Linear Programming Duality

Matthew J. Saltzman*

March 11, 1998

In this essay, we will "discover" the *dual problem* associated with an LP. We will see how to interpret the meanings of the dual decision variables in the context of the original problem, and we will present some theorems ("facts") about the relationship between the optimal primal and dual solutions that will lead us to the key ideas of the simplex method for solving LPs.

1 The Dual of an LP

Consider the fertilizer manufacturer's problem from Murty [1], Section 2.1:

$$\begin{array}{rrcrcr} \text{Maximize} & 15x_1 & + & 10x_2 & = & z \\ \text{subject to} & 2x_1 & + & x_2 & \leq & 1500 \\ & x_1 & + & x_2 & \leq & 1200 \\ & x_1 & & & \leq & 500 \\ & & & x_1, x_2 \geq 0 & & \end{array} \tag{1}$$

In this example, x_1 and x_2 represent the amount (in tons) of Hi-Ph and Lo-Ph fertilizer products, respectively, to be produced each day. The right-hand sides represent the amounts (in tons) of ingredients RM1, RM2, and RM3 available each day. The objective function coefficients represent the profit ($/ton) for each product, and the technology coefficients represent the amounts of each ingredient required to produce each product. Thus, it requires two tons of RM1 to produce one ton of Hi-Ph, etc. The units on these coefficients are in units of tons of ingredient per ton of product.

In our search for the optimal solution to this problem, we can generate lower bounds on the optimal value of z by generating feasible solutions to the problem. For example, if we propose the production of 500 tons of Hi-Ph and 500 tons of Lo-Ph, the solution is represented as $x = (500, 500)$ and the objective value is $z = 12500$. It is clear that any optimal solution must generate at least $12,500 in profit, because at least one feasible solution does. But it is difficult to tell for a general, large LP (although not too difficult for this example), how much more profit we might expect (or indeed if the optimal profit is even finite). It would be useful to be able to generate upper bounds on the optimal objective value as well.

We can generate such upper bounds by constructing an inequality (of the less-than-or-equal variety) satisfied by every feasible solution, and such that, for any feasible solution, the left-hand side is always at least as large as the objective function value and the right-hand side is a constant. If we had such an inequality $\alpha_1 x_1 + \alpha_2 x_2 \le z_U$, then we could conclude that, for any feasible solution,

$$z = 15x_1 + 10x_2 \le \alpha_1 x_1 + \alpha_2 x_2 \le z_U.$$

What conditions must our inequality satisfy? Since x_1 and x_2 are nonnegative, $\alpha_1 x_1 + \alpha_2 x_2$ will be larger than z if $\alpha_1 \ge 15$ and $\alpha_2 \ge 10$. In order to ensure that every feasible x satisfies the inequality, we must construct it from the original constraints in a way that preserves feasibility and the less-than-or-equal senses of the originals. For example, we get one such inequality if we multiply both sides of the first constraint by 10, yielding

$$20x_1 + 10x_2 \le 15000.$$

If we add five times the first inequality to five times the second, we get the better bound

$$15x_1 + 10x_2 \le 13500.$$

Let us consider linear combinations of the LP constraints, with multipliers π_1, π_2, and π_3, respectively. Then the coefficients of the generated inequality are

$$\begin{aligned}
\alpha_1 &= 2\pi_1 + \pi_2 + \pi_3, \\
\alpha_2 &= \pi_1 + \pi_2, \\
z_U &= 1500\pi_1 + 1200\pi_2 + 500\pi_3.
\end{aligned}$$

The multipliers we need must be nonnegative (to preserve the less-than-or-equal sense of the inequalities), and the resulting coefficients must be at least as large as the corresponding objective coefficients, to guarantee that the new left-hand side is at least as large as the objective value, for any feasible solution. The multipliers that give us the best upper bound are (remarkably) the solutions to the following LP:

$$\begin{array}{rrcrcrcl}
\text{Minimize} & 1500\pi_1 & + & 1200\pi_2 & + & 500\pi_3 & = & z_U \\
\text{subject to} & 2\pi_1 & + & \pi_2 & + & \pi_3 & \ge & 15 \\
 & \pi_1 & + & \pi_2 & & & \ge & 10 \\
 & & & \multicolumn{3}{c}{\pi_1, \pi_2, \pi_3 \ge 0} & &
\end{array} \qquad (2)$$

The LP (2) is called the *dual* of the LP (1). The LP (1) is called the *primal* problem. (The problem you start with the always the primal and the bounding problem is always the dual.)

