Advanced Microeconomics

Lecture 3: Welfare Review, Markets, and Market Power

Yihu Hou

Renmin University of China

Spring 2026

Outline

1. Revisit welfare measures: CV, EV, and why the two differ
2. Chapter 2: from individual demand to market equilibrium
3. Competition, monopoly, and basic oligopoly comparisons

Where CV and EV come from

Two value objects at a fixed price vector p

$$v(p, m) = \max_{x \in \mathbb{R}^n_+} \{u(x) : p \cdot x \leq m\}$$

$$e(p, \bar{u}) = \min_{x \in \mathbb{R}^n_+} \{p \cdot x : u(x) \geq \bar{u}\}$$

- $v(p, m)$ maps **money to utility**: "How much utility can I reach?"
- $e(p, \bar{u})$ maps **utility to money**: "How much money do I need?"
- CV and EV are both built from **the expenditure function**.

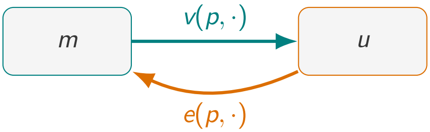

Inverse mapping idea

$$v\big(p, e(p, \bar{u})\big) = \bar{u}, \qquad e\big(p, v(p, m)\big) = m.$$

Reading CV and EV I: fix utility, compare prices

Suppose prices move from p^0 to p^1 and income stays at m. Let

$$u^0 = v(p^0, m), \qquad u^1 = v(p^1, m),$$

so for a harmful price change we typically have $u^1 < u^0$.

Both are expenditure differences at a common utility level

$$\mathrm{CV} = e(p^1, \boldsymbol{u^0}) - e(p^0, \boldsymbol{u^0}), \qquad \mathrm{EV} = e(p^1, \boldsymbol{u^1}) - e(p^0, \boldsymbol{u^1}).$$

- In both formulas, we compare **old prices** and **new prices** while holding utility fixed.
- CV uses the benchmark $\boldsymbol{u^0}$: how much more expensive is it to maintain the original utility after the price change?
- EV uses the benchmark $\boldsymbol{u^1}$: how much more expensive is it to finance the post-change utility at the new prices?

Reading CV and EV II: fix price, compare utilities

With the same notation, the key identity is

$$m = e(p^0, u^0) = e(p^1, u^1).$$

Both are expenditure differences at a common price vector

$$\begin{aligned} \mathrm{CV} &= e(\boldsymbol{p}^1, u^0) - e(\boldsymbol{p}^1, u^1), \\ \mathrm{EV} &= e(\boldsymbol{p}^0, u^0) - e(\boldsymbol{p}^0, u^1). \end{aligned}$$

- In both formulas, we compare **two utility levels** under the **same prices**.
- CV is evaluated at the **new prices** $\boldsymbol{p}^1$.
- EV is evaluated at the **old prices** $\boldsymbol{p}^0$.

Why both readings are true

$$\begin{aligned} \mathrm{CV} &= e(p^1, u^0) - m \\ &= e(p^1, u^0) - e(p^0, u^0) \\ &= e(p^1, u^0) - e(p^1, u^1), \\ \mathrm{EV} &= m - e(p^0, u^1) \\ &= e(p^1, u^1) - e(p^0, u^1) \\ &= e(p^0, u^0) - e(p^0, u^1). \end{aligned}$$

