Lagrangian Method and the Saddle Point

Toshihiko Mukoyama

May 2026

General Theory

The Constrained Problem

Consider a constrained optimization problem:

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

Let $\Omega \equiv \{(x_1,x_2) : g(x_1,x_2)=0\}$.

Equivalent maximization problem:

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

- Denote the solution by (x_1^*, x_2^*)
- Goal: find (x_1^*, x_2^*)

The Unconstrained Version

Consider first the unconstrained problem:

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

with optimum $(\tilde{x}_1, \tilde{x}_2)$.

Suppose the constraint is violated:

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

- This is without loss of generality; if $g(\tilde{x}_1, \tilde{x}_2) > 0$, redefine the constraint as $-g(x_1, x_2) = 0$
- The unconstrained optimum lies "outside" the feasible set
- Need a way to penalize violations of the constraint

Defining the Lagrangian

Define the Lagrangian:

$$L(x_1, x_2, \lambda) \equiv f(x_1, x_2) + \lambda\, g(x_1, x_2).$$

- Key observation: if $(x_1, x_2) \in \Omega$, then $g(x_1, x_2) = 0$, so

  $$L(x_1, x_2, \lambda) = f(x_1, x_2) \quad \text{for all } \lambda.$$

- For each λ, consider the **unconstrained Lagrangian problem**

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

  Denote its solution $(\hat{x}_1(\lambda), \hat{x}_2(\lambda))$ and value $\mathcal{L}(\lambda)$

An Important Inequality

Now consider the **constrained** Lagrangian problem:

$$\max_{x_1, x_2} L(x_1, x_2, \lambda) \quad \text{s.t.} \quad (x_1, x_2) \in \Omega.$$

- Solution to this constrained Lagrangian problem coincides with (x_1^*, x_2^*) for any λ (since $L = f$ on Ω)
- Adding a constraint cannot increase the maximum:

$$L(\hat{x}_1(\lambda), \hat{x}_2(\lambda), \lambda) \;\geq\; L(x_1^*, x_2^*, \lambda).$$

- Translating: $\mathcal{L}(\lambda) \geq f(x_1^*, x_2^*)$ for all λ

Key inequality:

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

Key Insight

- (x_1^*, x_2^*) delivers the same value of $L(x_1, x_2, \lambda)$ for **any** λ
- So $f(x_1^*, x_2^*)$ is always attainable as a value of the Lagrangian, regardless of λ
- Therefore $\mathcal{L}(\lambda)$ is bounded **below** by $f(x_1^*, x_2^*)$

Strategy from here:

- Vary λ to find one for which the unconstrained solution $(\hat{x}_1(\lambda), \hat{x}_2(\lambda))$ is feasible under the constraint
- At that λ, the unconstrained problem and the original constrained problem coincide

Sweeping λ

By the assumption that the unconstrained optimum violates the constraint, $g(\hat{x}_1(0), \hat{x}_2(0)) < 0$.

As λ increases, starting from 0:

- Weight on $g(x_1, x_2)$ in the Lagrangian increases
- At some point, the solver of the unconstrained Lagrangian problem prefers $g > 0$ rather than $g < 0$
- As $\lambda \to \infty$, f becomes negligible relative to λg, so the solver definitely wants $g > 0$

Hence $\exists\, \bar{\lambda}$ such that

$$g(\hat{x}_1(\lambda), \hat{x}_2(\lambda)) > 0 \quad \text{for all } \lambda \geq \bar{\lambda}.$$

By continuity (Theorem of the Maximum), there exists λ^* with

$$g(\hat{x}_1(\lambda^*), \hat{x}_2(\lambda^*)) = 0.$$

Why λ^* Solves the Original Problem

At $\lambda = \lambda^*$:

- $g(\hat{x}_1(\lambda^*), \hat{x}_2(\lambda^*)) = 0$, so $(\hat{x}_1(\lambda^*), \hat{x}_2(\lambda^*)) \in \Omega$
- The pair solves unconstrained Lagrangian problem
- The pair is feasible under the constraint, and thus solves the constrained Lagrangian problem
- Since the constrained Lagrangian problem has the same solution as the original problem:

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

Possible recipe: find λ^* first, then solve the unconstrained Lagrangian problem at λ^*. How do we find λ^*?

