

Course Companion

AP Microeconomics

Every topic handout in one volume

exercises.lol

Contents

1. Basic Economic Concepts	2
2. Supply and Demand	10
3. Production, Cost, and the Perfect Competition Model	21
4. Imperfect Competition	29
5. Factor Markets	34
6. Market Failure and the Role of Government	40

Basic Economic Concepts

AP Microeconomics

Scarcity



An outdoor market: markets allocate scarce resources among competing uses

Economics begins with one stubborn fact: people's wants are effectively unlimited, but the **resources** to satisfy them are limited. This gap is **scarcity**, and it is why economics exists – if everything were free and endless, there would be nothing to decide.

The resources (the **factors of production**) are **land** (natural resources), **labor** (human effort), **capital** (tools, machines, buildings), and **entrepreneurship** (organizing the other three and taking risks). Because these are scarce, every economy must answer three questions: **what** to produce, **how** to produce it, and **for whom** to produce.

Scarcity forces **choice**, and every choice has an **opportunity cost** – the value of the next-best alternative you give up. This single idea runs through the whole course: the true cost of anything is what you sacrifice to get it, not just the money you pay.

Resource Allocation and Economic Systems

An **economic system** is a society's way of answering the three questions.




Economic systems form a spectrum from a free market to full government planning

- In a **market economy**, private buyers and sellers decide through prices; self-interest and competition coordinate them (the “invisible hand”).
- In a **command economy**, the government owns resources and plans production.
- Real economies are **mixed** – mostly markets, with government providing some goods and correcting problems.

The key advantage of markets is that prices carry information and **incentives** automatically, with no central planner needed.

The Production Possibilities Curve


The **production possibilities curve (PPC)** shows the maximum combinations of two goods an economy can produce with its resources and technology fully and efficiently used.



Points on, inside, and outside the production possibilities curve

- Points **on** the curve are **efficient**; points **inside** are inefficient (wasted or unemployed resources); points **outside** are currently **unattainable**.
- The PPC's downward slope shows **opportunity cost**: making more of one good means making less of the other. The **slope** at a point is the opportunity cost of the good on the horizontal axis.
- A **bowed-out** (concave) PPC shows **increasing opportunity cost** – resources are not equally suited to both goods, so making more of one costs ever more of the other. A **straight-line** PPC means constant opportunity cost.
- **Economic growth** (more resources, better technology) shifts the whole curve **outward**.


Exam skill: be ready to read opportunity cost off a PPC, decide if a point is efficient/inefficient/unattainable, and explain what shifts the curve.



A production possibilities curve: efficient, inefficient, and unattainable points

Comparative Advantage and Gains from Trade

Two producers can both gain from trade by **specializing**.




Comparative advantage comes from differently-sloped production possibilities curves

- A producer has an **absolute advantage** in a good if it can make more of it with the same resources.
- A producer has a **comparative advantage** if it has the **lower opportunity cost** for that good. Comparative advantage – not absolute – determines who should specialize.
- Compute opportunity cost from output data: for each producer, one unit of good X costs (units of Y) ÷ (units of X). The producer who gives up less Y should make X.
- Both sides gain if they trade at a rate **between** their two opportunity costs – the **terms of trade**. Specializing by comparative advantage lets total output exceed what either could make alone, so consumption can lie **outside** each producer's own PPC.

Worked example. Anna can make 20 shirts **or** 10 tables a day; Ben can make 12 shirts **or** 12 tables. For Anna one table costs $\frac{20}{10} = 2$ shirts; for Ben it costs $\frac{12}{12} = 1$ shirt. Ben gives up fewer shirts, so **Ben specializes in tables** and Anna in shirts. Any terms of trade **between 1 and 2 shirts per table** – say 1.5 – leaves both better off than producing alone.

Exam skill: these calculations appear almost every year – practice finding comparative advantage from both output tables and input tables, and stating a valid terms-of-trade range.



Comparative advantage goes to the producer with the lower opportunity cost




Container shipping: comparative advantage and trade let countries specialise and gain

Cost-Benefit Analysis

Rational decision-makers weigh **benefits against costs**. An action is worthwhile only when its benefit is at least as large as its cost – including the opportunity cost. Money already spent and unrecoverable is a **sunk cost** and should be **ignored** in decisions: only future costs and benefits matter.

Marginal Analysis and Consumer Choice

Most real decisions are not “all or nothing” but “a little more or a little less.” **Marginal analysis** compares the **marginal benefit** (the benefit of one more unit) with the **marginal cost** (the cost of one more unit).