A picture: two equivalent readings of CV and EV

Fix utility, compare prices

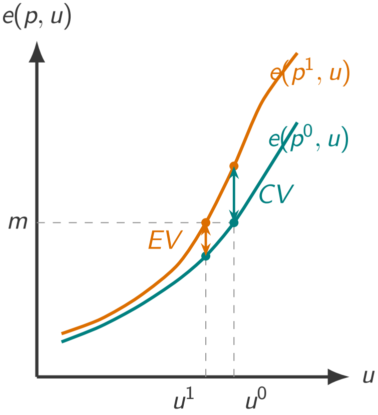

$\text{CV} = e(p^1, \boldsymbol{u^0}) - e(p^0, \boldsymbol{u^0})$

$\text{EV} = e(p^1, \boldsymbol{u^1}) - e(p^0, \boldsymbol{u^1})$

Fix price, compare utilities

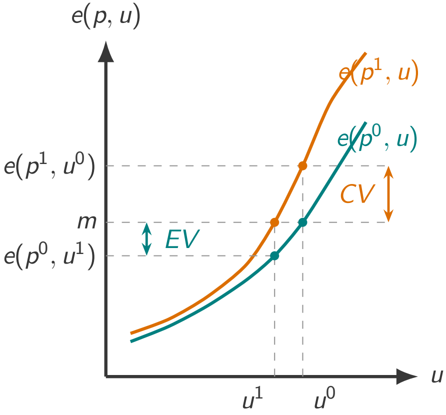

$\text{CV} = e(\boldsymbol{p^1}, u^0) - e(\boldsymbol{p^1}, u^1)$

$\text{EV} = e(\boldsymbol{p^0}, u^0) - e(\boldsymbol{p^0}, u^1)$

CV in the indifference-curve picture

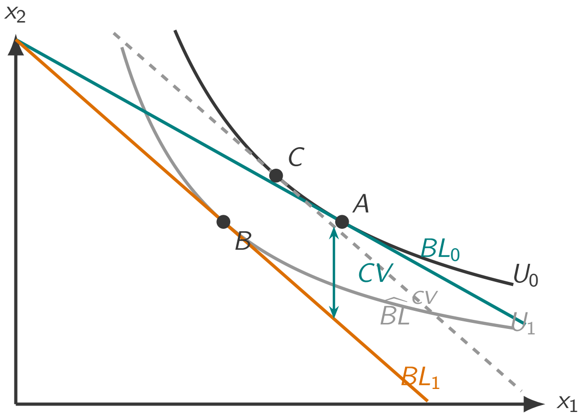

Read This Picture

- A is the original optimum on BL_0 and U_0.
- B is the new optimum after the price increase.
- C is the tangency on $\widehat{BL}^{CV}$ that restores U_0 at the new prices.
- CV is the upward shift from BL_1 to $\widehat{BL}^{CV}$.

EV in the indifference-curve picture

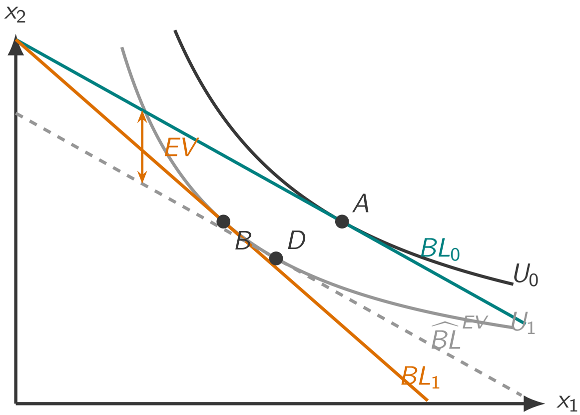

Read This Picture

- A is the original optimum on BL_0 and U_0.
- B is the new optimum after the price increase.
- D is the tangency on $\widehat{BL}^{EV}$ that preserves U_1 at the old prices.
- EV is the downward shift from BL_0 to $\widehat{BL}^{EV}$.

Which one is larger?

For a price increase of a normal good

$$\mathrm{CV} \geq \mathrm{EV} \geq 0.$$

- EV is measured **before** the price increase, when the consumer still faces the old prices.
- CV is measured **after** the price increase, when restoring the old utility is more expensive.
- If income effects are small (for example, quasi-linear preferences), then CV and EV are close; with no income effect, they coincide.

