



UNIVERSITÄT
BAYREUTH

Advanced Microeconomics I

Summer Term 2025

Organizational preliminaries

- Prof. Dr. Stefan Napel
 - ▶ Office hours: Monday, 2-4 pm;
please contact vw14@uni-bayreuth.de (Heidi Rossner-Schöpf)
- Downloads and information:
elearning.uni-bayreuth.de/course/view.php?id=43733
- Classes by Dr. Alex Mayer & Sebastian Schröter w/ one week delay to lectures
(start: May 12/13, 2025; use week before to work on *Mathematical Preliminaries* PDF by yourself)
- Optional Q&A sessions with student tutor Felix Brinkmann (start: May 16, 2025)
- One-open-book exam will be posed in English;
can be answered in English or German
(same for optional midterm exam on June 11, 2025)

Textbooks

- The reference (consider buying a used copy):
 - ▶ Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green (1995). *Microeconomic Theory*. New York, NY: Oxford University Press.
(≡ **MWG**)
- Other recommended textbooks:
 - ▶ Jehle, Geoffrey A., and Philip J. Reny (2011). *Advanced Microeconomic Theory*, 3rd ed. Amsterdam: Addison-Wesley.
 - ▶ Rubinstein, Ariel (2012). *Lecture Notes in Microeconomic Theory: The Economic Agent*, 2nd ed. Princeton, NJ: Princeton University Press.
[updated in 2023 and free: <http://arielrubinstein.tau.ac.il/>]
 - ▶ Varian, Hal R. (1992). *Microeconomic Analysis*, 3rd ed. New York, NY: W. W. Norton & Company.

Goals and structure

- Goals of this course:
 - ▶ Introduce key concepts of advanced microeconomic analysis
 - ▶ Aid the self study of MWG
 - ▶ Prepare for possible PhD studies:
we pick a level below a PhD program, but familiarize ourselves with the standard textbook
→ you may skip the small print and most proofs for now
 - ▶ Structure follows MWG

Prospective schedule for lectures

#	Date	Topic	Chs. in MWG
1	28.04.	Introduction	
2	05.05.	Preference and choice	1.A–D
3	12.05.	Consumer choice	2.A–F
4	19.05.	Classical demand theory	3.A–E, G
5	26.05.	Aggregate demand	3.I; 4.A–D
6	02.06.	Choice under risk	6.A–D, F
7	16.06.	Static games of complete information	7.A–E; 8.A–D, F
8	23.06.	Dynamic games of complete information	9.A–B; 12. App. A
9	30.06.	Games of incomplete information	8.E; 9.C
10	07.07.	Competitive markets	10.A–G
11	14.07.	Market power	12.A–F
12	21.07.	Question session for exam (\rightarrow t.b.a.)	

... blood, toil, tears, and sweat

- This course is different ...
 - ▶ Lectures will *not* provide a self-contained treatment of *all* material
 - ▶ Strenuous self-study cannot be avoided
(workload still much lower than in a UK/US research MPhil/MSc program;
NB: 8 ECTS points associated with 8 h of homework / week, plus 4 h attendance!)
- Mixture of slides and “chalk & talk”
- Optional midterm exam on June 11, 2025:
 - ▶ Two problems on topics of sessions #1 – #6,
each graded in a binary fashion (“+” or “○”)
 - ▶ Each “+” earns 5 bonus points for this term’s 60-point final and re-sit exams
(but not: next year’s exams)
- Most of what you learn in this course will be learned by *doing problems*,
i.e., preparing for weekly classes and exams

1. Introduction

- Microeconomics studies the behavior of (groups of) individuals or firms, how they interact and bring about collective outcomes
- We will look at models of
 - ▶ preferences, consumer choice, demand, choice under risk
 - ▶ strategic decision making (= game theory)
 - ▶ perfectly and imperfectly competitive markets
- Market failure, asymmetric information, and mechanism design
- General equilibrium
- Social choice and welfare

Models

- “Models” are simplified descriptions of a part of reality
- Their purposes in economics include
 - ▶ description per se
 - ▶ explanation and prediction
 - ▶ justification
 - ▶ decision support
- They can be represented in different ways, e.g.,
 - ▶ verbally
 - ▶ graphically or mechanically
 - ▶ mathematically
 - ▶ as software
- All representations boil down to a system of assumptions, axioms, premises, or initial conditions $\{A_1, \dots, A_n\}$


Models

- The system of assumptions, axioms, etc. should
 - ▶ be logically consistent, irreducible, and comprehensible
(A. Einstein: “... as simple as possible, but not simpler!”)
 - ▶ relevant for the model’s purpose, relate to reality, and have at least some empirical support
- The advantage of stating $\{A_1, \dots, A_n\}$ mathematically instead of in everyday language or software is that the model is particularly
 - ▶ concise and transparent
 - ▶ easy to check for consistency
 - ▶ amenable to formal manipulations and logical deduction
- Mathematical models are constructed with manipulability in mind;
this implies a delicate trade-off with realism
(Danger: “Searching where the light is ...”)

