

# Linear Programming Notes VI

## Duality and Complementary Slackness

### 1 Introduction

It turns out that linear programming problems come in pairs. That is, if you have one linear programming problem, then there is automatically another one, derived from the same data.

Start with an LP written in the form:

$$\max c \cdot x \text{ subject to } Ax \leq b, x \geq 0.$$

(We know from the study of problem transformations that you can write any LP in this form.) I will call this the **Primal**. It is useful to keep track of dimensions. Assume that there are  $n$  variables (components of  $x$ ) and  $m$  constraints. That means that  $c$  is  $n$ -dimensional,  $b$  is  $m$ -dimensional, and  $A$  is a matrix with  $m$  rows and  $n$  columns. The data of the problem are  $b$ ,  $c$ , and  $A$ . Using the same information, we can write down a new LP. This one has  $n$  constraints and  $m$  variables. The new problem, called the **Dual** has the form:

$$\min b \cdot y \text{ subject to } yA \geq c, y \geq 0.$$

You get the dual by “switching around” the parts of the Primal. The objective switches from max to min. The constraints switch from  $\leq$  to  $\geq$ . The number of constraints changes from  $m$  to  $n$ . The number of variables changes from  $n$  to  $m$ . The objective function coefficients switch roles with the right-hand side constants.

All we have done is switch symbols around. Here are two unrelated questions that might have come to mind. What’s the point? (This is a good, all-purpose, question.) Why stop at taking the dual of the primal, why not take the dual of the dual and the dual of the dual of the dual and so on? (This is not an all purpose question and probably would not occur to you unless you are really interested in logical manipulation.)

The real answer to the first question is that you will see. At least, I hope you will see. The Primal and the Dual are not just two linear programming problems formed using the same data. They are intimately related. Knowing something about one problem tells you something about the other. The mathematical relationship is described in what is called the Duality Theorem of Linear Programming. I will have a lot to say about this theorem.

The answer to the second question is simple. Before I give the answer, let me explain the question in more detail. I gave you a definition of the dual of a linear programming problem. Since any LP can be written in the standard form above, any LP has a dual. Since the dual of a LP is itself a LP, it has a dual. So we could keep on taking duals forever. Except that if you take the trouble to find the dual of the dual (in order to use the definition of duality you would need



### 3 The Duality Theorem

#### 3.1 Statement

In this subsection, I will state the theorem and try to explain what it implies.

**Theorem 1** *If problem (P) has a solution  $x^*$ , then problem (D) also has a solution (call it  $y^*$ ). Furthermore, the values of the problems are equal:  $c \cdot x^* = b \cdot y^*$ . If problem (P) is unbounded, then problem (D) is not feasible.*

*Similarly, if problem (D) has a solution  $y^*$ , then problem (P) also has a solution (call it  $x^*$ ). Furthermore, the values of the problems are equal:  $c \cdot x^* = b \cdot y^*$ . If problem (D) is unbounded, then problem (P) is not feasible.*

The Duality Theorem states that the problems (P) and (D) are intimately related. One way to think about the relationship is to create a table of possibilities.

P \ D	unbounded	has solution	not feasible
unbounded	no	no	possible
has solution	no	same values	no
not feasible	possible	no	possible

Each of the three rows represents one of three (exclusive and exhaustive) properties of the Primal. That is, (P) is either unbounded, has a solution, or is infeasible. The columns represent the same features of the Dual. If I grabbed two LPs at random, any one of the nine cells could happen. For example, the both problems could be unbounded. If the two problems are related by duality, then five of the nine boxes are impossible. Since one box must be true in each row and in each column, at least three of the boxes must be possible. What duality does, therefore, is rule out all but one possibility. If you are told that (P) is unbounded, then you know that (D) can't be feasible. If you are told that (P) has a solution, then you know that (D) has one too. Only when (P) is not feasible are you uncertain. Maybe (D) is infeasible too or maybe (D) is unbounded.

Much of this information can be summarized in a smaller table. Remember that every LP is either feasible or not feasible and that feasible problems either are unbounded or have a solution. The next table just divides problems into feasible and infeasible.

P \ D	feasible	not feasible
feasible	both have solutions; values equal	P unbounded
not feasible	D unbounded	possible

This table shows that if you know that both problems are feasible, then you know that neither problem is unbounded or, equivalently, that both have solutions. More than that, the values of the solutions are equal.

### 3.2 Why is the Duality Theorem True?

The Duality Theorem is a piece of mathematics. It requires a mathematical proof. I will spare you the details. You do not need to know the proof. One way to prove the theorem is to examine the simplex algorithm really carefully. It turns out that the algorithm solves (P) and (D) simultaneously. We will see that Excel spits out the solution to the dual whenever it solves a problem. This “proof” requires little more than high-school algebra and a willingness to tolerate a logical argument. That is, it is “elementary.” The other proof uses a tool from convex analysis called the separating hyperplane theorem. It is within the grasp of an undergraduate math major. Although it is not relevant for a Management Science major, ask and I’ll give you references.

