Math 3272: Linear Programming¹

Mikhail Lavrov

Lecture 5: The objective function

August 30, 2022

Kennesaw State University

1 Introducing objective functions

This time, we'll begin by posing a complete problem with an objective function to maximize:

Problem 1. A farmer splits 9 acres of $land^2$ between growing cotton $(x_1 \text{ acres})$, corn $(x_2 \text{ acres})$, and soy $(x_3 \text{ acres})$. Cotton is not regulated, but federal regulations require a balance in food crops sold: at most 75% of the total amount can be a single crop. However, an additional acre's worth of food can be sold in-state, where this regulation does not apply.

In units of hundreds of dollars, the farmer's profit is 2 per acre of cotton, 3 per acre of corn, and 4 per acre of soy. How can the farmer maximize profit?

To express the fictional federal regulations in this problem as linear constraints, we require that each of x_2, x_3 is at most three times the other, plus 1: $x_2 \leq 3x_3 + 1$ and $x_3 \leq 3x_2 + 1$. All the quantities in this problem must be nonnegative. This gives us the linear program we see below on the left; on the right, we've added slack variables to put it in equational form.

The first thing to realize is that when the equations hold, the objective function has many equivalent forms. Since $x_1 + x_2 + x_3 = 9$, for example, maximizing $2x_1 + 3x_2 + 4x_3$ is equivalent to maximizing $2x_1 + 3x_2 + 4x_3 + (x_1 + x_2 + x_3 - 9)$ or $3x_1 + 4x_2 + 5x_3 - 9$: if we maximize that objective function instead, we get the same solution.

We will give the expression $2x_1 + 3x_2 + 4x_3$ a name: we'll call it ζ (zeta).³ Writing down the equation $\zeta = 2x_1 + 3x_2 + 4x_3$ makes our lives somewhat easier: we now have 4 equations in 6 variables ζ , x_1 , x_2 , x_3 , w_1 , w_2 , and we are simply maximizing one of the variables.

This particular problem conveniently starts out row-reduced with x_1, w_1, w_2 as the basic variables; we can easily solve for them in terms of the non-basic variables x_2, x_3 . Out of our many representations for ζ , it is convenient to pick one that's *also* in terms of x_2, x_3 . Just subtract twice $(x_1 + x_2 + x_3 - 9)$ to get $\zeta = x_2 + 2x_3 + 18$.

¹This document comes from an archive of the Math 3272 course webpage: http://misha.fish/archive/ 3272-fall-2022

 $^{^2\}mathrm{An}$ unrealistically small amount to keep our numbers low.

³Many sources use z instead, and you may feel free to write z instead of ζ .

2 The dictionary

The **dictionary** is a representation of a linear program closely related to the systems of equations we were writing in the previous lecture. It has an additional equation for the objective value, which we traditionally separate as follows:

$$\frac{\max \zeta = 18 + x_2 + 2x_3}{x_1 = 9 - x_2 - x_3}$$
$$w_1 = 1 - x_2 + 3x_3$$
$$w_2 = 1 + 3x_2 - x_3$$

It is helpful to include "max" or "min" in the top left corner of the dictionary, to remind ourselves that we're maximizing or minimizing. The simplex method treats the two cases differently.

Each dictionary corresponds to a basic solution whose parameters we can read off from the column immediately after =. We have $x_1 = 9$, $w_1 = 1$, and $w_2 = 1$, while the nonbasic variables x_2, x_3 are set to 0; the objective value of this solution is $\zeta = 18$.

With the possible exception of ζ , all the numbers in this column should be nonnegative if we are looking at a basic *feasible* solution. If the dictionary has this property, we call it a **feasible dictionary** and, for the time being, we will not consider any other kind of dictionary.

There are as many feasible dictionaries as there are basic feasible solutions.⁴ The simplex method operates by moving from dictionary to dictionary until we arrive at one that gives us the optimal solution. The method of moving from dictionary to dictionary is the same as in the previous lecture; today, we will see how the objective value fits in.

3 Using the simplex method

3.1 The first pivoting step

Let's begin by bringing x_3 into the basis. This is an arbitrary choice for now, but we'll see what happens when we do this, and think about how we can make this choice more intelligently.

If x_3 is our entering variable, then we need to choose a leaving variable. This is not new; however, to be clear, we must never choose ζ to be our leaving variable. We will always keep ζ in the top left corner of our dictionary, so that we always know what the objective value is!

We choose the leaving variable in two steps:

- Out of x_1, w_1, w_2 , we reject w_1 immediately: it has a positive coefficient of x_3 in our dictionary, and we want a negative coefficient.
- x_1 's current value is 9 and it decreases at a rate of 1 as x_3 increases. Meanwhile, w_2 's current value is 1 and it decreases at a rate of 1 as x_3 increases. We choose the variable with the lowest ratio $\frac{\text{current value}}{\text{rate of decrease}}$; $\frac{9}{1} > \frac{1}{1}$, so we choose w_2 .

