

AP MICROECONOMICS

Unit 4: Imperfect Competition

Topics 4.1 – 4.5 | Lesson Reference Guide

AP Exam Weight: 22-25% of multiple-choice questions

What This Section Is About

Unit 4 steps away from the perfect-competition benchmark of Unit 3 and into the market structures that look like the real world: industries where firms have the power to set prices. Every imperfectly competitive firm faces a downward-sloping demand curve, which means marginal revenue sits below price. Set $MR = MC$, and the resulting price exceeds marginal cost — the condition that generates deadweight loss and defines the inefficiency of imperfect competition. This guide walks through all four non-PC structures: pure monopoly (one firm, blocked entry), price-discriminating monopoly (one firm, several prices), monopolistic competition (many firms, differentiated products, free entry), and oligopoly (a few interdependent firms, analyzed with game theory).

Topics covered:

- 4.1 Introduction to Imperfect Competition – four market structures classified by number of firms and product type
- 4.2 Monopoly – barriers to entry, $MR < P$, profit maximization, deadweight loss, natural monopoly regulation
- 4.3 Price Discrimination – three conditions, elastic/inelastic sorting, perfect PD and zero deadweight loss
- 4.4 Monopolistic Competition – short-run profit/loss, free-entry tangency to ATC, excess capacity
- 4.5 Oligopoly and Game Theory – payoff matrices, dominant strategies, Nash equilibrium, prisoner's dilemma, tacit collusion

◆ Numeric Models Used in This Guide

Topic 4.2 Monopoly (De Beers): $P = 1000 - 50Q$, $MR = 1000 - 100Q$, $MC = \$200$. At $Q = 9$, $P = \$550$; at $Q = 10$, $P = \$500$ — quantity effect +\$500, price effect -\$450, so $MR = \$50$. TR peaks at \$5,000 when $MR = 0$ ($Q = 10$).

Topic 4.2 Welfare: Competitive $Q_C = 100$ at $P_C = \$10$; monopoly $Q_M = 50$ at $P_M = \$20$. DWL triangle = $\frac{1}{2} \times 50 \times \$10 = \$250$.

Topic 4.3 Price Discrimination (Air Sunshine): $MC = ATC = \$200/\text{seat}$. 2,000 business travelers (WTP \$650) + 2,000 students (WTP \$250). Single-price profit \$900,000; two-price PD profit \$1,000,000.

Topic 4.5 Duopoly (dictionary apps): (Collude, Collude) → \$18M each; (Cheat, Cheat) → \$10M each; one-side cheat → cheater \$20M, colluder \$0. Nash equilibrium = (Cheat, Cheat).

TOPIC 4.1 — INTRODUCTION TO IMPERFECT COMPETITION

COLLEGE BOARD CED · ENDURING UNDERSTANDING IMC-1

Market power — the ability of a firm to influence the price it charges — generates the four classic market structures. In every structure other than perfect competition, price exceeds marginal cost at the profit-maximizing quantity, which produces allocative inefficiency.

Learning Objective IMC-1.A

Describe the principal types of imperfectly competitive markets — monopoly, oligopoly, and monopolistic competition.

Explain why a firm facing a downward-sloping demand curve must lower its price to sell additional units.

Explain why markets are allocatively inefficient when price exceeds marginal cost.

Essential Knowledge

IMC-1.A.1 Market power: a firm's ability to influence the price it charges.

IMC-1.A.2 Market structure is defined by the number of firms in an industry and whether their products are identical or differentiated.

IMC-1.A.3 Four market structures: perfect competition, monopolistic competition, oligopoly, and monopoly.

IMC-1.A.4 Perfectly competitive firms set $P = MC$ and are allocatively efficient. In imperfect competition, $P > MC$ at $Q^* \rightarrow$ under-production and deadweight loss.

IMC-1.A.5 Concentration ratio: the share of total industry sales accounted for by the largest N firms (commonly 4-firm or 8-firm).

From Perfect Competition to the Real World

Unit 3 ended with perfect competition — the benchmark market structure in which every firm is a price taker, products are identical, and long-run economic profit is zero. But most industries you encounter day to day look nothing like that. Verizon faces only three serious rivals. Coca-Cola and Pepsi are near-identical substitutes yet are clearly different products. The food court outside has a dozen vendors, each selling something slightly distinct. These are all forms of imperfect competition, and the key distinguishing feature is market power: the ability of a firm to influence the price at which it sells.

Market structure is determined by two variables: the number of firms in the industry and whether their products are identical (standardized/homogeneous) or differentiated. The two variables generate four combinations — the four market structures the AP exam asks about: perfect competition, monopolistic competition, oligopoly, and monopoly.

FIGURE 4.1-1 Types of Market Structure




Figure 4.1-1 — Types of Market Structure. The four structures emerge from crossing two questions: how many firms compete (One / Few / Many), and are the products differentiated? One firm $\rightarrow$ monopoly. Few firms $\rightarrow$ oligopoly (either way). Many firms $\rightarrow$ perfect competition (identical) or monopolistic competition (differentiated).

The Four Structures Side by Side

A monopoly is the extreme case of market power: a single firm is the only producer of a good with no close substitutes. When De Beers controlled virtually all of the world's diamond production in the 20th century, it was a monopolist. An oligopoly is an industry with only a small number of firms. No single firm has a monopoly, but each firm's output decisions noticeably affect market price, and firms must anticipate how rivals will respond. The world banana market, in which three firms produce 65% of output, is an oligopoly. Monopolistic competition combines features of monopoly and perfect competition: many firms compete, each sells a differentiated product, and there is free entry and exit in the long run. The mall food court is the textbook example — each vendor has some monopoly power from its distinctive menu but faces constant competition from the vendor two doors down.

The concentration ratio measures how dominated an industry is by its largest firms. A four-firm concentration ratio of 85%, for instance, means the top four firms account for 85% of sales. Higher concentration generally indicates oligopoly or monopoly; low concentration suggests monopolistic or perfect competition.