“Searching where the light is ...”

THE STREETLIGHT EFFECT

BY PEDRO ALBERTO ARROYO



WWW.BITSTRIPS.COM

Models and Economic Theory

- Early philosophers of science (Hempel, Oppenheim) argued that the distinctive feature of a *theory* (vs. a model) is:
at least some A_i is a universal law, i.e., a time and space-invariant, necessary connection between certain phenomena
- Such requirements would preclude any economic theory . . .
- Social scientists have to contend themselves with restricted regularities or mere tendencies (vs. laws of mechanics)
 - ▶ e.g., that individuals can usually decide between two available options and mostly do so in a consistent fashion
- Economics is harder than physics because it involves interpretation of a reality created by objects of study (individuals, firms, . . .) who themselves base their actions on individual interpretations of reality, possibly influenced by economic theory

Do Economics, not Mathematics

- Most microeconomic analysis uses mathematical language and techniques
- We need to do the maths (and data analysis) because even trained economic intuition is sometimes wrong:
 - ▶ One obtains a 'counter-intuitive' result doing the maths/econometrics, and only facing it realizes that some (ex post: intuitive) causal effects were overlooked
- Strive to focus on the economics in what you read and do, even though the maths may be more time-consuming
- A good intuition about agents' economic incentives is more useful than superb knowledge of Kuhn-Tucker conditions or semidefiniteness of matrices, even in optimization problems

1.1 Example

- Consider the following simple microeconomic problem:
 - ▶ Julian wants to buy spoons and forks
 - ▶ Each pair of one spoon and one fork gives Julian 1 unit of utility
 - ▶ A spoon not matched with a fork gives him only a units of utility, where $0 \leq a < 1/2$; a fork not matched with a spoon also gives a
 - ▶ Let p_1 be the price of spoons, p_2 the price of forks, and w the wealth that Julian plans to spend on spoons and forks
 - ▶ Assume he wants to get the highest possible utility for his money
- Find Julian's demand functions for spoons and forks!
Assume $0 < p_1 \leq p_2$ and perfect divisibility



UNIVERSITÄT
BAYREUTH

Advanced Microeconomics I

Summer Term 2025

Schedule for lectures

#	Date	Topic	Chs. in MWG
1	28.04.	Introduction	
2	05.05.	Preference and choice	1.A–D
3	12.05.	Consumer choice	2.A–F
4	19.05.	Classical demand theory	3.A–E, G
5	26.05.	Aggregate demand	3.I; 4.A–D
6	02.06.	Choice under risk	6.A–D, F
7	16.06.	Static games of complete information	7.A–E; 8.A–D, F
8	23.06.	Dynamic games of complete information	9.A–B; 12. App. A
9	30.06.	Games of incomplete information	8.E; 9.C
10	07.07.	Competitive markets	10.A–G
11	14.07.	Market power	12.A–F
12	21.07.	Question session for exam (→ t.b.a.)	

2. Preference and choice

- The basic constituent of most economic models is the neoclassical “economic man” or *homo economicus*
- He or she is a highly stylized model of real decision makers
- “*economicus*” refers to “the economic way” of decision making, not to the context of decisions
- Broadly speaking, *homo economicus* is assumed to
 - ▶ deliberately choose the most suitable means to his or her ends
 - ▶ evaluate options according to their anticipated consequences (decisions are made in the “shadow of the future”)
 - ▶ weigh the costs and the benefits of a particular choice...or rather behave “as if” he or she would be doing so

Preliminary remarks

- While the rationality embodied by *homo economicus* is the key assumption of most of modern economics, it should not be taken too literally
- Hardly any economist thinks that real people *are* as deliberate, future-oriented, and clever as is conventionally assumed
- Most would hold that people are behaving *as if* they were “economically rational” sufficiently often to derive useful conclusions from correspondingly pragmatic models
- See
 - ▶ Ariely, Dan (2008). *Predictably Irrational*. London: Harper Collins.
 - ▶ Kahneman, Daniel (2011). *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.