Reverse case

For a beneficial price decrease of a normal good, the ranking reverses:

$$\mathrm{EV} \geq \mathrm{CV} \geq 0.$$

Quasi-linear preferences: CV $=$ EV

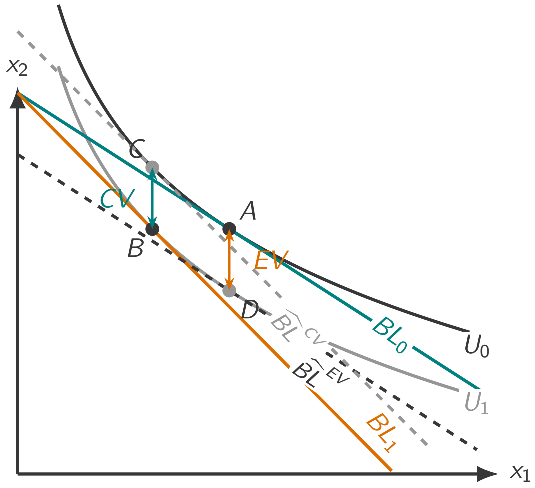

Why equality is visible

- Here $u(x_1, x_2) = 2 \ln x_1 + x_2$, so indifference curves are vertical translations.
- The optimal x_1 depends only on the slope (the price of x_1), not on income.
- So A and D share the same x_1, and B and C share the same x_1.
- The two vertical money shifts have the same length:

$$\mathrm{CV} = \mathrm{EV}.$$

From consumer theory to market analysis

Last lecture gave us individual demand

$x(p, m)$ from utility maximization.

This lecture moves one level up

many consumers' demands $+$ firms' supplies $\Rightarrow$ market price and quantity.

- Consumer theory tells us how one buyer reacts to prices and income.
- Market theory tells us how prices are determined when many buyers and sellers interact.

Supply, demand, and market equilibrium

Basic objects

$$Q^D = D(p), \qquad Q^S = S(p).$$

Equilibrium

$$D(p^*) = S(p^*).$$

- p^* clears the market.
- At p^*, planned purchases equal planned sales.

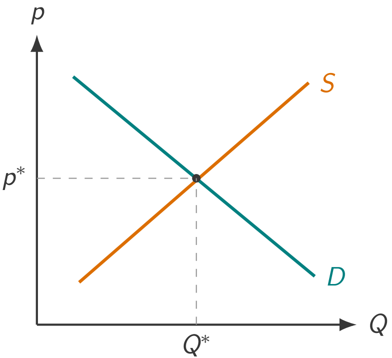

Movements along a curve vs. shifts of the curve

Movement along the demand curve

A change in **the good's own price** changes quantity demanded along the same curve.

Shift of the demand curve

A change in **income, tastes, expected future prices, or prices of substitutes/complements** changes the whole demand schedule.

Likewise for supply

Technology, input costs, regulation, or taxes shift supply; the market price then moves because the curve moved.

Reading rule

Always ask: did the *curve* move, or did we move *along* a fixed curve?

A per-unit tax on sellers

With a tax τ on sellers

$$D(p_b) = S(p_s), \qquad p_b = p_s + \tau.$$

- Buyers pay the gross price p_b.
- Sellers receive the net price p_s.
- Quantity falls because the tax creates a wedge between the two prices.

Incidence rule

The more **inelastic** side of the market bears more of the burden.

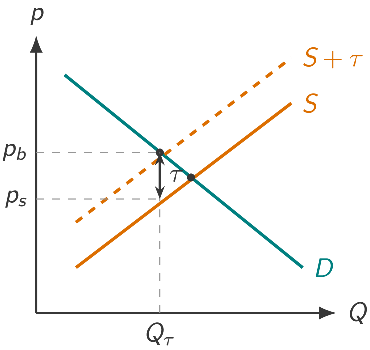

Price ceilings and price floors

Price ceiling: $\bar{p} < p^*$

- Quantity demanded exceeds quantity supplied.
- Shortage appears.
- Allocation is no longer determined by the competitive price mechanism.

Price floor: $\underline{p} > p^*$

- Quantity supplied exceeds quantity demanded.
- Surplus appears.
- Trade is reduced unless the government buys the excess or some rationing rule is used.

