

LP Methods.S2

Sensitivity Analysis

Generally speaking, the basic assumption that all the coefficients of a linear programming model are known with certainty rarely holds in practice. Moreover, it may be expedient to simplify causal relationships and to omit certain variables or constraints at the beginning of the analysis to achieve tractability. We have already explained how the dual simplex method can be used to reoptimize a model when a new constraint is added to the formulation. Adding a new variable can also be handled efficiently by simply pricing out the new column and seeing if its reduced cost is nonnegative. If so, the new variable has no effect on the optimum; if not, it becomes the entering variable and the algorithm continues until all reduced costs are nonnegative.

In this section, we deal implicitly with the issue of uncertainty in the data elements c_j , a_{ij} and b_i , by determining the bounds over which each such element can range without effecting a change in either the optimal solution or optimal basis. Such investigations fall under the heading of *sensitivity analysis*. For the most part, we deal with the simple case of perturbing one coefficient at a time. This allows us to derive closed-form solutions. When two coefficients are varied simultaneously, the analysis is much more complex because the “range” is described by a 2-dimensional polyhedron rather than an interval on the real line. When k coefficients are varied simultaneously, a k -dimensional polyhedron results. The only exception is *proportional ranging* which allows all elements of the original vector to vary simultaneously but in fixed proportion as defined by a second vector. This topic falls under the more general heading of *post-optimality analysis*.

Sensitivity to Variation in the Right-Hand Side

We have seen that for every basis \mathbf{B} associated with an LP, there is a corresponding set of m dual variables π , one for each row. The optimal values of the dual variables can be interpreted as prices. In this section, this interpretation is explored in further detail starting with an LP in standard equality form.

$$\text{maximize}\{\mathbf{c}\mathbf{x} : \mathbf{A}\mathbf{x} = \mathbf{b}, \mathbf{x} \geq \mathbf{0}\}$$

Suppose the optimal basis is \mathbf{B} with solution $(\mathbf{x}_B, \mathbf{0})$, where $\mathbf{x}_B = \mathbf{B}^{-1}\mathbf{b}$ and $\pi = \mathbf{c}_B\mathbf{B}^{-1}$ is unrestricted in sign. Now, assuming nondegeneracy, small changes in the vector \mathbf{b} will not cause the optimal basis to change. Thus for $\mathbf{b} + \Delta\mathbf{b}$ the optimal solution is

$$\mathbf{x} = (\mathbf{x}_B + \Delta\mathbf{x}_B, \mathbf{0})$$

where $\Delta\mathbf{x}_B = \mathbf{B}^{-1} \Delta\mathbf{b}$. Thus the corresponding increment in the objective function is

$$z = \mathbf{c}_B \mathbf{x}_B = \boldsymbol{\pi} \mathbf{b}.$$

This equation shows that $\boldsymbol{\pi}$ gives the sensitivity of the optimal payoff with respect to small changes in the vector \mathbf{b} . In other words, if a new problem were solved with \mathbf{b} changed to $\mathbf{b} + \Delta \mathbf{b}$, the change in the optimal value of the objective function would be $\boldsymbol{\pi} \Delta \mathbf{b}$.

For a maximization problem, this interpretation says that π_i directly reflects the change in profit due to a change in the i th component of the vector \mathbf{b} . Thus π_i may be viewed equivalently as the *marginal price* of b_i , since if b_i is changed to $b_i + \Delta b_i$, the value of the optimal solution changes by $\pi_i \Delta b_i$. When the constraints $\mathbf{Ax} = \mathbf{b}$ are written as $\mathbf{Ax} \leq \mathbf{b}$, the dual variables are nonnegative implying that for π_i positive, a positive change in b_i will produce an increase in the objective function value. In economic terms, it is common to refer to the dual variables as *shadow prices*.

Example 1

The shadow prices are associated with constraints but they are often used to evaluate prices or cost coefficients associated with variables of the primal problem. As an example, suppose we have an \mathbf{A} matrix representing the daily operation of an oil refinery, and a particular variable x_j representing the purchase of crude oil feedstock, with a cost of \$22.65/barrel ($c_j = 22.65$). The refinery wants to minimize its costs. There is an upper limit on the purchase of this oil of 50,000 barrels/day at this price. This is represented by the constraint

$$x_j + x_s = 50,000$$

where x_s is the associated slack variable. Assume at the optimum x_s has a reduced cost of $-\$2.17$ /barrel: what does this mean?

