EFR summary

Microeconomics, FEB11001X

2025-2026

Lectures 1 to 18

Weeks 1 to 7

Details

Subject: Microeconomics 2025-2026

Teacher: Prof.Benoit Crutzen

Date of publication: 13 December 2025

Microeconomics - IBEB
Lecture 1, week 1

What is (micro)economics

Economics is a science that studies how people, firms, and organisations behave and make choices when there is scarcity.
Microeconomics in particular is the study of individual choices and the study of group behaviour in individual markets.
Macroeconomics is the study of broader aggregations of markets.
Economists study **how** and **why** people make choices. Therefore economics is better at explaining individual markets than the entire economy.

The Cost-Benefit approach to decisions

Economists assume that choices are made based on the cost-benefit analysis.
The main question of the cost-benefit analysis is if you should do a particular activity:
Should I do activity X?
$B(x)$ = Benefit of activity X
$C(x)$ = Cost of activity X
If $B(x) > C(x)$ you **should** do activity X.
If $B(x) < C(x)$ you **shouldn't** do activity X.
If $B(x) = C(x)$ you are indifferent.

The reservation price of activity x = the price at which a person would be indifferent about doing x and not doing x.
If $B(x) - C(x) > B(y) - C(y)$:

- $B(x) - C(x) = B(y) - C(y) + r$ (reservation price)

So there is a reservation price for which you are indifferent about doing activities x and y.

The Cost-Benefit approach to decisions

Pitfall 1. Ignoring implicit costs

Implicit costs are costs that are not explicit. This is the loss of alternatives when one alternative is chosen.

An example of this is: if you spend 3 hours watching TikTok each day, you aren't able to do anything else. For example, go to work or hang out with your friends. You might have been able to earn 50€ in those 3 hours. This way something free like Tiktok will cost you money in the form of implicit costs.
Costs and benefits are reciprocal. It doesn't matter if you subtract those 50€ from the benefits or treat them as costs. Just make sure you don't count them twice!

Pitfall 2. Failing to ignore sunk costs.

Sunk costs are costs that are beyond recovery at the time a decision is made and so should be ignored. Because these costs are beyond recovery they are irrelevant.

An example of this is: Imagine you bought a concert ticket for 50€ and it isn't possible to resell this ticket. At the night of the concert you get invited to a free party which you would enjoy more than the concert. Should you go to the party or not?
In this case, you should definitely go to the party because the 50€ for the concert is beyond recovery and therefore should be ignored. This way you should go to the party.

Pitfall 3. Measuring costs and benefits as proportions rather than absolute monetary amounts

In your decision making you should only measure costs as absolute monetary amounts and not as proportions. So it shouldn't matter if you save 10€ on a TV of 500€, or 10€ on a shirt of 20€.

Pitfall 4. Failure to understand average vs. marginal distinctions

Marginal costs = Increase in total cost resulting from carrying out one additional unit of an activity.
Marginal benefit = Increase in total benefit that results from carrying out an additional unit of an activity.
You should increase your level of activity as long as **marginal B(x) >= marginal C(x)**

Average cost = Average cost of undertaking n units activity = **Total cost / n**
Average benefit undertaking n units of activity = **Total benefit / n**

The **optimal amount** of a continuously variable activity is when **MC = MB** (Marginal cost is equal to Marginal benefit).

Different approaches to choice behaviour

Positive approach: What do people choose, and how do we declare what they choose?
Normative approach: What ought or what should people choose?

The positive approach emphasizes declaring and understanding. This is used in Science, advertising, and managing people.

The normative approach helps people to make good decisions. It is important that it's good for those people.

Not all choices are good in the Cost-Benefit analysis

Economics assumes that people are rational. That doesn't mean that all the decisions are "good". People make mistakes.
The Homo-Economicus: Stereotypical decision maker in self-interest model.

Economic agent: Individual or group making choices. A group can also be a single agent. For example, if Apple increases their iPhone prices. In this case, Apple is a single agent.

Three principles of economists

1. People make choices by **optimising**: They try to make the best choices.
2. Lots of attention goes to the **equilibrium**: a situation where no one wants to change their choices.
3. **Empirical analysis**: Economists use data to test and prove their theories.

Causality: What causes what?

Observation: a statement based on something one has seen, heard or noticed.
For example: When the sun shines, there are a lot of people on the beach.

Microeconomics - IBEB

Lecture 2, week 1

The market

Definition: A market exists of all the buyers and suppliers of a good or service.
Markets come in all forms and sizes. Place, anonymity and time are all important factors of defining the market.
Economists are among others interested in:

1. Explaining the price of a good (P)
2. Explaining the quantity traded (Q)

The **demand curve** describes the relation between the quantity of a good that demanders want to buy and the price of that good.
Law of demand: Empirical observation that if the price of a product falls, the quantity demanded increases.
$Q = f(P)$
=> $\frac{dQ}{dP} = f'(P) < 0$

Two explanations for the law of demand:

1. Increase in price: which makes people look for alternatives: substitution-effect
2. Increase in price: something has to change: income-effect

The **supply curve** describes the relation between quantity of a good that suppliers want to sell and the price.
Law of supply: empirical observation that suppliers want to sell more if the price rises.

$$Q = f(P)$$
$$\frac{dQ}{dP} = f'(P) > 0$$

Two ways of reading the supply and demand curves:
Horizontal interpretation: The quantity that is offered/demanded at a certain price.
Vertical interpretation: The price that the market will move towards at a certain quantity offered/demanded.

The equilibrium

The equilibrium quantity and price is the intersection of the supply & demand curves. This is the point where all the participants are 'satisfied'. 'Satisfied' in this case means that buyers and sellers are buying/selling the amount they want to at that price.
So: $Qs(P) = Qd(P)$

If the price-quantity pair is above the equilibrium: excess supply
If the price-quantity pair is above the equilibrium: excess demand

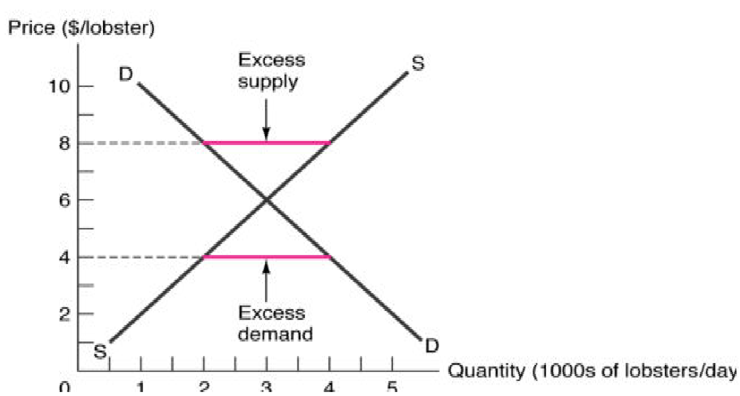

Determinants of demand/supply

Determinants of demand:

- Incomes
  - **Normal goods**: If income increases, quantity demanded increases.
  - **Inferior goods**: if income increases, quantity demanded decreases.
- Tastes
- Price of substitutes and complements
- Expectations
- Populations

Determinants of demand:

- Technology
- Factor prices (production prices)
- Expectations
- Weather

Government intervention

Government intervention often disrupts market equilibrium and ends up doing more harm than good.

An example is rental policy. In lots of cities the rent prices are too high for the poor. A reaction from politicians on this is the point system. The renting price is based on objective criteria instead of supply and demand. This results in excess demand. Another problem with this is that the poor people might prefer to use the extra money on something else, so the government would be better off giving the money to the people instead of cheaper houses.

Microeconomics - IBEB
Lecture 3, week 1

Rational choice model

A **model** is a simplified description of reality.

Structure of an economic model:

1. Description of possibilities of economic agent
2. Description of his goals.

A combination of both leads to an explanation of his behaviour.

A few assumptions need to be made for now to decide:

1. Complete information
2. One period
3. Other people don't matter

The budget curve (opportunity set)

The options of Charlie are depending on the price of beer, the price of pizza and his **monthly income**. Which makes: $M = PbB + PzZ$

Which can be rewritten as a line by writing $B = \frac{M-PzZ}{Pb}$

It's slope coëfficiënt = -(Pp/Pb) and can be interpreted as that if Charly wants to buy one extra pizza he should sacrifice Pp/Pb beers.

In reality instead of the choice being between two goods the choice might be between millions of goods. This can be written as:

$$M = \sum_{i=1}^{N} PiGi$$

Imagine instead of wanting to be looking at all goods we want to look at one good, and another good which consists of all other goods. This is the '**Composite good**'. The price of the composite good = 1.

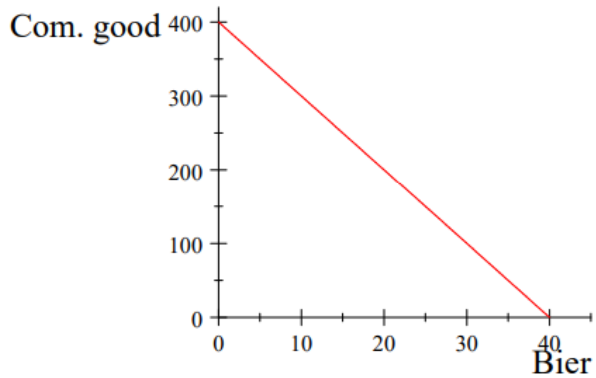

Description of his goals

We describe what an agent wants with help of preference orders: A scheme by which a person organises alternatives based on desirability.
For the preference orders there are a few assumptions:

1. **Completeness:**
   For any two bundles A and B, the consumer can compare them:
   A is preferred to B, or B is preferred to A, or the consumer is indifferent.
2. **More is Better (Monotonicity):**
   If a bundle A has *more of at least one good* and *no less of the other*, then A is preferred to C.
3. **Transitivity:**
   If A is preferred to C and C is preferred to D, then A is preferred to D.
4. **Continuity:**
   If a bundle A is only slightly different from B, preferences don't jump suddenly.
   Meaning: if A is preferred to B, then bundles *very close to A* are also preferred to B. (This ensures smooth indifference curves.)
5. **Convexity:**
   If the consumer is indifferent between A and C, then any mixture (combination) of A and C is at least as good as A or C.
   This is why indifference curves are **bowed inward** (people like **variety**).

A few exceptions on transitivity:

- Football games: if Feyenoord wins from Ajax, And Ajax wins from PSV it doesn't necessarily mean Feyenoord will win from PSV.
- Collective preference orders based on majority (votes)

The indifference curve

An indifference curve is a set of bundles between which a consumer is indifferent.

Marginal rate of substitution (MRS) is the slope coëfficient of the indifference curve: Which means how much of product A you are willing to give up for an extra product B.

Indifference map:

- Unlimited amounts of indifference curve
- The higher the curve, the higher the utility
- Different persons have different indifference curves and therefore also different indifference maps
- Indifference curves cannot cross each other.

Two ways to learn about indifference curves:

1. Asking questions about all sorts of bundles.
2. Statistics (How do income & prizes affect a choice)

How to maximise utility

The goal of the consumer is to maximise its utility (trying to reach the highest indifference curve)
The highest indifference curve is reached when the MRS (slope coëfficient of indifference curve) is equal to the slope coëfficient of the budget constraint.
So when: $MRS = \frac{Px}{Py}$

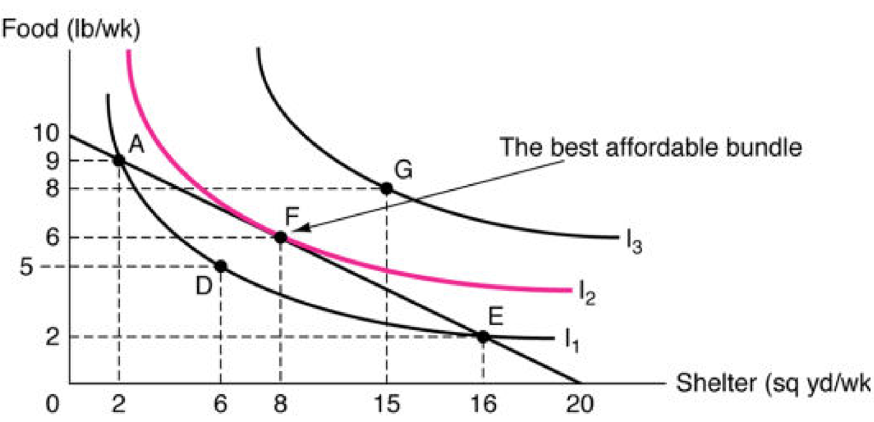

Microeconomics - IBEB
Lecture 4, week 2

Rational choice model into mathematics

The structure of the economic model can be written as:
max. U (x, y) s. t. y = f(x)
U (x, y) describes the indifference map (all the indifference curves)
y = f(x) describes the budget constraint.
Maximising this will make sure all the resources are used optimally.

The utility function is described as U = U(x, y) where U stands for utility.
Utility is something positive, so therefore we want to maximise our utility.
More-is-better implies:
$\frac{\partial U}{\partial Y} = UY > 0; \frac{\partial U}{\partial F} = UF > 0$
Which makes sure the marginal utility is greater than zero.

If you write the differential of U(F, S) it can be written as: $MRS = \frac{dF}{dS} = -\frac{Us}{Uf}$

Ordinal vs. marginal utility

Ordinal utility: People can say bundle A is better than bundle B but the number U (utility) has no meaning in itself.
Cardinal utility: The number in itself has meaning because you can compare people's utility number against each other.

Economics of happiness makes use of cardinal utility:
People from different countries were asked how happy they are on a scale of 4. An interpretation of this scale of 4 is that they were asking on which utility scale people are.
The answers got used to decide how happy people were, and it concluded that richer countries reported a higher utility than poorer countries.

