

Chapter 6: Sensitivity Analysis

Suppose that you have just completed a linear programming solution which will have a major impact on your company, such as determining how much to increase the overall production capacity, and are about to present the results to the board of directors. How confident are you in the results? How much will the results change if your basic data (e.g. profit per item produced, or availability of a component) is slightly wrong? Will that have a minor impact on your results? Will it give a completely different outcome, or change the outcome only slightly?

These are the kinds of questions addressed by sensitivity analysis. Formally, the question is this: is my optimum solution (both the values of the variables and the value of the objective function) sensitive to a small change in one of the original problem coefficients (e.g. coefficients of the variables in the objective function or constraints, or the right hand side constants in the constraints)? If Z or the x_i change when an original coefficient is changed, then we say that the LP is *sensitive*. We could ask, for example, if the Acme Bicycle Company solution is sensitive to a reduction in the availability of the metal finishing machine from 4 hours per day to only 3 (i.e. a change in the third constraint from $x_1+x_2 \leq 4$ to $x_1+x_2 \leq 3$).

This sort of examination of the impact of the input data on the output results is crucial. The procedures and algorithms of mathematical programming are important, but the problems that really bedevil you in practice are usually associated with data: getting it at all, and getting *accurate* data. Some data, necessary for your mathematical model, is inherently uncertain. Consider profit per item, for example, which is approximated from estimates of the fluctuating costs of raw materials, expected sales volumes, labour costs, etc. What you want to know from sensitivity analysis is which data has a significant impact on the results: then you can concentrate on getting accurate data for those items, or at least running through several scenarios with various values of the crucial data in place to get an idea of the range of possible outcomes.

There are several ways to approach sensitivity analysis. If your model is small enough to solve quite quickly, you can use a *brute force* approach: simply change the initial data and solve the model again to see what results you get. You can do this as many times as needed. At the opposite extreme, if your model is very large and takes a long time to solve, you can apply the formal methods of *classical sensitivity analysis*. The classical methods rely on the relationship between the initial tableau and any later tableau (in particular the optimum tableau) to quickly update the optimum solution when changes are made to the coefficients of the original tableau.

Between these two extremes is *computer-based ranging*. This is simple information about how much certain coefficients can change before the current optimum solution is fundamentally changed. Most commercial LP solvers provide such information. In this introductory book, we will concentrate on this form of sensitivity analysis.

One final observation on the state of the art in sensitivity analysis: you are typically limited to analyzing the impact of changing only *one* coefficient at a time. There are a few accepted techniques for changing several coefficients at once: the 100% rule, and parametric programming. The 100% rule is typically limited to changing only a few coefficients at once, with tight limits on how much they can change, and parametric programming changes all of the coefficients in ratio. What is really needed is a method that allows *all* of the coefficients to vary independently. Recent research by your humble author and graduate student Khaled Ramadan provides such a method: it allows all of the coefficients to be specified as intervals (e.g. the availability of the metal finishing machine is between 3 and 5 hours per day). Unfortunately this method is not yet implemented in commercial LP solvers.

Simple Computer-Based Sensitivity Analysis

Most commercial LP solvers return at least the following information:

- The objective function coefficients for the original variables at the optimum, called the *reduced costs*.
- The objective function coefficients for the slack and surplus variables at the optimum, called the *shadow prices* or *dual prices*.
- The ranges of the original objective function coefficients of the original variables for which the current basis remains optimal.
- The ranges of the right-hand-side constants for the constraints for which the current basis remains optimal.

What is missing from this list is any mention of the constraint coefficients. The real weakness of the simple computer-based sensitivity analysis is that it does not deal with changes to constraint coefficients.