⁴Note for the future: sometimes, unfortunately, there are slightly more—multiple feasible dictionaries for the same basic feasible solution! Don't worry about this for now.

To bring x_3 into the basis and w_2 out of the basis, we begin by solving w_2 's equation for x_3 , getting $x_3 = 1 + 3x_2 - w_2$. Now, we rewrite ζ , x_1 , and w_1 in terms of w_2 rather than x_3 , by substituting $1 + 3x_2 - w_2$ in place of x_3 wherever it occurs:

$\max \zeta = 18 + x_2 + 2(1 + 3x_2 - w_2)$	$\max\zeta = 20 + 7x_2 - 2w_2$
$x_1 = 9 - x_2 - (1 + 3x_2 - w_2)$	$x_1 = 8 - 4x_2 + w_2$
$w_1 = 1 - x_2 + 3(1 + 3x_2 - w_2)$	 $w_1 = 4 + 8x_2 - 3w_2$
$x_3 = 1 + 3x_2 - w_2$	$x_3 = 1 + 3x_2 - w_2$

We've obtained our new dictionary! The new values of our variables are $(x_1, x_2, x_3, w_1, w_2) = (8, 0, 1, 4, 0)$, and the objective value ζ has increased to 20.

3.2 How do we make progress?

Since ζ has gone from 18 to 20, apparently we've done something right. But it was a complete accident! Let's figure out what we did right so we can keep doing it.

How much did ζ go up? The change from 18 to 20 is the product of two things: the 2 which was the coefficient of x_3 (our entering variable) in the old equation for ζ , and the 1 which is the new value of x_3 .

Because we will always choose our leaving variable to get a feasible dictionary, the new value of our entering variable will always be positive. However, the coefficient of x_3 in the old equation for ζ could have been anything.

We conclude that:

• If we want ζ to increase, we should choose an entering variable with a *positive* coefficient in ζ 's equation.

That way, we multiply two positive numbers to compute the change in ζ .

• If we want ζ to decrease, we should choose an entering variable with a *negative* coefficient in ζ 's equation.

That way, we multiply a negative number by a positive number to compute the change in ζ .

There is an official term for the coefficient of x_3 in ζ 's equation; it is called the **reduced cost** of x_3 . Let me explain this term so it's easier to remember:

- The word "cost" actually comes from the fact that in the original equation $\zeta = 2x_1+3x_2+4x_3$, the 4 represented the cost at which the farmer can sell soy. Economic problems like this one are an important application of linear programs! The word cost has stuck around to be used in problems where the objective value has nothing to do with money.
- The word "reduced" has nothing to do with the fact that the number has gone down from 4 to 2; it could be larger or smaller. The term should properly be "row-reduced cost". After we've row-reduced our system of equations, the cost has changed: the row-reduced cost is its new value.

To summarize our rule: when we bring x_i into the basis, the change in ζ has the same sign as the reduced cost of x_i . Look for a positive reduced cost when maximizing ζ , and a negative reduced cost when maximizing.

3.3 One more pivot step

In our first pivot step, both reduced costs were positive, so any choice of entering variable would have been fine. In the next step, we are choosing between x_2 and w_2 . Only x_2 has a positive reduced cost, so it is the only valid choice of entering variable. (This makes sense: since w_2 just left the basis, pivoting around w_2 will return us to where we were previously, undoing our progress!)

In x_2 's column, only x_1 has a negative coefficient, so it is our only valid choice of leaving variable: we don't even have to compare ratios.

Solving the equation $x_1 = 8 - 4x_2 + w_2$ for x_2 , we get $x_2 = 2 - \frac{1}{4}x_1 + \frac{1}{4}w_2$. Now we are ready to substitute this in for x_2 in all the other rows of our dictionary:

$\max \zeta = 20 + 7(2 - \frac{1}{4}x_1 + \frac{1}{4}w_2) - 2w_2$		$\max \zeta = 34 - \frac{7}{4}x_1 - \frac{1}{4}w_2$
$x_2 = 2 - \frac{1}{4}x_1 + \frac{1}{4}w_2$	~~~	$x_2 = 2 - \frac{1}{4}x_1 + \frac{1}{4}w_2$
$w_1 = 4 + 8(2 - \frac{1}{4}x_1 + \frac{1}{4}w_2) - 3w_2$		$w_1 = 20 - 2x_1 - w_2$
$x_3 = 1 + 3(2 - \frac{1}{4}x_1 + \frac{1}{4}w_2) - w_2$		$x_3 = 7 - \frac{3}{4}x_1 - \frac{1}{4}w_2$

We can confirm that our objective value has increased from 20 to 34: the change is exactly equal to 7 (the reduced cost of x_2 in our previous dictionary) multiplied by 2 (the value of x_2 in our new feasible solution).

In full, our basic feasible solution is now $(x_1, x_2, x_3, w_1, w_2) = (0, 2, 7, 20, 0)$: we grow 2 acres of corn and 7 acres of soy, for a profit of \$100 $\cdot \zeta = 3400 . (This would seem more reasonable if the farm were a more realistic size!)

