Math 482: Linear Programming[1]	Mikhail Lavrov
Lecture 16: Sensitivity Analysis I	
March 2, 2020	*University of Illinois at Urbana-Champaign*

1 Applications of the dual simplex method

1.1 Finding a basic feasible solution (inequalities)

Suppose that we want to find some $\mathbf{x} \in \mathbb{R}^n$ satisfying a system of inequalities in the form $A\mathbf{x} \le \mathbf{b}, \mathbf{x} \ge \mathbf{0}$.

This can be done with the dual simplex method, by optimizing an artificial objective function; one possible choice is to minimize $x_1 + x_2 + \cdots + x_n$. For example, to solve the system of equations

$$\begin{cases} x_1 + 3x_2 \le 5 \\ x_1 - \ \ x_2 \ge 1 \\ x_1, x_2 \ge 0 \end{cases} \iff \begin{cases} \ \ x_1 + 3x_2 + s_1 = 5 \\ -x_1 + \ \ x_2 + s_2 = -1 \\ x_1, x_2, s_1, s_2 \ge 0 \end{cases}$$

we can set up the dual feasible tableau

	x_1	x_2	s_1	s_2	
s_1	1	3	1	0	5
s_2	−1	1	0	1	−1
$-z^a$	1	1	0	0	0

and then perform the dual simplex method until we find an optimal solution (x_1, x_2). We don't really care about minimizing $x_1 + x_2$, but in particular the optimal solution will satisfy all the constraints.

We can use this a part of a two-phase simplex method:

1. First, optimize an artificial objective function using the dual simplex method. The resulting tableau is, in particular, primal feasible for any objective function.
2. Second, use the primal simplex method starting from the final tableau of the previous step.

The original objective function can either be kept around and row-reduced at every pivot step of the first phase, or introduced and row-reduced at the beginning of the second phase.

1.2 Finding a basic feasible solution (equations)

The method above relies on having slack variables s_1 and s_2 around to serve as basic variables, but this isn't really critical.

If we are given constraints in the form $A\mathbf{x} = \mathbf{b}, \mathbf{x} \ge \mathbf{0}$, we can do the same thing. First, row-reduce the system $A\mathbf{x} = \mathbf{b}$ (ignoring the condition $\mathbf{x} \ge \mathbf{0}$) until you arrive at a basic solution. Then, make

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

up an artificial objective function that will turn what you have into a dual feasible tableau. (For example, take the sum of the nonbasic variables.) Finally, use the dual simplex method to make the tableau primal feasible.

In an example: the system on the left row-reduces to the system on the right.

$$\begin{cases} x_1 + 2x_2 - 3x_3 + x_4 = 5 \\ 2x_1 + 4x_2 + x_3 - 2x_4 = 2 \\ x_1, x_2, x_3, x_4 \geq 0 \end{cases} \rightsquigarrow \begin{cases} x_1 + 2x_2 \qquad\quad - \frac{5}{7}x_4 = \frac{11}{7} \\ \qquad\qquad\quad x_3 - \frac{4}{7}x_4 = -\frac{8}{7} \\ x_1, x_2, x_3, x_4 \geq 0 \end{cases}$$

We ended up with basic variables (x_1, x_3), so we make minimizing $x_2 + x_4$ be the dual objective function, giving us the tableau

	x_1	x_2	x_3	x_4	
x_1	1	2	0	$-5/7$	$11/7$
x_3	0	0	1	$-4/7$	$-8/7$
$-z^a$	0	1	0	1	0

From here, we can apply the dual simplex method to find a primal feasible tableau.

1.3 Adding constraints

Consider the linear program we solved in yesterday's lecture (with optimal tableau on the right):

$$\begin{aligned} \underset{x,y\in\mathbb{R}}{\text{minimize}} \quad & x + y \\ \text{subject to} \quad & 2x + y \geq 6 \\ & 3x + y \geq 7 \\ & x + 2y \geq 9 \\ & x, y \geq 0 \end{aligned}$$

	x	y	s_1	s_2	s_3	
x	1	0	$-2/3$	0	$1/3$	1
s_2	0	0	$-5/3$	1	$1/3$	0
y	0	1	$1/3$	0	$-2/3$	4
$-z$	0	0	$1/3$	0	$1/3$	-5

Suppose that we realize at the last minute that the linear program should have had another constraint: the constraint $x \geq 2$, or $-x + s_4 = -2$. We don't have to solve the linear program from scratch. Instead, we can add it as an additional row of the tableau (and then row-reduce to eliminate the nonzero coefficient of the basic variable x):

	x	y	s_1	s_2	s_3	s_4	
x	1	0	$-2/3$	0	$1/3$	0	1
s_2	0	0	$-5/3$	1	$1/3$	0	0
y	0	1	$1/3$	0	$-2/3$	0	4
s_4	-1	0	0	0	0	1	-2
$-z$	0	0	$1/3$	0	$1/3$	0	-5

$\rightsquigarrow$

	x	y	s_1	s_2	s_3	s_4	
x	1	0	$-2/3$	0	$1/3$	0	1
s_2	0	0	$-5/3$	1	$1/3$	0	0
y	0	1	$1/3$	0	$-2/3$	0	4
s_4	0	0	$-2/3$	0	$1/3$	1	-1
$-z$	0	0	$1/3$	0	$1/3$	0	-5

Then, we can use the dual simplex method to restore feasibility.