Even though I am not going to prove the theorem, I should make an attempt to tell you why it is true. On the surface it is really amazing. Why should the two problems be so closely related? A lazy answer is that they both use the same data, so they must be related somehow. Here is a slightly better answer.

By a **feasible value** for (P) I mean a number  $v$  with that property that there is some  $x$  such that  $Ax \leq b$  and  $x \geq 0$ , such that  $v = c \cdot x$ . In words, if  $v$  is a feasible value that means that there is some  $x$  that satisfies the constraints that yields objective function value  $v$ . I define a feasible value for (D) similarly.

I claim that any feasible value for (P) is less than or equal to any feasible value for (D). In symbols:

$$\text{if } Ax \leq b, x \geq 0, yA \geq c \text{ and } y \geq 0, \text{ then } c \cdot x \leq b \cdot y.$$

You have seen the proof of this already (in the discussion of the diet problem). All you do is multiply  $Ax \leq b$  on both sides by  $y$  and multiply  $yA \geq c$  on both sides by  $x$  to get

$$c \cdot x \leq yAx \leq b \cdot y.$$

For the time being, forget about the  $yAx$  in the middle (we’ll get back to it in the next section). We can conclude,

$$c \cdot x \leq b \cdot y.$$

This is just what I claimed.

This inequality tells you much of the duality theorem. Suppose that (P) was feasible. That means that there is an  $x$  that satisfies the constraints of (P). So the value of (D) is bounded below. That is, (D), a minimization problem, cannot be unbounded. Similarly, if (D) is feasible, then (P) cannot be unbounded. The Duality Theorem goes on to say two things. First, if one problem has a solution, then the other one does. Second, if the problems have solutions, then the values are equal. These facts are mysterious, although you may gain a greater intuition for them after you have interpreted some duals.

The purpose of this subsection was to provide a bit of insight into why the Duality Theorem is true. I did this by proving an easy fact, which takes you part of the way to the conclusion of the Duality Theorem.

### 3.3 Using the Duality Theorem

The Duality Theorem tells you that the behavior of one LP is related to the behavior of another LP. One useful way to employ the theorem is to conclude that since both primal and dual are feasible, both must have solutions. For example, take the diet problem. You can see that the diet problem is feasible without computation. Provided that it is possible to supply every nutrient (in symbols, this means for each nutrient  $i$  there exists a food  $j$  such that  $a_{ij} > 0$ ), you can satisfy all nutritional constraints simply by buying enough food. (In symbols: for every nutrient  $i$  select a food  $j$  that supplies  $i$  - so that  $a_{ij} > 0$  and buy  $\frac{b_i}{a_{ij}}$  of food  $j$ .) Similarly, when the prices of food are positive, it is clear that the dual problem (the Pill Problem) is feasible: the pill seller can satisfy all constraints by setting all nutrient prices to zero. From this the Duality Theorem tells you that both problems have solutions. (This kind of reasoning does not compute the solution, of course, but it gives you a strong clue about whether the problem is well posed.) The logic is general too. It depends on some intuitive general properties of the nutritional requirements and of prices, rather than specific information about the data of the problem.

The Duality Theorem can also be a useful way to identify whether a problem is unbounded or infeasible. Consider the following pair of problems:

$$\begin{array}{rllll} \max & x_1 & + & x_2 & \\ \text{subject to} & -3x_1 & + & 2x_2 & \leq -1 \\ & x_1 & - & x_2 & \leq 2 \\ & & & x & \geq 0 \end{array}$$

The dual is:

$$\begin{array}{rllll} \min & -y_1 & + & 2y_2 & \\ \text{subject to} & -3y_1 & + & y_2 & \geq 1 \\ & 2y_1 & - & y_2 & \geq 1 \\ & & & y & \geq 0 \end{array}$$

Naturally, you could graph each of these problems and easily determine their solution. You could solve them using the simplex algorithm (either by hand or with Excel). Let me point out what common sense and the Duality Theorem tell you.

First note that the primal is feasible. For example, the point  $(x_1, x_2) = (1, 0)$  satisfies the constraints. (How did I know this? I set  $x_2 = 0$  and observed that the constraints reduced to  $\frac{1}{3} \leq x_1 \leq 2$ . There are a lot of other feasible things. You could find them all by graphing the constraints.) The point is that for small problems it is often easy to figure out a point in the feasible set by common sense. In “abstract” problems formulated from economic principles (like the diet problem), it is also possible to determine feasibility from the logic of the problem. Now that we know that the primal is feasible, what do we do next?

Since this is an exercise in using the Duality Theorem, I propose that you look at the Dual. I claim that the Dual is infeasible. In general, it is painful

to confirm it. In this example, one bit of cleverness (or a graph) confirms this. Suppose that  $y = (y_1, y_2)$  actually satisfies the constraints of the dual. It would necessarily satisfy the constraints if I added them together. Do it. If you add the two resource constraints in the dual together you get:  $-y_1 \geq 2$ . This inequality implies that  $y_1 \leq -2$ , which is inconsistent with the non-negativity constraint. Hence the dual is infeasible. From this we can immediately conclude that the Primal is infeasible. (Why? The Duality Theorem says that if one problem is infeasible, the other problem is either infeasible or unbounded. We already checked that the primal was feasible, so it must be unbounded.)

