Sensitivity, Shadow Prices, Reduced Costs and Duality for Canonical LP

Sensitivity is what happens to a LP, when the vectors $\boldsymbol{b}$ and $\boldsymbol{c}$ undergo small changes. It is understood in terms of, as before, excess amounts, reduced costs, and shadow prices.

To do so, we'll look inside the simplex method. This will also enable us to prove the Strong Duality Theorem, as a by-product. It will also turn out that when we solve the primal problem, by little effort, we can immediately produce the solution for the dual. Adding this little effort, which is technically adding m more columns to the tableau is called the *Dual simplex method.*

Consider a LP in a Canonical form, call it *Primal*:

$$\begin{cases} \text{Max } [\text{Min}] \quad z = \boldsymbol{c}\cdot\boldsymbol{x}, \quad \text{s.}t. \\ \\ \boldsymbol{x} \geq 0, \qquad A\boldsymbol{x} = \boldsymbol{b}, \quad \boldsymbol{x}, \boldsymbol{c} \in \mathbb{R}^n,\ \boldsymbol{b} \in \mathbb{R}^m,\ n > m. \end{cases} \tag{1}$$

The discussion below pertains to the Max problem; the changes one should make for the Min problem are given in square brackets. Note that any LP can be reduced to a canonical one, so it remains true for all LPs. Recall that the canonical LP can be also cast as the Manufacturing or Diet problem, essentially by writing $A\boldsymbol{x} = \boldsymbol{b}$ as $A\boldsymbol{x} \leq \boldsymbol{b}$ and $-A\boldsymbol{x} \geq -\boldsymbol{b}$. Review the duality handout for the manufacturing-diet problem pair. Exposed below is in essence the same theory.

We have used the duality relationship between the Manufacturing/Diet problems (HW2) to derive that the dual problem to the canonical form turns out to be

$$\text{Min } \boldsymbol{b}\cdot\boldsymbol{y}, \quad \text{s.t.} \quad A^T\boldsymbol{y} \geq \boldsymbol{c}, \quad \text{where } \boldsymbol{b} \in \mathbb{R}^m. \tag{2}$$

We will further see how the dual problem arises in the depth of the simplex method.

Assume: *non-degeneracy*: $\boldsymbol{b}$ is not a linear combination of fewer than m columns of A. I.e. every BFS of (1) has exactly m nonzero components. Otherwise – the SM algorithm has to be made more complicated and will make things less patent.

Suppose, $\boldsymbol{x}^*$ is a BFS of (1). Then exactly m components $\boldsymbol{x}^*_{\text{b}} = (x^*_{i_1}, x^*_{i_2}, \ldots, x^*_{i_m})$, corresponding to the *basis* $\mathcal{B} = \{i_1, i_2, \ldots, i_m\} \subset \{1, 2, \ldots, n\}$, are nonzero. Without loss of generality, assume $\mathcal{B} = \{1, 2, \ldots, m\}$. Accordingly, decompose any $\boldsymbol{x} = (\boldsymbol{x}_{\text{b}}, \boldsymbol{x}_{\text{f}})$ and $\boldsymbol{c} = (\boldsymbol{c}_{\text{b}}, \boldsymbol{c}_{\text{f}})$ into *basic* and *free* components and $A = [A_{\text{b}}\, A_{\text{f}}]$ into basic and free sub-matrices. The basis $\mathcal{B}$ has been fixed once and for all.

Then $\boldsymbol{x}^* = (\boldsymbol{x}^*_{\text{b}}, \boldsymbol{0})$ is a BS in question. and thus $A_{\text{b}}\boldsymbol{x}^*_{\text{b}} = \boldsymbol{b}$, so

$$\boldsymbol{x}^*_{\text{b}} = A_{\text{b}}^{-1}\boldsymbol{b}. \tag{3}$$

Note: $\boldsymbol{x}^* \in \mathbb{R}^n_+$, while $\boldsymbol{x}^*_b \in \mathbb{R}^m$, and *positive*, by the non-degenerate assumption. Thus, a sufficiently small change $\boldsymbol{b} \to \boldsymbol{b} + \delta\boldsymbol{b}$ of the right hand side will not violate the feasibility of the basis $\mathcal{B}$. Note, A_{b}^{-1} *must* be invertible, or the non-degeneracy assumption gets violated!

Generally, a basis $\mathcal{B}$ is feasible for such values of $\boldsymbol{b}$, whenever $A_{\text{b}}^{-1}\boldsymbol{b} \geq 0$. Denoting $A_{\text{b}}^{-1}\boldsymbol{b} = \boldsymbol{v}$, we conclude that the basis is feasible, whenever $\boldsymbol{b} \in A_{\text{b}}\boldsymbol{v},\ \boldsymbol{v} \in \mathbb{R}^m_+$. This condition is a generalisation of what we drew in two dimensions in HW3.

