Pricing and Strategic Interaction

Day 4

Dr. Ernil Sabaj

Warwick Summer School
International Business and Finance

24 July 2026

Where we are

Mon 20 July	**Day 1**	The multinational firm: theories, frameworks, entry
Tue 21 July	**Day 2**	Countries: institutions, political risk, national advantage
Thu 23 July	**Day 3**	Strategy, organisation and market structure
Fri 24 July	**Day 4**	Pricing and strategic interaction

The final day. Days 1–3 built the firm, its environment, and its strategy. Today: how firms *set prices*, and how they *interact* when each move changes the others' payoff.

Transfer pricing picks up Tuesday's trapped-cash thread; the cola war closes the loop this afternoon's practical opens.

Reading

- Chapters 11–12, Collinson, Narula, Rugman & Qamar, *International Business*, 9th ed., Pearson, 2024 (production and marketing strategy).
- Lipczynski, Goddard & Wilson, *Industrial Organization*, 5th ed., Pearson, 2017 (pricing, ch. 5; game theory and collusion, chs. 6–10).
- Sloman et al., *Economics for Business*, 9th ed., Pearson, 2024 (price discrimination; oligopoly).

The pricing and game-theory mechanics are IO; the international overlay — transfer pricing, global rivalry — is Collinson.

Part I

Pricing and price discrimination

Pricing in Practice

Is there a unique equilibrium price? In most cases, no.

- **Competition**: monopoly and collusive oligopoly price very differently from competitive markets.
- **Information**: firms often barely know their own price elasticity of demand (the % fall in quantity when price rises 1%), let alone rivals'.
- **Aims**: profit, sales, growth — often several at once. **Life cycle**: launch pricing $\neq$ mature pricing.
- And plain **past practice**.

Theory gives the frontier; these frictions decide where on it firms actually sit.

Cost Plus Pricing Formula

- In a highly influential study, Hall and Hitch (1939) report the results of interviews with the managers of 38 businesses, 30 of whom reported the use of some form of cost plus pricing formula.

Under cost plus pricing, the firm calculates or estimates its AVC (average variable cost), and then sets its price by adding a percentage markup that includes a contribution towards the firm's fixed costs, and a profit margin:

$$\text{Price} = \text{AVC} + \%\text{markup}$$

or

$$P = (1 + m)\text{AVC}$$

where P denotes price, and the markup (expressed as a percentage) is $100 \times m$ percent.

How are prices determined? Bank of England survey

	Not important	Slightly important	Important	Very important
Price is primarily determined by your competitors' price	4.9	16.3	35.6	32.6
Price is based on direct cost per unit plus a percentage mark-up that varies	8.9	14.9	25.2	32.9
Price is based on direct cost per unit plus a fixed percentage mark-up	15.9	19.3	19.0	24.7
Price is primarily specified by your principal customer	22.5	20.5	17.7	9.5
Price is primarily determined in other ways	15.6	3.0	4.3	12.8

Source: Jennifer Greenslade and Miles Parker 'New insights into price-setting behaviour in the UK: Introduction and survey results' The Economic Journal, February 2012, Vol. 122, Issue 558

How UK firms set their prices

Pricing method	All	Manufacturing	Construction	Retail	Other services
Reference to market conditions	39	41	51	18	48
Competitor prices	25	26	11	30	23
Direct cost + variable markup	20	20	22	21	17
Direct cost + fixed markup	17	16	19	24	14
Customer set	5	6	3	0	6
Regulatory agency	2	1	0	0	3

Note: Data are percentages of sample firms reporting use of the method shown in the left-hand column. Percentages may exceed 100 per cent because firms are permitted to indicate more than one choice.

Source: Adapted from Hall, S., Walsh, M. and Yates, A. (1996) How do UK companies set prices? Bank of England Quarterly Bulletin, May, 36, 180–92, Table A, 13.

Why pricing gets the first chapter: the leverage arithmetic

A 1% price rise raises profit far more than a 1% cost cut or 1% extra sales. **Worked example.** A firm sells 100 units at £100. Revenue £10,000, costs £9,000, so operating profit is £1,000 — a margin of $m = 10\%$.

- **Raise the price by 1%** (to £101), selling the same units: revenue £10,100, costs unchanged, profit £1,100. A 1% price rise gave **+10% profit**.
- **Cut costs by 1%** instead: costs £8,910, profit £1,090 — only +9%. **Sell 1% more units**: profit £1,010 — just +1%, because the extra units bring their own costs.
- **The rule**: with margin m, a price rise of $\Delta p\%$ moves profit by roughly $\Delta p/m$. At a 10% margin that is a 10× multiplier; at a 5% margin, 20×.

So the scarce asset is the ability to raise price *without losing customers* — that is what "pricing power" means, and why it drives valuations.

Marn & Rosiello (1992), *HBR*.

