Math 3272: Linear Programming[1] Mikhail Lavrov

Lecture 17: Sensitivity analysis

October 13, 2022 *Kennesaw State University*

1 Sensitivity analysis of the costs

1.1 Intuition

Let's begin with a linear program we've already solved much earlier in the semester. Below is the linear program, along with a diagram of its feasible region:

$$\begin{aligned} \underset{x,y\in\mathbb{R}}{\text{maximize}} \quad & 2x + 3y \\ \text{subject to} \quad & -x + y \le 3 \\ & x - 2y \le 2 \\ & x + y \le 7 \\ & x, y \ge 0 \end{aligned}$$

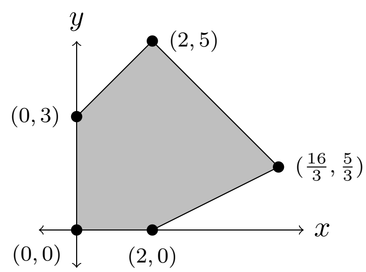

Today, we look at the following question: what happens to the optimal solution when we change the linear program slightly?

We will start by looking at one specific change in the objective function. Rather than maximizing $2x + 3y$, what happens when we maximize $(2 + \delta)x + 3y$ for some real number δ? (In an economic application, this could happen when the profitability of a product goes up or down.)

Because this is a small linear program with only five corner points, we can answer this question in the silliest way. At each corner point, we can compute the value of $(2 + \delta)x + 3y$, as a function of δ. We get the diagram on the left:

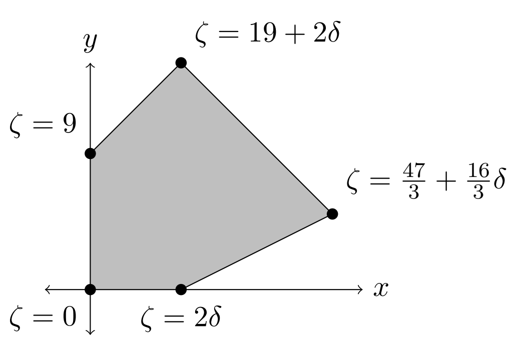

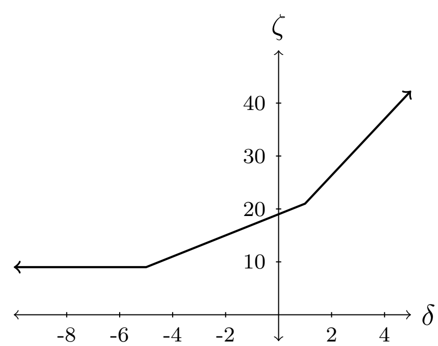

We know that the optimal solution is going to be the best of the corner points. Therefore, as a function of δ, the maximum value of $(2 + \delta)x + 3y$ is $\max\{0, 9, 19 + 2\delta, \frac{47}{3} + \frac{16}{3}\delta, 2\delta\}$. The diagram on the right plots the resulting function.

[1]This document comes from the Math 3272 course webpage: https://facultyweb.kennesaw.edu/mlavrov/courses/3272-fall-2022.php

What we see is a piecewise linear function whose slope increases from left to right. It has a segment where it is equal to $\zeta = 9$, a segment where it is equal to $\zeta = 19 + 2\delta$, and a segment where it acts as $\zeta = \frac{47}{3} + \frac{16}{3}\delta$. Those segments correspond to the ranges of δ where each of those corner points is optimal.

(There are no segments where the function is equal to $\zeta = 0$ or $\zeta = 2\delta$. That's because $(0, 0)$ and $(2, 0)$ will never be optimal as long as y has a positive coefficient in the objective function: $(0, 0)$ is always worse than $(0, 3)$ and $(2, 0)$ is always worse than $(2, 5)$.)

1.2 What we can compute

In larger examples, there will be too many corner points for this type of analysis to be reasonable. We do not want to go back to the naive way of solving linear programs, where we must find all of the corner points.

So let's ask a different question: what part of the information above can we compute from just knowing the optimal solution $(x, y) = (2, 5)$?