Characteristic	Perfect Competition	Monopolistic Competition	Oligopoly	Monopoly
Number of firms	Many	Many	Few	One
Product type	Identical	Differentiated	Identical or diff.	Unique
Entry barriers	None	None (free entry)	High	Very high
Price setting	Price-taker	Some power	Some power	Price-maker
Long-run profit	Zero	Zero	Positive possible	Positive
Allocatively efficient?	Yes ($P = MC$)	No ($P > MC$)	No ($P > MC$)	No ($P > MC$)

Why $P > MC$ Matters

The pivotal economic consequence of imperfect competition is the inequality that defines the rest of this unit: $P > MC$. Because an imperfectly competitive firm faces a downward-sloping demand curve, it must lower price to sell more units — which pulls marginal revenue below price. Setting $MR = MC$ to maximize profit therefore produces a price that exceeds marginal cost. The result is allocative inefficiency: units for which willingness to pay exceeds the cost of production are not produced, so deadweight loss appears. Allocative efficiency requires $P = MC$ on the last unit sold; whenever $P > MC$, there are buyers willing to pay more than it costs to produce the next unit — but it isn't produced. Those missed trades are the DWL of imperfect competition.

♪ **Industry examples by structure (AP canon)**

Perfect competition: agricultural commodities (wheat, corn, soybeans), foreign-currency trading, wholesale fish auctions — many sellers, standardized product.

Monopolistic competition: restaurants and fast food, clothing brands (Levi's, Gap, Uniqlo, H&M), hair salons, independent coffee shops, gas stations — many sellers, differentiated by menu, style, location, or quality.

Oligopoly: commercial aircraft (Boeing/Airbus duopoly), smartphones & operating systems (Apple/Samsung/Google), soft drinks (Coke/Pepsi), auto manufacturers, US airlines, wireless carriers, major banks — few firms, high barriers.

Monopoly: local public utilities (water, gas, electricity distribution), US Postal Service for first-class mail, patented pharmaceuticals during the patent window, De Beers historically for diamonds — single producer, very high barriers.

★ **AP Exam Tips — Topic 4.1**

Learn the “two-variable” grid, not memorized lists. Any FRQ asking you to classify an industry wants you to use number of firms & product type to pick one of the four. The industry name alone is not enough.

The $P > MC$ inequality unifies the whole unit. Every structure in Unit 4 (monopoly, oligopoly, monopolistic competition) is allocatively inefficient for exactly this reason. If you can state why $P > MC$ and why that creates DWL, you can answer the “why inefficient” question in any of them.

Concentration ratio = top-N firms’ share of industry sales. A high CR_4 (say 80%+) screams oligopoly or monopoly; a low CR_4 (say 20%) says MC or PC.

Key Terms — Topic 4.1

Market power	A firm's ability to influence the price at which it sells; absent in perfect competition, present in every imperfectly competitive structure.
Imperfect competition	Any market structure other than perfect competition — monopoly, oligopoly, or monopolistic competition. Firms face downward-sloping demand curves.
Monopoly	An industry with a single producer of a good that has no close substitutes, protected by barriers to entry.
Oligopoly	An industry dominated by a small number of firms whose output decisions are interdependent; may sell identical or differentiated products.
Monopolistic competition	An industry with many firms selling differentiated products with free entry and exit in the long run.
Concentration ratio	Share of total industry sales accounted for by the largest N firms; CR_4 and CR_8 are most common. Higher ratios signal more market power.
Allocative efficiency	Condition in which the quantity produced equates marginal benefit (price) with marginal cost. $P = MC \rightarrow$ efficient; $P > MC \rightarrow$ inefficient.
Deadweight loss (DWL)	The reduction in total surplus caused by under-production relative to the efficient quantity; shown as a triangle between $MR = MC$ output and $P = MC$ output.

TOPIC 4.2 — MONOPOLY

COLLEGE BOARD CED · ENDURING UNDERSTANDING IMC-2

A monopoly is the only producer of a good with no close substitutes, sustained by barriers to entry. A monopolist faces the entire downward-sloping market demand curve, so to sell more it must lower price — making marginal revenue less than price. Profit maximization at $MR = MC$ produces lower output, higher price, positive long-run profit, and deadweight loss relative to perfect competition.

Learning Objective IMC-2.A

Identify the types of barriers to entry that allow monopolies to persist.

Use a graph to show how a profit-maximizing monopolist chooses price and quantity.

Explain why monopolists charge a price that exceeds marginal cost.

Calculate consumer surplus, producer surplus, profit, and deadweight loss for a monopoly.

Essential Knowledge

IMC-2.A.1 A monopoly is the sole producer of a good with no close substitutes.

IMC-2.A.2 Barriers to entry: control of resources, economies of scale (natural monopoly), technological superiority (incl. network externalities), government-created barriers (patents, copyrights).

IMC-2.A.3 The monopolist's demand curve is the entire market demand curve — downward-sloping.

IMC-2.A.4 To sell more, the monopolist must lower price; this makes $MR < P$ at every quantity.

IMC-2.A.5 Profit maximization: produce where $MR = MC$; find price on the demand curve directly above that quantity.

IMC-2.A.6 Monopoly profit = $(P_M - ATC_M) \times Q_M$.

IMC-2.A.7 Monopoly creates deadweight loss because $Q_M < Q_{SE}$ (where $P = MC$).

Barriers to Entry

A monopoly is an industry controlled by a single producer selling a good with no close substitutes. For a monopoly to persist, something must prevent other firms from entering — a barrier to entry. The AP curriculum names four types.

Control of a scarce resource or input. If one firm owns essentially the entire supply of a key input, no one else can produce the good. De Beers achieved this with diamond mines in the 20th century.

Economies of scale. In some industries, average cost falls across the entire range of market demand. One large firm can produce the industry's output at a lower average cost than two or more smaller firms could. This is a natural monopoly — and most public utilities (electricity distribution, water, gas, cable internet) are examples because the fixed cost of laying infrastructure is enormous.

Technological superiority. A firm with technology competitors cannot replicate can dominate a market. In high-tech industries this often combines with network externalities — the value of the good to one user rises as more users adopt it. Microsoft Windows, Facebook, and WeChat all benefit from network externalities.

Government-created barriers. A patent grants an inventor exclusive rights to produce or license an invention for a limited period (typically 20 years in the US), creating a legal temporary monopoly. A copyright similarly grants exclusive rights to a literary or artistic work. Both exist to incentivize innovation by letting the creator capture returns from their work.