Each extra unit consumed gives less extra satisfaction (diminishing marginal utility)

- **Decision rule:** keep doing the activity while marginal benefit > marginal cost; stop where **MB = MC**. That point maximizes net benefit (total benefit minus total cost).
- Consumers get **utility** (satisfaction) from goods, but each extra unit usually gives less added satisfaction – the **law of diminishing marginal utility**
- A consumer maximizes total utility by the **utility-maximizing rule:** spend so that the **marginal utility per dollar** is equal across all goods, $\frac{MU_x}{P_x} = \frac{MU_y}{P_y}$, using up the whole budget. If one good gives more utility per dollar, buy more of it until the ratios equalize.

Worked example. A consumer is choosing between good X ($MU_x = 20$, $P_x = \$2$) and good Y ($MU_y = 15$, $P_y = \$1$). The utility per dollar is $\frac{20}{2} = 10$ for X but $\frac{15}{1} = 15$ for Y. Y delivers more satisfaction per dollar, so the consumer should buy **more Y and less X**; as they do, MU_y falls and MU_x rises until $\frac{MU_x}{2} = \frac{MU_y}{1}$ and the budget is spent.

Exam skill: the “MB = MC” logic and the “equal MU-per-dollar” rule are tested repeatedly – be able to find the optimal quantity from a table and explain *why* it is optimal.

Supply and Demand

AP Microeconomics


Demand



Shoppers in a market: demand shows how quantity demanded responds to price

A market works only where **property rights** are well defined – owners can use, exclude others from, and trade their resources; without them, prices cannot coordinate exchange. Within that framework, the two forces are demand and supply.


Demand is the relationship between the price of a good and the quantity buyers are willing and able to buy. The **law of demand** says price and quantity demanded move in **opposite** directions, so the demand curve slopes **downward**. Two reasons: the **substitution effect** (when a good gets pricier, buyers switch to alternatives) and the **income effect** (a higher price shrinks real purchasing power).



Rational choice: keep going while marginal benefit exceeds marginal cost, stopping where they are equal

Exam tips

- Opportunity cost is the next-best alternative given up; **ignore sunk costs** in any decision.
- Comparative advantage (lower opportunity cost) determines specialisation and the gains from trade.
- Apply marginal analysis: keep doing an activity while **MB > MC**, stopping where **MB = MC**.
- Maximise utility by equalising the **marginal utility per dollar** across goods.
- Read opportunity cost off the PPC slope; a bowed curve shows increasing cost.



The demand curve slopes down; income, tastes, and related prices shift it


Distinguish a **change in quantity demanded** (a move *along* the curve, caused only by the good's **own price**) from a **change in demand** (a *shift* of the whole curve). Demand shifts with the **determinants**, remembered as **TRIBE**:

- **Tastes/preferences**,
- **Related goods** – a rise in the price of a **substitute** raises demand; a rise in the price of a **complement** lowers it,
- **Income** – for a **normal good** more income raises demand; for an **inferior good** it lowers it,
- **Buyers** (number of consumers),
- **Expectations** of future price or income.

The **market demand curve** is the **horizontal summation** of every individual buyer's demand curve – at each price, add up the quantities all buyers want. That is why more buyers (the **B** above) shifts the whole market curve to the right.

Supply

Supply is the relationship between price and the quantity sellers will offer. The **law of supply**: price and quantity supplied move in the **same** direction, so the supply curve slopes **upward** (higher prices cover higher marginal costs and reward more output).



The supply curve slopes up; costs, technology, tax, and subsidy shift it

Again separate a **change in quantity supplied** (movement along the curve, from the good's own price) from a **change in supply** (a shift). Supply shifts with **input prices, technology, taxes and subsidies**, **prices of other goods** a firm could make, **number of sellers**, and **expectations**. Just like demand, the **market supply curve** is the horizontal sum of individual firms' supply curves – add the quantities all firms offer at each price – so more sellers shifts market supply to the right.


Exam skill: on the free-response section you must draw correctly labeled supply-and-demand graphs and shift the *right* curve in the *right* direction – decide first whether the event changes an “own price” (move along) or a determinant (shift).

Price Elasticity of Demand

Price elasticity of demand measures how responsive quantity demanded is to a price change:

$$E_d = \left| \frac{\% \Delta \text{quantity demanded}}{\% \Delta \text{price}} \right|$$

measures how responsive quantity de-



Inelastic versus elastic demand: a small versus large response to a price change

- $E_d > 1$: **elastic** (quantity responds strongly);
- $E_d < 1$: **inelastic** (quantity barely responds);
- $E_d = 1$: **unit elastic**.