2 General LPs, Their Duals, and Standard Form

The example above is a special case of LPs of the form

$$\begin{array}{rl} \text{Maximize} & cx \\ \text{subject to} & Ax \le b \\ & x \ge 0, \end{array} \tag{3}$$

where c is a row vector ($1 \times n$ matrix), A is $m \times n$, and b is $m \times 1$. By the reasoning above, the dual of such an LP is

$$\begin{array}{rl} \text{Minimize} & \pi b \\ \text{subject to} & \pi A \ge c \\ & \pi \ge 0, \end{array} \tag{4}$$

where π is a $1 \times m$ row vector.

It doesn't matter, though, whether the primal LP is in this particular form. We can construct the dual of any LP using reasoning similar to the above. For example, suppose the primal objective is to be maximized and a constraint is of the greater-than-or-equal variety. To use the constraint to get an upper bound, it must be multiplied by a non-positive scalar to change its sense. If the constraint is an equation, then its multiplier may have either sign. If a primal variable is unrestricted in sign, the corresponding coefficient in the bounding inequality must be equal to the objective coefficient. If the primal objective is to be minimized, then the dual objective is to be maximized, and the generated constraint must be of the greater-than-or-equal variety. These conclusions are summarized in Table 1.

Despite the fact that we can produce the dual for any LP, it is convenient when discussing the relationship between primal and dual problems to restrict our attention to problems with a particular form. Although some books discuss the primal-dual pair (3) and (4) (sometimes called *canonical form* or *symmetric form*), we will find it most useful to discuss the following primal-dual pair (usually called *standard form*):

$$\begin{array}{rl} \text{Minimize} & cx \\ \text{subject to} & Ax = b \\ & x \ge 0, \end{array} \tag{5}$$

Primal minimization/Dual maximization			
Primal Variables	Dual Constraints	Primal Constraints	Dual Variables
≥ 0	$\leq$	$\geq$	≥ 0
≤ 0	$\geq$	$\leq$	≤ 0
Unrestricted	$=$	$=$	Unrestricted
Primal maximization/Dual minimization			
Primal Variables	Dual Constraints	Primal Constraints	Dual Variables
≥ 0	$\geq$	$\geq$	≤ 0
≤ 0	$\leq$	$\leq$	≥ 0
Unrestricted	$=$	$=$	Unrestricted

Table 1: Relationship between primal and dual variables and constraints.

and

$$\begin{array}{ll} \text{Maximize} & \pi b \\ \text{subject to} & \pi A \leq c \\ & \pi \text{ unrestricted.} \end{array} \tag{6}$$

We can do this without loss of generality, because it is possible to transform any LP into an equivalent standard-form LP. We will discuss some key rules for doing this now, and put off a few others until later. Note that LP solvers work internally with a minor extension to standard form, where free variables and bounds on variables are handled directly. We won't discuss this version, because it complicates the description of the simplex method. Also note that modeling languages like LINGO handle the conversion to standard form automatically—the user can enter the constraints in any convenient form.

1. To convert a maximization problem to a minimization problem, multiply all coefficients by -1. Thus in our example problem, the objective becomes

$$\text{Minimize } -15x_1 - 10x_2.$$

2. To convert a less-than-or-equal constraint to an equation, include a *slack variable* and constrain it to be nonnegative. The first constraint in the example becomes

$$\begin{aligned} 2x_1 + x_2 + x_3 &= 1500 \\ x_3 &\geq 0. \end{aligned}$$

You need a different variable for each such constraint, so the remaining constraints are

$$\begin{aligned} x_1 + x_2 + x_4 &= 1200 \\ x_1 + x_5 &= 500 \\ x_4, x_5 &\geq 0. \end{aligned}$$

3. For a greater-than-or-equal constraint, subtract a nonnegative *surplus variable.* Thus, for example, the inequality

$$3x_1 - 2x_2 \geq 400$$

becomes

$$\begin{aligned} 3x_1 - 2x_2 - x_s &= 400 \\ x_s &\geq 0 \end{aligned}$$

4. For a variable with finite upper and lower bounds, incorporate the upper bound explicitly into the constraints, and treat the lower bound as described below.

