Math 482: Linear Programming[1] Mikhail Lavrov

Lecture 7: Pivoting Rules

February 5, 2020 *University of Illinois at Urbana-Champaign*

1 Degenerate tableaux

We say that a tableau is *degenerate* if some basic variables have value 0 in its basic feasible solution. For exampl, consider the problem

$$\begin{aligned} \underset{x_1,x_2\in\mathbb{R}}{\text{maximize}} \quad & x_1 + 4x_2 \\ \text{subject to} \quad & x_1 + x_2 \le 0 \\ & x_1 - 3x_2 \le 0 \\ & -2x_1 + x_2 \le 0 \\ & x_1, x_2 \ge 0. \end{aligned}$$

(As a side note, when we only have a bunch of ≤ 0 and ≥ 0 inequalities like this, there are two possible cases: either they cut off all escape routes from $(0,0)$ in any direction that maximizes $x_1 + 4x_2$, and $(0,0)$ is optimal, or they don't, and the LP is unbounded. The simplex method will be trying to determine which case we're in.)

After adding slack variables, here is the starting tableau, corresponding to the point $x_1 = x_2 = 0$:

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

We've already briefly seen why degenerate tableaux are a problem. In a degenerate tableau, we can pivot and not change the basic feasible solution. For example, say we pivot on x_1. Even if we rule out s_3 as a leaving variable because -2 is negative (and this is not yet entirely justified, because we *could* divide by -3 and get a feasible tableau in this case: the RHS remains 0) we still get a tie. The basic variables s_1 and s_2 both have ratios of $\frac{0}{1} = 0$, so they both have the smallest ratio, and either one can leave the basis.

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

Say we pick s_1 to leave the basis. We get the following tableau:

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

Pivot on x_2 next. Ruling out s_2 because $-4 < 0$, the ratios for x_1 and s_3 are $\frac{0}{1}$ and $\frac{0}{3}$, which is still a tie. If we choose s_3 to leave the basis, we get the following tableau:

	x_1	x_2	s_1	s_2	s_3	
x_1	1	0	$1/3$	0	$-1/3$	0
s_2	0	0	$5/3$	1	$4/3$	0
x_2	0	1	$2/3$	0	$1/3$	0
$-z$	0	0	-3	0	-1	0

The equation corresponding to $-z$'s row is $z = -3s_1 - s_3$, proving that $z = 0$ is optimal. But it's not clear how we could have gotten here on purpose. At each step, we were flailing around randomly. How do we direct our progress when all our ratios are 0?

2 Pivoting rules

If we're not careful, we'll encounter *cycling*: this is a phenomenon that happens when we loop forever through different representations of the same point. (For some examples of this happening, see p. 51 in **PS**, or p. 27 in **V**, or your homework assignment.) To avoid this, we need a *pivoting rule*: a strategy to make the decisions that the simplex method doesn't make for you.

There are two places where we still have freedom:

- When we pick an entering variable x_i, we choose one whose reduced cost has the appropriate sign: negative if we're minimizing, and positive if we're maximizing.

  Sometimes, there are multiple choices of entering variable that satisfy this condition.

- When we pick a leaving variable x_j, after we rule out the variables with a negative coefficient in x_i's column, we look at the ratios between the constants on the right-hand side and the coefficients in x_i's column. We choose the variable with the smallest such ratio.

  Sometimes, there's a tie, which is particularly significant when there are multiple ratios equal to 0.

Addressing the first bullet point but not the second, a common strategy is to pick an entering variable by the magnitude of reduced cost: highest reduced cost when maximizing, and most negative reduced cost when minimizing. This can lead to cycling if we choose leaving variables badly.

Otherwise, it's often good, though relies on unstated assumptions about units. Imagine you're solving a problem about buying rice and wheat, but you measure rice in tons and wheat in pounds. Then all rice-related variables will have much larger reduced costs, because a ton of something makes a much bigger difference than a pound. But this has nothing to do with whether rice or wheat is more profitable.

There are two complete pivoting strategies I want to mention briefly, though we won't discuss them too much.