Policy point

Price controls do not only change *price*; they also change *who is rationed* and *how much trade is lost*.

Perfect competition: the benchmark market structure

Key assumptions

- Many buyers and many sellers
- Homogeneous product
- Free entry and exit in the long run
- Each firm is too small to influence market price

Why it matters

Perfect competition is not a description of every real market. It is a **benchmark**: a clean case in which price-taking behavior and efficiency are easiest to see.

The competitive firm: profit maximization and supply

Firm problem

$$\max_{q \geq 0} \ \pi(q) = pq - C(q).$$

Interior condition

$$p = MC(q^*).$$

- A competitive firm takes p as given.
- In the short run, the firm's supply curve is the part of MC above AVC.
- If $p < AVC$, shutdown is optimal.

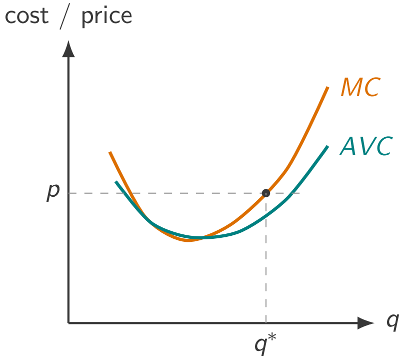

Short-run competitive equilibrium and producer surplus

Industry equilibrium in the short run

- Add up all individual firm supplies to get market supply.
- The intersection of market supply and market demand determines the short-run equilibrium price p^*.
- Each active firm then chooses q_i^* such that

$$p^* = MC(q_i^*).$$

Producer surplus

Producer surplus is the area *above* the supply curve and *below* the market price. In the short run it is revenue above variable cost, so it is not exactly the same as profit when fixed costs are present.

Aggregation example

If there are N identical firms and each firm has short-run supply $S_f(p)$, then

$$S(p) = N \cdot S_f(p).$$

Reading rule

Competitive equilibrium comes from two layers:

$$\text{firm optimization} \quad + \quad \text{market clearing.}$$

Short run vs. long run in perfect competition

Short run

- Number of firms is fixed.
- Firms may earn positive or negative economic profit.

Long run

- Entry and exit drive economic profit to zero.
- In equilibrium,

$$p^* = AC(q^*) = MC(q^*),$$

at the output that minimizes average cost.

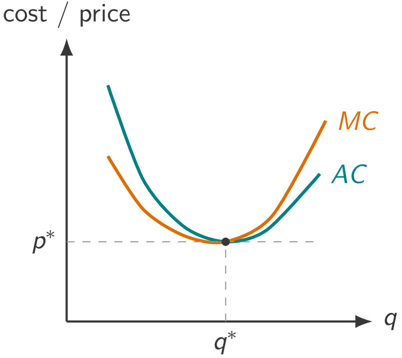

Entry/exit logic

If $p > \min AC$, entry pushes price down. If $p < \min AC$, exit pushes price up.

Perfect competition in reality: why the benchmark still helps

Real-world relevance

- Commodity markets such as wheat or corn often come closer to the competitive benchmark than markets with brands or network effects.
- In many other industries, the assumptions fail, but the model still provides a clean reference point.

Agricultural example

- Many small farmers produce nearly identical crops.
- Each individual farmer is too small to affect the global price.
- If price rises, more farmers plant next season, supply expands, and price is pushed back toward equilibrium.

Lesson

Price-taking behavior with many producers is the hallmark of perfect competition, even though real markets still involve weather shocks, subsidies, storage, and trade policy.

Why competition is a useful welfare benchmark

- In the competitive benchmark, each firm expands output until

$$p = MC.$$

- This means the value to consumers of the marginal unit equals the marginal social opportunity cost of producing it.
- Under ideal assumptions, no mutually beneficial trade is left unrealized.

What this benchmark is good for

It gives us a reference point for asking what market power changes:

price, quantity, consumer surplus, deadweight loss.

Why study imperfect competition?

