



Ciências
ULisboa

Faculdade
de Ciências
da Universidade
de Lisboa

Principles of Economics and Management I

Session 4 · Production, Costs and Market Structures

Introduction



Where We Came From

Sessions 2 and 3 built the buyer: preferences, the constrained optimum, market demand and its elasticity.

Today we build the other side. Supply comes out of **costs** and the objective of **maximizing profit**. 🏭

And then we ask what changes when the firm is large enough to face the whole demand curve by itself. 👑



Today's Map

1. 🏭 Production in the short and the long run, and why MC has a U shape
2. 🧱 Fixed, variable, average and marginal cost
3. 🎯 Profit maximization: $MR = MC$
4. 🏁 Perfect competition: $P = MC$, and where the supply curve comes from
5. ❤️ Producer surplus, shutdown, entry, and the elasticity of supply
6. 👑 Monopoly, the markup, and the link back to elasticity

By the end we hold a supply curve. Session 5 puts it against demand. 🔧

Part 1 · Production and Costs



Where Costs Come From

Costs come from a **production function** and the prices of inputs.

$$q = F(K, L)$$

where K is capital (the plant, the machines) and L is labour (hours of work).

So before we talk about costs, we need to know how output responds when the firm uses more K or more L .

 Source: OpenStax *Micro* 3e, §7.2 “Production in the Short Run”, tabulated for a firm adding workers one at a time.



Short Run and Long Run

Some inputs you can change this week, some you cannot.

You can call in two more workers for tomorrow morning. You cannot build a second factory by tomorrow morning. 🏗️

Short run: at least one input is **fixed**. **Long run:** every input can be changed.



Short Run and Long Run

In the short run the plant is given, $K = \bar{K}$, and the only choice left is L :

$$q = F(\bar{K}, L) = f(L)$$

This is why we write $f(L)$ and not $F(K, L)$: today the firm already has its plant.



Short Run and Long Run

The fixed input is where the **fixed cost** comes from. With $\bar{K}$ machines rented at r each, and workers paid the wage w :

$$FC = r\bar{K}, \quad VC = wL$$

The rent is paid whether you produce 0 units or 1000. The wage bill depends on how many people you call in.

In the long run there is no fixed cost: the firm can give back the plant, or build a bigger one. We come back to this at the end of Part 2. 



Marginal and Average Product

Hold $\bar{K}$ fixed and add workers one at a time.



Marginal product of labour: the extra output from one more worker, $MP_L = f'(L)$.

Average product of labour: output per worker, $AP_L = f(L)/L$. This is what we usually call **productivity**.

Same relation as marginal and average cost (we will see it in a moment): while $MP_L > AP_L$ the new worker pulls the average up, once $MP_L < AP_L$ the new worker pulls it down. So MP_L crosses AP_L at the **maximum** of AP_L .



Marginal and Average Product



MP_L rises up to the fifth worker and falls after that. AP_L peaks where MP_L crosses it.



Why MP Rises First, Then Falls

With two workers in a big workshop, each does everything: cut, assemble, pack, carry. A third one lets them split the tasks, each gets good at one thing and nobody walks between stations. This is the **gain from specialization**, and MP_L rises.



Adam Smith's pin factory (1776): ten workers splitting 18 operations made about 48 000 pins a day. One worker alone, he guessed, would not make twenty.

Why MP Rises First, Then Falls

But the plant is fixed. The tenth worker shares the same machines and the same floor with the other nine, and spends part of the day waiting for a machine. **Diminishing marginal returns:** from some point on, $f'' < 0$.

Diminishing returns need the fixed $\bar{K}$. In the long run, when the plant can grow with the workforce, nothing forces them.


Which Is Exactly Why MC Has a U Shape

To make one more unit you need $1/MP_L$ extra workers, each costing the wage w . So

$$MC = \frac{w}{MP_L}$$

While specialization makes MP_L rise, MC falls. Once diminishing returns set in, MP_L falls and MC rises.

And it tells you what shifts MC : the **wage** and **productivity**. Hold that list; it is exactly what shifts supply in session 5. 




Productivity and Average Cost

Same trick for the average. Variable cost is the wage bill,
 $VC = wL$, so


$$AVC = \frac{wL}{q} = \frac{w}{q/L} = \frac{w}{AP_L}$$

The more productive the average worker, the lower the cost per unit.