We can guess that for small values of δ, the optimal solution probably won't change. So if the objective value is $2 \cdot 2 + 3 \cdot 5 = 19$ right now, we expect to get $(2 + \delta) \cdot 2 + 3 \cdot 5 = 19 + 2\delta$, provided that δ is not too large.

We don't know what "not too large" is yet, and we don't know what happens for too-large δ. However, we can say a little bit about it. Even if the point $(2, 5)$ stops being optimal for large δ, it will still remain feasible. Therefore, for any value of δ, the maximum objective value possible will always be *at least* $19 + 2\delta$. (Our prediction is "pessimistic".)

Here is a rule that generalizes this logic:

Theorem 1. *Suppose that our linear program has optimal solution* $\mathbf{x}^*$ *with objective value* $\mathbf{c}^\mathsf{T}\mathbf{x}^* = \zeta^*$. *If we change the coefficient of* x_i *in the objective function from* c_i *to* $c_i + \delta$, *then the new objective value will be at least as good as* $\zeta^* + \delta x_i^*$. *(This is a lower bound when maximizing, and an upper bound when minimizing.)*

For small δ, *we can hope that the new objective value will be exactly* $\zeta^* + \delta x_i^*$.

As a special case, if a variable x_i is nonbasic, then we expect the objective value to stay the same when the cost of x_i changes by a small amount.

1.3 Ranging

Even a prediction for small values of δ can be more precise than this. From looking at the plot we got by comparing all corner points, we see that the prediction of $19 + 2\delta$ is exactly correct when $-5 \leq \delta \leq 1$. We can try to determine the interval where our prediction is correct: this is called **ranging**.

To do this, we'll need to look at the optimal dictionary, not just the optimal solution. The key idea is that to understand the effect of adding δx to the objective function, we can add δx to the equation for ζ in our optimal dictionary.

Of course, this is no longer a properly formed dictionary, because x is a basic variable and should never appear on the right-hand side. So we substitute the equation $x = 2 + \frac{1}{2}w_1 - \frac{1}{2}w_3$ into δx and simplify. This results in the dictionary on the right:

$$\begin{array}{rcl} \max \zeta &=& 19 - \frac{1}{2}w_1 - \frac{5}{2}w_3 + \delta x \\ \hline y &=& 5 - \frac{1}{2}w_1 - \frac{1}{2}w_3 \\ w_2 &=& 10 - \frac{3}{2}w_1 - \frac{1}{2}w_3 \\ x &=& 2 + \frac{1}{2}w_1 - \frac{1}{2}w_3 \end{array} \quad \rightsquigarrow \quad \begin{array}{rcl} \max \zeta &=& 19 + 2\delta + (-\frac{1}{2} + \frac{1}{2}\delta)w_1 + (-\frac{5}{2} - \frac{1}{2}\delta)w_3 \\ \hline y &=& 5 - \frac{1}{2}w_1 - \frac{1}{2}w_3 \\ w_2 &=& 10 - \frac{3}{2}w_1 - \frac{1}{2}w_3 \\ x &=& 2 + \frac{1}{2}w_1 - \frac{1}{2}w_3 \end{array}$$

We see that the new dictionary is optimal as long as $-\frac{1}{2} + \frac{1}{2}\delta \le 0$ and $-\frac{5}{2} - \frac{1}{2}\delta \le 0$. These inequalities simplify to $\delta \le 1$ and $\delta \ge -5$, so that's the range where $19 + 2\delta$ is guaranteed to predict the correct objective value.

As a shortcut for this calculation, we can compute ratios:

Theorem 2. *To determine the range where the prediction $\zeta^* + \delta x_i^*$ is guaranteed to be correct, when x_i is a basic variable, compute the values*

$$-\left(\frac{\textit{reduced cost of } x_j}{\textit{coefficient of } x_j \textit{ in the equation for } x_i}\right)$$

*for every nonbasic variable x_j. Negative values are lower bounds on δ and positive values are upper bounds on δ; the prediction is guaranteed to be correct if **all** the bounds hold.*

1.4 Nonbasic variables

The variable x was a bit special because it is a basic variable in our optimal solution. The case of nonbasic variables is different, and actually easier to handle.

