6. Duality of Linear Programming

6.1 The Duality Theorem

Here we formulate arguably the most important theoretical result about linear programs.

Let us consider the linear program

$$\begin{array}{lrcrcl} \text{maximize} & 2x_1 & + & 3x_2 & & \\ \text{subject to} & 4x_1 & + & 8x_2 & \leq & 12 \\ & 2x_1 & + & x_2 & \leq & 3 \\ & 3x_1 & + & 2x_2 & \leq & 4 \\ & \multicolumn{5}{l}{x_1, x_2 \geq 0.} \end{array} \tag{6.1}$$

Without computing the optimum, we can immediately infer from the first inequality and from the nonnegativity constraints that the maximum of the objective function is not larger than 12, because for nonnegative x_1 and x_2 we have

$$2x_1 + 3x_2 \leq 4x_1 + 8x_2 \leq 12.$$

We obtain a better upper bound if we first divide the first inequality by two:

$$2x_1 + 3x_2 \leq 2x_1 + 4x_2 \leq 6.$$

An even better bound results if we add the first two inequalities together and divide by three, which leads to the inequality

$$2x_1 + 3x_2 = \frac{1}{3}(4x_1 + 8x_2 + 2x_1 + x_2) \leq \frac{1}{3}(12 + 3) = 5,$$

and hence the objective function cannot be larger than 5.

How good an upper bound can we get in this way? And what does "in this way" mean? Let us begin with the latter question: From the constraints, we are trying to derive an inequality of the form

$$d_1x_1 + d_2x_2 \leq h,$$

where $d_1 \geq 2$, $d_2 \geq 3$, and h is as small as possible. Then we can claim that for all $x_1, x_2 \geq 0$ we have

$$2x_1 + 3x_2 \le d_1 x_1 + d_2 x_2 \le h,$$

and therefore, h is an upper bound on the maximum of the objective function. How can we derive such inequalities? We combine the three inequalities in the linear program with some nonnegative coefficients y_1, y_2, y_3 (nonnegativity is needed so that the direction of inequality is not reversed). We obtain

$$(4y_1 + 2y_2 + 3y_3)x_1 + (8y_1 + y_2 + 2y_3)x_2 \le 12y_1 + 3y_2 + 4y_3,$$

and thus $d_1 = 4y_1 + 2y_2 + 3y_3$, $d_2 = 8y_1 + y_2 + 2y_3$, and $h = 12y_1 + 3y_2 + 4y_3$.

How do we choose the best coefficients y_1, y_2, y_3? We must ensure that $d_1 \ge 2$ and $d_2 \ge 3$, and we want h to be as small as possible under these constraints. This is again a linear program:

$$\begin{array}{ll} \text{Minimize} & 12y_1 + 3y_2 + 4y_3 \\ \text{subject to} & 4y_1 + 2y_2 + 3y_3 \ge 2 \\ & 8y_1 + y_2 + 2y_3 \ge 3 \\ & y_1, y_2, y_3 \ge 0. \end{array}$$

It is called the linear program *dual* to the linear program (6.1) we started with. The dual linear program "guards" the original linear program from above, in the sense that every feasible solution (y_1, y_2, y_3) of the dual linear program provides an upper bound on the maximum of the objective function in (6.1).

How well does it guard? Perfectly! The optimal solution of the dual linear program is $\mathbf{y} = (\frac{5}{16}, 0, \frac{1}{4})$ with objective function equal to 4.75, and this is also the optimal value of the linear program (6.1), which is attained for $\mathbf{x} = (\frac{1}{2}, \frac{5}{4})$.

The duality theorem asserts that the dual linear program *always* guards perfectly. Let us repeat the above considerations in a more general setting, for a linear program of the form

$$\text{maximize } \mathbf{c}^T\mathbf{x} \text{ subject to } A\mathbf{x} \le \mathbf{b} \text{ and } \mathbf{x} \ge \mathbf{0}, \qquad \text{(P)}$$

where A is a matrix with m rows and n columns. We are trying to combine the m inequalities of the system $A\mathbf{x} \le \mathbf{b}$ with some nonnegative coefficients $y_1, y_2, \ldots, y_m$ so that

- the resulting inequality has the jth coefficient at least c_j, $j = 1, 2, \ldots, n$, and
- the right-hand side is as small as possible.

This leads to the **dual linear program**

$$\text{minimize } \mathbf{b}^T\mathbf{y} \text{ subject to } A^T\mathbf{y} \ge \mathbf{c} \text{ and } \mathbf{y} \ge \mathbf{0}; \qquad \text{(D)}$$

whoever doesn't believe this may write it in components. In this context the linear program (P) is referred to as the **primal linear program**.

From the way we have produced the dual linear program (D), we obtain the following result:

6.1.1 Proposition. *For each feasible solution* $\mathbf{y}$ *of the dual linear program* (D) *the value* $\mathbf{b}^T\mathbf{y}$ *provides an upper bound on the maximum of the objective function of the linear program* (P)*. In other words, for each feasible solution* $\mathbf{x}$ *of* (P) *and each feasible solution* $\mathbf{y}$ *of* (D) *we have*

$$\mathbf{c}^T\mathbf{x} \le \mathbf{b}^T\mathbf{y}.$$

In particular, if (P) *is unbounded,* (D) *has to be infeasible, and if* (D) *is unbounded (from below!), then* (P) *is infeasible.*

This proposition is usually called the *weak duality theorem*, weak because it expresses only the guarding of the primal linear program (P) by the dual linear program (D), but it doesn't say that the guarding is perfect. The latter is expressed only by the duality theorem (sometimes also called the *strong duality theorem*).

Duality theorem of linear programming

For the linear programs

$$\text{maximize } \mathbf{c}^T\mathbf{x} \text{ subject to } A\mathbf{x} \le \mathbf{b} \text{ and } \mathbf{x} \ge \mathbf{0} \qquad \text{(P)}$$

and

$$\text{minimize } \mathbf{b}^T\mathbf{y} \text{ subject to } A^T\mathbf{y} \ge \mathbf{c} \text{ and } \mathbf{y} \ge \mathbf{0} \qquad \text{(D)}$$

exactly one of the following possibilities occurs:

1. Neither (P) nor (D) has a feasible solution.
2. (P) is unbounded and (D) has no feasible solution.
3. (P) has no feasible solution and (D) is unbounded.
4. Both (P) and (D) have a feasible solution. Then both have an optimal solution, and if $\mathbf{x}^*$ is an optimal solution of (P) and $\mathbf{y}^*$ is an optimal solution of (D), then
$$\mathbf{c}^T\mathbf{x}^* = \mathbf{b}^T\mathbf{y}^*.$$
That is, the *maximum of* (P) *equals the minimum of* (D).

The duality theorem might look complicated at first encounter. For understanding it better it may be useful to consider a simpler version, called the Farkas lemma and discussed in Section 6.4. This simpler statement has several intuitive interpretations, and it contains the essence of the duality theorem.

Proving the duality theorem, which we will undertake in Sections 6.3 and 6.4, does take some work, unlike a proof of the weak duality theorem, which is quite easy.

The heart of the duality theorem is the equality $\mathbf{c}^T\mathbf{x}^* = \mathbf{b}^T\mathbf{y}^*$ in the fourth possibility, i.e., for both (P) and (D) feasible.

Since a linear program can be either feasible and bounded, or feasible and unbounded, or infeasible, there are 3 possibilities for (P) and 3 possibilities for (D), which at first sight gives 9 possible combinations for (P) and (D). The three cases "(P) unbounded and (D) feasible bounded," "(P) unbounded and (D) unbounded," and "(P) feasible bounded and (D) unbounded" are ruled out by the weak duality theorem. In the proof of the duality theorem, we will rule out the cases "(P) infeasible and (D) feasible bounded," as well as "(P) feasible bounded and (D) infeasible." This leaves us with the four cases listed in the duality theorem. All of them can indeed occur.

Once again: feasibility versus optimality. In Chapter 1 we remarked that finding a feasible solution of a linear program is in general computationally as difficult as finding an optimal solution. There we briefly substantiated this claim using binary search. The duality theorem provides a considerably more elegant argument: The linear program (P) has an optimal solution if and only if the following linear program, obtained by combining the constraints of (P), the constraints of (D), and an inequality between the objective functions, has a feasible solution:

$$\begin{array}{ll} \text{Maximize} & \mathbf{c}^T\mathbf{x} \\ \text{subject to} & A\mathbf{x} \leq \mathbf{b}, \\ & A^T\mathbf{y} \geq \mathbf{c}, \\ & \mathbf{c}^T\mathbf{x} \geq \mathbf{b}^T\mathbf{y}, \\ & \mathbf{x} \geq \mathbf{0}, \mathbf{y} \geq \mathbf{0}. \end{array}$$

(the objective function is immaterial here, and the variables are $x_1, \ldots, x_n$ and $y_1, \ldots, y_m$). Moreover, for each feasible solution $(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})$ of the last linear program, $\tilde{\mathbf{x}}$ is an optimal solution of the linear program (P). All of this is a simple consequence of the duality theorem.

6.2 Dualization for Everyone

The duality theorem is valid for each linear program, not only for one of the form (P); we have only to construct the dual linear program properly. To this end, we can convert the given linear program to the form (P) using the tricks from Sections 1.1 and 4.1, and then the dual linear program has the form (D). The result can often be simplified; for example, the difference of two nonnegative variables can be replaced by a single unbounded variable (one that may attain all real values).