Whose prices moved? Markups, 1980–2025

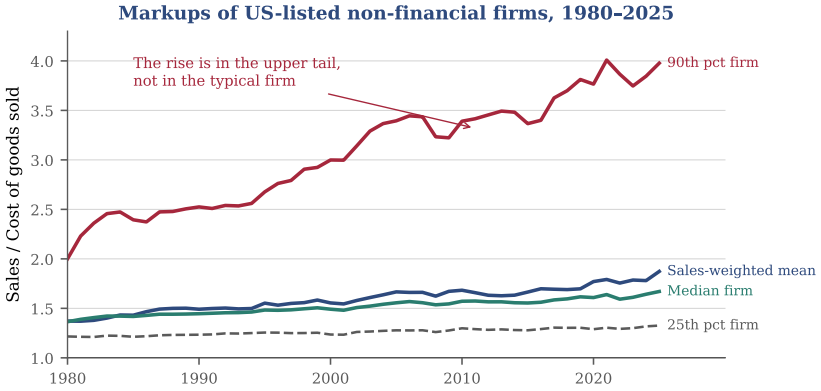

Four lines, highest to lowest: the **90th-percentile firm** (red, top) — the high-markup tail; the **sales-weighted mean**; the **median, i.e. typical, firm** (navy); the **25th-percentile firm** (teal, bottom). "Markup" here = sales ÷ cost of goods sold, so 1.5 means price is 1.5 times direct cost. **Read it as**: the typical firm barely moved (1.37 → 1.64) while the top tail roughly doubled — pricing power did not rise across the board, it *concentrated* in a minority of firms. This afternoon we measure it properly.

Own calculations, Compustat, sales/COGS, US-listed non-financials.

Pass-through: who pays a tariff, who eats a devaluation

Pass-through = how much of a cost shock a firm puts into its price, rather than absorbing it in its margin.

- The 2018–19 US tariffs were the cleanest test in decades: a sudden, precisely measured tax on imports from China.
- **At the border, pass-through was almost complete**: import prices rose by nearly the full tariff, so *US buyers* paid it — Chinese exporters did not cut their prices to absorb it (Amiti–Itskhoki–Konings; Cavallo et al.).
- **At the shop shelf it was only partial**: retailers spread the increase thinly over many products instead of raising the tariffed item by the full amount.
- **Exchange rates behave the same way**: when a currency falls, consumer prices rise by much less than the fall — firms *price to market*, setting a different margin in each country.

Who absorbs a shock reveals who has pricing power: a firm with a strong brand and few substitutes passes it on; a firm without one eats it.

Price discrimination

Price discrimination: meaning and the three degrees

- **Meaning of price discrimination**
- **Types of price discrimination**
  - first degree
    - charging consumers the maximum each is prepared to pay Whats the difficulty here?
  - third degree
    - characteristics, traits or attributes used to 'group' different consumers
  - second degree
    - offering a range of different pricing options for the consumer to choose from
- **Conditions for price discrimination**
  - **Firm must have some market power**
    - Faces a downward sloping demand curve
  - **Re-sale of the product between customers must be difficult/impossible**
    - I.e. consumed at time of purchase/perishable/high transaction costs
  - **Demand elasticity must vary between customers at any given price**

First-degree price discrimination in practice

- **First-degree price discrimination**
  - **perfect price discrimination**
    - difficult to establish the maximum people are willing to pay
    - asymmetric information
  - **personalised or person-specific pricing**
    - approaches pure first-degree price discrimination
    - more likely when there is scope for bargaining and the seller is a skilful haggler
    - **the impact of digital markets:** capturing consumers' characteristics; tracking browsing behaviour;
  - evidence of personalised pricing
  - constraints on the use of personalised pricing
- **a model of first-degree price discrimination**
  - revenue equals area under the demand curve
  - captures all the consumer surplus

First-degree Price Discrimination

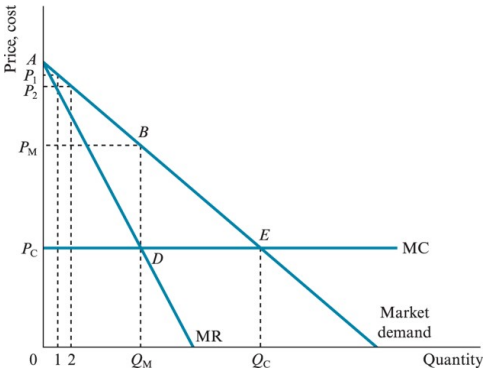

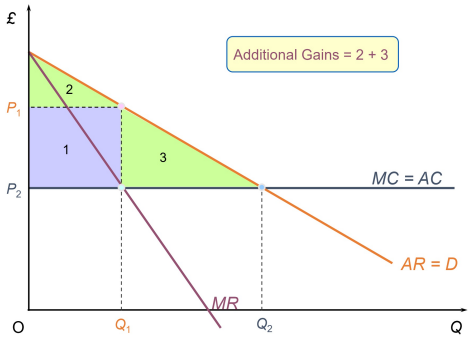

What the two panels show: on the left, ordinary *uniform* pricing — one price for everybody. On the right, *perfect* (first-degree) discrimination — the firm charges each buyer the most that buyer would pay, so the price line follows the demand curve itself and nobody keeps any surplus.

- **Area 1**: monopoly profit under uniform pricing — price P_1 above marginal cost MC, selling up to Q_1.
- **Areas 2 + 3**: what perfect discrimination adds — the consumer surplus captured, plus the deadweight loss recovered by expanding output to the efficient Q_2.

Examples of third-degree price discrimination

Characteristic	Example
Age	16–25 or senior rail card; half price children's tickets in the cinema.
Gender	'Ladies' night' in a bar or club where men pay the full price for drinks while women can get the same drinks at a discounted price.
Location	Pharmaceutical companies often charge different prices for the same medicine/drug in different countries. Consumers in the USA are often charged more than those from other countries.
Occupation	Apple, Microsoft and Orange provide price discounts to employees of educational institutions.
Business or individual	Publishers of academic journals charge much lower subscription rates to individuals than university libraries.
Past buying behaviour	Firms often charge new customers a lower price than existing customers for the same product or service as an 'introductory offer'.