So what is **important for utility**:

- Relative income position
- Marital status: A divorce is compensated by 100.000 Euros higher income.
- Work versus no work
- Security: Risk aversion
- Children makes you deeply unhappy

Solving the optimisation problem

We solve the optimisation problem with two methods:

1. The **Lagrange method**
2. The **substitution method**

The lagrange method:

1. Set up the lagrange function: $L(x, y) = U(x, y) - \lambda(ax + by - m)$
2. Set up the first order conditions:
   a. $Ux(x, y) - \lambda a = 0$
   b. $Uy(x, y) - \lambda b = 0$
   c. $ax + by = m$
3. Divide 2a and 2b: $\frac{Ux(x,y)}{Uy(x,y)} = \frac{a}{b}$
4. Solve the rest of the variables with this answer.

The substitution method:

1. Rewrite: $ax + by = m$ into for example: $y = \frac{m-ax}{b}$
2. Substitute y into U(x, y) so that makes $U(x, \frac{m-ax}{b})$
3. Take the derivative of U.
4. Optimise the derivative of U.

You can also solve problems with even more variables by setting up the first-order derivatives by putting all the first partial derivatives and the budget constraint together.

Microeconomics - IBEB
Lecture 5&6, week 2

Analysing the demand curve

Individual demand curve: how does the quantity demanded of a person change when prices differ?
Another name for the individual demand curve is the **price-consumption curve**(PCC). It shows for every price the best bundle (in relation to the different budget constraints and the indifference map).
A price increase will always make a person worse off, because he will reach a lower indifference curve.

Engel curve, Q = f(M), shows the relation between income and the quantity demanded.
With normal goods: f'(M) > 0
With inferior goods: f'(M) < 0

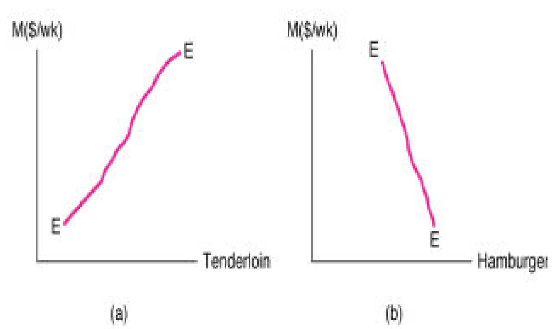

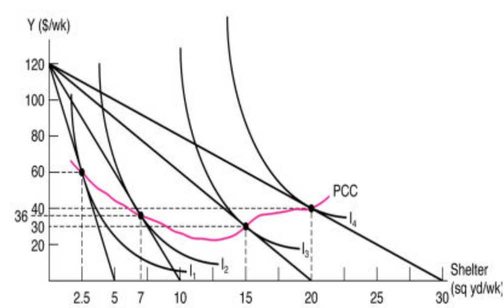

Effect when the price changes

There are two reasons why the quantity demanded changes when the price increases:

1. The **income-effect**: Lower real income (Engel curve)
2. **Substitution-effect**: Looking for alternatives.

To distinguish substitution effect and income effect: shift new budget equation to the point where the initial utility can be achieved.

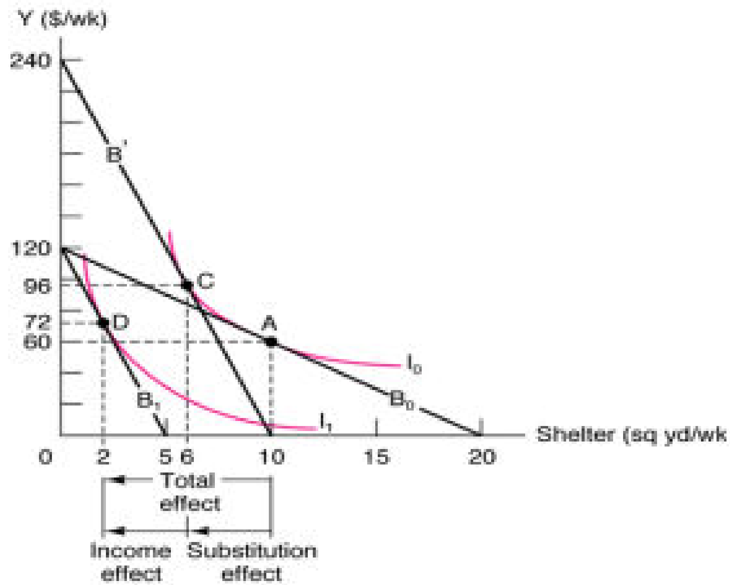

The difference between C and D in shelter is the income effect.
The difference between A and C in shelter is the substitution effect.

Government taxing

When the government set a tax on a product there will be changes in the budget curve of the consumer and will therefore change the utility the consumer is reaching.
We want to make the tax as efficient as possible:
A situation is inefficient if it is possible for the government to
achieve higher returns while the consumer receives the same utility.

When putting a tax on a product the budget curve will change:
$PxX + PyY = M \rightarrow (Px + t) + PyY = M$
When putting a tax on income the budget curve will change:
$PxX + PyY = M \rightarrow PxX + PyY = M - T$

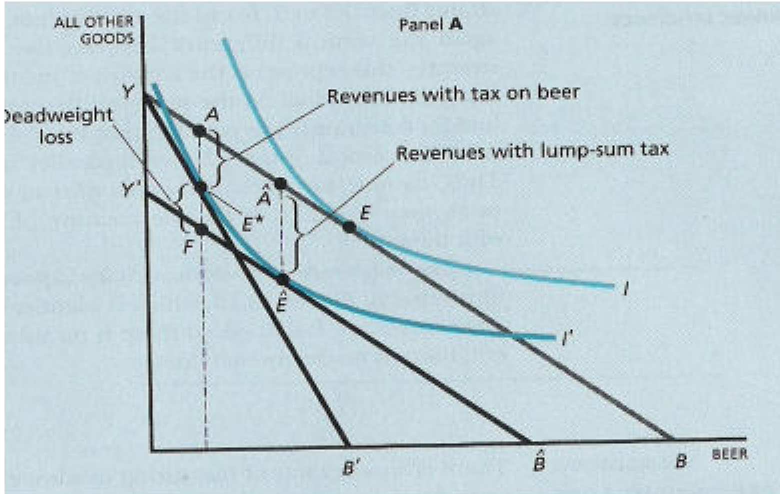

Income tax will reach higher returns while the consumer receives the same utility. Therefore not changing the VAT rate on a product and taxing income is more efficient. The tax on for example alcohol can be justified for other reasons, like health.

Individual curve to market curve

The market demand curve is the sum of the individual demand curves.
Summing the individual demand curves will give the market demand curve.
An example:

$P = 16 - 2Qa \; -> \; Qa = 8 - \frac{1}{2}p$

$P = 8 - 2Qb \; -> \; Qb = 4 - \frac{1}{2}p$

$P <= 8 \; -> \; \Sigma Qi = Q = 12 - P \; -> P = 12 - Q$

$P > 8 \; -> \; P = 16 - 2Q$

Price Elasticity (ε)

Managers would like to know how strong the demand reacts to the price.
Price elasticity = the resulting percentage change in quantity of a percent change in price.

$\varepsilon = \frac{dQ}{dP} * \frac{P}{Q}$

There are three possibilities for **price elasticity**:

$\varepsilon < \; -1$: **Elastic demand**, demand decreases more than one percent when price increases one percent.

$\varepsilon = \; -1$: **Unit-Elastic demand**, demand decreases one percent when price increases one percent.

$\varepsilon > -1$: **Elastic demand**, demand decreases less than one percent when price increases one percent.

Income elasticity = the resulting percentage change in quantity of a percent change in income.
$\eta = \frac{dQ}{dM} * \frac{M}{Q}$
Three possibilities for income elasticity:
$\eta < 0$ -> **inferior goods**
$0 <= \eta <= 1$ -> **Normal goods**
$\eta > 0$ -> **Luxury goods**

Price compensation

Price compensation means that you have so much more income that you can buy the same bundle you used to buy when prices were lower.
Price compensation will lead into higher utility, because the budget curve will go over the old indifference curve you used to reach. -> See illustration

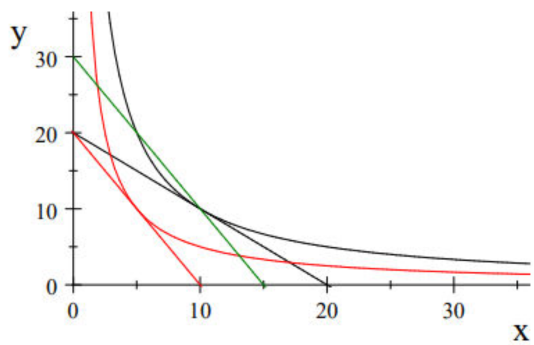

Microeconomics - IBEB
Lecture 7, week 3

The Intertemporal Consumption Choice Model

Up till now we made the assumption that the consumer spends his entire income. In reality consumers also save money and get loans. Therefore the question rises: how does a consumer distribute his consumption over time?

Therefore we make use of two periods: today and the future.

We also assume that there are no initial assets and no inheritances.

Intertemporal Choice

- Trade-off between consuming today (C_1) and consuming tomorrow (C_2).
- Consumer has income M_1 today and M_2 tomorrow (no inflation ⇒ prices constant).
- Consumption decisions involve saving or borrowing.
- Normalise $pC_2 = 1$.

Intertemporal Opportunity Cost

In chapter 4, the opportunity cost between goods x and y was $-Px/Py$.
Here, the cost of consuming today vs tomorrow is the interest rate **r**.

The gross opportunity cost of consuming today is $(1 + r)$, and thus, the slope of the budget line = $-(1 + r)$.

For $r \geq 0$, the line is steep → borrowing is costly.
To transform tomorrow's M euros in today's euros: $M \rightarrow M / (1 + r)$
To transform today's M euros in tomorrow's euros: $M \rightarrow M \times (1 + r)$

Intertemporal Budget Constraint

If today's consumption C_1 is on x-axis and C_2 on y-axis:

- If all income is consumed tomorrow: $C_2 = M_2 + M_1(1 + r)$
- If all income is consumed today: $C_1 = M_1 + M_2 / (1 + r)$
- Slope of budget constraint: $-(1 + r)$

We can read the budget constraint from either today's or tomorrow's perspective.

The intertemporal budget constraint in today's euros: $C_1 + C_2 / (1 + r) \leq M_1 + M_2 / (1 + r)$
The intertemporal budget constraint in tomorrow's euros: $(1 + r)C_1 + C_2 \leq (1 + r)M_1 + M_2$

Note: If the borrowing rate > lending rate, the constraint has a kink at the endowment (M_1, M_2).

Intertemporal Indifference Curves

Preferences over (C_1, C_2) are written as $U(C_1, C_2) = \bar{U}$.
The slope (MRS) is known as the Marginal Rate of Time Preference (MRTP):

$$MRTP = -U_{C_1}/U_{C_2}$$

Consumer maximises: $$\max_{C_1, C_2} U(C_1, C_2) - \lambda\left(C_1 + \frac{C_2}{1+r} - M_1 - \frac{M_2}{1+r}\right)$$

Optimum condition: $$MRTP(C_1^*, C_2^*) = -(1+r)$$

This condition is also known in macroeconomics as the Euler equation

Patience and Saving

If $C_1^* < M_1$ → consumer is not spending all his present income; He saves: $S_1 = M_1 - C_1$
Tomorrow, he can consume: $C_2 = M_2 + (1 + r)(M_1 - C_1)$.

- Patient if: MRTP < (1 + r).
- Impatient if: MRTP > (1 + r).

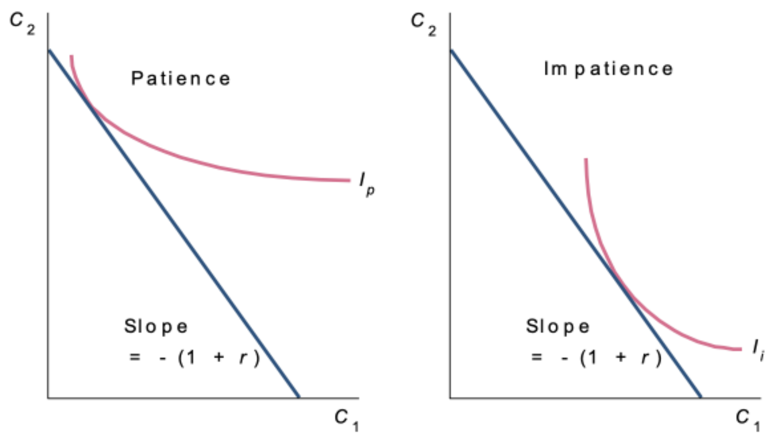

Modelling Patience in Utility

- Utility function with time-preference parameter β: $U(C_1, C_2) = u(C_1) + \beta u(C_2)$.
- u(·) increasing and concave
- β = gross rate of time preference = $1 / (1 + \delta)$; δ = net rate of time preference.
- Rewriting the Euler equation: $UC_1/UC_2 = \beta (1 + r) = (1 + r)/(1 + \delta)$.

Cases:

- 1. $r = \delta \Rightarrow C_1^* = C_2^*$ (equal consumption)
- 2. $r > \delta \Rightarrow C_1^* < C_2^*$ (patient consumer)
- 3. $r < \delta \Rightarrow C_1^* > C_2^*$ (impatient consumer)

Effects of Changes in r, M_1 and M_2

If r increases the opportunity cost of present consumption increases
If r decreases the opportunity cost of present consumption decreases
Substitution effect: $C_1 \downarrow$ (both borrowers and lenders).
Income effect: Borrower $C_1 \downarrow$, Lender $C_1 \uparrow$.

Therefore, the total effect is ambiguous.

Change in M_1 or M_2:

1. Change in current income (M_1) → C_1 and C_2 rise less than 1-for-1.
2. Change in future income (M_2) → C_2 and C_1 rise less than 1-for-1.
3. Permanent change ($\Delta M_1 = \Delta M_2$) → C_1 and C_2 increase almost equally (if $\delta \approx r$).

