

Linear Programming Notes VII

Sensitivity Analysis

1 Introduction

When you use a mathematical model to describe reality you must make approximations. The world is more complicated than the kinds of optimization problems that we are able to solve. Linearity assumptions usually are significant approximations. Another important approximation comes because you cannot be sure of the data that you put into the model. Your knowledge of the relevant technology may be imprecise, forcing you to approximate values in A , b , or c . Moreover, information may change. **Sensitivity analysis** is a systematic study of how sensitive (duh) solutions are to (small) changes in the data. The basic idea is to be able to give answers to questions of the form:

1. If the objective function changes, how does the solution change?
2. If resources available change, how does the solution change?
3. If a constraint is added to the problem, how does the solution change?

One approach to these questions is to solve lots of linear programming problems. For example, if you think that the price of your primary output will be between \$100 and \$120 per unit, you can solve twenty different problems (one for each whole number between \$100 and \$120).¹ This method would work, but it is inelegant and (for large problems) would involve a large amount of computation time. (In fact, the computation time is cheap, and computing solutions to similar problems is a standard technique for studying sensitivity in practice.) The approach that I will describe in these notes takes full advantage of the structure of LP programming problems and their solution. It turns out that you can often figure out what happens in “nearby” linear programming problems just by thinking and by examining the information provided by the simplex algorithm. In this section, I will describe the sensitivity analysis information provided in Excel computations. I will also try to give an intuition for the results.

2 Intuition and Overview

Throughout these notes you should imagine that you must solve a linear programming problem, but then you want to see how the answer changes if the problem is changed. In every case, the results assume that **only one thing about the problem changes**. That is, in sensitivity analysis you evaluate what happens when only one parameter of the problem changes.

¹OK, there are really 21 problems, but who is counting?

maybe you would want to set $x_2 = 0$ instead of $x_2 = 10.4$. On the other hand, a small reduction in x_2 's objective function coefficient would typically not cause you to change your solution. In contrast to the case of the non-basic variable, such a change will change the value of your objective function. You compute the value by plugging in x into the objective function, if $x_2 = 10.4$ and the coefficient of x_2 goes down from 4 to 2, then the contribution of the x_2 term to the value goes down from 41.6 to 20.8 (assuming that the solution remains the same).

If the coefficient of a basic variable goes up, then your value goes up and you still want to use the variable, but if it goes up enough, you may want to adjust x so that it x_2 is even possible. In many cases, this is possible by finding another basis (and therefore another solution). So, intuitively, there should be a range of values of the coefficient of the objective function (a range that includes the original value) in which the solution of the problem does not change. Outside of this range, the solution will change (to lower the value of the basic variable for reductions and increase its value of increases in its objective function coefficient). The value of the problem always changes when you change the coefficient of a basic variable.

2.2 Changing a Right-Hand Side Constant

We discussed this topic when we talked about duality. I argued that dual prices capture the effect of a change in the amounts of available resources. When you changed the amount of resource in a non-binding constraint, then increases never changed your solution. Small decreases also did not change anything, but if you decreased the amount of resource enough to make the constraint binding, your solution could change. (Note the similarity between this analysis and the case of changing the coefficient of a non-basic variable in the objective function.

Changes in the right-hand side of binding constraints always change the solution (the value of x must adjust to the new constraints). We saw earlier that the dual variable associated with the constraint measures how much the objective function will be influenced by the change.

2.3 Adding a Constraint

If you add a constraint to a problem, two things can happen. Your original solution satisfies the constraint or it doesn't. If it does, then you are finished. If you had a solution before and the solution is still feasible for the new problem, then you must still have a solution. If the original solution does not satisfy the new constraint, then possibly the new problem is infeasible. If not, then there is another solution. The value must go down. (Adding a constraint makes the problem harder to satisfy, so you cannot possibly do better than before). If your original solution satisfies your new constraint, then you can do as well as before. If not, then you will do worse.²

²There is a rare case in which originally your problem has multiple solutions, but only some of them satisfy the added constraint. In this case, which you need not worry about,

2.4 Relationship to the Dual

The objective function coefficients correspond to the right-hand side constants of resource constraints in the dual. The primal's right-hand side constants correspond to objective function coefficients in the dual. Hence the exercise of changing the objective function's coefficients is really the same as changing the resource constraints in the dual. It is extremely useful to become comfortable switching back and forth between primal and dual relationships.

3 Understanding Sensitivity Information Provided by Excel

Excel permits you to create a sensitivity report with any solved LP. The report contains two tables, one associated with the variables and the other associated with the constraints. In reading these notes, keep the information in the sensitivity tables associated with the first simplex algorithm example nearby.