Countries are the easiest way to separate customers — so firms price by passport.

- A price gap between countries survives only if buyers cannot buy cheap and resell dear. That resale is **arbitrage**, and the firm's real job is fencing it off.
- **Medicines**: the same drug is priced low in poorer countries and high in richer ones. Traders then buy low and re-export — *parallel imports*. Inside the EU this resale is legal, which is why European price gaps stay narrow.
- **Digital goods**: games and streaming services post a different price per country; customers use a VPN to pose as residents of a cheap one. How well the firm blocks that decides how big the gap can be.
- **Textbooks**: publishers sold cheap Asian editions until a student imported them and resold them in the US. The Supreme Court (*Kirtsaeng*, 2013) ruled that legal — and regional price differences narrowed immediately.

The rule: an international price gap can persist only up to the cost of moving the goods back across the border.

Which one of the following conditions is NOT necessary for (third-degree) price discrimination to take place?

1. The firm must be a monopoly.
2. Price elasticity of demand must differ in each market.
3. The firm must not be a price taker.
4. The markets must be separate.
5. People buying in the low-price market must not be able to sell to those buying in the high-price market.

Vevox poll. Which condition is the odd one out? Vote before we work through them.

Second-degree price discrimination: the idea

- **Second-degree price discrimination** — the price depends on *how much* you buy or *which version* you choose. The firm cannot tell the eager customer from the reluctant one, so it publishes a menu and lets customers **sort themselves**.
  - **discounts for buying more**
    - eg. quantity discounts
    - eg. block declining tariff — each successive block of units costs less per unit (electricity, water)
  - **everyday examples**: the multipack that is cheaper per item than the single; a rail season ticket against daily fares; gym membership against pay-per-visit; the large coffee that costs barely more than the small one.

Second-degree price discrimination: a quantity discount

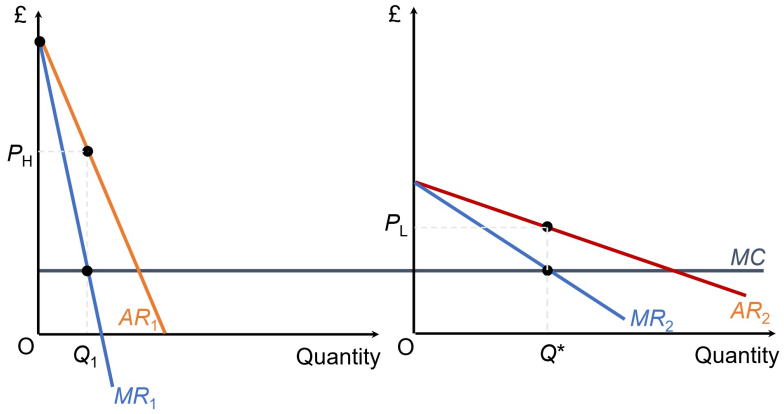

(a) Type 'I' consumer **(b)** Type 'E' consumer

The steps are a **block tariff**: the first block of units carries the highest price per unit, later blocks a lower one. A light user never leaves the first block and pays the high rate; a heavy user reaches the cheap blocks. Everyone faces the *same* price list — customers sort themselves by how much they want.

Bundling and versioning: second-degree discrimination at scale

Bundling works when the customers who value one item most value the other least.

- **That is what "negatively correlated valuations" means** (Adams–Yellen 1976). Ann would pay £90 for Word and £20 for Excel; Ben £20 for Word and £90 for Excel. Sold *separately* at £90 each, Ann buys Word and Ben buys Excel: revenue £180. Sold as *one bundle* at £110, both buy (each values the pair at £110): revenue **£220**. The bundle averages out their differing tastes.
- **Digital bundles** go further: add many goods that cost nothing to copy, and individual quirks cancel out until nearly every customer values the whole bundle at about the same amount — which the seller can then charge (Bakos–Brynjolfsson): streaming catalogues, software suites.
- **Versioning**: deliberately offer a worse version (fewer features, slower delivery) so the customers who care most pay to escape it.

For a dominant firm the same tool can tip into *foreclosure* — tying a must-have product to one facing competition (Microsoft).

Two-part tariffs in the wild: the Disneyland dilemma

Charge twice: a fixed fee F to get in, then a price p for each unit used (Oi 1971).

- F **is the entry or membership fee** — the Disneyland gate ticket, the Costco card. p **is the price per ride, item or unit** once you are inside. The classic recipe: keep p near marginal cost so people use a lot, and take the profit in F.
- **Costco**: goods sold at close to cost, with the membership fee as the profit engine.
- **Printers and razors reverse it**: a *low* F (cheap printer, cheap handle) and a *high* p (expensive ink, expensive blades). Heavy users buy more consumables and so pay more in total — the consumable acts as a **meter** of how much they use it.
- **Cloud and telecoms**: a platform fee plus usage charges; the regulatory arguments are about which of the two competition should be allowed to touch.

Design rule: put the charge where customers differ most — in *whether* they join (F), or in *how much* they use (p).

Behavioural pricing: small price changes, surprisingly large effects

Real buyers do not respond to price as smoothly as the textbook demand curve assumes.