Now, let us describe the optimality test for the BFS $\boldsymbol{x}^*$. To perform the optimality test, we need to compare the value of $\boldsymbol{x}^*$ to the value of all other feasible solutions $\boldsymbol{x}$. So, let $\boldsymbol{x}$ be *any* solution of the LP (1). Don't confuse it with the *chosen* $\boldsymbol{x}^*$! Split now $\boldsymbol{x}$ into basic and free components, as to the basis defined by $\boldsymbol{x}^*$: $\boldsymbol{x} = (\boldsymbol{x}_\mathrm{b},\, \boldsymbol{x}_\mathrm{f})$. $A\boldsymbol{x} = \boldsymbol{b}$ reads $[A_\mathrm{b}\, A_\mathrm{f}] \begin{bmatrix} \boldsymbol{x}_\mathrm{b} \\ \boldsymbol{x}_\mathrm{f} \end{bmatrix} = \boldsymbol{b}$. Multiplying it on the left by A_b^{-1} and using (3) yields

$$\boldsymbol{x}_\mathrm{b} = \boldsymbol{x}^*_\mathrm{b} - A_\mathrm{b}^{-1} A_\mathrm{f} \boldsymbol{x}_\mathrm{f}. \tag{4}$$

This corresponds to the upper rows in the (long) tableau, corresponding to the basic solution $\boldsymbol{x}^*$. Indeed, the above can be rewritten as $\mathrm{Id}\, \boldsymbol{x}_\mathrm{b} + (A_\mathrm{b}^{-1} A_\mathrm{f})\boldsymbol{x}_\mathrm{f} = \boldsymbol{x}^*_\mathrm{b}$, where Id is the $m \times m$ identity matrix. This *is* the tableau: the matrix $A_\mathrm{b}^{-1} A_\mathrm{f}$ provides the coefficients in the free columns, with $\boldsymbol{x}^*_\mathrm{b}$ being the value column.

Now use (4) to compute that the objective value – we'll use transposes rather than dot products to emphasize that what we get:

$$\boldsymbol{c}^T \boldsymbol{x} \;=\; \boldsymbol{c}_\mathrm{b}^T \boldsymbol{x}_\mathrm{b} + \boldsymbol{c}_\mathrm{f}^T \boldsymbol{x}_\mathrm{f} \;=\; \boldsymbol{c}_\mathrm{b}^T \boldsymbol{x}^*_\mathrm{b} \;-\; (\boldsymbol{c}_\mathrm{b}^T A_\mathrm{b}^{-1} A_\mathrm{f} \;-\; \boldsymbol{c}_\mathrm{f}^T)\boldsymbol{x}_\mathrm{f}, \tag{5}$$

is precisely the objective row in the tableau, corresponding to the basic solution $\boldsymbol{x}^*$. Indeed, it can be simply rewritten as

$$z + \boldsymbol{0} \cdot \boldsymbol{x}_\mathrm{b} + (\boldsymbol{c}_\mathrm{b}^T A_\mathrm{b}^{-1} A_\mathrm{f} - \boldsymbol{c}_\mathrm{f}^T)\boldsymbol{x}_\mathrm{f} = \boldsymbol{c}_\mathrm{b}^T \boldsymbol{x}^*_\mathrm{b},$$

and this is the objective row, the value of the BFS $\boldsymbol{x}^*$ sitting on the right, the coefficients in parentheses providing a row-vector, which may have nonzero components only in free variables' columns.

One may wish to denote

$$\boldsymbol{r}_\mathrm{f}^T = \boldsymbol{c}_\mathrm{b}^T A_\mathrm{b}^{-1} A_\mathrm{f} - \boldsymbol{c}_\mathrm{f}^T,$$

or equivalently

$$\boldsymbol{r}_\mathrm{f} = A_\mathrm{f}^T [A_\mathrm{b}^{-1}]^T \boldsymbol{c}_\mathrm{b} - \boldsymbol{c}_\mathrm{f}, \tag{6}$$

and define the n-vector of *reduce costs* $\boldsymbol{r} = (\boldsymbol{0}, \boldsymbol{r}_\mathrm{f})$, and *this* appears as the row of coefficients in the objective row.