for illuminating accounts of the “biases” of real decision makers

Choosing between several alternatives

- Consider an agent who needs to choose between several actions and suppose
 - ① each action is associated with a particular outcome, and
 - ② these outcomes are all that the agent cares about
 - Denote the set of all possible, mutually exclusive outcomes / options / alternatives by X
 - ▶ Options can be very concrete, like
$$X = \{\text{go to law school in Berlin, study economics in Bayreuth, ...}\},$$
or, for us, abstract placeholders like $X = \{x, y, z\}$
 - Economics presumes that whenever choosing from a subset $X' \subseteq X$, the agent picks an option $x \in X'$ which serves his or her goals best (whatever they may be ...)
- ⇒ If we observe that the agent chooses x from X' , we *conclude* that x was among the best options in X' for this agent (at least at the time of choosing)

2.1 Preferences vs. choice rules

- There are two main approaches to modeling choice behavior:
 - ▶ Binary preference relations
 - ▶ Choice rules
- Preference relations are less general, but more handy
(additional restrictions make them even more handy, e.g., allow representation by a utility function)
- Observing the choice of x when X' was available reveals that x is *weakly preferred* to any other element $y \in X'$ *when a choice must be made from X'*
- The preference approach entails the simplifying assumption:
 x is weakly preferred to y independently of the presence or absence of any other alternatives $z \in X$, i.e., also when a choice must be made from $X'' \neq X'$

Preference relations

- Given such context-independence, an agent's full choice behavior is well-defined by her choices from *binary* subsets $X' = \{x, y\}$
- When x is weakly preferred to y , we write: $x \succsim y$
- \succsim gives (some) *pairs* of elements $x, y \in X$ a specific connection; it is known mathematically as a *binary relation*
- A binary relation is formally just a subset of $X \times X$; some authors write $(x, y) \in \succsim$ instead of $x \succsim y$
(BTW: a function $f: X \rightarrow Y$ can similarly be viewed as a subset of $X \times Y$)

Other relations derived from \succsim

- If sometimes x and sometimes y is chosen out of $X' = \{x, y\}$, then the agent is said to be *indifferent* between x and y , i.e.,

$$x \succsim y \wedge y \succsim x \Leftrightarrow: x \sim y$$

- If the agent (weakly) prefers x over y and is not indifferent, he is said to *strictly prefer* x over y , i.e.,

$$x \succsim y \wedge \neg(y \succsim x) \Leftrightarrow: x \succ y$$

- $x \succ y$ is equivalent to saying:
“The agent never chooses y when x is available”

Rational preferences

- Economics does not care about *why* somebody prefers x to y ; neither does it proclaim which option the agent *should* prefer
- The common requirement for calling an individual *rational* is that her choices reflect preferences that are “complete” and “transitive”
- *Complete* means that for any two options $x, y \in X$, the agent either weakly prefers x or weakly prefers y or both, i.e.,

$$\forall x, y \in X : (x \succsim y) \vee (y \succsim x)$$

- Completeness reflects that the agent can reach a decision in any binary choice problem
- *Transitive* means that a preference for x over y together with a preference for y over z also entails a preference for x over z , i.e.,

$$(x \succsim y) \wedge (y \succsim z) \Rightarrow x \succsim z$$

- Transitivity rules out cycles that would, e.g., preclude a decision from $X' = \{x, y, z\}$

Violations of transitivity

- An argument against persistent intransitivity of real people is that one might (or the market would) ruin them with a *money pump*:
 - ▶ Suppose your colleague has intransitive preferences:
$$\text{apple} \succ \text{banana} \succ \text{citrus fruit} \succ \text{apple}$$
 - ▶ Give him an apple for free
 - ▶ Then offer to sell him a citrus fruit for the apple and, e.g., 1 cent; he will accept because he strictly prefers the citrus fruit
 - ▶ Next sell him a banana for the citrus fruit and 1 cent
 - ▶ Now sell him an apple for the banana and 1 cent, and repeat the cycle ...
- However, this ignores transaction costs, and the possibility that an intransitivity may be corrected (only) if someone exploits it
- Intransitivity is normal when alternatives are very finely graded:
$$\forall k \in \mathbf{N}_0: \text{coffee with } k \text{ grains of sugar} \sim \text{coffee with } k + 1 \text{ grains of sugar}$$
$$\Rightarrow \text{coffee without sugar} \sim \text{coffee with } 100\text{g of sugar?}$$