Simpler than doing this again and again is to adhere to the recipe below (whose validity can be proved by the just mentioned procedure). Let us assume that the primal linear program has variables $x_1, x_2, \ldots, x_n$, among which some may be nonnegative, some nonpositive, and some unbounded. Let the constraints be $C_1, C_2, \ldots, C_m$, where C_i has the form

$$a_{i1}x_1 + a_{i2}x_2 + \cdots + a_{in}x_n \left\{ \begin{array}{c} \leq \\ \geq \\ = \end{array} \right\} b_i.$$

(The nonnegativity or nonpositivity constraints for the variables are not counted among the C_i.) The objective function $\mathbf{c}^T\mathbf{x}$ should be *maximized.*

Then the dual linear program has variables $y_1, y_2, \ldots, y_m$, where y_i corresponds to the constraint C_i and satisfies

$$\left\{ \begin{array}{c} y_i \geq 0 \\ y_i \leq 0 \\ y_i \in \mathbb{R} \end{array} \right\} \text{ if we have } \left\{ \begin{array}{c} \leq \\ \geq \\ = \end{array} \right\} \text{ in } C_i.$$

The constraints of the dual linear program are $Q_1, Q_2, \ldots, Q_n$, where Q_j corresponds to the variable x_j and reads

$$a_{1j}y_1 + a_{2j}y_2 + \cdots + a_{mj}y_m \left\{ \begin{array}{c} \geq \\ \leq \\ = \end{array} \right\} c_j \text{ if } x_j \text{ satisfies } \left\{ \begin{array}{c} x_j \geq 0 \\ x_j \leq 0 \\ x_j \in \mathbb{R} \end{array} \right\}.$$

The objective function is $\mathbf{b}^T\mathbf{y}$, and it is to be *minimized.*

Note that in the first part of the recipe (from primal constraints to dual variables) the direction of inequalities is reversed, while in the second part (from primal variables to dual constraints) the direction is preserved.

Dualization Recipe

	Primal linear program	Dual linear program
Variables	$x_1, x_2, \ldots, x_n$	$y_1, y_2, \ldots, y_m$
Matrix	A	A^T
Right-hand side	$\mathbf{b}$	$\mathbf{c}$
Objective function	$\max \mathbf{c}^T\mathbf{x}$	$\min \mathbf{b}^T\mathbf{y}$
Constraints	ith constraint has $\leq$ $\geq$ $=$	$y_i \geq 0$ $y_i \leq 0$ $y_i \in \mathbb{R}$
	$x_j \geq 0$ $x_j \leq 0$ $x_j \in \mathbb{R}$	jth constraint has $\geq$ $\leq$ $=$

If we want to dualize a *minimization* linear program, we can first transform it to a maximization linear program by changing the sign of the objective function, and then follow the recipe.

In this way one can also find out that the rules work symmetrically "there" and "back." By this we mean that if we start with some linear program, construct the dual linear program, and then again the dual linear program, we get back to the original (primal) linear program; two consecutive dualizations cancel out. In particular, the linear programs (P) and (D) in the duality theorem are *dual to each other.*

A physical interpretation of duality. Let us consider a linear program

$$\text{maximize } \mathbf{c}^T\mathbf{x} \text{ subject to } A\mathbf{x} \leq \mathbf{b}.$$

According to the dualization recipe the dual linear program is

$$\text{minimize } \mathbf{b}^T\mathbf{y} \text{ subject to } A^T\mathbf{y} = \mathbf{c} \text{ and } \mathbf{y} \geq \mathbf{0}.$$

Let us assume that the primal linear program is feasible and bounded, and let $n = 3$. We regard $\mathbf{x}$ as a point in three-dimensional space, and we interpret $\mathbf{c}$ as the gravitation vector; it thus points downward.

Each of the inequalities of the system $A\mathbf{x} \leq \mathbf{b}$ determines a half-space. The intersection of these half-spaces is a nonempty convex polyhedron bounded from below. Each of its two-dimensional faces is given by one of the equations $\mathbf{a}_i^T\mathbf{x} = b_i$, where the vectors $\mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_m$ are the rows of the matrix A, but interpreted as column vectors. Let us denote the face given by $\mathbf{a}_i^T\mathbf{x} = b_i$ by S_i (not every inequality of the system $A\mathbf{x} \leq \mathbf{b}$ has to correspond to a face, and so S_i is not necessarily defined for every i).

Let us imagine that the boundary of the polyhedron is made of cardboard and that we drop a tiny steel ball somewhere inside the polyhedron. The ball falls and rolls down to the lowest vertex (or possibly it stays on a horizontal edge or face). Let us denote the resulting position of the ball by $\mathbf{x}^*$; thus, $\mathbf{x}^*$ is an optimal solution of the linear program. In this stable position the ball touches several two-dimensional faces, typically 3. Let D be the set of i such that the ball touches the face S_i. For $i \in D$ we thus have

$$\mathbf{a}_i^T\mathbf{x}^* = b_i. \tag{6.2}$$

Gravity exerts a force $\mathbf{F}$ on the ball that is proportional to the vector $\mathbf{c}$. This force is decomposed into forces of pressure on the faces touched by the ball. The force $\mathbf{F}_i$ by which the ball acts on face S_i is orthogonal to S_i and it is directed outward from the polyhedron (if we neglect friction); see the schematic two-dimensional picture below:

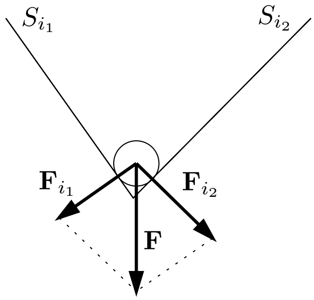

The forces acting on the ball are in equilibrium, and thus $\mathbf{F} = \sum_{i\in D} \mathbf{F}_i$. The outward normal of the face S_i is $\mathbf{a}_i$; hence $\mathbf{F}_i$ is proportional to $\mathbf{a}_i$, and for some nonnegative numbers y_i^* we have

$$\sum_{i\in D} y_i^* \mathbf{a}_i = \mathbf{c}.$$

If we set $y_i^* = 0$ for $i \notin D$, we can write $\sum_{i=1}^{m} y_i^* \mathbf{a}_i = \mathbf{c}$, or $A^T\mathbf{y}^* = \mathbf{c}$ in matrix form. Therefore, $\mathbf{y}^*$ is a feasible solution of the dual linear program.

Let us consider the product $(\mathbf{y}^*)^T(A\mathbf{x}^* - \mathbf{b})$. For $i \notin D$ the ith component of $\mathbf{y}^*$ equals 0, while for $i \in D$ the ith component of $A\mathbf{x}^* - \mathbf{b}$ is 0 according to (6.2). So the product is 0, and hence $(\mathbf{y}^*)^T\mathbf{b} = (\mathbf{y}^*)^T A\mathbf{x}^* = \mathbf{c}^T\mathbf{x}^*$.

We see that $\mathbf{x}^*$ is a feasible solution of the primal linear program, $\mathbf{y}^*$ is a feasible solution of the dual linear program, and $\mathbf{c}^T\mathbf{x}^* = \mathbf{b}^T\mathbf{y}^*$. By the weak duality theorem $\mathbf{y}^*$ is an optimal solution of the dual linear program, and we have a situation exactly as in the duality theorem. We have just "physically verified" a special three-dimensional case of the duality theorem.

We remark that the dual linear program also has an economic interpretation. The dual variables are called *shadow prices* in this context. The interested reader will find this nicely explained in Chvátal's textbook cited in Chapter 9.

6.3 Proof of Duality from the Simplex Method

The duality theorem of linear programming can be quickly derived from the correctness of the simplex method. To be precise, we will prove the following:

> *If the primal linear program* (P) *is feasible and bounded, then the dual linear program* (D) *is feasible (and bounded as well, by weak duality), with the same optimum value as the primal.*

Since the dual of the dual is the primal, we may interchange (P) and (D) in this statement. Together with our considerations about the possible cases after the statement of the duality theorem, this will prove the theorem.

The key observation is that we can extract an optimal solution of the dual linear program from the final tableau. We should recall, though, that proving the correctness of the simplex method, and in particular, the fact that one can always avoid cycling, requires considerable work.

Let us consider a primal linear program

$$\text{maximize } \mathbf{c}^T\mathbf{x} \text{ subject to } A\mathbf{x} \leq \mathbf{b} \text{ and } \mathbf{x} \geq \mathbf{0}. \tag{P}$$

After a conversion to equational via slack variables $x_{n+1}, \ldots, x_{n+m}$ we arrive at the linear program

$$\text{maximize } \overline{\mathbf{c}}^T\overline{\mathbf{x}} \text{ subject to } \bar{A}\overline{\mathbf{x}} = \mathbf{b} \text{ and } \overline{\mathbf{x}} \geq \mathbf{0},$$

where $\overline{\mathbf{x}} = (x_1, \ldots, x_{n+m})$, $\overline{\mathbf{c}} = (c_1, \ldots, c_n, 0, \ldots, 0)$, and $\bar{A} = (A \,|\, I_m)$. If this last linear program is feasible and bounded, then according to Theorem 5.8.1, the simplex method with Bland's rule always finds some optimal solution $\overline{\mathbf{x}}^*$ with a feasible basis B. The first n components of the vector $\overline{\mathbf{x}}^*$ constitute an optimal solution $\mathbf{x}^*$ of the linear program (P). By the optimality criterion we have $\mathbf{r} \leq \mathbf{0}$ in the final simplex tableau, where $\mathbf{r}$ is the vector in the z-row of the tableau as in Section 5.5. The following lemma and the weak duality theorem (Proposition 6.1.1) then easily imply the duality theorem.