Demand is more elastic when there are more **substitutes**, the good is a **luxury**, it takes a large share of income, or buyers have **more time** to adjust. Elasticity links to **total revenue** ($P \times Q$): if demand is elastic, a price cut *raises* revenue; if inelastic, a price cut *lowers* it. Along a straight-line demand curve, elasticity falls as you move down (elastic on top, inelastic at the bottom).


Worked example. A price rise from \$10 to \$12 (a +20% change) cuts quantity demanded from 100 to 85 (a -15% change). Then $E_d = \left| \frac{-15\%}{20\%} \right| = 0.75 < 1$, so demand is **inelastic**. Because it is inelastic, the price rise *raises* total revenue: from $10 \times 100 = \$1000$ to $12 \times 85 = \$1020$.

Price Elasticity of Supply

Price elasticity of supply

, $E_s = \frac{\% \Delta Q_s}{\% \Delta P}$, measures how responsive

sellers are. Supply is more elastic when firms can adjust output easily – spare capacity, storable goods, and above all **more time** (supply is more elastic in the long run than the short run).




The same price rise gives a small supply response when PES is below 1, a large one above 1

Other Elasticities

- **Cross-price elasticity** $= \frac{\% \Delta Q_x}{\% \Delta P_y}$: **positive** for substitutes, **negative** for complements – the sign tells you the relationship.
- **Income elasticity** $= \frac{\% \Delta Q}{\% \Delta \text{income}}$: **positive** for normal goods (and > 1 for luxuries), **negative** for inferior goods.

Worked example. When the price of coffee rises 10%, the quantity of tea demanded rises 6%: cross-price elasticity $= \frac{+6\%}{+10\%} = +0.6$. The **positive** sign confirms coffee and tea are **substitutes**. If instead a 5% rise in incomes raised restaurant-meal demand 10%, income elasticity $= \frac{10\%}{5\%} = +2$ – a normal good, and a **luxury** since it exceeds 1.




Cross elasticity: positive for substitutes, negative for complements

Exam skill: you are expected to compute an elasticity, classify it, and read a *sign* as evidence of substitute/complement or normal/inferior.

Market Equilibrium, Consumer Surplus, and Producer Surplus

Equilibrium is where supply meets demand: the **equilibrium price** clears the market, so quantity demanded equals quantity supplied and there is no shortage or surplus.



Consumer surplus and producer surplus at the market equilibrium


- **Consumer surplus** is the area **below the demand curve and above the price** – the gap between what buyers would pay and what they do pay.
- **Producer surplus** is the area **above the supply curve and below the price** – the gap between the price and sellers' minimum acceptable price.
- Their sum is **total surplus**, and the competitive equilibrium **maximizes** it – this is why free markets are called **allocatively efficient**.



A real market: many buyers and sellers meet, and price settles where the quantity wanted equals the quantity offered

Market Disequilibrium and Changes in Equilibrium

Away from equilibrium the market self-corrects: above the equilibrium price there is a **surplus** (excess supply) that pushes price down; below it a **shortage** (excess demand) that pushes price up.




Above equilibrium there is excess supply; below it, excess demand

When a determinant shifts a curve, the equilibrium moves predictably. A **single** shift gives a definite result for both price and quantity. When **both** curves shift, one of price or quantity is **indeterminate** (depends on the relative sizes of the shifts) – a favorite exam trap. Work these out by drawing the shifts.

The Effects of Government Intervention in Markets

Governments intervene, usually reducing total surplus:



A price ceiling causes a shortage; a price floor causes a surplus


- A **price ceiling** (a legal maximum, like rent control) below equilibrium causes a persistent **shortage**.
- A **price floor** (a legal minimum, like a minimum wage) above equilibrium causes a persistent **surplus**.
- An **excise tax** shifts supply up by the tax; the burden (**tax incidence**) falls more heavily on whichever side is **more inelastic**.
- A **subsidy** shifts supply down, raising quantity.

Each of these creates **deadweight loss** – mutually beneficial trades that no longer happen, shown as a triangle between the curves at the new quantity.

Exam skill: on a graph, be able to shade consumer surplus, producer surplus, government revenue from a tax, and deadweight loss, and say who bears the tax and why.

International Trade and Public Policy

When a country opens to trade at a **world price** :



A tariff raises domestic supply, cuts demand, and shrinks imports

- If the world price is **below** the domestic price, the country **imports** – consumers gain more than producers lose, so total surplus rises.
- If the world price is **above** the domestic price, the country **exports** – producers gain more than consumers lose.

A **tariff** (a tax on imports) or a **quota** (a limit on imports) raises the domestic price, helps domestic producers, hurts consumers more, and creates deadweight loss. Free trade maximizes total surplus, which is the core efficiency argument for it.