Why $MR < P$ for a Monopolist

The demand curve a monopolist faces is downward-sloping — it is literally the entire market demand curve. To sell more units, the monopolist must lower price. This is the critical difference

from perfect competition. A perfectly competitive firm faces a horizontal demand curve at the market price, so $MR = P$. A monopolist sees two effects every time it sells one more unit: the quantity effect (one more unit sold at the new, lower price) plus the price effect (every previously sold unit now sells at the lower price). Marginal revenue is the sum of these two effects, and because the price effect is negative, $MR < P$ at every quantity.

$$\text{MR} = \text{quantity effect} + \text{price effect} \quad \text{For a linear demand } P = a - bQ: \text{MR} = a - 2bQ$$

Consider De Beers facing a demand of $Q = 20 - P/50$ (equivalently $P = 1000 - 50Q$). At $Q = 9$, price = \$550 (point A), so total revenue is \$4,950. To sell a 10th diamond, price must drop to \$500 for all 10 diamonds (point B). The quantity effect adds \$500 (the 10th unit at \$500); the price effect subtracts \$450 (9 original units $\times$ \$50 price cut). Net marginal revenue is \$50 — far below the price of \$500. Mathematically, if demand is $P = a - bQ$, then $MR = a - 2bQ$, so MR has twice the slope of demand.

Price of diamond P	Quantity demanded Q	Total revenue TR = P $\times$ Q	Marginal revenue MR = $\Delta TR / \Delta Q$
\$1,000	0	\$0	—
\$950	1	\$950	\$950
\$900	2	\$1,800	\$850
\$850	3	\$2,550	\$750
\$800	4	\$3,200	\$650
\$750	5	\$3,750	\$550
\$700	6	\$4,200	\$450
\$650	7	\$4,550	\$350
\$600	8	\$4,800	\$250
\$550	9	\$4,950	\$150
\$500	10	\$5,000	\$50
\$450	11	\$4,950	–\$50
\$400	12	\$4,800	–\$150
\$350	13	\$4,550	–\$250
\$300	14	\$4,200	–\$350
\$250	15	\$3,750	–\$450
\$200	16	\$3,200	–\$550
\$150	17	\$2,550	–\$650
\$100	18	\$1,800	–\$750
\$50	19	\$950	–\$850
\$0	20	\$0	–\$950

Table 4.2-1 — Demand, Total Revenue, and Marginal Revenue for De Beers Diamond Monopoly. Highlighted rows ($Q = 9$, $Q = 10$) are the quantity-effect/price-effect move walked through above and visualized in Figure 4.2-2(a). TR peaks at \$5,000 when $Q = 10$ ($MR = 0$) and declines thereafter — the single-peaked TR curve of Figure 4.2-2(b).




Figure 4.2-2(a) — Demand and marginal revenue for De Beers. Moving from point A ($Q = 9$, $P = \$550$) to point B ($Q = 10$, $P = \$500$), the quantity effect (green) adds \$500 and the price effect (orange) subtracts \$450. Marginal revenue of the 10th diamond is \$50 (point C) — far below the \$500 price.




Figure 4.2-2(b) — Total revenue for De Beers. TR rises while the quantity effect dominates the price effect ($MR > 0$), reaches its maximum at \$5,000 when $Q = 10$ and $MR = 0$, then falls as the price effect dominates quantity effect ($MR < 0$). This single-peaked TR curve is the direct consequence of MR being linear and downward-sloping.

Profit Maximization — The Monopoly “Four-Step”

The profit-maximizing rule $MR = MC$ applies to a monopolist just as it does to a perfect competitor — but with one twist. Because $MR < P$ for a monopolist, the price it charges is not the MC value found at the intersection. The procedure is a four-step: (1) find the quantity Q_M where MR intersects MC (point A); (2) slide straight up to the demand curve to find P_M (point B); (3) read ATC at Q_M as height C; (4) the profit rectangle has width Q_M and height $(P_M - ATC_M)$.

$$\text{Monopoly profit} = (P_M - ATC_M) \times Q_M$$

Positive when $P_M > ATC_M$ · Zero when $P_M = ATC_M$ · Negative (loss) when $P_M < ATC_M$

The result: compared to a perfectly competitive industry producing the same product with the same cost structure, the monopolist (1) produces less ($Q_M < Q_C$), (2) charges more ($P_M > P_C$), and (3) earns positive economic profit that persists in the long run because entry is blocked.

The Simplest Textbook Case — Flat MC = ATC = \$200

Before the general U-shaped-ATC picture, work through the simplest textbook case: no fixed costs and marginal cost constant at \$200 per diamond, so $MC = ATC = \$200$. The flat cost line keeps the diagram clean so you can focus on demand and MR. Everything the general monopoly story requires is still here: $MR < P$, $MR = MC$ determines Q_M , P_M comes off demand, and the P - MC gap is the welfare cost of market power.

FIGURE 4.2-3 · The Monopolist's Profit-Maximizing Output and Price




Figure 4.2-3 — The Monopolist's Profit-Maximizing Output and Price. Flat $MC = ATC = \$200$. Monopolist produces $Q_M = 8$ at A ($MR = MC$), charges $P_M = \$600$ at B on the demand curve, earns green profit rectangle = \$3,200. Point C ($Q = 16$, $P = \$200$) is where a PC industry would produce.

Reading the three points

Point A ($Q = 8$). MR crosses MC at \$200. Up to the 8th diamond, $MR > MC$ — produce. Past the 8th, $MR < MC$ — stop. A is the exact balance point — the monopolist's profit-maximizing quantity.

Point B ($P = \$600$). The price the monopolist actually charges. Common exam error: students read price off the MR curve at $Q = 8$ (which gives \$200). Do not do this. Price is on the demand curve — slide straight up from A until you hit D. That height is \$600.

Point C ($Q = 16$, $P = \$200$). The perfectly competitive outcome, where D crosses MC. A PC industry would produce until $P = MC$. All 8 diamonds between $Q = 8$ and $Q = 16$ are units buyers value more than they cost to make, but under monopoly they are not produced. Those missed trades are the deadweight loss of monopoly.

Monopoly profit (green rectangle) = width $Q_M \times$ height $(P_M - ATC_M) = 8 \times (\$600 - \$200) = \$3,200$. Entry is blocked — profit persists in the long run.

Monopoly vs. PC at a glance: same costs, same demand. Monopoly produces half as much (8 vs. 16), charges three times the price (\$600 vs. \$200), and earns \$3,200 of profit that would be zero under PC. The green rectangle is a transfer from consumers to the monopolist; the deadweight loss is a separate loss that nobody captures.

FIGURE 4.2-4 The Monopolist's Profit




Figure 4.2-4 — The monopolist's profit. The monopolist produces Q_M where $MR = MC$ (point A), charges P_M from the demand curve (point B), and incurs ATC_M (point C). The green rectangle, $(P_M - ATC_M) \times Q_M$, is economic profit. Because entry is blocked, this profit persists in the long run.

