Farkas alternative and Duality Theorem

There are theorems, called *alternatives.* They say that there are two and only two possibilities for something, one of these possibilities must take place, and they can't happen together. An example of an alternative is, assuming that there is no state between life and death: a human being is alive or dead. One side of the alternative, say alive, is called the *obverse*, and the other, in this case dead, the *reverse.* In computer science, they use the word *xor* for exclusive *or.* Namely Mary *xor* Jane means Mary or Jane, but not the two of them.

This set of notes proves one such theorem, called the *Farkas alternative* and shows that, in fact, it underpins all the duality theory of linear programming. It underlies, in fact, most of optimisaiton, itself being a particular case of the Separating Hyperplane Theorem.

First, some definitions (strictly speaking unnecessary, but *de bon ton*).

Cones. A set $\mathcal{C} \subseteq \mathbb{R}^n$ is called a *cone* if for any $\boldsymbol{x} \in \mathcal{C}$, one has $\lambda \boldsymbol{x} \in \mathcal{C}$ for all $\lambda \geq 0$. This means that the origin O is in C and, geometrically, if any $\boldsymbol{x} \neq 0$ is in $\mathcal{C}$, then the whole ray $O\boldsymbol{x}$ is in $\mathcal{C}$. With such a general definition, a cone is not necessarily a closed or convex set.

Exercise: Later we shall deal with the notion of the *dual cone.* Namely, if $\mathcal{C}$ is any cone in $\mathbb{R}^n$, it dual cone C^* is defined as $\mathcal{C}^* = \{\boldsymbol{y} \in \mathbb{R}^n : \boldsymbol{y} \cdot \boldsymbol{x} \geq 0, \forall \boldsymbol{x} \in \mathcal{C}\}$. Geometrically $\mathcal{C}^*$ contains all vectors $\boldsymbol{y}$, such that the angle between $\boldsymbol{y}$ and *any* vector $\boldsymbol{x} \in \mathcal{C}$ is ninety degrees or less. From this point of view, it is clear that $(\mathcal{C}^*)^* = \mathcal{C}$, so the dual of the dual is primal. Show this by definition.

If A is an $m \times n$ matrix, consider the set

$$\mathcal{C}_A = \{\boldsymbol{y} \in \mathbb{R}^m : \boldsymbol{y} = A\boldsymbol{x}, \, \boldsymbol{x} \in \mathbb{R}^n_+\}.$$

Recall that $\mathbb{R}^n_+$ means $\boldsymbol{x} \geq 0$. This set represents a *closed convex* cone, which is built on the columns $\boldsymbol{a}^1, \ldots \boldsymbol{a}^n \in \mathbb{R}^m$ of A. The reason it is closed and convex is simply because $\mathbb{R}^n_+$ is a closed and convex set, and $\mathcal{C}_A$ is obtain from it via a linear transformation.

Now let A be a $m \times n$ matrix and $\boldsymbol{b} \in \mathbb{R}^m$.

Theorem (Farkas alternative): One and only one of the following two cases is always true: $A\boldsymbol{x} = \boldsymbol{b}$ has a solution $\boldsymbol{x} \in \mathbb{R}^n$, $\boldsymbol{x} \geq 0$, *xor* there exists $\boldsymbol{y} \in \mathbb{R}^m$, such that $A^T\boldsymbol{y} \geq 0$ and $\boldsymbol{y} \cdot \boldsymbol{b} < 0$.

Proof: Either $\boldsymbol{b} \in \mathcal{C}_A$ or not. If yes, then $\boldsymbol{b} = A\boldsymbol{x}$ for some $\boldsymbol{x} \geq 0$, by definition of $\mathcal{C}_A$.

If not, then we can apply the Separating Hyperplane Theorem. The two sets $\mathcal{C}_A$ and $\{\boldsymbol{b}\}$ are closed and convex and the latter set is bounded. Then there exists a hyperplane that strictly separates these two sets. I.e. for some $\boldsymbol{y} \in \mathbb{R}^m$, $\beta \in \mathbb{R}$, the equation of the hyperplane itself is $\boldsymbol{y} \cdot \boldsymbol{z} + \beta = 0$, and for all $\boldsymbol{z} \in \mathcal{C}_A$, one has $\boldsymbol{y} \cdot \boldsymbol{z} + \beta > 0$, while $\boldsymbol{y} \cdot \boldsymbol{b} + \beta < 0$. Note that $\boldsymbol{y}$ is the normal vector to the hyperplane, and $\boldsymbol{z} \in \mathbb{R}^m$ is a variable.