- Many real markets have a dominant firm or a small number of major players.
- Once firms have market power, they no longer behave like price takers.
- The central questions become:
  - how price and output differ from the competitive benchmark,
  - how much surplus is transferred,
  - and how much deadweight loss appears.

Goal of the next part

Compare monopoly and oligopoly outcomes to the competitive benchmark and identify what strategic power changes.

Monopoly: one firm chooses output strategically

Profit maximization

$$\max_{Q \geq 0} \; \pi(Q) = p(Q)\,Q - C(Q).$$

First-order condition

$$MR(Q) = MC(Q), \qquad MR(Q) = p(Q) + Qp'(Q).$$

- The monopolist recognizes that selling more lowers price on *all* units sold.
- Therefore $MR < p$ when demand slopes downward.

Monopoly vs. competition: price, quantity, and deadweight loss

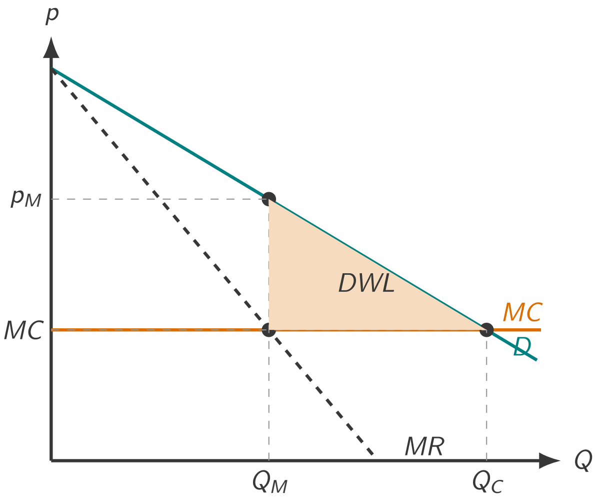

Comparison

- Monopoly output Q_M is below the competitive quantity Q_C.
- Monopoly price p_M is above marginal cost.
- Some surplus is transferred to the firm, and some disappears as deadweight loss.

Monopoly welfare and real-world examples

Welfare effects

- Consumer surplus falls because price is higher and quantity is lower.
- Part of the lost consumer surplus is transferred to the monopolist as profit.
- Another part disappears as deadweight loss.

Benchmark distortion

$$p > MC, \qquad Q < Q_C.$$

Examples

- Local utilities with strong scale economies
- Patent-protected pharmaceutical products
- Platforms with strong network effects

Typical policy responses

Antitrust enforcement, price regulation, franchise bidding, or public provision in extreme cases.

Price discrimination: the basic idea

- A firm with market power charges different prices to different buyers, groups, or units purchased.
- The goal is to extract more surplus or to expand output where uniform pricing would leave gains unrealized.

Three conditions

- The firm must have market power.
- It must be able to segment customers or induce self-selection.
- Resale or arbitrage must be limited.

First-degree price discrimination

Definition

The firm charges each buyer their maximum willingness to pay for each unit.

- In the textbook limit, all consumer surplus is extracted by the firm.
- In allocative terms, output can reach the efficient level because units are sold whenever willingness to pay exceeds marginal cost.
- Informationally, it is extremely demanding and therefore rare in pure form.

Example

Individual negotiation for cars, houses, or customized B2B contracts.

Second-degree and third-degree price discrimination

Second-degree

- The firm does not directly observe type.
- It offers a **menu** so that different buyers self-select.
- Typical forms: quantity discounts, versioning, premium vs. basic plans.

Third-degree

- The firm separates buyers into observable groups.
- It charges a higher price to the group with less elastic demand.
- Typical forms: student discounts, regional pricing, business vs. leisure fares.

Welfare aspects of price discrimination

- First-degree price discrimination can remove deadweight loss but transfers all surplus to the firm.
- Second-degree and third-degree price discrimination may increase total output relative to uniform monopoly pricing, but they usually transfer some surplus from consumers to the producer.
- The welfare evaluation depends on both:
  - how much extra trade is created,
  - and how the surplus is redistributed.

