

7. OPTIMIZATION: FINDING MAXIMA AND MINIMA

J.P. Bastos*

EPH 1100: Analytical Tools

Reading: Hoy et al., §6.1 (pp. 225–238), §6.2 (pp. 239–245). (§6.3 is not assigned.)

- **§6.1:** Necessary conditions (FOC); stationary values; monopoly profit maximization example.
- **§6.2 (pp. 239–245 only):** Sufficient conditions (SOC); classifying stationary values as maxima, minima, or inflection points. Stop after the Laffer curve example; the rest of §6.2 (monopsony, Pigou taxes, Taylor series) is not covered in this course.

1. The Optimization Problem

Almost every model in economics is built around an agent making the best choice available: a firm picks the output that maximizes profit, a consumer picks the bundle that maximizes utility, a government picks a tax rate that maximizes revenue.¹ The mathematics behind all of these is the same: find the value of a variable that makes some objective function as large or as small as possible.

Our recent classes on derivatives have shown how to find the direction of change and how to read the shape of a function from its derivatives, hinting at the tools we will be using for optimization. This class turns that into a systematic procedure for finding and confirming optimal choices, namely maxima and minima.

2. The First-Order Condition

Think about what must be true at the top of a hill. At the very peak, the ground is flat: you are neither climbing nor descending, which means the slope is zero. If the slope were positive, you would still be going up and the top hasn't been reached; if the slope were negative, you would already be past the peak.

The same logic applies to any differentiable function. At any interior maximum or minimum, the tangent line is horizontal, meaning the first derivative must equal zero.

*Assistant Professor of Economics, University of Austin (UATX). jbastos@uaustin.org

¹Even if, say, governments in the real world are not trying to strictly maximize revenues, we *model* it this way. Otherwise, we do not have a well-defined rule of behavior and the theory becomes predictionless – sometimes governments will do this, sometimes that. See Lazear (2000).

Definition. The requirement $f'(x^*) = 0$ is called the **first-order condition (FOC)** for an optimum. Any point x^* where $f'(x^*) = 0$ is called a **stationary point**, and $f(x^*)$ is the corresponding **stationary value**.

The FOC is a *necessary but not sufficient* condition: any interior maximum or minimum must satisfy it, but satisfying the FOC does **not** guarantee that x^* is a maximum or minimum. It may also be a point of inflection.

FOC as Marginal Reasoning

In economics, derivatives are marginal quantities, so the FOC always has a natural “marginal” interpretation. For a firm choosing output q to maximize profit $\pi(q) = R(q) - C(q)$, the FOC $\pi'(q^*) = 0$ becomes:

$$R'(q^*) = C'(q^*) \quad \Longleftrightarrow \quad \text{MR} = \text{MC}$$

At the profit-maximizing output, the extra revenue from selling one more unit must exactly equal the extra cost. If $\text{MR} > \text{MC}$, the firm is leaving profit on the table and should expand. If $\text{MR} < \text{MC}$, the last unit costs more than it earns and output should fall. The familiar rule $\text{MR} = \text{MC}$ is simply the FOC written in economic language.

4. The Second-Order Condition

Setting $f'(x^*) = 0$ locates a stationary point, but a zero slope can occur for three different reasons: a point of inflection, a local maximum, or a local minimum. To distinguish between them, we need the second derivative. The second derivative $f''(x^*)$ tells us how the slope is changing at the stationary point:

Second-Order Conditions (SOC). Suppose $f'(x^*) = 0$.

- If $f''(x^*) < 0$, then x^* is a **local maximum**.
- If $f''(x^*) > 0$, then x^* is a **local minimum**.
- If $f''(x^*) = 0$, the test is inconclusive; further analysis is needed.

