Introduction to duality

Consider the manufacturing problem max $\boldsymbol{c}\cdot\boldsymbol{x}$, s.t. $A\boldsymbol{x}\le\boldsymbol{b}$, $\boldsymbol{x},\mathbb{R}^n_+$, $\boldsymbol{c}\in\mathbb{R}^n$, $\boldsymbol{b}\in\mathbb{R}^m$, where the variables x_j, $j=1,\ldots,n$ are the amounts (not necessarily integer) of goods j to be produced, out of m raw materials (resources) $i=1,\ldots,m$. (Notation-wise, further instead of dot-products like $\boldsymbol{c}\cdot\boldsymbol{x}$, $y\cdot\boldsymbol{b}$ one may use the matrix multiplication notation $\boldsymbol{c}^T\boldsymbol{x}=\boldsymbol{x}^T\boldsymbol{c}$, $\boldsymbol{y}^T\boldsymbol{b}=\boldsymbol{b}^T\boldsymbol{y}$, always meaning that the vector notations $\boldsymbol{x},\boldsymbol{c},\boldsymbol{y},\boldsymbol{b}$ are also used for matrices consisting of a single column, while their transposes are matrices consisting just of a single row.)

The inequalities in the manufacturing problem reflect the fact that there is a maximum of b_i units of the raw material i available, the entry a_{ij} of the matrix A is the amount of raw material i it takes to manufacture a unit of the product j. The component c_j of $\boldsymbol{c}$ is the market price, at which a unit of the good j is to be sold. The manufacturer's objective is to find the optimal solution $\boldsymbol{x}$ of the above LP, given $(A,\boldsymbol{b},\boldsymbol{c})$. We do not require non-negativity of $\boldsymbol{b},\boldsymbol{c}$. In fact, in the stock market, one can have negative asset amounts by short selling.

The above manufacturing problem is further referred to as *primal*. To it we will assigned a *dual* problem, which is a diet problem, which is obtained from the primal by replacing A with A^T and swapping $\boldsymbol{b}$and $\boldsymbol{c}$. Here is some "economic" motivation for the dual problem. Consider a market that gives one no guaranteed opportunity to make money. If the manufacturer is to buy the raw materials at the market from some seller, at a unit price $y_i\ge0$, they may expect that the vector $\boldsymbol{y}\in\mathbb{R}^m_+$ will satisfy the following inequalities: $A^T\boldsymbol{y}\ge\boldsymbol{c}$, $\boldsymbol{y}=(y_1,\ldots,y_m)$. Indeed, if $y_i\ge0$ is the unit price of the raw material i, then the left-hand-side of the jth inequality of $A^T\boldsymbol{y}\ge\boldsymbol{c}$, is simply how much it will cost to buy an exact basket of raw materials (the raw materials are an abstraction, they may include labour, salaries, etc.) in order to produce a unit of good j. Had the jth inequality gone the other way, it would mean a guaranteed profit for the manufacturer after manufacturing the "underpriced" jth good and selling it. (Economists call a guaranteed profit *arbitrage* and believe that real market prices close arbitrage opportunity fairly quickly, since as soon as there is an opportunity to make something out of nothing many people discover it and the opportunity disappears. So if it were $a_{1j}y_1+a_{2j}y_2+\ldots+a_{mj}y_m<c_j$, everyone would start buying the raw materials in the right proportion, in order to manufacture this (underpriced) good and sell it, this would cause the selling price c_j go down and thereby close the arbitrage opportunity. Out theory though considers $\boldsymbol{c}$ and $\boldsymbol{b}$ as fixed. Practically though, everyone is looking for short-term market arbitrage opportunity, and the economical aspect of the duality theory is at least that it enables one to identify such an opportunity.)