Policy point

Price discrimination is not automatically good or bad; the key question is whether the output expansion outweighs the redistribution and fairness concerns.

When there are only a few firms, supply-demand alone is not enough

Oligopoly changes the logic

With only a few firms, each firm's best action depends on what rivals do.

- A supply curve is often no longer something we can take as given.
- We need a **strategic model**: quantity competition, price competition, or sequential move games.
- This is the bridge from Chapter 2 to game theory.

Cournot model: quantity competition

- Firms choose quantities $q_1, \ldots, q_n$ at the same time.
- Market price is determined by inverse demand

$$p = p(Q), \qquad Q = \sum_{i=1}^{n} q_i.$$

- Firm i solves

$$\max_{q_i \geq 0} \ \pi_i = p(Q)q_i - C_i(q_i),$$

taking rivals' quantities as given.

Equilibrium concept

Each firm chooses its profit-maximizing quantity given the quantities chosen by its rivals.

Cournot model: a basic 2-firm example

Suppose inverse demand is

$$p(Q) = a - Q,$$

and each firm has constant marginal cost c.

Firm 1's problem

$$\max_{q_1 \geq 0} \big(a - q_1 - q_2 - c\big) q_1.$$

Profit-maximizing quantity given the other firm's choice

$$q_1^*(q_2) = \frac{a - c - q_2}{2}, \qquad q_2^*(q_1) = \frac{a - c - q_1}{2}.$$

Solving the two equations gives

$$q_1^* = q_2^* = \frac{a - c}{3}, \qquad Q^* = \frac{2(a - c)}{3}.$$

Cournot model: intuition and comparative statics

- Each firm expands output, but it knows that more total output pushes the market price down.
- Therefore the Cournot outcome is less aggressive than perfect competition:

$$p^* > c.$$

- At the same time, quantity is higher than under monopoly, because firms still compete against each other.

Main comparison

With n identical firms, inverse demand $p(Q) = a - Q$, and marginal cost c, the symmetric Cournot outcome is

$$q_i^* = \frac{a-c}{n+1}, \qquad Q^* = \frac{n(a-c)}{n+1}, \qquad p^* = \frac{a+nc}{n+1}.$$

As n grows, Q^* rises toward $a - c$ and p^* falls toward c, so the outcome moves closer to perfect competition.

Bertrand model: price competition

- Firms choose prices at the same time for a homogeneous product.
- Consumers buy from the firm with the lowest price.
- With identical marginal cost c, any price above c can be profitably undercut by a rival.

Core result

For homogeneous goods and no capacity constraint,

$$p^* = c.$$

Even a duopoly can generate the competitive price.

Extensions

Capacity constraints or product differentiation soften price competition and push Bertrand away from the competitive benchmark.

Bertrand model: why the result is so strong

- If one firm charges any price above c, a rival can undercut it by a tiny amount and capture the whole market.
- This is why homogeneous-good price competition is often much more aggressive than quantity competition.
- The textbook result is sometimes called the **Bertrand paradox**: two firms can generate the competitive price.

Why the paradox weakens in reality

Capacity constraints, product differentiation, switching costs, and repeated interaction all make undercutting less decisive.

Stackelberg model: sequential quantity competition

- One firm chooses quantity first.
- The other firm observes that choice and then picks its own quantity.
- The first firm therefore takes into account how the second firm will adjust when choosing its own output.

First-mover advantage

Moving first can change the other firm's output choice, often allowing the first mover to earn more profit than in Cournot.

Stackelberg example: solving from the second mover back to the first

Keep the same inverse demand $p(Q) = a - Q$ and constant marginal cost c.

Step 2: the second firm's optimal quantity given the first firm's choice

$$q_F^*(q_L) = \frac{a - c - q_L}{2}.$$

This is the same quantity rule as in the basic 2-firm Cournot example, with q_L playing the role of the other firm's output.