- Theory says demand falls *gradually* as price rises. In real shop data it sometimes falls **abruptly** at particular prices — that is all "kinks" means here.
- **Left-digit effect**: raising a price from £4.99 to £5.00 — one penny — can cut sales noticeably, because shoppers read the first digit and file it as "five pounds". Hence the endless 9-endings.
- **Reference points**: "was £120, now £79" sells better than £79 on its own, because the crossed-out price supplies the comparison. Adding a deliberately overpriced third option (a *decoy*) makes the middle one look sensible.
- **Drip pricing**: advertise a low headline price and add fees at the checkout. Effective enough that regulators in the US and the EU have moved against it.

Keep both readings: firms exploit predictable mistakes — *and* a firm that prices honestly while rivals hide fees simply looks expensive, and loses.

Second-degree price discrimination: the other forms

- Beyond the quantity discounts we have just worked through, the same self-selection logic gives:
  - **coupons and vouchers** — the discount is open to everyone, but only shoppers willing to spend the *time and effort* to clip, search or enter a code actually claim it, and those are exactly the price-sensitive ones. The effort is the sorting device.
  - **intertemporal pricing** — the same product costs more early and less later, so buyers sort themselves by *how long they are willing to wait*: airline tickets, hardback then paperback, a new phone then its discount.

Intertemporal price discrimination (based on willingness to pay at a time period)

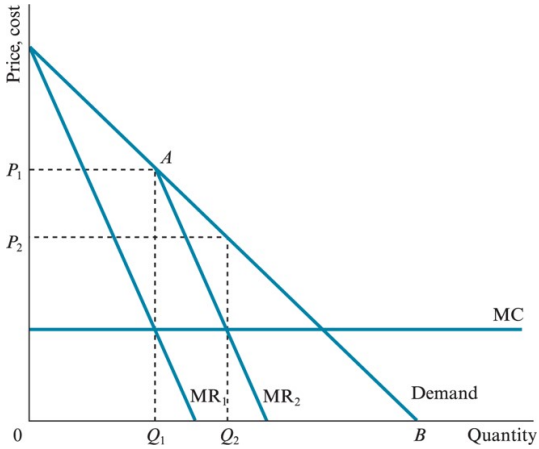

Two sales in sequence: the firm first sets the high price P_1 and sells Q_1 to impatient, high-willingness-to-pay buyers; *later* it drops the price to P_2 and picks up the patient, price-sensitive buyers it left out at the start. (MR_1 is the marginal revenue from that first group.)

Dynamic and algorithmic pricing — and the collusion nobody signed

When software sets the prices, discrimination becomes continuous — and can drift into collusion.

- **Dynamic pricing** is intertemporal discrimination run by computer: airlines, hotels and ride-hailing apps change prices constantly with demand, time remaining, and who is asking.
- **The experiment**: Calvano et al. (2020) let simple self-learning pricing programs compete. They learned to hold prices *above* the competitive level and to punish each other with price cuts when one undercut — with no communication and no agreement.
- **Why it works**: a high price needs only repetition, the ability to watch rivals' prices, and patience — software has all three.
- **The legal difficulty**: competition law forbids an *agreement* — but if two machines merely learn the same behaviour, what was agreed? Cases so far target the *shared pricing software* rival firms all subscribe to.

Auctions: pricing when you don't know the demand curve

When willingness to pay is unknown and stakes are lumpy, make buyers reveal it.

- **Spectrum auctions**: "spectrum" means the radio frequencies mobile networks must have in order to operate. Governments own them and, since nobody knows what a licence is really worth, they *auction* them — the modern designs (Milgrom and Wilson, Nobel 2020) are built to stop bidders colluding and to stop the winner overpaying.
- **Online advertising**: every advertising slot you are shown was auctioned in the milliseconds while the page loaded.
- **The winner's curse**: when everyone is guessing the same unknown value, the winner is whoever guessed *highest* — so winning is itself evidence you over-estimated. Experienced bidders deliberately bid below their own estimate.

For the MNE, entry is often a bid, not a list price: privatisations, licences, takeover battles.

Part II

Transfer pricing and the life cycle

Transfer Pricing

A transfer price is what one part of a company charges another part of the same company.

- **Example**: the group's German plant builds an engine for £600 and "sells" it to the group's UK factory, which puts it in a car sold for £20,000. The engine never left the company — but a price still had to be written on it. That price is the transfer price.
- It matters because it **decides where the group's profit shows up**: charge £600 and the profit lands in the UK; charge £3,000 and it lands in Germany. Total group profit is unchanged — *where it appears* is not.
- **The organisational problem comes first**: if the selling division prices like a monopolist, the buying division sees an expensive input, makes too little, and the group as a whole earns less.

An internal accounting question — until the two divisions sit in countries with different tax rates. That is the next slide.

Transfer pricing is a capital-structure decision too

The transfer price is one channel for moving profit. There are two more.

- **Internal debt**: the group's treasury in a low-tax country lends to an operating company in a high-tax one; the interest is a cost there and income here.
- **IP (intellectual property) royalties**: register the patent or brand in a low-tax country, then charge the other units to use it. Europe's taxable profit falls, the low-tax country's rises — and nothing physical has moved.
- **That is what BEPS means**: *base erosion* = shrinking the profit declared in a country; *profit shifting* = moving it somewhere taxed less.
- **The policy answer**: the **OECD** (the club of mostly rich countries that writes international tax standards) runs the BEPS programme — country-by-country reporting plus a **15% global minimum tax** (Pillar Two, from 2024), topping profit up to 15% wherever it was parked.