Exam tips

- Separate a **movement along** a curve (own price) from a **shift** (a determinant); a single shift gives a definite result, a double shift leaves one variable indeterminate.
- Compute elasticity $E_d = |\% \Delta Q \div \% \Delta P|$; elastic (> 1) means a price cut raises total revenue, inelastic (< 1) lowers it.
- Read the sign of cross-price elasticity (substitutes $+$, complements $-$) and income elasticity (normal $+$, inferior $-$).

- Shade **consumer/producer surplus** and the **deadweight loss** from a tax or price control; incidence falls on the more **inelastic** side.
- A price ceiling below equilibrium causes a shortage; a price floor above causes a surplus.

Production, Cost, and the Perfect Competition Model

AP Microeconomics

The Production Function



A steel plant: production functions and costs shape how firms supply goods

The **production function** links inputs to output. In the **short run** at least one input (usually **capital**) is **fixed**; in the **long run** all inputs vary.

- **Total product** is the total output.
- **Marginal product** (MP) is the extra output from one more unit of a variable input (usually labor): $MP = \frac{\Delta \text{ output}}{\Delta \text{ labor}}$.
- **Average product** is output per worker.


As you add workers to fixed capital, MP eventually falls – the **law of diminishing marginal returns**. This shape of MP is what drives the shape of the cost curves below.



The production function turns inputs — labour and capital like this robot — into output

Short-Run Production Costs

Costs split into **fixed** and **variable**:



The short-run marginal, average, and average variable cost curves

- **Fixed cost** (FC) does not change with output (rent on the factory).
- **Variable cost** (VC) rises with output (materials, labor).
- **Total cost** $TC = FC + VC$.

The **per-unit** costs matter most:


- **Average fixed cost** $AFC = FC/Q$ falls continuously (spreading fixed cost).
- **Average variable cost** $AVC = VC/Q$ and **average total cost** $ATC = TC/Q$ are **U-shaped**.
- **Marginal cost** $MC = \frac{\Delta TC}{\Delta Q}$ is the cost of one more unit.

Two key facts: **MC is the mirror image of MP** (when MP rises, MC falls, and vice versa), and **MC cuts both AVC and ATC at their minimum points** (when marginal is below average, the average falls; when above, it rises).

Worked example. A firm has fixed cost $FC = \$100$. At $Q = 10$ its total cost is \$300, so variable cost is \$200: then $ATC = \frac{300}{10} = \$30$, $AVC = \frac{200}{10} = \$20$, and $AFC = \frac{100}{10} = \$10$ (note $AFC + AVC = ATC$). If raising output to $Q = 11$ pushes total cost to \$325, the marginal cost of that 11th unit is $\frac{\Delta TC}{\Delta Q} = \frac{325 - 300}{1} = \25 .

Long-Run Production Costs

In the long run all inputs – including plant size – can change, so only the **long-run average total cost (LRATC)** curve matters. Its shape shows **returns to scale**:



Economies and diseconomies of scale on the long-run average cost curve

- **Economies of scale** : LRATC falls as output grows (bulk buying, specialization).
- **Constant returns to scale**: LRATC flat.
- **Diseconomies of scale** : LRATC rises (a firm too big to manage well).

Exam skill: know that “diminishing returns” is a **short-run** idea (one fixed input) while “economies of scale” is a **long-run** idea (all inputs vary) – examiners test the distinction.

Types of Profit

Economists count **all** opportunity costs, including the owner’s:

- **Explicit costs** are out-of-pocket payments; **implicit costs** are the opportunity costs of owner-supplied resources (forgone salary, forgone interest).

- **Accounting profit** = revenue – explicit costs.
- **Economic profit** = revenue – explicit – implicit costs.

Because economic profit subtracts more, it is smaller. **Zero economic profit** (a **normal profit**) still means the owner is earning exactly the opportunity cost of their resources – a perfectly acceptable outcome.

Worked example. A shopkeeper takes in \$200,000 of revenue and pays \$120,000 in explicit costs, so **accounting profit** is \$80,000. But she gave up a \$70,000 salary elsewhere and \$5,000 of interest on the savings she invested – \$75,000 of **implicit costs**. Her **economic profit** is $80,000 - 75,000 = \$5,000$: still positive, so staying in business genuinely beats her next-best alternative.

Profit Maximization

Every firm, in every market structure, maximizes profit at the quantity where **marginal revenue equals marginal cost**:

$$MR = MC.$$

Produce each unit whose MR exceeds its MC; stop when they are equal. Producing less leaves profitable units unmade; producing more adds units that cost more than they earn.