As we shall see, the shadow price on the constraint is also $-\$2.17$ /barrel. This does not mean, however, that we should only pay \$2.17 for another barrel of crude. It means we should be prepared to pay another \$2.17/barrel for an opportunity to purchase extra supplies *given* that any further purchases would cost \$22.65/barrel; i.e., the objective function will decrease by \$2.17 for every extra barrel we can purchase at the price c_j already in the cost row. This means we should be prepared to bid up to $22.65 + 2.17 = \$25.82$ /barrel on the spot market for extra supplies of that crude. Note that \$25.82/barrel is the *breakeven price*, in that we decrease our objective function z if we can purchase a barrel for less than this price, increase z if we purchase for more, and make no change at all to z if we purchase for exactly \$25.82/barrel.

Reduced Cost

The reduced cost of a nonbasic variable at its lower bound is often referred to as the *opportunity cost* of that variable. If management made the (nonoptimal) decision of increasing that nonbasic variable from its lower bound, for a maximization problem the reduced cost gives the decrease in z per unit increase in the variable (for a certain range). This represents the opportunity loss in departing from the optimal solution.

Ranging

For reasons that practitioners understand implicitly, it is often said that postoptimality analysis is the most important part of the LP calculations. The majority of the coefficients that appear in an LP are rarely known with certainty and so have to be estimated from historical or empirical data. Under these circumstances we would like to know the range of variation of these coefficients for which the optimal solution remains optimal; i.e., the basis does not change. Three categories are investigated below: objective coefficients c_j , right-hand-side terms b_i , and matrix coefficients a_{ij} .

Changes in the objective row

(a) *Nonbasic variable* The change in the objective coefficient of a nonbasic variable affects the reduced cost of that variable only, and the change is in direct proportion. If Δc_q is a perturbation associated with the original objective coefficient c_q , for some $q \in Q$, then at optimality we can write the reduced cost coefficient of nonbasic variable x_q as $\bar{c}_q(\Delta c_q) = \pi \mathbf{A}_q - (c_q + \Delta c_q)$. In order for the current basis \mathbf{B} to remain optimal, we must have $\bar{c}_q(\Delta c_q) \geq 0$. This means

$$\pi \mathbf{A}_q - c_q = \bar{c}_q$$

Not surprisingly, there is no lower bound on Δc_q . Reducing the value of an objective coefficient associated with a nonbasic variable cannot make the variable more attractive.

The reduced costs of all the other variables are independent of c_q and so will remain nonnegative. If a Δc_q is chosen that violates the above inequality, x_q would be identified as the entering variable and we would continue the application of the simplex method until an optimal basis for the modified problem was found.

It is worth mentioning that in most commercial LP codes there is a second range given as well — the range over which x_q can be increased from zero before a change of basis occurs. When $\bar{c}_q = \bar{c}_q$, the reduced cost is zero implying that x_q can be increased without affecting the value of the objective function; alternatively, it would imply that there are multiple optima. The maximum value x_q can take without effecting a change in basis is given by $\min_i \bar{b}_i / \bar{a}_{iq} : \bar{a}_{iq} > 0$ which is the minimum ratio test in Step 4 of the revised simplex algorithm.

(b) *Basic variable* A change in the objective coefficient of a basic variable may affect the reduced cost of all the nonbasic variables. Let \mathbf{e}_i be the i th unit vector of length m and suppose we increment the objective coefficient of the i th basic variable $x_{B(i)}$ by Δ ; i.e., $\mathbf{c}_B \leftarrow \mathbf{c}_B + \Delta \mathbf{e}_i^T$. This gives $\pi(\cdot) = (\mathbf{c}_B + \Delta \mathbf{e}_i^T) \mathbf{B}^{-1}$ so the dual vector is an affine function of Δ . The reduced cost of the q th nonbasic variable is now

$$\begin{aligned} \bar{c}_q(\Delta) &= (\mathbf{c}_B + \Delta \mathbf{e}_i^T) \mathbf{B}^{-1} \mathbf{A}_q - c_q \\ &= \mathbf{c}_B \mathbf{B}^{-1} \mathbf{A}_q + \Delta \mathbf{e}_i^T \mathbf{B}^{-1} \mathbf{A}_q - c_q \\ &= \bar{c}_q + \Delta \bar{a}_{iq} \end{aligned}$$

where $\bar{a}_{iq} = (\mathbf{B}^{-1} \mathbf{A}_q)_i$ is the i th component of the updated column of \mathbf{A}_q . This value can be found for the nonbasic variable x_q by solving $\mathbf{B}^T \mathbf{y} = \mathbf{e}_i$ for \mathbf{y} , then computing $\bar{a}_{iq} = \mathbf{y}^T \mathbf{A}_q$. (Obviously, if $\bar{a}_{iq} = 0$ for any x_q , the reduced cost does not change.)