2.2 Utility representation

- If the set of alternatives X is *finite* (or *countably infinite*) and the agent has a complete and transitive preference relation \succsim over it, then the agent's preferences over X can be *represented* by a *utility function* $u: X \rightarrow \mathbf{R}$, i.e., we can find real numbers $u(x)$ such that

$$x \succsim y \Leftrightarrow u(x) \geq u(y)$$

- Note that if $u(\cdot)$ represents the agent's preferences, then so does any $v(\cdot)$ which is a strictly increasing transformation of $u(\cdot)$
- The latter implies that differences or ratios between utility levels for x and y do not mean anything:
 $u(\cdot)$ only allows conclusions about the order of x and y , and is therefore called an *ordinal* utility function

Utility representation

- If the set of alternatives X is *uncountably infinite*, then completeness and transitivity of a preference relation are not sufficient to guarantee existence of a utility representation
- In particular, *lexicographic preferences* \succsim_L over bundles $(x_1, x_2) \in \mathbf{R}_+^2$ of two goods defined by

$$(x_1, x_2) \succsim_L (y_1, y_2) :\Leftrightarrow x_1 > y_1 \vee \{x_1 = y_1 \wedge x_2 > y_2\},$$

and

$$(x_1, x_2) \sim_L (y_1, y_2) :\Leftrightarrow x_1 = y_1 \wedge x_2 = y_2$$

do not possess a utility representation

Utility representation

- A preference relation \succsim is called *continuous* : \Leftrightarrow
whenever $x^k \succsim y$ (resp. $y \succsim x^k$) holds for all elements x^k of a sequence $\{x^k\}_{k=1,2,\dots}$
with limit point x^* then $x^* \succsim y$ (resp. $y \succsim x^*$) holds too
- Continuity rules out that minimal changes flip the ordering of two options:
 - ▶ Lexicographic preferences rank $x^k = (2 + 1/k, 1)$ strictly higher than $y = (2, 2)$ for every $k = 1, 2, \dots$
 - ▶ The limit point $x^* = (2, 1)$, however, is ranked strictly lower than $(2, 2)$
- A key result in decision theory:
If \succsim is a complete, transitive and continuous preference relation on an arbitrary set of outcomes X , then
 - ▶ \succsim can be represented by an ordinal utility function $u: X \rightarrow \mathbf{R}$
 - ▶ $u(\cdot)$ can be chosen to be continuous
(but not necessarily also differentiable, or even C^1 , C^2 , etc.)

Remarks

- Economic rationality itself does not require existence of a utility representation of an agent's preferences
- Only for convenience is economic rationality sometimes equated with utility-maximizing behavior, but inaccurately so
- In any case, assuming utility maximization does *not* require agents to “know their utility function” and “try to maximize”;
as it happens, if their preferences satisfy completeness and transitivity (+ continuity), they act exactly *as if* they did ...
- Use of a particular utility function (e.g., $u(x_1, x_2) = x_1 + x_2$) amounts to an additional assumption *on top* of that of a *homo economicus*

Further remarks

- Preferences are individual characteristics that economists take as given and fixed
- We tend to ignore preferences'
 - ▶ origin or causes
 - ▶ intensity
 - ▶ possible dynamics
- There is, however, also economic research that investigates preference saliences or patterns and drivers of preference change
- The key challenge in the context of changing / reference-dependent preferences is the welfare interpretation of outcomes

2.3 Choice structures and choice rules

- Recall that the move from observing choice x from X' towards a binary preference relation entailed a presumption of context-independence regarding greater desirability of x than $y \in X'$
- If one does not want to impose this restriction, one can work with so-called *choice structures*
- A choice structure $(\mathcal{B}, C(\cdot))$ has two ingredients:
 - ▶ $\mathcal{B} \subseteq 2^X$ is a family of nonempty subsets of X ;
elements $B \in \mathcal{B}$ are called *budget sets*,
 \mathcal{B} is meant to describe all choice experiments that could be posed to the decision maker, or on which we have data
 - ▶ The so-called *choice rule* or *choice correspondence* $C(\cdot)$ maps each budget set $B \in \mathcal{B}$ to a nonempty subset $C(B) \subseteq B$;
it lists all alternatives that the decision maker might choose from B
(i.e., finds equally acceptable from B)