Why Monopoly Is Allocatively Inefficient

Monopoly is not just a transfer of surplus from consumers to the monopolist — it is a net loss for society. Panel (a) of Figure 4.2-5 shows the perfectly competitive outcome: price $P_C = \$10$, quantity $Q_C = 100$, consumer surplus is the full triangle under demand above P_C , producer surplus is the triangle between MC and P_C . Panel (b) shows the monopoly outcome: $Q_M = 50$ is smaller, $P_M = \$20$ is higher. Consumer surplus (blue) shrinks dramatically. The green rectangle (and the triangle beneath it) is producer surplus — most of it is pure profit. The yellow triangle between Q_M and Q_C , bounded above by demand and below by MC , is deadweight loss: trades that would have happened under perfect competition but don't under monopoly — buyers whose willingness to pay exceeded MC but who no longer get to buy at the higher P_M .

FIGURE 4.2-5 Monopoly Causes Inefficiency




Figure 4.2-5 — Monopoly causes inefficiency. Panel (a): under perfect competition, output is $Q_C = 100$ and price is $P_C = \$10$; the entire blue triangle is consumer surplus (no DWL). Panel (b): the monopoly cuts output to $Q_M = 50$ and raises price to $P_M = \$20$; CS_M shrinks, PS_M (green) is captured as profit, and the yellow triangle ($DWL = \frac{1}{2} \times 50 \times \$10 = \$250$) is deadweight loss.

Natural Monopoly and Price Regulation

For most monopolies, the policy response is antitrust — prevent monopoly formation, break up existing monopolies, or regulate conduct. For natural monopolies (where high fixed costs make a single producer the lowest-cost option), breaking the firm into competitors would just raise average costs for everyone. The two main policy responses are public ownership (the government owns the firm, like USPS or many foreign utilities), and price regulation — the regulator caps the price.

The textbook regulated price is where $P = ATC$ (so-called “fair-return” pricing): the firm earns zero economic profit while quantity expands beyond Q_M . The socially efficient regulated price is $P = MC$, but that would force the firm into losses whenever $ATC > MC$ — a problem for any natural monopoly whose ATC is still falling at the quantity that would clear the market at $P = MC$.

FIGURE 4.2-6 Unregulated and Regulated Natural Monopoly




Figure 4.2-6 — Unregulated and regulated natural monopoly. Panel (a): unregulated, the natural monopolist produces Q_M at P_M , earning profit (green) but creating deadweight loss (yellow). Panel (b): under “fair-return” price regulation, the regulator sets $P_R = ATC$ where the ATC curve crosses demand — output expands to Q_R , CS expands, and the firm earns zero economic profit.

★ AP Exam Tips — FRQ Graphing Protocol for Monopoly

A graded monopoly graph needs: (1) correctly labeled axes (“Price, cost” on y-axis; “Quantity” on x-axis), (2) downward-sloping D , (3) MR below D with twice the slope and the same y-intercept, (4) U-shaped ATC and swoosh- or U-shaped MC , (5) MC through min ATC , (6) Q_M at $MR = MC$, (7) P_M on the demand curve directly above Q_M . Dot and label every intersection.

Profit: shade the rectangle between P_M and ATC_M out to Q_M . Loss: shade the rectangle between ATC_M (above) and P_M (below) out to Q_M . DWL: shade the triangle from (Q_M, P_M) down to (Q_M, MC) and over to the $P = MC$ quantity.

Linear demand trick: if D is $P = a - bQ$, then $MR = a - 2bQ$. Same y-intercept as D ; twice the slope. Draw MR with its x-intercept at exactly half of D 's x-intercept.

Common FRQ trap: at Q_M , students read P off the MR curve instead of D — that's where $MR = MC$, not where price lies. P_M is always on the demand curve.

Key Terms — Topic 4.2

Monopoly	An industry with a single producer of a good that has no close substitutes, protected by barriers to entry.
Barrier to entry	Any factor that prevents new firms from entering an industry — control of resources, economies of scale, technological superiority / network externalities, or government-created barriers (patents, copyrights, licenses).
Natural monopoly	An industry in which a single firm's average cost is lower than any smaller firm's would be, because economies of scale extend across the entire range of market demand.
Network externality	The condition where the value of a good to one user rises as more people use it (e.g., social networks, operating systems); can be a source of monopoly power.
Patent / Copyright	Government-granted temporary legal monopolies on inventions (patent, ~20 years) or original creative works (copyright); intended to incentivize innovation.
Marginal revenue (MR)	The additional total revenue from selling one more unit. For a monopolist facing linear demand $P = a - bQ$, $MR = a - 2bQ$, so $MR < P$ for every $Q > 0$.
Quantity effect	The positive contribution to marginal revenue from selling one additional unit at the new (lower) price.
Price effect	The negative contribution to marginal revenue from having to lower the price on all previously sold units in order to sell one more; it is why $MR < P$ for a monopolist.
Monopoly profit	$(P_M - ATC_M) \times Q_M$, read off the monopoly graph as a rectangle with width Q_M and height $(P_M - ATC_M)$. Positive and persistent because entry is blocked.
Deadweight loss of monopoly	The reduction in total surplus caused by under-production; the triangle between Q_M and Q_C bounded by demand (above) and MC (below).
Fair-return pricing	A natural-monopoly price regulation that sets $P = ATC$. The firm earns zero economic profit, output expands beyond Q_M , and consumer surplus rises; CS gain is partial but DWL is not fully eliminated.

TOPIC 4.3 — PRICE DISCRIMINATION

COLLEGE BOARD CED · ENDURING UNDERSTANDING IMC-3

A firm with market power can earn additional profit by charging different prices to different buyers for the same good. Price discrimination requires (i) market power, (ii) the ability to distinguish buyers by willingness to pay, and (iii) the ability to prevent resale. In the limit, perfect price discrimination eliminates consumer surplus and deadweight loss — the entire area between demand and MC becomes profit.

Learning Objective IMC-3.A

Define price discrimination and distinguish it from a single-price monopoly.

Explain why price discrimination allows a monopolist to capture additional consumer surplus.

Show how perfect price discrimination eliminates deadweight loss.

Essential Knowledge

IMC-3.A.1 A single-price monopolist charges every buyer the same price.

IMC-3.A.2 A firm engages in price discrimination when it charges different prices to different buyers for the same good.