- **Bland's pivoting rule.** Give the variables an order (such as $x_1, x_2, \dots, x_n, s_1, s_2, \dots, s_m$; we can pick any order, we just need to be consistent). Whenever there is a tie for entering or leaving variable, the variable earlier in the order wins.

  This rule prevents cycling. (We won't prove this; if you're curious, see Theorem 5.8.1 in **GM**, Theorem 2.9 in **PS**, or Theorem 3.3 in **V**.) Unfortunately, it's very slow: when cycling is not an issue, it usually takes more steps to get to the optimal solution.

- **Random pivoting rule.** Break ties randomly. Statistically, this avoids cycling forever, unless you're infinitely unlucky.

3 The lexicographic pivoting rule

The lexicographic pivoting rule is the interesting one.

As motivation for this rule: cycling is only a concern when many basic feasible solutions happen to describe the same exact point. If we randomly adjusted the right-hand sides of our constraints by tiny values, then those basic feasible solutions would move apart in different directions by tiny amounts, so we'd never encounter degenerate tableaux.

That's not quite what we want to do: if we just introduced random changes, then our final solution would also be slightly off from the correct solution. The lexicographic pivoting rule simulates making these tiny changes, in a way that we can forget about when we're done.

3.1 The basic method

Imagine very tiny numbers $\epsilon_1, \epsilon_2, \dots, \epsilon_m$ such that

$$1 \gg \epsilon_1 \gg \epsilon_2 \gg \epsilon_3 \gg \cdots \gg \epsilon_m > 0$$

where "$\gg$" means "this stays true up to factors of any number that will ever appear in our tableau". Then add ϵ_i to the right-hand side of the i^{th} constraint, making that constraint looser by a very tiny margin. This prevents ties between rows from *ever* occurring.

(Adding the ϵ's only helps us choose the leaving variable. For the entering variable, we can choose the one with the highest/lowest reduced cost, which is a pretty good heuristic and means that this method won't be as slow as Bland's rule.)

Our example problem becomes

$$\begin{aligned}
\underset{x_1, x_2 \in \mathbb{R}}{\text{maximize}} \quad & x_1 + 4x_2 \\
\text{subject to} \quad & x_1 + \ \ x_2 \le \epsilon_1 \\
& x_1 - 3x_2 \le \epsilon_2 \\
& -2x_1 + \ \ x_2 \le \epsilon_3 \\
& x_1, x_2 \ge 0.
\end{aligned}$$

There are two common ways to write the tableau to incorporate the ϵ's:

	x_1	x_2	s_1	s_2	s_3	
s_1	1	1	1	0	0	ϵ_1
s_2	1	-3	0	1	0	ϵ_2
s_3	-2	1	0	0	1	ϵ_3
$-z$	1	4	0	0	0	0

	x_1	x_2	s_1	s_2	s_3	1	ϵ_1	ϵ_2	ϵ_3
s_1	1	1	1	0	0	0	1	0	0
s_2	1	-3	0	1	0	0	0	1	0
s_3	-2	1	0	0	1	0	0	0	1
$-z$	1	4	0	0	0	0	0	0	0

We can simply write out the expression with $\epsilon_1, \dots, \epsilon_m$ in the rightmost column, or we can subdivide that column into the coefficient of 1, coefficient of ϵ_1, coefficient of ϵ_2, and so on.

The highest reduced cost rule for the entering variable would tell us to bring x_2 into the basis. The ratios for the possible leaving variables are $\frac{\epsilon_1}{1}$, $\frac{\epsilon_2}{-3}$, and $\frac{\epsilon_3}{1}$. Note that the second of these ratios is negative, so we can't pick it! (Before we introduced the ϵ's, it would have been tied for 0.)