To get more info about $\boldsymbol{y}$ and β, let us try now some special points $\boldsymbol{z}$ from $\mathcal{C}_A$. First off, $\mathbf{0} \in \mathcal{C}_A$, so try $\boldsymbol{z} = \mathbf{0}$. This implies, $\boldsymbol{y} \cdot \mathbf{0} + \beta > 0$, hence $\beta > 0$. Therefore, $\boldsymbol{y} \cdot \boldsymbol{b} < 0$. Now try $\boldsymbol{z} = M\mathbf{a}^j$, where $\mathbf{a}^j$ is the jth column of A, and $M > 0$ is a huge real. (For what $\boldsymbol{x} \geq 0$ do we have $M\mathbf{a}^j = A\boldsymbol{x}$?) It follows that for all j:

$$\boldsymbol{y} \cdot \mathbf{a}^j \geq -\frac{\beta}{M}, \qquad \text{for any } M > 0.$$

Passing to the limit $M \to \infty$, we have that for all j, $\boldsymbol{y} \cdot \mathbf{a}^j \geq 0$, which is the matrix notation is written exactly as $A^T\boldsymbol{y} \geq 0$. $\square$

What if one replaces $A\boldsymbol{x} = \boldsymbol{b}$ in the obverse of the Farkas alternative by $A\boldsymbol{x} \leq \boldsymbol{b}$? The answer is easy: add the slack variables. This augments the matrix A to $[A \; I]$, where I is the $m \times m$ identity matrix. On the reverse side $A^T\boldsymbol{y} \geq 0$ should now apply to the augmented matrix, so it becomes $A^T\boldsymbol{y} \geq 0$ and $I\boldsymbol{y} \geq 0$. In other words, here is one more formulation.

Farkas alternative, inequality formulation: One and only one of the following two cases is always true: $A\boldsymbol{x} \leq \boldsymbol{b}$ has a solution $\boldsymbol{x} \in \mathbb{R}^n_+$, *xor* there exists $\boldsymbol{y} \in \mathbb{R}^m_+$, such that $A^T\boldsymbol{y} \geq 0$ and $\boldsymbol{y} \cdot \boldsymbol{b} < 0$.

There is yet another interesting formulation that we'll meet later speaking about Lagrange multipliers. It bears a name of its own, the *Fredholm alternative.* It is obtained by changing the obverse of Farkas, removing the non-negativity claim on $\boldsymbol{x}$.

Corollary (Fredholm alternative): $A\boldsymbol{x} = \boldsymbol{b}$ has a solution $\boldsymbol{x} \in \mathbb{R}^n$, *xor* there exists $\boldsymbol{y} \in \mathbb{R}^m$, such that $A^T\boldsymbol{y} = 0$ and $\boldsymbol{y} \cdot \boldsymbol{b} \neq 0$.

Motivation. Suppose A is a square matrix, so $m = n$. If $\det A \neq 0$, then $A\boldsymbol{x} = \boldsymbol{b}$ has a solution for any $\boldsymbol{b}$. But otherwise the set $\{\boldsymbol{y} = A\boldsymbol{x},\ \boldsymbol{x} \in \mathbb{R}^n\}$ is a sub-space L of $\mathbb{R}^n$ of dimension less than n, called the Rank of A. Fredholm now tells us that either $\boldsymbol{b}$ lies in the subspace L, or we can find a vector $\boldsymbol{y}$ such that it is orthogonal to every element of L, but not orthogonal to $\boldsymbol{b}$.

Proof. In the obverse side of Farkas, let $\boldsymbol{x} = \boldsymbol{u} - \boldsymbol{v}$, where both $\boldsymbol{u}, \boldsymbol{v} \geq 0$. This augments the matrix to $\tilde{A} = [A \ \ -A]$, and the unknown to $\tilde{\boldsymbol{x}} = (\boldsymbol{u}, \boldsymbol{v})$. So the obverse side of the alternative is $\tilde{A}\tilde{\boldsymbol{x}} = \boldsymbol{b}$ has a solution $\tilde{\boldsymbol{x}} \geq 0$. The reverse side then is: there exists $\boldsymbol{y} \in \mathbb{R}^m$, such that $\boldsymbol{y} \cdot \boldsymbol{b} < 0$, and $\tilde{A}^T\boldsymbol{y} \geq 0$. But the latter means $A^T\boldsymbol{y} \geq 0$ and $A^T\boldsymbol{y} \leq 0$, which means $A^T\boldsymbol{y} = 0$. For the latter $-\boldsymbol{y}$ is just as good as $\boldsymbol{y}$, and to embrace it $\boldsymbol{y} \cdot \boldsymbol{b} < 0$ gets "generalised" to $\boldsymbol{y} \cdot \boldsymbol{b} \neq 0$. □

Farkas alternative implies strong duality theorem. We have seen two different interpretations of duality: one as the manufacturing-diet pair, the other via the canonical form. Both are equivalent and can be transformed into one another by augmenting. The Canonical form was introduced specifically to enable the Simplex Method to run, and one cannot do without it there. But to develop general duality theory the two formulations are equally good. As a matter of fact the "old" MP/DP formulation in "nicer" to produce a general result, because there is more symmetry built into it.

