2 The linear one-dimensional case.

2.1 Extremizers are not in the interior.

Let $c \in \mathbb{R}$ be given and assume that $c \neq 0$. The collection of nontrivial continuous linear functions is given by $f_c : \mathbb{R} \to \mathbb{R}$ defined by $f_c(x) = cx$. All $x \in \mathbb{R}$ are in the interior of $\mathbb{R}$. The derivative $f_c'(x) = c$ is never zero. It follows that there are no extremizers. If $C \subset \mathbb{R}$ and extremizers of $f_c : C \to \mathbb{R}$ are sought, then the only possibility are points in C that are not in the interior of C.

2.2 Linear constraints.

A standard linear constraint in one variable has the form $ax \leq b$, where $a,b \in \mathbb{R}$ and $a \neq 0$. The corresponding constrained set is given by

$$C = \left\{x \in \mathbb{R} \middle| \, ax \leq b\right\} = \begin{cases} \left\{x \in \mathbb{R} \middle| \, x \leq b/a\right\} \text{ if } a > 0 \\ \left\{x \in \mathbb{R} \middle| \, x \geq b/a\right\} \text{ if } a < 0 \end{cases} = \begin{cases} (-\infty, b/a] \text{ if } a > 0 \\ [b/a, +\infty) \text{ if } a < 0 \end{cases}.$$

The set C is a closed. The only $x \in C$ that is not in the interior of C is given by $x = b/a$. Let $m \in \mathbb{N}$ be given. Write $a_1 x \leq b_1, \cdots, a_m x \leq b_m$ when m linear constraints are imposed at once. Denote the corresponding sets by $C_1, \cdots, C_m$. Imposing these constraints at once means $x \in \mathbb{R}$ must satisfy $a_j x \leq b_j$ for every single $j \in \{1, \ldots, m\}$. The constrained set is therefore given by $C = C_1 \cap \cdots \cap C_m$, which is a closed set.

2.3 The primal problem.

The *primal* problem deals only with maximizers of functions $f_c : \mathbb{R} \to \mathbb{R}$ given by $f_c(x) = cx$ with $c \neq 0$. It is assumed that $x \geq 0$, and $a_j x \leq b_j$ for $j \in \{1, \ldots, m\}$ and some given $m \in \mathbb{N}$. The following notation is used:

$$\max cx \text{ when } \begin{cases} \mathrm{a}_1 x \leq b_1 \\ \vdots \\ a_m x \leq b_m \\ x \geq 0 \end{cases}.$$

With the help of the column matrices

$$A = \begin{bmatrix} a_1 \\ \vdots \\ a_m \end{bmatrix}, \; b = \begin{bmatrix} b_1 \\ \vdots \\ b_m \end{bmatrix},$$

this becomes

$$\max cx \text{ when } \begin{cases} Ax \leq b \\ x \geq 0 \end{cases}.$$

The inequality sign between two matrices of the same order indicates that each pair of corresponding entries satisfies the inequality. The collection of all $x \in \mathbb{R}$ such that $Ax \le b$ and $x \ge 0$ is known as the *feasible set* of the primal problem. If $a_j > 0$ and $b_j < 0$ for some j, then the constraints are *inconsistent* and the feasible set is empty. Suppose the constraints are consistent with $a_j > 0$ and $b_j = 0$ for some j, then the case is *trivial* since $x = 0$ is the only point in the feasible set. If $a_j < 0$ and $b_j \ge 0$ for some j, then the constraint is *redundant* since it is subsumed by the constraint $x \ge 0$. In order to have a nonempty nontrivial feasible set with no redundant constraints it must be that $a_j > 0$ and $b_j > 0$, or $a_j < 0$ and $b_j < 0$ for all j. Let $\alpha = \max\left\{b_j / a_j \,\middle|\, a_j < 0, b_j < 0\right\}$, or $\alpha = 0$ if there is no j such that $a_j < 0$ and $b_j < 0$. Let $\beta = \min\left\{b_j / a_j \,\middle|\, a_j > 0, b_j > 0\right\}$, or $\beta = +\infty$ if there is no j such that $a_j > 0$ and $b_j > 0$. There are three possibilities. If $\alpha > \beta$, then the feasible set is empty. If $\alpha = \beta$, then $x = \alpha = \beta$ is the only number in the feasible set. If $\alpha < \beta$, then the feasible set is equal to $[\alpha, \beta]$ when $\beta < +\infty$, and $[\alpha, +\infty)$ if $\beta = +\infty$. If $c > 0$, then $x = \alpha$ is a minimizer. If $c < 0$, then $x = \alpha$ is a maximizer. If $\beta < +\infty$, then $x = \beta$ is a maximizer if $c > 0$, and a minimizer if $c < 0$.