Did you think that the trick of adding the constraints in the dual was too magical? Maybe you are right. You have options: you can solve problems directly (graphing, Excel, . . .); you can practice (sometimes magic is what you call the unfamiliar); or, in this case, you can try to use common sense in a different way.

Go back to the primal. You have decided that it is feasible. Could it be unbounded (remember: we are assuming that you have not already solved the problem or figured out that the dual is infeasible)? Yes, if you can make either one or both of the variables in the objective function arbitrarily large. Can you do that? Maybe. Although  $x_1$  enters the first constraint with a positive sign and  $x_2$  enters the second constraint with a positive sign, in both constraints you can subtract the other variable. That is, maybe you can make  $x_1$  large if you make  $x_2$  large at the same time. Indeed, this is true. There are many ways to see this, but one is to imagine that  $x_1 = x_2$ . Under this assumption, the objective function is  $\max 2x_1$  and the constraints simplify to  $-x_1 \leq -1$  and  $x_1 \geq 0$  (the second resource constraint becomes  $0 \leq 2$ , which is always true). The conclusion is that the primal is unbounded. You can make the value of the problem equal to  $2K$  by setting  $x_1 = x_2 = K$ . This choice of  $x$  is feasible (provided  $K \geq 1$ ) and there is no limit to the value of the objective function you can get. The discussion in this paragraph said nothing about the dual. It was a direct, “common sense” explanation of why the original problem is unbounded. On the basis of this discussion we can conclude (using the Duality Theorem) that the dual is infeasible.

The insight is that the Duality Theorem allows you to infer something that may not be obvious. There are three kinds of inference.

1. Observe that both primal and dual are feasible. Conclude: both have solution.
2. Observe that primal is feasible and dual is not. Conclude: primal is unbounded.
3. Observe that primal is unbounded. Conclude: dual is infeasible.

(The second and third observations remain true if you interchange the words primal and dual.)

## 4 Complementary Slackness

The Duality Theorem implies a relationship between the primal and dual that is known as complementary slackness. I will try to explain the term first. Recall that the number of variables in the dual is equal to the number of constraints in the primal and the number of constraints in the dual is equal to the number of variables in the primal. This correspondence suggests that variables in one problem are **complementary** to constraints in the other. We talk about a constraint having slack if it is not binding. For an inequality constraint, the constraint has slack if the slack variable is positive. For a variable constrained to be non-negative, there is slack if the variable is positive. The term **complementary slackness** refers to a relationship between the slackness in a primal constraint and the slackness (positivity) of the associated dual variable. (Remember, this paragraph was only designed to give you an idea of where the terminology comes from.)

**Theorem 2 Complementary Slackness** *Assume problem (P) has a solution  $x^*$  and problem (D) has a solution  $y^*$ .*

1. *If  $x_j^* > 0$ , then the  $j$ th constraint in (D) is binding.*
2. *If the  $j$ th constraint in (D) is not binding, then  $x_j^* = 0$ .*
3. *If  $y_i^* > 0$ , then the  $i$ th constraint in (P) is binding.*
4. *If the  $i$ th constraint in (P) is not binding, then  $y_i^* = 0$ .*

The theorem identifies a relationship between variables in one problem and associated constraints in the other problem. Specifically, it says that if a variable is positive, then the associated dual constraint must be binding. It also says that if a constraint fails to bind, then the associated variable must be zero. The statement really is about “complementary slackness” in the sense that it asserts that there cannot be slack in both a constraint and the associated dual variable. I will outline a proof of this statement in the next section; once you have the duality theorem it is really easy to prove. It is extremely important to note that the result says that you cannot have slack in two associate places at the same time (primal variable, dual constraint) or (primal constraint, dual variable). So: it **is** possible for a primal constraint to be binding while the associated dual variable is equal to zero (that is, no slack in two places), but it is not possible for the primal constraint to have slack (to be non-binding) and the associated dual variable be positive. While I listed four statements in the theorem, there really are only two. The second two statements have precisely the same content as the first two statements, they just switch around the roles of primal and dual. (Remember, if you started with the Dual, then its dual is the original Primal.)

The theorem on Complementary Slackness is useful because it helps you interpret dual problems and dual variables, because it enables you to solve (easily) the dual of an LP knowing the solution to the primal, and because it enables you to check whether a feasible “guess” is a solution to a LP.

## 4.1 Why the Complementary Slackness Condition is True

The theorem on Complementary Slackness is really the Duality Theorem in disguise. The math behind the result is simple. You do not need to know the details, but some of you may find this subsection useful.

