Sensitivity Analysis: A Sample LINDO Output

Consider the linear program:

$$
\begin{array}{llllllll}
\text{Minimize} & z = & 84x_1 & +72x_2 & +60x_3 & & & \\
\text{Subject to:} & & & & & & & \\
 & & 90x_1 & +20x_2 & +40x_3 & \geq & 200 & (1) \\
 & & 30x_1 & +80x_2 & +60x_3 & \geq & 180 & (2) \\
 & & 10x_1 & +20x_2 & +60x_3 & \geq & 150 & (3)
\end{array}
$$

$$x_1,\, x_2,\, x_3 \geq 0\,.$$

This problem comes from Exercise 3.6-4. Briefly, the scenario that gives rise to this formulation is as follows. A farmer wishes to determine the quantities of the available types of feed, which are corn, tankage, and alfalfa, that should be given to each pig. Since pigs will eat any mix of these feed types, the objective is to determine which mix will meet a given set of basic nutritional requirements at a minimum cost. The decision variables are:

x_1 = kilograms of corn,

x_2 = kilograms of tankage, and

x_3 = kilograms of alfalfa.

The basic nutritional ingredients are: carbohydrates, protein, and vitamins; and their respective requirements are expressed in constraints (1)–(3) above. Finally, the costs per kilogram of feed types, 84, 72, and 60, are in cents.

Now, launch LINDO and enter the problem as:

```
MIN 84 X1 + 72 X2 + 60 X3
ST
90 X1 + 20 X2 + 40 X3 > 200
30 X1 + 80 X2 + 60 X3 > 180
10 X1 + 20 X2 + 60 X3 > 150
END
```

where X1, X2, and X3 correspond, respectively, to x_1, x_2, and x_3.

Next, click on the “Solve” menu and then select “Solve”. Click on “Yes” in the “Do Range (Sensitivity) Analysis?” dialog box. The following report should now be displayed in the “Reports Window”:

```
LP OPTIMUM FOUND AT STEP      2

        OBJECTIVE FUNCTION VALUE

        1)      241.7143

  VARIABLE        VALUE          REDUCED COST
        X1         1.142857          0.000000
        X2         0.000000         17.714285
        X3         2.428571          0.000000


       ROW   SLACK OR SURPLUS     DUAL PRICES
        2)         0.000000         -0.771429
        3)         0.000000         -0.485714
        4)         7.142857          0.000000

 NO. ITERATIONS=       2


 RANGES IN WHICH THE BASIS IS UNCHANGED:

                           OBJ COEFFICIENT RANGES
 VARIABLE         CURRENT        ALLOWABLE        ALLOWABLE
                   COEF          INCREASE         DECREASE
       X1        84.000000       51.000000        37.199997
       X2        72.000000        INFINITY        17.714285
       X3        60.000000       11.272726        22.666666

                           RIGHTHAND SIDE RANGES
      ROW         CURRENT        ALLOWABLE        ALLOWABLE
                    RHS          INCREASE         DECREASE
        2       200.000000       24.999998        80.000000
        3       180.000000      120.000000         6.000000
        4       150.000000        7.142857         INFINITY
```

We will interpret this output from top to bottom.

The Simplex algorithm went through two iterations to produce the optimal solution. The optimal objective-function value is 241.7143 (cents). LINDO automatically introduces slack or surplus variables to convert inequality constraints into equalities. The numbering of the rows starts with the number 1, as opposed to 0. Thus, ROW 1 refers to the objective-function row, ROW 2 refers to (functional) constraint (1), and so on; correspondingly, the slack or surplus variables will also be numbered this way. For this problem, three surplus variables are introduced, and they are named SLK 2, SLK 3, and SLK 4 (LINDO uses SLK to denote both slack and surplus variables). In our standard notation, these correspond to s_1, s_2, and s_3; therefore, the optimal solution is $(x_1, x_2, x_3, s_1, s_2, s_3) = (1.142857, 0, 2.428571, 0, 0, 7.142857)$.