2.4 Simple primal problems.

The simplest case is given by

$$\max cx \text{ when } x \ge 0.$$

If $c > 0$, then there is no maximizer. If $c < 0$, then $x = 0$ is a maximizer. The second simplest case is given by

$$\max cx \text{ when } \begin{cases} \mathrm{ax} \le \mathrm{b} \\ x \ge 0 \end{cases}.$$

If $a > 0$ and $b < 0$, then there is no maximizer. If $a > 0$ and $b = 0$, then $x = 0$ is the only point in the feasible set and therefore a maximizer. If $a > 0$ and $b > 0$, then the feasible set is given by $[0, b/a]$. In this case $c > 0$ implies that $x = b/a$ is a maximizer, and $c < 0$ implies that $x = 0$ is a maximizer. If $a < 0$ and $b \ge 0$, then the feasible set is equal to $[0, +\infty)$. This time there is no maximizer if $c > 0$, and $x = 0$ is a maximizer if $c < 0$. Finally, if $a < 0$ and $b < 0$, then the feasible set is given by $[b/a, +\infty)$. Again, there is no maximizer if $c > 0$, and $x = b/a$ is a maximizer if $c < 0$.

2.5 Duality.

To get a concrete interpretation of the primal problem, assume x is the mass in kilograms of the product to be produced and ultimately sold. Let c be the selling price in dollars per kilogram. The efficiency at different stages j of the production is given by a_j in hours per kilogram. The available resources at stage j are given by b_j in hours. The problem is to maximize the income

cx in dollars, without exceeding the available resources. It follows that the hours $a_j x$ it takes to produce the product at stage j must be less than or equal to the available hours b_j for each j. It is assumed that there exist buyers of the produced product at the price c. From a different point of view, competitors are interested in the resources available to the producer. Assume a competitor is willing to offer λ_j in dollars per hour for the resources used at stage j. The expression $\lambda_1 b_1 + \cdots \lambda_m b_m$ is the competitor's cost in dollars for the resources available to the producer. The expression $\lambda_1 a_1 + \cdots \lambda_m a_m$ is the cost in dollars per kilogram of the product as calculated by the competitor. If $\lambda_1 a_1 + \cdots + \lambda_m a_m < c$, then it is no longer clear that there are buyers of the produced product at the price c. The competitor would like to minimize $\lambda_1 b_1 + \cdots \lambda_m b_m$ subject to the constraints $\lambda_1 a_1 + \cdots \lambda_m a_m \geq c$ and $\lambda_j \geq 0$ for all j. This leads to the *dual* problem:

$$\min \lambda^T b \text{ when} \begin{cases} A^T \lambda \geq c \\ \lambda \geq 0 \end{cases}.$$

The transposes $\lambda^T = \begin{bmatrix} \lambda_1 & \cdots & \lambda_m \end{bmatrix}$ and $A^T = \begin{bmatrix} a_1 & \cdots & a_m \end{bmatrix}$ are here used together with the matrix product. The 1×1 matrix $\lambda^T b$ is identified with the value of its entry and similarly for c. This form of the dual problem has several syntactical advantages made evident when the primal and the dual problem are analyzed together. Incidentally, the primal function $c^T x$ is referred to as the *primal objective*. The dual function $\lambda^T b$ is of course the *dual objective*.

Exercise

Set-up but do not solve the dual problem to

$$\max 2x \text{ when} \begin{cases} 3x \leq 5 \\ 7x \leq 11 \\ x \geq 0 \end{cases}.$$

Exercise

Solve the primal problem corresponding to

$$\min 2\lambda_1 + 3\lambda_2 + 5\lambda_3 \text{ when} \begin{cases} 7\lambda_1 + 11\lambda_2 + 13\lambda_3 \geq 17 \\ \lambda_1, \lambda_2, \lambda_3 \geq 0 \end{cases}.$$

2.6 Slack variables.

In the primal problem

$$\max cx \text{ when} \begin{cases} a_1 x \leq b_1 \\ \vdots \\ a_m x \leq b_m \\ x \geq 0 \end{cases},$$

let $y_j = b_j - a_j x$ be the j'th *slack variable*. The j'th constraint is satisfied if and only if

$y_j \geq 0$. Let

$$\begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} x \\ y_1 \\ \vdots \\ y_m \end{bmatrix}, \begin{bmatrix} A & I \end{bmatrix} = \begin{bmatrix} a_1 & 1 & 0 & \cdots & 0 & 0 \\ a_2 & 0 & 1 & & 0 & 0 \\ \vdots & \vdots & & \ddots & & \vdots \\ a_{m-1} & 0 & 0 & & 1 & 0 \\ a_m & 0 & 0 & \cdots & 0 & 1 \end{bmatrix}.$$