6.3.1 Lemma. *In the described situation the vector* $\mathbf{y}^* = (\overline{\mathbf{c}}_B^T \bar{A}_B^{-1})^T$ *is a feasible solution of the dual linear program* (D) *and the equality* $\mathbf{c}^T\mathbf{x}^* = \mathbf{b}^T\mathbf{y}^*$ *holds.*

Proof. By Lemma 5.5.1, $\overline{\mathbf{x}}^*$ is given by $\overline{\mathbf{x}}^*_B = \bar{A}_B^{-1}\mathbf{b}$ and $\overline{\mathbf{x}}^*_N = \mathbf{0}$, and so

$$\mathbf{c}^T\mathbf{x}^* = \overline{\mathbf{c}}^T\overline{\mathbf{x}}^* = \overline{\mathbf{c}}_B^T\overline{\mathbf{x}}_B^* = \overline{\mathbf{c}}_B^T(\bar{A}_B^{-1}\mathbf{b}) = (\overline{\mathbf{c}}_B^T\bar{A}_B^{-1})\mathbf{b} = (\mathbf{y}^*)^T\mathbf{b} = \mathbf{b}^T\mathbf{y}^*.$$

The equality $\mathbf{c}^T\mathbf{x}^* = \mathbf{b}^T\mathbf{y}^*$ thus holds, and it remains to check the feasibility of $\mathbf{y}^*$, that is, $A^T\mathbf{y}^* \geq \mathbf{c}$ and $\mathbf{y}^* \geq \mathbf{0}$.

The condition $\mathbf{y}^* \geq \mathbf{0}$ can be rewritten to $I_m\mathbf{y}^* \geq \mathbf{0}$, and hence both of the feasibility conditions together are equivalent to

$$\bar{A}^T\mathbf{y}^* \geq \overline{\mathbf{c}}. \tag{6.3}$$

After substituting $\mathbf{y}^* = (\overline{\mathbf{c}}_B^T\bar{A}_B^{-1})^T$ the left-hand side becomes $\bar{A}^T(\overline{\mathbf{c}}_B^T\bar{A}_B^{-1})^T = (\overline{\mathbf{c}}_B^T\bar{A}_B^{-1}\bar{A})^T$. Let us denote this $(n{+}m)$-component vector by $\mathbf{w}$. For the basic components of $\mathbf{w}$ we have

$$\mathbf{w}_B = (\overline{\mathbf{c}}_B^T\bar{A}_B^{-1}\bar{A}_B)^T = (\overline{\mathbf{c}}_B^T I_m)^T = \overline{\mathbf{c}}_B,$$

and thus we even have equality in (6.3) for the basic components. For the nonbasic components we have

$$\mathbf{w}_N = (\overline{\mathbf{c}}_B^T \bar{A}_B^{-1} \bar{A}_N)^T = \overline{\mathbf{c}}_N - \mathbf{r} \geq \overline{\mathbf{c}}_N$$

since $\mathbf{r} = \overline{\mathbf{c}}^N - (\overline{\mathbf{c}}_B^T \bar{A}_B^{-1} \bar{A}_N)^T$ by Lemma 5.5.1, and $\mathbf{r} \leq \mathbf{0}$ by the optimality criterion. The lemma is proved. □

6.4 Proof of Duality from the Farkas Lemma

Another approach to the duality theorem of linear programming consists in first proving a simplified version, called the *Farkas lemma*, and then substituting a skillfully composed matrix into it and thus deriving the theorem. A nice feature is that the Farkas lemma has very intuitive interpretations.

Actually, the Farkas lemma comes in several natural variants. We begin by discussing one of them, which has a very clear geometric meaning.

6.4.1 Proposition (Farkas lemma). *Let A be a real matrix with m rows and n columns, and let $\mathbf{b} \in \mathbb{R}^m$ be a vector. Then exactly one of the following two possibilities occurs:*

(F1) *There exists a vector $\mathbf{x} \in \mathbb{R}^n$ satisfying $A\mathbf{x} = \mathbf{b}$ and $\mathbf{x} \geq \mathbf{0}$.*
(F2) *There exists a vector $\mathbf{y} \in \mathbb{R}^m$ such that $\mathbf{y}^T A \geq \mathbf{0}^T$ and $\mathbf{y}^T\mathbf{b} < 0$.*

It is easily seen that both possibilities cannot occur at the same time. Indeed, the vector $\mathbf{y}$ in (F2) determines a linear combination of the equations witnessing that $A\mathbf{x} = \mathbf{b}$ cannot have any nonnegative solution: All coefficients on the left-hand side of the resulting equation $(\mathbf{y}^T A)\mathbf{x} = \mathbf{y}^T\mathbf{b}$ are nonnegative, but the right-hand side is negative.

The Farkas lemma is not exactly a difficult theorem, but it is not trivial either. Many proofs are known, and we will present some of them in the subsequent sections. The reader is invited to choose the "best" one according to personal taste.

A geometric view. In order to view the Farkas lemma geometrically, we need the notion of convex hull; see Section 4.3. Further we define, for vectors $\mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_n \in \mathbb{R}^m$, the **convex cone** generated by $\mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_n$ as the set of all linear combinations of the $\mathbf{a}_i$ with nonnegative coefficients, that is, as

$$\left\{ t_1\mathbf{a}_1 + t_2\mathbf{a}_2 + \cdots + t_n\mathbf{a}_n : t_1, t_2, \ldots, t_n \geq 0 \right\}.$$

In other words, this convex cone is the convex hull of the rays $p_1, p_2, \ldots, p_n$, where $p_i = \{t\mathbf{a}_i : t \geq 0\}$ emanates from the origin and passes through the point $\mathbf{a}_i$.

6.4.2 Proposition (Farkas lemma geometrically). *Let $\mathbf{a}_1$, $\mathbf{a}_2$,..., $\mathbf{a}_n, \mathbf{b}$ be vectors in $\mathbb{R}^m$. Then exactly one of the following two possibilities occurs:*

(F1′) *The point $\mathbf{b}$ lies in the convex cone C generated by $\mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_n$.*

(F2′) *There exists a hyperplane h passing through the point* $\mathbf{0}$, *of the form*

$$h = \{\mathbf{x} \in \mathbb{R}^m : \mathbf{y}^T\mathbf{x} = 0\}$$

for a suitable $\mathbf{y} \in \mathbb{R}^m$, *such that all the vectors* $\mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_n$ *(and thus the whole cone* C*) lie on one side and* $\mathbf{b}$ *lies (strictly) on the other side. That is,* $\mathbf{y}^T\mathbf{a}_i \geq 0$ *for all* $i = 1, 2, \ldots, n$ *and* $\mathbf{y}^T\mathbf{b} < 0$.

A drawing illustrates both possibilities for $m = 2$ and $n = 3$:

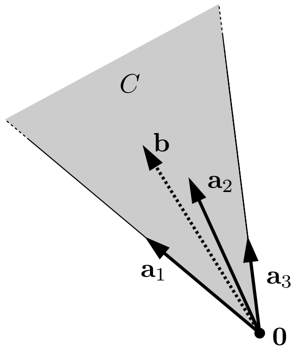

(F1′)

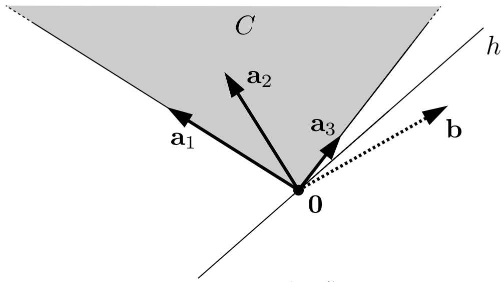

(F2′)

To see that Proposition 6.4.1 and Proposition 6.4.2 really tell us the same thing, it suffices to take the columns of the matrix A for $\mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_n$. The existence of a nonnegative solution of $A\mathbf{x} = \mathbf{b}$ can be reexpressed as $\mathbf{b} = t_1\mathbf{a}_1 + t_2\mathbf{a}_2 + \cdots + t_n\mathbf{a}_n$, $t_1, t_2, \ldots, t_n \geq 0$, and this says exactly that $\mathbf{b} \in C$. The equivalence of (F2) and (F2′) hopefully doesn't need any further explanation.

This result is an instance of a *separation theorem* for convex sets. Separation theorems generally assert that disjoint convex sets can be separated by a hyperplane. There are several versions (depending on whether one requires strict or nonstrict separation, etc.) and several proof strategies. Separation theorems in infinite-dimensional Banach spaces are closely related to the Hahn–Banach theorem, one of the cornerstones of functional analysis. In Section 6.5 we prove the Farkas lemma along these lines, viewing it as a geometric separation theorem.

Variants of the Farkas lemma. Proposition 6.4.1 provides an answer to the question, "When does a system of linear *equalities* have a *nonnegative* solution?" In part (i) of the following proposition, we restate Proposition 6.4.1 (in a slightly different, but clearly equivalent form), and in parts (ii) and (iii), we add two more variants of the Farkas lemma. Part (ii) answers the question, "When does a system of linear *inequalities* have a *nonnegative* solution?" and part (iii) the question, "When does a system of linear *inequalities* have *any* solution at all?"

6.4.3 Proposition (Farkas lemma in three variants). *Let* A *be a real matrix with* m *rows and* n *columns, and let* $\mathbf{b} \in \mathbb{R}^m$ *be a vector.*

(i) *The system* $A\mathbf{x} = \mathbf{b}$ *has a nonnegative solution if and only if every* $\mathbf{y} \in \mathbb{R}^m$ *with* $\mathbf{y}^T A \geq \mathbf{0}^T$ *also satisfies* $\mathbf{y}^T\mathbf{b} \geq \mathbf{0}$.

(ii) *The system* $A\mathbf{x} \leq \mathbf{b}$ *has a nonnegative solution if and only if every nonnegative* $\mathbf{y} \in \mathbb{R}^m$ *with* $\mathbf{y}^T A \geq \mathbf{0}^T$ *also satisfies* $\mathbf{y}^T\mathbf{b} \geq 0$.