Consider the solution output returned by LINDO solver for the Acme Bicycle Company problem, for example:

```

LP OPTIMUM FOUND AT STEP      2

      OBJECTIVE FUNCTION VALUE
  1)      50.000000

      VARIABLE            VALUE            REDUCED COST
      X1                   2.000000             .000000
      X2                   2.000000             .000000

      ROW  SLACK OR SURPLUS      DUAL PRICES
  2)              .000000             5.000000
  3)              1.000000             .000000
  4)              .000000            10.000000

```

Note that LINDO refers to the objective function as “row (1)”, and the constraints as rows (2)–(4). LINDO also provides the simple ranging information that can be used for sensitivity analysis:

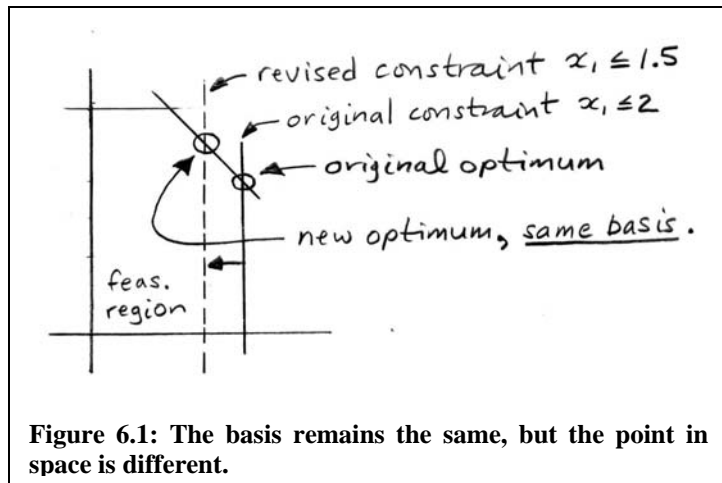
RANGES IN WHICH THE BASIS IS UNCHANGED:

VARIABLE	CURRENT COEF	OBJ COEFFICIENT RANGES	
		ALLOWABLE INCREASE	ALLOWABLE DECREASE
x1	15.000000	INFINITY	5.000000
x2	10.000000	5.000000	10.000000

ROW	CURRENT RHS	RIGHTHAND SIDE RANGES	
		ALLOWABLE INCREASE	ALLOWABLE DECREASE
2	2.000000	2.000000	1.000000
3	3.000000	INFINITY	1.000000
4	4.000000	1.000000	2.000000

There are several things to observe about this output data. First, note that the reduced costs for x_1 and x_2 are both zero in the top table: why? Well, the reduced costs are the objective function coefficients of the original variables, and since both original variables (x_1 and x_2) are basic at the optimum, their objective function coefficients must be zero when the tableau is put into proper form. This is always true: either the variable is zero (nonbasic), or the reduced cost or dual price is zero. You can see that the pattern holds for the slack and surplus variables too. The dual prices for rows (2) and (4) are nonzero at the optimum because they correspond to the two active constraints at the optimum, hence their slack variables are nonbasic (value is zero), so the dual prices can be nonzero. When *both* the variable and the associated reduced cost or dual price are zero, then you have either degeneracy if the variable is basic, or multiple optima if the variable is nonbasic, just as we would expect from the tableau.

In the ranging information, the “allowable increase” and “allowable decrease” refer to the maximum *changes* from the “current coefficient” or “current right hand side” which will keep the optimum solution at the same basis. Remember that the basis is the division of the variables into the basic and nonbasic sets. Now because the nonbasic variables identify the constraints that are active, this amounts to saying the



following: if you do not change a coefficient more than the amounts in the “allowable increase” or “allowable decrease”, then the optimum will still be at the intersection of the same constraints. But note this carefully: *even if the basis does not change, the point in space $[(x_1, x_2)]$ for Acme may change!* This also means that the optimum value of the objective function may change too. Figure 6.1 shows how this can happen.