For a solution to remain optimal, we must have $\bar{c}_q(\Delta) \geq 0$, or

$$\bar{c}_q + \Delta \bar{a}_{iq} \geq 0 \text{ for all } q \in Q \quad (1)$$

where \bar{c}_q is the reduced cost at the current optimum. This constraint produces bounds on Δ . For a basic variable, the *range* over which c_i can vary and the current solution remain optimal is given by $c_i + \Delta$, where

$$\max_{q \in Q} \frac{-\bar{c}_q}{\bar{a}_{iq}} : \bar{a}_{iq} > 0 \qquad \min_{q \in Q} \frac{-\bar{c}_q}{\bar{a}_{iq}} : \bar{a}_{iq} < 0$$

since this is the range for which (1) is satisfied. If there is no $\bar{a}_{iq} > 0$, then $\bar{c}_q < 0$; likewise if there is no $\bar{a}_{ij} < 0$, then $\bar{c}_q > 0$.

Note that perturbing the value c_i for the i th basic variable has no effect on the reduced costs of any of the basic variables. All reduced costs will remain zero. This can be seen from the definition $\bar{\mathbf{c}}_B = \mathbf{c}_B \mathbf{B}^{-1} \mathbf{B} - \mathbf{c}_B = \mathbf{0}$ at optimality so if any of the components of \mathbf{c}_B are perturbed by δ , the effect cancels itself out.

Example 2

Suppose we have an optimal solution to an LP given in tableau form with attached variables

$$\begin{aligned} \text{Maximize } z &= 4.9 - 0.1x_3 - 2.5x_4 - 0.2x_5 \\ \text{subject to } x_1 &= 3.2 - 0.5x_3 - 1.0x_4 - 0.6x_5 \\ x_2 &= 1.5 + 1.0x_3 + 0.5x_4 - 1.0x_5 \\ x_6 &= 5.6 - 0.5x_3 - 2.0x_4 - 1.0x_5 \end{aligned}$$

The set of nonbasic variables $Q = \{3, 4, 5\}$ so the current basis remains optimal as long as $\bar{c}_q \leq 0$ for all $q \in Q$. When $q = 3$, for instance, this means $\bar{c}_3 \leq 0.1$. If the original coefficient $c_3 = 1$, then the current basis remains optimal for $c_3 \leq 1.1$.

If the objective coefficient of the basic variable x_2 becomes $c_2 + \delta$, the reduced costs of the nonbasic variables become

$$x_3: \bar{c}_3(\delta) = 0.1 + (-1.0)\delta$$

$$x_4: \bar{c}_4(\delta) = 2.5 + (-0.5)\delta$$

$$x_5: \bar{c}_5(\delta) = 0.2 + (+1.0)\delta$$

Note that $x_{B(i)} = \bar{b}_i - \sum_{j \in \{3,4,5\}} \bar{a}_{ij} x_j$ for $i = 1, 2, 6$, so \bar{a}_{ij} is the negative of the

number appearing in the above equations. The range that δ can take is given by

$$\max \frac{-0.2}{1.0} \quad \min \frac{-2.5, -0.1}{-0.5 \quad -1.0}$$

$$-0.2 \quad 0.1$$

When c_2 assumes one of the limits of its range, a reduced cost becomes zero. In this example, for $c_2 = 0.1$ the reduced cost of x_3 is zero, so that if the objective coefficient of x_2 increases by more than 0.1 it becomes advantageous for x_3 to become active. The minimum ratio test: $\min \frac{3.2}{0.5}, \frac{5.6}{0.5} = 6.4$, indicates that x_3 can be increased to 6.4 before x_1 becomes zero and a change of basis is required. Analogously, for $c_2 = -0.2$, the reduced cost for x_5 is zero and any further decrease in c_2 will require a basis change to remain optimal. In this case, the ratio test indicates that x_2 would be the leaving variable.