(iii) *The system* $A\mathbf{x} \leq \mathbf{b}$ *has a solution if and only if every nonnegative* $\mathbf{y} \in \mathbb{R}^m$ *with* $\mathbf{y}^T A = \mathbf{0}^T$ *also satisfies* $\mathbf{y}^T\mathbf{b} \geq 0$.

The three parts of Proposition 6.4.3 are mutually equivalent, in the sense that any of them can easily be derived from any other. Having three forms at our disposal provides more flexibility, both for applying the Farkas lemma and for proving it.

The proof of the equivalence (i)⇔(ii)⇔(iii) is easy, using the tricks familiar from transformations of linear programs to equational form. We will take a utilitarian approach: Since we will use (ii) in the proof of the duality theorem, we prove only the implications (i)⇒(ii) and (iii)⇒(ii), leaving the remaining implications to the reader.

Proof of (i)⇒(ii). In (ii) we need an equivalent condition for $A\mathbf{x} \leq \mathbf{b}$ having a nonnegative solution. To this end, we form the matrix $\bar{A} = (A \,|\, I_m)$. We note that $A\mathbf{x} \leq \mathbf{b}$ has a nonnegative solution if and only if $\bar{A}\,\overline{\mathbf{x}} = \mathbf{b}$ has a nonnegative solution. By (i), this is equivalent to the condition that all $\mathbf{y}$ with $\mathbf{y}^T\bar{A} \geq \mathbf{0}^T$ satisfy $\mathbf{y}^T\mathbf{b} \geq 0$. And finally, $\mathbf{y}^T\bar{A} \geq \mathbf{0}^T$ says exactly the same as $\mathbf{y}^T A \geq \mathbf{0}^T$ and $\mathbf{y} \geq \mathbf{0}$, and hence we have the desired equivalence. □

Proof of (iii)⇒(ii). Again we need an equivalent condition for $A\mathbf{x} \leq \mathbf{b}$ having a nonnegative solution. This time we form the matrix $\bar{A}$ and the vector $\overline{\mathbf{b}}$ according to

$$\bar{A} = \left(\frac{A}{-I_n}\right), \quad \overline{\mathbf{b}} = \left(\frac{\mathbf{b}}{\mathbf{0}}\right).$$

Then $A\mathbf{x} \leq \mathbf{b}$ has a nonnegative solution if and only if $\bar{A}\mathbf{x} \leq \overline{\mathbf{b}}$ has any solution. The latter is equivalent, by (iii), to the condition that all $\overline{\mathbf{y}} \geq \mathbf{0}$ with $\overline{\mathbf{y}}^T\bar{A} = \mathbf{0}^T$ satisfy $\overline{\mathbf{y}}^T\overline{\mathbf{b}} \geq 0$. Writing

$$\overline{\mathbf{y}} = \left(\frac{\mathbf{y}}{\mathbf{y}'}\right),$$

$\mathbf{y}$ a vector with m components, we have

$$\overline{\mathbf{y}} \geq \mathbf{0}, \;\; \overline{\mathbf{y}}^T\bar{A} = \mathbf{0}^T \quad \text{exactly if} \quad \mathbf{y} \geq \mathbf{0}, \;\; \mathbf{y}'^T = \mathbf{y}^T A \geq \mathbf{0}^T$$

and

$$\overline{\mathbf{y}}^T\overline{\mathbf{b}} = \mathbf{y}^T\mathbf{b}.$$

From this and our chain of equivalences, we deduce that $A\mathbf{x} \leq \mathbf{b}$ has a nonnegative solution if and only if all $\mathbf{y} \geq \mathbf{0}$ with $\mathbf{y}^T A \geq \mathbf{0}^T$ satisfy $\mathbf{y}^T\mathbf{b} \geq 0$, and this is the statement of (ii). □

Remarks. A reader with a systematic mind may like to see the variants of the Farkas lemma summarized in a table:

	The system $A\mathbf{x} \leq \mathbf{b}$	The system $A\mathbf{x} = \mathbf{b}$
has a solution $\mathbf{x} \geq \mathbf{0}$ iff	$\mathbf{y} \geq \mathbf{0}, \mathbf{y}^T A \geq \mathbf{0}$ $\Rightarrow \mathbf{y}^T\mathbf{b} \geq 0$	$\mathbf{y}^T A \geq \mathbf{0}^T$ $\Rightarrow \mathbf{y}^T\mathbf{b} \geq 0$
has a solution $\mathbf{x} \in \mathbb{R}^n$ iff	$\mathbf{y} \geq \mathbf{0}, \mathbf{y}^T A = \mathbf{0}$ $\Rightarrow \mathbf{y}^T\mathbf{b} \geq 0$	$\mathbf{y}^T A = \mathbf{0}^T$ $\Rightarrow \mathbf{y}^T\mathbf{b} = 0$

We had three variants of the Farkas lemma, but the table has four entries. We haven't mentioned the statement corresponding to the bottom right corner of the table, telling us when a system of linear *equations* has *any* solution. We haven't mentioned it because it doesn't deserve to be called a Farkas lemma—the proof is a simple exercise in linear algebra, and there doesn't seem to be any way of deriving the Farkas lemma from this variant along the lines of our previous reductions. However, we will find this statement useful in Section 6.6, where it will serve as a basis of a proof of a "real" Farkas lemma.

Let us note that, similar to "dualization for everyone," we could also establish a unifying "Farkas lemma for everyone," dealing with a system containing both linear equations and inequalities and with some of the variables nonnegative and some unrestricted. This would contain all of the four variants considered above as special cases, but we will not go in this direction.

A logical view. Now we explain yet another way of understanding the Farkas lemma, this time variant (iii) in Proposition 6.4.3. We begin with something seemingly different, namely, deriving new linear inequalities from old ones. From two given inequalities, say

$$4x_1 + x_2 \leq 4 \quad \text{and} \quad -x_1 + x_2 \leq 1,$$

we can derive new inequalities by multiplying the first inequality by a *positive* real number α, the second one by a *positive* real number β, and adding the resulting inequalities together (we must be careful so that both inequality signs have the same direction!); we have already used this many times. For instance, for $\alpha = 3$ and $\beta = 2$ we derive the inequality $10x_1 + 5x_2 \leq 14$. More generally, if we start with a system of several linear inequalities, of the form $A\mathbf{x} \leq \mathbf{b}$, we can derive new inequalities by repeating this operation for various pairs, which may involve both the original inequalities and new ones derived earlier. So if we start with the system

$$4x_1 + x_2 \leq 4, \ -x_1 + x_2 \leq 1, \quad \text{and} \quad -2x_1 - x_2 \leq -3,$$

we can first derive $10x_1 + 5x_2 \leq 14$ from the first two as before, and then we can add to this new inequality the third inequality multiplied by 5. In this case both of the coefficients on the left-hand side cancel out, and we get the inequality $0 \leq -1$. This last inequality *obviously* never holds, and so the original triple of inequalities cannot be satisfied by any $(x_1, x_2) \in \mathbb{R}$ either (as is easy to check using a picture).

The Farkas lemma turns out to be equivalent to the following statement:

> Whenever a system $A\mathbf{x} \leq \mathbf{b}$ of finitely many linear inequalities is **inconsistent**, that is, there is no $\mathbf{x} \in \mathbb{R}^n$ satisfying it, we can derive the (obviously inconsistent) inequality $0 \leq -1$ from it by the above procedure.

A little thought reveals that each inequality derived by the procedure (repeated combinations of pairs) has the form $(\mathbf{y}^T A)\mathbf{x} \leq \mathbf{y}^T\mathbf{b}$ for some nonnegative vector $\mathbf{y} \in \mathbb{R}^m$, and thus, equivalently, we claim that whenever $A\mathbf{x} \leq \mathbf{b}$ is inconsistent, there exists a vector $\mathbf{y} \geq \mathbf{0}$ with $\mathbf{y}^T A = \mathbf{0}^T$ and $\mathbf{y}^T\mathbf{b} = -1$. This is clearly equivalent to part (iii) of Proposition 6.4.3.

The reader may wonder why we have bothered to consider repeated pairwise combinations of inequalities, instead of using a single vector $\mathbf{y}$ specifying a combination of all of the inequalities right away. The reason is that the "pairwise" formulation makes the statement more similar to a number of important and famous statements in various branches of mathematics. In logic, for example, theorems are derived (proved) from axioms by repeated application of certain simple derivation rules. In the first-order propositional calculus, there is a *completeness theorem*: Any true statement (that is, a statement valid in every model) can be derived from the axioms by a finite sequence of steps, using the appropriate derivation rules, such as *modus ponens*. In contrast, the celebrated Gödel's first incompleteness theorem asserts that in Peano arithmetic, as well as in any theory containing it, there are statements that are true but *cannot* be derived.

In analogy to this, we can view the inequalities of the original system $A\mathbf{x} \leq \mathbf{b}$ as "axioms," and we have a single derivation rule (derive a new inequality from two existing ones by a positive linear combination as above). Then the Farkas lemma tells us that any inconsistent system of "axioms" can be refuted by a suitable derivation. (This is a "weak" completeness theorem; we could also consider a more general "completeness theorem," stating that whenever a linear inequality is valid for all $\mathbf{x} \in \mathbb{R}^n$ satisfying $A\mathbf{x} \leq \mathbf{b}$, then it can be derived from $A\mathbf{x} \leq \mathbf{b}$, but we will not go into this here.) Such a completeness result means that the theory of linear inequalities is, in a sense, "easy." Moreover, the simplex method, or also the Fourier–Motzkin elimination considered in Section 6.7, provide ways to construct such a derivation.

