2 Linear Programming

2.1 Introductory examples

2.1.1 Problem 1

Let's start with a problem from the text. Farmer Bob has 625 acres of land and 1000 acre-feet of water. Available labor is 300 hours/week. The following table summarizes cost and revenue components for his crop choices.

	$Corn = x$	$Wheat = y$	$Oats = z$
$irrigation(feet)$	3.0	1.0	1.5
$labor(hours/week$	0.8	0.2	0.3
$yield(dollars/acre)$	400	200	250

Farmer Bob naturally wants to maximize his profit

$$p(x, y, z) = 400x + 200y + 250z,$$

subject to the land constraint

$$x + y + z \leq 625,$$

the water constraint

$$3x + y + 1.5z \leq 1000,$$

and the labor constraint

$$0.8x + 0.2y + 0.3z \leq 300.$$

Although this is a simple linear programming problem, it is already complex enough that the basic geometric structure is not clear. To help clarify the geometric structure, and to see some of the issues that might arise, we consider a number of simpler problems.

2.1.2 Problem 2

First let's look at a 2-crop version of the problem. Again, the table summarizes cost and revenue components for Bob's crop choices.

	$Corn = x$	$Wheat = y$
$irrigation(feet)$	3.0	1.0
$labor(hours/week$	0.8	0.2
$yield(dollars/acre)$	400	200

Now the profit is

$$p(x, y) = 400x + 200y,$$

subject to the land constraint

$$x + y \leq 250,$$

the water constraint

$$3x + y \leq 450,$$

and the labor constraint

$$0.8x + 0.2y \leq 100.$$

The feasible set of $x - y$ values is just the set of nonnegative pairs x, y lying in the intersection of the regions satisfying the constraints.

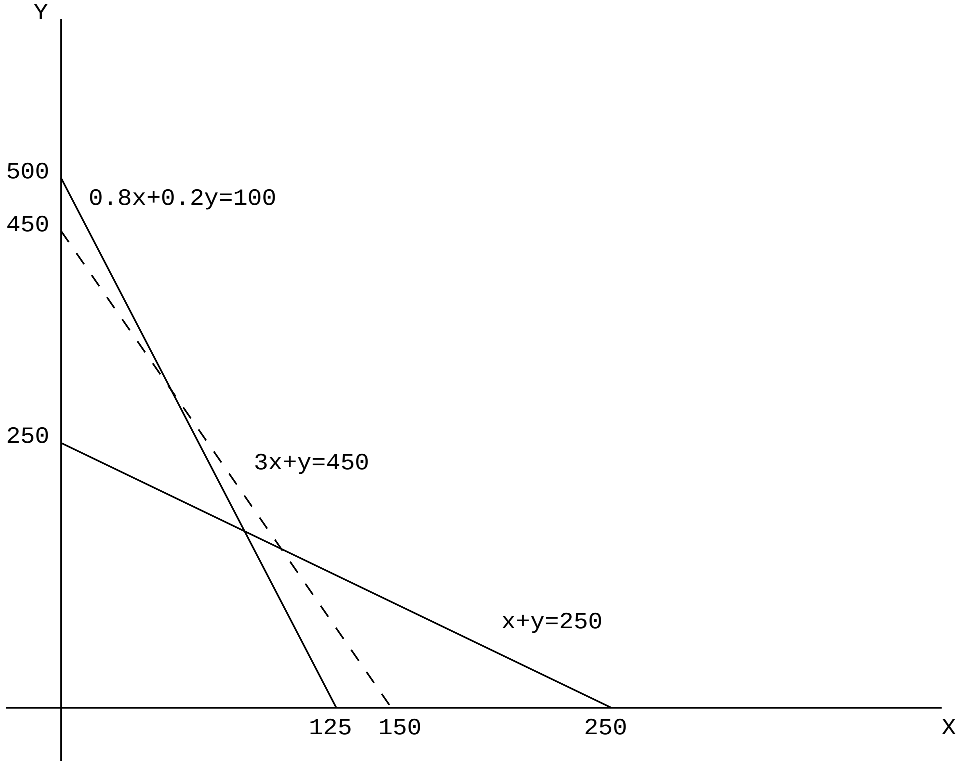

Figure 2.a: Feasible region for problem 2

As is typical for linear programming problems,

$$\nabla p(x, y) = (400, 200)$$

is a nonzero constant, so there is no interior maximum in the feasible region. Also notice that the water constraint is irrelevent. If the other constraints are satisfied, the water constraint is satisfied too.

Finally, when x and y are constrained to lie on the boundary lines of the feasible region, then $p(x, y)$ is an affine function (polynomial of degree at most 1, or linear plus constant) of one variable. Any maximum must be at an endpoint, that is a corner of the feasible region. Checking the values of the profit at the corners, we find the best choice is when $x = 83$, $y = 167$, and the resulting profit is $\$66,600$.