Put the two side by side: $MC = w/MP_L$ and $AVC = w/AP_L$.
 MP_L crosses AP_L at the **maximum** of AP_L , so MC crosses
 AVC at the **minimum** of AVC .



Productivity and Average Cost



The cost curves on the right are the product curves on the left, turned upside down ($w = 30$).



From a Production Function to a Cost Curve

Take $q = \sqrt{L}$, so $L = q^2$, with wage w and a fixed cost F .

$$TC(q) = F + w q^2, \quad MC(q) = 2wq, \quad AC(q) = \frac{F}{q} + wq$$

MC rises because $f'' < 0$; AC falls first (the fixed cost is spread over more units) and rises later (marginal cost drags it up).

Here $MP_L = \frac{1}{2\sqrt{L}}$ falls from the very first worker: no specialization phase, only diminishing returns. The algebra is simpler and the logic is the same. 

Types of Cost

Fixed cost (FC): does not depend on quantity (rent, insurance). It comes from the fixed input.

Variable cost (VC): grows with output (wages, materials, energy).

Total cost: $TC = FC + VC$.

Average cost: $AC = TC/q$.

Average variable cost: $AVC = VC/q$.

+ Marginal Cost

Marginal cost (MC): the cost of producing **one more** unit.


$$MC = \frac{dTC}{dq} = \frac{dVC}{dq}.$$

The fixed cost does not change with q , so it drops out of the derivative.

Typically U shaped, as we just saw: it falls first (gains from specialization), then rises (diminishing returns).



The Cost Curves




Note: MC crosses AVC and AC at their **minima**, and it has to.

$AC' = (MC - AC)/q$, so AC falls exactly while $MC < AC$ and rises once $MC > AC$ (the same holds for AVC). 

 Source: OpenStax *Micro* 3e, §7.3 “Costs in the Short Run”, “Fixed and Variable Costs” and “Where do marginal and average costs meet?”



Short Run and Long Run



Each orange curve is the short run AC of one plant size $\bar{K}$. In the long run the firm picks the plant, so it sits on the lowest curve for the q it wants: the long run AC is their **lower envelope**.



Short Run and Long Run

Read the long run AC from left to right.

Where it falls, a bigger firm makes each unit more cheaply: **economies of scale**. Where it rises, the firm has become too big to run well: **diseconomies of scale**.

If it falls over the whole range the market needs, one firm can serve everyone more cheaply than two. Keep this for Part 3. 👑

From here until the long run supply, we stay in the **short run**: the plant $\bar{K}$ is fixed, and the cost curves are the ones of the previous plot.

The Golden Rule of Profit

The firm chooses q to maximize profit $\pi(q) = R(q) - TC(q)$.

First order condition: $\pi'(q) = 0$, that is

$$MR = MC.$$

Produce up to the point where **marginal revenue** equals **marginal cost**.

Check the second order condition too: $\pi'' = MR' - MC' < 0$, so MC must cut MR **from below**. Of the two solutions, only the **rising** one is a maximum. ⚠️

? Multiple Choice 1

A firm's marginal cost curve is U shaped and it faces a given price. Profit is maximized where $P = MC$:


- A. At either of the two quantities where the curves cross.
 - B. At the lower of the two, where MC is falling.
 - C. Only if the firm is also minimizing average cost.
 - D. At the higher of the two, where MC is rising.
- ✓ **D.** The second order condition needs $MC' > MR' = 0$. On the falling branch, $P = MC$ is a profit **minimum**, and producing one more unit strictly improves things.

? Multiple Choice 2

Marginal cost is currently **below** average cost. Then average cost is:

- A. Falling.
- B. Rising.
- C. At its minimum.
- D. Constant.

✓ **A.** From $AC' = (MC - AC)/q$, the sign of AC' is the sign of $MC - AC$. It is the same arithmetic as a grade average: a mark below your average pulls it down.



Numerical Question

A firm has $q = \sqrt{L}$, a wage of $w = 2$ euros and a fixed cost of $F = 18$ euros.

- a) Write $TC(q)$, $MC(q)$ and $AC(q)$.
- b) At which quantity is average cost minimized?
- c) Check that MC equals AC there.
- d) What is the minimum average cost?