Strong duality theorem: Consider the primal problem

$$\max \boldsymbol{c} \cdot \boldsymbol{x}, \quad A\boldsymbol{x} \leq \boldsymbol{b}, \quad \boldsymbol{x} \geq 0$$

and its dual

$$\min \boldsymbol{b} \cdot \boldsymbol{y}, \quad A^T\boldsymbol{y} \geq \boldsymbol{c}, \quad \boldsymbol{y} \geq 0,$$

where A is an $m \times n$ matrix. Then exactly one of the following cases occurs:

(i) Both the primal and the dual have optimal solutions $\boldsymbol{x}$ and $\boldsymbol{y}$ with equal values $\boldsymbol{c} \cdot \boldsymbol{x} = \boldsymbol{b} \cdot \boldsymbol{y}$.
(ii) The dual is unfeasible, and the primal is unbounded, i.e. there are feasible $\boldsymbol{x}$ with $\boldsymbol{c} \cdot \boldsymbol{x} \to +\infty$.
(iii) The primal is unfeasible, and the dual is unbounded, i.e. there are feasible $\boldsymbol{y}$ with $\boldsymbol{b} \cdot \boldsymbol{y} \to -\infty$.
(iv) Both primal and dual are unfeasible.

Proof. The proof will use two things: weak duality[1] and Farkas. It follows from a single application of the Farkas alternative to the following system of inequalities as its obverse:

$$\left\{ \begin{array}{rrl} A\boldsymbol{x} & & \leq \boldsymbol{b}, \\ & -A^T\boldsymbol{y} & \leq -\boldsymbol{c}, \\ -\boldsymbol{c}^T\boldsymbol{x} & +\boldsymbol{b}^T\boldsymbol{y} & \leq 0. \end{array} \right. \tag{1}$$

Indeed, the first two lines are constraints on feasibility of $\boldsymbol{x}, \boldsymbol{y}$. The third inequality, by weak duality, can only be satisfied when $\boldsymbol{x}, \boldsymbol{y}$ are optimal (in which case it is satisfied tightly).

So (i) is the obverse. Suppose, it does not take place. Then it's the reverse. Let $\tilde{A}$ be the whole augmented matrix

$$\tilde{A} = \begin{bmatrix} A & 0 \\ 0 & -A^T \\ -\boldsymbol{c}^T & \boldsymbol{b}^T \end{bmatrix}$$

in (1). Note that the transposes of $\boldsymbol{c}, \boldsymbol{b}$ appear in the last row imply to testify that this is just a single iequality. The right-hand side is $\tilde{\boldsymbol{b}} = (\boldsymbol{b}, -\boldsymbol{c}, 0) \in \mathbb{R}^{m+n+1}$.

The reverse side: there exists $\tilde{\boldsymbol{y}} \in \mathbb{R}_+^{m+n+1}$, such that $\tilde{A}^T\tilde{\boldsymbol{y}} \geq 0$, and $\tilde{\boldsymbol{y}} \cdot \tilde{\boldsymbol{b}} < 0$. Let $\tilde{\boldsymbol{y}} = (\boldsymbol{p}, \boldsymbol{q}, r)$, with $\boldsymbol{p} \in \mathbb{R}_+^m,\ \boldsymbol{q} \in \mathbb{R}_+^n,\ r \in \mathbb{R}_+$.

Then $\tilde{\boldsymbol{y}} \cdot \tilde{\boldsymbol{b}} < 0$ reads

$$\boldsymbol{c} \cdot \boldsymbol{q} > \boldsymbol{b} \cdot \boldsymbol{p}. \tag{2}$$

[1]Recall: if $\boldsymbol{x}, \boldsymbol{y}$ are feasible for, respectively, the prima and dual, then $\boldsymbol{c} \cdot \boldsymbol{x} \leq \boldsymbol{b} \cdot \boldsymbol{y}$. This is proven just by, say, taking the dot product of the primal inequalities with $\boldsymbol{y}$, and this procedure will be repeated explicitly in the forthcoming proof a few times, without being referred to as weak duality.