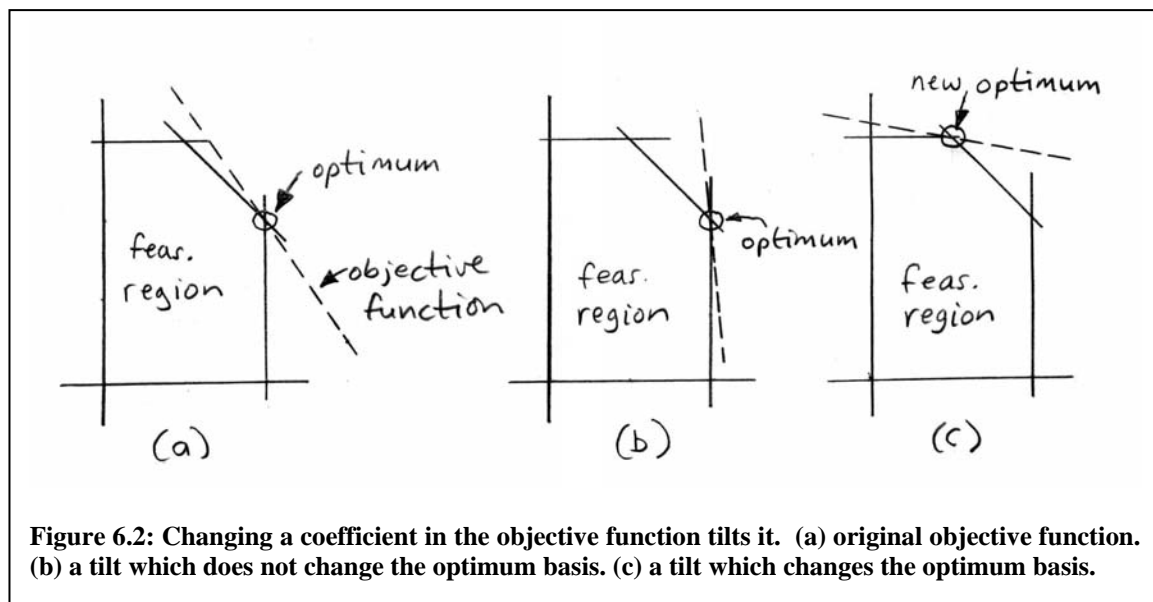
So what is the advantage of staying at the same optimum basis if both the point and the objective function value might change? The advantage lies in the fact that you can still use the solution that you already have to determine what the new point and Z will be after the coefficients have changed, as we will see below. You do not need to re-solve the problem from scratch, which can be a huge advantage for very large problems.

Now let's look at changing the various coefficients as permitted in simple computer-based ranging.

Changing Objective Function Coefficients

Here is the first important observation about changing the coefficients of the objective function: this does not affect the feasible region! For this reason, the optimum point that you found during the original solution of the problem will remain a feasible cornerpoint. The worst that can happen is that the original optimum cornerpoint will no longer be optimum after the objective function is changed. In that case, you can probably restart the solution at the original optimum cornerpoint and continue iterating until you reach the new optimum. The simple ranging analysis will tell you whether you have to do this or not.

Figure 6.2 shows that the effect of changing a coefficient of the objective function is to tilt it. What the allowable range tells you is essentially the maximum tilt in any direction before the optimum moves to a different cornerpoint. For example, looking at the LINDO ranging data above, the objective function coefficient of x_2 can increase from 10 to as much as 15 before the optimum moves to a different basis. The objective function coefficient of x_2 can also decrease from 10 to as little as zero before the optimum moves to a different basis.



If you are investigating a change to an objective function coefficient, just compare the change to the ranging table. If the change is more than allowed, then you have no option

but to re-solve the problem to find the new optimum cornerpoint. Most commercial solvers allow you to re-start from any basis, so if you have previously saved the optimum basis, then restart there since the new optimum is likely to be only a few pivots away. Or you can restart the model from scratch in the worst case.

But what if the proposed change to the objective function coefficient is within the allowable range? Then how do you determine the new values of Z and the variables? This is straightforward:

- The new optimum variable values will be the same as in the original solution because the constraints are not changed by a modification to the objective function, and we are at the same basis, so we will be at the same place in space.
- You can calculate the new value of the objective function easily: just substitute the old values of the variables into the new objective function. Alternatively, since only one objective function coefficient at a time is changed, you can calculate the change in the objective function value by calculating the difference due to the changed coefficient.

For example, is the Acme Bicycle Company solution sensitive if the objective function coefficient of x_1 is changed from 15 to 25? As you can see from the tables above, the allowable increase for the objective function coefficient for x_1 is infinity, so the basis will remain the same. The optimum solution will still be $(x_1, x_2) = (2, 2)$, but the new Z will be given by $25x_1 + 10x_2 = 25 \times 2 + 10 \times 2 = 70$. You could also calculate the new Z by looking only at the changed coefficient: the change in Z will be $50 + 2 \times [(new\ coefficient) - (old\ coefficient)] = 50 + 2 \times (25 - 15) = 70$.