5. To convert a variable $x_j \geq l_j$ with a nonzero lower bound to one with a zero lower bound, define $y_j = x_j - l_j$. It is clear that whenever $x_j \geq l_j$, then $y_j \geq 0$. Make the substitution $x_j = y_j + l_j$ in the LP, and collect the constant column $l_j A_{\cdot j}$ on the right hand side. Thus if our fertilizer problem had the minimum production requirement $x_1 \geq 100$, we would define $y_1 = x_1 - 100$. Substituting $x_1 = y_1 + 100$ into the problem gives

$$\begin{array}{lrcrcrcrcrcr} \text{Maximize} & 15(y_1+100) & + & 10x_2 & & & & & & & = & z \\ \text{subject to} & 2(y_1+100) & + & x_2 & + & x_3 & & & & & = & 1500 \\ & (y_1+100) & + & x_2 & & & + & x_4 & & & = & 1200 \\ & (y_1+100) & & & & & & & + & x_5 & = & 500 \\ & & & & & y_1, x_2, \ldots, x_5 \geq 0 & & & & & & \end{array} \tag{7}$$

or equivalently,

$$\begin{array}{lrcrcrcrcrcr} \text{Maximize} & 15y_1 & + & 10x_2 & & & & & & & = & z - 1500 \\ \text{subject to} & 2y_1 & + & x_2 & + & x_3 & & & & & = & 1300 \\ & y_1 & + & x_2 & & & + & x_4 & & & = & 1100 \\ & y_1 & & & & & & & + & x_5 & = & 400 \\ & & & & & y_1, x_2, \ldots, x_5 \geq 0 & & & & & & \end{array} \tag{8}$$

Note the constant adjustment to the objective function—the result of solving the LP with y_1 will be \$1500 too low, since y_1 only gives the excess production of HiPh over 100 tons.

6. To convert a variable $x_j \le u_j$ with a finite upper bound and no lower bound, define $y_j = u_j - x_j$ and make the substitution $x_j = u_j - y_j$, as above.

7. For unrestricted variables, Murty describes a process of using equality constraints to substitute them out of the problem. (Note that if we run out of equations before we run out of free variables, then the problem is unbounded.) A technique that does somewhat less violence to the formulation is to define two new variables for each free variable: $x_j = x_j^+ - x_j^-$, where $x_j^+, x_j^- \ge 0$. Now whatever value x_j takes on, there are many possible values for x_j^+ and x_j^-. Of particular interest are the ones where exactly one of x_j^+ and x_j^- is zero. For example, if $x_j = 3$, then $x_j^+ = 3$ and $x_j^- = 0$, and if $x_j = -2$ then $x_j^+ = 0$ and $x_j^- = 2$. In the simplex method, it is guaranteed that the final values of x_j^+ and x_j^- will have one of these forms.

The conversion of the example problem to standard form gives:

$$
\begin{array}{lrcrcrcrcrcr}
\text{Minimize} & -15x_1 & - & 10x_2 & & & & & & & = & z \\
\text{subject to} & 2x_1 & + & x_2 & + & x_3 & & & & & = & 1500 \\
 & x_1 & + & x_2 & & & + & x_4 & & & = & 1200 \\
 & x_1 & & & & & & & + & x_5 & = & 500 \\
 & & & & & x_1, \dots, x_5 \ge 0 & & & & & &
\end{array} \tag{9}
$$