Example

- Suppose that $X = \{BT, KU, N\}$ and $\mathcal{B} = \{\{KU, N\}, \{BT, KU, N\}\}$
- A possible choice structure is $(\mathcal{B}, C_1(\cdot))$, where
 - ▶ $C_1(\{KU, N\}) = \{KU\}$
 - ▶ $C_1(\{BT, KU, N\}) = \{KU\}$
- Kulmbach is preferred location no matter what other alternatives are in the budget set
- A possible choice structure is $(\mathcal{B}, C_2(\cdot))$, where
 - ▶ $C_2(\{KU, N\}) = \{N\}$
 - ▶ $C_2(\{BT, KU, N\}) = \{KU\}$
- She prefers the location in the budget set which is second-closest to Bayreuth

Weak Axiom

- A common restriction on choice structures $(\mathcal{B}, C(\cdot))$, which rules out behavior of the latter kind, is the *weak axiom of revealed preference (WARP or WA)*:
 - ▶ If x is chosen for a $B \in \mathcal{B}$ that also contains y ,
and y is chosen for another $B' \in \mathcal{B}$ that also contains both,
then x must be equally acceptable in B' , i.e.,

$$x, y \in B, x \in C(B) \text{ and } x, y \in B', y \in C(B') \Rightarrow x \in C(B')$$

- We interpret the existence of a budget set $B \ni x, y$ with $x \in C(B)$ as:
“ x is revealed weakly preferred to y (for some budget set)”
- So WARP more simply says:
If x is revealed weakly preferred to y , then y cannot be revealed strictly preferred to x

Relation between preferences and choice rules (1)

- Two natural questions arise about WARP:
 - ① If a decision maker has a rational preference ordering \succsim , do her choices – when facing budget sets in \mathcal{B} – necessarily satisfy WARP?
 - ② If an individual's choice behavior for budget sets \mathcal{B} is captured by a structure $(\mathcal{B}, C(\cdot))$ that satisfies WARP, does a rational preference relation \succsim exist which is consistent with these choices (i.e., which 'rationalizes' $C(\cdot)$ relative to \mathcal{B})?

Relation between preferences and choice rules (2)

- Both questions can basically be answered affirmatively:
 - ① A choice structure which is generated by a rational preference ordering \succsim automatically satisfies WARP
 - ② That a choice structure $(\mathcal{B}, C(\cdot))$ satisfies WARP is sufficient for existence of a (unique) preference relation \succsim that rationalizes it *if* \mathcal{B} includes all subsets $X' \subseteq X$ with $|X'| \leq 3$ (only then does WARP guarantee transitivity)
- So, if choices are defined on all subsets of X and satisfy WARP, then the preference and choice rule-based approaches to modeling behavior are equivalent
- Preview: consumer decisions described by a demand function $x(\mathbf{p}, w)$ are defined only for special subsets of X ; then stronger properties than WARP are needed to guarantee that choices are rationalizable (in the microeconomic sense)



UNIVERSITÄT
BAYREUTH

Advanced Microeconomics I

Summer Term 2025

Schedule for lectures

#	Date	Topic	Chs. in MWG
1	28.04.	Introduction	
2	05.05.	Preference and choice	1.A–D
3	12.05.	Consumer choice	2.A–F
4	19.05.	Classical demand theory	3.A–E, G
5	26.05.	Aggregate demand	3.I; 4.A–D
6	02.06.	Choice under risk	6.A–D, F
7	16.06.	Static games of complete information	7.A–E; 8.A–D, F
8	23.06.	Dynamic games of complete information	9.A–B; 12. App. A
9	30.06.	Games of incomplete information	8.E; 9.C
10	07.07.	Competitive markets	10.A–G
11	14.07.	Market power	12.A–F
12	21.07.	Question session for exam (→ t.b.a.)	

3. Choice-based demand theory

- Now study *homo economicus* as a consumer in a competitive market economy; adopt a choice-based perspective first (\leftrightarrow preference-based in 4.)
- Choice of quantities of goods or services provided by the market, called *commodities*, is subject to physical and economic constraints
- Any particular quantity combination (x_1, x_2, \dots, x_L) of L different commodities corresponds to a point \mathbf{x} in *commodity space* \mathbf{R}^L
- Definition of the relevant commodities comes with great flexibility: same good delivered at different points in time, different locations, or distinct ‘states of the world’ are just different commodities
- Physical restrictions on bundles that the individual can consume are reflected by restricting \mathbf{R}^L to a *consumption set* $X \subseteq \mathbf{R}^L$