Early I showed that if  $x$  is feasible for (P) and  $y$  is feasible for (D), then

$$c \cdot x \leq yAx \leq b \cdot y.$$

Furthermore, if  $x^*$  solves (P) and  $y^*$  solves (D), then the first and last terms are equal, so both inequalities must really be equations:

$$c \cdot x^* = y^*Ax^* = b \cdot y^*.$$

You can write the first equation as  $(c - y^*A) \cdot x^* = 0$ . This expression can be written in more detail as  $\sum_{j=1}^n (c - y^*A)_j x_j^* = 0$  where I use  $(c - y^*A)_j$  to represent the  $j$ th component of  $y^*A$ . For each  $j$ , we know that  $(c - y^*A)_j$  is a nonnegative number (this follows because  $y^*$  is feasible for the Dual) and  $x_j^*$  is nonnegative. So when we multiply them, we must get a non-negative number. The only way that you can add up a bunch of non-negative numbers and get zero is if each one of the non-negative numbers is zero. Consequently, for each  $j = 1, \dots, n$ ,  $(c - y^*A)_j x_j^* = 0$ . This expression says that when you multiply two numbers ( $(c - y^*A)_j$  and  $x_j^*$ ) you get zero. This can only happen if at least one of the numbers is zero. And this is precisely what the complementary slackness condition says: If  $x_j^* > 0$ , then the  $j$ th dual constraint binds (that is,  $(c - y^*A)_j = 0$ ) and if the  $j$ th dual constraint does not bind (that is,  $(c - y^*A)_j > 0$ , then  $x_j^* = 0$ . You can derive the other CS conditions by applying the same kind of reasoning to the equation  $y^*Ax^* = b \cdot y^*$ .

## 4.2 Complementary Slackness and Interpretation of Dual

The most important part about the dual problem is that the dual variables provide information about quantities relevant to the original problem. You get different kinds of information. I will discuss the relationship between Complementary Slackness and the dual in this subsection. Later in these notes I will have more to say about interpretations of duals.

We have formulated a problem and its dual already. In the first week of class I presented the diet problem. The problem was to minimize the amount spent on food subject to meeting nutritional constraints. The given information for the problem were the costs of food, the nutritional requirements, and the nutrient content in each of the foods. For this problem, we also formulated the dual. I presented a rather artificial story in which the problem was to find prices of nutrient pills to maximize the cost of the nutrients needed subject to the constraints that nutrients supplied in any food can be purchased more cheaply by buying pills than by buying the food. I want you to think about the solution to the dual as providing prices of nutrients in the sense that they

represent how much the food buyer (the person in charge of the cafeteria) is willing to pay to get the nutrient directly.

Consider what the complementary slackness conditions mean in this context. Suppose that when you solve the diet problem you find that it is part of the solution to buy a positive amount of the first food. The complementary slackness condition then implies that the first dual constraint must be binding. That means that the price of (one unit of) the first food is exactly equal to the cost of (one unit of) the nutrients supplied by the food. Why should this be true? Imagine that the pill seller is really there, offering nutrients at the prices that solve the dual. If you really buy a positive amount of the first food, then it must be no more expensive to do that than to get the nutrients through pills. What if the first dual constraint is not binding? The first dual constraint states that the price of one unit of the first food should be at least as great as the cost of the nutrients inside the food. If the constraint does not bind, then it means that the food is actually more expensive than the nutrients inside it. That is, the food is strictly more expensive than its nutrient content. If this happens, then it makes no sense to buy the food:  $x_1^* = 0$ .

If a constraint in the primal does not bind, then there is some nutrient that it oversupplies when you solve the diet problem. Surely this can happen. Imagine - just for a theoretical example - that every time that you find Vitamin C in a food you also find an equal amount of Vitamin E (this is not really true!), but that you only need half as much Vitamin E as Vitamin C. In this case, if you satisfy the Vitamin C constraint, then you automatically more than satisfy (or satisfy with slack) the Vitamin E constraint. You cannot possibly satisfy all nutrient constraints as equations. Now if you found yourself in this situation (where you get more than enough Vitamin E just meeting your Vitamin C requirement), how much would you be willing to pay for a Vitamin E pill? Complementary Slackness says that the answer is zero. If your food supplies more of a nutrient than you need, then the price of that nutrient - that is, how much you are willing to pay for the nutrient - is zero. Now suppose that a dual variable is positive. In that case, you are willing to pay a strictly positive amount for nutrient supplied by the pill. Complementary Slackness says that (at a solution) it must be the case that you are supplying exactly the amount of the nutrient you need (not anything extra).

The complementary slackness conditions guarantee that the values of the primal and dual are the same. In the diet problem, the pill seller guarantees that pills are no more expensive than food. CS guarantees that when you solve the problem, pills cost exactly the same as food for those foods that you actually buy. (By the way, nothing rules out the possibility that there is a food that costs exactly as much as its nutritional content, but that you don't buy the food. If this happens, you have no slack in both primal variable and dual constraint.)