Deciding to Produce, Enter, or Exit


Even at $MR = MC$, a firm must check whether to operate:

- **Short run – the shutdown rule:** keep producing only if price covers **average variable cost** ($P \geq AVC$). If $P < AVC$, shut down and lose only fixed cost. So the firm's short-run supply curve is its **MC curve above minimum AVC**.
- **Profit check:** compare P to ATC at the profit-maximizing quantity – $P > ATC$ means economic profit, $P < ATC$ means a loss, $P = ATC$ means break-even.
- **Long run – entry and exit:** economic profits attract **entry**; losses cause **exit**. This continues until economic profit is zero.

Perfect Competition

A **perfectly competitive** market has many small firms, an identical (homogeneous) product, free entry and exit, and perfect information. Each firm is a **price taker** – too small to affect the price – so it faces a **horizontal** demand curve at the market price, and **price = marginal revenue = average revenue** ($P = MR = AR$).


This is why economists prize perfect competition. In **long-run equilibrium** it is both **allocatively efficient** and **productively efficient**. Allocative efficiency holds because the firm produces where $P = MC$: the price (society's marginal benefit from the last unit) equals its marginal cost, so exactly the right amount is made. Productive efficiency holds because free entry drives price down to the minimum of average total cost, so goods are made as cheaply as possible. The market price is itself a **signal** that communicates everyone's marginal costs and benefits.



short-run $P > ATC$ attracts entry $\Rightarrow S$ right $\Rightarrow P \rightarrow \min ATC$ (long-run zero profit)

The market sets the price; the price-taking firm produces where $MC = P$, here with short-run profit

The side-by-side panel is the graph the exam wants: the **market** (left) fixes the price P^* ; the **firm** (right) takes it as a horizontal line and produces where $MC = P$. When $P > ATC$ there is a short-run profit (the shaded box); that profit attracts entry, which shifts market supply right and pushes the price down.



A perfectly competitive firm makes only normal profit in the long run

- **Short run:** the firm can earn profit, break even, or take a loss, producing where $P = MC$.
- **Long run:** entry and exit drive the firm to **zero economic profit**, producing where $P = MC = ATC$ at **minimum ATC**. This double efficiency – **productive efficiency** (lowest-cost output) and **allocative efficiency** ($P = MC$, so the last unit's value equals its cost) – is the benchmark against which every other market structure is judged.

Exam skill: be able to draw the side-by-side market and single-firm graphs, show a short-run profit or loss, then show how entry/exit restores long-run equilibrium at minimum ATC.

Exam tips

- Know the cost curves: AFC falls, ATC/AVC are U-shaped, and **MC cuts ATC and AVC at their minimums**.
- Every firm maximises profit at **MR = MC**; check the shutdown rule ($P \geq AVC$ in the short run).
- Distinguish **accounting profit** (revenue – explicit costs) from **economic profit** (also minus implicit/opportunity costs); zero economic profit is normal.
- “Diminishing returns” is short-run (one fixed input); “economies of scale” is long-run (all inputs vary).

- In perfect competition $P = MR = AR$; entry/exit drive long-run economic profit to zero at minimum ATC.

Imperfect Competition

AP Microeconomics

Imperfect Competition and Market Power



A car assembly line: market power and firm scale differ under imperfect competition

Most real markets are **imperfectly competitive** : firms have some **market power** – the ability to set price rather than take it. The source of that power is a **downward-sloping demand curve** facing the firm, which means that to sell one more unit the firm must **lower the price on all units**.




Market structures range from perfect competition to monopoly

Because of this, **marginal revenue is below price** ($MR < P$) for every firm with market power. On a straight-line demand curve the MR curve has the same intercept but **twice the slope**, lying below demand. This one fact – $MR < P$ – separates every imperfect market from perfect competition and explains why they produce **less** and charge **more**.

Worked example. A firm with market power faces demand $P = 100 - Q$. Selling 1 unit gives $P = \$99$ and total revenue \$99; selling a 2nd forces the price down to \$98 on **both** units, so revenue rises to $2 \times 98 = \$196$. The marginal revenue of that 2nd unit is $196 - 99 = \$97$ – below its \$98 price. In general $MR = 100 - 2Q$, falling twice as fast as demand.

Monopoly

A **monopoly** is a single seller of a product with no close substitutes, protected by **barriers to entry** (economies of scale – a **natural monopoly** – control of a key resource, patents, or government license).