This afternoon: the footprint in the accounts — 21.3% vs 6.3% effective tax for same-country firms, by legal address.

Marginal-cost transfer prices fix the internal problem — and open the external one.

- Cross-border internal prices face no market test ⇒ set so **profits land where taxes are lowest**.
- Starbucks, Apple, Amazon, Coca-Cola: high internal charges for logos, brands and services 'provided' from Luxembourg or the Caymans.
- OECD answer: **arm's-length pricing** — as if between independent firms. Hard to police: the 'subsidiary' is often one office, one employee.

This afternoon: the footprint in the accounts — 21.3% vs 6.3% effective tax, by legal address.

Pricing and the product life cycle

The stages in a product's life cycle

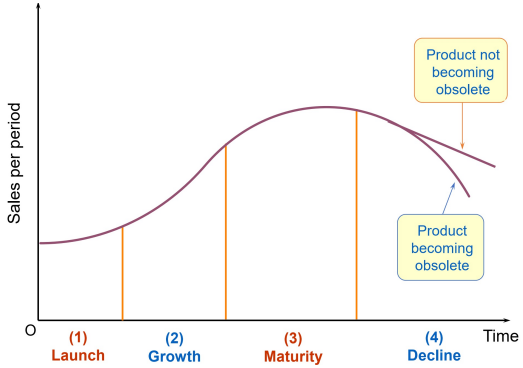

Sales start slowly, accelerate, flatten, then fall. The point for us is that the *right price* changes with the stage, because the competition does.

Pricing through the life cycle: what firms actually do at each stage

Stage	Typical pricing	Example
Introduction	*skimming* (start high) or *penetration* (start low)	new iPhone at launch; a £1 streaming trial
Growth	hold the price, add versions	mid-range and "Pro" models appear
Maturity	discounts, promotions, bundles	supermarket own-label vs brand; contract deals
Decline	harvest, or exit	DVDs, printed encyclopaedias

- **Skimming** sells first to impatient enthusiasts at a high price, then cuts to reach everyone else — that is exactly the intertemporal discrimination from Part I, spread over a product's life.
- **Penetration** does the opposite: price low to build the customer base quickly, then earn from renewals, consumables or add-ons. It is the right choice where switching costs or network effects keep customers once they arrive.

Same product, opposite strategies — the choice depends on whether customers are locked in later, or only impatient now.

Part III

Strategic interaction: games

What is Game Theory?

- Game theory is a tool for studying strategic behavior which considers the expected behavior of others and the mutual recognition of interdependence.
- It shows how situations can arise where firms take decisions that may appear rational individually but lead to suboptimal outcomes for all combined.
- It is an approach to decision-making where two or more decision-makers or players face choices between a number of possible actions at any stage of the game.

Applications of Game Theory

- Oligopoly
- Sports
- Chess
- Military strategy and nuclear deterrence
- International Business
- Economics
- Rock Paper Scissors Game
- Bidding at Auction
- Poker
- Market Shares and Stockholders

Dominant strategies

Dominant strategies: the easiest games to solve

A dominant strategy is one that is best for you whatever the other player does.

- **Strictly dominant**: strictly better than every alternative, against every choice the rival could make. If you have one, you need not predict the rival at all — just play it.
- **Weakly dominant**: never worse, and strictly better against at least one of the rival's choices. You would still play it, but the argument is weaker — against some rival choices you only tie.
- **Dominated**: the mirror image — a strategy that some other option always beats. Rational players never use one, and crossing dominated strategies out repeatedly often solves the whole game.
- Most games have *no* dominant strategy. Then we need the Nash equilibrium — which is where we go next.

Always look for dominance first: it is the cheapest thing to check, and when it exists the answer is immediate.

Entry Game with a Weakly Dominant Strategy

Weak dominance: at least as good against everything, strictly better against something.

- Setup: incumbent **A** vs potential entrant **B**, moving simultaneously.
- B chooses *Enter* / *Stay out*; A pre-commits *Accommodate* / *Fight*.
- A's capacity to fight requires **irreversible** steps before B decides — costless if B stays out.

The figure: A has a weakly dominant strategy; B has none. Find them before the next slide.

Figure 2. Entry game with a weakly dominant strategy

		Firm B's strategies	
		Entry	*No entry*
Firm A's strategies	**Accommodate**	**2**, *4*	**3**, *3*
	Fight	**1**, *1*	**3**, *3*

Entry game with a weakly dominant strategy

Best responses, line by line.

- **A**: if B stays out, payoff 3 either way; if B enters, Accommodate (2) beats Fight (1) $\Rightarrow$ *Accommodate* weakly dominates.
- **B**: against Accommodate, Enter (4) beats Stay out (3); against Fight, Stay out (3) beats Enter (1) $\Rightarrow$ no dominant strategy.
- Equilibrium: **(Accommodate, Enter)**.

The threat to fight is empty — unless it is made *sunk*. That was this morning's deterrence slide.