The above analysis can be generalized without too much difficulty to allow proportional changes in the vector \mathbf{c} rather than changes in only one coefficient at a time. To perform proportional ranging, we must stipulate an n -dimensional row vector \mathbf{c}^* and consider the new vector $\mathbf{c}(\lambda) = \mathbf{c} + \lambda \mathbf{c}^*$. The analysis would proceed as in part (b) above but with \mathbf{e}_i replaced by \mathbf{c}^* (see Exercise 3).

Changes in the right-hand-side vector

We wish to investigate the effect of a change $b_i \leftarrow b_i + \Delta b_i$ for some $1 \leq i \leq m$. It is usual to consider the case where b_i is the right-hand side of an *inequality* constraint, which therefore has a slack variable associated with it. The goal is to determine the range over which the current solution remains optimal. If the constraint is an equality, it can be analyzed by regarding its associated artificial variable as a positive slack (which must be nonbasic for a feasible solution).

(a) *Basic slack variable* If the slack variable associated with the i th constraint is basic the constraint is not binding at the optimum. The analysis is simple: the value of the slack gives the range over which the right-hand-side b_i can be reduced for a \leq constraint or increased for a \geq constraint. The solution remains feasible and optimal for the range $b_i + \Delta b_i$, where

$$-\hat{x}_s < \Delta b_i < \hat{x}_s \quad \text{for } \leq \text{ type constraint}$$

$$- < \hat{x}_s \text{ for } \text{type constraint}$$

where \hat{x}_s is the value of the associated slack variable.

(b) *Nonbasic slack variable* If a slack variable is nonbasic at zero, then the original inequality constraint is binding at the optimum. At first glance it would seem that because the constraint is binding, there is no possibility of changing the right-hand-side term, particularly in *decreasing* the value of b_i (for \leq type constraints). It turns out that by changing the

vector \mathbf{b} we also change \mathbf{x}_B ($= \mathbf{B}^{-1}\mathbf{b} = \bar{\mathbf{b}}$) so there is a range over which \mathbf{x}_B remains nonnegative. For the associated values we still retain an optimal feasible solution in the sense that the basis does not change. (Note that both \mathbf{x}_B and $z = \mathbf{c}_B\mathbf{x}_B$ change value.)

Consider the k th constraint

$$a_{k1}x_1 + a_{k2}x_2 + \dots + a_{kn}x_n + x_s = b_k$$

where x_s is the slack variable. If the right-hand side becomes $b_k + \Delta$, rearranging this equation gives

$$a_{k1}x_1 + a_{k2}x_2 + \dots + a_{kn}x_n + (x_s - \Delta) = b_k \quad (2)$$

so that $(x_s - \Delta)$ replaces x_s . Thus if x_s is nonbasic at zero in the final tableau, we have the expression

$$\mathbf{x}_B = \bar{\mathbf{b}} - \bar{\mathbf{A}}_s(-\Delta)$$

where $\bar{\mathbf{A}}_s$ is the updated column in the tableau corresponding to x_s .

Because \mathbf{x}_B must remain nonnegative, we have $\bar{\mathbf{b}} + \bar{\mathbf{A}}_s\Delta \geq \mathbf{0}$ which is used to solve for the range over which Δ can vary.

$$\max_i \frac{\bar{b}_i}{-\bar{a}_{is}} : \bar{a}_{is} > 0 \qquad \min_i \frac{\bar{b}_i}{-\bar{a}_{is}} : \bar{a}_{is} < 0$$

If there is no $\bar{a}_{is} > 0$, then $\Delta > -\infty$; if there is no $\bar{a}_{is} < 0$, then $\Delta < \infty$.

For \leq type constraints, Δ changes sign. This follows because we can analyze $\sum_{j=1}^n a_{ij}x_j \leq b_i$ in the form $-\sum_{j=1}^n a_{ij}x_j \geq -b_i$, so that $-(x_s + \Delta)$ replaces $(x_s - \Delta)$ in Eq. (2). Another way of seeing this is to consider the change to the right-hand side in the form

$$\mathbf{b}(\delta) = \mathbf{b} + \delta \mathbf{e}_k$$

Thus the new value of \mathbf{x}_B is given by

$$\begin{aligned} \mathbf{x}_B(\delta) &= \mathbf{B}^{-1}\mathbf{b}(\delta) = \mathbf{B}^{-1}\mathbf{b} + \mathbf{B}^{-1}\delta\mathbf{e}_k \\ &= \bar{\mathbf{b}} + \delta \mathbf{B}^{-1}\mathbf{e}_k \end{aligned}$$