IMC-3.A.3 Price discrimination requires: (1) market power, (2) ability to distinguish buyers by willingness to pay, and (3) ability to prevent resale.

IMC-3.A.4 Buyers with more inelastic demand are charged higher prices; more elastic buyers get lower prices.

IMC-3.A.5 Under perfect price discrimination, each buyer is charged exactly their willingness to pay — consumer surplus = 0 and DWL = 0 (all of it becomes profit).

Three Conditions — and How Firms Satisfy Them

A single-price monopolist charges the same price to every customer, as the graphs in Topic 4.2 assumed. But many monopolists (and many non-monopolists with some market power) can do better by charging different prices to different buyers for the same product — the practice known as price discrimination. Three conditions must hold for a firm to price-discriminate successfully: (1) it must have some market power (price-takers cannot discriminate), (2) it must be able to distinguish buyers who have different willingness to pay, and (3) it must be able to prevent resale from low-price buyers to high-price buyers.

Consider Air Sunshine, which serves 2,000 business travelers (willing to pay up to \$650 per ticket) and 2,000 students (willing to pay up to \$250). Marginal cost per seat is \$200. If Air Sunshine must charge a single price, it chooses between \$650 (serves only business travelers, earns $2,000 \times \$450 = \$900,000$) and \$250 (serves everyone, earns $4,000 \times \$50 = \$200,000$). It picks \$650 and leaves the student market unserved. If it can distinguish the two groups — say, by requiring student ID or by selling tickets through different channels — it can charge \$650 to business travelers and \$250 to students, earning $\$900,000 + \$100,000 = \$1,000,000$. The student market is served, and total surplus rises.

In practice, airlines rarely charge different fares by group label (which is often illegal). Instead they design pricing mechanisms that sort customers by willingness to pay. The AP exam names three common techniques.

Advance-purchase restrictions. Tickets bought weeks in advance cost less than tickets bought the day before departure. Business travelers (less price-sensitive, book last-minute) pay more; leisure travelers (more price-sensitive, book early) pay less.

Volume discounts. Lower per-unit price for buyers purchasing in bulk. Buyers of large quantities tend to be more price-sensitive than one-off shoppers.

Two-part tariffs. A fixed membership or access fee plus a per-unit price (Costco's annual membership plus item prices). The fixed fee captures consumer surplus from heavy users without

pricing out light users.

The general principle: more price-elastic buyers get lower prices; less elastic buyers get higher prices. This matches the inverse elasticity rule of monopoly pricing — markups should be larger where elasticity is smaller.




Figure 4.3-1(a,b) — Air Sunshine, two groups of buyers. Charging \$650 to business travelers and \$250 to students, Air Sunshine captures profit from both groups: solid green area $B = 2,000 \times (\$650 - \$200) = \$900,000$ from business travelers; hatched area $S = 2,000 \times (\$250 - \$200) = \$100,000$ from students. Total profit = \$1,000,000, versus \$900,000 under a single \$650 price. $MC = ATC = \$200$.

Perfect Price Discrimination — The Efficient Extreme

The extreme case is perfect price discrimination: the monopolist charges each buyer exactly their willingness to pay — not a penny more and not a penny less. The firm then produces every unit for which willingness to pay $\geq MC$ — exactly the socially efficient quantity. There is no consumer surplus (every buyer pays their maximum) and no deadweight loss (every mutually beneficial trade occurs). The entire area between the demand curve and MC becomes profit.

This is a famous exam trap: perfect price discrimination is allocatively efficient ($Q = Q_C$, no DWL) even though it is also maximally unequal (all surplus goes to the firm). Students often check “efficient = fair.” Be careful: price discrimination transfers surplus; it doesn't destroy it.




Figure 4.3-1(c) — Perfect price discrimination. Each consumer is charged their exact willingness to pay. The firm's profit becomes the entire green triangle between the demand curve and $MC = ATC$. Consumer surplus = 0 and deadweight loss = 0 — the firm captures every mutually beneficial trade. Output extends to Q where $P = MC$ (the perfectly competitive quantity).

★ **AP Exam Tips — Topic 4.3**

If asked to compare single-price monopoly vs. perfect PD: perfect PD has higher output (= competitive Q), zero CS, zero DWL, and higher profit. The industry is allocatively efficient but all surplus is captured by the firm.

The three conditions must all appear when an FRQ asks “what is required for price discrimination?”: (1) market power, (2) separate buyer groups by willingness to pay, (3) prevent resale. Missing any one and the strategy collapses.

Match the direction to elasticity: more elastic $\rightarrow$ lower price; less elastic $\rightarrow$ higher price. Students sometimes write this backwards — remember the firm charges the least to the buyers who are most ready to walk away.

Perfect PD vs. perfect competition: same output and same $P-MC$ equality on the last unit, but opposite surplus split (all CS $\rightarrow$ producer under PD).

Key Terms — Topic 4.3

Single-price monopolist	A monopoly that charges the same price to every buyer. The default assumption in the Topic 4.2 graphs.
Price discrimination	Charging different prices to different buyers for the same good. Requires market power, buyer sorting, and no resale.
Advance-purchase pricing	Selling tickets or units at a lower price when bought far in advance; sorts price-sensitive early planners from price-insensitive last-minute buyers.
Volume discount	A lower per-unit price when a buyer purchases in bulk; sorts large (more elastic) buyers from small (less elastic) buyers.
Two-part tariff	A pricing scheme with a fixed access/membership fee plus a per-unit price (e.g., Costco membership + item prices); captures surplus from heavy users without pricing out light users.
Perfect price discrimination	The extreme case where every buyer is charged exactly their willingness to pay; $CS = 0$, $DWL = 0$, output = competitive Q , and the entire area between D and MC becomes profit.
Inverse elasticity rule	The result that profit-maximizing markups are higher in markets with less elastic demand; the engine behind every PD scheme.

TOPIC 4.4 — MONOPOLISTIC COMPETITION

COLLEGE BOARD CED · ENDURING UNDERSTANDING IMC-4

Monopolistic competition features many firms, differentiated products, and free entry/exit in the long run. Each firm faces its own downward-sloping demand curve (so $MR < P$), but entry drives long-run profit to zero at the tangency of demand with ATC. At that tangency $P > MC$ (allocatively inefficient) and Q is below the minimum-ATC quantity (excess capacity).

Learning Objective IMC-4.A

Determine price and profit for a monopolistically competitive firm in both the short run and the long run.