Friedman's Permanent Income Hypothesis

- Permanent income: total income in present value, $Y_p = M_1 + M_2/(1 + r)$.
- Permanent income changes → reflected one-for-one in consumption.
- Temporary income changes → mostly saved.
- We consume permanent income but save temporary income.

Microeconomics - IBEB
Lecture 8, week 3

Probability Theory

Deterministic and Random Variables

- Deterministic variable: has a fixed, certain value (e.g. income = €1000).
- Random variable: value not known ex-ante; represents a lottery or gamble.

Examples:

- Coin flip → Heads/Tails ($p = ½$ each)
- Dice roll → 1 to 6 ($p = 1/6$ each)
- Lottery draw → 1 to N ($p = 1/N$ each)

Events and Support

Event: specific outcome of a random variable.
Support: all possible outcomes.

Discrete: finite set (e.g. dice, coin).
Continuous: infinite set within an interval (e.g. GDP, income).

Probabilities and Probability Distributions

- Probability $p = Pr(X = x)$ lies between 0 and 1.
- The sum of all probabilities = 1 ($\Sigma p_i = 1$).
- For a continuous variable, $Pr(X = x)=0$ because there are infinite outcomes.
- Use a Probability Density Function (PDF) to describe likelihoods.

Probability Density Function (PDF)

- Describes how likely each value of X is.
- For continuous X: $\int f(x)dx = 1$.

Standard Normal Distribution: Bell-shaped, symmetric around 0

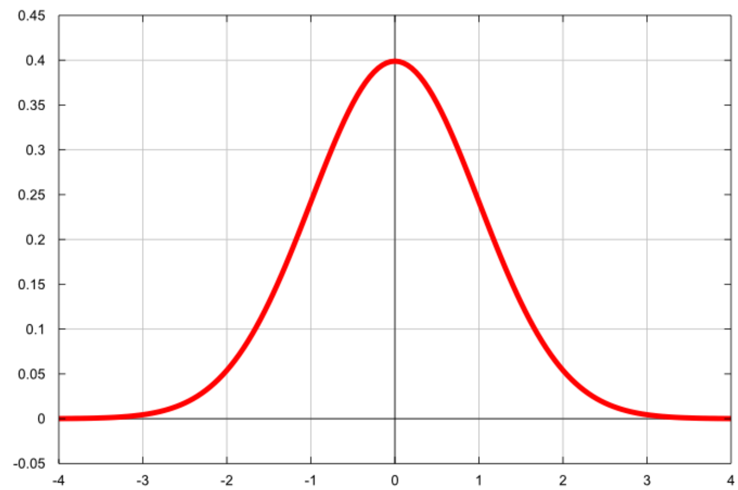

Discrete Distributions

- Represented by frequency histograms showing how often each value occurs.
- Height of each bar = probability or frequency

Cumulative Distribution Function (CDF)

- Gives probability that $X \leq x$: $F(x) = Pr(X \leq x)$.
- Always increasing; $F(-\infty)=0$, $F(+\infty)=1$.
- Example (dice): $Pr(X<3)=1/3$.
- Standard Normal CDF = S-shaped curve from 0 to 1

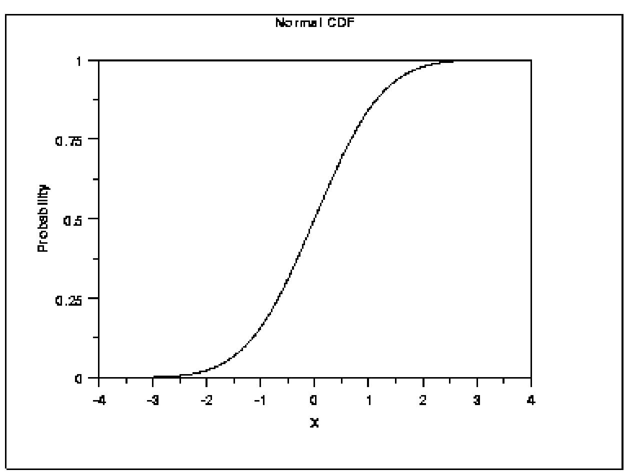

Expected Value (Mean)

- Weighted average of all possible outcomes.
- For discrete X: $E(X)=\Sigma\, p_i x_i$
- For continuous X: $E(X) = \int x\, f(x)dx$

Examples:

- Dice → E(X)=3.5

Uniform Distribution

A uniform distribution is a continuous distribution where all outcomes are equally likely.

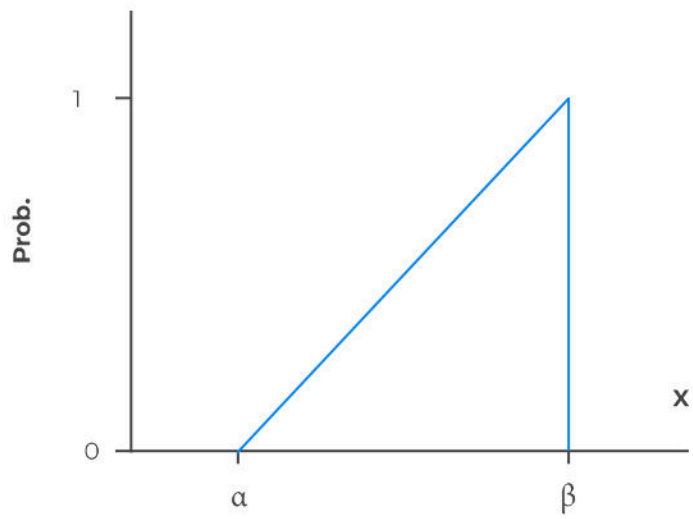

Support = [a, b]

- PDF: $f(x)=1/(b-a)$
- CDF: $F(x)=(x-a)/(b-a)$
- Mean: $E(X)=(a+b)/2$

Conditional Expectations (Uniform Distribution)

If X ~ Uniform[a, b]:

- $E(X \mid X < x)=(a+x)/2$
- $E(X \mid X > x)=(x+b)/2$

Key Notes

- Uncertainty is represented by random variables.
- PDFs and CDFs describe distributions of possible outcomes.
- Expected value summarises average outcome.
- The uniform and normal distributions are foundations for analysing expected utility and decision-making under risk in upcoming lectures.

Microeconomics - IBEB
Lecture 9, week 4

Uncertainty in economics

Preferences under Uncertainty

In the real world, individuals face uncertain outcomes — wealth, income, or returns can vary.

Von Neumann–Morgenstern Expected Utility Model

To model behavior under uncertainty, economists use the Von Neumann Morgenstern (VNM) expected utility model.

The VNM utility function assigns utility to each possible outcome of a random variable (a gamble).

Note: Agents rank uncertain prospects by their expected utility, not by expected money.

Diminishing marginal utility ($U'' < 0$) is *not always assumed* — it depends on risk attitude.

The Expected utility of a gamble (E(U)): expected value of utility over all possible outcomes of the gamble.

$$E[U(x)] = \sum p_i U(x_i)$$

Attitudes Toward Risk

1. Risk Aversion

- Utility: increasing and concave (U′ > 0, U″ < 0).
- Prefers the utility of the *average outcome* to the average utility of outcomes:

$$U[E(X)] > E[U(X)]$$

- Always refuses a fair gamble.
- Graph: curved upward, flattening as wealth increases

Interpretation: prefers certainty; will accept a lower but guaranteed return to avoid risk.

2. Risk Neutrality

- Utility: linear in money, U(M)=αM
- Indifferent between certain outcomes and fair gambles.

$$U[E(X)] = E[U(X)]$$

- Typical of firms or investors focused on *expected profits only*.

3. Risk-Loving (Risk-Seeking)

- Utility: increasing and convex (U″ > 0).
- Prefers extremes to average

$$U[E(X)] < E[U(X)]$$

- Accepts fair gambles since upside is valued more than downside is disliked.

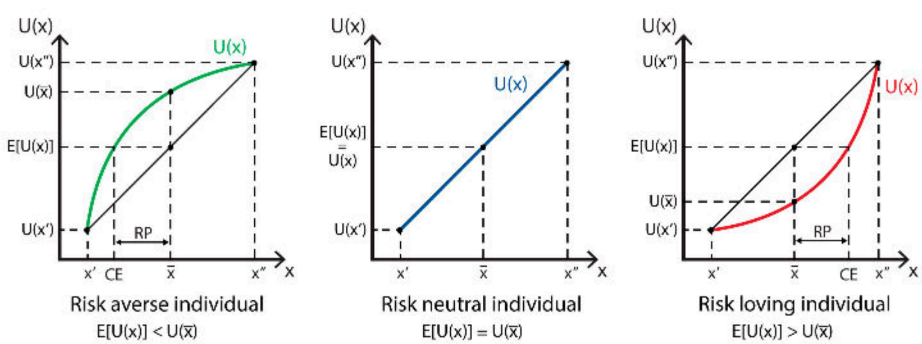

Certainty Equivalent (CE) and Risk Premium

For a risky lottery, the certainty equivalent (CE) is the guaranteed income level yielding the same utility as the expected utility of the gamble:

$$U(CE) = E[U(X)]$$

For a risk-averse agent: CE < Expected Value (EV).

The risk premium is:

$$\pi = EV - CE$$

It measures how much one is willing to pay to remove risk.

Insurance

- Risk-averse individuals are willing to pay a premium to avoid uncertainty.
- Full insurance removes risk completely: income becomes independent of accidents.

An example of what you could be insuring is your health. Let's say there's a ⅔ chance you stay healthy and a ⅓ chance you get sick. Your expected utility will be EU=⅔ x 600+ ⅓ x 150 = 500. We are also assuming you are risk averse. The amount you are willing to pay for insurance = M, if you stay healthy, - M corresponding to the expected utility which lies on the same indifference curve.

The Value of Information – Advisors and Experts

- Risk-averse individuals value information because it reduces uncertainty.
- They are willing to pay for advice that improves decision quality.

Example:
A must choose between:

- Becoming a tennis player: 0.01×€250M if successful, 0.99×€0.2M if not.
- Opening a pub: guaranteed €1M

Utility: U(M)= $\sqrt{M}$

Without advice:

$$EU(Tennis) = 0.01\sqrt{250M} + 0.99\sqrt{0.2M} < \sqrt{1M}$$

→ Chooses the pub.

If a perfect advisor reveals the true outcome beforehand:

$$EU(with\ advice) = 0.01\sqrt{250M - P} + 0.99\sqrt{1M - P}$$

Setting $EU(with\ advice) = 1000$ gives $P_{max} \approx €276,690.$

He'd pay up to that much for perfect information.

If advisor only charges when news is good:

Theoretical $P_{max} \approx 249M$ -> showing that "pay on success" can extract almost all returns.

Takeaway: Risk-averse people will pay for information, but they should beware of biased incentives.

Bias and Advice Example – Parents and Dating

- You consider dating someone; the outcome's quality qqq is uncertain (distributed over [-μ, μ]).
- Your parent knows *q* but has a **bias** (cost *c*<0) against you dating.
- They advise "yes" if c + q > 0 → q > -c

Interpretation:

- A *positive* recommendation from a biased parent is highly credible.
- A *negative* one may cause you to miss good opportunities.

Moral Hazard and Adverse Selection

Moral Hazard

Occurs **after** a contract is made — one party's behavior changes because the other bears the risk.

Examples:

- Drivers being less careful once insured.
- Workers being lazy when effort isn't monitored.

Adverse Selection

Occurs before a contract — one side hides private information about risk or quality. Example: unhealthy people are more likely to buy health insurance ("market for lemons").

Signaling and Information Disclosure

When preferences are misaligned, signaling helps reveal private information.

1. Costly-to-Fake Principle:
    - For a signal to be credible, it must be costly or difficult to imitate.
    - Examples: advertising, elite education, reputation, luxury goods.

2. Full-Disclosure Principle:

    - Silence implies the worst — individuals disclose even unfavorable information to avoid suspicion.
    - Examples:
        - Firms offering long warranties (Toyota 5 years, Kia 7 years).
        - Applicants disclosing facts that might otherwise appear hidden.

Microeconomics - IBEB
Lecture 10, week 4

Adverse Selection

Adverse selection = problems that arise when one party in a market transaction has more information than the other (hidden information).
This often leads to market failure.

Applications

A. Professional Guilds

- Professions like doctors, lawyers, accountants, and teachers often have barriers to entry (licenses, degrees, guilds).
- These barriers reduce adverse selection by ensuring quality — screening out unqualified individuals.

B. Finding a Partner

- Dating involves hidden information: people don't fully know each other's quality or preferences.
- People use signals to reduce uncertainty — joining dating platforms or revealing info (education, job, hobbies).
- Being on a dating site itself is a signal.
- Universities also act as matching markets — e.g., graduates of law programs are much more likely to marry each other (Eika et al., 2019).

Search Theory

A crucial question in search theory is when do you stop looking and when do you go into action.

Imagine you are looking for a job and every job pays between 100 and 200 euros. Looking for another job costs you 10 euros each time, and the salary is uniform divided. So 1 euro extra from 100 euros is worth as much as 1 euro extra from 199 euros.

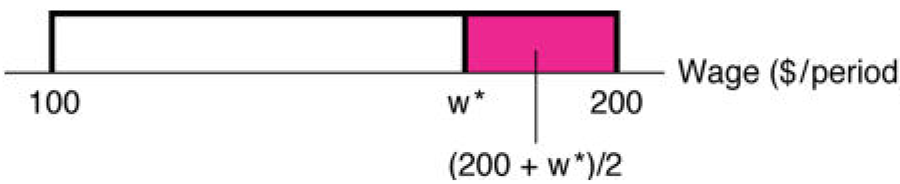

The benefit of looking for a new job is: $Pr(W > W^*)[E(W|W > W^*) - W^*]$. Where Pr(W>W*) stands for the chance that W* > W, so the chance that you find a higher result by searching.