To explain the columns under REDUCED COST and DUAL PRICES, we need to refer to the final tableau. To generate the final tableau, click on the title bar at the top of the problem window (to return focus to that window), click on the "Reports" menu, and then click "Tableau". The final tableau now appears at the bottom of the previous "Reports Window". It is pasted below.

THE TABLEAU

ROW	(BASIS)	X1	X2	X3	SLK 2	SLK 3
1	ART	0.000	17.714	0.000	0.771	0.486
2	X3	0.000	1.571	1.000	0.007	-0.021
3	X1	1.000	-0.476	0.000	-0.014	0.010
4	SLK 4	0.000	69.524	0.000	0.286	-1.190

ROW	SLK 4	
1	0.000	-241.714
2	0.000	2.429
3	0.000	1.143
4	1.000	7.143

This tableau is in our standard format, except that it overflows into a second line. Here, the variable ART refers to z. (Recall that z was artificially created for convenience.) The basic variables are X3 (for ROW 2), X1 (for ROW 3), and SLK 4 (for ROW 4); and they assume the corresponding values listed in the RHS column. These values are consistent (modulo rounding differences) with the optimal solution reported above. Note, however, that the RHS constant in ROW 1 is listed as -241.714. The negative sign here is the consequence of converting a minimization problem into a maximization problem, which LINDO does automatically. (Working with maximization problems only simplifies interpretation.) Therefore, this negative sign should be reversed in the final report, which, again, the program does automatically. In ROW 1, the coefficients of X2, SLK 2, and SLK 3, the nonbasic variables, are all positive; this indicates that what we have is indeed the (unique) final tableau.

We now return to the report above. The values listed in the REDUCED COST column are taken from the coefficients of X1, X2, and X3 in ROW 1, in the final tableau. In other words, X1 and X3 have a reduced cost of 0, whereas X2 has a reduced cost of 17.714. Formally (for a maximization problem), the *reduced cost* for a nonbasic variable is defined as the amount by which the value of z will decrease if we increase the value of that nonbasic variable by 1 (while holding all other nonbasic variables at 0). The adjective "reduced" is used because such a cost is relative to a specific tableau, i.e., is from the viewpoint of the particular current basic feasible solution. The 1-unit increment in the nonbasic variable is nominal, in that we are only contemplating an increase, even when it is not feasible to do so. The reduced cost for a basic variable is defined as 0. Mechanically, this is because basic variables always have a coefficient of 0 in the objective-function row; and conceptually, this is because the basic variables are already "participating" in the current solution (and therefore we do not attempt to bring them into the basis).

Thus, the final tableau tells us that it is not optimal to include any tankage in the mix. Moreover, if we insist on having tankage in the mix, then the cost per kilogram of tankage is 17.714 cents. Like the concept of shadow price, this cost is relative to our current optimal solution, i.e., it has nothing to do with the "market" cost of tankage (which is at 72 cents per kilogram).

Next, we move on to the values listed in the DUAL PRICES column. The term "dual prices" is equivalent to shadow prices. (Every linear program has an associated dual linear program, and the concept of dual prices originates from the dual linear program. We will not discuss the dual of a linear program, as it is a more-advanced topic.) Recall (from parts (a) and (b) in the previous example) that the shadow price associated with the RHS constant of an original constraint (or with the availability of a resource) is defined as the amount by which the optimal objective-function value will improve if we increase the value of that constant by 1. (Again, this 1-unit increment is nominal.) In the current problem, the original functional constraints are of the "$\geq$" type. For such a constraint, an increase in the RHS constant corresponds to a tightening of that constraint; hence, the increase will (typically) result in a degradation of the optimal objective-function value. Indeed, the reported dual prices for ROW 2 and ROW 3 are negative.