We have $\frac{\epsilon_3}{1} < \frac{\epsilon_1}{1}$, so s_3 leaves the basis. We get:

	x_1	x_2	s_1	s_2	s_3	
s_1	3	0	1	0	-1	$\epsilon_1 - \epsilon_3$
s_2	-5	0	0	1	3	$\epsilon_2 + 3\epsilon_3$
x_2	-2	1	0	0	1	ϵ_3
$-z$	9	0	0	0	-4	$-4\epsilon_3$

	x_1	x_2	s_1	s_2	s_3	1	ϵ_1	ϵ_2	ϵ_3
s_1	3	0	1	0	-1	0	1	0	-1
s_2	-5	0	0	1	3	0	0	1	3
x_2	-2	1	0	0	1	0	0	0	1
$-z$	9	0	0	0	-4	0	0	0	-4

There is only one positive reduced cost, so we must pivot on x_1. Just for practice, we can compare the right-hand sides: we have

$$\epsilon_1 - \epsilon_3 > \epsilon_2 + 3\epsilon_3 > \epsilon_3.$$

But the comparison is irrelevant, because of the ratios $\frac{\epsilon_1 - \epsilon_3}{3}$, $\frac{\epsilon_2 + 3\epsilon_3}{-5}$, and $\frac{\epsilon_3}{-2}$, only the first ratio is positive. Therefore x_1 must replace s_1 in the basis, giving the tableau below:

	x_1	x_2	s_1	s_2	s_3	
x_1	1	0	$1/3$	0	$-1/3$	$\frac{\epsilon_1 - \epsilon_3}{3}$
s_2	0	0	$5/3$	1	$4/3$	$\frac{5\epsilon_1 + 3\epsilon_2 + 4\epsilon_3}{3}$
x_2	0	1	$2/3$	0	$1/3$	$\frac{2\epsilon_1 + \epsilon_3}{3}$
$-z$	0	0	-3	0	-1	$-3\epsilon_1 - \epsilon_3$

	x_1	x_2	s_1	s_2	s_3	1	ϵ_1	ϵ_2	ϵ_3
x_1	1	0	$1/3$	0	$-1/3$	0	$1/3$	0	$-1/3$
s_2	0	0	$5/3$	1	$4/3$	0	$5/3$	1	$4/3$
x_2	0	1	$2/3$	0	$1/3$	0	$2/3$	0	$1/3$
$-z$	0	0	-3	0	-1	0	-3	0	-1

All of our reduced costs are negative, so we've maximized z. The largest possible value is $3\epsilon_1 + \epsilon_3$, which rounds off to 0 for the original problem.

Lexicographic pivoting has the nice properties that:

- If combined with the highest reduced cost rule (or lowest reduced cost, if minimizing) then there are no choices to make: one variable is always better than the others. There are no ties.
- When $\epsilon_1, \epsilon_2, \dots$ are taken into account, the objective value is always improving, so we are always making progress and can't possibly end up looping back.

Therefore lexicographic pivoting prevents cycling.

3.2 Shortcuts

You may have noticed that the coefficients of $\epsilon_1, \epsilon_2, \epsilon_3$ always match the coefficients of s_1, s_2, s_3. This is a convenient shortcut when using slack variables: when two columns, such as the ϵ_1 and s_1 column, start out identical, they will remain identical no matter how we pivot. In such a case, we don't even have to write out the ϵ_i columns, and just use the slack variable coefficients to break ties.

In general, whether we use slack variables or not, we can avoid writing out $\epsilon_1, \epsilon_2, \dots, \epsilon_m$ as follows. Mark the columns which formed our initial basis; say these columns are $C_1, C_2, \dots, C_m$. Whenever you get a tie between some basic variables by taking the ratio

$$\frac{\text{number on the RHS}}{\text{number in entering variable's column}},$$

break that tie between those variables by computing the ratio

$$\frac{\text{number in column } C_1}{\text{number in entering variable's column}}.$$

If that's also a tie, use the numbers in column C_2 as a tiebreaker, then the numbers in column C_3, and so on. This will be equivalent to using the ϵ's.

Whether or not we use this shortcut, we can definitely skip writing out the ϵ's in the objective value: the bottom right cell of the tableau. Those coefficients aren't going to affect what happens in the rest of the matrix, and we're going to round them away at the end. (In the example above, it never matters that z goes from 0 to $4\epsilon_3$ to $3\epsilon_1 + \epsilon_3$.)