The expected gain in the case if W* > W is: [E(W|W > W*)-W*).
If you are looking for the amount where you should stop searching is: $\frac{200-W^*}{2} * \frac{200-W^*}{200-100} = 10(price\ of\ looking)$. A more general formula can be written as $\frac{max-W^*}{2} * \frac{max-W^*}{max-min} = price$

Statistical Discrimination

Employers make judgments based on group averages when individual productivity is uncertain.

Example:

Group productivity is uniformly distributed between 10 and 30 → average 20.
An employer offers wage = average productivity = €20/hour.

If a worker takes a test that is correct only half the time:

- Test score = 24 → estimated productivity = (24 + 20)/2 = 22
- Test score = 16 → estimated productivity = (16 + 20)/2 = 18

The employer updates expectations between the test score and the group mean. Two workers with identical scores but from different groups (different average abilities) get different offers.

This creates statistical discrimination — differences in expected productivity across groups lead to wage gaps.

This explains why university reputation and program choice affect pay — your "signal" acts as insurance against uncertainty.

However, above-average individuals are initially underpaid, and below-average ones are overpaid.

Producer theory

The definition of products is anything which supplies utility, now or in the future. (not only physical goods).

Definitions of production are: a process that creates utility, now or in the future. Or a process that production factors (inputs) turns into products (outputs)

Examples of inputs are: labor, capital, land, energy, raw materials, entrepreneurship, and knowledge. Output is anything that supplies utility, now or in the future.

The production function: Input -> Business/production function -> output
In this course the only inputs (production factors) we are gonna use are labour, which we describe as L, and capital, which we describe as K.

We can write the production function as Q = F(K, L), where Q is the amount of products we produce. Function F is the production technology.
An example of a production function is the Cobb-Douglas production function: $Q = mK^a L^b$ with $a,\ b \epsilon [0,\ 1)$ and $m > 0$

Long vs. short run

The producer chooses the production factors (inputs). Some choices are able to be changed quickly, these are variable inputs, other choices aren't possible to change quickly, these are fixed inputs.

In the **short-run** there is at least 1 production factor 'fixed'.
In the **long-run all** production factors are 'variable'.

In this course Labor 'L' is always variable and Capital 'K' is fixed.

Production in the short-run

If we have a production function F(K, L)=KL and capital is fixed in the short run to K = K_0 = 1 we can illustrate that (see illustration). Later we change K = K_1 = 2, therefore the production function also changes.

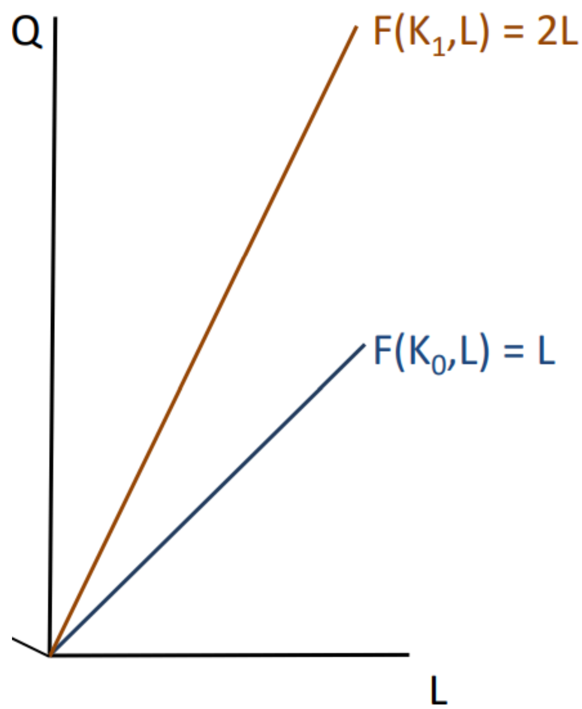

The short-run production function always goes through the origin.

Total, average and marginal product

Total product (TP):

- How much gets produced: Q (production function)

Average product (AP):

- Output per unit variable input
- The average product of labour: $AP_L = \frac{Q}{L} = \frac{F(K,L)}{L}$

Marginal product (MP):

- How much changes the output with the change of 1 unit of an input.
- The marginal product of labour: $MP_L = \frac{dQ}{dL} = \frac{\partial F(K,L)}{\partial L}$

Effects of labour in the short-run:

If there is one person he will do everything. If there are multiple people there will be specialisation (Adam Smith), and if there are too many people they get in each other's way. This effect is called the law of diminishing returns.

A more formal definition of the law of diminishing returns is: a principle stating that profits or benefits gained from something will represent a proportionally smaller gain as more money or energy is invested in it.

The law of diminishing returns, works for every short-run input although it doesn't work in the long-run. For example if we multiply our restaurant (long-run because it also affects fixed costs) and place it on the other side of town it won't necessarily give less returns than the first restaurant.

Properties of the marginal product of labour (MP_L):

- Slope coefficient of the total product (TP)
- Rises for L if L is small (specialisation)
- Decreases if L is bigger (law of diminishing returns)
- TP: The inflection point is at the start of diminishing returns
- The MP can be negative for L if L is big (people get in each other's way)
- The start of negative returns is when $MP_L = 0$

See the illustration for these properties illustrated.

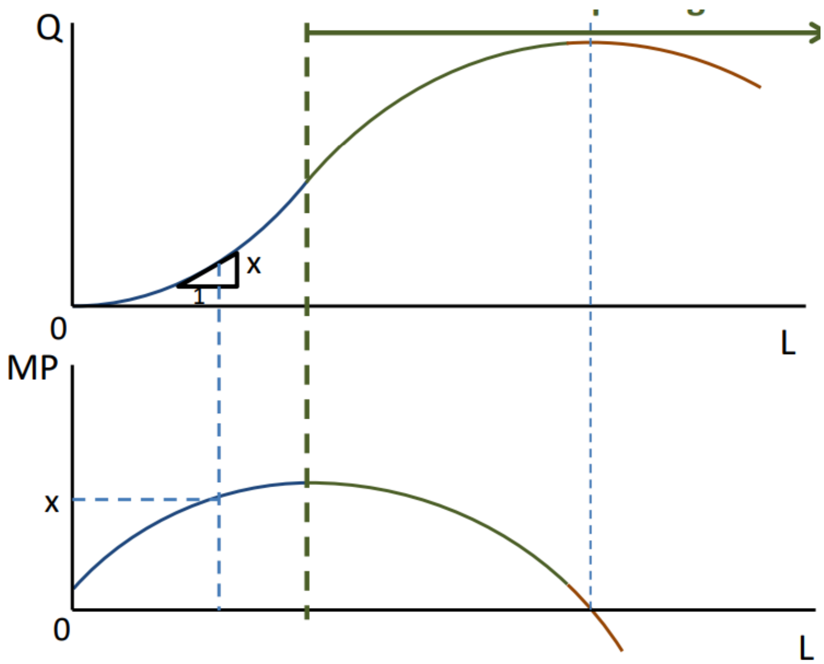

Properties of the average product:

- The coefficient of the origin to a point on the production curve
- When L -> 0: $AP_L = MP_L$
- MP > AP => AP increases
- Maximum AP_L : $AP_L = MP_L$ (not the point in the origin)
- MP < AP => AP decreases

If we want to maximise production in the short-run with certain amounts of labours we should make sure $MP_{L1} = MP_{L2}$

Production in the long-run

In the long-run all production factors are variable. The production function will be described as Q = F(K, L). We can show this graphically with an isoquant map. This isoquant map consists of unlimited amounts of isoquant curves for which the amount of production is the same, regardless of the distribution of the production factors.

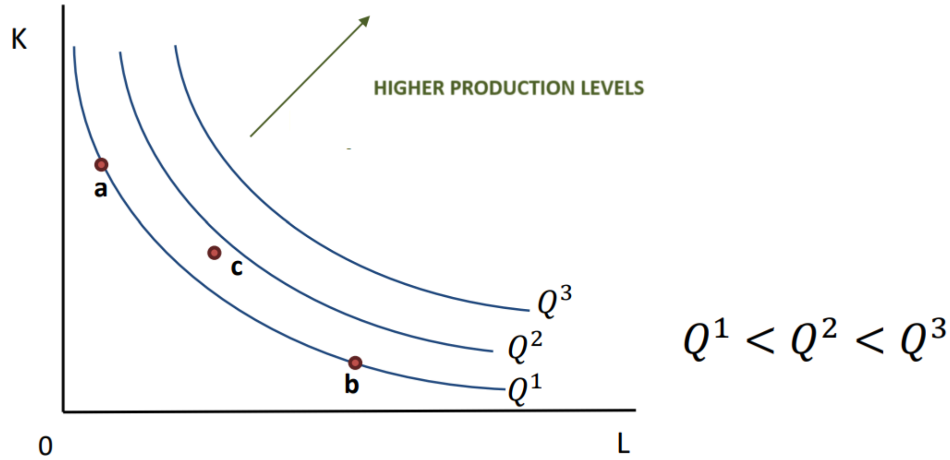

Perfect substitutes and complements also work the same for the isoquant map. These types of functions are illustrated below:

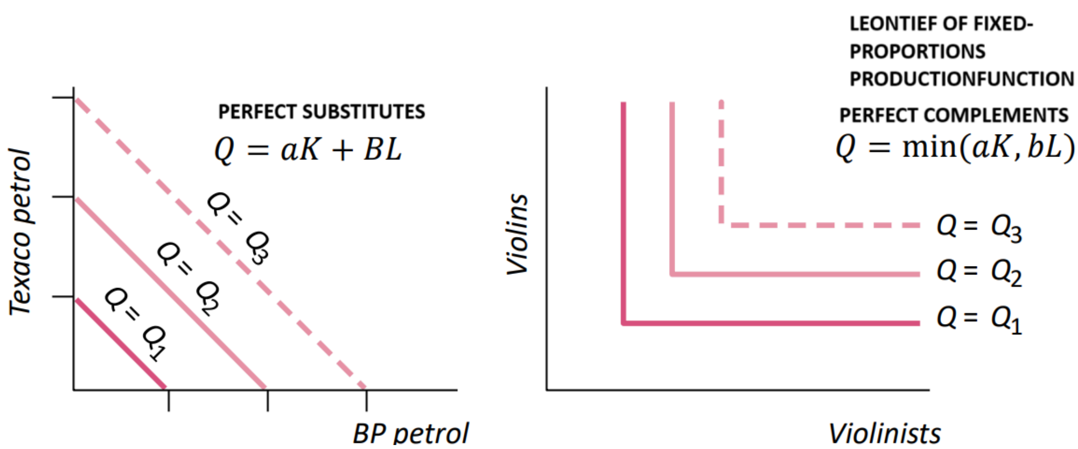

The marginal rate of technical substitution (MRTS): is the absolute value of the slope coefficient of the isoquant.

MRTS = $|\frac{dK}{dL}|$

The economic interpretation of the MRTS is the ratio to which capital can be exchanged for labour without changing the production quantity

Returns to scale

Returns to scale is about what happens with the production when you increase the production factors (inputs) proportionally. So both with the same proportion

- Increasing returns to scale
  F (cK, cL) > cF(K, L)
  This can happen because of specialisation, or the law of big numbers.

- Constant returns to scale
  F(cK, cL) = cF(K, L)
- Decreasing returns to scale
  F(cK, cL) < cF(K, L)
  This can happen when people for example get in each other's way while working. This is not the same as the law of diminishing returns.

Example exercise: Decide if the production function $F(K, L) = K^{1/4}L^{1/2}$ has increasing, constant or decreasing returns to scale.
Solution: F(cK, cL) = $c^{1/4}K^{1/4}c^{1/2}L^{1/2} = c^{3/4}K^{1/4}L^{1/2} < cF(K,\ L)$ so this function has decreasing returns to scale.

Microeconomics - IBEB
Lecture 11, week 5

Short-run costs

In the short term we have:

- **Fixed costs (FC)**, a synonym for this is overhead costs.
  Where the cost of capital is defined as 'r' and fixed capital defined as K_0: FC = rK_0. You pay fixed costs even if you produce nothing.
- **Variable costs (VC)**, depends on how much you produce. The hourly wage we define as 'w', the hours worked we define as L_1 , and the output is defined as Q_1.
  $VC_{Q1} = wL_1$
- **Total costs (TC)**, this is the sum of fixed and variable costs
  $TC_{Q1} = rK_0 + wL_1$

How do we calculate the fixed costs. Let's say we have $F(K, L) = K^{1/2}L^{1/2}$, K_0 = 4, r = 2
The fixed costs are FC = rK = 4*2 = 8.
The quantity of labour: $Q(L) = K^{1/2}L^{1/2} = 2L^{1/2} <=> L(Q) = (Q/2)^2$
$TC_Q = FC + VC_Q$ = 8 + Q^2

Marginal costs is the slope coefficient of the total costs.
$MC(Q) = \frac{dTC(Q)}{dQ} = \frac{d(FC+VC(Q))}{dQ} = \frac{dVC(Q)}{dQ}$

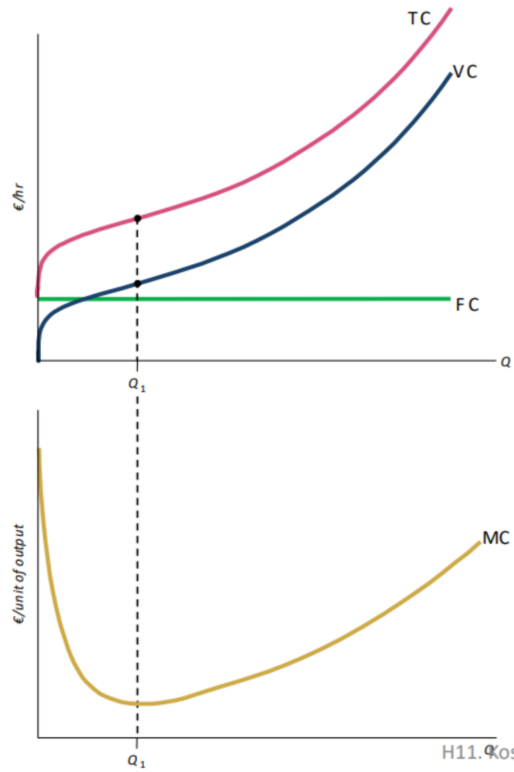

So we see that the MC is also the slope coefficient of the variable costs. If we want to illustrate this in a picture we can see that the inflection point of the TC and VC is the minimum point of the MC.