### 4.3 Using Complementary Slackness to Solve Duals

Let's take a familiar example (the problem we solved using the simplex algorithm in the third set of lecture notes).

$$\begin{array}{rllllll}
\max & 2x_1 & + & 4x_2 & + & 3x_3 & + & x_4 \\
\text{subject to} & 3x_1 & + & x_2 & + & x_3 & + & 4x_4 & \leq & 12 \\
& x_1 & - & 3x_2 & + & 2x_3 & + & 3x_4 & \leq & 7 \\
& 2x_1 & + & x_2 & + & 3x_3 & - & x_4 & \leq & 10 \\
& & & & & & & x & \geq & 0
\end{array}$$

We know that the solution to this problem is  $x_0 = 42, x_1 = 0; x_2 = 10.4; x_3 = 0; x_4 = .4$ . We can use this information (and Complementary Slackness conditions) to solve the dual. In order to do this, we first need to write the dual:

$$\begin{array}{rllll}
\min & 12y_1 & + & 7y_2 & + & 10y_3 \\
\text{subject to} & 3y_1 & + & y_2 & + & 2y_3 & \geq & 2 \\
& y_1 & - & 3y_2 & + & y_3 & \geq & 4 \\
& y_1 & + & 2y_2 & + & 3y_3 & \geq & 3 \\
& 4y_1 & + & 3y_2 & - & y_3 & \geq & 1 \\
& & & & & y & \geq & 0
\end{array}$$

The Complementary Slackness conditions typically allow you to reduce solving the dual to solving a system of equations with one unknown per equation. The equations are the constraints that correspond to positive primal variables. In this example,  $x_2$  and  $x_4$  are positive. That means that in the solution to the dual the second and fourth constraints hold as equations. That leaves us two equations and three unknowns,  $y_1, y_2, y_3$ . But there is another CS condition: if a primal constraint is not binding, then the associated dual variable must be zero. In this example, the first and third primal constraints bind, but the second one does not. Hence  $y_2 = 0$ . So let's solve the second and fourth dual constraints using  $y_1$  and  $y_3$ :

$$y_1 + y_3 = 4 \text{ and } 4y_1 - y_3 = 1.$$

It follows that  $y_1 = 1$  and  $y_3 = 3$ . At this point you can check that the vector  $y$  we have constructed,  $y = (1, 0, 3)$ , is actually feasible for the dual (that is,  $y \geq 0$  and the first and third constraints, which we have ignored up until now, are really satisfied). Moreover, the value of the dual is 42, which is equal to the value of the primal. Hence we can conclude that  $(1, 0, 3)$  is the solution to the dual.

This exercise required knowing a solution to the primal. Using the solution, we wrote down the dual and then used complementary slackness to obtain an easy problem (solving two equations and two unknowns) that led to the dual's solution. On one hand, this is an unimportant exercise: We could have solved the dual directly. On the other hand, it underlines the power of duality. Knowing the solution to the primal makes it easier to find a solution to the dual.

#### 4.4 Using Complementary Slackness to Check Optimality

In the previous section we started with a solved problem and used Complementary Slackness to find the solution to the dual. If you have a linear program,



It could happen if one of the primal guesses was zero. For example, if I had guessed that  $x = (2, 0)$  solved the primal, then I would take  $y_1 = y_3 = 0$  (first and third primal constraints not binding) and solve the first dual constraint as an equation (since  $x_1 > 0$ ) to conclude  $y_2 = 1$ . Hence I obtain a non-negative guess for the dual,  $y = (0, 1, 0)$ . This guess is not feasible, however, because it does not satisfy the second dual constraint.

## 5 Interpretation of the Dual

The Duality Theorem and the associated Theorem on Complementary Slackness give us strong reason to believe that the primal and dual are related problems. The story that I told about the diet problem and the pill problem suggest that you can construct a story that goes along with the dual that relates it the primal. I will call such a story an interpretation of the dual.

Finding interpretations of dual programs is an art rather than a science in that I cannot provide unambiguous mechanical rules for performing the interpretation. It is, however, a valuable art, and clear guidelines are available.

The main reason why interpreting the dual is useful is that dual variables provide valuable information about the primal. The information is largely what I have already described in the discussions of the Duality Theorem and Complementary Slackness, but one other insight is necessary.

In this section, I will explain the additional insight, give some general advice about how to go about interpreting the dual, and then interpret the dual of a general problem. In order to have a sense of how to interpret a dual, you need to practice. An ambitious exercise would be to take every formulation problem you have seen in the course and write and interpret the associated dual problem.

### 5.1 Dual Variables as Prices

Dual Variables are so important that they have many names. Mathematicians will call them dual variables. In more general contexts, they may call them Lagrange multipliers (or just multipliers). Economists may call them dual prices, shadow prices, or implicit prices. These terms all have “price” in them. Let me try to explain why.

In the diet problem, I interpreted dual variables as prices. Specifically, the  $i$ th dual variable was the price of the  $i$ th nutrient. It is a price in a rather peculiar way.  $y_i$  represented how much the pill seller would charge for one unit of the  $i$ th nutrient. This price depends only on the context of the problem. It may not be a market price. (The pill seller may not exist, so that the market for nutrients may not really exist.) So the prices are implicit because they may describe transactions that you cannot really make. In the problem, nutrients are important because you must satisfy nutritional constraints. You can only satisfy them using available foods. The cost of satisfying them depends on the price of the different foods. The value of nutrients would change if the price of a nutrient rich food rose (this would tend to raise nutrient prices); if a new

nutrient rich food were discovered (this might lower nutrient prices); or if the nutritional requirements changed (this is not enough information to predict the direction of change in the value of nutrients).

