1 Introducing objective functions

Let's take an example linear program in equational form, conveniently formulated so that it starts out row-reduced:

$$\begin{aligned}
\underset{\mathbf{x} \in \mathbb{R}^5}{\text{maximize}} \quad & x_1 + x_2 + x_3 + x_4 + x_5 \\
\text{subject to} \quad & x_1 \qquad\quad + 2x_4 + x_5 = 1 \\
& \quad x_2 \quad - 3x_4 + x_5 = 2 \\
& \qquad x_3 + x_4 - 3x_5 = 3 \\
& x_1, x_2, x_3, x_4, x_5 \geq 0.
\end{aligned}$$

The first thing to realize is that when the equations hold, the objective function has many equivalent forms. Since $x_1 + 2x_4 + x_5 = 1$, for example, if we subtract $x_1 + 2x_4 + x_5$ and add 1, we can maximize $x_2 + x_3 - x_4 + 1$ instead (and get the same solution).

Life gets a bit more convenient if we are only adding and subtracting equations. So we will slightly rephrase the problem as

$$\begin{aligned}
\underset{\mathbf{x} \in \mathbb{R}^5}{\text{maximize}} \quad & z \\
\text{subject to} \quad & x_1 + x_2 + x_3 + x_4 + x_5 = z \\
& x_1 \qquad\quad + 2x_4 + x_5 = 1 \\
& \quad x_2 \quad - 3x_4 + x_5 = 2 \\
& \qquad x_3 + x_4 - 3x_5 = 3 \\
& x_1, x_2, x_3, x_4, x_5 \geq 0.
\end{aligned}$$

Which of the many equivalent expressions for z should we use? Well, one of the nice things about the row-reduced equations is that we can set $x_4 = x_5 = 0$ and read off a basic solution. We can eliminate x_1, x_2, x_3 from our expression for z, and then be able to read off the objective value of the basic solution, too. To do this, subtract all other equations (or, in general, their multiples) from the equation for z, getting:

$$\begin{aligned}
\underset{\mathbf{x} \in \mathbb{R}^5, z \in \mathbb{R}}{\text{maximize}} \quad & z \\
\text{subject to} \quad & x_4 + 2x_5 = z - 6 \\
& x_1 \qquad\quad + 2x_4 + x_5 = 1 \\
& \quad x_2 \quad - 3x_4 + x_5 = 2 \\
& \qquad x_3 + x_4 - 3x_5 = 3 \\
& x_1, x_2, x_3, x_4, x_5 \geq 0.
\end{aligned}$$

---

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

The basic solution $\mathbf{x} = (1, 2, 3, 0, 0)$ obtained by setting $x_4 = x_5 = 0$ satisfies $0 = z - 6$, so the objective value is 6.

2 The simplex tableau

At this point, it's convenient to switch away from writing out an entire system of equations every single time: after all, when solving systems of equations in linear algebra, we just write down grids of numbers. This notation for the simplex method is called a *simplex tableau* and includes a bit more information:

	x_1	x_2	x_3	x_4	x_5	
x_1	1	0	0	2	1	1
x_2	0	1	0	-3	1	2
x_3	0	0	1	1	-3	3
$-z$	0	0	0	1	2	-6

Some comments:

- The columns are simply labeled with the name of the variable whose coefficients are written in that column. This should really be done in ordinary Gaussian elimination as well, but people are lazy.

- The first three rows are labeled with the basic variable "in charge of" that row. When you set the nonbasic variables x_4 and x_5 to 0, we find the value of that basic variable, in the rightmost column. Here, $x_1 = 1$, $x_2 = 2$, and $x_3 = 3$.

- The final row is the row with the objective function in it. It is labeled $-z$, because the rightmost column of -6 contains the *negative* of z's value when $\mathbf{x} = (1, 2, 3, 0, 0)$.

  Why does it contain the negative of z's value? Because the row is short for the equation "$x_4 + 2x_5 = z - 6$".

  Why don't we just write 6 instead of -6 and z instead of $-z$? Because we want to set it up so that we can just subtract any row from any other row, and end up with true equations.

3 Pivoting and the simplex tableau

Let's see what happens when we do a pivot step to bring x_4 into the basis.

Some review: when x_5 stays 0, but x_4 changes, the three equations tell us that x_1, x_2, x_3 change according to the rules

$$x_1 = 1 - 2x_4 \qquad x_2 = 2 + 3x_4, \qquad x_3 = 3 - x_4.$$

As we increase x_4, x_1 will reach 0 when $x_4 = \frac{1}{2}$, x_2 will never reach 0, and x_3 will reach 0 when $x_4 = 3$. The first variable to reach 0 is x_1, so x_1 leaves the basis and x_4 enters.