Solution

- a)** $L = q^2$, so $TC = 18 + 2q^2$, $MC = 4q$ and $AC = 18/q + 2q$.
- b)** $AC' = -18/q^2 + 2 = 0 \Rightarrow q^2 = 9$, so $q = 3$.
- c)** $MC(3) = 12$ and $AC(3) = 6 + 6 = 12$. They meet, exactly as $AC' = (MC - AC)/q$ requires. 
- d)** 12 euros per unit. Note this is the long run break-even price: below it the firm cannot cover its full cost at any quantity.

Part 2 · Competitive Supply

Perfect Competition

Many small firms, a homogeneous product, buyers and sellers who know the prices, and **free entry and exit**. Each firm is a **price taker**.

It faces the polar case from session 3: a **perfectly elastic** demand for its own output. Charge a cent more and it sells nothing.

If it is a price taker, selling one more unit always brings in P : hence $MR = P$, and the profit rule becomes $P = MC$. 

 Source: OpenStax *Micro* 3e, §8.2 “How Perfectly Competitive Firms Make Output Decisions”, worked through for a raspberry farm, and “Marginal Cost and the Firm’s Supply Curve”.



Free Entry and Exit

Of all the assumptions, the one that does the most work is **free entry and exit**.

Anyone can open a firm in this market, and anyone can close one: no licence, no patent, no secret recipe, no cost of getting in that you cannot get back when you leave.

It is what makes profits temporary. If firms make money, others come in; if they lose money, some leave. We use it at the end of this part.

And it is exactly what a monopoly does not have. Part 3 starts from a **barrier to entry**. 



Where the Supply Curve Comes From

The firm takes the price as given and chooses q so that


$$MR(q) = MC(q) \implies P = MC(q)$$

Now change the price. At the new price the firm solves the same equation again, and picks a new q .

🤔 Do this for every possible price. What do you get?



Where the Supply Curve Comes From



Each dot is the optimal choice at one price. Join the dots and you have the firm's **supply curve**: the rising part of MC .



Where the Supply Curve Comes From

Read $P = MC(q)$ as a relation between the price and the quantity the firm chooses.

Written this way, with the price alone on the left, it gives the price at which the firm wants to sell q units: the **inverse supply** $P(q) = MC(q)$.

Now put q alone on one side and everything else (including P) on the other. That is the **supply** $q(P)$: how much the firm sells at each price.

With $q = \sqrt{L}$ from Part 1, $MC = 2wq$, so

$$P = 2wq \implies q(P) = \frac{P}{2w}$$

A higher price, more output. A higher wage, less output at every price.



The Seller's Side, Same Logic

Read $P = MC$ the way we read demand in session 3. A seller hands over the good and receives the price, so they sell when

$$p - c \geq 0$$

where c is the cost of producing **that** unit.

Reservation price of a seller: the lowest price at which they still sell. It equals their **marginal cost** c .

Sort those costs and you get the **supply curve**: whoever can produce below the price sells. 



A Warning About That Cost

c is the cost of producing **one more** unit.

Costs the firm pays whether or not it produces that unit (the rent on the factory, the licence, the loan already taken) are **not** in c .

So the supply curve is built out of **variable** cost only. We need this again for producer surplus.

+ From the Firm to Market Supply


Same aggregation as demand, in the other direction: at each price, ask every firm how much it wants to produce and **add the quantities**.

$$Q_s(p) = \sum_{j=1}^m q_j(p)$$

where $q_j(p)$ is firm j 's supply, solved from $p = MC_j(q_j)$.

Each firm's supply is its rising marginal cost, so market supply slopes **up**. Horizontal aggregation again. 

+ From the Firm to Market Supply



Firm 1: $q_1 = P/2$. Firm 2 only starts at $P = 2$: $q_2 = P - 2$. At $P = 8$ the market supplies $4 + 6 = 10$. The kink at $P = 2$ is where firm 2 enters.

Producer Surplus, Carefully


Producer surplus (PS): revenue minus the **variable** cost of what was produced. The area **below the price and above the supply curve**.