Explain how free entry drives economic profit to zero in long-run equilibrium.

Explain why monopolistic competition leads to excess capacity and allocative inefficiency.

Discuss advertising and product differentiation as competitive tools.

Essential Knowledge

IMC-4.A.1 Monopolistic competition: many firms, differentiated products, free entry and exit.

IMC-4.A.2 Each firm faces a downward-sloping demand curve $\rightarrow MR < P \rightarrow$ profit max at $MR = MC$.

IMC-4.A.3 Short run: a firm may earn positive, zero, or negative economic profit.

IMC-4.A.4 Long run: entry (if profit) or exit (if loss) shifts each firm's D until D is tangent to ATC at the profit-maximizing quantity $\rightarrow$ zero economic profit.

IMC-4.A.5 At the long-run equilibrium: $P > MC$ (allocatively inefficient) and $Q < \text{min-ATC quantity}$ (excess capacity).

IMC-4.A.6 Firms differentiate by style/type, location, and quality, and compete through advertising and brand names.

The Short Run — Looks Like a Mini-Monopoly

Monopolistic competition is the hybrid market structure in which many firms sell differentiated versions of a similar product with free entry and exit. The food court, hair salons, independent coffee shops, and gas stations are all canonical examples. Each firm has some monopoly power — your haircut from stylist A is not literally the same as your haircut from stylist B — so each firm faces a downward-sloping demand curve. But because there are many near-substitute options, that demand curve is relatively elastic, and new firms can enter easily.

In the short run, a monopolistically competitive firm behaves exactly like a mini-monopoly: it finds Q where $MR = MC$, charges P on its demand curve above Q , and compares P to ATC to determine profit or loss. If the firm's demand D_p sits above ATC at the profit-maximizing quantity, it earns a positive economic profit (green rectangle). If the firm's demand D_u lies entirely below ATC — meaning no price yields non-negative profit — it operates at a loss (orange rectangle), minimizing the loss by producing at $MR = MC$ where loss is smallest.

FIGURE 4.4-1 The Monopolistically Competitive Firm in the Short Run




Figure 4.4-1 — The monopolistically competitive firm in the short run. Panel (a): a profitable firm chooses the profit-maximizing quantity Q_P where $MR_P = MC$, charges P_P on D_P , and earns the green profit rectangle $(P_P - ATC_P) \times Q_P$. Panel (b): an unprofitable firm faces ATC above D_U at every quantity; it minimizes the loss (orange rectangle) at Q_U where $MR_U = MC$.

The Long Run — Entry Drives Profit to Zero

The long run is where monopolistic competition differs sharply from pure monopoly. Because entry is free, positive economic profit attracts new firms. New entrants poach customers from existing firms, which shifts each existing firm's demand curve to the left (and MR with it). Symmetrically, losses cause exit — each remaining firm's demand shifts rightward. Entry (or exit) continues until the typical firm's demand curve has shifted to be tangent to its ATC curve at the profit-maximizing quantity. At that tangency, $P_{MC} = ATC_{MC}$, so economic profit = 0.

FIGURE 4.4-2 Entry and Exit Shift Existing Firms' Demand Curves and Marginal Revenue Curves




Figure 4.4-2 — Entry and exit shift each firm's demand. Panel (a): when existing firms are profitable, new firms enter and each existing firm's demand (and MR) shifts leftward as customers are divided among more sellers. Panel (b): when existing firms are losing money, exit shifts each remaining firm's demand (and MR) rightward. The process continues until economic profit = 0 — the tangency of D with ATC .

Allocative Inefficiency and Excess Capacity

Long-run equilibrium under monopolistic competition therefore has two defining features that distinguish it from both monopoly and perfect competition.

First, zero economic profit. Free entry and exit force the long-run tangency of D with ATC . Firms in monopolistic competition cannot earn positive economic profit in the long run, just like firms in perfect competition.

Second, allocative inefficiency and excess capacity. Because the firm faces a downward-sloping D , the tangency with ATC must occur on the left (downward-sloping) portion of ATC , not at the minimum. So at the profit-maximizing quantity Q_{MC} , two things are true: (i) $P > MC$ (allocatively inefficient — same $P > MC$ story as monopoly), and (ii) $Q_{MC} < \text{the min-ATC quantity}$, which is called excess capacity. The firm is operating at a point where producing more would lower its average cost.

FIGURE 4.4-4 Comparing Long-Run Equilibrium in Perfect Competition and Monopolistic Competition




Figure 4.4-4 — Long-run equilibrium: perfect competition vs. monopolistic competition. Panel (a) PC: the firm produces at the minimum-cost output Q_{PC} , where $P_{PC} = MC_{PC} = ATC_{PC}$. Panel (b) MC: the firm produces Q_{MC} where D_{MC} is tangent to ATC — $P_{MC} = ATC_{MC}$ (zero profit) but $P_{MC} > MC_{MC}$, and Q_{MC} is below the minimum-cost output. The gap is excess capacity.

Why Firms Differentiate — and Advertise

Why do firms in monopolistic competition try so hard to differentiate their products? Because differentiation is what preserves the downward slope of their demand curve — the only thing that gives them any pricing power. The AP curriculum identifies three forms of product differentiation: by style or type (a Thai food stall and a burger stall at the food court sell imperfect substitutes), by location (the gas station just off the freeway exit charges more than one two miles into town), and by quality (consumers may be willing to pay more for a product they believe is of higher quality, whether or not it objectively is).

Firms reinforce these differences with advertising and brand names. Advertising serves two economic functions: direct information (announcing a product's features) and indirect signaling (a large ad budget signals that the firm is confident enough in its product to spend heavily). Brand names create recognition that reduces consumer search costs — but also allow firms to charge brand premiums over generic substitutes.

Is monopolistic competition a problem? Economists are split. On one hand, excess capacity and $P > MC$ create deadweight loss just like monopoly does. On the other hand, consumers value

variety — the fact that there are 15 types of toothpaste rather than one has real welfare benefits that the standard surplus analysis ignores. The AP exam position is that long-run monopolistic competition is allocatively inefficient ($P > MC$) and productively inefficient ($Q < \min ATC$) — but may still be the best market structure available when products are naturally varied.

★ AP Exam Tips — Topic 4.4

Two long-run facts memorized together: (i) $P = ATC$ at Q^* (zero profit), (ii) $P > MC$ at Q^* ($P > MC$ still holds $\rightarrow$ DWL), and (iii) $Q^* < \min\text{-}ATC$ quantity (excess capacity). The AP exam often asks you to label all three on a graph.