In general, treat dual variables as answers to the question: If the amount of the constant on the right of some constraint changes, how does the value of my objective function change? In the diet problem, the question becomes: If the nutritional requirements change, so that people are required to consume one more unit of Vitamin E each day, how does the cost of the cost minimizing diet change? This question is interesting and it does not require a story about pill sellers to ask. The answer to the question is given by the dual variable associated with the Vitamin E constraint. It gives the cost of Vitamin E from the context of the rest of the information in the diet problem.

This interpretation of the dual stems from the “surprising” part of the Duality Theorem: The fact that if you have a solution to both primal and dual, then values are equal. Suppose that you start with an LP in standard form. You find its solution and the solution to its dual. Call this the old problem and denote the solutions  $x^{old}$  and  $y^{old}$  and the associated value  $V^{old}$ . It follows that

$$V^{old} = c^{old} \cdot x^{old} = b^{old} \cdot y^{old}.$$

Create a new problem and find the solution to it. Call this the new problem and denote the solutions  $x^{new}$  and  $y^{new}$  and the associated value  $V^{new}$ . Again you have

$$V^{new} = c^{new} \cdot x^{new} = b^{new} \cdot y^{new}.$$

Now make two assumptions. First, assume that you get from the old problem to the new problem by changing one of the resource constants,  $b_i$ , by adding  $\Delta$  to it ( $\Delta$  could be positive or negative). Second, assume that when you make this change, the solution to the dual does not change, that is  $y^{new} = y^{old}$ . You are free to invent any new problem that you wish. So the first assumption is not really an assumption, it is a description of how you changed the problem. The second assumption really is an assumption. It turns out that when you change the resource constants in the primal, you usually do not change the solution to the dual (but you do typically change the *value* of the solution). On the other hand, changing the resource constraints typically do change the solution to the primal. For example, if you learned that the government requires that you eat less Vitamin C, you may buy fewer oranges.

Finally, go back to the expressions above. The change in the value of the problems,  $V^{new} - V^{old}$ , is equal to

$$b^{new} \cdot y^{new} - b^{old} \cdot y^{old}.$$

We have assumed that the dual variables do not change. We can write  $y^{new} = y^{old} = y$ . Therefore,

$$V^{new} - V^{old} = (b^{new} - b^{old}) \cdot y = \Delta y_i.$$

The second equation above comes from the assumption that you get  $b^{new}$  from  $b^{old}$  by adding  $\Delta$  to the  $i$ th component of  $b^{old}$ , while leaving the other com-

ponents unchanged. The last equation is the algebraic expression of my interpretation of dual variables. The difference in value as you go from old to new problem is equal to the change in the resource constant times the associated dual variable. Mathematically, what is neat about this is that you can figure out the change in value without solving the new LP. Mathematically, the important assumption is that when you change the LP, you do not change the solution to the dual. One expects that “small” changes in  $b_i$  do not change the solution to the dual. Hence the restriction that the dual solution does not change means that the interpretation of dual variables remains valid only for “small” changes in the level of resources available. When Excel solves an LP, it will tell you just how much you can change the amounts of available resources without changing the solution to the dual.

## 5.2 Interpreting the Dual: Watch your Units

I cannot give you precise rules for interpreting a dual LP. One feature is common to all problems. Since the values of primal and dual are equal, they have the same units. Usually, values come in units of money. In the diet problem, the value of the primal is the amount you spend on food. The value of the dual is the amount the pill seller can earn selling pills. In both cases, the units are monetary units. You arrive at the units differently. You are typically given the data of the problem ( $A$ ,  $b$ , and  $c$ ). This means that you know the units of these quantities. You define the primal variables,  $x$ . These variables must have units that are compatible with the rest of the problem. In the diet problem, the objective function coefficients are in the units of price per quantity of food (for example, hamburger sells for \$2.00 per pound). In order to translate this into monetary units, you must multiply it by a quantity of food (so that  $x_1$  might be a number of pounds of hamburger). When you formulate a problem, you define the primal variables, but the problem itself “forces” you to choose the fundamental unit. The same insight applies to the dual variables. You pick them, but the problem itself has determined their units. In the diet problem you know that the nutrient constraints involve requirements  $b_i$  that are given in units like “grams of Vitamin C.” Since  $b_i y_i$  must be an amount of money (as it is a term in the objective function), it must be the case that  $y_i$  is in units of “dollars per unit of Vitamin C.” Keeping this in mind does not provide a description of the dual, but it does give you a start. It also allows you to recognize nonsense (inappropriate units) when you see them.