First of all, if a variable is nonbasic in the optimal solution, its value in the optimal solution is 0, so our prediction says that the objective value will not change when the cost of the variable changes.

For how long will this be true? Well, adding δ to the cost of a basic variable is the same as adding δ to its reduced cost. So we have just one restriction: addition of δ can't change the sign of the reduced cost.

To see this in our example, we'll have to do something a bit unnatural and see what happens when we add a δw_1 term to the objective function. (Usually, slack variables don't show up at all in the objective function, but we'll make an exception here.) This changes the reduced cost: $\zeta = 19 - \frac{1}{2}w_1 + \frac{5}{2}w_3$ becomes $\zeta = 19 + (\delta - \frac{1}{2})w_1 + \frac{5}{2}w_3$. Therefore our solution remains optimal as long as $\delta - \frac{1}{2} \le 0$, or $\delta \le \frac{1}{2}$. (In particular, making δ an arbitrarily big negative number will never change anything.)

2 Sensitivity analysis of the constraints

2.1 The dual variables as "shadow costs"

What if we make a different kind of change: a change in one of our constraints? For example, suppose we change the constraint $-x + y \le 3$ in our linear program to $-x + y \le 3 + \delta$. What happens to the objective value?

This is much harder to understand by looking at a plot of the feasible region. The problem is that as we move the constraint $-x+y \leq 3$, not only do some corner points themselves begin to move, but some of them vanish, and new ones appear! This is not as prominent in a 2-dimensional problem, because every boundary of the feasible region just has two corner points on it. In a 3-dimensional feasible region, even the number of corner points on the boundary of the moving constraint can vary!

The trick is to look at the dual solution. Remember: if the primal program has constraints $A\mathbf{x} \leq \mathbf{b}$ or $A\mathbf{x} = \mathbf{b}$, the dual program has objective function $\mathbf{u}^\mathsf{T}\mathbf{b}$. This means that even though changing the vector $\mathbf{b}$ does weird unpredictable things to primal solutions, it leaves the dual solutions entirely unchanged—all that changes is their objective values.

As a result, the previous rules for changes in the objective function give us rules about changes in the constraints, provided that we use the dual solution instead.

Theorem 3. *Suppose that our linear program has an optimal solution with objective value ζ^*, and a dual optimal solution $\mathbf{u}^*$. If we change the bound in the j^{th} constraint from b_i to $b_i + \delta$, then the new objective value will be no better than $\zeta^* + \delta u_i^*$. (This is an upper bound when maximizing, and a lower bound when minimizing.)*

For small δ, we can hope that the new objective value will be exactly $\zeta^* + \delta u_i^*$.

In the example, we are looking at a maximization problem with $A\mathbf{x} \leq \mathbf{b}$ constraints, so we can use a simple method for finding the dual solution: it is the negative of the reduced costs of the slack variables. Looking at our optimal dictionary again, we see that w_1 has reduced cost $-\frac{1}{2}$, w_2 has reduced cost 0, and w_3 has reduced cost $-\frac{5}{2}$. Therefore the optimal dual solution is $\mathbf{u} = (\frac{1}{2}, 0, \frac{5}{2})$.

In particular, when the constraint $-x + y \leq 3$ changes to $-x + y \leq 3 + \delta$, we predict that the objective value will change to $19 + \frac{1}{2}\delta$. At least, we suspect that this is true when δ is not too large.

This theorem has one detail we did not mention: it is an "optimistic" prediction. Intuitively, for large changes in δ, one of two things will happen:

- If there is a very large positive change, say to $-x + y \leq 10000$, what can happen is that the constraint might stop being relevant. In this problem, even if the $-x + y \leq 3$ constraint is removed entirely, the optimal solution will just be $(0, 7)$ with an objective value of 21. Therefore at some point
- If there is a very large negative change, say to $-x + y \leq -10000$, what can happen is that the problem might become infeasible. In this case, we consider the maximum objective value to be $-\infty$, which is infinitely worse than the prediction $19 + 2\delta$.