3.4 The end of the simplex method

If we look at the latest dictionary we've gotten, and try to pick an entering variable, it looks at first like we have a problem. Both x_1 and w_2 have a negative reduced cost: both of them would decrease ζ if we brought them into the basis.

This means we can't improve our objective value by one step of the simplex method. Should we worry that we're trapped at a local optimum that isn't as good as some far-away corner? No! In fact, we can prove that our current basic feasible solution is optimal.

The top equation of the dictionary says $\zeta = 34 - \frac{7}{4}x_1 - \frac{1}{4}w_2$. Remember: this is a universal equation that holds for *every* feasible solution of our linear program, because we deduced it from combining $\zeta = 2x_1 + 3x_2 + 4x_3$ with our constraints.

All our variables are nonnegative; in particular, $x_1 \ge 0$ and $w_2 \ge 0$. So we are taking 34 and subtracting two nonnegative values from it. It follows that at all feasible solutions, $\zeta \le 34$.

But we've just seen that our current basic feasible solution achieves an objective value of $\zeta = 34$ exactly. We conclude that our basic feasible solution is optimal: it's impossible for the farmer to

make a larger profit. This is the universal rule for when the simplex method halts:

- When maximizing ζ , stop when all reduced costs are at least 0.
- When minimizing ζ , stop when all reduced costs are at most 0.

In both cases, it is fine to see a reduced cost of 0. What happens when we pivot on a variable that has a reduced cost of 0? The objective value always changes by the product of the old reduced cost and the new value of the entering variable. In this case, that product will always be 0, because the reduced cost is 0. So pivoting on such a variable will never change the objective function.

However, a reduced cost of 0 indicates that there may be multiple optimal solutions: we can get other solutions with the same objective value by pivoting on such a variable.

By contrast, in our case, being given $\zeta = 34 - \frac{7}{4}x_1 - \frac{1}{4}w_2$ tells us that we must have $x_1 = 0$ and $w_2 = 0$ in any optimal solution. (If one of these variables were positive, we'd subtract a positive number from 34, and get a smaller objective value.) The feasible solution we have is actually the only solution possible when $x_1 = 0$ and $w_2 = 0$, so it is the *unique* optimal solution to our problem.

4 Optional: dictionaries and tableaux

There are two main ways that people have come up with to represent intermediate steps in the simplex method. Following the textbook, we are using dictionaries, which were introduced by Chvátal in his 1983 textbook on linear programming.

A tableau is another way of representing the same information. Fundamentally, they are based on writing the same equations with *all* variables on the right, and with constants and the objective value on the left. For example, here is how one of our dictionaries would appear in this form:

$\zeta = 20 + 7x_2 - 2w_2$		$\zeta - 20 =$	$7x_2$	$-2w_2$
$x_1 = 8 - 4x_2 + w_2$	~~~	$8 = x_1 +$	$4x_2$	$- w_2$
$w_1 = 4 + 8x_2 - 3w_2$	• /	4 =	$-8x_{2}$	$+ w_1 + 3w_2$
$x_3 = 1 + 3x_2 - w_2$		1 =	$-3x_2 + x_3$	$+ w_2$

This is closer to the way we write things when we do Gaussian elimination. It has more columns, but the advantage is that it is easier to put in a table, without having to write the variables every time. Here, the *simplex tableau* would be:

ī

		x_1	x_2	x_3	w_1	w_2
$-\zeta$	-20	0	7	0	0	-2
x_1	8	1	4	0	0	-1
w_1	4	0	-8	0	1	3
x_3	1	0	-3	1	0	1

We annotate the columns with the variables whose coefficients are in those columns; we annotate the rows with the basic variable in that row. We write $-\zeta$ in the objective row to remind ourselves that with this method, -20 is the *negative* of the objective value. Iterations of the simplex method are just ordinary row reduction with this grid of numbers.

Having the negative of the objective value appear in the tableau is a bit weird, so you might also see tableaux written with the top equation ordered differently: as $\zeta - 7x_2 + 2w_2 = 20$. Then, the tableau could look like the following:

		ζ	x_1	x_2	x_3	w_1	w_2
ζ	20	1	0	-7	0	0	2
x_1	8	0	1	4	0	0	-1
w_1	4	0	0	-8	0	1	3
x_3	1	0	0	-3	1	0	1

This way, you can read off the current solution and the objective value from the left of the tableau. The downside of **this** approach is that the reduced costs in this version are the negatives of the reduced costs we're used to seeing! There's nothing wrong with that—provided we reverse our rules of dealing with the reduced costs—but it means that building the initial tableau is a little bit weird. The numbers we'll have to put in the top row of the tableau will be the *negatives* of the coefficients in the objective function, because we rewrite $\zeta = c_1 x_1 + c_2 x_2 + \cdots + c_n x_n$ as $\zeta - c_1 x_1 - c_2 x_2 - \cdots - c_n x_n = 0$.

There are other variants of the tableau, with the rows and columns rearranged in minor ways. This makes it extra important to keep the rows and columns labeled with variables, so that we can interpret them more easily.