## 5.3 An Example: Production Problem

A standard kind of LP is a problem that asks you to find a production plan that maximizes profit from a variety of processes given available resources. You have seen versions of this general problem. One of the practice formulation problems asks you to mix together varieties of nuts to make different blends, that can be sold at different prices. On the first homework assignment you figured out how many of different varieties of coffin to make. In general, imagine that there

are  $n$  different production activities. If you operate activity  $j$  at unit level (for example, you make one can of the deluxe mixture of nuts), then you can sell it for  $c_j$ . Your production processes use  $m$  different basic resources. You have the amount  $b_i$  of resource  $i$ . When you operate activity  $j$  at unit level, you use up some of the basic resources. The technology matrix  $A$  describes this information. The entry  $a_{ij}$  of the matrix  $A$  is the amount of basic resource  $i$  needed to operate the  $j$ th activity at unit level. The problem is find a production plan that maximizes profit using only the available resources. Assuming that the technology exhibits constant returns to scale and additivity, this becomes a standard kind of LP. Find  $x = (x_1, \dots, x_j, \dots, x_n)$  to solve:

$$\max c \cdot x \text{ subject to } Ax \leq b, x \geq 0.$$

The objective function just adds up the profits earned from operating each of the activities. That is, if you operate activity 1 at level  $x_1$ , then you earn  $c_1x_1$  from it. Total profit comes from adding the profit of each of the activities. The resources constraints state that the production plan does not use up more of any resource than is available. If you follow the production plan  $x$ , then

$$\sum_{j=1}^n a_{ij}x_j$$

is the amount of the  $i$ th resources consumed. So you need to have

$$\sum_{j=1}^n a_{ij}x_j \leq b_i$$

for all  $i = 1, \dots, m$ . The non-negativity constraint states that you cannot operate a production process at a negative level.

Once you have formulated the problem, you can ask whether it is feasible and, if so, whether the LP has a solution. In this case, feasibility is satisfied if (as seems sensible) all resources are available in non-negative quantities (that is,  $b \geq 0$ ). In this case,  $x = 0$  is an element of the feasible set. (The feasible set may be non-empty even if one or more of the  $b_i$  are negative. Feasibility would be harder to check in this case. Besides, this case doesn't fit with the story.) The LP will have a solution unless it is unbounded. In order for the LP to be unbounded, you would have to be able to produce something in arbitrarily large amounts (that is, one of the components of  $x$  would be able to grow to infinity). A simple way to rule this out is to assume that every entry in the matrix  $A$  is strictly positive. There are weaker assumptions that would be sufficient to guarantee that the problem is bounded.

Since the production problem is an LP in standard form, you know that the mathematical form of the dual is

$$\min b \cdot y \text{ subject to } yA \geq c, y \geq 0.$$

The point of this section is to give this problem an interpretation. Imagine that you are the owner of a firm that can operate the  $n$  production processes,

can sell output at prices  $c$ , and has available the resources  $b$ . If you choose to use your production process, then you should pick a production plan  $x$  to solve the original LP. Now imagine that someone offers to buy out your inventory of basic resources ( $b$ ). If you sell, then you can't operate your production process. How might this outsider convince you to sell out? One way to do it is if the potential buyer set prices for each of these resources. Call the unit price of the  $i$ th resources  $y_i$ . The buyout artist comes to you and says: "I want to buy your complete inventory. I will pay you  $y_i$  for each unit of the  $i$ th resource. Moreover, I will set my prices high enough so that you earn at least as much selling your resources to me as you would turning your resources into final products (using your production technology) and then selling them at the prices  $c$ ." The buyout artist then argues that her prices for resources are sufficiently high that the price of a unit of final output,  $c_j$  is always less than the price she'll pay for the ingredients used in that output. Algebraically, she guarantees that for each  $j = 1, \dots, n$ ,

$$\sum_{i=1}^m a_{ij}y_i \geq c_j.$$

Of course, this constraint is just a representative constraint of the dual. Hence the dual of the problem can be interpreted as the buyout artist's problem. Set prices (of basic resources)  $y = (y_1, \dots, y_i, \dots, y_m)$  to minimize what the buyout artist pays to acquire the resources ( $y \cdot b$ ) subject to the constraint that the buyout artist always pays at least as much for the resources as you could get from transforming the resources into a final product.

This kind of story should sound familiar. The tale I told about the dual of the diet problem was almost the same. The nature of the story should tell you things that you know from the Duality Theorem. The value of the original problem must be no greater than the value of the dual. The surprising implication of the Duality Theorem is that when you solve these problems the values are equal.

The dual variables in this formulation are prices of the basic inputs. Although they are prices, they may have nothing to do with market supply and demand. They depend only on the basic information of the problem ( $A, b, c$ ). The sense in which they are prices is that they value the inputs to the production process. That is, they provide answers to questions like: How much is the first ingredient worth to you (if you have access to the technology that turns inputs into outputs according to the matrix  $A$ ,  $b$  is your list of available inputs, and  $c$  are the prices at which you can sell final outputs). If  $A$ ,  $b$ , or  $c$  changes, then you would expect the "price" of ingredients to change too.

The buyout story is contrived, but it is accurate. If someone offers to buy a little bit of the first input, then you would accept if you are offered  $y_1$  or more (per unit). Similarly, if you could buy a more of the first input at less than  $y_1$ , it would be profitable for you to do so. The identity between the value of the primal and the value of the dual tells you that the ingredients, when evaluated according to the prices  $y$ , are worth exactly as much as the final product.