We can rewrite the marginal cost function also as $MC(Q) = \frac{w}{MPL}$. Here we can see that the slope of the TC and VC is inversely proportional with the slope of the TP. The inflection point of the functions will be the same, although the TP will diminishingly increase and the total cost will increase more than before. This works with the rules of diminishing returns. Which we talked about in the previous lecture.

Average costs:

- **Average Total Costs (ATC)**: $ATC(Q) = \frac{TC(Q)}{Q}$
- **Average Variable Costs (AVC)**: $AVC(Q) = \frac{VC(Q)}{Q}$
- **Average Fixed Costs (AFC)**: $AFC(Q) = \frac{FC}{Q}$

The relation between marginal/average costs:
$MC < ATC$: ATC $decreases$
$MC > ATC$: ATC $increases$
MC $hits$ ATC in its $minimum$

$MC < AVC$: AVC $decreases$
$MC > AVC$: AVC $increases$
MC $hits$ AVC in its $minimum$

Optimal allocation of short-term costs (minimizing costs):
The optimum for short-term cost allocation is when **MC_1=MC_2**

I will illustrate this with an example: imagine you have two factories with two total cost functions. $TC_1(Q_1)=10(Q_1)^2 + 10$, $TC_2(Q_2)^2=5(Q_2) + 20Q_2 + 3$, $Q_1 + Q_2 = 31$

We can rewrite the last function as: $Q_1 = 31 - Q_2$
If we solve $TC_1(31-Q_2)=TC_2(Q_2)$ we will reach our answer. You can try this for yourself at home. The answer will be $Q_1 = 11$ and $Q_2 = 20$

Relation Between Productivity and Marginal Costs

$$MC = \frac{\partial VC}{\partial Q}, \quad MPL = \frac{\partial Q}{\partial L}$$

If labour is the only variable input:

$$VC = wL \Rightarrow MC = w\frac{\partial L}{\partial Q} = \frac{w}{MPL}$$

Long-run costs

All inputs are variable in the long run:
Firms choose the optimal mix of capital (K) and labour (L) that minimizes cost for a given output or maximizes output for a given cost.

Isocost lines are like the budget curves of firms. Isocost line:

$$C = rK + wL$$

The slope coefficient of the line will be *-w/r.* It works almost the same as the consumer's budget curve, see the illustration below.

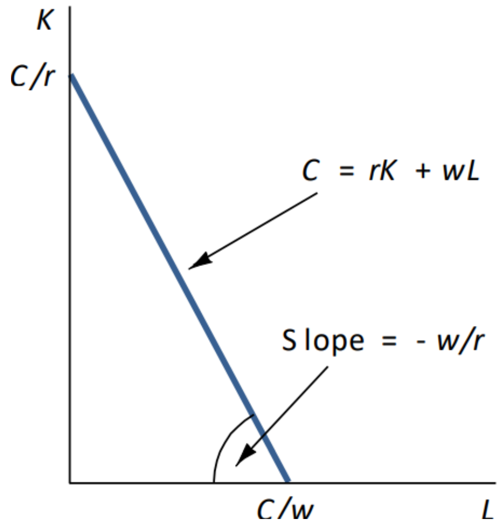

Isoquant: combinations of (K, L) yielding the same output.

The **maximum output for a certain costs** is where $MRTS = \frac{w}{r}$. This equation can be rewritten as $\frac{MPl}{MPk} = \frac{w}{r}$. This works exactly the same when we are looking for the **minimum cost for a certain output**. It occurs where the isoquant is tangent to the isocost line: See the difference in the illustration below.

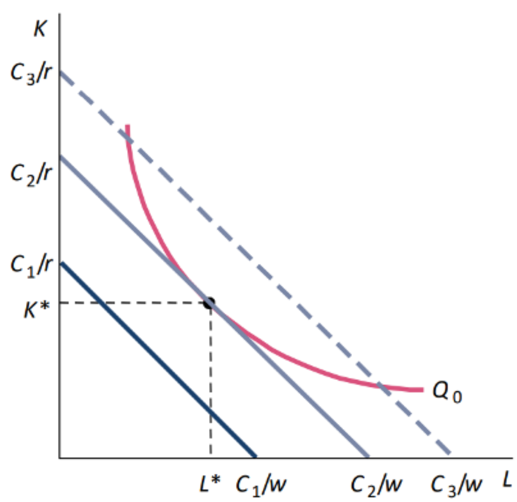

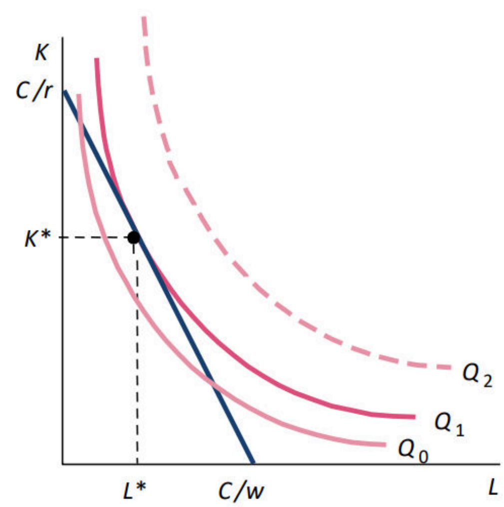

The left illustration is the minimum cost for a certain output and the right illustration is the maximum output for a certain cost. Keep in mind that solving this works the same as solving for example the consumer problem. If you struggle with this I recommend looking back to that part of the summary.

The **output expansion curve** is the curve which measures the optimal cost allocation (most Q, quantity) for each amount of total costs.

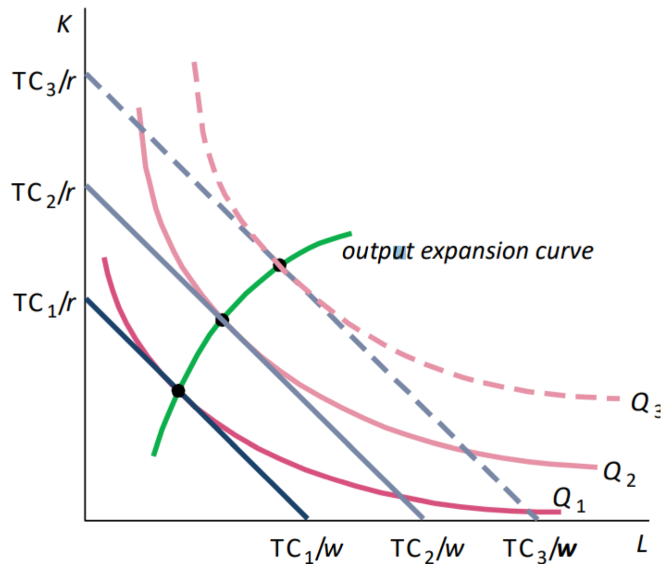

- The **long-term total costs (LTC)** will always go through the origin because in the long-term all the costs are variable.
- The **long-term marginal costs**: $LMC = \frac{dLTC}{dQ}$ (slope of LTC)
- The **long-term average costs**: $LAC = \frac{dLTC}{dQ}$ (avg. cost per unit)

Returns to Scale and Cost Curves

Returns to Scale	**LTC Shape**	**Market Type**
CRS	Straight line through origin	No optimal size (many firm sizes possible)
IRS	Concave curve	Natural monopoly
DRS	Convex curve	Many small firms

LAC is the lower envelope of all **ATC** curves for different fixed capital levels.

Recap of Firm Problems

Production Maximisation:

The firm wants to choose K and L to maximise output given a spending budget C_0

$$\max_{K,L} F(K,L) - \lambda(C_0 - rK - wL)$$

The first-order conditions give the tangency condition:

$$\frac{MPL}{MPK} = \frac{w}{r}$$

Cost Minimisation

The firm wants to choose K and L to minimise total cost while still producing a required output Y_0:

$$\min_{K,L} (rK + wL) - \mu(Y_0 - F(K,L))$$

The first-order conditions again give:

$$\frac{MPL}{w} = \frac{MPK}{r} \iff \frac{MPL}{MPK} = \frac{w}{r}$$

Both problems yield the exact same optimal ratio of K and L.

Even though one problem maximises output with a cost constraint, and the other minimises cost with an output constraint, they both lead to the identical condition.

What is a monopoly?

A monopoly is a market where there is a **Monopolist**.
A **Monopolist** is the only supplier of a good and all close substitutes.
What are causes of a monopoly?

1. **Exclusive access to production inputs** (For example railways)
2. **Increasing returns to scale.** Therefore the LAC curve is downward shaped. If a firm produces more their average cost will go down. A **Natural Monopoly** is a market where it is most efficient in terms of costs to let one firm serve the entire market.

3. **Patents**: Patents help protect inventions during a few years/decennia. This helps to promote innovation. If there weren't any patents it would be possible to steal other firms' ideas.
4. **Network Effects**: A product gets more valuable if more people use it. For example operating systems or social media. This is strictly seen a form of returns to scale.
5. **Monopoly Licenses granted by the government**: For example restaurants next to the motorways.

A monopoly faces a **downward-sloping demand curve** → price-maker.

Profit Maximising by the monopolist

Remember profit is $\pi(Q) = TR(Q) - TC(Q)$. In the case of the monopolist, they can change the price, so therefore the price isn't the market price anymore. Now we define the P in TR(Q) as the market demand. Therefore:

$$\pi = TR - TC = P(Q)Q - C(Q)$$

Profit is maximised where MR(Q)=MC(Q) & P>AVC (remember from perfect competition). P<AVC is the firms shutdown condition.

In the figure below is illustrated that the marginal cost and marginal revenue curve intersect twice. So where is profit maximised. For this we have the second order condition that the MR intersects the MC from above. This can mathematically be written as $\frac{dMR(Q)}{dQ} < \frac{dMC(Q)}{dQ}$.

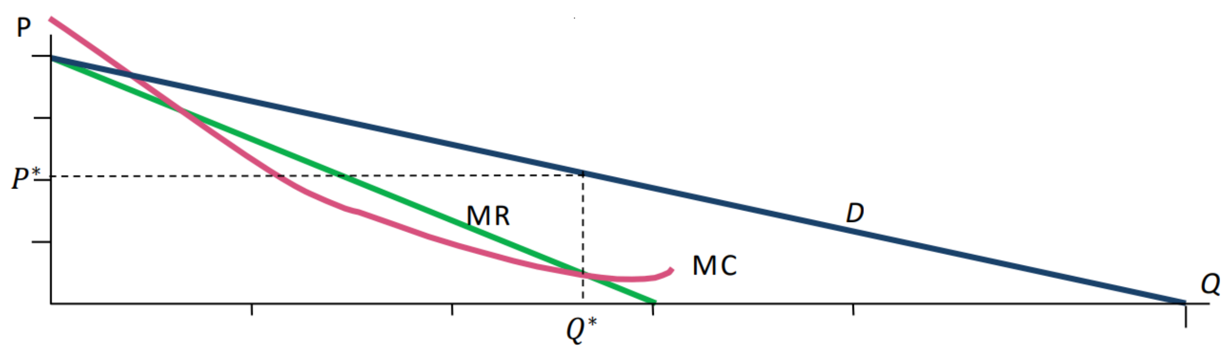

Below I illustrated an exercise for profit maximising.
Demand curve: P(Q)=100-2Q
Total costs: $TC(Q) = \frac{1}{2}Q^2$

What is the profit for an optimising monopolist? What is the price?

Answer:
$TR(Q) = 100Q - 2Q^2$
$MR(Q) = 100 - 4Q$
$MC(Q) = Q$
$MR(Q) = MC(Q) \Leftrightarrow 100 - 4Q = Q \Leftrightarrow Q = 20$
$P(20) = 100 - 2 * 20 = 60$
$\pi = 100x20 - 2x20^2 - (\frac{20^2}{2}) = 1000$

MR in terms of Elasticity
The marginal revenue curve can be written as $MR(Q) = P(Q) + \frac{dP}{dQ}Q$.
This can be rewritten as $MR(Q) = P(Q) + P(Q)\frac{dP}{dQ}\frac{Q}{P(Q)} = P(Q) * (1 - \frac{1}{|\varepsilon|})$. Where the ε is the elasticity of demand. This can help us understand that $0 = MR(Q) \Leftrightarrow \varepsilon =- 1$. Thus, the monopolist never produces on the inelastic part of demand.

The mark-up of a monopolist: $\frac{P-MC}{P}$. How much higher are the prices in comparison to the marginal costs? This says something about the market power a monopolist has. In perfect competition the markup is 0, because the producers have no profit.

Microeconomics - IBEB
Lecture 12, week 5

Price Discrimination & Perfect Competition

Monopolists do not always stop at simple MR = MC profit maximisation.
If possible, they attempt to increase profits further by charging different prices to different consumers or for different quantities.
This behaviour is called price discrimination.
It doesn't always harm consumers — sometimes welfare may increase

First-Degree Price Discrimination (Perfect PD)

Here, the monopolist charges each consumer their exact willingness to pay for every unit. The monopolist extracts all consumer surplus.

The marginal revenue curve becomes identical to the demand curve.

The optimality condition becomes:

$$D(Q) = MC(Q) \quad \Rightarrow \quad p^* = MC$$

This is the generalisation of the two-part tariff strategy from earlier in the course.
But: in reality, firms don't know each consumer's WTP → mechanism design is needed.