This view makes the Farkas lemma a (small) cousin of various completeness theorems of logic and of other famous results, such as

Hilbert's Nullstellensatz in algebraic geometry. Computer science also frequently investigates the possibility of deriving some object from given initial objects by certain derivation rules, say in the theory of formal languages.

Proof of the duality theorem from the Farkas lemma. Let us assume that the linear program (P) has an optimal solution $\mathbf{x}^*$. As in the proof of the duality theorem from the simplex method, we show that the dual (D) has an optimal solution as well, and that the optimum values of both programs coincide.

We first define $\gamma = \mathbf{c}^T\mathbf{x}^*$ to be the optimum value of (P). Then we know that the system of inequalities

$$A\mathbf{x} \le \mathbf{b},\ \mathbf{c}^T\mathbf{x} \ge \gamma \tag{6.4}$$

has a nonnegative solution, but for any $\varepsilon > 0$, the system

$$A\mathbf{x} \le \mathbf{b},\ \mathbf{c}^T\mathbf{x} \ge \gamma + \varepsilon \tag{6.5}$$

has *no* nonnegative solution. If we define an $(m+1)\times n$ matrix $\hat{A}$ and a vector $\hat{\mathbf{b}}_\varepsilon \in \mathbb{R}^m$ by

$$\hat{A} = \begin{pmatrix} A \\ -\mathbf{c}^T \end{pmatrix}, \quad \hat{\mathbf{b}}_\varepsilon = \begin{pmatrix} \mathbf{b} \\ -\gamma - \varepsilon \end{pmatrix},$$

then (6.4) is equivalent to $\hat{A}\mathbf{x} \le \hat{\mathbf{b}}_0$ and (6.5) is equivalent to $\hat{A}\mathbf{x} \le \hat{\mathbf{b}}_\varepsilon$.

Let us apply the variant of the Farkas lemma in Proposition 6.4.3(ii). For $\varepsilon > 0$, the system $\hat{A}\mathbf{x} \le \hat{\mathbf{b}}_\varepsilon$ has no nonnegative solution, so we conclude that there is a nonnegative vector $\hat{\mathbf{y}} = (\mathbf{u}, z) \in \mathbb{R}^{m+1}$ such that $\hat{\mathbf{y}}^T\hat{A} \ge \mathbf{0}^T$ but $\hat{\mathbf{y}}^T\hat{\mathbf{b}}_\varepsilon < 0$. These conditions boil down to

$$A^T\mathbf{u} \ge z\mathbf{c},\ \mathbf{b}^T\mathbf{u} < z(\gamma + \varepsilon). \tag{6.6}$$

Applying the Farkas lemma in the case $\varepsilon = 0$ (the system has a nonnegative solution), we see that the very same vector $\hat{\mathbf{y}}$ must satisfy $\hat{\mathbf{y}}^T\hat{\mathbf{b}}_0 \ge 0$, and this is equivalent to

$$\mathbf{b}^T\mathbf{u} \ge z\gamma.$$

It follows that $z > 0$, since $z = 0$ would contradict the strict inequality in (6.6). But then we may set $\mathbf{v} := \frac{1}{z}\,\mathbf{u} \ge \mathbf{0}$, and (6.6) yields

$$A^T\mathbf{v} \ge \mathbf{c},\ \mathbf{b}^T\mathbf{v} < \gamma + \varepsilon.$$

In other words, $\mathbf{v}$ is a feasible solution of (D), with the value of the objective function smaller than $\gamma + \varepsilon$.

By the weak duality theorem, every feasible solution of (D) has value of the objective function at least γ. Hence (D) is a feasible and bounded linear program, and so we know that it has an optimal solution $\mathbf{y}^*$ (Theorem 4.2.3). Its value $\mathbf{b}^T\mathbf{y}^*$ is between γ and $\gamma + \varepsilon$ for every $\varepsilon > 0$, and thus it equals γ. This concludes the proof of the duality theorem. □

6.5 Farkas Lemma: An Analytic Proof

In this section we prove the geometric version of the Farkas lemma, Proposition 6.4.2, by means of elementary geometry and analysis. We are given vectors $\mathbf{a}_1, \ldots, \mathbf{a}_n$ in $\mathbb{R}^m$, and we let C be the convex cone generated by them, i.e., the set of all linear combinations with nonnegative coefficients. Proving the Farkas lemma amounts to showing that for any vector $\mathbf{b} \notin C$ there exists a hyperplane separating it from C and passing through $\mathbf{0}$. In other words, we want to exhibit a vector $\mathbf{y} \in \mathbb{R}^m$ with $\mathbf{y}^T\mathbf{b} < 0$ and $\mathbf{y}^T\mathbf{x} \geq 0$ for all $\mathbf{x} \in C$.

The plan of the proof is straightforward: We let $\mathbf{z}$ be the point of C nearest to $\mathbf{b}$ (in the Euclidean distance), and we check that the vector $\mathbf{y} = \mathbf{z} - \mathbf{b}$ is as required; see the following illustration:

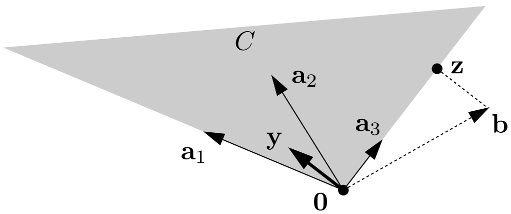

The main technical part of the proof is to show that the nearest point $\mathbf{z}$ exists. Indeed, in principle, it might happen that no point is the nearest (for example, such a situation occurs for the point 0 on the real line and the open interval $(1, 2)$; the interval contains points with distance to 0 as close to 1 as desired, but no point at distance exactly 1).

6.5.1 Lemma. *Let C be a convex cone in $\mathbb{R}^m$ generated by finitely many vectors $\mathbf{a}_1, \ldots, \mathbf{a}_n$, and let $\mathbf{b} \notin C$ be a point. Then there exists a point $\mathbf{z} \in C$ nearest to $\mathbf{b}$ (it is also unique but we won't need this).*

Proof of Proposition 6.4.2 assuming Lemma 6.5.1. As announced, we set $\mathbf{y} = \mathbf{z} - \mathbf{b}$, where $\mathbf{z}$ is a point of C nearest to $\mathbf{b}$.

First we check that $\mathbf{y}^T\mathbf{z} = 0$. This is clear for $\mathbf{z} = \mathbf{0}$. For $\mathbf{z} \neq \mathbf{0}$, if $\mathbf{z}$ were not perpendicular to $\mathbf{y}$, we could move $\mathbf{z}$ slightly along the ray $\{t\mathbf{z} : t \geq 0\} \subseteq C$ and get a point closer to $\mathbf{b}$. More formally, let us first assume that $\mathbf{y}^T\mathbf{z} > 0$, and let us set $\mathbf{z}' = (1 - \alpha)\mathbf{z}$ for a small $\alpha > 0$. We calculate $\|\mathbf{z}' - \mathbf{b}\|^2 = (\mathbf{y} - \alpha\mathbf{z})^T(\mathbf{y} - \alpha\mathbf{z}) = \|\mathbf{y}\|^2 - 2\alpha\mathbf{y}^T\mathbf{z} + \alpha^2\|\mathbf{z}\|^2$. We have $2\alpha\mathbf{y}^T\mathbf{z} > \alpha^2\|\mathbf{z}\|^2$ for all sufficiently small $\alpha > 0$, and thus $\|\mathbf{z}' - \mathbf{b}\|^2 < \|\mathbf{y}\|^2 = \|\mathbf{z} - \mathbf{b}\|^2$. This contradicts $\mathbf{z}$ being a nearest point. The case $\mathbf{y}^T\mathbf{z} < 0$ is handled similarly.

To verify $\mathbf{y}^T\mathbf{b} < 0$, we recall that $\mathbf{y} \neq \mathbf{0}$, and we compute $0 < \mathbf{y}^T\mathbf{y} = \mathbf{y}^T\mathbf{z} - \mathbf{y}^T\mathbf{b} = -\mathbf{y}^T\mathbf{b}$.

Next, let $\mathbf{x} \in C$, $\mathbf{x} \neq \mathbf{z}$. The angle $\angle\mathbf{bzx}$ has to be at least 90 degrees, for otherwise, points on the segment $\mathbf{zx}$ sufficiently close to $\mathbf{z}$ would lie closer to $\mathbf{b}$ than $\mathbf{z}$; equivalently, $(\mathbf{b}-\mathbf{z})^T(\mathbf{x}-\mathbf{z}) \le 0$ (this is similar to the above argument for $\mathbf{y}^T\mathbf{z} = 0$ and we leave a formal verification to the reader). Thus $0 \ge (\mathbf{b}-\mathbf{z})^T(\mathbf{x}-\mathbf{z}) = -\mathbf{y}^T\mathbf{x}+\mathbf{y}^T\mathbf{z} = -\mathbf{y}^T\mathbf{x}$. The Farkas lemma is proved. □

It remains to prove Lemma 6.5.1. We do it in several steps, and each of them is an interesting little fact in itself.