What about changing the objective function coefficient of x_2 from 10 to 25? As you can see from the tables, this is a change of 15 units, which is greater than the allowable increase of 5. This means that the basis will change, so the only way to find out what the variable values and Z will be at the new optimum is to restart the solution and continue iterating until you reach the new optimum point.

Changing a Right Hand Side Constant

Right hand side constants normally represent a limitation on a resource, and are likely to change in practice as business conditions change. The allowable increase and allowable decrease listed in the ranging tables again show by how much a right hand side can change before the basis changes. Look again at Figure 6.1 to see how the variable values and Z might change at the optimum even when the basis does not change. On the other hand, these things might *not* change if a change is made to a constraint that is not active at the current optimum. Why? Because the change may not affect the optimum point at all, as shown in Figure

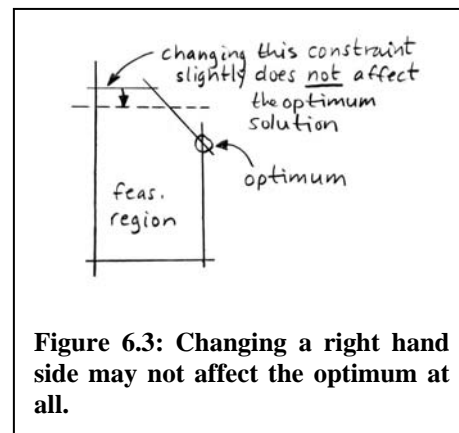



Figure 6.3: Changing a right hand side may not affect the optimum at all.

6.3.

Figure 6.4, a picture of our favourite Acme Bicycle Company problem, shows more clearly why there is a limited range in which the right hand side can change before the basis changes. Note that a change to the right hand side is the same as a parallel shift of the constraint. In Figure 6.4, the basis remains the same (intersection of 1st and 3rd constraints) while the right hand side of the 1st constraint is changed, until it has been shifted enough that a different constraint becomes active, so the basis changes. As the right hand side of the 1st constraint is gradually increased, the basis does not change until the right hand side reaches a value of 4 (the 1st constraint is now $x_1 \leq 4$). If the right hand side moves beyond 4, then the nonnegativity constraint on x_2 becomes active, and this defines a different basis (3rd constraint and x_2 nonnegativity active). If the right hand side is gradually decreased, then the basis does not change until the right hand side reaches a value of 1 (1st constraint is now $x_1 \leq 1$). Beyond here, the 2nd constraint ($x_2 \leq 3$) becomes active, and the basis is defined by the intersection of the 2nd and 3rd constraints.



Again, you can see in Figure 6.4 how the Z and variable values may change even though the basis is the same, as long as we propose changes to the coefficients that are within the allowable ranges.

Here is the overall procedure for examining proposed changes to the right hand sides of constraints. First check whether the proposed change is within the allowable range of changes for the right hand side of the constraint. If it is not within the allowable range, then you must re-start the solution process and iterate to the new optimum basis and solution. If the proposed change is within the allowable range, then the values of Z and the variables at the optimum are recovered as follows:

- To calculate the new value of Z :
 - If it is a maximization problem: $Z_{\text{new}} = Z_{\text{original}} + (\text{dual price}) \times \Delta b$, where Δb represents the change in the right hand side value ($b_{\text{new}} - b_{\text{original}}$).
 - If it is a minimization problem: $Z_{\text{new}} = Z_{\text{original}} - (\text{dual price}) \times \Delta b$.
- Getting the values of the variables is harder. If the solver allows you to recover the “basis inverse” matrix, denoted by \mathbf{B}^{-1} , then you can calculate the matrix product $\mathbf{B}^{-1} \mathbf{b}_{\text{new}}$, where \mathbf{b}_{new} denotes the new right hand side vector. This will give the values of the basic variables. Or, knowing the basis, you can solve the matrix obtained by setting the nonbasic variables to zero.