We can quickly determine this by taking the ratios of the right-hand-side constants over the coefficients of x_4. (We skip the ratio in the x_2 row, because the coefficient of x_4 is negative.) For

convenience, let's record these in the tableau:

	x_1	x_2	x_3	x_4	x_5		x_4's ratios
x_1	1	0	0	$\boxed{2}$	1	1	$^1/_2$
x_2	0	1	0	-3	1	2	—
x_3	0	0	1	1	-3	3	$^3/_1$
$-z$	0	0	0	1	2	-6	

To find the variable that leaves the basis, we look for the smallest of these ratios, which is x_1 in this case. So we want to make the boxed entry of the tableau, x_4's entry in x_1's row, be the new pivot for that row. This is just row reduction. We divide that row by 2, then add or subtract appropriate multiples of that row from other rows, to make every other entry in x_4's column 0. The result:

	x_1	x_2	x_3	x_4	x_5	
x_4	$^1/_2$	0	0	1	$^1/_2$	$^1/_2$
x_2	$^3/_2$	1	0	0	$^5/_2$	$^7/_2$
x_3	$-^1/_2$	0	1	0	$-^7/_2$	$^5/_2$
$-z$	$-^1/_2$	0	0	0	$^3/_2$	$-^{13}/_2$

The new basic feasible solution is $\mathbf{x} = (0, \frac{7}{2}, \frac{5}{2}, \frac{1}{2}, 0)$ with objective value $\frac{13}{2}$.

4 Pivoting and the objective function

After this pivot step, the objective value increased from 6 to $\frac{13}{2} = 6.5$. This is good, because we're trying to maximize z, but it's a complete accident, because we didn't know what we were doing. How can we predict what will happen?

We affect the value of z when we zero out x_4's coefficient in the objective row. Specifically,

- If that coefficient is positive, then we will subtract a multiple of x_4's row from the objective row. The rightmost value in x_4's row must be nonnegative: it's the value of x_4 in the new basic feasible solution. So we'll subtract a nonnegative value from the current $-z$, decreasing $-z$ and increasing z.

- If that coefficient is negative, then we will add a multiple of x_4's row to the objective row. This will increase $-z$, so it will decrease z.

The coefficient of a variable x_i in the objective row is called x_i's *reduced cost*. (It is a "cost" by analogy with an economic problem when the objective value is the total cost of doing something. It is "reduced" because it has been changed by row reduction.) So the rule is:

The change in z from pivoting has the same sign as the reduced cost.

So if we're solving a maximization problem, we want to pivot on variables with positive reduced cost. If we're solving a minimization problem, we want to pivot on variables with a negative reduced cost.

More is true: the value of the reduced cost is actually the *rate* at which z changes as we increase the variable we're pivoting on. So if the reduced cost is very large (in absolute value) then the

objective value is changing very quickly. As a result, one possible pivoting strategy is to choose the pivot with the largest possible reduced cost (of the correct sign).

This isn't foolproof, however, and you might also just choose the first pivot whose reduced cost has the correct sign. Ways to make this choice are called "pivoting rules", and we will talk about the details later.

5 The end of the simplex method

If we do another step of the simplex method, our options are to pivot on x_1 and x_5. We shouldn't pivot on x_1, because it has a negative reduced cost (not surprising, since pivoting on x_1 would bring us back to where we were before). Pivoting on x_5 makes x_4 leave the basis, and we get the tableau

	x_1	x_2	x_3	x_4	x_5	
x_5	1	0	0	2	1	1
x_2	-1	1	0	-5	0	1
x_3	3	0	1	7	0	6
$-z$	-2	0	0	-3	0	-8

The new basic feasible solution is $\mathbf{x} = (0, 1, 6, 0, 1)$, with objective value 8.

Now what? All the reduced costs are negative, so no single pivot step will improve our situation. This by itself might not be a guarantee of anything: could the actual optimal solution be some other point which we can't get to by a single pivot step?

No, it couldn't, and hidden in this simplex tableau is a proof that 8 is the best possible objective value. Translated back into equation form, the objective row of the tableau reads

$$-2x_1 - 3x_4 = z - 8 \iff z = 8 - 2x_1 - 3x_4.$$

Both $2x_1$ and $3x_4$ are guaranteed to be nonnegative, because $\mathbf{x} \geq \mathbf{0}$ is a constraint of our linear program. So this tells us that:

- The maximum value of z, the objective function, can't possibly be higher than 8, because we're starting at 8 and subtracting some nonnegative stuff.

- Since our current basic feasible solution $\mathbf{x} = (0, 1, 6, 0, 1)$ achieves $z = 8$, it is the optimal solution.

- Any optimal solution must have $x_1 = x_4 = 0$: if either of these variables is strictly positive, then $z < 8$. Conversely, any solution with $x_1 = x_4 = 0$ must have $z = 8$, so it will be optimal.

  (In this case, the solution we've found is the only optimal solution, but this is useful to know in problems where there are many optimal solutions.)

This argument works whenever we use the simplex method. So if no more pivot steps are available that move z in the right direction, then the simplex method is finished and we've arrived at an optimal solution.