6.5.2 Lemma. *Let $X \subseteq \mathbb{R}^m$ be a nonempty closed set and let $\mathbf{b} \in \mathbb{R}^m$ be a point. Then X has (at least one) point nearest to $\mathbf{b}$.*

Proof. This is simple but it needs basic facts about compact sets in $\mathbb{R}^d$. Let us fix an arbitrary $\mathbf{x}_0 \in X$, let $r = \|\mathbf{x}_0 - \mathbf{b}\|$, and let $K = \{\mathbf{x} \in X : \|\mathbf{x}-\mathbf{b}\| \le r\}$. Clearly, if K has a nearest point to $\mathbf{b}$, then the same point is a point of X nearest to $\mathbf{b}$. Since K is the intersection of X with a closed ball of radius r, it is closed and bounded, and hence compact. We define the function $f\colon K \to \mathbb{R}$ by $f(\mathbf{x}) = \|\mathbf{x}-\mathbf{b}\|$. Then f is a continuous function on a compact set, and any such function attains a minimum; that is, there exists $\mathbf{z} \in K$ with $f(\mathbf{z}) \le f(\mathbf{x})$ for all $\mathbf{x} \in K$. Such a $\mathbf{z}$ is a point of K nearest to $\mathbf{b}$. □

So it remains to prove the following statement:

6.5.3 Lemma. *Every finitely generated convex cone is closed.*

This lemma is not as obvious as it might seem. As a warning example, let us consider a closed disk D in the plane with $\mathbf{0}$ on the boundary. Then the cone generated by D, that is, the set $\{t\mathbf{x} : \mathbf{x} \in D\}$, is an open half-plane plus the point $\mathbf{0}$, and thus it is not closed. Of course, this doesn't contradict to the lemma, but it shows that we must use the finiteness somehow.

Let us define a *primitive cone* in $\mathbb{R}^m$ as a convex cone generated by some $k \le m$ linearly independent vectors. Before proving Lemma 6.5.3, we deal with the following special case:

6.5.4 Lemma. *Every primitive cone P in $\mathbb{R}^m$ is closed.*

Proof. Let $P_0 \subseteq \mathbb{R}^k$ be the cone generated by the vectors $\mathbf{e}_1, \ldots, \mathbf{e}_k$ of the standard basis of $\mathbb{R}^k$. In other words, P_0 is the nonnegative orthant, and its closedness is hopefully beyond any doubt (for example, it is the intersection of the closed half-spaces $x_i \ge 0, i = 1, 2, \ldots, k$).

Let the given primitive cone $P \subseteq \mathbb{R}^m$ be generated by linearly independent vectors $\mathbf{a}_1, \ldots, \mathbf{a}_k$. We define a linear mapping $f\colon \mathbb{R}^k \to \mathbb{R}^m$ by $f(\mathbf{x}) = x_1\mathbf{a}_1 + x_2\mathbf{a}_2 + \cdots + x_k\mathbf{a}_k$. This f is injective by the linear independence of the $\mathbf{a}_j$, and we have $P = f(P_0)$. So it suffices to prove the following claim: *The image $P = f(P_0)$ of a closed set P_0 under an injective linear map $f\colon \mathbb{R}^k \to \mathbb{R}^m$ is closed.*

To see this, we let $L = f(\mathbb{R}^k)$ be the image of f. Since f is injective, it is a linear isomorphism of $\mathbb{R}^k$ and L. A linear isomorphism f has a linear inverse map $g = f^{-1}: L \to \mathbb{R}^k$. Every linear map between Euclidean spaces is continuous (this can be checked using a matrix form of the map), and we have $P = g^{-1}(P_0)$. The preimage of a closed set under a continuous map is closed by definition (while the *image* of a closed set under a continuous map need not be closed in general!), so P is closed as a subset of L. Since L is closed in $\mathbb{R}^m$ (being a linear subspace), we get that P is closed as desired. $\square$

Lemma 6.5.3 is now a consequence of Lemma 6.5.4, of the fact that the union of finitely many closed sets is closed, and of the next lemma:

6.5.5 Lemma. *Let C be a convex cone in $\mathbb{R}^m$ generated by finitely many vectors $\mathbf{a}_1, \ldots, \mathbf{a}_n$. Then C can be expressed as a union of finitely many primitive cones.*

Proof. For every $\mathbf{x} \in C$ we are going to verify that it is contained in a primitive cone generated by a suitable set of linearly independent vectors among the $\mathbf{a}_i$. We may assume $\mathbf{x} \neq \mathbf{0}$ (since $\{\mathbf{0}\}$ is the primitive cone generated by the empty set of vectors).

Let $I \subseteq \{1, 2, \ldots, n\}$ be a set of minimum possible size such that $\mathbf{x}$ lies in the convex cone generated by $A_I = \{\mathbf{a}_i : i \in I\}$ (this is a standard trick in linear algebra and in convex geometry). That is, there exist nonnegative coefficients α_i, $i \in I$, with $\mathbf{x} = \sum_{i \in I} \alpha_i \mathbf{a}_i$. The α_i are even strictly positive since if some $\alpha_i = 0$, we could delete i from I. We now want to show that the set A_I is linearly independent. For contradiction, we suppose that there is a nontrivial linear combination $\sum_{i \in I} \beta_i \mathbf{a}_i = \mathbf{0}$, where not all β_i are 0. Then there exists a real t such that all the expressions $\alpha_i - t\beta_i$ are nonnegative and at least one of them is zero. (To see this, we can first consider the case that some β_i is strictly positive, we start with $t = 0$, we let it grow, and see what happens. The case of a strictly negative β_i is analogous with t decreasing from the initial value 0.) Then the equation

$$\mathbf{x} = \sum_{i \in I} (\alpha_i - t\beta_i)\mathbf{a}_i$$

expresses $\mathbf{x}$ as a linear combination with positive coefficients of fewer than $|I|$ vectors. $\square$

6.6 Farkas Lemma from Minimally Infeasible Systems

Here we derive the Farkas lemma from an observation concerning minimally infeasible systems. A system $A\mathbf{x} \leq \mathbf{b}$ of m inequalities is called *minimally infeasible* if the system has no solution, but every subsystem obtained by dropping one inequality does have a solution.

6.6.1 Lemma. *Let $A\mathbf{x} \leq \mathbf{b}$ be a minimally infeasible system of m inequalities, and let $A^{(i)}\mathbf{x} \leq \mathbf{b}^{(i)}$ be the subsystem obtained by dropping the ith inequality, $i = 1, 2, \ldots, m$. Then for every i there exists a vector $\tilde{\mathbf{x}}^{(i)}$ such that $A^{(i)}\tilde{\mathbf{x}}^{(i)} = \mathbf{b}^{(i)}$.*

Let us set $\mathbf{a}_i = (a_{i1}, a_{i2}, \ldots, a_{in})$ and write the ith inequality as $\mathbf{a}_i^T\mathbf{x} \leq b_i$. Here is an illustration for an example in the plane ($n = 2$) with $m = 3$ inequalities:

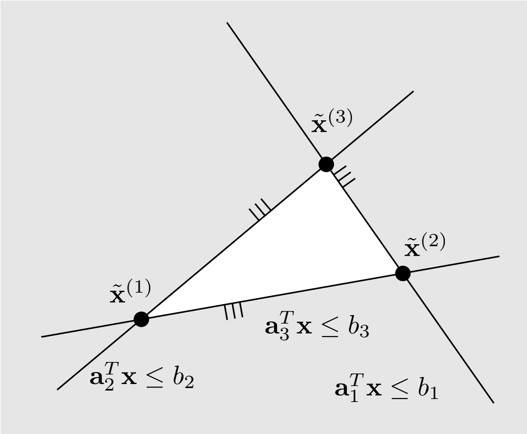

Proof. We consider the linear program

$$\begin{array}{ll} \text{minimize} & z \\ \text{subject to} & A\mathbf{x} \leq \mathbf{b} + z\mathbf{e}_i, \end{array} \qquad (\mathrm{LP}^{(i)})$$

where $\mathbf{e}_i$ is the ith unit vector. The idea of $(\mathrm{LP}^{(i)})$ is to translate the half-space $\{\mathbf{x} : \mathbf{a}_i^T\mathbf{x} \leq b_i\}$ by the minimum amount necessary to achieve feasibility. For the example illustrated above and $i = 3$, this results in the following picture:

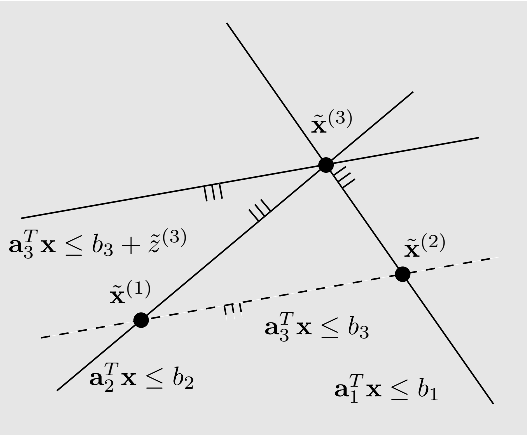

To show formally that (LP$^{(i)}$) has an optimal solution, we first argue that it has a feasible solution. Indeed, by the assumption, the system $A^{(i)}\mathbf{x} \leq \mathbf{b}^{(i)}$ has at least one solution. Let us fix an arbitrary solution of this system and denote it by $\overline{\mathbf{x}}$. We put $\overline{z} = \mathbf{a}_i^T\overline{\mathbf{x}} - b_i$, and we note that the vector $(\overline{\mathbf{x}}, \overline{z})$ is a feasible solution of the linear program (LP$^{(i)}$).

Next, we note that (LP$^{(i)}$) is also bounded, since $A\mathbf{x} \leq \mathbf{b}$ has no solution. Therefore, the linear program has an optimal solution $(\tilde{\mathbf{x}}^{(i)}, \tilde{z}^{(i)})$ with $\tilde{z}^{(i)} > 0$ by Theorem 4.2.3.