Divisibility and price taking

- For simplicity, we consider \mathbf{R}_+^L as agents' consumption set; this is a *convex* set, i.e., we assume *perfect divisibility*
- We also assume a *complete market*, i.e., every commodity $i = 1, \dots, L$ is traded (i.e., property rights are well-defined for every relevant good)
- The considered consumer is presumed to be a *price taker*, i.e., his or her individual decisions do not affect prices
- Suppliers use linear pricing schemes, i.e., sell at a constant unit price, e.g., because there is perfect competition (vs. *non-linear pricing*: quantity discounts, two-part tariffs, ...)
- For convenience, let the price of any good i be positive, i.e., $p_i > 0$ for $i = 1, \dots, L$

3.1. Walrasian budget sets

- The economic constraint faced by the agent is that he or she must afford the selected commodity bundle $\mathbf{x} \in \mathbf{R}_+^L$, i.e., for a given price vector $\mathbf{p} \in \mathbf{R}_+^L$ total expenditure

$$\mathbf{p} \cdot \mathbf{x} := p_1 x_1 + \cdots + p_L x_L$$

cannot exceed wealth $w > 0$

- The set of affordable, physically feasible bundles for given \mathbf{p} and w is the consumer's *Walrasian* or *competitive budget set*

$$B_{\mathbf{p},w} := \{\mathbf{x} \in \mathbf{R}_+^L : \mathbf{p} \cdot \mathbf{x} \leq w\}$$

- The consumer's choice problem is thus:
"Choose a consumption bundle \mathbf{x} from $B_{\mathbf{p},w}$ "

Budget hyperplane

- The set $\{\mathbf{x} \in \mathbf{R}^L : \mathbf{p} \cdot \mathbf{x} = w\}$ is known as the *budget line*; or for $L > 2$ as the *budget hyperplane*; it is the upper boundary of $B_{\mathbf{p},w}$
- Its respective intercepts are w/p_i , i.e., the maximal affordable quantity if only good i is purchased
- The fact that $\mathbf{p} \cdot \mathbf{x} = w$ and $\mathbf{p} \cdot \mathbf{x}' = w$ for any two points \mathbf{x} and \mathbf{x}' on the budget hyperplane implies that \mathbf{p} is orthogonal to it
[recall that the *dot product* of any vectors $\mathbf{x}, \mathbf{y} \in \mathbf{R}^L$ satisfies $\mathbf{x} \cdot \mathbf{y} = |\mathbf{x}| \cdot |\mathbf{y}| \cdot \cos(\theta)$ where θ is the angle between \mathbf{x} and \mathbf{y} ; in particular, $\mathbf{x} \cdot \mathbf{y} = 0$ iff \mathbf{x} and \mathbf{y} are orthogonal]

3.2 Walrasian demand

- Set $\mathcal{B}^{\mathcal{W}} = \{B_{\mathbf{p},w} : \mathbf{p} \in \mathbf{R}_{++}^L \wedge w > 0\}$ is just a particular family of budget sets
- At least in principle, we can observe a consumer's choices $C(B) \subseteq B$ for any budget set $B = B_{\mathbf{p},w} \in \mathcal{B}^{\mathcal{W}}$
- These choices are called the (*Walrasian*) *demand* of the consumer and we refer to

$$x(\mathbf{p}, w) := C(B_{\mathbf{p},w})$$

as the consumer's *Walrasian demand correspondence*

- We often focus on cases in which $C(B_{\mathbf{p},w})$ is singleton-valued, i.e., the consumer picks a unique element in any Walrasian budget set
- $x(\mathbf{p}, w)$ is then called the Walrasian demand *function* (w/o brackets around $\{\mathbf{x}^*\}$)

Homogeneity of Walrasian demand

- A function $f: X \rightarrow Y$ (or a correspondence $f: X \rightrightarrows Y$) between vector spaces X and Y is called *homogeneous of degree r* $:\Leftrightarrow \forall \lambda > 0: \forall \mathbf{x} \in X: f(\lambda \mathbf{x}) = \lambda^r \cdot f(\mathbf{x})$
 - Demand is *homogeneous of degree zero* iff $x(\lambda \mathbf{p}, \lambda w) \equiv x(\mathbf{p}, w)$,
i.e., when prices and wealth all change by the same factor then demand does not change (\rightarrow only relative prices matter)
 - We will assume that the individual cares only about the commodities, and doesn't suffer any "money illusion"
- \Rightarrow Choice depends only on which bundles are affordable;
so the fact that $B_{\mathbf{p}, w} \equiv B_{\lambda \mathbf{p}, \lambda w}$ implies $x(\lambda \mathbf{p}, \lambda w) \equiv x(\mathbf{p}, w)$