Thus, in a market with no arbitrage opportunities, to be able to buy the raw materials, in order to have at least b_i of each (to be able to satisfy *any* feasible production strategy for the manufacturing problem above), one should be ready to pay the amount y_i per unit of raw material i, which is feasible for the *dual* inequalities $A^T\boldsymbol{y}\ge\boldsymbol{c}$, $\boldsymbol{y}\in\mathbb{R}^m_+$, and target the objective of minimizing $\boldsymbol{b}\cdot\boldsymbol{y}$. This is a diet problem. We call it *dual* to the original manufacturing problem.

Taking the transpose of the above (dual) system of inequalities, as well as the (dual) objective, leads to the equivalent form min $\boldsymbol{y}^T\boldsymbol{b}$, s.t. $\boldsymbol{y}^TA\ge\boldsymbol{c}^T$, $\boldsymbol{y}\in\mathbb{R}^m_+$. The latter is useful to write some proofs, like weak duality below, but certainly they can all be written via dot products only, the key to translating one notation into the other is the identity $\boldsymbol{y}\cdot(A\boldsymbol{x})=(A^T\boldsymbol{y})\cdot\boldsymbol{x}$. Observe that the constraints above $\boldsymbol{y}^TA\ge\boldsymbol{c}^T$ are written in a row – one after the other – rather than a column – one under the other. The latter equivalent form is $A^T\boldsymbol{y}\ge\boldsymbol{c}$, after taking the transpose. Recall that taking transposes reverses the order of matrix multiplication.

Let us now study the mathematical relation between the two problems. So, we have a problem max $\boldsymbol{c}\cdot\boldsymbol{x}$, s.t. $A\boldsymbol{x}\le\boldsymbol{b}$ and assign to it the dual problem min $\boldsymbol{b}\cdot\boldsymbol{y}$, s.t. $A^T\boldsymbol{y}\ge\boldsymbol{c}$ by the following rule: $\boldsymbol{b}$ and $\boldsymbol{c}$ get swapped, the matrix is transposed, the feasibility inequalities get reversed.

- Involution property of the duality relation: The dual of the dual is the primal: verify this by rewriting the dual as max $-\boldsymbol{b}\cdot\boldsymbol{y}$, s.t. $-A^T\boldsymbol{y}\le -\boldsymbol{c}$, $\boldsymbol{y}\in\mathbb{R}^m_+$, (a manufacturing problem) and then looking at its dual by the above rule. Indeed, writing the dual (of the dual) as the diet problem form, we replace $-A^T$ with its transpose, and swap $-\boldsymbol{c}$ and $-\boldsymbol{b}$. This yields $-A\boldsymbol{x}\ge -\boldsymbol{b}$, min $-\boldsymbol{c}\cdot\boldsymbol{x}$, which is equivalent to the original manufacturing problem.

- Weak duality – by the dual inequalities, while making each individual good the manufacturer will at best break even: If $\boldsymbol{x}$ is feasible for the primal and $\boldsymbol{y}$ is feasible for the dual, then $\boldsymbol{y}\cdot\boldsymbol{b}\ge\boldsymbol{c}\cdot\boldsymbol{x}$. That is a haphazard pair of feasible strategies $(\boldsymbol{x},\boldsymbol{y})$ of manufacturing goods/buying resources would most likely lead to a loss. Indeed, one can multiply the primal inequalities by $\boldsymbol{y}^T$ on the left, getting $\boldsymbol{y}^TA\boldsymbol{x}\le\boldsymbol{y}^T\boldsymbol{b}$. This preserves the inequality sign, because $\boldsymbol{y}\ge 0$. Then, as $\boldsymbol{y}^TA\ge\boldsymbol{c}^T$ ($\boldsymbol{y}$ is feasible for the dual) and $\boldsymbol{x}\ge 0$, one has $\boldsymbol{c}^T\boldsymbol{x}\le\boldsymbol{y}^TA\boldsymbol{x}\le\boldsymbol{y}^T\boldsymbol{b}$, and it does the job. In the sequel, let $V(\boldsymbol{x})\equiv\boldsymbol{c}^T\boldsymbol{x}$, $V(\boldsymbol{y})=\boldsymbol{y}^T\boldsymbol{b}$ be referred to as *values* of $\boldsymbol{x}$ and $\boldsymbol{y}$.