A common example of an inflection point is $f(x) = x^3$. Its derivative is $f'(x) = 3x^2$, so $f'(0) = 0$: the origin satisfies the FOC. Yet f has no maximum or minimum at $x = 0$; the function is increasing for any x . The function momentarily levels off exactly at $x = 0$, and around that point the function changes from increasing at decreasing rates ($f'(x) > 0$; $f''(x) < 0$) to increasing at increasing rates ($f'(x) > 0$; $f''(x) > 0$). Thus, notice two things: first, the slope itself ($f'(x)$) does not change sign (it is positive always), and the change of the second derivative occurs right before and right after, *but not at* $x = 0$. This is why a second condition evaluated at $f'(x) = 0$ is needed.

Figure 1 puts all three outcomes side by side: the same FOC, $f'(0) = 0$, but three different verdicts once the SOC is applied.




Figure 1: Inflection Point, Minimum, and Maximum.

Notes: All three functions satisfy the FOC at $x = 0$ (marked point). $y = x^3$ has $f''(0) = 0$: the slope does not change sign around $x = 0$, so this is neither a max nor a min, but an inflection point. $y = x^2$ has $f''(0) = 2 > 0$: the slope is rising through zero, confirming a local minimum. $y = -x^2$ has $f''(0) = -2 < 0$: the slope is falling through zero, confirming a local maximum.

- If $f''(x^*) < 0$: the slope is *falling* through zero. The function was rising to the left and is now falling to the right. The function is concave at x^* : it bends like the top of a hill. Combined with the FOC, this confirms a **local maximum**.
- If $f''(x^*) > 0$: the slope is *rising* through zero. The function was falling to the left and is now rising to the right. The function is convex at x^* : it bends like the bottom of a bowl. Combined with the FOC, this confirms a **local minimum**.

Local vs. global optima. These conditions identify *local* optima: the best point within some neighborhood of x^* . In many economic models there is only one stationary point, so it is automatically global. A stronger guarantee: if $f''(x) < 0$ for all x (the function is globally concave), any local maximum is also the unique global maximum. The same applies symmetrically for globally convex functions and minima.

A function can have more than one stationary point, each with its own verdict. Figure 2 previews the general case: both marked points satisfy the FOC, $f'(x) = 0$, which is exactly why they are candidates in the first place. The SOC is what tells them apart.

5. The Optimization Workflow

To find and classify the optimum of $f(x)$:

1. **FOC:** Set $f'(x) = 0$ and solve for the stationary point(s) x^* .
2. **SOC:** Evaluate $f''(x^*)$.
 - $f''(x^*) < 0$: confirmed local maximum.
 - $f''(x^*) > 0$: confirmed local minimum.
3. **Report:** Substitute x^* back into f to find the optimal value $f(x^*)$.




Figure 2: Local Maximum and Minimum.

Notes: $f(x) = x^3 - 9x^2 + 24x$ has two stationary points, $x = 2$ and $x = 4$, both satisfying the FOC $f'(x) = 0$. The SOC tells them apart: $f''(x) = 6x - 18$, so $f''(2) = -6 < 0$ (local max) and $f''(4) = 6 > 0$ (local min). This is exactly the two-step procedure formalized below.

Example 1: A monopolist faces inverse demand $p = 100 - Q$ and total cost $C(Q) = 25Q$. Find the profit-maximizing output, price, and profit.

Step 1: Set up the profit function. Revenue is price times quantity sold, $R(Q) = P(Q) \cdot Q$; substituting the inverse demand curve $P(Q) = 100 - Q$ gives:

$$R(Q) = (100 - Q) \cdot Q$$

Profit is revenue minus cost:

$$\pi(Q) = R(Q) - C(Q) = (100 - Q)Q - 25Q = 75Q - Q^2$$

Step 2: Apply the FOC.

$$\pi'(Q) = 75 - 2Q = 0 \implies Q^* = 37.5$$

In terms of marginal quantities: $MR = 100 - 2Q$ and $MC = 25$. Setting $MR = MC$ gives $100 - 2Q = 25$, which yields the same answer. The two approaches are equivalent.

Step 3: Apply the SOC.

$$\pi''(Q) = -2 < 0 \text{ for all } Q$$

Since π'' is negative everywhere, π is globally concave and $Q^* = 37.5$ is the unique global maximum.

