Lagrangian method and saddle point

Toshihiko Mukoyama

1 General theory

Consider the following constrained optimization problem:

$$\max_{x_1,x_2} \quad f(x_1,x_2) \qquad \text{subject to } g(x_1,x_2)=0.$$

Let the set of combinations of (x_1,x_2) that satisfies $g(x_1,x_2)=0$ be Ω. Then the problem can also be written as:

$$\max_{x_1,x_2} \quad f(x_1,x_2) \qquad \text{subject to } (x_1,x_2)\in\Omega.$$

Denote the solution for this constrained problem (that is, what we are going to look for) as (x_1^*,x_2^*).

Let us consider the following unconstrained-version of the problem:

$$\max_{x_1,x_2} \quad f(x_1,x_2),$$

and suppose that at the unconstrained optimum $(\tilde{x}_1,\tilde{x}_2)$, the constraint is not satisfied: $g(\tilde{x}_1,\tilde{x}_2)\neq 0$. Instead, without loss of generality, suppose that

$$g(\tilde{x}_1,\tilde{x}_2)<0 \tag{1}$$

holds. (If not, use $-g(x_1,x_2)=0$ as the constraint.)

Define a Lagrangian

$$L(x_1,x_2,\lambda)\equiv f(x_1,x_2)+\lambda g(x_1,x_2). \tag{2}$$

Note that $L(x_1,x_2,\lambda)=f(x_1,x_2)$ if $(x_1,x_2)\in\Omega$, because $g(x_1,x_2)=0$ holds when $(x_1,x_2)\in\Omega$.

Now, for a given λ, consider the following unconstrained Lagrangian problem:

$$\max_{x_1,x_2} \quad L(x_1,x_2,\lambda). \tag{3}$$

Call the maximum value of $L(x_1, x_2, \lambda)$ as $\mathcal{L}(\lambda)$. And call the solution of the problem (3) as $(\hat{x}_1(\lambda), \hat{x}_2(\lambda))$ for each λ. Now consider a constrained-version of the same problem:

$$\max_{x_1, x_2} \quad L(x_1, x_2, \lambda) \qquad \text{subject to } (x_1, x_2) \in \Omega. \tag{4}$$

Of course, the solution to the constrained Lagrangian problem (4) is the same as the original constrained problem (x_1^*, x_2^*) for any value of λ, because $L(x_1, x_2, \lambda) = f(x_1, x_2)$ when $(x_1, x_2) \in \Omega$. Also note that

$$L(\hat{x}_1(\lambda), \hat{x}_2(\lambda), \lambda) \geq L(x_1^*, x_2^*, \lambda) \text{ for all } \lambda$$

holds, because adding a constraint won't make things better (i.e. in the problem (3), one can choose from a wider range of (x_1, x_2) than in the problem (4)). The left-hand side is equal to $\mathcal{L}(\lambda)$ and the right-hand side is equal to $f(x_1^*, x_2^*)$, thus

$$\mathcal{L}(\lambda) \geq f(x_1^*, x_2^*) \text{ for all } \lambda. \tag{5}$$

The important insight here is that (x_1^*, x_2^*) always gives the same value of $L(x_1, x_2, \lambda)$ for *any* λ, and thus the optimal value of the original problem, $f(x_1^*, x_2^*)$, is *always* attainable as a value of the Lagrangian no matter what the value of λ is.

Recall that, from Assumption (1), $g(\hat{x}_1(0), \hat{x}_2(0)) < 0$ holds. Now let us gradually increase λ and consider what happens. As λ increases, the weight on $g(x_1, x_2)$ function in (2) keeps increasing, which means that the solution to the problem (3) would at some point want to choose a combination of (x_1, x_2) that makes $g(x_1, x_2)$ positive rather than negative. (As λ approaches to infinity, we should ignore $f(x_1, x_2)$ and only care about $g(x_1, x_2)$ in the problem (3).) Then, there exists a sufficiently large λ, $\bar{\lambda}$, that makes $g(\hat{x}_1(\lambda), \hat{x}_2(\lambda)) > 0$ for all $\lambda \geq \bar{\lambda}$.

Suppose that the functions f and g are such that $(\hat{x}_1(\lambda), \hat{x}_2(\lambda))$ moves smoothly with λ (see the Theorem of Maximum). Then, as we move λ from 0 to $\bar{\lambda}$, we can find a value of λ, λ^*, that makes $g(\hat{x}_1(\lambda^*), \hat{x}_2(\lambda^*)) = 0$.

From the definition of constraint set Ω, the fact that $g(\hat{x}_1(\lambda^*), \hat{x}_2(\lambda^*)) = 0$ means that $(\hat{x}_1(\lambda^*), \hat{x}_2(\lambda^*)) \in \Omega$. This in turn means that $(\hat{x}_1(\lambda^*), \hat{x}_2(\lambda^*))$ is feasible for the constrained Lagrangian problem (4), and therefore the solution of the problem (4). (Otherwise, it cannot be the solution for the *unconstrained* problem (3) when $\lambda = \lambda^*$.) This, in turn, implies that we found the solution for the original constrained problem: $(\hat{x}_1(\lambda^*), \hat{x}_2(\lambda^*)) = (x_1^*, x_2^*)$.