The dual of this problem is

$$
\begin{array}{lrcrcrcr}
\text{Maximize} & 1500\pi_1 & + & 1200\pi_2 & + & 500\pi_3 & = & z_U \\
\text{subject to} & 2\pi_1 & + & \pi_2 & + & \pi_3 & \le & -15 \\
 & \pi_1 & + & \pi_2 & & & \le & -10 \\
 & \pi_1 & & & & & \le & 0 \\
 & & & \pi_2 & & & \le & 0 \\
 & & & & & \pi_3 & \le & 0
\end{array} \tag{10}
$$

which is easily seen to be equivalent to (2). (Note that in this version the dual variables are restricted to be nonpositive. They are just the negatives of the variables in (2). Also note that the interpretation described below now applies to the negatives of these variables.)

3 Economic Interpretation of Dual Variables

LPs that represent problems from the real world have coefficients that correspond in a meaningful way to real quantities being modeled, and they have units associated with them. We can use this information to impose an interpretation on the dual multipliers. In our example, the dual constraints are

$$2\pi_1 + \pi_2 + \pi_3 \geq 15$$

and

$$\pi_1 + \pi_2 \geq 10.$$

From the original problem, the coefficients have units "tons of ingredient RMi per ton of product" and the right hand sides have units "\$ per ton of product". This forces the π_is to have units "\$ per ton of ingredient RMi", in order to make the values on both sides of the inequalities comparable. The dual objective

$$1500\pi_1 + 1200\pi_2 + 500\pi_3$$

has the units "\$", just as the primal objective does. We can interpret the entire example dual problem from the point of view of an *investor* interested in purchasing resources from the *producer* who is solving the primal problem. The dual variables represent premiums (over cost) to be paid for each unit of resource bought by the investor. We call the dual variables *shadow prices* for the resources. The dual constraints ensure that the premium paid for a "package" of resources used to produce a unit of (say) Hi-Ph is at least equal to the profit for that unit. If a proposed set of prices violates the dual constraint corresponding to Hi-Ph, then the producer would refuse to sell the package of resources, preferring instead to use them to produce Hi-Ph. The investor's objective is to minimize the amount offered for the entire set of resources.

It seems clear that the investor will have to offer prices such that the total value is at least equal to the maximum profit that the producer could make using the resources to manufacture product. In fact, the relationship between primal and dual solutions ensures that the best prices give a minimum total value for the resources exactly equal to the maximum profit to be made by using the resources in production. These and other relationships between primal and dual solutions are explored below.

4 Duality, Optimality, and Complementarity

The following theorems describe the relationship between optimal solutions to a primal-dual pair of LPs. The results are presented for problems in standard form, but similar results hold for any primal-dual pair.

Theorem 1 *The dual of (6) is (5).*

To see that this is true, we can convert (6) to standard form:

1. Begin with (6):

$$\begin{array}{rc} \text{Maximize} & \pi b \\ \text{subject to} & \pi A \leq c \\ & \pi \text{ unrestricted.} \end{array}$$

2. Substitute for the unrestricted variables $\pi = \pi^+ - \pi^-$ where $\pi^+, \pi^- \geq 0$, add dual slacks, and convert to minimization form:

$$\begin{array}{rc} \text{Minimize} & -\pi^+ b + \pi^- b \\ \text{subject to} & \pi^+ A - \pi^- A + \bar{c} = c \\ & \pi^+, \pi^-, \bar{c} \geq 0. \end{array}$$

3. Take the dual:

$$\begin{array}{rc} \text{Maximize} & cx \\ \text{subject to} & Ax \leq -b \\ & -Ax \leq b \\ & x \leq 0. \end{array}$$

4. Combine the inequalities, substitute $\hat{x} = -x$, and convert to minimization form:

$$\begin{array}{rc} \text{Minimize} & c\hat{x} \\ \text{subject to} & A\hat{x} = b \\ & \hat{x} \geq 0. \end{array}$$

Alternatively, we could have constructed the dual directly using the rules in Table 1.