Step 4: Report the results.

$$p^* = 100 - 37.5 = \$62.50, \quad \pi^* = 75(37.5) - (37.5)^2 = \$1,406.25$$

Example 2: A government levies a specific tax of \$ t per unit on a market. As the tax rises, quantity traded falls: $Q(t) = 120 - 4t$. Total tax revenue is:

$$T(t) = t \cdot Q(t) = 120t - 4t^2$$

What tax rate maximizes revenue?

Step 1: Apply the FOC.

$$T'(t) = 120 - 8t = 0 \implies t^* = 15$$

Step 2: Apply the SOC.

$$T''(t) = -8 < 0$$

The function is globally concave, confirming $t^* = \$15$ is the global revenue maximum.

Step 3: Report the result.

$$T^* = 120(15) - 4(15)^2 = 1,800 - 900 = \$900$$

The concavity of $T(t)$ has a concrete policy implication. Starting from $t = 0$, raising the tax rate initially increases revenue: each dollar of tax applies to many units sold. But as the tax climbs, the quantity effect grows: fewer and fewer units are traded. Beyond $t^* = \$15$, higher taxes actually reduce total revenue because the volume loss outweighs the rate gain. The government is in self-defeating territory: it collects less by taxing more. This inverted-U shape is sometimes called the **Laffer curve**, shown in Figure 3.




Figure 3: The Laffer curve: tax revenue as a function of the tax rate.

Notes: $T(t) = 120t - 4t^2$. Revenue rises with the tax rate up to $t^* = 15$, where $T'(t) = 0$ and $T''(t) = -8 < 0$ confirms a maximum. Beyond t^* , the shrinking tax base outweighs the higher rate and revenue falls.

Practice Problems

- Find all stationary points of each function. Apply the SOC to classify each as a local maximum, local minimum, or state that the test is inconclusive.

(a) $f(x) = -x^2 + 8x - 7$

(b) $f(x) = x^3 - 6x^2 + 9x + 2$

(c) $f(x) = x^4 - 8x^2 + 3$

- A firm sells in a competitive market at price $p = 10$. Its cost function is $C(q) = q^2 + 2q + 5$.

(a) Write the profit function $\pi(q)$.

(b) Find the profit-maximizing output using the FOC. Confirm it is a maximum using the SOC.

(c) What is the maximum profit? Show that $p = MC$ at the optimum.

- A monopolist faces inverse demand $p = 60 - 2Q$ and cost function $C(Q) = 20Q$.

(a) Find the profit-maximizing output and price.

(b) Verify the SOC.

(c) Find MR as a function of Q and confirm that $MR = MC$ at the optimum found in (a).

- A government sets an income tax rate $\tau \in [0, 1]$. In a simple model, total income before tax is $Y = \$800$, and the amount of taxable activity (hours worked, transactions made) falls with the tax rate: effective taxable income is $(1 - \tau) \cdot Y$. Tax revenue is therefore:

$$T(\tau) = \tau \cdot (1 - \tau) \cdot 800 = 800\tau - 800\tau^2$$

(a) Find the revenue-maximizing tax rate τ^* .

(b) Verify the SOC.

(c) What does this model imply about tax rates above τ^* ? Is a 100% tax rate a good idea?

- A workshop has a fixed stock of $\bar{K} = 10$ machines. Daily output depends on the number of workers L scheduled onto those machines:

$$Q(L) = 60L - 2L^2$$

Beyond some point, adding another worker just means more people waiting for a free machine or getting in each other's way, and output starts to fall.

(a) Find the number of workers that maximizes daily output using the FOC.

(b) Verify it is a maximum using the SOC.

- (c) What is the maximum daily output?
- (d) In plain language, why does output eventually decline as more workers are added, even though nothing is broken or removed?

References

Lazear, E. P. (2000). Economic imperialism. *The Quarterly Journal of Economics*, 115(1):99–146.