Now comes the conclusion: $\boldsymbol{x}^*$ is *not optimal* if and only if at least one component of $\boldsymbol{r}$ is negative [positive for Min]. Otherwise, $\boldsymbol{x}^*$ *not optimal*, i.e. the tableau is *final.*

Suppose now until the end that $\boldsymbol{x}^*$, associated to the basis $\mathcal{B}$ is an optimal solution of (1). Then $\boldsymbol{r}_\mathrm{f} \geq [\leq]\, 0$ (component-wise). In this case, a single component of the reduced cost vector equals the maximum amount, by which one can increase [decrease] the corresponding free variable's objective component $\boldsymbol{c}_\mathrm{f}$, without violating the optimality of the solution $\boldsymbol{x}^*$.

In general, will a *small* perturbation of the objective $\boldsymbol{c}$ violate the optimality of $\boldsymbol{x}^*$ and the basis $\mathcal{B}$? No, provided that $\boldsymbol{c}$ is such that $\boldsymbol{r}_\mathrm{f}, > [<]\, 0$, that is the reduced cost of each free variable is *strictly* positive [negative]. I.e. if we don't have the alternative solutions scenario. This is a "dual analogue" of the non-degeneracy assumption on $\boldsymbol{b}$.

To make more sense out of (6) and bring the dual problem into the game, let us denote a row-vector

$$\boldsymbol{y}^* \;=\; [A_\mathrm{b}^{-1}]^T \boldsymbol{c}_\mathrm{b}. \tag{7}$$

This is a "dual analogue" of the formula (3). The vector $\boldsymbol{y}^* \in \mathbb{R}^m$ is called the *shadow price* of the constraint vector $\boldsymbol{b}$, for reasons to become apparent.

Clearly, $\boldsymbol{y}^*$ satisfies the following system of inequalities: $A^T \boldsymbol{y}^* \geq \boldsymbol{c}$. [$\leq$ for the Min problem. From now on let us stop talking about the Min problem.] Indeed for the basic components it follows from (7) which defines $\boldsymbol{y}^*$; for the free components, this is the final tableau condition that the free variables' coefficients (6) in the objective row of the free variables be non-negative.

In addition,

$$\boldsymbol{y}^* \cdot \boldsymbol{b} = \boldsymbol{b} \cdot [A_{\rm b}^{-1}]^T \boldsymbol{c}_{\rm b} = A_{\rm b}^{-1} \boldsymbol{b} \cdot \boldsymbol{c}_{\rm b} = \boldsymbol{x}^* \cdot \boldsymbol{c},$$

by (3).

So, the dual problem (2) has popped up on its own: we know now that $\boldsymbol{y}^*$ is feasible for it and has the same value as $\boldsymbol{x}^*$ for the primal. Moreover, the above formula is a strong duality claim! We will be done now by showing – easily – weak duality. This will imply that $\boldsymbol{y}^*$ is *optimal* for the dual.

`Theorem:  Duality`

1. `weak duality:` If $\boldsymbol{x}$, $\boldsymbol{y}$ is a pair of feasible solutions for the primal (1) and the dual (2) LPs, one always has
$$\boldsymbol{c} \cdot \boldsymbol{x} \;\leq\; \boldsymbol{y} \cdot \boldsymbol{b}. \tag{8}$$
Hence, if the primal (dual) is unbounded, the dual (primal) is unfeasible.

2. `strong duality:` If $\boldsymbol{x}^* = (\boldsymbol{x}_{\rm b}^*, 0)$, where $\boldsymbol{x}_{\rm b}^*$ is given by (3) is optimal for the primal, then $\boldsymbol{y}^*$, defined by (7) is optimal for the dual, with equal *equilibrium value*
$$\mathrm{V} \;=\; \boldsymbol{c} \cdot \boldsymbol{x}^* \;=\; \boldsymbol{y}^* \cdot \boldsymbol{b}. \tag{9}$$

`Proof:` We only have to prove weak duality, the rest has been shown already. As $\boldsymbol{x} \geq 0$ and $A^T \boldsymbol{y} \;\geq\; \boldsymbol{c}$, we have $\boldsymbol{c} \cdot \boldsymbol{x} \;\leq\; \boldsymbol{y} \cdot A\boldsymbol{x} = \boldsymbol{y} \cdot \boldsymbol{b}$: the objective value for any feasible solution for the primal is $\leq$ the objective value for any feasible solution for the dual. This proves (8). It follows that if primal is unbounded, dual (whose value should be still greater) has to be unfeasible. And the other way around.