The constraints in the primal problem have the new form

$$\begin{aligned} \begin{bmatrix} A & I \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} &= b \\ \begin{bmatrix} x \\ y \end{bmatrix} &\geq 0 \end{aligned}.$$

The dimension of the problem has increased from 1 to $m+1$. In the dual problem

$$\min \lambda^T b \text{ when } \begin{cases} A^T \lambda \geq c \\ \quad \lambda \geq 0 \end{cases},$$

let $\mu = A^T \lambda - c$ be the one slack variable. The first constraint is satisfied if and only if $\mu \geq 0$. The constraints have the new form

$$\begin{aligned} A^T \lambda - \mu &= c \\ \begin{bmatrix} \lambda \\ \mu \end{bmatrix} &\geq 0 \end{aligned}.$$

The dimension of the problem has increased from m to $m+1$.

2.7 Tableau.

With the slack variables in place, the data of the primal problem is recorded in the following *initial tableau*

$$\begin{array}{cccccc|c} a_1 & 1 & 0 & \cdots & 0 & 0 & b_1 \\ a_2 & 0 & 1 & \cdots & 0 & 0 & b_2 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ a_{m-1} & 0 & 0 & \cdots & 1 & 0 & b_{m-1} \\ a_m & 0 & 0 & \cdots & 0 & 1 & b_m \\ \hline c & 0 & 0 & \cdots & 0 & 0 & 0 \end{array}.$$

The data above the horizontal line corresponds to the equality constraints. The bottom row has the coefficients of the function to be maximized to the left of the vertical line, and the negative of the value of the function to the right. To find a solution that is not in the interior of the feasible set it suffices to let one of the variables be equal to zero. It is convenient to choose x to be zero so that $s_j = b_j$ for all j, because $s_j = b_j$ is readily seen in the tableaux. The columns corresponding to the slack variables each look like a standard basis vector. The corresponding variable is referred to as a *basic variable*, and all other variables are *non-basic* variables. The non-basic variables will always have value zero.

2.8 Pivoting rule.

Suppose x is a non-basic variable in the feasible set. It follows that all b_j are positive and it may be assumed that all a_j are positive as well. Assume $c > 0$ so that it is clear that it is beneficial to increase x. Using elementary row operations the x column will be transformed so that there is a single 1 in row j and 0 in all other rows. The next tableau is given by

$$\begin{array}{cccccc|c}
0 & 1 & 0 & -\dfrac{a_1}{a_j} & 0 & 0 & b_1 - a_1 \dfrac{b_j}{a_j} \\
\vdots & 0 & 1 & \cdots & 0 & 0 & \vdots \\
1 & \vdots & \vdots & \dfrac{1}{a_j} & \vdots & \vdots & \dfrac{b_j}{a_j} \\
\vdots & 0 & 0 & \cdots & 1 & 0 & \vdots \\
0 & 0 & 0 & -\dfrac{a_m}{a_j} & 0 & 1 & b_m - a_m \dfrac{b_j}{a_j} \\
\hline
0 & 0 & 0 & -\dfrac{c}{a_j} & 0 & 0 & -c\dfrac{b_j}{a_j}
\end{array}.$$

The subsequent value of x is b_j / a_j. Examine the column to the right of the vertical line in the tableau. Since all variables are required to be greater than or equal to zero, it follows that it is necessary to choose j so that $b_j / a_j = \min\left\{a_k / b_k \,\middle|\, k \in \{1, \dots, m\}\right\}$. This choice is referred to as the *pivoting rule*. The j'th slack variable is equal to zero. The problem is transformed, and now the function

$$s_j \mapsto -\frac{c}{a_j} s_j$$

is to be maximized. Since $s_j \geq 0$ is required and $-c / a_j < 0$, it follows that there is nothing to be gained by increasing s_j from zero, which is the current value. The current solution $x = b_j / a_j$ is optimal. The new tableau is in fact the *final tableau*.

2.9 Duality and the final tableau.

If $\lambda_k = 0$ for all $k \neq j$ and $\lambda_j = c / a_j$, then $\lambda^T b = c\dfrac{b_j}{a_j}$ and $c = A^T \lambda \geq c$. It follows that the final tableau has produced a feasible λ for the dual problem with the property that the value of the dual objective is equal to the value of the primal objective. It is shown below how this forces λ to be optimal in the dual problem.