The underlying reason that this prediction is "optimistic" while our previous predictions were "pessimistic" is that the dual program is the reverse of the primal: it minimizes when the primal maximizes, and vice versa.

The dual variables are also called **shadow costs** due to an economic application of this analysis.

Suppose that our objective value $2x + 3y$ measures the profit we make from a particular solution. Then the prediction "when the constraint $-x + y \leq 3$ changes to $-x + y \leq 3 + \delta$, we predict that the objective value will change to $19 + \frac{1}{2}\delta$" means that an increase by δ in this constraint is worth $\frac{1}{2}\delta$ dollars to us. In other words, we should be willing to pay up to 50 cents for each unit increase in this upper bound.

The term "shadow cost" refers to the idea that we're putting an inferred value on something that might not have a clear inherent value to us. For example, if the i^{th} constraint is given by the number of hours our employees can work, then the dual variable u_i tells us the price of labor: a limit on how much we should be willing to pay one of them to work an additional hour.

2.2 Ranging with slack variables

Just as when analyzing the costs, we can put a range on the values of δ for which this prediction is exact. To avoid getting bogged down in calculations, we will only do this for cases like our example: when we have a problem with constraints $A\mathbf{x} \leq \mathbf{b}$ which we put into equational form as $A\mathbf{x} + \mathbf{w} = \mathbf{b}$.

In a situation like this, we have a different way to do the calculation. In equational form, changing $-x + y + w_1 = 3$ to $-x + y + w_1 = 3 + \delta$ is equivalent to changing it to $-x + y + (w_1 - \delta) = 3$, and we can track the effects of this change by just replacing w_1 by $w_1 - \delta$ everywhere in the dictionary:

$$\begin{array}{r@{\;=\;}l}
\max \zeta & 19 - \frac{1}{2}(w_1 - \delta) - \frac{5}{2}w_3 \\ \hline
y & 5 - \frac{1}{2}(w_1 - \delta) - \frac{1}{2}w_3 \\
w_2 & 10 - \frac{3}{2}(w_1 - \delta) - \frac{1}{2}w_3 \\
x & 2 + \frac{1}{2}(w_1 - \delta) - \frac{1}{2}w_3
\end{array}
\quad \rightsquigarrow \quad
\begin{array}{r@{\;=\;}l}
\max \zeta & (19 + \frac{1}{2}\delta) - \frac{1}{2}w_1 - \frac{5}{2}w_3 \\ \hline
y & (5 + \frac{1}{2}\delta) - \frac{1}{2}w_1 - \frac{1}{2}w_3 \\
w_2 & (10 + \frac{3}{2}\delta - \frac{3}{2}w_1 - \frac{1}{2}w_3 \\
x & (2 - \frac{1}{2}\delta) + \frac{1}{2}w_1 - \frac{1}{2}w_3
\end{array}$$

This gives us the same conclusion: that the objective value changes to $19+\frac{1}{2}\delta$. However, we can also see what happens to the basic variables. What does this tell us about the limits on δ? Well, this prediction stops being valid if our basic solution stops being feasible: if any of the basic variables become negative. So we get the following constraints on δ:

$$5 + \frac{1}{2}\delta \geq 0, \qquad 10 + \frac{3}{2}\delta \geq 0, \qquad 2 - \frac{1}{2}\delta \geq 0.$$

This gives us a lower bound $\delta \geq -10$, another lower bound $\delta \geq -\frac{20}{3}$, and an upper bound $\delta \leq 4$. Therefore our prediction is valid for δ in the range $[-\frac{20}{3}, 4]$.

For a basic variable like w_2, the same method works, but we can work things out more intuitively. Seeing $w_2 = 10$ in the dictionary tells us that the constraint $x - 2y \leq 2$ is not even tight at the moment: $x - 2y$ is 10 lower than its upper bound. So changing the right-hand side by a small amount will not affect the objective value, and this stays true, provided that we don't reduce it by more than 10.