Besides, $\tilde{A}^T\tilde{\boldsymbol{y}} \geq 0$ reads now $A\boldsymbol{q} \leq r\boldsymbol{b}$ and $A^T\boldsymbol{p} \leq r\boldsymbol{c}$. Let us use weak duality. Namely, take the dot product of the first one of the latter two inequalities with $\boldsymbol{p} \geq 0$:

$$r\boldsymbol{b}\cdot\boldsymbol{p} \geq \boldsymbol{p}\cdot A\boldsymbol{q} = A^T\boldsymbol{p}\cdot\boldsymbol{q} \geq r\boldsymbol{c}\cdot\boldsymbol{q}.$$

This, to come to terms with (2), implies $r = 0$.

Hence, we have (2) together with

$$\begin{cases} A\boldsymbol{q} \leq 0, \\ A^T\boldsymbol{p} \geq 0. \end{cases} \tag{3}$$

Let us see what possibilities for the primal/dual pair

$$\begin{cases} A\boldsymbol{x} \leq \boldsymbol{b}, \\ A^T\boldsymbol{y} \geq \boldsymbol{c}. \end{cases} \tag{4}$$

this leaves.

First off, both the primal and the dual cannot be feasible. For if there were feasible $\boldsymbol{x}, \boldsymbol{y}$, then we could take any huge real $M > 0$, add to (4) the M-multiple of (3) and get

$$\begin{cases} A(\boldsymbol{x} + M\boldsymbol{q}) \leq \boldsymbol{b}, \\ A^T(\boldsymbol{y} + M\boldsymbol{p}) \geq \boldsymbol{c}. \end{cases}$$

In other words, $\boldsymbol{x}_M = \boldsymbol{x} + M\boldsymbol{q}$ and $\boldsymbol{y}_M = \boldsymbol{y} + M\boldsymbol{p}$ would also be feasible solutions. By weak duality then $\boldsymbol{b}\cdot\boldsymbol{y}_M \geq \boldsymbol{c}\cdot\boldsymbol{x}_M$. This means

$$\boldsymbol{b}\cdot\boldsymbol{p} - \boldsymbol{c}\cdot\boldsymbol{q} \geq \frac{1}{M}(\boldsymbol{c}\cdot\boldsymbol{x} - \boldsymbol{b}\cdot\boldsymbol{y}).$$

Taking the limit $M \to \infty$ implies $\boldsymbol{b}\cdot\boldsymbol{p} - \boldsymbol{c}\cdot\boldsymbol{q} \geq 0$, in contradiction with (2).

So, if (i) in the Theorem's formulation does not occur, both primal and dual cannot be feasible. One possibility, of course, is (iv) – both unfeasible.

Let us now show that if the primal is feasible (in which case. as we already know, the dual is unfeasible) the primal is in fact, unbounded. Suppose $A\boldsymbol{x} \leq \boldsymbol{b}$ for some $\boldsymbol{x}$. Take a dot product with $\boldsymbol{p} \geq 0$:

$$\boldsymbol{p}\cdot\boldsymbol{b} \geq \boldsymbol{p}\cdot A\boldsymbol{x} = \boldsymbol{x}\cdot A^T\boldsymbol{p} \geq \boldsymbol{x}\cdot 0 = 0.$$

Then, by (2) one has $\boldsymbol{c}\cdot\boldsymbol{q} > 0$. So, $\boldsymbol{x}_M = \boldsymbol{x} + M\boldsymbol{q}$ is feasible for the primal, with $\boldsymbol{c}\cdot\boldsymbol{x}_M = \boldsymbol{c}\cdot\boldsymbol{x} + M\boldsymbol{c}\cdot\boldsymbol{q} \to +\infty$. Unbounded.

Finally, suppose $A^T\boldsymbol{y} \geq \boldsymbol{c}$ for some $\boldsymbol{y}$. Then, as we already know, the primal is unfeasible. Let us show that the dual is, in fact, unbounded. Take a dot product with $\boldsymbol{q} \geq 0$:

$$\boldsymbol{q}\cdot\boldsymbol{c} \leq \boldsymbol{q}\cdot A^T\boldsymbol{y} = \boldsymbol{y}\cdot A\boldsymbol{q} \leq \boldsymbol{y}\cdot 0 = 0.$$

Then, by (2) one has $\boldsymbol{b}\cdot\boldsymbol{p} < 0$. Thus $\boldsymbol{y}_M = \boldsymbol{y} + M\boldsymbol{p}$ is feasible for the dual, with $\boldsymbol{b}\cdot\boldsymbol{y}_M = \boldsymbol{b}\cdot y + M\boldsymbol{b}\cdot\boldsymbol{p} \to -\infty$. Unbounded.

$\square$