The Saddle Point

At λ^*:

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

Because $\mathcal{L}(\lambda) \geq \mathcal{L}(\lambda^*)$ for all λ, λ^* **minimizes** $\mathcal{L}(\lambda)$:

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

- This is the **saddle point** of the Lagrangian
- Inner $\max$ over (x_1, x_2), outer $\min$ over λ
- Solution simultaneously gives (x_1^*, x_2^*) and λ^*

Visualizing $\mathcal{L}(\lambda)$

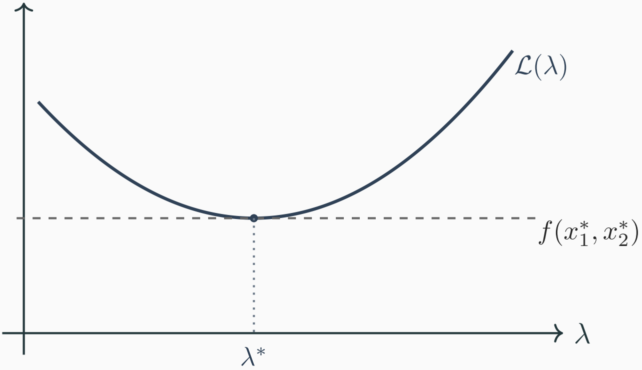

- $\mathcal{L}(\lambda) \geq f(x_1^*, x_2^*)$ for all λ (the dashed floor)
- $\mathcal{L}(\lambda^*) = f(x_1^*, x_2^*)$ — the floor is attained at λ^*
- Hence $\lambda^* = \arg\min_\lambda \mathcal{L}(\lambda)$

First-Order Conditions

When f and g are differentiable, the saddle point satisfies:

$$\begin{aligned}\frac{\partial}{\partial x_1}L(x_1,x_2,\lambda)&=0,\\ \frac{\partial}{\partial x_2}L(x_1,x_2,\lambda)&=0,\\ \frac{\partial}{\partial \lambda}L(x_1,x_2,\lambda)&=0.\end{aligned}$$

- The third condition is just $g(x_1, x_2) = 0$, that is, the original constraint
- These are the standard textbook Lagrangian FOCs
- See Fryer and Greenman (1987) for a graphical exposition

Interpretation: A Consumer Example

A Consumer at the Grocery Store

Consumer buys x_1 and x_2 at prices p_1 and p_2, with income I:

$$p_1 x_1 + p_2 x_2 = I.$$

- Utility $u(x_1, x_2)$ with usual properties
- The cash register is **lazy**: doesn't check the basket at the checkout
- Consumer pays I regardless of the value of items taken
- Without a penalty, consumer would happily go over budget

The Gatekeeper

At the exit, a **gatekeeper** checks the basket. He charges a fee

$$F(x_1, x_2, \lambda) \equiv -\lambda(I - p_1 x_1 - p_2 x_2)$$

in **utility units**, with $\lambda \geq 0$.

- If consumer goes over budget ($I - p_1 x_1 - p_2 x_2 < 0$): the fee is positive (punishment)
- If consumer goes under budget ($I - p_1 x_1 - p_2 x_2 > 0$): the fee is negative (reward)

Consumer's utility upon exit:

$$L(x_1, x_2, \lambda) \equiv u(x_1, x_2) + \lambda(I - p_1 x_1 - p_2 x_2).$$

The Gatekeeper's Choice of λ

The gatekeeper wants the budget constraint to hold **voluntarily**.