A monopoly maximises profit where $MC = MR$ and can earn supernormal profit

- The monopolist maximizes profit at $MR = MC$, then charges the **highest price the demand curve allows** at that quantity (read price up to the demand curve, not the MR curve).
- Compared with perfect competition it produces **less** and charges **more**, earning long-run economic profit because entry is blocked.
- This under-production creates **deadweight loss** – the monopoly is **allocatively inefficient** because $P > MC$: the last unit is worth more than it costs, yet is not made.

Governments may respond with **antitrust** law, or regulate a natural

monopoly's price toward $P = MC$ (**socially optimal**, but may cause a loss) or $P = ATC$ (**fair-return**, break-even).

Exam skill: the monopoly graph is a staple – show quantity at $MR = MC$, price up on the demand curve, and shade profit (between P and ATC) and deadweight loss (the triangle to the competitive quantity).



A large utility plant: monopoly and market power mean a single firm can set price above marginal cost

Price Discrimination

Price discrimination is charging different buyers different prices for the same good. It requires **market power**, the ability to **separate** buyers by willingness to pay, and the ability to **prevent resale**.

Under **perfect (first-degree) price discrimination**, the firm charges each buyer their maximum price. The result: the firm produces the **efficient** quantity (where $P = MC$, no deadweight loss), but converts **all** consumer surplus into **producer surplus** (profit). Real examples – student discounts, airline fares, coupons – are partial versions.

Monopolistic Competition

Monopolistic competition has **many** firms selling **differentiated** products with **easy entry and exit** (restaurants, salons, clothing brands). Product differentiation gives each firm a **downward-sloping** demand curve and thus some price-setting power.

- **Short run:** like a monopoly – produce at $MR = MC$, and profit or loss is possible.
- **Long run:** easy entry competes away profit, so each firm earns **zero economic profit**, producing where demand is **tangent** to ATC .
- Because it still charges $P > MC$ and does **not** produce at minimum ATC , it has **excess capacity** and is neither allocatively nor productively efficient. The trade-off society accepts is **product variety** and choice.

Oligopoly and Game Theory

An **oligopoly** is a market with a **few** large, **interdependent** firms (cars, airlines, phones). Because each firm's best move depends on rivals' moves, we analyze them with **game theory**.

- A **payoff matrix** shows each firm's profit for every combination of choices.
- A **dominant strategy** is a firm's best choice **regardless** of what the rival does.
- A **Nash equilibrium** is an outcome where **no** firm can do better by unilaterally changing its choice.
- The classic **prisoners' dilemma** shows why firms often fail to cooperate: both would gain by **colluding** (acting like a monopoly), but each has an incentive to cheat, so they end up at the competitive-ish, lower-profit outcome. **Collusion** and **cartels** are unstable for the same reason (and usually illegal).

Worked example. Two firms each choose a **High** or **Low** price. Profits (A, B): both High $\rightarrow (10, 10)$; A Low, B High $\rightarrow (14, 4)$; A High, B Low $\rightarrow (4, 14)$; both Low $\rightarrow (6, 6)$. For A: if B plays High, Low (14) beats High (10); if B plays Low, Low (6) beats High (4) – so **Low is A's dominant strategy**, and by symmetry B's too. The Nash equilibrium is **both Low** (6, 6), worse for both than the cooperative (10, 10) – the prisoners' dilemma in action.

Factor Markets

AP Microeconomics

Introduction to Factor Markets



A garment factory: factor markets determine wages for labour and other inputs

A **factor market** is where firms buy the **factors of production** – labor, land, capital. The roles reverse: **firms demand** and **households supply**. The most-tested factor is **labor**.

		Firm B		(A, B) payoffs
		High price	Low price	
Firm A	High price	10 / 10	4 / 14	
	Low price	4 / 14	6 / 6	

Nash equilibrium

Each firm's dominant strategy is Low — so both end at (6, 6), worse than (10, 10).


Nash: each plays a best response given the other · here Low is dominant for both

A prisoners' dilemma: each firm's dominant strategy is Low, giving the Nash equilibrium

Exam skill: be able to find each firm's dominant strategy and the Nash equilibrium from a 2×2 payoff matrix, and explain why the cooperative outcome is hard to sustain.

Exam tips

- A firm with market power faces a downward-sloping demand, so $MR < P$ (MR has twice the slope).
- A monopoly produces where $MR = MC$, then charges the price **up on the demand curve**; it makes deadweight loss because $P > MC$.
- Under perfect price discrimination the firm produces the efficient quantity but captures all consumer surplus.
- Monopolistic competition earns zero long-run profit with excess capacity (product variety is the trade-off).
- In an oligopoly find the **dominant strategy** and **Nash equilibrium** from the payoff matrix; the cooperative outcome is hard to sustain.