Thus, for every unit of increase in the nutritional requirement for carbohydrates, the cost of the *optimal* mix will increase by 0.771429 cents. Similarly, the corresponding increase in cost associated with protein is 0.485714. That the dual price for vitamins equals 0 is a consequence of the fact that the optimal mix already exceeds the vitamins requirement by a margin of 7.142857 (that is, the surplus variable SLK 4 equals 7.142857 in the optimal solution).

Interestingly, the reported dual prices are the negatives of the coefficients of SLK 2, SLK 3, and SLK 4 in ROW 1, in the final tableau. This is not a coincidence, and it can be attributed to the fact that the coefficients in the surplus-variable columns are negatives of those in the artificial-variable columns, which are introduced as part of Phase I of the solution process. We will not pour over the somewhat technical details.

Next, we examine the RANGES IN WHICH THE BASIS IS UNCHANGED. Here, THE BASIS refers to the optimal basis; and two sets of ranges are displayed, one for the original objective-function coefficients and one for the original RHS constants.

Consider the range for X2 first. Recall (from part (d) in the previous example) that if the objective-function coefficient of a variable is revised to a new value, then the revision is directly reflected in the coefficient of that variable in ROW 1 of the final tableau. In the example here, this means that if the cost of tankage is revised by δ, to $72 + \delta$, then the coefficient of X2 in ROW 1 of the final tableau will be revised to $17.714+\delta$. Moreover, since X2 is nonbasic, no further revision in the final tableau is necessary. It follows that δ has to be less than -17.714 to make the new coefficient negative, i.e., to make it desirable to include tankage in the mix. Indeed, LINDO reports that the objective-function coefficient of X2 should stay inside the interval $[72 - 17.714285, 72 + \infty)$, or $[54.285715, \infty)$, in order for the current basis, and solution, to remain optimal

Interestingly, the maximum possible decrease of 17.714 is identical to the reduced cost of X2. That these two values are in agreement is expected, as they are derived from two (different-sounding but) equivalent viewpoints.

Next, suppose the cost of corn is revised by δ, from 84 to $84 + \delta$. A similar argument shows that this will imply a corresponding revision of the coefficient of X1 in ROW 1 to $0 + \delta$. Since X1 is basic, doing so will disqualify X1 as the basic variable for ROW 3. Hence, as in part (f) of the previous example, one needs to execute a pivot in the X1-column to eliminate the new entry δ (assuming it is nonzero). This pivot will result in a new set of coefficients in ROW 1; and therefore, we need to limit the scope of δ to ensure that these new coefficients remain nonnegative (i.e., to ensure that the current basis remains optimal). A little bit of calculations (which we leave out) now shows that δ should stay within the interval $[-37.19997, 51.000]$; and this corresponds to the LINDO statement for the range of the objective-function coefficient of X1.

The specified range for the objective-function coefficient of X3 is interpreted similarly: If the cost of alfalfa is revised from 60 to $60 + \delta$, then δ should stay within the interval $[-22.666666, 11.272726]$ in order for the current tableau (and solution) to remain optimal.

Finally, we examine the ranges for the original RHS constants. For ROW 2, LINDO reports that if the RHS constant 200, which is the carbohydrates requirement, is revised to $200 + \delta$, then the current basis will remain optimal for all δ inside the interval $[-80.000000, 24.999998]$. Similarly, the scopes of δ for ROW 3 and ROW 4, or for protein and vitamins, are reported as $[-6.000000, 120.000000]$ and $[-\infty, 7.142857]$, respectively. The supporting calculations for these statements are similar to what we did in parts (a) and (b) in the previous example.

It is important to realize, however, that although the optimal basis stays the same as long as δ is within these individual ranges, the specific *values* of the optimal basic variables do depend on δ. Again, the details are similar to parts (a) and (b) in the previous example.

From a practical viewpoint, the reported ranges for both the objective-function coefficients and the RHS constants are rather wide. This "insensitivity" of the optimal basis with respect to changes in the input parameters is remarkable.