Saying "we realize at the last minute that we forgot a constraint" is a contrived example, but there are many cases in which we want to solve a linear program, then add additional constraints. We'll see examples later on in this course.

2 Shadow costs

Suppose that we solve a linear program to optimality, and then make some small changes. Specifically, imagine starting with

$$\begin{aligned} \underset{\mathbf{x}\in\mathbb{R}^n}{\text{maximize}} \quad & \mathbf{c}^\mathsf{T}\mathbf{x} \\ \text{subject to} \quad & A\mathbf{x} \le \mathbf{b} \\ & \mathbf{x} \ge \mathbf{0} \end{aligned}$$

and making a small change either to the coefficients $\mathbf{c}$ in the objective function, or to the right-hand side $\mathbf{b}$. Can we predict what effect that change will have on the optimal solution?

Consider the example

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

The optimal solution is $(x, y) = (2, 5)$ with an objective value of $2x + 3y = 19$.

2.1 Changes to the objective function

Now we change "maximize $2x + 3y$" to "maximize $2.1x + 3y$".

1. **What do we expect to happen?**

   A small change like this will probably not affect optimality. The point $(2, 5)$ is likely to remain optimal. We go from an objective value of $2 \cdot 2 + 3 \cdot 5 = 19$ to an objective value of $2.1 \cdot 2 + 3 \cdot 5 = 19.2$.

2. **What can we guarantee will happen?**

   Without doing more work, we can't know for sure that the point $(2, 5)$ remains optimal. However, we know that it remains feasible, because we haven't changed the constraints at all. If another point becomes optimal, it's because a different point is better for the new objective function $2.1x + 3y$.

   So we can guarantee that after the change, the optimal objective value will increase from 19 to *at least* 19.2; the actual improvement might actually be greater.

The same happens if we decrease a coefficient slightly. In general, if we change $2x+3y$ to $(2+\delta)x+3y$, this changes the objective value at $(2, 5)$ from 19 to $19+2\delta$, and this is a lower bound for the optimal objective value. (We expect it to be correct when $|\delta|$ is small, but that is just a guess.)

This linear prediction is a lower bound when we're solving a maximization problem; for a minimization problem, we'll get an upper bound instead (because now, other points might be better at minimizing the objective value). In other words, whether we're minimizing or maximizing, this prediction is always a *pessimistic* prediction.

2.2 Changes to the RHS

Let's consider a different change: we change the constraint $x + y \le 7$ to $x + y \le 6$. This is no longer quite as easy to deal with, because our previous argument doesn't work: the point $(2, 5)$ is no longer feasible.

Instead, we can look at the dual program, which is

$$\begin{aligned} \underset{u,v,w\in\mathbb{R}}{\text{minimize}} \quad & 3u + 2v + 7w \\ \text{subject to} \quad & -u + v + w \ge 2 \\ & u - 2v + w \ge 3 \\ & u, v, w \ge 0 \end{aligned}$$

To understand the effect of changing 7 to 6, we need to know the dual optimal solution, which is $(u, v, w) = (\frac{1}{2}, 0, \frac{5}{2})$.

1. **What do we expect to happen?**

   Once again, for a small change, we expect the point $(\frac{1}{2}, 0, \frac{5}{2})$ to remain dual optimal. The new objective value is $3 \cdot \frac{1}{2} + 2 \cdot 0 + 6 \cdot \frac{5}{2} = \frac{33}{2} = 16.5$.

2. **What can we guarantee will happen?**

   It's possible that a different solution will become dual optimal after we make this change. If so, the new objective value will be smaller than the prediction above.

   Again, this varies depending on whether the primal is a minimization or a maximization problem. But the general pattern is that (from the point of view of the primal) this prediction is *optimistic*: the actual objective value might be worse than what we predict in this way.

If we change the primal constraint to $x+y \le 7\pm\delta$, the prediction sayst that our objective value will change to $19 \pm \frac{5}{2}\delta$. In other words, the value of the dual variable $w = \frac{5}{2}$ represents the sensitivity of the objective value to changes in the corresponding constraint.

Economists call the dual variables "shadow costs" for this reason: $\frac{5}{2}$ is the marginal cost of a change in the constraint $x + y \le 7$. If the function we were maximizing represents our profit, then we should be willing to pay up to $\frac{5}{2}\delta$ for a change of $+\delta$ in this constraint. Or, if someone were willing to pay us for a change of $-\delta$ in this constraint, we should ask for a price of at least $\frac{5}{2}\delta$.

2.3 Looking ahead

In both cases, we've made a prediction that's valid "for small δ". We haven't been specific about the word "small", because the prediction might have an arbitrarily small range in which it's valid. (In rare cases, the prediction will not be exact for any nonzero δ.)

In the next lecture, we'll refine our prediction to also give an interval in which it's valid. For example, we can say that if we change the objective function $2x + 3y$ to $(2 + \delta)x + 3y$, then the predicted objective value of $19 + 2\delta$ will be correct as long as $\delta \in [-5, 1]$. Outside this range, it is only an upper bound.