Homogeneity of demand and numeraire good

- Given that we can scale prices and wealth up or down by $\lambda > 0$ without affecting demand, it is often convenient to normalize such that $w = 1$ or such that $p_i = 1$ for some good i
- In the latter case, all prices and wealth are expressed in units of good i , which is then called the *numeraire good*

Walras' Law

- We say that a Walrasian demand function (or correspondence) $x(\mathbf{p}, w)$ satisfies *Walras' law* or is *budget balancing* iff it is an element of the budget hyperplane for all \mathbf{p} and w , i.e.,

$$\mathbf{x} = x(\mathbf{p}, w) \Rightarrow \mathbf{p} \cdot \mathbf{x} = w$$

(or $\mathbf{x} \in x(\mathbf{p}, w)$)

- Walras' law says that the consumer fully expends his or her wealth
- When understood in a broad way (e.g., as applying to the entire lifetime of an agent, with bequests viewed as commodities, too), this is not very restrictive

3.3 Comparative statics w.r.t. wealth

- How do observed choices vary with changes in wealth and prices?
- Examination of outcome changes caused by a change in underlying economic parameters is known as *comparative statics* analysis
- The *wealth effect* for good i at (\mathbf{p}, w) is simply $\partial x_i(\mathbf{p}, w) / \partial w$
- Commodity i is *normal* at (\mathbf{p}, w) if the wealth effect for it is positive, i.e., demand increases in wealth;
 i is *inferior* at (\mathbf{p}, w) if the wealth effect is negative
- If all commodities are normal at all (\mathbf{p}, w) , demand is called normal
- If we fix prices \mathbf{p}' then $x(\mathbf{p}', w)$ is called the consumer's *Engel function* and $x_i(\mathbf{p}', w)$ is his or her *Engel curve* for good i ;
the image of $x(\mathbf{p}', w)$ is known as the *wealth expansion path*

Comparative statics w.r.t. prices

- Derivative $\partial x_i(\mathbf{p}, w)/\partial p_k$ is the price effect of p_k on demand for good i at (\mathbf{p}, w) ; the Jacobian matrix $D_{\mathbf{p}}x(\mathbf{p}, w)$ collects these in a compact format
- Good i is said to be a *Giffen good* at (\mathbf{p}, w) if $\partial x_i(\mathbf{p}, w)/\partial p_i > 0$, i.e., an increase (drop) in i 's price raises (reduces) the demand for it
- Preview: under WARP and Walras' law, a commodity can only be Giffen if it is also (very) inferior, e.g., a very low-quality good purchased by a poor consumer
- We commonly plot $x_i(\mathbf{p}, w)$ as a function of p_i for fixed \mathbf{p}_{-i} and w ; the image of $x_i(\mathbf{p}, w)$ in, e.g., x_1 - x_2 -space when only p_i is varied is known as an *offer curve*

3.4 Minimal condition for rationalizing demand

- $\mathcal{B}^{\mathcal{W}} = \{B_{\mathbf{p}, w} : \mathbf{p} \in \mathbf{R}_{++}^L \wedge w > 0\}$ and $x(\mathbf{p}, w)$ define a choice structure
- If $x(\mathbf{p}, w)$ is single-valued, i.e., a function, then WARP becomes:

$$x(\mathbf{p}, w) \neq x(\mathbf{p}', w') \quad \wedge \quad \mathbf{p} \cdot x(\mathbf{p}', w') \leq w \quad \Rightarrow \quad \mathbf{p}' \cdot x(\mathbf{p}, w) > w'$$

- That is:
If $x(\mathbf{p}', w')$ is affordable in price-wealth situation (\mathbf{p}, w) but ignored, then choice of $x(\mathbf{p}', w')$ at (\mathbf{p}', w') requires that $x(\mathbf{p}, w)$ would blow the budget in situation (\mathbf{p}', w') (if $x(\mathbf{p}, w)$ is *revealed preferred* to $x(\mathbf{p}', w')$ then (\mathbf{p}', w') mustn't be *revealed preferred* to $x(\mathbf{p}, w)$... but choice of $x(\mathbf{p}', w')$ at (\mathbf{p}', w') would reveal so if $x(\mathbf{p}, w)$ were also affordable at (\mathbf{p}', w'))
- NB:
WARP is not sufficient to conclude that demand can be rationalized by a preference relation over commodity bundles (why?)