2.1.3 Problem 3

Let's return to the original three crop problem, but assume that only the water constraint applies. There is no advantage to using less than the full allotment of water, so the constraint now is the equality constraint

$$3x + y + 1.5z = 1000.$$

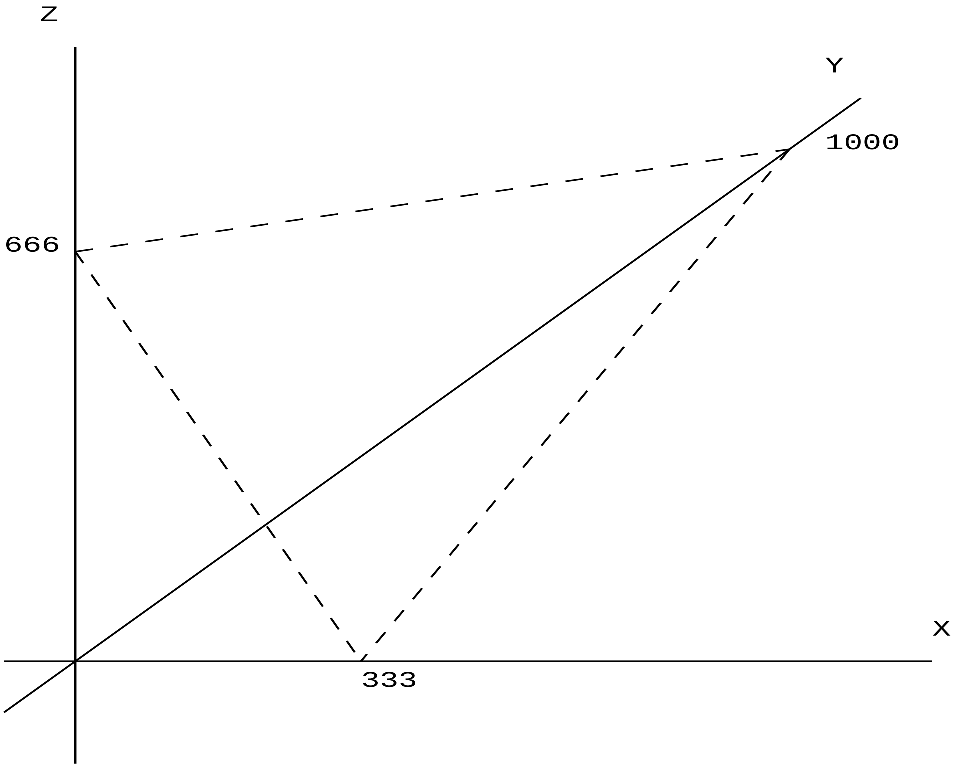

Figure 2b: Feasible region for problem 3

Once again the profit is

$$p(x, y, z) = 400x + 200y + 250z.$$

Once again any maximum of p on one of the boundary lines must occur at a corner. These are the points

$$E_1 : x = 333\frac{1}{3}, y = 0, z = 0,$$

$$E_2 : x = 0, y = 1000, z = 0,$$

$$E_3 : x = 0, y = 0, z = 666\frac{2}{3}.$$

Checking the profit values at these points, we find that the optimal choice is E_2.

2.1.4 Problem 4

As a final introductory problem we consider two modifications of problem 3. First, we add one equality contraint, the land constraint. Then we consider two profit functions. Thus we have two problems, to maximize

$$p_1(x, y, z) = 400x + 200y + 250z,$$

subject to

$$3x + y + 1.5z = 1000.$$

and

$$x + y + z = 625,$$

and then to maximize

$$p_2(x, y, z) = 300x + 200y + 250z,$$

subject to the same constraints.

These constraints may be rewritten as

$$4x + z = 750,$$

$$y = 625 - x - z.$$

The intersection of the solutions to these two equations are a line, given parameterically by