Thus, we found that if we want to look for (x_1^*, x_2^*), we can instead look for λ^* first and then calculate $(\hat{x}_1(\lambda^*), \hat{x}_2(\lambda^*))$ by solving the unconstrained problem (3). How can we find λ^*? First, note that

$$\mathcal{L}(\lambda^*) = L(\hat{x}_1(\lambda^*), \hat{x}_2(\lambda^*), \lambda^*) = f(\hat{x}_1(\lambda^*), \hat{x}_2(\lambda^*)) = f(x_1^*, x_2^*).$$

The first equality is the definition of $\mathcal{L}$ function, the second is because $g(\hat{x}_1(\lambda^*), \hat{x}_2(\lambda^*)) = 0$, and the third is from the previous paragraph. From the relationship (5), this means that

$$\mathcal{L}(\lambda) \geq \mathcal{L}(\lambda^*) \text{ for all } \lambda.$$

In other words, λ^* can be found by minimizing the $\mathcal{L}(\lambda)$ function with respect to λ. From the definition of $\mathcal{L}(\lambda)$, we can find λ^* by solving

$$\min_{\lambda} \max_{x_1, x_2} \quad L(x_1, x_2, \lambda),$$

that is, we look for the *saddle point* of the L function.

When f and g are differentiable, first-order conditions for this problem are

$$\frac{\partial}{\partial x_1} L(x_1, x_2, \lambda) = 0,$$

$$\frac{\partial}{\partial x_2} L(x_1, x_2, \lambda) = 0,$$

and

$$\frac{\partial}{\partial \lambda} L(x_1, x_2, \lambda) = 0,$$

where the last one is the same as imposing $g(x_1, x_2) = 0$. This is the regular procedure of the textbook Lagrangian method.

Note: See Fryer and Greenman (1987) for a graphical exposition.

2 Interpretation with an example

Suppose that a consumer is purchasing two goods, good 1 and good 2, at a grocery store. The prices of these goods are p_1 and p_2, and the quantities purchased are denoted as x_1 and x_2. The consumer's income is I. Thus the consumer's budget constraint is

$$p_1 x_1 + p_2 x_2 = I. \tag{6}$$

Suppose that the consumer can purchase whatever amount in the grocery store—above or below the budget; the cash register at the supermarket is very lazy and doesn't check exactly how many goods the consumer bought, and the consumer just pays I. The consumer has the utility function $u(x_1, x_2)$ with usual properties, so if there is no penalty, she would like to go above budget.

Now, suppose that at the exit of the grocery store, after paying I at the cash register (who is too lazy to check how many goods the consumer has in the bag), the consumer has to meet with a gatekeeper. The gatekeeper asks how many of each goods the consumer has purchased (and the consumer has to answer honestly), and thus he can compare between I and $p_1x_1 + p_2x_2$. The gatekeeper would want the consumer to keep the budget constraint (6). The gatekeeper, however, has only one method to achieve this. He can collect a fee of

$$F(x_1, x_2, \lambda) \equiv -\lambda(I - p_1x_1 - p_2x_2)$$

utils (that is, the fee is measured in terms of the consumer's utility) from the consumer, where $\lambda \geq 0$. When the consumer goes over the budget ($I - p_1x_1 - p_2x_2 < 0$), the fee $F(x_1, x_2, \lambda)$ is positive. Thus, the gatekeeper can punish the consumer when the consumer goes over the budget. If $I - p_1x_1 - p_2x_2 > 0$, the gatekeeper has to pay the fee to the consumer because $F(x_1, x_2, \lambda)$ becomes a negative value. Under this assumption, the consumer's utility upon exit is

$$L(x_1, x_2, \lambda) \equiv u(x_1, x_2) - F(x_1, x_2, \lambda) = u(x_1, x_2) + \lambda(I - p_1x_1 - p_2x_2). \tag{7}$$

Suppose that the gatekeeper wants the consumer (who chooses x_1 and x_2 freely) to voluntarily satisfy the budget constraint (6). The gatekeeper wants to achieve this by choosing the level of fee λ appropriately. As is mentioned above, if λ is very close to zero, the consumer would prefer to go above budget: $I - p_1x_1 - p_2x_2 < 0$. If λ is very large, the consumer would choose to collect the fee instead by going under the budget: $I - p_1x_1 - p_2x_2 > 0$. Thus somewhere in between, there should be an appropriate level of λ that makes the budget constraint (6) hold with equality. Call this λ as λ^*.

How can the gatekeeper find such a λ? What would be the appropriate level of punishment, so that the consumer won't go over-budget or under-budget? From the consumer's perspective, by satisfying the budget constraint (6) with equality and choosing x_1 and x_2 to maximize $u(x_1, x_2)$ (call such a combination of x_1 and x_2 as x_1^* and x_2^*), she can always achieve the value of $L(x_1^*, x_2^*, \lambda) = u(x_1^*, x_2^*)$ no matter what λ

is. This fact means that if there is a λ that makes the consumer to *voluntarily* choose x_1 and x_2 so that the budget constraint (6) is satisfied (we called it λ^*), λ^* is the one that minimizes (across different λ) what the consumer can achieve by trying her best in maximizing $L(x_1, x_2, \lambda)$ by choosing x_1 and x_2. This result follows because another value of λ can potentially give a better value of $L(x_1, x_2, \lambda)$ to the consumer while λ^* can provide at most (the always-achievable) $u(x_1^*, x_2^*)$ to the consumer as a value of L.

Thus, the gatekeeper can induce the consumer to satisfy the budget constraint voluntarily by choosing a λ to minimize $\max_{x_1,x_2} L(x_1, x_2, \lambda)$. That is, when λ and (x_1, x_2) are the solution of

$$\min_{\lambda} \max_{x_1,x_2} \quad L(x_1, x_2, \lambda),$$

where the L function is defined as (7), the budget constraint is satisfied and $u(x_1, x_2)$ is maximized within the budget constraint. This is what we derived in the general theory above. And what is the λ^* here? The value of λ^* is just the amount of the punishment (in utility term) that the consumer is willing to take for going over budget for one dollar. In other words, it is the value of one dollar in utility term.

References

[1] Fryer, M. J. and J. V. Greenman (1987). *Optimisation Theory,* Edward Arnold.