  Here is the above proof of Weak duality with the dot product notation: $\boldsymbol{y}\cdot A\boldsymbol{x}\le\boldsymbol{y}\cdot\boldsymbol{b}$, now $\boldsymbol{y}\cdot A\boldsymbol{x}=\boldsymbol{x}\cdot A^T\boldsymbol{y}\ge\boldsymbol{c}\cdot\boldsymbol{x}$.

  An immediate consequence of week duality is that if one of the manufacturing/diet problems pair is unbounded – either there are feasible $\boldsymbol{y}$'s so that $V(\boldsymbol{y})\to-\infty$ or feasible $\boldsymbol{x}$'s so that $V(\boldsymbol{x})\to+\infty$, then the other is unfeasible.

- Strong duality – if both problems are feasible, then whatever $\boldsymbol{b}$ and $\boldsymbol{c}$ the manufacturer always has the best strategy to break even: If $\boldsymbol{x}$ is optimal for the primal (a maximizer; 'a', because there can be more than one in the alternative solutions case) and $\boldsymbol{y}$ is optimal for the dual (a minimizer), then the manufacturer breaks even: $\boldsymbol{y}\cdot\boldsymbol{b}=\boldsymbol{c}\cdot\boldsymbol{x}\equiv V$, the value of the LP. *This, although intuitively clear, takes quite a bit of work to prove and will be done later in the course. For now, we assume the Strong Duality Theorem.*

  Strong duality, paraphrase: Suppose, $\boldsymbol{x}$ is feasible for the primal, $\boldsymbol{y}$ is feasible for the dual, and $V(\boldsymbol{y})=V(\boldsymbol{x})$. Then $(\boldsymbol{x},\boldsymbol{y})$ are optimal for the primal/dual.

Complementary slackness

It turns out that the dual problem provides an *optimality test* enabling one to conclude whether a given $\boldsymbol{x}$ (a manufacturer hires a consultant who recommends the manufacturing strategy $\boldsymbol{x}$ ... is he right?) is optimal for the primal problem or not. This follows from so-called *Complementary Slackness theorem* which is a practical equivalent of the Strong Duality theorem.

Its formulation is as follows. *Let $\boldsymbol{x},\boldsymbol{y}$ be optimal for the primal and dual, respectively. Then if $x_j>0$, then $\boldsymbol{y}$ satisfies the jth dual constraint tightly, i.e as equation, not inequality. And if $y_i>0$, then $\boldsymbol{x}$ satisfies the ith primal constraint tightly, i.e as equation, not inequality.*

This makes it practical: we know how to solve systems of linear equations with any number of unknowns!

Next come the proof of Complementary slackness, using Weak duality and assuming Strong duality, as well as a number of equivalent formulations, with more LP slang.

Let $\boldsymbol{x},\boldsymbol{y}$ be *optimal* solutions for the primal/dual, respectively.

1. Suppose, the component x_j of the optimizer $\boldsymbol{x}$ is not zero (that is some positive amount of the good j is to be manufactured). Then the good is *not underpriced*, that is the corresponding dual inequality is *tight*: one has $y_1a_{1j}+y_2a_{2j}+\ldots+y_ma_{mj}=c_j$ (rather than $>c_j$). We further say that a non-strict ($\le$ or $\ge$) inequality is *tight* if it is satisfied as an equation $=$ and *slack* if it is satisfied as a strict inequality ($<$ or $>$).