A competitive labour market sets the wage where labour demand meets labour supply

Factor demand is a **derived demand** – firms want labor not for itself but for the output (and revenue) it produces. So the demand for workers depends on the demand for the product they make.

Changes in Factor Demand and Factor Supply

Labor demand shifts when:

- the **product price** or product demand changes (more valuable output $\Rightarrow$ more labor demanded),
- **worker productivity** changes (better training or capital),
- the **price of other inputs** changes (a substitute input like machines, or a complement).

Labor supply shifts with the **number of qualified workers**, **immigration**, workers' **preferences** (wages needed to attract them), and the wages available in **other jobs**. The market wage is set where labor demand meets labor supply.

Profit-Maximizing Behavior in Perfectly Competitive Factor Markets

A firm decides how many workers to hire with the same **marginal logic** as everything else – compare the marginal benefit and marginal cost of one more worker.

- The marginal benefit of a worker is the **marginal revenue product** (MRP): the extra revenue from one more worker, $MRP = MP \times$ product price (in a competitive product market, where $P = MR$). MRP falls as more workers are hired, because of **diminishing marginal product**.
- The marginal cost of a worker is the **wage** (in a competitive labor market the firm is a wage **taker**, so marginal factor cost = the wage).
- **Hiring rule:** hire workers up to where $MRP =$ wage (more generally, $MRP = MFC$). The MRP curve is the firm's labor **demand** curve.

Worked example. A worker's marginal product is 8 units and the firm sells each in a competitive market for \$5, so $MRP = 8 \times \$5 = \40 . If the wage is \$30, the firm hires this worker because $MRP >$ wage. As more workers join, diminishing marginal product pulls MRP down; the firm keeps hiring until a worker whose $MP = 6$ gives $MRP = 6 \times \$5 = \30 , exactly the wage – the profit-maximizing stopping point.

To choose the **least-cost combination** of two inputs, set the **MP-per-dollar equal** across inputs: $\frac{MP_L}{P_L} = \frac{MP_K}{P_K}$ – the same equal-bang-per-buck rule as consumer choice.

Exam skill: compute MRP from a table (marginal product $\times$ price) and find how many workers to hire at a given wage – a common free-response task.




A car assembly line: firms hire labour until the marginal revenue product equals the wage

Monopsonistic Markets

A **monopsony** is a market with a **single buyer** of a factor – for example, the only large employer in a town.

- Because it is the whole market, the monopsonist faces an **upward-sloping labor supply** curve: to hire more workers it must raise the wage **for everyone**.
- So the **marginal factor cost** (MFC) is **above** the wage (like $MR < P$ for a monopoly, in reverse), and the MFC curve lies above the supply curve.
- The monopsonist hires where $MRP = MFC$, then pays the **lower wage** the supply curve allows at that quantity.
- Result: compared with a competitive labor market, a monopsony hires **fewer** workers at a **lower** wage – the mirror image of monopoly.



A monopsony hires where $MFC = MRP$ and pays a lower wage than a competitive market

Exam skill: draw the monopsony graph (supply, MFC above it, MRP as demand), mark quantity at $MRP = MFC$ and the wage down on the supply curve, and contrast it with the competitive wage and quantity.

Exam tips

- Factor demand is **derived demand**; hire labour up to where **MRP = wage** (MRP = marginal product $\times$ price).
- MRP falls as more workers are hired (diminishing marginal product).
- For least-cost input use, equalise the **marginal product per dollar** across inputs.
- A **monopsony** faces an upward supply, so its marginal factor cost is **above** the wage; it hires fewer workers at a lower wage.
- Draw the monopsony graph (supply, MFC above it, MRP as demand) and contrast with the competitive outcome.

Market Failure and the Role of Government

AP Microeconomics

Socially Efficient and Inefficient Market Outcomes

A market outcome is **socially efficient** when it maximizes total surplus – when the **marginal social benefit** equals the **marginal social cost** of the last unit. Perfect competition reaches this on its own. But sometimes markets fail to, producing too much or too little. This is **market failure**, and it is the main economic justification for government action.

productive efficiency
*making goods at the
lowest possible cost*

allocative efficiency
*making the goods
people most want*


Productive efficiency is lowest cost; allocative efficiency is making what people want

Externalities



Cigarettes: externalities mean private markets may not reach socially efficient outcomes

An **externality** is a cost or benefit that falls on a **third party** not involved in the transaction – so the private decision-makers ignore it.



