Math 482: Linear Programming[1] Mikhail Lavrov

Lecture 17: Sensitivity Analysis II

March 4, 2020 *University of Illinois at Urbana-Champaign*

1 What we know so far

Suppose we have a linear program (**P**) with a dual (**D**) as follows:

$$(\mathbf{P}) \begin{cases} \underset{\mathbf{x}\in\mathbb{R}^n}{\text{maximize}} & \mathbf{c}^\mathsf{T}\mathbf{x} \\ \text{subject to} & A\mathbf{x} \le \mathbf{b} \\ & \mathbf{x} \ge \mathbf{0} \end{cases} \qquad (\mathbf{D}) \begin{cases} \underset{\mathbf{u}\in\mathbb{R}^m}{\text{minimize}} & \mathbf{u}^\mathsf{T}\mathbf{b} \\ \text{subject to} & \mathbf{u}^\mathsf{T}A \ge \mathbf{c}^\mathsf{T} \\ & \mathbf{u} \ge \mathbf{0} \end{cases}$$

Let their optimal solutions be $\mathbf{x}^*$ and $\mathbf{u}^*$, respectively, with (by strong duality) equal objective value $z^* = \mathbf{c}^\mathsf{T}\mathbf{x}^* = {\mathbf{u}^*}^\mathsf{T}\mathbf{b}$. Then we know that, for some $\delta \in \mathbb{R}$ which is "sufficiently small" in absolute value,

- If we change c_i to $c_i + \delta$, we expect $\mathbf{x}^*$ to remain primal optimal, and so z^* will change by δx_i. This is guaranteed to be a lower bound.
- If we change b_i to $b_i + \delta$, we expect $\mathbf{u}^*$ to remain dual optimal, and so z^* will change by δu_i. This is guaranteed to be an upper bound.

(If we switch things up so that (**P**) is a minimization problem and (**D**) is a maximization problem, then the guarantees also switch: in that case, we'd get an upper bound in the first case and a lower bound in the second.)

We want to know: how large can δ get before things go wrong, and the prediction is no longer exact?

To answer this, we will consider the example below, with the optimal tableau given to its right:

$$\begin{array}{ll} \underset{\mathbf{x}\in\mathbb{R}^3}{\text{maximize}} & 5x_1 - x_2 + 7x_3 \\ \text{subject to} & x_1 + x_2 + x_3 \le 5 \\ & 3x_1 + 4x_3 \le 7 \\ & 2x_1 + x_2 + 3x_3 \le 5 \\ & x_1, x_2, x_3 \ge 0 \end{array}$$

	x_1	x_2	x_3	s_1	s_2	s_3	
s_1	0	2	0	1	-1	1	3
x_1	1	-4	0	0	3	-4	1
x_3	0	3	1	0	-2	3	1
$-z$	0	-2	0	0	-1	-1	-12

The optimal primal solution is $\mathbf{x}^* = (1, 0, 1)$, and the optimal dual solution is $\mathbf{u}^* = (0, 1, 1)$. These let us predict the rate of change in $z^* = 6$ when $\mathbf{c} = (5, -1, 7)$ or $\mathbf{b} = (5, 7, 5)$ are changed; we just need to know how far we can trust these predictions.

[1] This document comes from the Math 482 course webpage: https://faculty.math.illinois.edu/~mlavrov/courses/482-spring-2020.html

2 Easy cases

There are two cases in which a modification to the linear program is easy to analyze. They are exactly the cases in which the rules above predict no change at all.

2.1 Right-hand sides of slack constraints

Right now, $\mathbf{x}^*$ satisfies $x_1 + x_2 + x_3 = 2 < 5$. So the constraint $x_1 + x_2 + x_3 \leq 5$ is slack. What happens if we change the 5 to $5 + \delta$?