However, $\bar{\mathbf{A}}_s = \mathbf{B}^{-1}\mathbf{e}_k$ for a \leq type constraint

and $\bar{\mathbf{A}}_s = -\mathbf{B}^{-1}\mathbf{e}_k$ for a \geq type constraint

since the column corresponding to the slack variable is $+\mathbf{e}_k$ for a \leq constraint and $-\mathbf{e}_k$ for a \geq constraint. Thus we have

$$\bar{\mathbf{b}} - \bar{\mathbf{A}}_s(\delta) \geq \mathbf{0} \quad \text{for a } \leq \text{ type constraint, and}$$

$$\bar{\mathbf{b}} - \bar{\mathbf{A}}_s(\delta) \leq \mathbf{0} \quad \text{for a } \geq \text{ type constraint.}$$

Example 3

Consider Example 2 again and suppose x_4 represents a slack variable for constraint 1 (\leq type). If the coefficient b_1 is varied by an amount δ , for the following data

$$\bar{\mathbf{b}} = (3.2, 1.5, 5.6)^T \text{ and } \bar{\mathbf{A}}_s = \bar{\mathbf{A}}_4 = (1.0, -0.5, 2.0)^T$$

we have

$$x_1(\delta) = 3.2 - 1.0(\delta)$$

$$x_2(\delta) = 1.5 + 0.5(\delta)$$

$$x_6(\delta) = 5.6 - 2.0(\delta)$$

Thus

$$x_1(\delta) \geq 0 \text{ for } 3.2 - 1.0(\delta) \geq 0, \text{ that is, } \frac{3.2}{1.0}$$

$$x_2(\delta) \geq 0 \text{ for } 1.5 + 0.5(\delta) \geq 0, \text{ that is, } \frac{1.5}{0.5}$$

$$x_6(\theta) \geq 0 \text{ for } 5.6 - 2.0(-\theta) \geq 0, \text{ that is, } \frac{5.6}{-2.0}$$

Therefore, θ can vary in the range

$$\max \frac{3.2}{-1.0}, \frac{5.6}{-2.0} \quad \min \frac{1.5}{0.5}$$

$$-2.8 \quad 3.0$$

When it is desirable to perform proportional ranging on the vector \mathbf{b} , the above analysis is the same but the unit vector \mathbf{e}_k is replaced with \mathbf{b}^* giving $\mathbf{b}(\theta) = \mathbf{b} + \theta \mathbf{b}^*$. The current basis \mathbf{B} will remain optimal as long as $\mathbf{x}_B(\theta) = \mathbf{B}^{-1}\mathbf{b}(\theta) \geq 0$. This inequality gives rise to a range on θ within which primal feasibility is preserved (see Exercise 4).

Changes in matrix coefficients

The structural coefficients a_{ij} are usually known with much more certainty than the objective row or right-hand-side vector, since they customarily represent some physical interaction between variables and are not subject to the same market fluctuations as costs and demands. We shall consider changes to the coefficients of nonbasic variables only; changes to basic variables coefficients alter the basis matrix \mathbf{B} and are rather complicated to analyze (see Murty 1983).

Consider the j th nonbasic variable with corresponding column \mathbf{A}_j . If the i th element of \mathbf{A}_j is changed by an amount θ , this affects the reduced cost \bar{c}_j as follows.

$$\text{If } \mathbf{A}_j(\theta) = \mathbf{A}_j + \theta \mathbf{e}_i$$

$$\text{then } \bar{c}_j(\theta) = \pi(\mathbf{A}_j + \theta \mathbf{e}_i) - c_j$$

$$= \bar{c}_j + \theta \pi \mathbf{e}_i$$

$$= \bar{c}_j + \theta \pi_i$$

where $\pi (= \mathbf{c}_B \mathbf{B}^{-1})$ is the dual vector. Thus the solution remains optimal as long as $\bar{c}_j(\theta) \geq 0$. The corresponding range for θ is

$$-\frac{\bar{c}_j}{i} \text{ for } i > 0$$

$$-\frac{\bar{c}_j}{i} \text{ for } i < 0$$