Tangency is on the downward-sloping portion of ATC. If a student draws the tangency at the ATC minimum, they've drawn perfect competition by mistake. The tangency must occur before ATC hits its minimum.

MC vs. PC long-run comparison (likely FRQ): both earn zero profit, both face free entry; but PC has $P = MC = \min ATC$, while MC has $P > MC$ and $Q < \min ATC$. Any FRQ asking “how does MC differ from PC?” expects you to name these two inefficiencies.

Advertising is a competitive tool, not a market failure — unless the AP question is specifically about asymmetric information. Keep it in-unit: differentiation preserves the downward slope of D.

Key Terms — Topic 4.4

Monopolistic competition	Many firms selling differentiated products with free entry and exit; each firm faces a downward-sloping (but relatively elastic) demand curve.
Product differentiation	Any strategy that makes a firm's product imperfectly substitutable for rivals' products — by style/type, location, or quality.
Short-run equilibrium (MC firm)	Each firm sets Q where $MR = MC$; may earn positive, zero, or negative economic profit depending on where ATC sits at Q^* .
Long-run equilibrium (MC firm)	Entry/exit shifts each existing firm's demand until D is tangent to ATC at Q^* . $P = ATC$ (zero profit), $P > MC$ (inefficient), $Q < \min\text{-}ATC$ (excess capacity).
Excess capacity	The gap between a monopolistically competitive firm's long-run Q^* and the quantity that would minimize its ATC; a sign that the firm could lower its average cost by producing more.
Advertising	Spending to inform consumers of a product's features (direct) and to signal the firm's confidence in the product (indirect); supports and reinforces differentiation.
Brand name	A product identifier that creates recognition, reduces consumer search cost, and often permits a price premium over generic substitutes.

TOPIC 4.5 — OLIGOPOLY AND GAME THEORY

COLLEGE BOARD CED · ENDURING UNDERSTANDING IMC-5

An oligopoly is a market dominated by a small number of interdependent firms. Each firm's profit depends on rivals' actions, so firms behave strategically. Game theory tools — payoff matrices, dominant strategies, Nash equilibrium, and the prisoner's dilemma — predict both one-shot outcomes and, through repeated play (tit-for-tat), the tacit collusion that sustains above-competitive profits.

Learning Objective IMC-5.A

Explain how oligopolists are interdependent and why they have incentives to collude.

Analyze two-player games using payoff matrices.

Identify dominant strategies and Nash equilibria.

Explain how repeated interaction can sustain tacit collusion.

Essential Knowledge

IMC-5.A.1 Oligopolists are interdependent: each firm's profit depends on rivals' actions.

IMC-5.A.2 Collusion: firms coordinate to raise joint profit. A cartel is a formal collusive agreement.

IMC-5.A.3 Noncooperative behavior: each firm acts in its own self-interest.

IMC-5.A.4 A payoff matrix shows every player's payoff for every combination of actions.

IMC-5.A.5 A dominant strategy is a player's best action regardless of what others do.

IMC-5.A.6 A Nash equilibrium is an outcome where no player can improve their payoff by changing strategy, given the others' choices. A game may have 0, 1, or multiple NE.

IMC-5.A.7 The prisoner's dilemma: a game in which each player has a dominant strategy to defect, but both playing their dominant strategy yields a worse outcome than mutual cooperation would.

IMC-5.A.8 Tit-for-tat in repeated games can sustain tacit collusion by rewarding cooperation and punishing defection.

Interdependence — The Defining Feature

An oligopoly is an industry dominated by a small number of firms. The defining feature is interdependence: because each firm's output or pricing decision noticeably affects the market, each firm must anticipate how the others will react. In perfect competition, no firm is big enough to matter. In monopoly, there is no rival to worry about. In oligopoly, every move invites a counter-move — and that is why oligopoly analysis uses game theory: the formal study of strategic decisions under interdependence.

The simplest case is a duopoly — two firms. Consider two developers selling essentially identical dictionary apps. If they collude and both charge the monopoly price of \$6, they jointly earn \$36M (each firm \$18M). If one cheats and undercuts to \$5, the cheating firm captures the whole market and the loyal firm earns nothing. If both cheat down to \$5, they split the enlarged market but earn far less than they would have under collusion. The four outcomes form a payoff matrix with exactly the structure of the prisoner's dilemma.

The Prisoner's Dilemma — and Why Oligopolists Face It

FIGURE 4.5-2 The Prisoner's Dilemma

		Justin	
		Don't confess	Confess
Richard	Don't confess	Justin gets 5-year sentence. Richard gets 5-year sentence.	Justin gets 2-year sentence. Richard gets 20-year sentence.
	Confess	Justin gets 20-year sentence. Richard gets 2-year sentence.	Justin gets 15-year sentence. Richard gets 15-year sentence.

Figure 4.5-2 — The Prisoner's Dilemma. If neither confesses, each gets 5 years. If one confesses while the other stays silent, the confessor gets 2 years and the silent one gets 20. If both confess, each gets 15 years. “Confess” is a dominant strategy for each suspect, so the Nash equilibrium is (Confess, Confess) → 15 years each — strictly worse than (Don't confess, Don't confess) → 5 years each.

The logic: each suspect reasons, “If the other stays silent, confessing cuts my sentence from 5 years to 2. If the other confesses, confessing cuts my sentence from 20 years to 15. Either way, confess.” Confess is a dominant strategy — strictly better regardless of what the other does. Both confess; both get 15 years. The outcome is a Nash equilibrium: given the other's choice to confess, neither can improve by changing. Yet both would prefer the (Don't confess, Don't confess) outcome. That is the prisoner's dilemma: dominant-strategy play is worse for both players than mutual cooperation.

The duopolist's situation is structurally identical. “Collude” plays the role of “Don't confess”; “Cheat” plays the role of “Confess.” The one-shot Nash equilibrium is both firms cheating. For the dictionary-app duopoly: if a rival colludes at \$6, I earn \$20M by cheating vs. \$18M by colluding. If a rival cheats at \$5, I earn \$10M by cheating vs. \$0 by sticking with \$6. Cheat is dominant for both firms → Nash equilibrium (Cheat, Cheat) at (\$10M, \$10M), strictly worse than (Collude, Collude) at (\$18M, \$18M).