A quick way to answer that question is just to make the change in the tableau. Because s_1 is the pivot variable for the first row, we know that that row represents starting with the equation $x_1 + x_2 + x_3 + s_1 = 5$, and then adding or subtracting some other rows. So changing the 5 to $5 + \delta$ would change the final tableau to the following:

	x_1	x_2	x_3	s_1	s_2	s_3	
s_1	0	2	0	1	-1	1	$3+\delta$
x_1	1	-4	0	0	3	-4	1
x_3	0	3	1	0	-2	3	1
$-z$	0	-2	0	0	-1	-1	-12

This tableau still has the same optimal solution $\mathbf{x}^* = (1, 0, 1)$, and the same optimal value $z^* = 12$. It will also remain optimal (that is, dual feasible) for any value we choose. However, if $3 + \delta < 0$ (if $\delta < -3$), then this tableau will no longer be feasible. So if we predict that the objective value will not change, we are only confident in that prediction for $\delta \in [-3, \infty)$.

That was exactly what we expected based on looking at the dual solution, in which $u_1^* = 0$. We knew that for "sufficiently small" δ, the optimal value should change by $\delta u_1^* = 0 \cdot \delta = 0$, and now we know that "sufficiently small" means "in the range $[-3, \infty)$".

2.2 Costs of nonbasic variables

Right now, $x_2^* = 0$ in the optimal solution, so small changes to its coefficient $c_2 = -1$ shouldn't affect the optimal value. (In other words, the optimal objective value changes by $\delta x_2^* = 0 \cdot \delta = 0$.)

To analyze the magnitude of this change, we can consider the effect on the tableau. We never multiply the reduced cost row by any scalars, we only add and subtract from it. So adding δ to c_2 has the same effect in the optimal tableau: it changes to

	x_1	x_2	x_3	s_1	s_2	s_3	
s_1	0	2	0	1	-1	1	3
x_1	1	-4	0	0	3	-4	1
x_3	0	3	1	0	-2	3	1
$-z$	0	$-2+\delta$	0	0	-1	-1	-12

For any $\delta \in (-\infty, 2]$, the reduced cost $-2 + \delta$ will remain nonpositive, indicating that the tableau is still optimal. This means that we can change c_2 by any δ in that range, and the prediction that z^* doesn't change will be valid.

If $(\mathbf{P})$ were a minimization problem, we'd want nonnegative reduced costs, so the inequality would be reversed. To avoid making a mistake, make sure the interval you write down includes $\delta = 0$.

3 Hard cases

When the rules *do* predict a change in the optimal objective value at some rate, it is still possible to say when that rate of change will remain valid, but it's a bit harder.

3.1 Costs of basic variables

Suppose that we change c_1 to $c_1 + \delta$. As before, we know that we can make this change in the optimal tableau:

	x_1	x_2	x_3	s_1	s_2	s_3	
s_1	0	2	0	1	-1	1	3
x_1	1	-4	0	0	3	-4	1
x_3	0	3	1	0	-2	3	1
$-z$	δ	-2	0	0	-1	-1	-12

Now, however, the tableau is not fully row-reduced. We must subtract δ times x_1's row from the reduced cost row, getting

	x_1	x_2	x_3	s_1	s_2	s_3	
s_1	0	2	0	1	-1	1	3
x_1	1	-4	0	0	3	-4	1
x_3	0	3	1	0	-2	3	1
$-z$	0	$-2+4\delta$	0	0	$-1-3\delta$	$-1+4\delta$	$-12-\delta$

Note that the new optimal value here is $12 + \delta$, exactly as predicted by our rule: a change by $\delta x_1^* = \delta$. However, that prediction is only valid if the new tableau above is still optimal: if the reduced costs are all nonnegative. This requires

$$\begin{cases} -2 + 4\delta \leq 0 \\ -1 - 3\delta \leq 0 \\ -1 + 4\delta \leq 0 \end{cases} \implies \begin{cases} \delta \leq \frac{1}{2} \\ \delta \geq -\frac{1}{3} \\ \delta \leq \frac{1}{4} \end{cases}$$

All of these must be true, so we take the tightest upper and lower bound, and conclude that the allowable change is $\delta \in [-\frac{1}{3}, \frac{1}{4}]$.