The supply curve is built from marginal costs, and marginal cost contains **only variable** cost. Fixed costs never entered it, so they are not in PS either.

$$PS(q) = p q - VC(q)$$

$$\pi(q) = p q - VC(q) - FC = PS(q) - FC$$

♥ Producer Surplus, Carefully



$MC = 1 + q$, $P = 7$, so $q^* = 6$. Revenue $7 \times 6 = 42$ is the whole rectangle. Below MC is the variable cost, 24. Above it, up to the price, is $PS = 18$.

So PS Is Not Profit

A firm can have a perfectly healthy producer surplus and still **lose money**, once the fixed costs are paid.

That gap is exactly why a firm keeps operating at a loss in the short run: as long as $PS > 0$, shutting down would lose even more.

Two different questions: “is this trade worth doing?” (PS) and “was this business worth starting?” (profit).

Should It Produce at All?

$P = MC$ says **how much**. It does not say **whether**. For that, compare with average cost.

 **Short run:** fixed costs are sunk, so ignore them (session 1). Produce if $P \geq AVC$; shut down if not.

 **Long run:** everything is variable. Stay in the market only if $P \geq AC$; otherwise exit.


Between AVC and AC the firm makes a **loss and still operates**, because it covers its variable cost and contributes something toward the fixed one.

That gap is exactly the difference between **producer surplus** and **profit** we saw above. Now you know where it lives. 

 Source: OpenStax *Micro* 3e, [§8.2](#), "The Shutdown Point"; the exit decision is [§8.3](#) "Entry and Exit Decisions in the Long Run".




Three Prices



$P = 11$: the firm produces $q^* = 7.50$, where $AC = 7.67$. Profit $= (P - AC) q^* = 25.00$ euros, the green rectangle. ●




Three Prices



$P = 6$, between min AVC and min AC. At $q^* = 5.97$ the firm loses 8.96 (red). Shut down, it would lose the whole fixed cost of 20. The blue part, 11.04, is fixed cost paid out of revenue. Keep producing. ●




Three Prices



$P = 3$, below min AVC . Even at its best $q^* = 4.17$ the firm loses 24.63: the fixed cost of 20, plus 4.63 because the price does not even pay the workers (dark red). Shut down and lose only 20. ●



The Firm's Supply Curve



$$q(P) = \begin{cases} 0 & \text{if } P < \min AVC \\ \text{the } q \text{ that solves } P = MC(q) & \text{if } P \geq \min AVC \end{cases}$$



In the Long Run, Entry Eats the Profit

Suppose price sits above minimum average cost, so firms are making money.


Nothing stops new firms entering. Market supply expands, and the price falls.

It stops falling only where profit is zero, that is at $P = \min AC$. Which is why long run competitive price equals minimum average cost.

“Zero profit” means capital earns exactly what it would earn elsewhere. The opportunity cost is already inside the cost. 💡



Long Run Market Supply



Demand rises from D to D' . In the short run the number of firms is fixed and the price jumps along the steep short run supply (A to B). Profits bring in new firms until the price is back at $\min AC$ (C). Joining A and C, the long run supply is **horizontal**.



Elasticity of Supply

Now that we hold a supply curve, we can measure it the way we measured demand.

Price elasticity of supply: the percentage change in quantity supplied over the percentage change in price.

$$\varepsilon_s = \frac{dQ_s}{dP} \cdot \frac{P}{Q_s} > 0.$$

It is positive, because supply slopes up.

It depends above all on **time** 🕒. Overnight a firm can only run its existing plant harder, *MC* climbs steeply and quantity barely moves: **inelastic**. Over years it can build another plant, and the same price rise brings out much more: **elastic**.


Session 5 needs this number. Together with ε on the demand side, it decides who pays a tax.



Source: OpenStax *Micro* 3e, §5.1, "Price Elasticity of Supply".





Elasticity of Supply



The same price rise, from 5 to 6. With the plant fixed, quantity goes from 50 to 52.5 ($\epsilon_s = 0.25$). With time to build, it goes to 70 ($\epsilon_s = 2$).



Elasticity of Supply: The Two Extremes



Vertical: the quantity does not move whatever the price.

Horizontal: at that price the firms supply any quantity, below it nothing.



Elasticity of Supply: The Two Extremes

$\varepsilon_s = 0$: the quantity cannot respond at all. The paintings of Amadeo de Souza-Cardoso (he died in 1918). The seats on the 8:00 Lisbon to Porto flight next Monday. Towels on one beach in August.

