FORK1005
Preparatory Course in Mathematics 2015/16
Lecture 6: Constrained Optimization & Lagrange Multipliers

August 17, 2015

1 Introduction to Constrained Optimization

In going from *unconstrained optimization* problems to *constrained optimization* problems, we go from the problem

$$\text{Maximize:} \quad f(x,y)$$

to the problem

$$\begin{aligned} \text{Maximize:} &\quad f(x,y) \\ \text{subject to:} &\quad g(x,y) = c \end{aligned}$$

So far, we have looked at *unconstrained optimization* problems, where we are asked to maximize a function of one or several variables, without any constraints on the variables.

Example 1.1. For example, in Lecture 5, we wanted to maximize a company's profit function

$$P(x,y) = -x^2 - 2y^2 + 30x + 15y - 50,$$

where x was millions of USD spent on research, and y millions of USD spent on advertising. By calculating partial derivatives, we found the maximization point $(x,y) = (15, 3.75)$.

In this problem, there were no constraints on the variables x and y. Any point (x,y) that maximizes the function P is a solution. But what if the company could only spend a total of 10 million USD on research and advertising? Then we have the constraint

$$x + y = 10.$$

With this constraint, our unconstrained maximizer $(15, 3.75)$ is no longer a solution, so we need to find a new maximization point. However, we can't just take partial derivatives of the profit function P as before, since $x+y=10$ is an *external* constraint and is not captured by the dynamics or slopes of P.

The maximization problem

$$\begin{aligned} \text{Maximize:} \quad & P(x,y) = -x^2 - 2y^2 + 30x + 15y - 50, \\ \text{subject to:} \quad & x + y = 10. \end{aligned}$$

is an example of a *constrained optimization* problem. Problems of this nature are of huge relevance in economics: Consumers optimize their consumption decisions, constrained by their budget; Companies that want to maximize their profits, might have production constraints, due to limitations on their labour force.

So the main question of this lecture is, how can we incorporate external constraints into an optimization problem, so that all solutions reflect this constraint? We will take a look at the more intuitive, but impractical, approach, before we do the *Lagrange Multiplier method.*

2 Direct Substitution

A consumer has decided to spend all her available income on two goods, X and Y. The price of good X is p_x, and the price of good Y is p_y. x and y denotes how much is purchased of goods X and Y respectively. The consumer's available income is m. So we have the *budget constraint*

$$xp_x + yp_y = m.$$

The consumer's ranking of a purchase is described by the *utility function* $U(x,y)$. For two different purchasing options (x_1, y_1) and (x_2, y_2), the inequality

$$U(x_1, y_1) > U(x_2, y_2)$$

means exactly that the consumer prefers the purchase (x_1, y_1) to the purchase (x_2, y_2). So when deciding how much to buy of good X and of Y, the consumer wants to maximize $U(x,y)$, and so she faces the constrained optimization problem

$$\begin{aligned} \text{Maximize:} \quad & U(x,y), \\ \text{subject to:} \quad & xp_x + yp_y = m. \end{aligned}$$

One simple way of solving such a constrained optimization problem is to use the constraint to write y in terms of x:

$$xp_x + yp_y = m$$

becomes

$$y = \frac{m - xp_x}{p_y}.$$

Then we can plug this into the function we want to maximize, $U(x,y)$ to get the unconstrained maximization problem in one variable:

$$\text{Maximize:} \quad U\left(x, \frac{m - xp_x}{p_y}\right)$$

Example 2.1. Suppose the consumer has 200 kroner to spend on apples and oranges. The apples cost $p_x = 1$ and the oranges cost $p_y = 2$. Let x denote the number of apples purchased, and y the number of oranges purchased. The utility function is

$$U(x, y) = xy,$$

so we need to solve

$$\begin{aligned} \text{Maximize:} \quad & U(x, y) = xy, \\ \text{subject to:} \quad & x + 2y = 200. \end{aligned}$$

We rewrite the constraint to get

$$y = \frac{200 - x}{2} = 100 - \frac{x}{2}.$$

We then plug this into the utility function:

$$U\left(x, 100 - \frac{x}{2}\right) = x\left(100 - \frac{x}{2}\right) = 100x - \frac{x^2}{2}.$$