Let me repeat the Complementary Slackness exercise. If when you solve the production process, you find that you don't use all of your supply of the first

resource, that means that you would gladly sell a bit of that resource at *any* positive price. On the other hand, you won't not pay a positive amount for more of this resource. Hence  $y_1 = 0$ . Notice that this doesn't mean that you would never place a positive value on the first resource. If you have 100 units of the resource and your production plan uses 90, then you would be willing to give away (sell for the price  $y_1 = 0$ ) ten units, but further sales would interfere with your production plan. When the first input becomes scarce enough, its price becomes positive. Similar, you would expect that if prices of outputs change, then the value of the inputs would change as well. Suppose that the first input is the primary ingredient of the  $n$ th output. If the price of the  $n$ th output goes up, you might want to change your production plan and make more of the  $n$ th output. This would increase your demand for the first input. If demand increases enough to make the first resource constraint bind, then  $y_1$  would become positive.

You can tell similar stories for the other Complementary Slackness conditions. For example, suppose that you produce positive amounts of the  $n$ th good when you solve the problem ( $x_n > 0$ ). The dual constraint says that the amount you can sell the  $n$ th good for is no greater than the value of its ingredients. If you really could sell the ingredients, then you would do so (instead of setting  $x_n > 0$  unless the  $n$ th constraint of the dual is binding. Conversely, if a dual constraint fails to bind, then the corresponding production level in the primal must be zero.

It is useful to think of the dual variables as prices. The prices tell you the value of the resources from the point of view of the technology and output prices that describe the primal. Knowing the dual variables tells you how much you gain (lose) from an increase (decrease) in one of the basic resources.

## 6 Hints on Writing Duals

This section is optional. Knowing its contents gives you insight into duality, but you can do everything I want you to be able to do without knowing its contents. (The next paragraph explains why.)

You know how to write the dual of any LP written in the form:

$$\max c \cdot x \text{ subject to } Ax \leq b, x \geq 0.$$

Furthermore, you know that you can write any LP in this form. Hence, in theory, there is no problem: you can find the dual of any LP. These general tricks are good to know, and they enable you to write the dual of any problem, but in practice it is useful to recognize a few short cuts.

Imagine that you have a linear programming problem in which the objective is to maximize something and the constraints are either equations or "less than or equal to a constant ( $\leq b_i$ )." Some variables may be explicitly constrained to be non-negative. Others may be free (unconstrained). In this case, you'll be able to write the dual using one variable for each resource constraint in the

primal. The dual objective will be to minimize  $\sum_{i=1}^m b_i y_i$  (where there are  $m$  resource constraints in the primal; the  $b_i$  are the right-hand side constants; and the  $y_i$  are the dual variables). The constraints of the dual will look the same as usual, except that unconstrained variables in the primal give rise to equality constraints in the dual (constrained primal variables give rise to  $\geq c_j$  constraints as before). Further, the dual variables corresponding to inequality constraints in the primal are explicitly constrained to be non-negative, while the dual variables corresponding to equality constraints in the primal are not constrained at all.

Primal	Dual
max	min
equality constraint	unconstrained variable
$\leq$ constraint	unconstrained variable
non-negative variable	$\geq$ constraint
unconstrained variable	equality constraint

You can confirm these claims by taking the original problem, transforming it to the standard form, taking the dual of that problem, and then simplifying what you get. These steps are tedious, but a useful exercise in writing the dual.

It is useful to try to understand the transformation knowing what we know about dual variables and complementary slackness. Take a constraint of the form “blah, blah, blah  $\leq b_i$ ” in a maximization problem. You know several things about the associated dual variable,  $y_i$ . First,  $y_i \geq 0$ . Second,  $y_i$  measures the increase in value of the objective function in the primal associated with an increase in the right-hand side constant  $b_i$ . Third, by complementary slackness, if  $y_i > 0$ , then the constraint binds and if the constraint does not bind, then  $y_i = 0$ . These properties are related. The second one is really the most powerful. If you accept that  $y_i$  measures how the primal objective function will change if  $b_i$  increases, then you know that  $y_i$  must be nonnegative. Increasing  $b_i$  (when the  $i$ th constraint in of the form “blah, blah, blah  $\leq b_i$ ” must not lower the value of the objective function (you can use the original solution). If changing  $b_i$  changes your solution, then your value goes up. If changing  $b_i$  does not change your solution, then your value stays the same. Going down is not possible. So  $y_i \geq 0$ . Now consider the third property. If when you solve the original problem, the  $i$ th constraint does not bind, then increasing  $b_i$  more is not going to help you. Hence  $y_i = 0$ . On the other hand, if  $y_i > 0$ , then it must be the case that a little bit more  $b_i$  would help you and a little bit less  $b_i$  would hurt you. So, you must be using all of the  $i$ th resource.

Consider these statements when the  $i$ th constraint is an equation. Continue to assume that  $y_i$  measures how the primal objective function will change if  $b_i$  increases. Now you have no way of knowing whether an increase in  $b_i$  is good or bad. That is, there is nothing in the problem that tells you what sign  $y_i$  should have. Hence it makes sense that the dual variable associated with an equality constraint should be unconstrained. Also, you know that the complementary

slackness condition always holds no matter what  $y_i$  is equal to (the primal constraint **must** be binding).