Not Every Game Is a Prisoner's Dilemma

Nash equilibria can exist without any dominant strategies, and some games have multiple Nash equilibria. In the game of avoidance below, two chain stores prefer to locate in different towns than compete in the same one. There is no dominant strategy — the best choice depends on what the rival does. There are two Nash equilibria: whichever chain commits first typically picks its preferred town; the other then picks the other town.

FIGURE 4.5-3 A Game of Avoidance




Figure 4.5-3 — A game of avoidance. Each chain earns more when it is alone in a town. The two Nash equilibria (circled) are (Chain 1 in Ayville, Chain 2 in Beeville) → \$6M / \$4M, and (Chain 1 in Beeville, Chain 2 in Ayville) → \$5M / \$8M. No dominant strategy — each firm's best choice depends on the rival's.

Repeated Interaction and Tacit Collusion

The prisoner's dilemma logic suggests oligopolists should compete away all their joint profit. In practice, they often don't. Why? Because real oligopolists interact repeatedly, not once. The strategy of tit-for-tat — cooperate on the first round, then mimic whatever the rival did last round — rewards sustained cooperation and immediately punishes defection. If both firms expect the other to play tit-for-tat, both can rationally sustain cooperation indefinitely. This is called tacit collusion: firms coordinate on the monopoly outcome without ever meeting to agree.




Figure 4.5-6 — Tit-for-tat vs. always cheat (repeated game). If both firms play tit-for-tat, each earns \$180M every year — mutual cooperation sustained. If both play always cheat, each earns \$160M every year — the one-shot Nash. The one-period defection payoff (\$200M in year 1) is overwhelmed by the stream of lower payoffs thereafter, so tit-for-tat sustains tacit collusion.

Several forces push against tacit collusion: (1) antitrust laws make explicit collusion illegal and tacit collusion risky; (2) the more firms in the industry, the harder it is for all of them to coordinate; (3) complex products and pricing make it hard to tell whether a rival is cheating; and

(4) differences in firm interests (different cost structures, different market shares) make the “fair” cooperative price hard to agree on.

★ AP Exam Tips — Topic 4.5

Finding a Nash equilibrium — the quick procedure. For each cell, ask: given the column player's choice, could the row player do strictly better by switching rows? If no, mark that cell for the row player. Repeat for the column player. Any cell marked by both players is a Nash equilibrium.

Dominant strategy vs. Nash. A dominant strategy is the same action in every response row/column. A Nash equilibrium is a cell that's a mutual best response. Every dominant-strategy outcome is a Nash equilibrium, but not every Nash equilibrium has a dominant strategy behind it.

Prisoner's dilemma shape check: (i) each player has a dominant strategy to defect; (ii) the (defect, defect) cell is strictly worse for both than (cooperate, cooperate). If both conditions hold, it's a PD. If only one, it's just a one-shot game.

Why tacit collusion survives: repeated interaction + credible threat of punishment. Tit-for-tat is the AP-canonical trigger strategy. If an FRQ asks “why don't oligopolists compete away all profit?”, this is the answer.

Key Terms — Topic 4.5

Oligopoly	An industry dominated by a small number of firms whose decisions are interdependent.
Duopoly	The simplest oligopoly: exactly two firms.
Interdependence	The feature of oligopoly in which each firm's profit depends on the decisions of its rivals; motivates strategic reasoning and game theory.
Game theory	The formal study of strategic decisions by rational agents who anticipate each other's actions; the analytical framework for oligopoly.
Payoff matrix	A table listing each player's payoff for every combination of players' strategy choices.
Dominant strategy	A strategy that gives a player a higher payoff than any alternative, regardless of what the other player does.
Nash equilibrium	An outcome in which each player's choice is a best response to the others' choices; no player can improve by unilaterally changing. A game may have zero, one, or multiple NE.
Prisoner's dilemma	A game in which each player has a dominant strategy to defect, but mutual defection yields a worse outcome for both players than mutual cooperation would.
Collusion / Cartel	Coordination among firms to restrict output and raise price toward the monopoly outcome. A cartel is a formal collusive agreement; collusion is often illegal under antitrust law.
Tacit collusion	Coordination on cooperative pricing without any explicit agreement; sustained in repeated interaction by strategies like tit-for-tat.
Tit-for-tat	A repeated-game strategy in which a player cooperates in the first round, then mimics the rival's previous action; rewards cooperation and punishes defection, sustaining tacit collusion.

UNIT 4 QUICK REFERENCE SUMMARY

Topic	Key Idea	Core Graph / Outcome	What P & Q Look Like
4.1 Introduction to Imperfect Competition	Market power = ability to influence price. Four structures sorted by (# firms) × (identical vs. differentiated).	Four-cell matrix: PC, MC, Oligopoly, Monopoly. In every structure other than PC, $P > MC$.	PC: $P = MC$. All others: $P > MC \rightarrow$ under-production and DWL.
4.2 Monopoly	Single producer, blocked entry. $MR < P$ because selling more requires lowering price on all units.	Q_M where $MR = MC$; P_M on D above Q_M . Profit rectangle = $(P_M - ATC_M) \times Q_M$.	$Q_M < Q_C$, $P_M > P_C$, long-run profit > 0 , DWL triangle between Q_M and Q_C .
4.3 Price Discrimination	Charge different prices to different buyers. Needs market power + buyer sorting + no resale.	Perfect PD: profit = entire area between D and MC. CS = 0, DWL = 0.	Output expands from Q_M toward Q_C ; all CS captured as profit; allocatively efficient but unequal.
4.4 Monopolistic Competition	Many firms, differentiated products, free entry $\rightarrow$ long-run tangency of D with ATC.	Long-run: D tangent to ATC at Q^* ; $P = ATC$ (zero profit) but $P > MC$ and $Q < \text{min-ATC } Q$.	Zero LR economic profit (like PC) but $P > MC$ and excess capacity — allocatively and productively inefficient.
4.5 Oligopoly & Game Theory	Interdependent few. Payoff matrix $\rightarrow$ dominant strategies, Nash equilibrium, prisoner's dilemma. Repeated play $\rightarrow$ tacit collusion via tit-for-tat.	2×2 payoff matrix: find dominant strategies, mark best responses, intersect for NE.	One-shot NE often (Cheat, Cheat) — PD outcome. Repeated games can sustain (Collude, Collude) $\rightarrow$ near-monopoly P and Q.

End of Unit 4 Lesson Reference Guide. Pair with the Unit 4 Learning Hub ([unit4-apmicro.html](#)) and the Unit 4 PowerPoint for complete coverage.