Two-Part Tariffs with Heterogeneous Consumers

If the monopolist **knows** there are high-demand and low-demand types, but **does not know which consumer is which**, direct perfect PD is impossible.

Solution: Offer two different contracts

- One designed to **squeeze the low-demand consumer** as much as possible
- One designed so the **high-demand consumer prefers their own contract** over the low-demand one (incentive compatibility)

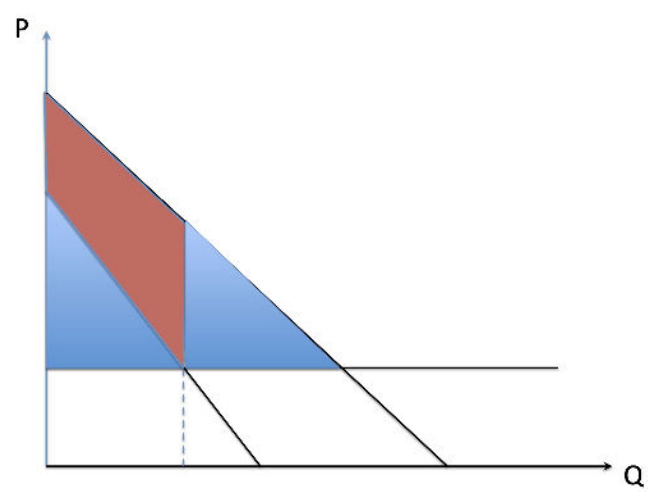

- Blue area = surplus extracted from low-demand type
- Red area = surplus extracted extra from high-demand type
- Low-demand side gets squeezed — high-demand must get a “good deal” to self-select.

Second-Degree Price Discrimination (Quantity-Based Pricing)

Also called **nonlinear pricing**:

The monopolist offers a **menu of prices** for different quantity ranges

: $$0 < Q < Q_1 : P_1 > P_2 > P_3$$

Used when the monopolist cannot identify consumer types, but can vary pricing by quantity.

Third-Degree Price Discrimination (Market-Based Pricing)

The monopolist sells to **different markets** at **different prices**:

- Different countries
- Different cities

Assume two markets with demands P1(Q1), P2(Q2), and one cost function C(Q1+Q2).

Profit maximisation problem:

$$\max_{Q_1,Q_2} P_1(Q_1)Q_1 + P_2(Q_2)Q_2 - C(Q_1 + Q_2)$$

FOCs:

$$MR_1(Q_1^*) = MC(Q_T^*)$$

$$MR_2(Q_2^*) = MC(Q_T^*)$$

Thus:

$$MR_1 = MR_2 = MC$$

(Price is higher in the market with more inelastic demand (lower |ε|)).

Example:

Market 1: $P_1 = 100 - Q_1$
Market 2: $P_2 = 80 - Q_2$
Cost: $C = \frac{1}{2}(Q_1 + Q_2)^2$

Solving FOCs gives:

- $Q_1^* = 55/2$
- $Q_2^* = 35/2$
- $P_1^* = 145/2$
- $P_2^* = 125/2$
- Profit = 2075

If discrimination not allowed, markets merge: Profit falls to 2025.

Price discrimination increases monopoly profit.

Why market segmentation matters

If two markets have **identical demand**, discrimination gives **no gain**.
If they differ, discrimination raises profits.

Hurdle Model of Price Discrimination

Discrimination via **quality tiers** or "hurdles" consumers must cross:

- Airline seat classes
- Hotel room tiers
- Train ticket types
- Coupons & rebate vouchers
- Slow refund processes

Low-elasticity customers **don't bother** with the hurdle → pay more.
High-elasticity customers **jump the hurdle** → pay less.

Transition to Perfect Competition

After discrimination, the lecture shifts fully into the opposite market structure: perfect competition.

Conditions for Perfect Competition

1. Standardised product
2. Firms are price takers
3. Perfect mobility of production factors in long run
4. Perfect information for firms & consumers

Short-Run Profit Maximisation in Perfect Competition

Since firms take price as given:

$$TR = P \cdot Q \quad \Rightarrow \quad MR = P$$

Profit maximisation:

$$P = MC$$

Shutdown Condition

Profit: $\pi = Q(P - ATC)$

Firm should continue producing even with losses as long as:
$P > AVC$

Shutdown if:
$P < AVC$

Breakeven:
$P = ATC$

Firm’s Short-Run Supply Curve

A perfectly competitive firm’s supply curve is: MC above AVC

Industry supply = horizontal sum of individual supply curves.

Short-Run Competitive Equilibrium

Market supply *S* intersects market demand *D* to give:

Market price: P^*
Total quantity Q^*
Each firm produces where $P^* = MC(q_i^*)$

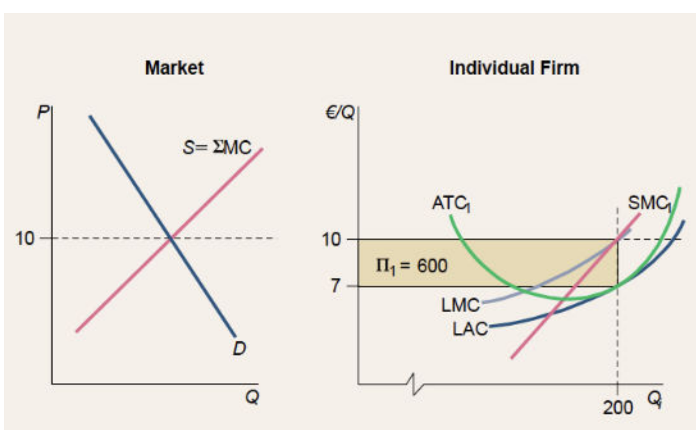

Long-Run Competitive Adjustment

If firms make losses:

Exit market
Market supply ↓
Price ↑
Losses disappear

If firms make profits:

Entry
Market supply ↑
Price ↓
Profits return to 0

Long-run equilibrium conditions:

$$P = MC = \min LAC$$

All surviving firms are identical; inefficient firms exit.

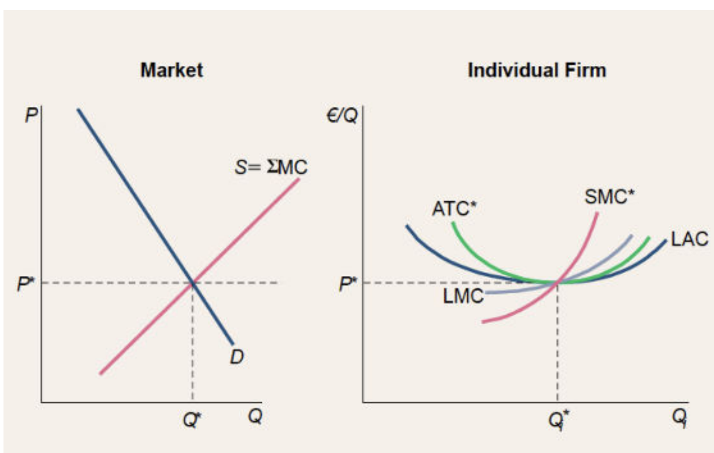

Price discrimination increases monopoly profit by tailoring prices to demand segments.

Perfect competition is the polar opposite of monopoly: firms take price as given and earn zero economic profit in the long run.

The long-run equilibrium in perfect competition is ultra efficient.

Microeconomics - IBEB
Lecture 13, week 5

Producer Surplus (PS)

Producer Surplus = the monetary benefit a firm gets from producing:

$$PS = TR - VC$$

- PS ≠ profit unless Fixed Costs = 0.
- Graphically: area above the supply curve (= MC) and below the price line.

At the market level, aggregate PS is shown as the lower triangle under price, above the market supply curve.

Comparing Perfect Competition and Monopoly

Illustration: firm with constant marginal cost.
Under monopoly, output is where MR = MC, which is lower than competitive output.
This generates a deadweight loss (Harberger triangle) — labelled W in the slides.

Key point:

- Monopoly reduces total surplus relative to perfect competition.

- But first-degree price discrimination eliminates the deadweight loss, because the monopolist produces the efficient quantity (even though consumers get no surplus).

Government Regulation of Monopolies

Two ways mentioned:

A. Theoretical (but unrealistic) solution

1. Allow first-degree price discrimination
2. Charge a lump-sum tax
3. Redistribute to society

B. Realistic policy

1. Estimate monopolist's cost curve
2. Impose a regulated price
3. Let the firm choose its output

Regulating digital/IT companies is more complex than classic natural monopolies.

Why price supports are a bad idea

A price support (a government-set minimum price):

- Allows inefficient firms to survive
- Prevents market forces from reallocating resources
- Leads to misallocation and excess supply
- Better policy: help workers/firms transition rather than freeze market prices

Efficiency effects of taxes

A unit tax (per-unit tax) creates:

- Higher consumer price
- Lower producer price
- Lower quantity
- Deadweight loss (the triangle between S and D that disappears when tax is imposed)

The slides show multiple layers of the triangle to illustrate revenue vs lost surplus.

If D or S is perfectly inelastic, deadweight loss disappears → reason to tax goods with inelastic demand/supply.

Introduction to General Equilibrium

Up until now = Partial equilibrium
→ We fixed prices, ignored interactions across markets.

General equilibrium examines:

- Multiple consumers
- Multiple goods
- Market-clearing prices
- Allocation efficiency

Prices now perform allocative (not just rationing) functions.

Simple Exchange Economy (2 × 2 × 2 Model)

- 2 consumers: A and B
- 2 goods: Clothing (C) and Food (F)
- No production → economy is pure exchange
- Consumers are endowed with bundles (“manna from heaven”)

Example endowment R = (CA, FA) for A
Then B’s endowment is what remains.

Edgeworth Box

A graphical tool to analyse exchange:

- Width = total clothing
- Height = total food
- A’s origin = bottom-left
- B’s origin = top-right
- Any point = allocation between A and B

Indifference Curves in the Edgeworth Box

- A’s indifference curves are convex (as usual) toward A’s origin.
- B’s indifference curves are also convex, but toward B’s origin (top right).
- Further from each origin = higher utility.

Gains from Trade

If A and B start at endowment R:

- If their indifference curves intersect at R → trade is possible.
- Example: A trades +10 food for +15 clothing and moves to T.
  - A is happier (higher indifference curve).
  - B is happier (also higher indifference curve).

They continue trading until no further mutually beneficial exchange exists.

Tangency = Pareto Optimum

Mutually beneficial trades continue until A's and B's indifference curves are tangent → point M (Pareto efficient).

Characteristics:

- Both consumers strictly prefer M to the initial endowment R.
- No further trade can make one better off without making the other worse off.

Microeconomics - IBEB
Lecture 14 - week 6

Partial vs general equilibrium

The **partial equilibrium analysis** focuses on the market for one product. This market is isolated from the rest of the economy.

The study of broader interactions in the economy is called the **general equilibrium analysis**. For example a cycle in which businesses demand labour and capital supplied by households and supply products to the households. The households supply labour and capital, and demand products.

We'll firstly focus on a trade economy without production.
We assume there is a production with only consumers.

- 2 Consumers: A & B
- 2 Consumption Goods: C(clothing) and F(food)
- $C_A+C_B=C$; $F_A+F_B=F$

The **allocation** is every combination of C & F in possession of A & B.
For this we huge the **Edgeworths box**. We will set the initial endowments $R_A=(70, 75)$, $R_B=(130, 25)$ up in the Edgeworths box just as illustrated below.

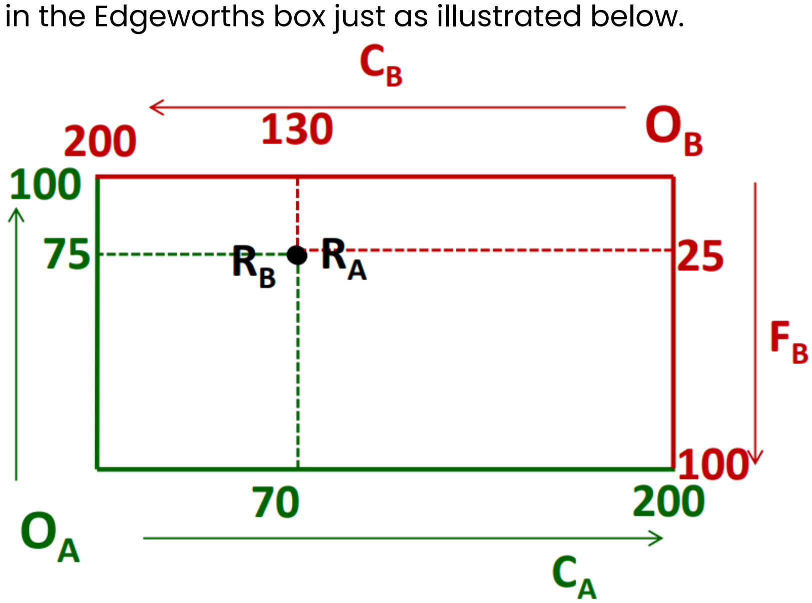

In this box is every possible allocation of F & C between A & B.
Now we can also set the indifference curves of both consumers in the edgeworth boxes. Although they will be mirrored because of the box's nature. These indifference curves will have the same assumptions as talked about in an earlier lecture (for example, convexity, transivity etc).