Proof: Suppose this is not the case, that is $x_j > 0$ and $y_1 a_{1j} + y_2 a_{2j} + \ldots + y_m a_{mj} - c_j = r_j > 0$. Then $\boldsymbol{y}$ is still *feasible* for the problem min $\boldsymbol{y} \cdot \boldsymbol{b}$, s.t. $A^T \boldsymbol{y} \geq \tilde{\boldsymbol{c}}$, where $\tilde{\boldsymbol{c}}$ is the same as $\boldsymbol{c}$, except that its jth component has been increased to $c_j + r_j$. Modify the primal problem accordingly, by changing c_j to $c_j + r_j$ in its objective function. Clearly, $\boldsymbol{x}$ is still *feasible* for this modified primal. So if before modifying $\boldsymbol{c}$ to $\tilde{\boldsymbol{c}}$ we had $\boldsymbol{y} \cdot \boldsymbol{b} = \boldsymbol{x} \cdot \boldsymbol{c}$, now we have $\boldsymbol{y} \cdot \boldsymbol{b} < \boldsymbol{x} \cdot \tilde{\boldsymbol{c}}$, which contradicts weak duality for the modified primal-dual pair.

   Equivalently – just logic – if a good j is underpriced, that is $y_1 a_{1j} + y_2 a_{2j} + \ldots + y_m a_{mj} - c_j = r_j > 0$, then it is not manufactured, that is $x_j = 0$.

2. In the same fashion, suppose, the component y_i of the optimizer $\boldsymbol{y}$ is not zero (that is the material i actually costs some money). Then one cannot afford to buy the excess amount of this material (everything that is bought must be used up): the corresponding primal inequality must be *tight*: $a_{i1} x_1 + a_{i2} x_2 + \ldots + a_{in} x_n = b_i$ (rather than $<$).

   The proof repeats the previous one word by word (as there is no preference in the primal/dual relation as to which problem – the manufacturing or diet – comes first). Indeed, assume to the contrary that $y_i > 0$ and $b_i - (a_{i1} x_1 + a_{i2} x_2 + \ldots + a_{in} x_n) = e_i > 0$. Then $\boldsymbol{x}$ is still feasible for the modified primal problem max $\boldsymbol{c} \cdot \boldsymbol{x}$, s.t. $A\boldsymbol{x} \leq \tilde{\boldsymbol{b}}$, where $\tilde{\boldsymbol{b}}$ is the same as $\boldsymbol{b}$, except that its ith component is reduced to $b_i - e_i$. Modify the dual problem, just by changing b_i to $b_i - e_i$ in its objective function. Clearly, $\boldsymbol{y}$ is still feasible for the modified dual. Look at the objective values, delivered by $\boldsymbol{x}$ and $\boldsymbol{y}$ for the modified primal-dual pair. While $\boldsymbol{x}$ delivers the same objective value $V = \boldsymbol{c} \cdot \boldsymbol{x}$ for the modified primal problem, $\boldsymbol{y}$ delivers the value $\boldsymbol{y} \cdot \tilde{\boldsymbol{b}}$ for the dual, which is *smaller* than $\boldsymbol{c} \cdot \boldsymbol{x} = \boldsymbol{b} \cdot \boldsymbol{y}$ by $y_i e_i > 0$. This contradicts weak duality for the modified primal-dual pair.

   Conversely, if a material i is bought in excess, that is $b_i - (a_{i1} x_1 + a_{i2} x_2 + \ldots + a_{in} x_n) = e_i > 0$, then it is available for free, that is $y_i = 0$.

To this effect, there is some terminology.

1. In the primal optimal strategy $\boldsymbol{x}$ a component j is called *basic* if $x_j > 0$. In the dual optimal strategy $\boldsymbol{y}$, a component i is called basic if $y_i > 0$. Non-basic components are called *free*.

2. The quantity $y_1 a_{1j} + y_2 a_{2j} + \ldots + y_m a_{mj} - c_j = r_j$ is called the *reduced cost* of the good j. If j is basic, then by complementary slackness, $r_j = 0$. Conversely, if $r_j > 0$, then j is free.