- λ small $\Rightarrow$ consumer overspends ($I - p_1x_1 - p_2x_2 < 0$)
- λ large $\Rightarrow$ consumer underspends to collect the fee
- Somewhere in between: a λ^* that makes the budget constraint hold with equality

At λ^*, the consumer's choice (x_1^*, x_2^*) also solves

$$\max_{x_1,x_2} u(x_1, x_2) \quad \text{s.t.} \quad p_1x_1 + p_2x_2 = I.$$

Any feasible improvement over $u(x_1^*, x_2^*)$ (under the binding budget constraint) would also raise $L(\cdot, \lambda^*)$, contradicting that (x_1^*, x_2^*) maximizes L at λ^*.

How the Gatekeeper Finds λ^*

Key fact: by satisfying the budget constraint and choosing (x_1^*, x_2^*), the consumer can **always** attain

$$L(x_1^*, x_2^*, \lambda) = u(x_1^*, x_2^*) \quad \text{for any } \lambda.$$

- So $u(x_1^*, x_2^*)$ is a **floor** on what the consumer can achieve
- With λ other than λ^*, the consumer may do strictly better
- The right λ^* is the one that **minimizes** the consumer's best L, bringing it to the floor

Hence the gatekeeper solves

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

which is exactly the saddle-point problem from the general theory.

Interpretation of λ^*

At the saddle point:

- The budget constraint is satisfied
- Utility $u(x_1, x_2)$ is maximized within the budget
- λ^* is the punishment (in utility units) the consumer is just willing to pay for going over budget by one dollar

λ^* **is the value of one dollar in utility terms.**

The Meaning of the Lagrange Multiplier

Marginal Value of Income

Let us build on this intuition to formally see the meaning of the Lagrange multiplier

Write the optimal choice as $x_1^*(I), x_2^*(I)$ (suppressing p_1, p_2).

The value of one extra dollar in utility terms is

$$\frac{d}{dI}\, u\big(x_1^*(I), x_2^*(I)\big) = x_1^{*\prime}(I)\, u_1\big(x_1^*(I), x_2^*(I)\big) + x_2^{*\prime}(I)\, u_2\big(x_1^*(I), x_2^*(I)\big).$$

- u_i denotes the partial derivative of u w.r.t. its i-th argument
- Primes denote derivatives with respect to I

Using the First-Order Conditions

The FOCs for the Lagrangian:

$$\frac{\partial}{\partial x_i} L(x_1^*, x_2^*, \lambda^*) = u_i\big(x_1^*(I), x_2^*(I)\big) - \lambda^* p_i = 0.$$

Differentiating the budget constraint $x_1^*(I)p_1 + x_2^*(I)p_2 = I$:

$$x_1^{*\prime}(I)p_1 + x_2^{*\prime}(I)p_2 = 1.$$

Substituting $p_i = u_i/\lambda^*$ from the FOC:

$$x_1^{*\prime}(I)\frac{u_1\big(x_1^*(I), x_2^*(I)\big)}{\lambda^*} + x_2^{*\prime}(I)\frac{u_2\big(x_1^*(I), x_2^*(I)\big)}{\lambda^*} = 1.$$

λ^* is the Marginal Utility of Income

Combining with the marginal-value formula:

$$\frac{d}{dI}\,u\big(x_1^*(I), x_2^*(I)\big) = \lambda^*.$$

- λ^* is the **marginal utility of an additional dollar of income**
- Equivalent to the consumer's willingness to pay (in utils) to relax the budget constraint by one dollar
- Confirms the gatekeeper's interpretation

More generally:

The Lagrange multiplier is the marginal gain in the objective when the constraint is loosened by one unit.

Summary

Key Takeaways

1. **Constrained $\to$ unconstrained**: replace the constraint with a penalty $\lambda g(x_1, x_2)$
2. **Saddle point**: solve $\min_\lambda \max_{x_1,x_2} L(x_1, x_2, \lambda)$
3. **FOCs**: the textbook Lagrangian recipe is just the saddle-point conditions
4. **Economic meaning**: λ^* is the marginal value of relaxing the constraint
    - In the consumer example: marginal utility of income
5. **Intuition via the gatekeeper**: a fee per unit of constraint violation, set so the consumer chooses to satisfy the constraint

References

- Fryer, M. J. and J. V. Greenman (1987). Optimisation Theory. Edward Arnold.