A negative externality: the market over-produces, causing a welfare loss

- A **negative externality** (pollution) means **marginal social cost** > **marginal private cost**. The market **over-produces**; the deadweight loss is a triangle. Fixes: a **Pigovian tax** or regulation to make producers “internalize” the cost (shifting private cost up to social cost).
- A **positive externality** (vaccination, education) means **marginal social benefit** > **marginal private benefit**. The market **under-produces**. Fixes: a **subsidy** to consumers or producers.
- **Coase theorem** : if property rights are clear and bargaining is costless, private parties can solve an externality themselves – but that is rare with many people involved.

Worked example. A factory’s private marginal cost of a chemical is \$8 per unit, but each unit also dumps \$3 of pollution on neighbors, so the marginal **social** cost is \$11. Ignoring the \$3, the market over-produces. A **Pigovian tax of exactly \$3 per unit** raises the firm’s private cost to \$11 – now equal to the social cost – so the firm cuts output to the efficient quantity and the externality is “internalized.”

Exam skill: identify whether an externality is positive or negative, show on a graph that private and social curves diverge, mark the deadweight loss, and name the correcting tax or subsidy.



Industrial pollution: a negative externality where social cost exceeds private cost

Public and Private Goods

Goods differ along two traits: **rivalry** (does one person’s use reduce another’s?) and **excludability** (can non-payers be kept out?).

	Excludable	Non-excludable
Rival	Private good food, clothes	Common resource ocean fish, air
Non-rival	Club good cinema, toll road	Public good street light, defence

Goods classified by whether they are rival and excludable

- A **private good** is rival and excludable (a sandwich) – markets handle these well.

- A **public good** is **non-rival and non-excludable** (national defense, a lighthouse). Because non-payers cannot be excluded, people **free-ride** – enjoy it without paying – so markets **under-provide** it, and government usually supplies it.
- **Common resources** are rival but non-excludable (ocean fish); over-use ruins them – the **tragedy of the commons**.

The Effects of Government Intervention in Different Market Structures

Government tools change outcomes differently by structure:


- **Regulating a monopoly's price** toward $P = MC$ can *increase* efficiency (unlike a price control in a competitive market, which creates shortage/surplus).
- **Antitrust laws** break up or block anti-competitive mergers and collusion.
- **Taxes and subsidies** shift output toward the socially efficient quantity when externalities are present.

The key insight: the **same** policy can help or hurt efficiency depending on whether the market was already efficient (competitive) or already failing (monopoly, externality).

Per-unit vs lump-sum. A **per-unit tax** adds to marginal cost, so it shifts supply up, raises price, cuts quantity, and creates deadweight loss. A **lump-sum tax** (a fixed charge like a licence fee) does **not** change marginal cost or marginal benefit, so in the **short run** the profit-maximising price and quantity are **unchanged** – only the firm's profit falls; in the long run it changes entry and exit. For a **natural monopoly**, regulating price down to $P = MC$ (allocatively efficient) forces a loss, so a **lump-sum subsidy** is needed to keep it producing that efficient quantity.

Inequality

Markets reward productivity, but the resulting **income distribution** can be very unequal. Economists measure it with the **Lorenz curve** (share of income against share of population; the further it bows from the 45° line, the more unequal) and the **Gini coefficient** (0 = perfect equality, 1 = perfect inequality).



The Lorenz curve and Gini coefficient measure income inequality

Worked example. Suppose the poorest 20% of households earn just 5% of income, while the richest 20% earn 50%. Under perfect equality each 20% would earn exactly 20%, so the Lorenz curve would be the 45° line. Because the bottom fifth is well below 20% and the top fifth well above, the curve bows far beneath the diagonal – a large gap that means a **high Gini coefficient** and substantial inequality.

Sources of inequality include differences in **human capital** (skills, education), inheritance, and discrimination. Governments redistribute through **progressive taxes** (higher rates on higher incomes) and **transfer payments** (welfare, subsidies). There is a debated **efficiency–equity trade-off**

: redistribution can reduce work incentives, so societies balance fairness against the size of the total pie.



Wealth and poverty side by side: markets alone may leave large income inequality unaddressed

Exam tips

- A **negative externality** makes the market over-produce; correct with a **Pigovian tax** equal to the external cost. A positive externality under-produces; correct with a subsidy.
- A **public good** is non-rival and non-excludable, so free-riding leads markets to under-provide it.
- Classify goods by rivalry and excludability; common resources (rival, non-excludable) suffer the tragedy of the commons.
- The same policy can help or hurt efficiency depending on whether the market was already efficient or failing.
- Measure inequality with the **Lorenz curve** and **Gini coefficient** (0 = equality, 1 = inequality).