Theorem 2 (Weak Duality) *If $\bar{x}$ is a feasible solution to (5) and $\bar{\pi}$ is a feasible solution to (6), then*

$$c\bar{x} \geq \bar{\pi} b.$$

This theorem implies that any feasible shadow prices value the resources at a level at least equal to the profit to be made by turning the resources into product in any feasible mix.

Dual	Primal		
	Feasible	Infeasible	Unbounded
Feasible	Yes	No	No
Infeasible	No	Yes	Yes
Unbounded	No	Yes	No

Table 2: Possible relationships between primal and dual solutions.

Theorem 3 (Strong Duality) *If (5) has an optimal solution x^*, then (6) has an optimal solution π^*, and*

$$cx^* = \pi^* b.$$

This theorem implies that the best shadow prices value the resources at a level exactly equal to the profit for the best production plan. It also has some consequences in the case where one or the other LP is unbounded or infeasible. Table 2 lays out the possible combinations of primal and dual solutions.

If we include slack variables in (6), we get the equivalent dual problem

$$\begin{array}{ll} \text{Maximize} & \pi b \\ \text{subject to} & \pi A + \bar{c} = c \\ & \pi \text{ unrestricted}, \ \bar{c} \geq 0, \end{array} \tag{11}$$

and the dual slacks are defined to be

$$\bar{c} = c - \pi A.$$

In our example problem (10):

$$\bar{c}_1 = -15 - 2\pi_1 - \pi_2 - \pi_3 \tag{12}$$
$$\bar{c}_2 = -10 - pi_1 - \pi_2 \tag{13}$$
$$\bar{c}_3 = 0 - \pi_1 \tag{14}$$
$$\bar{c}_4 = 0 - \pi_2 \tag{15}$$
$$\bar{c}_5 = 0 - \pi_3. \tag{16}$$

These dual slacks are also called *reduced costs* or *relative costs* (if the primal problem is a minimization) or *relative profits* (if the primal problem is a maximization). The slack corresponding to the jth constraint indicates the amount by which the valuation of the package of resources used in product j exceeds the profit to be made on a unit of product j.

The following theorem describes the relationship between primal decision variables and dual slacks at optimality.[1]

Theorem 4 (Complementary Slackness (Karush-Kuhn-Tucker)) *If x^* and π^* are solutions to (5) and (6), respectively, then they are optimal solutions if and only if*

1. *x^* is primal feasible;*
2. *π^* is dual feasible (i.e., $\bar{c} \geq 0$); and*
3. *At least one of each pair $(x_j, \bar{c}_j)$ is zero for $j = 1, \ldots, n$ (alternatively, $x_j\bar{c}_j = 0$ for $j = 1, \ldots, n$).*

For the fertilizer example in standard form, consider the primal solution $\bar{x} = (300, 900, 0, 0, 200)^T$ and the dual solution $\bar{\pi} = (-5, -5, 0)$. Then $\bar{c} = (0, 0, 5, 5, 0)$. It is clear that $\bar{c} \geq 0$ and $\bar{x}_j\bar{c}_j = 0$ for $j = 1, \ldots, 5$, so this primal-dual pair of solutions is optimal. Both primal and dual solutions have objective value -13500.

4.1 Complementarity and Optimality

The KKT theorem suggests a strategy for testing if a proposed primal solution $\bar{x}$ is optimal, namely:

1. Construct all complementary dual solutions.
2. If any complementary dual solution is dual feasible, then the primal solution is optimal.

For example, consider the solution $\bar{x} = (500, 500, 0, 200, 0)^T$ to (9). A complementary dual solution must satisfy the equations

$$\begin{array}{rcrcrcr} 2\pi_1 & + & \pi_2 & + & \pi_3 & = & -15 \\ \pi_1 & + & \pi_2 & & & = & -10 \\ & & \pi_2 & & & = & 0 \end{array}$$

The unique solution to these equations is $\bar{\pi} = (-10, 0, 5)$. This solution is not dual feasible, since it violates the last inequality in (10).