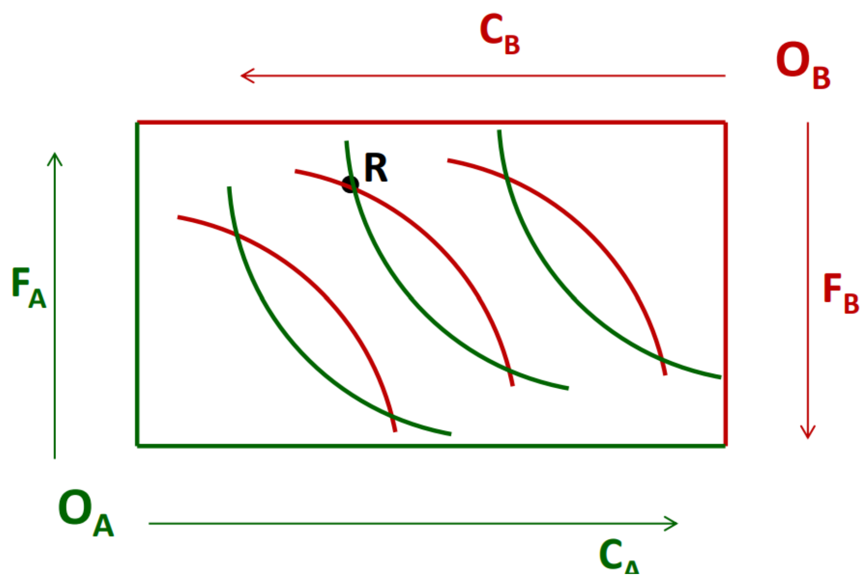

At the moment R is the current allocation of F & C. You can see the indifference curves crossing R make some sort of ellipse shape. Every point in this ellipse shape is **Pareto preferred**. This ellipse is called the "**eye of Pareto**".

Pareto superior = Allocation that at least one individual prefers and the others like at least as well.
Pareto preferred = The movement to a pareto superior allocation.

The two consumers will trade so that the point R will move in the eye of pareto. They can't trade outside of the eye of pareto, because one player would gain a disadvantage and they wouldn't trade. This continues until there is no eye of pareto left.
This is where $MRS_A = MRS_B$. So where the indifference curves touch eachother. This point is the **Pareto optimal** distribution.

The **contract curve** is a set of all Pareto optimal distributions, so a line of all touchpoints of the indifference curves touch.

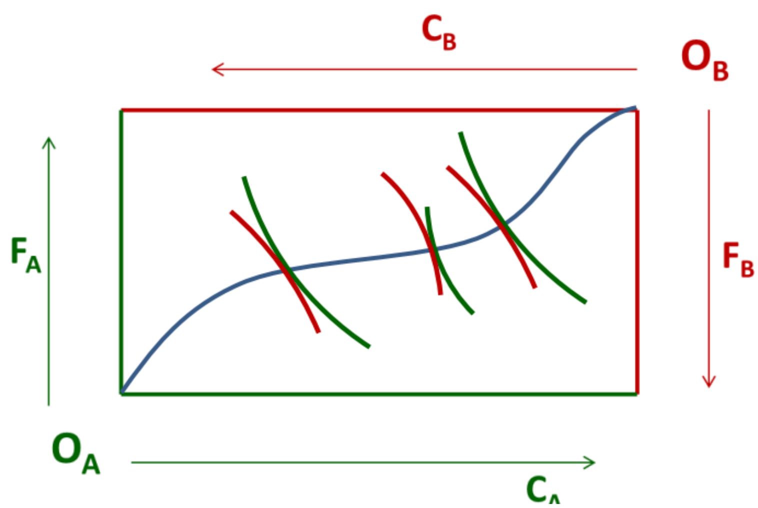

So the contract curve is a curve of all the points where $MRS_A=MRS_B$ with $C_A+C_B=C$ and $F_A+F_B=F$. When an allocation isn't on the contract curve both will want to trade until they reach the contract curve.

Consumers get on the contract curve by:

- Negotiation and exchange
- Using a third good: money
- Trade for given relative prices

1. We can insert the budget curve in the Edgeworths box. The curve of the Edgeworths box is $-(P_C{}^0/P_F{}^0)$
2. Then we will move the budget curve until $P_C/P_F=MRS_A=MRS_B$.
3. Keep in mind the prices are relative as we can change the budget curve.

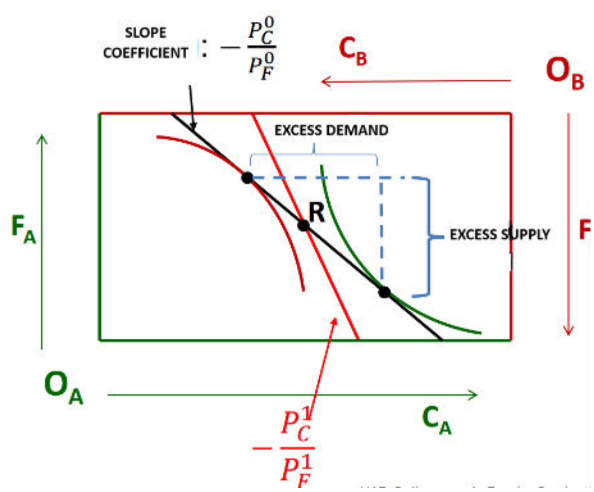

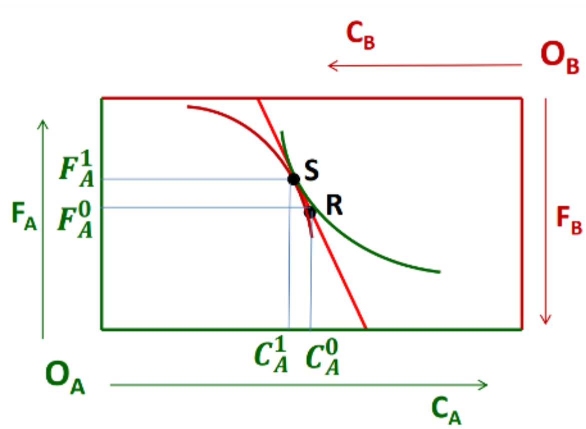

The two welfare theorems

The **first welfare theorem**: If there are complete markets (for every good which agents care about there is a market) and all agents are price-takers, then the market equilibrium is always Pareto-efficiënt.
The **second welfare theorem**: If there are complete markets and all agents are price-takers, then every allocation on the contract curve can be reached as equilibrium.

If consumer A got all the clothing and all the food in the system it also is a pareto efficient equilibrium. No consumer can reach a better position without making one other worse off.

The general equilibrium with production

This model is called the 'Simple' model. It is the model we talked about in the start of this lecture.

- Two price-taking firms: 1 & 2
  - 1 produces F (food)
  - 2 produces C (clothing)
- Two production-inputs: K(capital) & L(labour)
- Two consumers: A & B
  - A in possession of K_A & L_A
  - B in possession of K_B & L_B
  - $K_A+K_B=K$ & $L_A+L_B=L$
- Prices depend on preferences, scarcity, technology
- Price: w(labour), r(capital), P_C(clothing), P_F(food)
- $M_A=wL_A+rK_B$
- $M_A=wL_A+rK_B$

We can set up an Edgeworths box with isoquants for the producers.

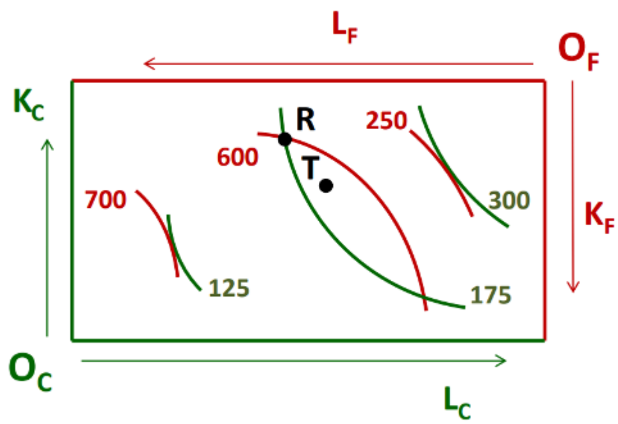

We can set up a contract curve for the producers as well. This is where the MRTS(Marginal rate of technical substitution/$\frac{MPl}{MPk}$) of both firms touch each other on the entire map.

The optimal mix for each business is where MRTS = w/r as illustrated a few lessons before. All the combinations on the contract curve lead to the **production possibility frontier**. This is each level of production for each point on the contract curve. See the illustration below:

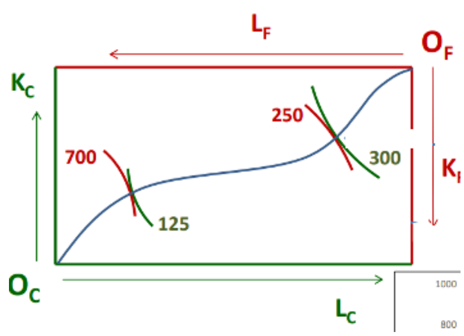

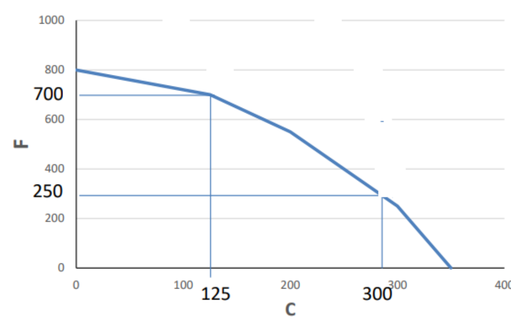

The slope of the possibilities frontier is the **marginal rate of transformation (MRT):**
MRT=$\left|\frac{dF}{dC}\right| = \frac{MCc}{MCf}$

In the most efficient production mix there are a few conditions:

- $MRT = \frac{MCc}{MCf}$
- $Pc = MCc$, P_F=MC_F
- MRS=P_C/P_F
- MRS=MRT
- MRTS=w/r
- P_{Prod}=MR_{prod}=MC_{prod}
- MRT=P_C/P_F

Microeconomics - IBEB
Lecture 15&16, week 6

Game theory

Remember from one of the first lectures on Pareto efficiency from Adam Smith: "Striving to self-interest leads to Pareto-efficiency". His main argument for this is that competition between producers leads to better products, cheaper production, higher profits, decreasing prices. This proces repeats itself. Competition out of self-interest is good for consumers.
A condition for Pareto efficiency is that all the agents are priceaccepters.

External costs and benefits are costs and benefits which are not carried by the person undertaking the activity.

Modelling strategic interactions via gametheory is very useful we will see.
The building blocks of a game:

- **Players** (N >= 2): They make decisions
- **Rules**: Who can and when choose what?
- **Payoffs**: What do players receive when everyone makes the choices.
- **Information**: Who knows what and when. We assume complete information.

A classic example of game theory is the prisoner's dilemma:

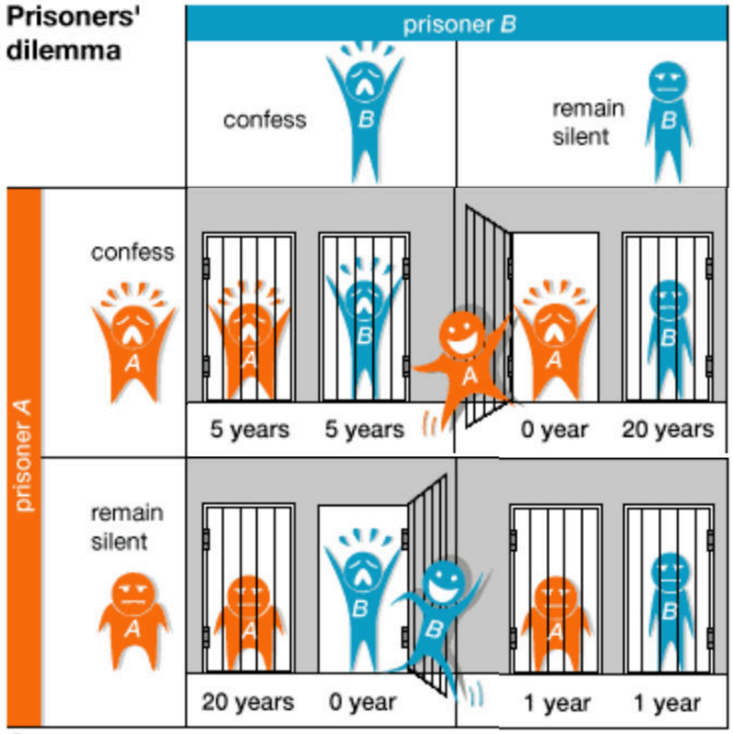

In gametheory we often write this as:

Game		Prisoner A	Prisoner A
		Confess	Remain Silent
Prisoner B	Confess	-5, 5	0, 20
Prisoner B	Remain Silent	-20, 0	-1, =1

Where the number on the left always is for the player on the left and the result on the right is for the player on top.
We can see that the combination (Confess, Confess) isn't **Pareto efficient**. Because both players can better their position without harming anothers position, see the outcome for (Remain silent, remain silent). Although confessing is the **Dominant Strategy** for both. It is individual optimal to betray in both situations (-5>20 and 0>-1).

A **dominant strategy** gives a higher outcome regardless of the opponents strategy. A strategy gets dominated if another strategy is always better.
A **social dilemma** is a type of game where individual incentives conflict with efficiency. The reason for this is external costs and benefits (see the prisoners dilemma). This problem can disappear if we look at perfect competitive markets with lots of agents, where the individual has no impact.

We can also have a 3x3 game where we can use **iterative elimination** of strictly dominated strategies to find the equilibrium.

			Player B	
		Left (L)	Center (C)	Right (R)
	Top(T)	12,20	16,0	-10,1
Player A	Middle(M)	9, 10	15, 100	-11, 3
	Bottom(B)	15, 1	12, 2	-2,3

Row M gets strictly dominated by row T. (12>9; 16>15; -10>-11) therefore player A never chooses row M. Now we can compare column C and column R. Player B will never choose column C (1>0; 3>2). Therefore we can eliminate that row. Now we have four options left ((T, L); (T, R); (B, L); (B, R)). Player A will never choose row T (15>11; -2>-10). THerefore player B will choose R because 3>1. Therefore the equilibrium is (B, R)

Repeated interaction

When we play multiple times the same game strategies can help.
An example of an strategy is the tit-for-tat strategy:

- mn
- From then you repeat the action of your opponent in their previous turn (coöperate/defect)
- You should let your opponent know that this is your strategy.