Nash equilibrium

Nash equilibrium

- In a **Nash equilibrium**, no firm can improve its payoff by changing strategy *on its own*, given that the other firm keeps its strategy unchanged.
- Named after **John Nash** (1928–2015), the mathematician who shared the 1994 Nobel Prize in Economics for this idea.
- His life was filmed as *A Beautiful Mind* (2001), which won four Oscars — among them Best Supporting Actress for Jennifer Connelly.

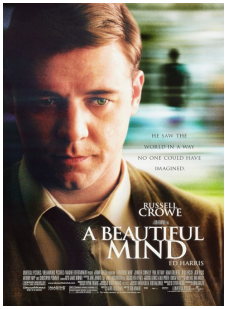

To read: `https://www.economist.com/schools-brief/2016/08/20/prison-breakthrough`

Figure 3. Payoff matrix for the advertising budgets of firms A and B

- In **Figure 3**, firms A and B must decide simultaneously their advertising expenditures.
- They have a choice between three levels of expenditure: *Low*, *Medium* or *High*.
- Both firms' payoffs from the advertising campaign depend on their own expenditure and on the expenditure of the other firm.

		Firm B's budget		
		Low	*Medium*	*High*
Firm A's budget	**Low**	**40** *40*	**35** *45*	**10** *25*
	Medium	**35** *35*	**45** *30*	**15** *20*
	High	**30** *25*	**25** *15*	**20** *30*

Payoff matrix for the advertising budgets of firms A and B

As before, consider firm A's choices:

- If B chooses *Low*, A's best response is *Low*.
- If B chooses *Medium*, A's best response is *Medium*.
- If B chooses *High*, A's best response is *High*.

Similarly, consider firm B's choices:

- If A chooses *Low*, B's best response is *Medium*.
- If A chooses *Medium*, B's best response is *Low*.
- If A chooses *High*, B's best response is *High*.

So the Nash equilibrium is (High, High) — the only pair where the two lists agree: A's best reply to *High* is *High*, and B's best reply to *High* is *High*, so neither firm can gain by moving on its own. At every other cell at least one firm would want to change.

Prisoner's dilemma

Figure 4: Production game with prisoner's dilemma structure

Firm A's strategies	Firm B's strategies: *Low*	Firm B's strategies: *High*
Low	**3** *3*	**1** *4*
High	**4** *1*	**2** *2*

Production game with prisoner's dilemma structure

Both play the dominant strategy High — and both wish they hadn't.

- (High, High) is the dominant-strategy equilibrium *and* the Nash equilibrium: payoffs (2, 2).
- Yet (Low, Low) pays (3, 3): **jointly better, individually unreachable**.
- Unilateral deviation from (Low, Low) pays — which is exactly why it cannot hold.

That gap between equilibrium and optimum is the entire economics of cartels — and of advertising wars.

Conflict vs. Cooperation in Prisoner's Dilemma

Relax the assumptions and cooperation can survive.

- **Communication**: frequent contact lets firms exchange information and monitor each other.
- **Reaction lag**: the faster the punishment, the smaller the temptation to defect.
- **Repetition**: in continuous rivalry, players can learn that cooperation beats aggression.

Hold these three — they become the collusion checklist later this morning.

Another example: OPEC

Prisoner's Dilemma: OPEC Oil Production

- **Scenario:** members of **OPEC** (the Organization of the Petroleum Exporting Countries) must decide how much oil to produce.
- **Payoff Structure:**
  - If all members limit production, prices stay high — everyone benefits.
  - If one country cheats (overproduces), it gains extra revenue short-term.
  - But if all cheat, prices collapse — *everyone loses*.
- **Prisoner's Dilemma Logic:**
  - *Dominant strategy* is to cheat — but leads to a worse outcome collectively.
  - Cooperation requires trust, monitoring, or repeated interaction.
- **International Business Insight:** Cartel behavior in commodities often faces dilemma dynamics — repeated games, penalties, and reputational concerns are key to sustaining cooperation.

Vevox poll. Cooperate or defect? Predict where OPEC members land.

Sequential games

Now one player moves first and the other observes: map it as a game tree.

- Two cereal makers each launch one product: *crunchy* or *fruity*.
- Crunchy is the more popular taste — so who moves first matters.
- Solve trees **backwards**: anticipate the follower, then choose.

Commitment value comes from moving first *visibly and irreversibly* — next slide.

Figure 7: Breakfast cereals game — the payoff matrix ("strategic form")

Firm A's strategies \ Firm B's strategies	*Crunchy*	*Fruity*
Crunchy	**3** *3*	**5** *4*
Fruity	**4** *5*	**2** *2*

If both choose at the same time there are two Nash equilibria — (Crunchy, Fruity) and (Fruity, Crunchy). In each, the two firms pick *different* flavours, and neither can gain by switching alone; both landing on the same flavour is never an equilibrium. Which of the two happens is exactly what moving first decides.

The same game as a decision tree, when Firm A moves first

Firm A chooses first, Firm B sees the choice and replies. Payoffs are (A, B). Solve **backwards**: work out B's best reply to each branch, then pick A's best opening move.