$$t \to (t, 3t - 125, 750 - 4t).$$

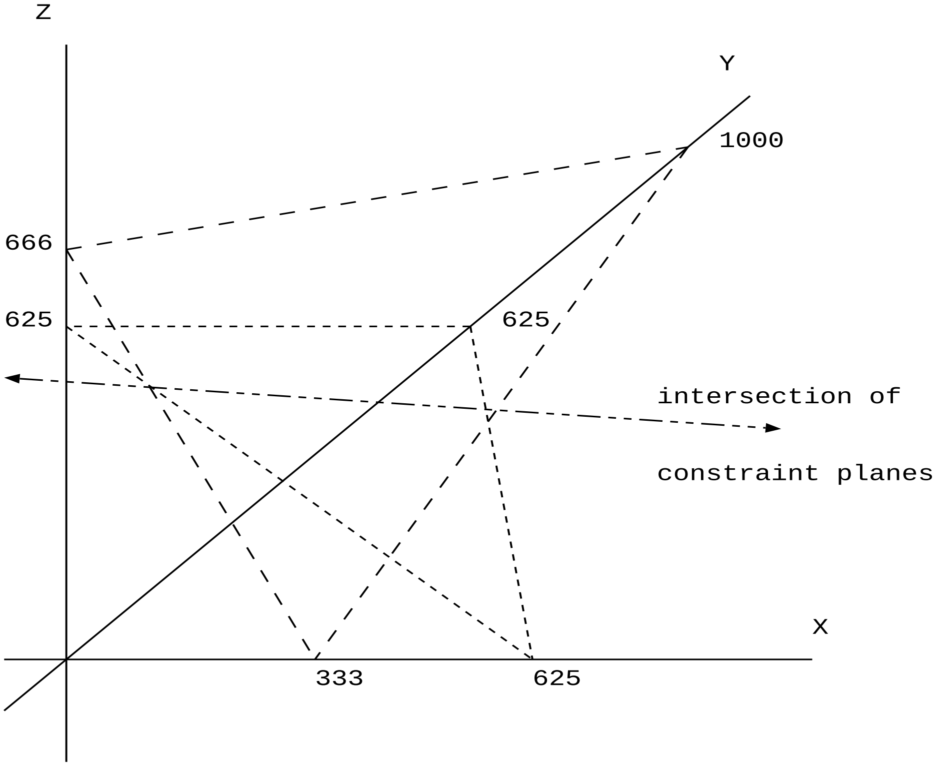

Figure 2c: Feasible region for problem 4

Along the line of points satisfying both constraints we find

$$p_1(t) = 400t + 200(3t - 125) + 250(750 - 4t) = 187,500.$$

The profit is a constant along the line. This situation is unusual, but not impossible, as this example shows. More typically, the second profit has the form

$$p_2(t) = 300t + 200(3t - 125) + 250(750 - 4t) = 187,500 - 100t.$$

Since the variables x, y, z must all be nonnegative, we find

$$t \geq 0, \quad 3t - 125 \geq 0, \quad 750 - 4t \geq 0.$$

That is,

$$41\frac{2}{3} \le t \le 187\frac{1}{2}.$$

The maximum for p_2 is achieved when $t = 41\frac{2}{3}$, that is when

$$x = 41\frac{2}{3}, \quad y = 0, \quad z = 583\frac{1}{3}.$$

2.2 First steps

Many practical optimization problems are examples of linear programming. The book Linear Programming, by H. Karloff (Birkhäuser 1991) looks like a good reference on the subject. It appears to be concise, well written, and up to date. In addition to covering the commonly used simplex algorithm, this book also treats the more recent, and quite different algorithms by Khachiyan (a theoretical breakthrough) and Karmarkar (a practical competitor to simplex).

The typical linear programming problem asks for the maximization of a linear (affine) function subject to a system of linear inequalities

$$\text{maximize} \quad g(x_1, \ldots, x_N) = \sum_{n=1}^{N} c_n x_n \tag{2.1}$$

$$\text{with constraints} \quad \sum_{n=1}^{N} \alpha_{ln} x_n \leq \beta_l, \quad l = 1, \ldots, L,$$

$$x_n \geq 0, \quad n = 1, \ldots, N.$$

The constraints define the *feasible region* for the problem. This problem will be called an *Inequality Linear Program* or ILP.

Notice that $\nabla g = (c_1, \ldots, c_N)$, which is not zero in any interesting case. Thus we never see the maximum of g in the interior of the feasible region.

There are a variety of essentially equivalent forms for these problems. It is useful to observe that the problem as stated can be recast in an alternative form, the *equality linear program* (ELP), where the first set of constraints has the form

$$\sum_{n=1}^{N} a_{ln} x_n = b_l, \quad l = 1, \ldots, L.$$

Suppose we start with a problem with inequality constraints, as originally posed. Rewrite the constraints with L new *slack variables*

$$\sum_{n=1}^{N} \alpha_{ln} x_n + x_{N+l} = \beta_l, \quad l = 1, \ldots, L,$$

where the slack variables $x_{N+1}, \ldots, x_{N+L}$ are also nonnegative.