In general, when modifying the coefficient c_i, we consider every nonbasic variable x_j (or s_j) and compute the ratio

$$\frac{\text{reduced cost of that variable}}{\text{coefficient of that variable in } x_i\text{'s row}}$$

in the optimal tableau. Positive ratios give upper bounds and negative ratios give lower bounds.

(If you think about it, the same thing works whether the program is a minimization or a maximization problem.)

3.2 Right-hand sides of tight constraints

Suppose that we change the right-hand side b_2: the right-hand side of the constraint $3x_1 - x_2 - x_3 \leq 2$, which is tight for the current optimal solution $\mathbf{x}^*$.

The problem is that there is no s_2 row in the tableau, so we can't just add δ to the right-hand side of the s_2 row and understand the effect of changing b_2 to $b_2 + \delta$. We could imagine making that change in the *initial* tableau, giving us the tableau on the left below. Then, in theory, we could repeat all the row-reduction steps that brought us to the final tableau (on the right), and see the effect of δ on the RHS:

	x_1	x_2	x_3	s_1	s_2	s_3	
s_1	1	1	1	1	0	0	5
s_2	3	0	4	0	1	0	$7+\delta$
s_3	2	1	3	0	0	1	5
$-z$	5	-1	7	0	0	0	0

$\rightsquigarrow$

	x_1	x_2	x_3	s_1	s_2	s_3	
s_1	0	2	0	1	-1	1	$3-\delta$
x_1	1	-4	0	0	3	-4	$1+3\delta$
x_3	0	3	1	0	-2	3	$1-2\delta$
$-z$	0	-2	0	0	-1	-1	$-12-\delta$

However, there is a better way to obtain the final tableau on the right. Notice that in the initial tableau, the following rule holds for every row: the number in s_2's column is equal to the coefficient of δ in the final column. This rule is not going to be affected by row reduction: we will be doing the same operations to both columns. So in the final tableau, we can compute the coefficients of δ on the RHS just by copying the numbers in s_2's column.

This prediction will be valid for as long as the resulting tableau remains optimal. The reduced costs stay the same, but the values of the basic variables don't: and if those values become negative, we lose primal feasibility, and need to update the tableau. So for the prediction to remain valid, we want to satisfy the conditions

$$\begin{cases} 3-\delta \geq 0 \\ 1+3\delta \geq 0 \\ 1-2\delta \geq 0 \end{cases} \implies \begin{cases} \delta \leq 3 \\ \delta \geq -\frac{1}{3} \\ \delta \leq \frac{1}{2} \end{cases}$$

therefore any change $\delta \in [-\frac{1}{3}, \frac{1}{2}]$ is fine.

In general, if we are considering changing b_i to $b_i + \delta$, we should compute the *negative* of the ratio of RHS entries over entries in s_i's column (in the optimal tableau). Positive results will give us upper bounds on δ, and negative results will give us lower bounds.

4 Summary

In general, we can summarize the rules above as a list of one or more limits on δ. The limits we get can be positive (giving us upper bounds on δ) or negative (giving us lower bounds).

- For the RHS of a slack constraint whose slack variable is s_i, the only limit on δ is the negative of the RHS of s_i's row.
- For the cost of a nonbasic variable x_i, the only limit on δ is the negative of x_i's reduced cost.
- For the cost of a basic variable x_i, the limits are the ratios $\frac{\text{reduced cost entry}}{\text{entry in } x_i\text{'s row}}$ computed for every column of a nonbasic variable.
- For the RHS of a tight constraint whose slack variable is s_i, the limits are the negative ratios $-\frac{\text{entry in RHS column}}{\text{entry in } s_i\text{'s column}}$ computed for every row (except the reduced cost row).