The only problem we face is that a given $\bar{x}$ may not lead to a unique complementary $\bar{\pi}$. There may be no complementary dual solution, in which case the primal x

[1] If you studied constrained optimization at all in Calculus, you will recognize the shadow prices as Lagrange multipliers. The KKT theorem describes the optimality conditions for a constrained optimization problem, here specialized to the case of linear objective and linear constraints.

is not optimal, or there may be many such solutions, in which case additional effort is required to determine if *any* of them are dual feasible. To avoid having to deal with these problems, we turn our attention to a systematic method of generating complementary primal-dual pairs.

5 Bases and Basic Solutions

Consider the system of constraint equations $Ax = b$ from the standard-form LP (5). We can assume without loss of generality that A is $m \times n$, with $m \leq n$, and that $\text{rank}(A) = m$. Recall that $\text{rank}(A) = m$ means that we can find up to m linearly independent columns in A. Suppose we select such a set and use them to form the matrix B. With the leftover columns, form the matrix N. Then we can reorder the columns in A so that the basic columns are at the left, and reorder the components of c x the same way, and call the basic components c_B and x_B and the nonbasic components c_N and x_N. The LP can be rewritten as follows:

$$\begin{array}{rc} \text{Minimize} & cx \\ \text{subject to} & Ax = b \\ & x \geq 0, \end{array}$$

$$\begin{array}{rc} \text{Minimize} & c_B x_B + c_N x_N \\ \text{subject to} & B x_B + N x_N = b \\ & x_B \geq 0, \quad x_N \geq 0. \end{array}$$

The dual can also be rewritten as follows:

$$\begin{array}{rc} \text{Maximize} & \pi b \\ \text{subject to} & \pi A \leq c \\ & \pi \text{ unrestricted.} \end{array}$$

$$\begin{array}{rc} \text{Maximize} & \pi b \\ \text{subject to} & \pi B \leq c_B \\ & \pi N \leq c_N \\ & \pi \text{ unrestricted.} \end{array}$$

Since the columns of B are linearly independent (B is $m \times m$ with $\text{rank}(B) = m$), B must be nonsingular, with inverse B^{-1}. Multiplying both sides of the primal constraint equation by B^{-1} and solving for x_B gives

$$x_B = B^{-1}b - B^{-1}N x_N,$$

the formula for values of the basic variables in terms of the nonbasic variables. These equations, together with the equations $x_N = 0$ (from the nonnegativity constraints), form a system of n linearly independent equations in n unknowns, and has a unique solution, namely

$$x_B = B^{-1}b; \quad x_N = 0.$$

This solution is called the *basic primal solution* corresponding to the basis B. If $x_B \geq 0$, then this solution x is feasible for the LP, and it is called a *basic feasible solution* or *BFS*. (But for a particular choice of B, the corresponding basic primal solution may not be feasible.)

If $x_B > 0$, then the complementarity conditions would require that the dual slacks corresponding to B and c_B be satisfied with equality. Again, we can multiply both sides by B^{-1} (on the right this time) to get the solution:

$$\begin{aligned} \pi B &= c_B \\ \pi &= c_B B^{-1} \end{aligned}$$

This π is called the *basic dual solution* corresponding to the basis B.

Since these solutions are complementary, the KKT theorem applied to the solutions corresponding to a basis B says that if $x_B \geq 0$ and $\pi N \leq c_N$ then x and π are an optimal primal-dual pair.

In our example problem, consider the basis formed by columns 1, 2, and 4 (in that order). Then we have

$$\begin{aligned} x_B &= (x_1, x_2, x_4)^T & \qquad x_N &= (x_3, x_5)^T \\ c_B &= (-15, -10, 0) & c_N &= (0, 0) \\ B &= \begin{bmatrix} 2 & 1 & 0 \\ 1 & 1 & 1 \\ 1 & 0 & 0 \end{bmatrix} & N &= \begin{bmatrix} 1 & 0 \\ 0 & 0 \\ 0 & 1 \end{bmatrix} \end{aligned}$$

Then

$$B^{-1} = \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & -2 \\ -1 & 1 & 1 \end{bmatrix}$$

and the corresponding primal and dual solutions are $x_B = (500, 500, 200)^T$, $x_N = (0,0)^T$), or $x = (500, 500, 0, 200, 0)^T$, and $\pi = (-10, 0, 5)$. The primal solution is feasible, hence a BFS, but the dual solution is not, so we conclude that this is not an optimal primal-dual pair.