`Corollary:` Given the basis $\mathcal{B}$, $\boldsymbol{x}^* = (\boldsymbol{x}_{\rm b}^*, 0)$, where $\boldsymbol{x}_{\rm b}^*$ is given by (3) is optimal for the primal if and only if $\boldsymbol{y}^*$, given by (7) is *feasible* (and then automatically optimal) for the dual. I.e all the reduced costs $\boldsymbol{r}_{\rm f}$ in (6), corresponding to free variables, must be positive for the maximum and negative for the minimum problem. Otherwise at least one of the inequalities $A_{\rm f}^T \boldsymbol{y} \;\geq\; \boldsymbol{c}_{\rm f}$ gets violated.

Under both non-degeneracy assumptions, made along the way, the sensitivity questions are now easy to answer. A sufficiently small change of the constraint vector $\boldsymbol{b} \to \boldsymbol{b} + \delta\boldsymbol{b}$ or the objective vector $\boldsymbol{c} \to \boldsymbol{c} + \delta\boldsymbol{c}$ will not change the optimal basis, hence the formulae (3,7,9) will remain valid. Therefore, the change of the optimal value will change as

$$\mathrm{V} \;\to\; \mathrm{V} \;+\; \boldsymbol{y}^* \cdot \delta\boldsymbol{b} \;+\; \delta\boldsymbol{c}^T \boldsymbol{x}^*. \tag{10}$$

Thus, the shadow price $\boldsymbol{y}^*$ is a proportionality coefficient in the amount, by which the objective value will change due to a small variation of the constraint vector $\boldsymbol{b}$, while $\boldsymbol{x}^*$ reflects the same apropos of variations of the objective vector $\boldsymbol{c}$. The quantities $\boldsymbol{x}^*$ and $\boldsymbol{y}^*$ are often called cost and constraint *sensitivities.* In economics they tell a financier how the profit of their enterprise will react upon the market price and supply and demand variations.

Dual simplex method

So solving the dual problem inside the Simplex method does not require too much work: $\boldsymbol{y}^*$ must satisfy (7). One can in effect solve the problem by SM, identify the optimal basis, invert the matrix A_{b} and use the formula (7). Equivalently, the inversion of the matrix can be incorporated into the tableau calculus. Recall that a square matrix can be inverted by the Gauss-Jordan method (See HW1, Problem 4 in the review group. In short, to get M^{-1}, for an $m \times m$ matrix M, one needs to augment M to a $m \times 2m$ matrix by writing the identity matrix to the right and then do, if possible, a sequence of pivots that would turn M into the identity matrix. What is now on the right is M^{-1}.)

So, all one needs to do, whether it is the two-phase method or not, is to add the $m \times m$ unit matrix to the right of the initial tableau (long or short), put down the components of $\boldsymbol{c}$, usually zeroes, (especially if this is the two-phase method) corresponding to the starting solution in the last row to complete the bottom row, and proceed as normal. (If this is the two-phase method, do not erase anything to the right of the value column when the artificial variables are being deleted.) Than at the end of the calculation, the numbers in the bottom row to the right of the value column will yield $\boldsymbol{y}^*$. There is still a shorter form of doing this, see e.g. 1.10, 1.11 in Franklin's book.

If the original problem is the Manufacturing problem and has slack variables, this is redundant: the reduced costs of slack variables in the final tableau will be equal to the constraints' shadow prices. This is, in fact, no more than a tautology (as is arguably the whole of maths, yet outs is an easy one). To see this, consider the Canonical form

$$\text{Max}\, \boldsymbol{c} \cdot \boldsymbol{x} + \boldsymbol{0} \cdot \boldsymbol{s}, \quad \text{s.t.} \quad [A \;\; \text{Id}] \begin{bmatrix} \boldsymbol{x} \\ \boldsymbol{s} \end{bmatrix} = \boldsymbol{b}, \quad \boldsymbol{x}, \boldsymbol{s} \geq 0,$$

which has come from the Manufacturing problem `Max` $\boldsymbol{c} \cdot \boldsymbol{x}$, s.t. $A\boldsymbol{x} \leq \boldsymbol{b}$, $\boldsymbol{x} \geq 0$ after adding slack variables. The dual (to the Canonical form) is

$$\text{Min}\, \boldsymbol{b} \cdot \boldsymbol{y}, \quad \text{s.t.} \quad A^T \boldsymbol{y} \geq \boldsymbol{c}, \; \text{Id}^T \boldsymbol{y} \geq \boldsymbol{0}.$$

If $\boldsymbol{y}^*$ solves the latter, then the *reduced costs* of the slack variables are, by definition, the differences between the left- and right-hand-side of the second group of inequalities, which is just $\boldsymbol{y}^*$.

So, if one does the Manufacturing problem, the final tableau entries in the bottom row, slack columns *are* $\boldsymbol{y}^*$. But in general, to get $\boldsymbol{y}^*$ one needs to do a bit of extra work described above.