3.1 Sensitivity Information on Changing (or Adjustable) Cells

The top table in the sensitivity report refers to the variables in the problem. The first column (Cell) tells you the location of the variable in your spreadsheet; the second column tells you its name (if you named the variable); the third column tells you the final value; the fourth column is called the reduced cost; the fifth column tells you the coefficient in the problem; the final two columns are labeled "allowable increase" and "allowable decrease." Reduced cost, allowable increase, and allowable decrease are new terms. They need definitions.

The allowable increases and decreases are easier. I will discuss them first. The allowable increase is the amount by which you can increase the coefficient of the objective function without causing the optimal basis to change. The allowable decrease is the amount by which you can decrease the coefficient of the objective function without causing the optimal basis to change.

Take the first row of the table for the example. This row describes the variable x_1 . The coefficient of x_1 in the objective function is 2. The allowable increase is 9, the allowable decrease is "1.00E+30," which means 10^{30} , which **really** means ∞ . This means that provided that the coefficient of x_1 in the objective function is less than $11 = 2 + 9 = \text{original value} + \text{allowable increase}$, the basis does not change. Moreover, since x_1 is a non-basic variable, when the basis stays the same, the value of the problem stays the same too. The information in this line confirms the intuition provided earlier and adds something new. What is confirmed is that if you lower the objective coefficient of a non-basic variable, then your solution does not change. (This means that the allowable decrease will always be infinite for a non-basic variable.) The example also demonstrates

your value will stay the same.

that increasing the coefficient of a non-basic variable may lead to a change in basis. In the example, if you increase the coefficient of x_1 from 2 to anything greater than 9 (that is, if you add more than the allowable increase of 7 to the coefficient), then you change the solution. The sensitivity table does not tell you how the solution changes, but common sense suggests that x_1 will take on a positive value. Notice that the line associated with the other non-basic variable of the example, x_3 , is remarkably similar. The objective function coefficient is different (3 rather than 2), but the allowable increase and decrease are the same as in the row for x_1 . It is a coincidence that the allowable increases are the same. It is no coincidence that the allowable decrease is the same. We can conclude that the solution of the problem does not change as long as the coefficient of x_3 in the objective function is less than or equal to 10.

Consider now the basic variables. For x_2 the allowable increase is infinite while the allowable decrease is 2.69 (it is $2\frac{9}{13}$ to be exact). This means that if the solution won't change if you increase the coefficient of x_2 , but it will change if you decrease the coefficient enough (that is, by more than 2.7). The fact that your solution does not change no matter how much you increase x_2 's coefficient means that there is no way to make $x_2 > 10.4$ and still satisfy the constraints of the problem. The fact that your solution does change when you increase x_2 's coefficient by enough means that there is a feasible basis in which x_2 takes on a value lower than 10.4. (You knew that. Examine the original basis for the problem.) The range for x_4 is different. Line four of the sensitivity table says that the solution of the problem does not change provided that the coefficient of x_4 in the objective function stays between 16 (allowable increase 15 plus objective function coefficient 1) and -4 (objective function coefficient minus allowable decrease). That is, if you make x_4 sufficiently more attractive, then your solution will change to permit you to use more x_4 . If you make x_4 sufficiently less attractive the solution will also change. This time to use less x_4 . Even when the solution of the problem does not change, when you change the coefficient of a basic variable the value of the problem will change. It will change in a predictable way. Specifically, you can use the table to tell you the solution of the LP when you take the original constraints and replace the original objective function by

$$\max 2x_1 + 6x_2 + 3x_3 + x_4$$

(that is, you change the coefficient of x_2 from 4 to 6), then the solution to the problem remains the same. The value of the solution changes because now you multiply the 10.4 units of x_2 by 6 instead of 4. The objective function therefore goes up by 20.8.

The reduced cost of a variable is the smallest change in the objective function coefficient needed to arrive at a solution in which the variable takes on a positive value when you solve the problem. This is a mouthful. Fortunately, reduced costs are redundant information. The reduced cost is the negative of the allowable increase for non-basic variables (that is, if you change the coefficient of x_1 by -7 , then you arrive at a problem in which x_1 takes on a positive

value in the solution). This is the same as saying that the allowable increase in the coefficient is 7. The reduced cost of a basic variable is always zero (because you need not change the objective function at all to make the variable positive). Neglecting rare cases in which a basis variable takes on the value 0 in a solution, you can figure out reduced costs from the other information in the table: If the final value is positive, then the reduced cost is zero. If the final value is zero, then the reduced cost is negative one times the allowable increase. Remarkably, the reduced cost of a variable is also the amount of slack in the dual constraint associated with the variable. With this interpretation, complementary slackness implies that if a variable that takes on a positive value in the solution, then its reduced cost is zero.