$\varepsilon_s \rightarrow \infty$: the long run market supply we drew above. With free entry, any price above $\min AC$ brings in as much output as the market wants.

✨ The Efficiency of Competition

Under perfect competition, $P = MC$: the price equals the cost of producing the last unit.

The buyer's value of the last unit equals its cost. No trade worth doing is left undone, and none that destroys value is done.

Everything in Part 3 is measured **against** this benchmark.

? Multiple Choice 3

In the short run a competitive firm finds $AVC < P < AC$. It should:


- A. Keep producing, even though it is making a loss.
- B. Shut down at once, since it is losing money.
- C. Raise its price until it covers AC.
- D. Produce where $P = AC$ instead.

✓ **A.** Fixed costs are **sunk** in the short run, so they should not enter the decision. Since the price covers variable cost, operating loses less than shutting down. Option C is not available: the firm is a price taker.

? Multiple Choice 4

Producer surplus is:

- A. Revenue minus total cost, including fixed cost.
 - B. Revenue minus the variable cost of what was produced.
 - C. The firm's accounting profit for the year.
 - D. The area above the demand curve and below the price.
- ☒ **B.** The supply curve is made of marginal costs, which contain no fixed cost. Subtract fixed cost from PS and only then do you have profit.



Numerical Question

Total cost $TC = q^2 + 10$ (euros), in a competitive market at a price of 12 euros.

- a) Find the optimal quantity.
- b) Compute profit.
- c) Compute producer surplus, and reconcile it with (b).
- d) Below which price would the firm shut down in the short run?



Solution

- a) $MC = 2q$, and $P = MC$ gives $12 = 2q$, so $q^* = 6$. MC is rising, so it really is the maximum.
- b) $\pi = 12 \times 6 - (36 + 10) = 72 - 46 = 26$ euros.
- c) $VC = q^2 = 36$, so PS = $72 - 36 = 36$ euros. And $PS - FC = 36 - 10 = 26$, the profit.



- d) $AVC = q^2/q = q$, which is minimized as $q \rightarrow 0$. So $AVC \leq P$ holds for any positive price: this firm **never** shuts down in the short run. The fixed cost of 10 is sunk and irrelevant to that call.

Part 3 · Market Power

Monopoly

A single seller. The firm is **not** a price taker: it faces the **entire** demand curve.

To sell more, it has to **cut** the price on all units. That is why $MR < P$.

The profit rule ($MR = MC$) still holds, but now MR lies **below** the price.



Where Market Power Comes From

A monopoly needs a barrier that keeps rivals out.



Costs: a natural monopoly, where average cost falls over the whole relevant range, so one firm serves the market more cheaply than two.



Law: patents, licences, concessions. Deliberately granted, usually to pay for the invention.




Control of an input, or a network that gets more valuable the more users it has.

 Source: OpenStax *Micro* 3e, [§9.1](#) “How Monopolies Form: Barriers to Entry”.



The Monopolist's Choice





It picks the q where $MR = MC$, and charges the price P_m **on the demand curve** (above MC).

 Source: OpenStax *Micro* 3e, §9.2 “How a Profit-Maximizing Monopoly Chooses Output and Price”, built numerically around a drug called HealthPill.



The Inefficiency of Monopoly



Demand $P = 10 - Q$, $MC = 2$. Competition: $Q_c = 8$ at $P = 2$.
Monopoly: $Q_m = 4$ at $P_m = 6$.



The Inefficiency of Monopoly

The monopolist produces **less** and charges **more** than competition would.

The yellow rectangle was consumer surplus under competition. Under monopoly it becomes the firm's profit, 16 euros. This part is a **transfer** from buyers to the seller, total surplus does not change because of it.

Between $Q_m = 4$ and $Q_c = 8$ buyers value the good above what it costs to make, $P > MC$, and those trades do **not** happen.

Deadweight loss: the surplus destroyed by producing below the efficient level. The red triangle, $\frac{1}{2}(6 - 2)(8 - 4) = 8$ euros.

Check the total: $32 = 8 + 16 + 8$. Competition gave all of it to buyers; monopoly splits it and loses a quarter.

 Source: OpenStax Micro 3e, [§9.2](#), "The Inefficiency of Monopoly".



