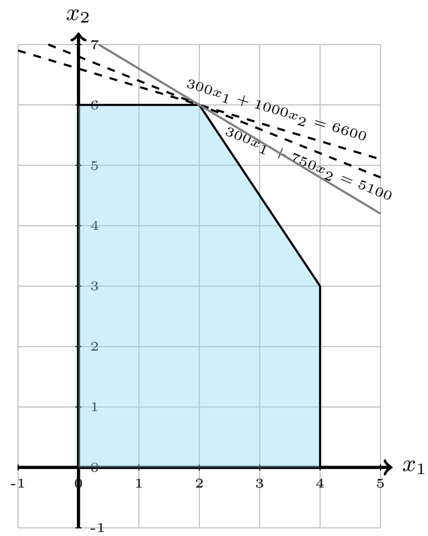

Figure 10: The objective function value projection lines plotted against the feasible region. Each dashed line represents a larger profit margin made on the sale of one window at the current optimal solution. The solid gray line represents the current objective function projection line associated with the optimal solution $300x_1+500x_2 = 3600$

Now we can see that our slope of the objective function line is the cost coefficient for the first variable, or the profit margin associated with making one door, divided by the cost coefficient for the second variable, or the profit margin associated with making one window. If we want to determine how much we can increase c_2 before our optimal solution changes we set the slope of the objective function equal to the slope of the constraint and solve for c_2.

$$-\frac{300}{c_2} = 0$$

However, in this instance, when we go to solve for c_2 we should notice that if we multiply both sides by c_2 we get $300 = 0$, which is not true. Let's think deeper here. What would we get if we set $c_2 = \$2000$?

$$-\frac{300}{2000} = -\frac{3}{20} = -0.15$$

What if we let $c_2 = \$10000$?

$$-\frac{300}{10000} = -\frac{3}{100} = -0.03$$

As c_2 increases, the slope approaches 0. So as c_2 gets infinitely larger the slope will get closer and closer to 0. However, this means it will never actually equal 0.

Interpretation: The current solution $(2, 6)$ remains optimal as long as the window profit c_2 stays above $200. There is no upper limit; increasing c_2 only strengthens the case for producing more windows under the current plan.

Problem 1 walks through how to algebraically solve for the range of the parameter c_1 to maintain the optimal solution of $(2,6)$.

Problem 3.4.1: Range of Optimality for c_1

Using the same context for the Wyndor Glass Co. problem. What is the range of profit made off of one door that maintains the optimal solution of producing two doors and six windows per week?

Solution:

We know the optimal solution $(2,6)$ is at the vertex of the constraints $3x_1 + 2x_2 \leq 18$ and $2x_2 \leq 12$. Therefore, the slope of the objective function line, must lie on or between the slopes of these two constraints.

First, let's put the constraints in y-intercept form to identify their slopes.

$$\begin{aligned} 3x_1 + 2x_2 &\leq 18 \\ 2x_2 &\leq -3x_1 + 18 \\ x_2 &\leq -\frac{3}{2}x_1 + 18 \end{aligned} \qquad \begin{aligned} 2x_2 &\leq 12 \\ x_2 &\leq 6 \end{aligned}$$

So the slope of our objective function must be on the interval $[-\frac{3}{2}, 0]$.

We know the slope of our objective function line from Equation 1 is $-\frac{c_1}{c_2}$. So we set up two equations and solve for c_1, with $c_2 = 500$ (held constant as part of what-if analysis, since we only change one parameter at a time).

$$\begin{aligned} -\frac{c_1}{500} &= -\frac{3}{2} \\ c_1 &= \frac{1500}{2} \\ c_1 &= 750 \end{aligned} \qquad \begin{aligned} -\frac{c_1}{500} &= 0 \\ c_1 &= 0 \end{aligned}$$

Interpretation: Therefore, the range of profit to be made on producing one door can be between \$0 and \$750 to keep the current solution optimal, if c_2 is held constant at \$500.

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.