   If j is not basic, one can only say that $r_j \geq 0$ (although it is most likely going to be strictly positive). The meaning of the reduced cost r_j is simple. It is by how much the good is underpriced, that is by how much its market price should increase, so that it would make sense to manufacture it. Indeed, suppose $r_j > 0$. Suppose, the price of the good j is increasing from its original value c_j to $c_j + r_j$. Then $\boldsymbol{x}$ is always feasible for the (modified) primal, while $\boldsymbol{y}$ is always feasible for the (modified) dual. Moreover, the objective values $\tilde{V}(\boldsymbol{x})$ and $\tilde{V}(\boldsymbol{y})$ that $\boldsymbol{x}$ and $\boldsymbol{y}$ deliver respectively for the modified primal and dual problems remain the same (because $x_j = 0$, so the change of c_j does not matter for the objective value in the primal), equal to $V = \boldsymbol{c}^T \boldsymbol{x} = \boldsymbol{y}^T \boldsymbol{b}$. This means, by strong duality, that the strategies $(\boldsymbol{x}, \boldsymbol{y})$ all this time remain optimal. But this is no longer true if c_j is increased by more than r_j : indeed, $\boldsymbol{y}$ will no longer be feasible for the modified dual. Thus, one will have to change the optimal strategies and most likely start manufacturing the good j.

   Complementary slackness: basic goods have zero reduced cost. Or, nonzero reduced cost $\Rightarrow$ the good is non-basic.

3. In the same fashion, the quantity $b_i - (a_{i1} x_1 + a_{i2} x_2 + \ldots + a_{in} x_n) = e_i \geq 0$ is called the *excess amount* of the material i. Repeating the last block of reasoning, e_i is the maximum amount of the

material i, which can be thrown away without affecting the optimal strategy or the value of the problem.

Complementary slackness: basic resources come in zero excess amount. Or, positive excess amount $\Rightarrow$ the resource is free.

4. The quantity y_i is called the *shadow price* of the constraint i. The reason is the above economical interpretation: y_i is the price the manufacturer might be ready to pay for a unit of the resource i, in order to break even.

   Suppose, the available amount of the raw material i decreases by a (small) quantity ϵ_i. Clearly, $\boldsymbol{y}$ remains feasible for the (modified) dual, with its objective value having decreased by $\epsilon_i y_i$. This means (by weak duality) that the value of the primal problem will decrease by *at least* the same amount $\epsilon_i y_i$. It will be exactly the same amount, provided that $\boldsymbol{y}$ remains *optimal* for the modified dual (which is usually the case for a small ϵ_i: in this case, having lost ϵ_i units of the raw material i, the manufacturer shall be able to adjust the production strategy $\boldsymbol{x}$ in order to ensure that the sales do not decrease by more than $\epsilon_i y_i$ (which is just y_i if ϵ_i is 1, thus the term "shadow price").

Note: if a constraint (resource) has a zero shadow price (is free), it does not *necessarily* come in the excess amount. Just the same, if a good is not basic (not manufactured), it is not always true that it must have a nonzero reduced cost. However, the situations when the above is not true are exceptional (degenerate) and correspond to the case of alternative solutions, which will be briefly mentioned later in the course.

Complementary slackness as optimality test

The easiest way to see this is to consider an example. Take a problem of weekly manufacturing Desks, Tables and Chairs, selling respectively for £60, 30 and 20, provided that: (i) manufacturing a single Desk takes 8 square feet of wood, 4 hours of carpentry, and 2 hours of finishing; (ii) manufacturing a single Table takes 6 square feet of wood, 2 hours of carpentry, and 1.5 hours of finishing; (iii) manufacturing a single Chair takes 2 square feet of wood, 1.5 hours of carpentry, and .5 hours of finishing. The amounts of 48 square feet of wood, 20 carpentry hours and 8 finishing hours are available per week. Besides, a weekly demand for Tables is at most 5.