Markup and Elasticity

Recall the formula from session 3: $MR = P \left(1 + \frac{1}{\varepsilon}\right)$. With $MR = MC$:

$$\frac{P - MC}{P} = -\frac{1}{\varepsilon} = \frac{1}{|\varepsilon|}.$$

Lerner index: the gap between price and marginal cost, **as a share of the price**. It is **larger** the **more inelastic** demand is. Market power is the power to set price above cost.



Sanity check: perfect competition is $|\varepsilon| \rightarrow \infty$, giving a Lerner index of zero and $P = MC$. Competition is the limit case.

A Monopolist Never Prices in the Inelastic Region

From the same first order condition:

$MC \geq 0$ and $MC = MR = P \left(1 + \frac{1}{\varepsilon}\right)$, with $P > 0$. So $1 + 1/\varepsilon \geq 0$, which forces $|\varepsilon| \geq 1$.

So a profit-maximizing monopolist always sits on the **elastic** part of demand. Never the inelastic part.

The intuition from session 3: where demand is inelastic, raising the price raises **revenue** and cuts **output**, so it raises revenue and cuts cost at the same time. No firm leaves that on the table.





Financial Applications



Understanding **market power** is essential for valuing firms: high margins signal competitive advantage, and the Lerner index says where they come from.

Sector analysis (how concentrated the market is, how elastic demand is) feeds directly into asset valuation. 

A margin that no barrier explains will not survive competition.

? Multiple Choice 5

Under monopoly, at the optimum:

A. $P = MC$.

B. $MR > P$.

C. $P > MC$.

D. $P = MR$.

☒ **C.** Since $MR < P$ and $MR = MC$ at the optimum, it follows that $P > MC$: the price is above marginal cost.

? Multiple Choice 6

A profit-maximizing monopolist with $MC \geq 0$ always operates where demand is:


A. Inelastic, since that is where the margin is largest.

B. Unit elastic, since that maximizes revenue.

C. At whatever elasticity the market happens to have.

D. Elastic, since $MR = MC \geq 0$ requires $|\varepsilon| \geq 1$.

✓ **D.** From $MC = P(1 + 1/\varepsilon) \geq 0$. Option A confuses the Lerner index, which says a *less* elastic demand allows a bigger margin, with the claim that the firm ends up in the inelastic **region**. It never does.



Numerical Question

A monopoly with demand $P = 20 - Q$ and constant $MC = 4$ euros. Note $MR = 20 - 2Q$.

- a) Find the monopoly quantity and price.
- b) Compute the Lerner index and the elasticity at that point.
- c) What would a competitive industry with the same costs produce?
- d) Compute the deadweight loss.



Solution

- a) $20 - 2Q = 4 \Rightarrow Q_m = 8$, and $P_m = 20 - 8 = 12$ euros.
- b) Lerner = $(12 - 4)/12 = 2/3$, so $|\varepsilon| = 3/2 > 1$: elastic, as it must be.
- c) $P = MC$ gives $20 - Q = 4$, so $Q_c = 16$ at a price of 4 euros. The monopolist produces **half** as much.
- d) The lost triangle between Q_m and Q_c : $\frac{1}{2}(12 - 4)(16 - 8) = 32$ euros of surplus that simply disappears. 

Wrap-Up

What to Take From This Session

 In the short run the plant is fixed and $q = f(L)$. Specialization first, diminishing returns later, so $MC = w/MP_L$ has a U shape and $AVC = w/AP_L$.

 MC cuts AC at its **minimum**: $AC' = (MC - AC)/q$ is an identity, true for any cost function.

The firm produces where $MR = MC$, on the **rising** branch. Under price taking that is $P = MC(q)$, the inverse supply; solve it for q and you have the supply $q(P)$.

 **PS is not profit**: it leaves out fixed cost, which is what makes the short run shutdown rule $P \geq AVC$ rather than $P \geq AC$. In the long run, entry pushes the price to $\min AC$.

 Market power is $P > MC$, and the Lerner index ties the margin to $1/|\varepsilon|$. Competition is the case $|\varepsilon| \rightarrow \infty$.

Next Session

Market Equilibrium and Taxes.

We now have both curves. Next they meet: what sets the price, why that outcome maximizes surplus, and what a tax does to it.



See you next week. 