What we have done is to derive a new function in only one variable,

$$f(x) = 100x - \frac{x^2}{2}$$

which describes the utility function $U(x, y)$ for all values of x that satisfy the budget constraint $x + 2y = 200$. This function f can easily be maximized with standard optimization techniques:

1. We differentiate f:
$$f'(x) = 100 - x$$
2. We set $f'(x) = 0$ to find stationary points:
$$100 - x = 0 \quad \Longrightarrow \quad x = 100.$$
3. We find the second-order derivative to check that $x_* = 100$ maximizes the function:
$$f''(x) = -1 \quad \Longrightarrow \quad f''(100) = -1 < 0$$
so by the second derivative test, $x_* = 100$ maximizes the function.

Therefore, the consumer optimizes her utility when she purchases $x = 100$ apples. Since

$$y = 100 - \frac{x}{2},$$

we conclude that $y = 100 - 100/2 = 50$. **So the optimal purchase is 100 apples and 50 oranges.**

While this method is simple and intuitive, it is not good for more complicated constraint functions:

Example 2.2. Consider the constrained optimization problem

$$\begin{aligned} \text{Maximize:} \quad & f(x,y) = x^2 + y^2, \\ \text{subject to:} \quad & x^2 + xy + y^2 = 12. \end{aligned}$$

Now it is more difficult to use the constraint function $x^2 + xy + y^2 = 12$ to write y in terms of x. It is in fact a quadratic expression, with the solutions

$$y = \frac{-x + \sqrt{x^2 - 4(x^2 - 12)}}{2} \quad \text{and} \quad y = \frac{-x - \sqrt{x^2 - 4(x^2 - 12)}}{2},$$

but trying to plug these solutions into $f(x,y)$ to maximize it would probably be a very tiresome process.

Other times, it might be impossible to separate the variables x and y in the constraint function, so that you cannot write y in terms of x. The Lagrange multiplier method, on the other hand, overcomes a lot of these difficulties.

3 The Lagrange Multiplier Method

Example 3.1. Consider again the consumer maximization problem

$$\begin{aligned} \text{Maximize:} \quad & U(x,y) = xy, \\ \text{subject to:} \quad & x + 2y = 200. \end{aligned}$$

This problem was straightforward to solve by substitution but now we will use a different method called the *Lagrange multiplier method*.

1. We define the **Lagrangian function**

$$\mathcal{L}(x,y,\lambda) = U(x,y) + \lambda(x + 2y - 200) = xy + \lambda(x + 2y - 200).$$

   The new variable λ is called the **Lagrangian multiplier**.

2. We differentiate $\mathcal{L}$ with respect to x, y and λ to get the three partial derivatives

$$\begin{aligned} \mathcal{L}'_x(x,y,\lambda) &= y + \lambda, \\ \mathcal{L}'_y(x,y,\lambda) &= x + 2\lambda, \\ \mathcal{L}'_\lambda(x,y,\lambda) &= x + 2y - 200. \end{aligned}$$

3. We set each partial derivative equal to zero, and solve for x, y and λ:

$$\begin{aligned} y + \lambda = 0 \quad &\Longrightarrow \quad y = -\lambda \\ x + 2\lambda = 0 \quad &\Longrightarrow \quad x = -2\lambda \\ x + 2y - 200 = 0 \quad &\Longrightarrow \quad x + 2y = 200. \end{aligned}$$

4. This gives us a system of 3 equations, and we can plug the first two equations into the third to get

$$\begin{aligned} -2\lambda + 2(-\lambda) &= 200 \\ -4\lambda &= 200 \\ \lambda &= -50. \end{aligned}$$

   This in return gives us the solution

$$x = 100, \quad y = 50.$$

   Note that this is the same solution that we got when using substitution.

The Lagrange Multiplier Method

We want to maximize $f(x, y)$ given the constraint $g(x, y) = c$.

1. Define the Lagrangian function

$$\mathcal{L}(x, y, \lambda) = f(x, y) + \lambda \left(g(x, y) - c \right).$$

2. Compute partial derivatives $\mathcal{L}'_x$, $\mathcal{L}'_y$ and $\mathcal{L}'_\lambda$ and set them equal to zero (the equation $\mathcal{L}'_\lambda(x, y, \lambda) = 0$ is always equivalent to the constraint $g(x, y) = c$).