3.2 Sensitivity Information on Constraints

The second sensitivity table discusses the constraints. The cell column identifies the location of the left-hand side of a constraint; the name column gives its name (if any); the final value is the value of the left-hand side when you plug in the final values for the variables; the shadow price is the dual variable associated with the constraint; the constraint R.H. side is the right hand side of the constraint; allowable increase tells you by how much you can increase the right-hand side of the constraint without changing the basis; the allowable decrease tells you by how much you can decrease the right-hand side of the constraint without changing the basis.

Complementary Slackness guarantees a relationship between the columns in the constraint table. The difference between the “Constraint Right-Hand Side” column and the “Final Value” column is the slack. (So, from the table, the slack for the three constraints is $0 (= 12 - 12)$, $37 (= 7 - (-30))$, and $0 (= 10 - 10)$, respectively. We know from Complementary Slackness that if there is slack in the constraint then the associated dual variable is zero. Hence CS tells us that the second dual variable must be zero.

Like the case of changes in the variables, you can figure out information on allowable changes from other information in the table. The allowable increase and decrease of non-binding variables can be computed knowing final value and right-hand side constant. If a constraint is not binding, then adding more of the resource is not going to change your solution. Hence the allowable increase of a resource is infinite for a non-binding constraint. (A nearly equivalent, and also true, statement is that the allowable increase of a resource is infinite for a constraint with slack.) In the example, this explains why the allowable increase of the second constraint is infinite. One other quantity is also no surprise. The allowable decrease of a non-binding constraint is equal to the slack in the constraint. Hence the allowable decrease in the second constraint is 37. This means that if you decrease the right-hand side of the second constraint from its original value (7) to anything greater than -30 you do not change the optimal basis. In fact, the only part of the solution that changes when you do this is that the value of the slack variable for this constraint changes. In this paragraph, the point is only this: If you solve an LP and find that a constraint is not binding,

then you can remove all of the unused (slack) portion of the resource associated with this constraint and not change the solution to the problem.

The allowable increases and decreases for constraints that have no slack are more complicated. Consider the first constraint. The information in the table says that if the right-hand side of the first constraint is between 10 (original value 12 minus allowable decrease 2) and infinity, then the basis of the problem does not change. What these columns do not say is that the solution of the problem does change. Saying that the basis does not change means that the variables that were zero in the original solution continue to be zero in the new problem (with the right-hand side of the constraint changed). However, when the amount of available resource changes, necessarily the values of the other variables change. (You can think about this in many ways. Go back to a standard example like the diet problem. If your diet provides exactly the right amount of Vitamin C, but then for some reason you learn that you need more Vitamin C. You will certainly change what you eat and (if you aren't getting your Vitamin C through pills supplying pure Vitamin C) in order to do so you probably will need to change the composition of your diet - a little more of some foods and perhaps less of others. I am saying that (within the allowable range) you will not change the foods that you eat in positive amounts. That is, if you ate only spinach and oranges and bagels before, then you will only eat these things (but in different quantities) after the change. Another thing that you can do is simply re-solve the LP with a different right-hand side constant and compare the result.

To finish the discussion, consider the third constraint in the example. The values for the allowable increase and allowable decrease guarantee that the basis that is optimal for the original problem (when the right-hand side of the third constraint is equal to 10) remains obtain provided that the right-hand side constant in this constraint is between -2.3333 and 12. Here is a way to think about this range. Suppose that your LP involves four production processes and uses three basic ingredients. Call the ingredients land, labor, and capital. The outputs vary use different combinations of the ingredients. Maybe they are growing fruit (using lots of land and labor), cleaning bathrooms (using lots of labor), making cars (using lots of labor and a bit of capital), and making computers (using lots of capital). For the initial specification of available resources, you find that you want to grow fruit and make cars. If you get an increase in the amount of capital, you may wish to shift into building computers instead of cars. If you experience a decrease in the amount of capital, you may wish to shift away from building cars and into cleaning bathrooms instead.

As always when dealing with duality relationships, the the "Adjustable Cells" table and the "Constraints" table really provide the same information. Dual variables correspond to primal constraints. Primal variables correspond to dual constraints. Hence, the "Adjustable Cells" table tells you how sensitive primal variables and dual constraints are to changes in the primal objective function. The "Constraints" table tells you how sensitive dual variables and primal constraints are to changes in the dual objective function (right-hand side constants in the primal).

4 Example

In this section I will present another formulation example and discuss the solution and sensitivity results.

Imagine a furniture company that makes tables and chairs. A table requires 40 board feet of wood and a chair requires 30 board feet of wood. Wood costs \$1 per board foot and 40,000 board feet of wood are available. It takes 2 hours of skilled labor to make an unfinished table or an unfinished chair. Three more hours of labor will turn an unfinished table into a finished table; two more hours of skilled labor will turn an unfinished chair into a finished chair. There are 6000 hours of skilled labor available. (Assume that you do not need to pay for this labor.) The prices of output are given in the table below:

Product	Price
Unfinished Table	\$70
Finished Table	\$140
Unfinished Chair	\$60
Finished Chair	\$110

We want to formulate an LP that describes the production plans that the firm can use to maximize its profits.