This method of finding the new Z works because the dual price is the objective function coefficient for the slack variable associated with the changed constraint. It thus gives the

change in Z per unit change in the slack variable for the constraint, and the change in the slack variable is identical to any change in the right hand side.

As an example, let us consider whether the Acme Bicycle Company solution is sensitive to a change in the availability of the metal finishing machine from 4 hours per day to 3.5 hours per day (this is a change to the right hand side of the 3rd constraint, row 4 in the LINDO tables). The tables show that this decrease of $\frac{1}{2}$ unit in the right hand side of the 3rd constraint is within the allowable range. Hence the basis will remain unchanged. The objective function value in this maximization problem will change as follows: $Z_{\text{new}} = 50 + 10 \times (3.5 - 4) = 45$. We will not recover the new values of the variables. So the solution is sensitive to the proposed change in the availability of the metal finishing machine.

Is the Acme Bicycle Company solution sensitive to a change in the right hand side of the 2nd constraint from $x_2 \leq 3$ to $x_2 \leq 4$? From the tables, this proposed change is within the allowable range, and the constraint dual price is zero (constraint is not active), so no, the solution is not sensitive to this change: the Z and variable values will be unaffected.

Buying Extra Resources

As mentioned above, the right hand side often represents the amount of a given resource that is available. Given that you are currently limited to a maximum amount b of some resource, how much would you willingly pay for an extra unit of b if you could get it? This question often comes up when you observe that any possible improvement in your objective function value is being limited by the availability of some resource. For example, suppose that the available pool of labour is limiting your increase in profits: how much would you pay per unit of extra labour (e.g. by paying overtime, or by hiring temporary workers)? If extra labour can be had cheaply, then you may be able to improve your optimum Z considerably, but if extra labour is expensive, then paying for it may in fact worsen your optimum Z . The idea is to determine the upper limit on how much to pay before you have a worsening effect on the objective function value.

Suppose that the Acme Bicycle Company could rent time on another metal finishing machine for \$8 per day: would this be a good idea? Here is how the reasoning goes:

- Acme currently pays nothing per hour for the metal finishing machine because they own it.
- The dual price of the metal finishing constraint is \$10 from the LINDO solution. This means that each unit increase in the right hand side of the metal finishing machine constraint (3rd constraint) increases Z by \$10.
- Hence, we could pay up to \$10 for an extra hour of metal finishing machine capacity and still increase our overall profit.

So, yes, we would happily pay \$8 for another hour's worth of use of a metal finishing machine.

The general rule is that you would pay up to (price you are paying now for the resource)+(dual price) for each extra unit of a resource. Note, though, that this analysis

only holds if the amount of extra resource that you are planning to buy remains within the allowable range as shown in the tables. Why? Because buying extra resources is the same as changing a right hand side coefficient. If you buy too much of the extra resource, then the basis changes and you are no longer sure if the purchase at the given price is a good idea.

For example, consider buying extra units of capacity for making mountain bikes (i.e. a change to the right hand side of the 1st constraint). In the LINDO tables above, we see that the dual price of the 1st constraint (row 2 in the table) is \$5 per unit. Thus if offered the chance to buy extra units at a price of \$3 per unit, Acme would accept. But how many units should they buy? The ranging tables show that the maximum increase in the right hand side of the 1st constraint is just two units before the basis changes. You can see why in Figure 6.4.

Suppose that Acme was offered extra units of racer production at a bargain rate of \$1 per unit. Should they accept? This is equivalent to a change in the right hand side of the 2nd constraint, and the tables above show that the dual price of the 2nd constraint is 0. Hence Acme should *not* buy any extra units of this resource. Look again at Figure 6.3 to see why. The 2nd constraint is not limiting production now, so a slight increase or decrease in its right hand side value has no impact on the optimum solution. There is no point in buying any extra racer production capacity.

A Note on Terminology

As I've defined it here, a model is "sensitive" if a proposed change to an original coefficient causes a change to the optimum objective function value or variable values. These are the outcomes that materially affect the use of the solution in practice. However, to some authors, a model is "sensitive" only if the basis will be changed by the proposed change to a coefficient. As we have seen, the objective function value and the variable values can both change even when the basis does *not* change.

So beware of miscommunication when talking with others about the sensitivity of a linear program!