3. This gives us the (not necessarily linear) system of equations

$$\begin{aligned} f'_x(x, y) + \lambda g'_x(x, y) &= 0 \\ f'_y(x, y) + \lambda g'_y(x, y) &= 0 \\ g(x, y) &= c \end{aligned}$$

4. Solve these for x, y and λ.

Note: Any solution to the constrained optimization problem will also solve the above system of equations, but not all solutions to the system of equations is a solution to the constrained optimization problem. It is therefore important to check that your answer makes sense: Make sure the solution is a maximizer, not a minimizer, and if you get more than one solution, you pick the one among those that maximize $f(x, y)$.

The following theorem is a formal statement of how the Lagrange multiplier works:

Theorem 3.2. *Suppose that the maximization problem* 'maximize $f(x, y)$ when $g(x, y) = c$', *has a solution* (x^*, y^*). *Then there exists a Lagrange multiplier* λ^* *such that* (x^*, y^*, λ^*) *solves* *the system of equations*

$$\begin{aligned} f'_x(x,y) + \lambda g'_x(x,y) &= 0 \\ f'_y(x,y) + \lambda g'_y(x,y) &= 0 \\ g(x,y) &= c \end{aligned}$$

So if any solution to the constrained problem exists, then it can be located as a solution to the above system of equations.

Example 3.3. Now we'll try the problem

$$\begin{aligned} &\text{Maximize:} & f(x,y) &= x^2 + y^2, \\ &\text{subject to:} & x^2 + xy + y^2 &= 12. \end{aligned}$$

1. We define the Lagrangian function

$$\mathcal{L}(x,y,\lambda) = x^2 + y^2 + \lambda(x^2 + xy + y^2 - 12).$$

2. We differentiate $\mathcal{L}$ with respect to x, y and λ to get the three partial derivatives

$$\begin{aligned} \mathcal{L}'_x(x,y,\lambda) &= 2x + 2x\lambda + y\lambda, \\ \mathcal{L}'_y(x,y,\lambda) &= 2y + 2y\lambda + x\lambda, \\ \mathcal{L}'_\lambda(x,y,\lambda) &= x^2 + xy + y^2 - 12. \end{aligned}$$

3. We set each partial derivative equal to zero, and solve for x, y and λ:

$$\begin{aligned} 2x + 2x\lambda + y\lambda = 0 \quad &\Longrightarrow & \lambda &= \frac{-2x}{2x+y} \\ 2y + 2y\lambda + x\lambda = 0 \quad &\Longrightarrow & \lambda &= \frac{-2y}{x+2y} \\ x^2 + xy + y^2 - 200 = 0 \quad &\Longrightarrow & x^2 + xy + y^2 &= 12. \end{aligned}$$

4. We combine the first two equations to get

$$\begin{aligned} \frac{-2x}{2x+y} &= \frac{-2y}{x+2y} \\ -2x(x+2y) &= -2y(2x+y) \\ x^2 + 2xy &= y^2 + 2xy \\ x^2 &= y^2 \\ x &= \pm y \end{aligned}$$

5. So we have two possibilities: $y = x$ and $y = -x$.

- We try the first case: If $y = x$, the constraint $x^2 + xy + y^2 = 12$ becomes

$$x^2 + x^2 + x^2 = 12 \quad \implies \quad x^2 = 4 \quad \implies \quad x = \pm 2.$$

Note that $f(-x, -y) = f(x, y)$ so we don't need to distinguish between the answers $(2, 2)$ and $(-2, -2)$. So this gives us the solution $(2, 2)$. We have

$$f(2, 2) = 4 + 4 = 8.$$

- Then we try the second case: If $y = -x$, the constraint $x^2 + xy + y^2 = 12$ becomes

$$x^2 - x^2 + x^2 = 12 \quad \implies \quad x^2 = 12 \quad \implies \quad x = \pm\sqrt{12}.$$

So this gives us the solution $(\sqrt{12}, -\sqrt{12})$. We have

$$f(\sqrt{12}, -\sqrt{12}) = 12 + 12 = 24.$$

6. We conclude that $(\sqrt{12}, -\sqrt{12})$ (and $(-\sqrt{12}, \sqrt{12})$) solves our constrained maximization problem.