Qualities of the tit-for-tat strategy is that it is friendly, strict and forgiving.
Although there is a very big problem with the tit-for-tat strategy: let's say I know that this is the last interaction I will choose to defect. Let's say I know this is the penultimate interaction I will choose to betray and so on. Repeated interaction only leads to efficiëncy with an unknown end of the game.

Nash equilibrium

The definition of the **Nash** equilibrium is a combination of strategies, where no player has incentive to deviate from, given the strategy of the other player. It is possible to have multiple Nash equilibriums:

Game		Player 2	Player 2
		Option A	Option B
Player 1	Option A	**10, 10**	0, 0
Player 1	Option B	0,0	**8, 8**

Keep in mind that the Nash-equilibrium is a combination of strategies and not payoffs. Therefore in this example you should write (Option A, Option A) and (Option B, Option B).
If there is 1 Nash-equilibrium it's called a unique equilibrium.

Commitment problem: How do you make a threat credible? Let's say you are gonna choose A either way. The problem then is how do you make it believable to your opponent that you are gonna choose A either way. The commitment is the way you make it credible.
You deal with commitment problems via commitment devices, government coordination or alternative equilibriums.
An alternative equilibrium is the Maximin equilibrium: Maximise the minimal outcome. For example choosing A will give you 10000 or -1 and choosing B will give you 1 or 2. You will have to choose B.

Sequential games

In the game below A will choose first and then B will choose. You notate the strategy for B as: (Choice if T, Choice if O). B chooses from four strategies: (L, L), (L, R), (R, L), (R, R).

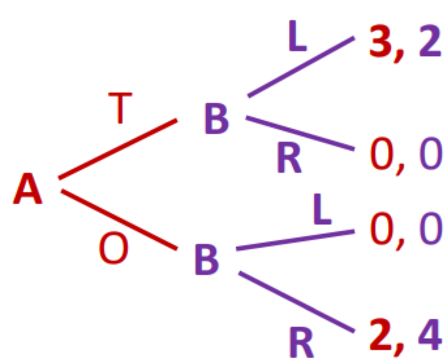

We can also write this as a game matrix

	(L, L)	(L, R)	(R, L)	(R, R)
T	3,2	3,2	0,0	0,0
O	0,0	2,4	0,0	2,4

The nash equilibriums here are (T, (L, L)), (T, (L, R)) and (O, (R, R)). Although (O, (R, R)) is not believable. When B chooses the choice of A is already known so why would A still choose for T.

To bypass this problem you can use backwards induction. Start at the last decision point and work towards the start. If we do this we will end up with the **'Subgame perfect'** equilibrium: (T, (L, R))

Microeconomics - IBEB

Lecture 17, week 7

Nash Equilibrium vs Subgame-Perfect Equilibrium

- Sometimes a Nash equilibrium is not credible because it depends on behavior that is not optimal in future subgames.
- To solve dynamic games, use backward induction to find the Subgame-Perfect Equilibrium (SPE).
- We learn SPE because Stackelberg is a dynamic game and requires backward induction.

Oligopolies

A **Oligopoly** is a market form where there are a few small producers.
A **Duopoly** is a market form where there are two producers.

Firms have direct influence on the market equilibrium (unlike perfect competition) and therefore each others decisions. This leads to strategic interaction (game theory).

Cournot Model: Duopoly

There are two firms in the cournot model. $Firm\ i\ with\ i \in (1, 2)$.
These firms sell perfect substitutes in the same market.
Firms **choose simultaneous** the **produced quantity**.

Each firm will look for the best quantity (Qi) to produce given the quantity produced by the opponent. This is most relevant when production decisions are made in advance.

The market demand is $P(Q) = a - bQ\ with\ Q = Q1 + Q2$
The market demand per firm is $P_i(Q_i)$=a-bQ_1-bQ_2
Therefore the TR curve is $TR_i(Q_i)$=Q_i*P(Q_1+Q_2)

And this shows the marginal revenue curve can be written as MR_i=a-bQ_j-2bQ_i
We want to maximize profit. We can do this by setting $MR_i(Q_i)$=$MC_i(Q_i)$ or maximizing π_i by taking the derivative and setting it to zero.

If we do this for both the quantity of firm 1 and firm 2. We get a scheme with 2 functions. These functions are the reaction functions: How much does firm i produce for every Q firm j produces.

$Q_1=(a-bQ_2-c_1)/2b$; $Q_2=(a-bQ_1-c_2)/2b$.

Solving this will give the best outcomes this will also lead to the Nash-equilibrium. No firm has incentive to change the quantity they produce. The invisible hand to the equilibrium is illustrated below.

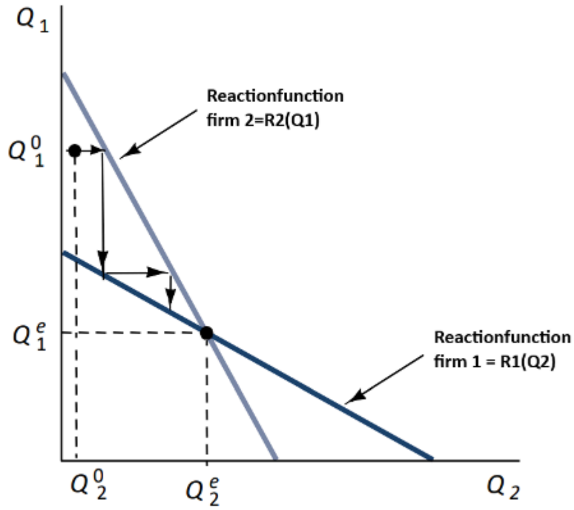

In the cournot model firms produce individually less than in a monopoly. Although together they produce more than a monopoly. Therefore the price in a cournot duopoly is lower than in a monopoly.

$Q_1=Q_2<Q^{mon}$ and $Q_1+Q_2>Q^{mon}$

Cournot firms produce more than a monopolist but less than in perfect competition..

$$Q^{Cournot} > Q^{Monopoly}, \quad Q^{Cournot} < Q^{Perfect\ Competition}$$

Bertrand Model: Duopoly

Bertrand criticized the assumption that every player the quantity produced of the opponent as a given takes in the Cournot model. In the bertrand model the firms choose the prizes. Every firm takes the price of the opponent as a given.

In the base model there are 2 firms i=1, 2. The marginal costs are equal $MC_1=MC_2=c$. The goods sold are perfect substitutes. The firms choose the prices. THe consumers only buy the lowest prices. The market demand is $Q(\min\{P_1,P_2\})$. The firm with the highest price doesn't sell anything, the firm with the lowest price sells the whole market demand and for equal prices both firms sell 50% of the market demand.

Let's say $P_1>C$(Marginal cost). The best response for firm 2 is to set it's price a tiny bit lower than firm 1 -> $c<P_2<P_1$. This situation is unstable and firm 1 will lower it's prices again. This will happen until P = MC (just like perfect competition). The nash-equilibrium is $P_1=P_2=c=MC$

If $P_1 > c$, firm 2 can profitably set $P_2=P_1-\varepsilon$ and capture the entire market.

This undercutting continues until **P = MC**, which is the **unique Nash equilibrium**.

Even with only 2 firms, Bertrand competition yields the **competitive price**.

Collusion in economics is the collaboration between companies that seek to gain an extensive competitive advantage in the marketplace. Collusion is illegal. This means that if firms conspire together, they will reach a higher profit. They could for example set the market price to the monopoly price(/Quantity). This reaches the highest profit for both of them.

Stackelberg Model: Duopoly

In the stackelberg model a firm is the market leader and the other firms identify as followers. An example of this is Albert Heijn for Dutch supermarkets. Firms choose their quantities although and both firms try to maximize their profit. Although in this case it is a sequential game:

1. First the leader chooses his quantity
2. The follower then sees what the leader does, and reacts by choosing its own quantity.
3. The total quantity is sold and the market decides the price.
4. Firms receive their profit and the game ends.

The game has the same nature as the Cournot model although in this case it's a sequential game. Remember for sequential games we can find the subgame equilibrium by using backward induction. Therefore we start with the decisions for the follower who sees Q_L (Quantity of leader) as a given. The follower wil try to maximize it's profit: $\Pi_v=P(Q_L+Q_v)Q_v-c_vQ_v=(a-b(Q_L+Q_v))Q_v-c_vQ_v$
The optimum is $Q_v=(a-bQ_L-c_v)/2b$. The leader knows the reaction function of the follower and will therefore try to maximize it's own profit. $\Pi_L=P(Q_L+R_v(Q_L))Q_L-c_LQ_L$.
Maximizing this will give the quantity the leader is gonna produce and we can substitute this in the quantity of the follower and so find the solution to this problem.

Stackelberg leadership:

- Every player wants to be the "Stackelberg Leader"
- Being the leader requires commitment: The leader announces the quantity it's gonna produce and sticks to it.
- The follower believes the commitment and reacts.

- Therefore the commitment should be believable.
- An example of a commitment is a sunk investment in extra production capacity.

Below is a scheme with all the outcomes for the different models with P(Q)=a-bQ and $TC_i=cQ_i$.

Model	Industry output Q	Price P	Industryprofit $\boldsymbol{\Pi}$
Shared monopoly (cartel)	$Q^{mon}=\frac{a-c}{2b}$	$P = \frac{a+c}{2}$	$\boldsymbol{\Pi}^{mon}$
Cournot	$4/3*Q^{mon}$	$P = \frac{a+2c}{3}$	$8/9*\boldsymbol{\Pi}^{mon}$
Stackelberg	$3/2*Q^{mon}$	$P = \frac{a+3c}{4}$	$3/4*\boldsymbol{\Pi}^{mon}$
Bertrand	$2Q^{mon}$	c	0
Perfect competition	$2Q^{mon}$	c	0

Intuition:

Monopoly: lowest output, highest price, highest profits
Cournot: higher output than monopoly, lower price
Stackelberg: leader produces more, industry output higher than Cournot
Bertrand: P = MC, zero profits
Firms would prefer to collude (cartel) to replicate monopoly profits if it were legal.

Microeconomics - IBEB
Lecture 18, week 7

Motivation: Spatial Competition

Traditional oligopoly models (Cournot, Bertrand) assume:

- Firms are symmetric
- Consumers treat firms symmetrically

But in many real markets, consumers choose the nearest firm → firms compete in location, not only price or quantity.

Examples: supermarkets, petrol stations, pharmacies.

This leads to spatial competition models, starting with Hotelling (1929).

Hotelling’s Line Model

Assumptions:
Firms sell an identical homogeneous product
Price is the same for all firms (e.g., Bertrand P regulated or identical)
Consumers are uniformly distributed on a line of length 1
Each consumer buys 1 unit
They buy from the nearest firm (transportation cost determines choice)

If only 1 firm exists:
It can locate anywhere on the line — it is a monopolist.

Hotelling: 2 Firms
Big result:
Both firms locate at the center (½) → “Principle of Minimum Differentiation”.

Explanation:

If A moves slightly left of ½, B captures the entire right half.
If B moves slightly right of ½, A captures the entire left half.
At ½, each has 50% market share and no profitable deviation exists
→ Nash equilibrium.

Consumer Welfare Problem

Is ½–½ best for consumers?

No.
Consumers prefer:
Firms located at ¼ and ¾

Why?

- Average travel distance is smaller
- Total welfare improves without harming firms (they still get 50% market share each)

But:

(¼, ¾) is not a Nash equilibrium.
Each firm has incentive to move toward the center.

→ Hotelling NE is not socially optimal.

This is a classic example of market failure from strategic behavior.

Two Problems Identified

Hotelling reveals:

1. **NE is not welfare-maximising**
   → Potential role for government intervention in location regulation.
2. **The line model is flawed**
   → The "race to the center" is partly due to endpoints.

Solution: switch to a **circle** (Salop).

Salop's Circular City Model

Interpretation:

- Location now represents variety or type of the product
- Consumers are uniformly distributed on a circle
- Transportation cost = taste distance (differentiation)
- Firms choose locations (varieties)
- With N firms, the equilibrium outcome is
- All firms spaced equidistantly on the circle

Key contrast with Hotelling:

- Hotelling → no differentiation (firms cluster at the center)
- Salop → maximum possible differentiation (even spacing)

Policy Application: Optimal Number of Pharmacies (Salop Model)

• City is a circle of circumference 1.
• Consumers uniformly distributed (L people).
• Travel cost t per km. Each consumer buys 1 unit.
• Shops sell identical product at identical price.
• Shop cost: TC(B) = F + B → ATC decreases as B increases.
• Government chooses N (number of shops) to minimize total cost.

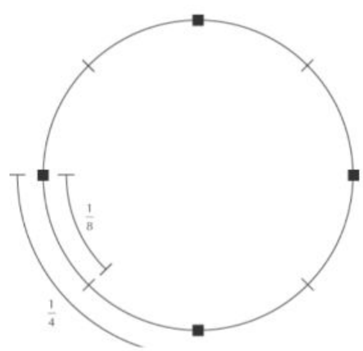

(eg. for n = 4)

Transportation Costs with N Shops
• Shops located evenly around the circle.
• Maximum distance to nearest shop = 1/(2N).
• Average distance = 1/(4N).
• Total transport cost: $TC_{transport} = tL / (2N)$.

Production Costs

- Each shop serves L/N customers.
- Total production cost: $TC_{production} = NF + L$.

Total Cost Minimization

Total cost: $TC(N) = tL/(2N) + NF + L$

FOC:

$-tL/(2N^2) + F = 0$

So:

$N^2 = tL/(2F)$

Optimal number of shops:

$N^* = \sqrt{(tL\ /\ (2F))}$

Interpretation

- If travel cost t increases → more shops optimal.
- If population L increases → more shops optimal.
- If fixed cost F increases → fewer shops optimal.
- Shows how economics guides policy design.