Step 1: the first firm chooses quantity

$$\max_{q_L \geq 0} \left[a - q_L - q_F^*(q_L) - c\right] q_L.$$

Substituting the second firm's quantity choice yields

$$q_L^* = \frac{a - c}{2}, \qquad q_F^* = \frac{a - c}{4}.$$

Interpretation

Compared with Cournot, the firm that moves first produces more, while the other firm produces less, and total output is higher.

Stackelberg model: interpretation

- Leader's output:
$$q_L^* = \frac{a-c}{2}.$$
- Follower's output:
$$q_F^* = \frac{a-c}{4}.$$
- Total output:
$$Q^* = \frac{3(a-c)}{4},$$
which is above Cournot but still below perfect competition.

Economic message

Moving first matters: it changes how the other firm adjusts, and this can give the first mover an advantage.

Three benchmark oligopoly models

Model	Firms choose	Timing	Main message
Cournot	Quantities	At the same time	Output lies between monopoly and competition
Bertrand	Prices	At the same time	With homogeneous goods, price can be driven down to marginal cost
Stackelberg	Quantities	One after the other	A firm that moves first can gain an advantage

Bridge to the next lecture

These models study how firms make choices when each firm's payoff depends on what the others do.

A basic ranking to remember

Typical quantity ranking

$$Q_{\text{monopoly}} < Q_{\text{Cournot}} < Q_{\text{Stackelberg}} < Q_{\text{Bertrand}} = Q_{\text{competition}}.$$

Typical price ranking

$$p_{\text{monopoly}} > p_{\text{Cournot}} > p_{\text{Stackelberg}} > p_{\text{Bertrand}} = p_{\text{competition}}.$$

Cournot with n identical firms

In the linear case, $Q_{\text{Cournot}} = \frac{n(a-c)}{n+1}$ and $p_{\text{Cournot}} = \frac{a+nc}{n+1}$, so more firms push the outcome toward competition.

Welfare comparison: competition, monopoly, oligopoly

- **Perfect competition**: under ideal assumptions, $p = MC$ and there is no deadweight loss.
- **Monopoly**: output is restricted, price rises above marginal cost, and deadweight loss appears.
- **Oligopoly**: outcomes are usually in between, and the intensity of competition depends on whether firms compete in quantities, prices, or one after the other.

Policy implication

Competition policy asks whether competition is effective, whether firms are colluding, or whether entry barriers keep the market far from the competitive benchmark.

Real-world oligopoly cases

Airlines and rail

- A few carriers dominate a route.
- Pricing reacts strategically to rivals' capacity and schedules.

Telecom and platforms

- Large fixed costs and network effects limit the number of major firms.
- Pricing, bundling, and product differentiation matter.

Autos and consumer brands

- Firms compete in price, quantity, advertising, and product design.
- Pure Bertrand or pure Cournot is too simple, but both offer useful intuition.

Lesson

Real markets mix quantity competition, price competition, capacity constraints, and repeated interaction.

End of Lecture 3

- CV and EV are money-metric welfare measures built from the expenditure function.
- Market equilibrium combines demand and supply; taxes and price controls create wedges and rationing problems.
- Perfect competition is a benchmark; monopoly and oligopoly show how market power changes price, quantity, and welfare.

Next lecture

We will study models in which each person's choice depends on what others do, and compare cases where people move at the same time with cases where they move one after another.

References

Selected references for Lecture 3

- Mas-Colell, A., Whinston, M. D., and Green, J. R. (1995). *Microeconomic Theory*. Oxford University Press, esp. Chs. 10, 12, 16, 22.
- Varian, H. R. (1992). *Microeconomic Analysis*. 3rd ed. W. W. Norton, esp. Chs. 6, 10, 13, 14, 16, 22.
- Jehle, G. A., and Reny, P. J. (2011). *Advanced Microeconomic Theory*. 3rd ed. Financial Times/Prentice Hall, esp. Chs. 4, 6, 7.
- Tirole, J. (1988). *The Theory of Industrial Organization*. MIT Press, esp. Chs. 1, 3, 5–7.