3.1 Things to Watch Out For

Sometimes the set of points (x, y) you want to maximize over have more restrictions than just constraint $g(x, y) = c$. For example, we often assume that x and y are non-negative ($x \geq 0$, $y \geq 0$). Then, however, in the case where the constrained optimization problem has a solution located on the boundary of these restrictions (i.e. $x = 0$ or $y = 0$, then the Lagrange method might not give you these solutions.

Example 3.4. To demonstrate this, we consider the maximization problem of maximizing $f(x, y) = x^2 + y^2$ subject to the constraint $x + y = 10$. Let us first allow for negative values of x and y, so for example $(-1, 11)$ is an admissible point. Then it is clear that there is no upper bound for $f(x, y)$ under this constraint, because no matter how large x is, you can set $y = 10 - x$, and you get

$$f(x, 10 - x) = x^2 + (10 - x)^2 = 2x^2 + 100 - 20x$$

which can be made arbitrarily large. So even under constraints, there is no solution to the maximization problem.

Now let's add the restrictions $x \geq 0$ and $y \geq 0$. These make sense for example if x and y represent amount of goods to purchase, or something physical that cannot take negative values. The restriction $x + y = 10$ then means that x and y cannot be larger than 10, so $f(x, y) = x^2 + y^2$ is bounded. However, if we apply the Lagrange method, we get the system of equations

$$\begin{aligned} 2x &= -\lambda \\ 2y &= -\lambda \\ x + y &= 10 \end{aligned}$$

which has the unique solution

$$x = y = 5.$$

However, $(5, 5)$ is not a maximizer, but a minimizer. If we plug it into f, we get

$$f(5, 5) = 25 + 25 = 50.$$

On the other hand, $(10, 0)$ also satisfies the constraint, and if we plug it in, we get

$$f(10, 0) = 100.$$

The maximizers of this constrained optimization problem are in fact $(10, 0)$ and $(0, 10)$. The Lagrange method did not find these points because they were on the boundary of the constraints $x \geq 0$ and $y \geq 0$ respectively.

Remark 3.5. Other scenarios where the Lagrange multiplier method doesn't work:

- If $g'_x(x^*, y^*) = g'_y(x^*, y^*) = 0$, it might not work.
- Both g'_x and g'_y have to be continuous and well-defined around the solution.

Example 3.6. Suppose you want to maximize

$$f(x, y) = x + 4y,$$

subject to the constraint

$$\sqrt{x} + \sqrt{y} = 5.$$

We set up the Lagrangian function

$$\mathcal{L}(x, y, \lambda) = x + 4y + \lambda(\sqrt{x} + \sqrt{y} - 5)$$

We differentiate to get the system of equations

$$\begin{aligned}
\mathcal{L}'_x(x, y, \lambda) &= 1 + \frac{\lambda}{2\sqrt{x}} = 0 \\
\mathcal{L}'_y(x, y, \lambda) &= 4 + \frac{\lambda}{2\sqrt{y}} = 0 \\
\mathcal{L}'_\lambda(x, y, \lambda) &= \sqrt{x} + \sqrt{y} - 5.
\end{aligned}$$

Rearranging the first two equations and combining them, we get

$$-\lambda = 2\sqrt{x} = 8\sqrt{y}$$

which means that

$$x = 4^2 y = 16y.$$

It also means that

$$\sqrt{x} = 4\sqrt{y},$$

so we can plug this into the constraint equation to get

$$4\sqrt{y} + \sqrt{y} = 5 \qquad \Longrightarrow \qquad \sqrt{y} = 1 \qquad \Longrightarrow \qquad y = 1.$$

So we have the solution

$$x = 16, \qquad y = 1.$$

Plugging this into f gives us

$$f(16, 1) = 16 + 4 \cdot 1 = 20.$$

But $(x, y) = (0, 25)$ also solves the constraint, and plugging it into f gives us

$$f(0, 25) = 4 \cdot 25 = 100.$$

So $(16, 1)$ is not a solution to the maximizing problem. The real solution $(0, 25)$ is not found by the Lagrange method, because the function $g'_x(x, y) = \dfrac{1}{2\sqrt{x}}$ is not defined at $x = 0$.