We claim that the just defined $\tilde{\mathbf{x}}^{(i)}$ satisfies $A^{(i)}\tilde{\mathbf{x}}^{(i)} = \mathbf{b}^{(i)}$. We already know that $A^{(i)}\tilde{\mathbf{x}}^{(i)} \leq \mathbf{b}^{(i)}$. Let us suppose for contradiction that $\mathbf{a}_j^T\tilde{\mathbf{x}}^{(i)} = b_j - \varepsilon$ for some $j \neq i$ and $\varepsilon > 0$. We will show that then $(\tilde{\mathbf{x}}^{(i)}, \tilde{z}^{(i)})$ cannot be optimal for (LP$^{(i)}$). To this end, let us consider an optimal solution $(\tilde{\mathbf{x}}^{(j)}, \tilde{z}^{(j)})$ of (LP$^{(j)}$). The idea is that by moving the point $(\tilde{\mathbf{x}}^{(i)}, \tilde{z}^{(i)})$ slightly towards $(\tilde{\mathbf{x}}^{(j)}, \tilde{z}^{(j)})$, we remain feasible for (LP$^{(i)}$), but we improve the objective function of (LP$^{(i)}$). More formally, for a real number $t \geq 0$, we define $\tilde{\mathbf{x}}(t) = (1-t)\tilde{\mathbf{x}}^{(i)} + t\tilde{\mathbf{x}}^{(j)}$. It follows that

$$\begin{aligned} \mathbf{a}_j^T\tilde{\mathbf{x}}(t) &\leq b_j - (1-t)\varepsilon + t\tilde{z}^{(j)}, \\ \mathbf{a}_i^T\tilde{\mathbf{x}}(t) &\leq b_i + (1-t)\tilde{z}^{(i)}, \\ \mathbf{a}_k^T\tilde{\mathbf{x}}(t) &\leq b_k, \text{ for all } k \neq i, j. \end{aligned}$$

Thus for t sufficiently small, namely, for $0 < t \leq (1-t)\varepsilon/\tilde{z}^{(j)}$, the pair $(\tilde{\mathbf{x}}(t), (1-t)\tilde{z}^{(i)})$ is a feasible solution of (LP$^{(i)}$) with objective function strictly smaller than $\tilde{z}^{(i)}$, contradicting the assumed optimality of $(\tilde{\mathbf{x}}^{(i)}, \tilde{z}^{(i)})$. Thus, $A^{(i)}\tilde{\mathbf{x}}^{(i)} = \mathbf{b}^{(i)}$ and the lemma is proved. □

We need another lemma that proves an "easy" variant of the Farkas lemma, concerned with *arbitrary* solutions of systems of *equalities.*

This lemma establishes the implication in the bottom right corner of the table of Farkas lemma variants on page 92.

6.6.2 Lemma. *The system* $A\mathbf{x} = \mathbf{b}$ *has a solution if and only if every* $\mathbf{y} \in \mathbb{R}^m$ *with* $\mathbf{y}^T A = \mathbf{0}^T$ *also satisfies* $\mathbf{y}^T\mathbf{b} = 0$.

Proof. One direction is easy. If $A\mathbf{x} = \mathbf{b}$ has some solution $\tilde{\mathbf{x}}$, and if $\mathbf{y}^T A = \mathbf{0}^T$, then $0 = \mathbf{0}^T\tilde{\mathbf{x}} = \mathbf{y}^T A\tilde{\mathbf{x}} = \mathbf{y}^T\mathbf{b}$.

If $A\mathbf{x} = \mathbf{b}$ has no solution, we need to find a vector $\mathbf{y}$ such that $\mathbf{y}^T A = \mathbf{0}^T$ and $\mathbf{y}^T\mathbf{b} \neq 0$. Let us define $r = \operatorname{rank}(A)$ and consider the $m \times (n+1)$ matrix $(A \,|\, \mathbf{b})$. This matrix has rank $r + 1$ since the last column is not a linear combination of the first n columns. For the very same reason, the matrix

$$\left(\begin{array}{c|c} A & \mathbf{b} \\ \hline \mathbf{0}^T & -1 \end{array}\right)$$

has rank $r+1$. This shows that the row vector $(\mathbf{0}^T \mid -1)$ is a linear combination of rows of $(A \mid \mathbf{b})$, and the coefficients of this linear combination define a vector $\mathbf{y} \in \mathbb{R}^m$ with $\mathbf{y}^T A = \mathbf{0}^T$ and $\mathbf{y}^T \mathbf{b} = -1$, as desired. $\square$

Now we proceed to the proof of the Farkas lemma. The variant that results most naturally is the one with an arbitrary solution of $A\mathbf{x} \le \mathbf{b}$, that is, Proposition 6.4.3(iii).

Proof of Proposition 6.4.3(iii). As in Lemma 6.6.2, one direction is easy: If $A\mathbf{x} \le \mathbf{b}$ has some solution $\tilde{\mathbf{x}}$, and if $\mathbf{y} \ge \mathbf{0}$, $\mathbf{y}^T A = \mathbf{0}^T$, we get $0 = \mathbf{0}^T \tilde{\mathbf{x}} = \mathbf{y}^T A \tilde{\mathbf{x}} \le \mathbf{y}^T \mathbf{b}$. The interesting case is that $A\mathbf{x} \le \mathbf{b}$ has no solution. Our task is then to construct a vector $\mathbf{y} \ge \mathbf{0}$ satisfying $\mathbf{y}^T A = \mathbf{0}^T$ and $\mathbf{y}^T \mathbf{b} < 0$.

We may assume that $A\mathbf{x} \le \mathbf{b}$ is minimally infeasible, by restricting to a suitable subsystem: A vector $\mathbf{y}$ for this subsystem can be extended to work for the original system by inserting zeros at appropriate places.

Since $A\mathbf{x} \le \mathbf{b}$ has no solution, the system $A\mathbf{x} = \mathbf{b}$ has no solution either. By Lemma 6.6.2, there exists a vector $\mathbf{y} \in \mathbb{R}^m$ such that $\mathbf{y}^T A = \mathbf{0}^T$ and $\mathbf{y}^T \mathbf{b} \neq 0$. By possibly changing signs, we may assume that $\mathbf{y}^T \mathbf{b} < 0$. We will show that this vector also satisfies $\mathbf{y} \ge \mathbf{0}$, and this will finish the proof. To this end, we fix $i \in \{1, 2, \ldots, m\}$ and consider the vector $\tilde{\mathbf{x}}^{(i)}$ as in Lemma 6.6.1 above. With the terminology of the lemma, we have $A^{(i)} \tilde{\mathbf{x}}^{(i)} = \mathbf{b}^{(i)}$, and using $\mathbf{y}^T A = \mathbf{0}^T$, we can write

$$y_i(\mathbf{a}_i^T \tilde{\mathbf{x}}^{(i)} - b_i) = \mathbf{y}^T (A\tilde{\mathbf{x}}^{(i)} - \mathbf{b}) = -\mathbf{y}^T \mathbf{b} > 0.$$

Proposition 6.4.3(iii) is proved. $\square$

This proof of the Farkas lemma is based on the paper

M. Conforti, M. Di Summa, and G. Zambelli: Minimally infeasible set partitioning problems with balanced constraints, *Mathematics of Operations Research*, to appear.

The proof given there is even more elementary than ours in the sense that it does not use linear programming. We have chosen the linear programming approach since we find it somewhat more transparent.

6.7 Farkas Lemma from the Fourier–Motzkin Elimination

When explaining the "logical view" of the Farkas lemma in Section 6.4, we started with a system of 3 inequalities and combined pairs of inequalities together, until we managed to eliminate all variables and obtained the obviously unsatisfiable inequality $0 \le -1$. The Fourier–Motzkin elimination is a systematic procedure for eliminating all variables from an arbitrary system $A\mathbf{x} \le \mathbf{b}$ of linear inequalities. If the final inequalities with no variables hold, we can reconstruct a solution of the original system by tracing the computations backward, and if one of the final inequalities does not hold, it certifies that the original system has no solution.

The Fourier–Motzkin elimination is similar in spirit to Gaussian elimination for systems of linear equations, and it is just as simple. As in Gaussian elimination, variables are removed one at a time, but there is a price to pay: To get rid of one variable, we typically have to introduce many new inequalities, so that the method becomes impractical already for moderately large systems. The Fourier–Motzkin elimination can be considered as a simple but inefficient alternative to the simplex method. For the purpose of proving statements about systems of inequalities, efficiency is not a concern, so it is the simplicity of the Fourier–Motzkin elimination that makes it a very handy tool.

As an example, let us consider the following system of 5 inequalities in 3 variables:

$$\begin{array}{rcrcrcr} 2x & - & 5y & + & 4z & \le & 10 \\ 3x & - & 6y & + & 3z & \le & 9 \\ 5x & + & 10y & - & z & \le & 15 \\ -x & + & 5y & - & 2z & \le & -7 \\ -3x & + & 2y & + & 6z & \le & 12. \end{array} \tag{6.7}$$

In the first step we would like to eliminate x. For a moment let us imagine that y and z are some fixed real numbers, and let us ask under what conditions we can choose a value of x such that together with the given values y and z it satisfies (6.7). The first three inequalities impose an upper bound on x, while the remaining two impose a lower bound. To make this clearer, we rewrite the system as follows:

$$\begin{array}{rcrcrcr} x & \le & 5 & + & \frac{5}{2}y & - & 2z \\ x & \le & 3 & + & 2y & - & z \\ x & \le & 3 & - & 2y & + & \frac{1}{5}z \\ x & \ge & 7 & + & 5y & - & 2z \\ x & \ge & -4 & + & \frac{2}{3}y & + & 2z. \end{array}$$