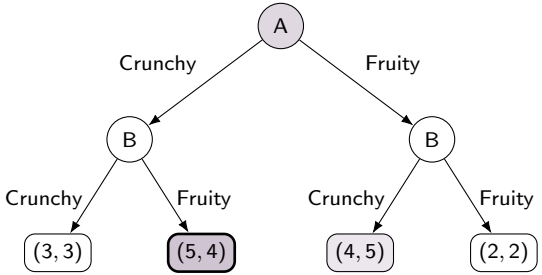

If A picks **Crunchy**, B compares 3 (crunchy) with 4 (fruity) and picks **Fruity** $\Rightarrow$ A gets **5**. If A picks **Fruity**, B compares 5 (crunchy) with 2 (fruity) and picks **Crunchy** $\Rightarrow$ A gets 4. So A opens with **Crunchy** and the outcome is **(5, 4)**, shaded. Moving first let A take the popular flavour and leave B the niche — the first-mover advantage, in one tree.

The tree assumed A could move first. Should it want to?

- Moving first only works if the move is **irreversible and visible** — if A could quietly switch flavours later, B would ignore the opening move and the advantage disappears.
- But irreversibility has a cost: you commit before you know how the market turns out. **Waiting has value** when the future is uncertain and the decision can be reversed (Dixit–Pindyck call this the *option value* of waiting).
- **So: pre-empt** where the first-mover advantage is durable (scarce sites, a standard, an installed base); **wait** where uncertainty is high and you can still change course cheaply.
- **Reading it for international business**: staged entry — export, then joint venture, then a wholly-owned subsidiary — is not indecision; it is buying information before committing capital.

Strategy is choosing which is worth more: the commitment, or the option to change your mind.

Exit is a game too: who leaves a dying industry?

When an industry is shrinking, whoever stays last can be profitable — so leaving becomes a game.

- As demand falls there is no longer room for everyone. If a rival exits, the survivor inherits its customers and may do well; if nobody exits, both lose money. This waiting contest is called a **war of attrition**.
- **The surprising result** (Ghemawat–Nalebuff): the *larger* firm usually exits first. Its bigger plant needs more volume to cover its fixed costs, so as the market shrinks it becomes unprofitable *sooner* than the small specialist.
- **What decides it**: how costly exit is (redundancy, clean-up, contracts), whether capacity can be closed in pieces rather than all at once, and **commitment** — publicly ruling out exit makes the rival's position hopeless.
- **Where you see it**: coal, printed newspapers, older chemicals plants — and every "last one standing" pitch to investors.

Monday's data: half the worst-performing firms are gone within five years. *How* they go is a strategic choice.

Part IV

Repeated games and tacit collusion

Repeated Games

- In the previous discussion of single-period prisoner's dilemma and other games, it is assumed that the game is played only once.
- However, some games may be played repeatedly by the same players. In this case, it is more likely that cooperative behaviour will evolve as the two firms observe and learn from each other's behaviour.
- In a **repeated or multiple-period game**, each firm may attempt to influence its rival's behaviour by sending signals that **promise to reward cooperative behaviour and threaten to punish non-cooperative behaviour**.

Let's look again at Figure 4: prisoner's dilemma example

		Firm B's strategies: *Low*	Firm B's strategies: *High*
Firm A's strategies	**Low**	**3** *3*	**1** *4*
	High	**4** *1*	**2** *2*

Tit-for-Tat Strategy in Repeated Games

One shot: (High, High) pays (2, 2), though (Low, Low) would pay (3, 3). Repeat indefinitely, and A can teach B to cooperate. The strategy:

- Period 1: A chooses **Low** (cooperate).
- If B chose Low in $t-1$: A chooses Low in t.
- If B chose High in $t-1$: A chooses High in t.

Why this strategy is famous: in 1980 the political scientist **Robert Axelrod** invited scholars worldwide to submit computer strategies for a repeated prisoner's dilemma and played them against each other in a tournament. The winner was the *simplest* entry submitted — tit-for-tat, just four lines long. It wins by being **nice** (never defects first), **retaliatory** (punishes immediately), **forgiving** (returns to cooperating as soon as the rival does) and **clear** (the rival can work out the rule and see that cooperating pays). The cola chart on the next slide is forty years of exactly this, played by two real firms.

Forty years of the game, in the accounts

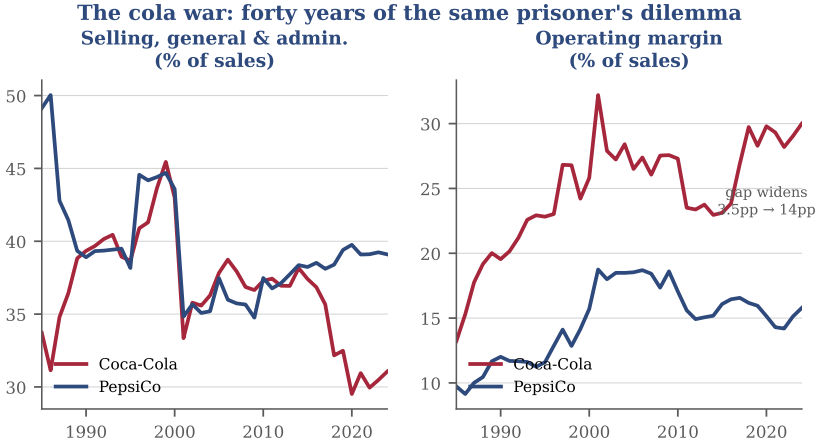

Coke and Pepsi held SG&A at ∼37–39% of sales for three decades. Neither ever defected to "don't advertise" — exactly as the repeated game predicts.

Own calculations, Compustat, fiscal 1985–2024.

When does the cooperative outcome actually survive? A checklist

Tit-for-tat says cooperation can hold. Whether it does depends on punishment being easy.