Suppose, you hire a consultant who says that the optimal strategy is to produce 2 Desks, 0 Tables and 8 Chairs a week. Shall we trust him? And if yes, what is the minimum increase in the market price of Tables, sufficient to make one start manufacturing them, the rest of the data remaining unchanged.

Solution:

1. First, the problem itself: if x_1 is the number of desks, x_2 of tables, and x_3 of chairs to be manufactured, then max $60x_1 + 30x_2 + 20x_3$, s.t. $\boldsymbol{x} \geq 0$ and $\begin{cases} 8x_1 + 6x_2 + 2x_3 \leq 48 \\ 4x_1 + 2x_2 + 1.5x_3 \leq 20, \\ 2x_1 + 1.5x_2 + .5x_3 \leq 8, \\ x_2 \leq 5. \end{cases}$

   Equivalently $\mathsf{Max}\ \boldsymbol{c} \cdot \boldsymbol{x}$, s.t. $\boldsymbol{x} \geq 0, \quad A\boldsymbol{x} \leq \boldsymbol{b}$.

2. Dual: $\boldsymbol{y} = (y_1, y_2, y_3, y_4)$ and $\mathsf{Min}\ \boldsymbol{y} \cdot \boldsymbol{b}$, s.t. $\boldsymbol{y} \geq 0, \quad A^T \geq \boldsymbol{c}$. Write it out explicitly!

3. First, let us see if $\boldsymbol{x} = (2, 0, 8)$ is at all feasible. Plug it into the primal problem. Feasible it is: $32 < 48$, $20 = 20$, $8 = 8$, $0 < 5$. The first and the last constraints are slack, having the excess amount of 16 and 5, respectively. The second and third are tight. Let us now go the complementary slackness trail, and as long as we do not run into contradiction, $\boldsymbol{x}$ is optimal.

Namely, if $\boldsymbol{x}$ is optimal, the shadow price of the slack constraints equals zero, by complementary slackness. I.e. the minimizer $\boldsymbol{y} = (y_1, y_2, y_3, y_4)$ for the dual has $y_1 = y_4 = 0$. Furthermore, as x_1 and x_3 are *basic* (i.e. positive), the first and the third inequalities in the dual must be tight. That is (y_2, y_3) are to satisfy

$$\begin{cases} 4y_2 + 2y_3 = 60, \\ 1.5y_2 + .5y_3 = 20. \end{cases}$$

So $y_2 = y_3 = 10$.

4. So far so good. And by construction the value $\boldsymbol{b} \cdot \boldsymbol{y}$ of $\boldsymbol{y}$ equals the value $\boldsymbol{c} \cdot \boldsymbol{x}$ of $\boldsymbol{x}$ equals 280. So, by strong duality $\boldsymbol{x}$ above (as well as $\boldsymbol{y}$) will be optimal, provided that $\boldsymbol{y}$ is, in fact, *feasible* for the dual – we have only verified the dual inequalities that we concluded by looking at $\boldsymbol{x}$ must be tight, but how about the rest?

   Substituting $\boldsymbol{y} = (0, 10, 10, 0)$ into the remaining second dual inequality yields $35 > 30$, which is true. So. $\boldsymbol{x}$ is optimal for the primal, as we have found a feasible solution $\boldsymbol{y}$ of the dual that has the same value, and this may only hap[pen when both $\boldsymbol{x}$ and $\boldsymbol{y}$ are optimal.

   Finally, the reduced cost of Tables is by how much the right hand-side in the second inequality of the dual problem increases a Table's market price: $r_2 = 2 * 10 + 1.5 * 10 - 30 = 5$. So, as a Table price climbs up to £35 or more, one should start manufacturing them. Note: the case when it is exactly £35 would lead to the alternative solutions situation.

   Hence – general strategy for optimality test: $\boldsymbol{x}$ is *optimal* for the primal if the $\boldsymbol{y}$ computed from $\boldsymbol{x}$ using complementary slackness is *feasible* for the dual.