So given y and z, the admissible values of x are exactly those in the interval from $\max(7+5y-2z, -4+\frac{2}{3}y+2z)$ to $\min(5+\frac{5}{2}y-2z, 3+2y-z, 3-2y+\frac{1}{5}z)$. If this interval happens to be empty, there is no admissible x. So the inequality

$$\begin{aligned} &\max(7 + 5y - 2z, -4 + \tfrac{2}{3}y + 2z) \\ &\quad \le \min(5 + \tfrac{5}{2}y - 2z, 3 + 2y - z, 3 - 2y + \tfrac{1}{5}z) \end{aligned} \tag{6.8}$$

is *equivalent* to the existence of x that together with the considered y and z solves (6.7). The key observation in the Fourier–Motzkin elimination is that (6.8) can be rewritten as a system of linear inequalities in the variables y and z. The inequalities simply say that *each of the lower bounds is less than or equal to each of the upper bounds*:

$$\begin{aligned}
7 + 5y - 2z &\le 5 + \tfrac{5}{2}y - 2z\\
7 + 5y - 2z &\le 3 + 2y - z\\
7 + 5y - 2z &\le 3 - 2y + \tfrac{1}{5}z\\
-4 + \tfrac{2}{3}y + 2z &\le 5 + \tfrac{5}{2}y - 2z\\
-4 + \tfrac{2}{3}y + 2z &\le 3 + 2y - z\\
-4 + \tfrac{2}{3}y + 2z &\le 3 - 2y + \tfrac{1}{5}z.
\end{aligned}$$

If we rewrite this system in the usual form $A\mathbf{x} \le \mathbf{b}$, we arrive at

$$\begin{array}{rcrcr}
\tfrac{5}{2}y & & & \le & -2\\
3y & - & z & \le & -4\\
7y & - & \tfrac{11}{5}z & \le & -4\\
-\tfrac{11}{6}y & + & 4z & \le & 9\\
-\tfrac{4}{3}y & + & 3z & \le & 7\\
\tfrac{8}{3}y & + & \tfrac{9}{5}z & \le & 7.
\end{array} \tag{6.9}$$

This system has a solution exactly if the original system (6.7) has one, but it has one variable fewer. The reader is invited to continue with this example, eliminating y and then z. We note that (6.9) gives 4 upper bounds for y and 2 lower bounds, and hence we obtain 8 inequalities after eliminating y.

For larger systems the number of inequalities generated by the Fourier–Motzkin elimination tends to explode. This wasn't so apparent for our small example, but if we have m inequalities and, say, half of them impose upper bounds on the first variable and half impose lower bounds, then we get about $m^2/4$ inequalities after eliminating the first variable, about $m^4/16$ after eliminating the second variable (again, provided that about half of the inequalities give upper bounds for the second variable and half lower bounds), etc.

Now we formulate the procedure in general.

Claim. *Let* $A\mathbf{x} \le \mathbf{b}$ *be a system with* $n \ge 1$ *variables and* m *inequalities. There is a system* $A'\mathbf{x}' \le \mathbf{b}'$ *with* $n-1$ *variables and at most* $\max(m, m^2/4)$ *inequalities, with the following properties:*

(i) $A\mathbf{x} \le \mathbf{b}$ *has a solution if and only if* $A'\mathbf{x}' \le \mathbf{b}'$ *has a solution, and*

(ii) *each inequality of* $A'\mathbf{x}' \le \mathbf{b}'$ *is a positive linear combination of some inequalities from* $A\mathbf{x} \le \mathbf{b}$.

Proof. We classify the inequalities into three groups, depending on the coefficient of x_1. We call the ith inequality of $A\mathbf{x} \le \mathbf{b}$ a *ceiling* if $a_{i1} > 0$, and we call it a *floor* if $a_{i1} < 0$. Otherwise (if $a_{i1} = 0$), it is a *level*. Let $C, F, L \subseteq \{1, \ldots, m\}$ collect the indices of ceilings, floors, and levels. We may assume that

$$a_{i1} = \begin{cases} 1 & \text{if } i \in C\\ -1 & \text{if } i \in F\\ 0 & \text{if } i \in L. \end{cases} \tag{6.10}$$

This situation can be reached by multiplying each inequality in $A\mathbf{x} \leq \mathbf{b}$ by a suitable positive number, which does not change the set of solutions.

Now we can eliminate x_1 between all pairs of ceilings and floors, by simply adding up the two inequalities for each pair.

If $\mathbf{x}'$ is the (possibly empty) vector $(x_2, \ldots, x_n)$, and $\mathbf{a}'_i$ is the (possibly empty) vector $(a_{i2}, \ldots, a_{in})$, then the following inequalities are implied by $A\mathbf{x} \leq \mathbf{b}$:

$$\mathbf{a}'^T_j \mathbf{x}' + \mathbf{a}'^T_k \mathbf{x}' \leq b_j + b_k, \quad j \in C, \, k \in F. \tag{6.11}$$

The level inequalities of $A\mathbf{x} \leq \mathbf{b}$ can be rewritten as

$$\mathbf{a}'^T_\ell \mathbf{x}' \leq b_\ell, \quad \ell \in L. \tag{6.12}$$

So if $A\mathbf{x} \leq \mathbf{b}$ has a solution, then the system of $|C| \cdot |F| + |L|$ inequalities in $n-1$ variables given by (6.11) and (6.12) has a solution as well. Conversely, if the latter system has a solution $\tilde{\mathbf{x}}' = (\tilde{x}_2, \ldots, \tilde{x}_n)$, we can determine a suitable value $\tilde{x}_1$ such that the vector $(\tilde{x}_1, \tilde{x}_2, \ldots, \tilde{x}_n)$ solves $A\mathbf{x} \leq \mathbf{b}$. To find $\tilde{x}_1$, we first observe that (6.11) is equivalent to

$$\mathbf{a}'^T_k \mathbf{x}' - b_k \leq b_j - \mathbf{a}'^T_j \mathbf{x}', \quad j \in C, \, k \in F.$$

This in particular implies

$$\max_{k \in F} \left(\mathbf{a}'^T_k \tilde{\mathbf{x}}' - b_k \right) \leq \min_{j \in C} \left(b_j - \mathbf{a}'^T_j \tilde{\mathbf{x}}' \right).$$

We let $\tilde{x}_1$ be any value between these bounds. It follows that

$$\begin{aligned} \tilde{x}_1 + \mathbf{a}'^T_j \tilde{\mathbf{x}}' &\leq b_j, \quad j \in C, \\ -\tilde{x}_1 + \mathbf{a}'^T_k \tilde{\mathbf{x}}' &\leq b_k, \quad k \in F. \end{aligned}$$

By our assumption (6.10), we have a feasible solution of the original system $A\mathbf{x} \leq \mathbf{b}$. We note that this argument also works for $C = \emptyset$ or $F = \emptyset$, with the usual convention that $\max_{t \in \emptyset} f(t) = -\infty$ and $\min_{t \in \emptyset} f(t) = \infty$.

Now we can prove the Farkas lemma. The variant that results most naturally from the Fourier–Motzkin elimination is (as in Section 6.6) the one with an arbitrary solution of $A\mathbf{x} \leq \mathbf{b}$, that is, Proposition 6.4.3(iii).

Proof of Proposition 6.4.3(iii). One direction is easy. If $A\mathbf{x} \leq \mathbf{b}$ has some solution $\tilde{\mathbf{x}}$, and $\mathbf{y} \geq \mathbf{0}$ satisfies $\mathbf{y}^T A = \mathbf{0}^T$, we get $0 = \mathbf{0}^T \tilde{\mathbf{x}} = \mathbf{y}^T A \tilde{\mathbf{x}} \leq \mathbf{y}^T \mathbf{b}$. If $A\mathbf{x} \leq \mathbf{b}$ has no solution, then our task is to construct a vector $\mathbf{y}$ satisfying

$$\mathbf{y} \geq \mathbf{0}, \quad \mathbf{y}^T A = \mathbf{0}^T, \quad \text{and } \mathbf{y}^T \mathbf{b} < 0. \tag{6.13}$$

To find such a witness of infeasibility, we use induction on the number of variables. Let us first consider the base case in which the system $A\mathbf{x} \le \mathbf{b}$ has no variables, meaning that it is of the form $\mathbf{0} \le \mathbf{b}$ with $b_i < 0$ for some i. We set $\mathbf{y} = \mathbf{e}_i$ (the ith unit vector), and this clearly satisfies the requirements for $\mathbf{y}$ being a witness of infeasibility (the condition $\mathbf{y}^T A = \mathbf{0}^T$ is vacuous, since A has no column).

If $A\mathbf{x} \le \mathbf{b}$ has at least one variable, we perform a step of the Fourier–Motzkin elimination. This yields an infeasible system $A'\mathbf{x}' \le \mathbf{b}'$, consisting of the inequalities (6.11) and (6.12). Because the latter system has one variable fewer, we inductively find a witness of infeasibility $\mathbf{y}'$ for it. We recall that all inequalities of $A'\mathbf{x}' \le \mathbf{b}'$ are positive linear combinations of original inequalities; equivalently, there is an $m \times m$ matrix M with all entries nonnegative and

$$(\mathbf{0} \,|\, A') = MA, \quad \mathbf{b}' = M\mathbf{b}.$$

We claim that $\mathbf{y} = M^T\mathbf{y}'$ is a witness of infeasibility for the original system $A\mathbf{x} \le \mathbf{b}$. Indeed, we have $\mathbf{y}^T A = \mathbf{y}'^T MA = \mathbf{y}'^T(\mathbf{0} \,|\, A') = \mathbf{0}^T$ and $\mathbf{y}^T\mathbf{b} = \mathbf{y}'^T M\mathbf{b} = \mathbf{y}'^T\mathbf{b}' < 0$, since $\mathbf{y}'$ is a witness of infeasibility for $A'\mathbf{x}' \le \mathbf{b}'$. The condition $\mathbf{y} \ge \mathbf{0}$ follows from $\mathbf{y}' \ge \mathbf{0}$ by the nonnegativity of M. □