On the other hand, consider the basis formed by columns 1, 2, and 5 (in that order). Then we have

$$x_B = (x_1, x_2, x_5)^T \qquad x_N = (x_3, x_5)^T$$
$$c_B = (-15, -10, 0) \qquad c_N = (0, 0)$$
$$B = \begin{bmatrix} 2 & 1 & 0 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix} \qquad N = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix}$$

Then

$$B^{-1} = \begin{bmatrix} 1 & -1 & 0 \\ -1 & 2 & 0 \\ -1 & 1 & 1 \end{bmatrix}$$

and the corresponding primal and dual solutions are $x_B = (300, 900, 200)^T$ and $x_N = (0, 0)^T$, or $x = (300, 900, 0, 0, 200)^T$, and $\pi = (-5, -5, 0)$. The primal solution is feasible, hence a BFS, and the dual solution is as well, so we conclude that this is an optimal primal-dual pair.

A Complementarity for Canonical Problems

Comparison of primal problems (1) and (9) and dual problems (2) and (10) suggests that, when slack variables are added to convert a canonical primal problem to standard form, the corresponding dual constraints simply impose the sign constraints on the dual variables. More generally, we can write the conversion steps to get from primal to dual as follows:

1. Begin with the canonical primal problem:

$$\begin{array}{rc} \text{Maximize} & cx \\ \text{subject to} & Ax \leq b \\ & x \geq 0. \end{array}$$

2. Convert to standard form by adding slack variables s:

$$\begin{array}{rc} \text{Minimize} & -cx - 0s \\ \text{subject to} & Ax + Is = b \\ & x, s \geq 0. \end{array}$$

3. Take the dual:

$$\begin{array}{rl} \text{Maximize} & \pi b \\ \text{subject to} & \pi A \leq -c \\ & \pi I \leq 0. \end{array}$$

4. Substitute $\pi = -\hat{\pi}$:

$$\begin{array}{rl} \text{Maximize} & -\hat{\pi} b \\ \text{subject to} & -\hat{\pi} A \leq -c \\ & -\hat{\pi} I \leq 0. \end{array}$$

5. Switch from maximization to minimization, multiply constraints through by -1, cancel the multiplication by I:

$$\begin{array}{rl} \text{Minimize} & \hat{\pi} b \\ \text{subject to} & \hat{\pi} A \geq c \\ & \hat{\pi} \geq 0. \end{array}$$

By exploiting this knowledge of the structure of canonical problems, and observing that the definitions of the dual slacks associated with the sign constraints in standard form reduce to $\bar{c}_j = -\pi_j$ (so $\bar{c}_j = 0$ if and only if $pi = 0$), we can write a KKT theorem for canonical problems.

Theorem 5 (Complementary Slackness—Canonical Form) *A primal-dual pair of solutions x^* and π^* to (3) and (4), respectively, is optimal if and only if*

1. *x^* is feasible;*
2. *π^* is feasible;*
3. *whenever $x^*_j > 0$, then $\pi^* A_{\cdot j} = c_j$; and*
4. *whenever $\pi^* A_{\cdot j} > c_j$, then $x^*_j = 0$.*

For the fertilizer example in canonical form, consider the primal solution $\bar{x} = (300, 900)^T$ and the dual solution $\bar{\pi} = (-5, -5, 0)$. Then the slacks for the primal constraints are $\bar{s} = (0, 0, 200)^T$ and the slacks for the dual constraints are $\bar{c} = (0, 0)$. It is clear that $\bar{x}_j \bar{c}_j = 0$ for $j = 1, 2$ and $s_i \pi_i = 0$ for $i = 1, 2, 3$, so this primal-dual pair of solutions is optimal. Both primal and dual solutions have objective value 13500.

References

[1] K. G. Murty, *Operations Research: Deterministic Optimization Models*, Prentice Hall, 1995.