Collusion easier when...	**because...**
few firms, symmetric	one defector is obvious, punishment coordinated
transparent prices, frequent interaction	cheating detected fast, punished fast
stable demand, high entry barriers	the prize for patience is durable
multimarket contact	retaliation available everywhere at once

- Policy runs the list in reverse: **leniency** pays the first confessor — a designed defection.
- OPEC scores high on transparency and repetition, low on symmetry — its production history in one row.

Cartels in the wild: the checklist, with convictions

Three real cartels — each one ticks the boxes on the previous slide.

- **Lysine (1990s)**: lysine is an additive in animal feed — a standardised product made by a handful of firms, so cheating is easy to spot. Executives of the US agribusiness ADM met rivals to fix prices; an FBI informant recorded them saying *"our competitors are our friends; our customers are the enemy"*. Executives went to prison. It is the classic "few firms, identical product" case.
- **Vitamins (1999)**: the world's vitamin producers divided up markets and volumes for roughly a decade. The Swiss firm Roche paid a US fine of about $500m — a record at the time. Long-lived because demand was stable and the members few.
- **EU trucks (2016–17)**: Europe's major truckmakers coordinated prices for **fourteen years**; the European Commission fined them around EUR 4bn. It surfaced only when one member applied for **leniency** — confessing first in exchange for immunity.

Note the pattern: cartels rarely collapse from competition. They collapse because a member defects — to the regulator.

Multimarket contact: why global rivals pull their punches

MNEs meet the same rivals in dozens of countries — pooling the punishment budget.

- Undercut me in Brazil; I retaliate in Brazil, Germany and Japan at once (Bernheim–Whinston 1990).
- Evidence: airfares *higher* on routes flown by high-contact carrier pairs (Evans–Kessides 1994).
- Spheres of influence: competition softest where I am strong and you are weak.

Entering a rival's home market = acquiring a hostage for every market you share.

Example: Coca-Cola vs. Pepsi (Global Pricing & Advertising)

- Duopoly competing on price, marketing, and product variety in global markets.
- Long-term interaction leads to **tacit collusion** — avoidance of aggressive price wars.
- Example: Avoidance of deep price cuts in mature markets (e.g., US, UK).

Vevox poll. Do Coke and Pepsi ad budgets move together? Guess the sign of the correlation.

References I

- Marn, M. & Rosiello, R. (1992), 'Managing Price, Gaining Profit', *Harvard Business Review* 70(5), 84–94.
- Phlips, L. (1983), *The Economics of Price Discrimination*, Cambridge University Press.
- OECD (2021–), *Two-Pillar Solution to Address the Tax Challenges of the Digitalisation of the Economy* (Pillar Two).
- von Neumann, J. & Morgenstern, O. (1944), *Theory of Games and Economic Behavior*, Princeton University Press.
- Nash, J. (1950), 'Equilibrium Points in n-Person Games', *PNAS* 36(1), 48–49.
- Schelling, T. (1960), *The Strategy of Conflict*, Harvard University Press.
- Axelrod, R. (1984), *The Evolution of Cooperation*, Basic Books.
- Dixit, A. & Pindyck, R. (1994), *Investment under Uncertainty*, Princeton University Press.
- Dixit, A. & Nalebuff, B. (2008), *The Art of Strategy*, W.W. Norton.
- Sloman, J. et al. (2024), *Economics for Business*, 9th ed., Pearson. Chs on pricing and business strategy.
- Lipczynski, Goddard & Wilson (2017), *Industrial Organization*, 5th ed., Chs 5–7, 13.
- Adams, W. & Yellen, J. (1976), 'Commodity Bundling and the Burden of Monopoly', *QJE* 90(3), 475–498.

References II

- Bakos, Y. & Brynjolfsson, E. (1999), 'Bundling Information Goods', *Management Science* 45(12), 1613–1630.
- Calvano, E., Calzolari, G., Denicolò, V. & Pastorello, S. (2020), 'Artificial Intelligence, Algorithmic Pricing, and Collusion', *AER* 110(10), 3267–3297.
- Milgrom, P. (2004), *Putting Auction Theory to Work*, Cambridge University Press.
- Ivaldi, M., Jullien, B., Rey, P., Seabright, P. & Tirole, J. (2003), *The Economics of Tacit Collusion*, report for DG Competition.
- Bernheim, B.D. & Whinston, M. (1990), 'Multimarket Contact and Collusive Behavior', *RAND Journal of Economics* 21(1), 1–26.
- Evans, W. & Kessides, I. (1994), 'Living by the "Golden Rule": Multimarket Contact in the U.S. Airline Industry', *QJE* 109(2), 341–366.
- Ghemawat, P. & Nalebuff, B. (1985), 'Exit', *RAND Journal of Economics* 16(2), 184–194.
- Amiti, M., Itskhoki, O. & Konings, J. (2019), 'The Impact of the 2018 Tariffs on Prices and Welfare', *JEP* 33(4), 187–210.
- Cavallo, A., Gopinath, G., Neiman, B. & Tang, J. (2021), 'Tariff Pass-Through at the Border and at the Store', *AER: Insights* 3(1), 19–34.
- Oi, W. (1971), 'A Disneyland Dilemma: Two-Part Tariffs for a Mickey Mouse Monopoly', *QJE* 85(1), 77–96.
- Connor, J. (2001), *Global Price Fixing*, Springer.