Proposition 2.1. *Every linear program in ILP form can be rewritten as an equivalent problem in ELP form.*

Proof. It is easy to check that if we maximize g subject to these equality constraints, then the first N coordinates will be feasible, and must also be maximizing for the inequality constraints. Conversely, if we maximize g subject to the original inequality constraints, then we can choose values for the slack variables by taking

$$x_{N+l} = \beta_l - \sum_{n=1}^{N} \alpha_{ln} x_n, \quad l = 1, \dots, L,$$

and the vectors $(x_1, \dots, x_{N+L})$ will optimize the ELP. □

For most of the remaining discussion, the linear programs are considered in the (ELP) form: maximize

$$g(x_1, \dots, x_N) = \sum_{n=1}^{N} c_n x_n \tag{2.2}$$

subject to the constraints

$$\sum_{n=1}^{N} a_{ln} x_n = b_l, \quad l = 1, \dots, L,$$

$$x_n \geq 0, \quad n = 1, \dots, N.$$

Notice that the matrix A and the other variables do not have to come from an inequality linear program, despite the use of the same notation. For notational convenience we define rows $A_1, \dots, A_L$ for our constraint matrix $A = (a_{ln})$. To avoid problems do to redundancy in the constraints, we assume that $rank(A) = L$, that is the vectors $A_1, \dots, A_L$ are linearly independent. In order to have any feasible points there can't be too many independent constraints, so we must have

$$N \geq L.$$

Define

$$B = \begin{pmatrix} b_1 \\ \vdots \\ b_L \end{pmatrix}.$$

Before turning to the important geometric aspects of linear programming, consider the system

$$AY = B.$$

Since A has L independent rows, the system has a solution for every B. Moreover the null space of A, which is the set of solutions of the homogeneous system

$$AY = 0, \quad Y = \begin{pmatrix} y_1 \\ \vdots \\ y_N \end{pmatrix}.$$

is a subspace of dimension $N - L$. Suppose X_p is any solution of $AX = B$. The general solution of $AX = B$ always has the form $X = X_p + Y$, where Y is any solution of $AY = 0$. A set of the form $X_p + \mathcal{V}$, where $\mathcal{V}$ is a vector subspace (of $\mathbb{R}^N$), is called an affine subspace. Thus we have an initial result,

Proposition 2.2. *The feasible region for a linear program in ELP form is the intersection of the affine subspace of solutions to* $AX = B$, *which has dimension* $N - L$, *with the positive sector* $x_n \geq 0$ *in* $\mathbb{R}^N$.

Let's take a look at linear programs from the point of view of Lagrange multipliers. The constraints $AX = B$ are a system of L equations which can taken to be the simultaneous equality constraints

$$h_1(X) = a_{11}x_1 + \cdots + a_{1N}x_N = b_1, \quad \ldots, \quad h_L(X) = a_{L1}x_1 + \cdots + a_{LN}x_N = b_L.$$

Lagrange multiplier theory says that at local maximum of a function g subject to the equality constraints, we must have

$$\nabla g = \lambda_1 \nabla h_1 + \cdots + \lambda L \nabla h_L. \tag{2.3}$$

In a linear programming problem, the objective function g is affine and the constraints h_l are linear, so g and the constraint functions h_l have constant gradients. Thus the condition (2.3) on gradients either holds nowhere in the affine subspace, or everywhere in the subspace. This suggests the following result, although a careful proof will use a different method.

Proposition 2.3. *Suppose* g *is an affine function defined on an affine subspace* $\mathcal{W} \subset \mathbb{R}^N$. *Then* g *has a local maximum in* $\mathcal{W}$ *if and only if* g *is constant on* $\mathcal{W}$.

Proof. Suppose $\mathcal{V}$ is a subspace of $\mathbb{R}^N$, X_p is a vector in $\mathbb{R}^N$, and $\mathcal{W}$ is the affine subspace $\mathcal{W} = X_p + \mathcal{V}$. Suppose g has a local maximum at $Y = X_p + V_1 \in \mathcal{W}$, and there is another vector $Z = X_p + V_2 \in \mathcal{W}$ with $g(Z) < g(Y)$. Here $V_1, V_2 \in \mathcal{V}$

Notice that for $t \in \mathbb{R}$ the line

$$(1-t)Y + tZ = (1-t)(X_p + V_1) + t(X_p + V_2) = X_p + (1-t)V_1 + tV_2, \quad t \in \mathbb{R},$$

lies in $\mathcal{W}$. Define the function

$$\widetilde{g}(t) = g((1-t)Y + tZ),$$

which is just the restriction of $g : \mathcal{W} \to \mathbb{R}$ to the line joining Y and Z. Each coordinate x_k of $(1-t)Y + tZ$ is first order in t,

$$x_k(t) = (1-t)y_k + tz_k,$$

so $\widetilde{g}(t)$ is affine as a function of t, that is a polynomial of degree at most 1 in t. Also, the function $\widetilde{g}$ has a local maximum at $t = 0$, where $\widetilde{g}(0) = g(Y)$.

Now we obtain a contradiction, since an affine function of t with a local maximum must be constant (use the derivative test), but we also have $\widetilde{g}(1) = g(Z) < \widetilde{g}(0) = g(Y)$. Thus Z can't exist, and g is constant on $\mathcal{W}$.

□