The relevant variables are the number of finished and unfinished tables, I will call them T_F and T_U , and the number of finished and unfinished chairs, C_F and C_U . The revenue is (using the table):

$$70T_U + 140T_F + 60C_U + 110C_F,$$

, while the cost is $40T_U + 40T_F + 30C_U + 30C_F$ (because lumber costs \$1 per board foot).

The constraints are:

1. $40T_U + 40T_F + 30C_U + 30C_F \leq 40000$.
2. $2T_U + 5T_F + 2C_U + 4C_F \leq 6000$.

The first constraint says that the amount of lumber used is no more than what is available. The second constraint states that the amount of labor used is no more than what is available.

Excel finds the answer to the problem to be to construct only finished chairs (1333.333 - I'm not sure what it means to make a sell $\frac{1}{3}$ chair, but let's assume that this is possible). The profit is \$106,666.67.

Here are some sensitivity questions.

1. What would happen if the price of unfinished chairs went up? Currently they sell for \$60. Because the allowable increase in the coefficient is \$50, it would not be profitable to produce them even if they sold for the same amount as finished chairs. If the price of unfinished chairs went down, then certainly you wouldn't change your solution.

2. What would happen if the price of unfinished tables went up?

Here something apparently absurd happens. The allowable increase is greater than 70. That is, even if you could sell unfinished tables for more than finished tables, you would not want to sell them. How could this be? The answer is that at current prices you don't want to sell finished tables. Hence it is not enough to make unfinished tables more profitable than finished tables, you must make them more profitable than finished chairs. Doing so requires an even greater increase in the price.

3. What if the price of finished chairs fell to \$100? This change would alter your production plan, since this would involve a \$10 decrease in the price of finished chairs and the allowable decrease is only \$5. In order to figure out what happens, you need to re-solve the problem. It turns out that the best thing to do is specialize in finished tables, producing 1000 and earning \$100,000. Notice that if you continued with the old production plan your profit would be $70 \times 1333\frac{1}{3} = 93,333\frac{1}{3}$, so the change in production plan was worth more than \$6,000.
4. How would profit change if lumber supplies changed? The shadow price of the lumber constraint is \$2.67. The range of values for which the basis remains unchanged is 0 to 45,000. This means that if the lumber supply went up by 5000, then you would continue to specialize in finished chairs, and your profit would go up by $\$2.67 \times 5000 = \$10,333$. At this point you presumably run out of labor and want to reoptimize. If lumber supply decreased, then your profit would decrease, but you would still specialize in finished chairs.
5. How much would you be willing to pay an additional carpenter? Skilled labor is not worth anything to you. You are not using the labor than you have. Hence, you would pay nothing for additional workers.
6. Suppose that industrial regulations complicate the finishing process, so that it takes one extra hour per chair or table to turn an unfinished product into a finished one. How would this change your plans?

You cannot read your answer off the sensitivity table, but a bit of common sense tells you something. The change cannot make you better off. On the other hand, to produce 1,333.33 finished chairs you'll need 1,333.33 extra hours of labor. You do not have that available. So the change will change your profit. Using Excel, it turns out that it becomes optimal to specialize in finished tables, producing 1000 of them and earning \$100,000. (This problem differs from the original one because the amount of labor to create a finished product increases by one unit.)

7. The owner of the firm comes up with a design for a beautiful hand-crafted cabinet. Each cabinet requires 250 hours of labor (this is 6 weeks of full time work) and uses 50 board feet of lumber. Suppose that the company can sell a cabinet for \$200, would it be worthwhile? You could solve this

problem by changing the problem and adding an additional variable and an additional constraint. Note that the coefficient of cabinets in the objective function is 150, which reflects the sale price minus the cost of lumber. I did the computation. The final value increased to 106,802.7211. The solution involved reducing the output of unfinished chairs to 1319.727891 and increasing the output of cabinets to 8.163265306. (Again, please tolerate the fractions.) You could not have guessed these figures in advance, but you could figure out that making cabinets was a good idea. The way to do this is to value the inputs to the production of cabinets. Cabinets require labor, but labor has a shadow price of zero. They also require lumber. The shadow price of lumber is \$2.67, which means that each unit of lumber adds \$2.67 to profit. Hence 50 board feet of lumber would reduce profit by \$133.50. Since this is less than the price at which you can sell cabinets (minus the cost of lumber), you are better off using your resources to build cabinets. (You can check that the increase in profit associated with making cabinets is \$16.50, the added profit per unit, times the number of cabinets that you actually produce.) I attached a sheet where I did the same computation assuming that the price of cabinets was \$150. In this case, the additional option does not lead to cabinet production.