Slutsky wealth compensation

- A price change has two effects:
 - ① It alters the relative price of different commodities
 - ② It changes the consumer's real wealth (affordability)
- Weak axiom restricts demand changes in response to price changes when taking affordability into account
- One can isolate the effect of relative price changes by adjusting the budget in a way that keeps the baseline bundle just affordable, i.e., consider $w' = \mathbf{p}' \cdot \mathbf{x}(\mathbf{p}, w)$
- This adjustment is known as a *Slutsky wealth compensation*, resulting in *Slutsky compensated price changes*

WARP \approx compensated law of demand

- If the Walrasian demand function $x(\mathbf{p}, w)$ is homogeneous of degree zero and satisfies Walras' law, WARP is equivalent to the *compensated law of demand (CLD)*:

$x(\mathbf{p}, w)$ satisfies WARP

\Leftrightarrow For any **compensated** price change from (\mathbf{p}, w) to

$$(\mathbf{p}', w') = (\mathbf{p}', \mathbf{p}' \cdot x(\mathbf{p}, w))$$

we have

$$(\mathbf{p}' - \mathbf{p}) \cdot [x(\mathbf{p}', w') - x(\mathbf{p}, w)] \leq 0$$

with strict inequality whenever $x(\mathbf{p}', w') \neq x(\mathbf{p}, w)$

- As a special case we have:

$$\Delta \mathbf{p} = (\mathbf{p}' - \mathbf{p}) = (0, \dots, 0, \Delta p_i, 0, \dots, 0) \text{ implies } \Delta p_i \Delta x_i \leq 0;$$

so price p_i and *compensated* demand x_i move in opposite directions

- Question: Should the same be true for uncompensated demand?

Substitution and income effects

- Let us fix a reference bundle $\mathbf{z} = \mathbf{x}(\mathbf{p}^0, w^0)$ and look at the *Slutsky compensated demand function* $x^s(\mathbf{p}, \mathbf{z}) \equiv x(\mathbf{p}, \mathbf{p} \cdot \mathbf{z})$
- As prices vary, $x^s(\mathbf{p}, \mathbf{z})$ changes; this change reflects a pure *substitution effect*: the consumer responds to new relative prices, while his or her real wealth has stayed constant (in the sense of \mathbf{z} still being affordable)
- A change Δx_i in uncompensated demand can be decomposed into such a (virtual) substitution effect $\Delta x_i^{\text{sub.}}$ and the income effect $\Delta x_i^{\text{inc.}}$ from the (virtual) change in income from $\mathbf{p} \cdot \mathbf{z}$ to w^0
- Looking at marginal price changes, i.e., taking the derivative of $x_i^s(\mathbf{p}, \mathbf{z}) \equiv x_i(\mathbf{p}, \mathbf{p} \cdot \mathbf{z})$ w.r.t p_k at \mathbf{p}_0 , one obtains the *Slutsky equation*

$$\underbrace{\frac{\partial x_i^s(\mathbf{p}^0, \mathbf{z})}{\partial p_k}}_{s_{ik} :=} = \frac{\partial x_i(\mathbf{p}^0, w^0)}{\partial p_k} + \frac{\partial x_i(\mathbf{p}^0, w^0)}{\partial w} \cdot x_k(\mathbf{p}^0, w^0)$$

Slutsky matrix

- These pure substitution effects (of a change in p_k on demand for commodity i) can be collected in an $L \times L$ -matrix, known as the *substitution* or *Slutsky matrix* $S(\mathbf{p}, w)$ [= $D_p x^s(\mathbf{p}, \mathbf{z})$ with $\mathbf{z} = x(\mathbf{p}, w)$]
- Multiplying $s_{ik} = \partial x_i^s(\mathbf{p}, \mathbf{z}) / \partial p_k$ with the change Δp_k for $k = 1, \dots, L$ and adding these changes up, we obtain the total change Δx_i caused by a compensated price change $\Delta \mathbf{p}$ (infinitesimal units)
- Doing this for all $i = 1, \dots, L$, we get the change in compensated demand $\Delta \mathbf{x} = S(\mathbf{p}, w) \Delta \mathbf{p}$ caused by price change $\Delta \mathbf{p}$
- The compensated law of demand, namely $\Delta \mathbf{p} \cdot \Delta \mathbf{x} \leq 0$, thus requires that

$$\Delta \mathbf{p} \cdot S(\mathbf{p}, w) \Delta \mathbf{p} \leq 0$$

holds for any $\Delta \mathbf{p} \in \mathbf{R}^L$

Negative semidefiniteness of Slutsky matrix

- So the assumptions of Walras' law, homogeneity of degree zero, and WARP (\Leftrightarrow CLD) imply that above quadratic form is never positive, i.e., $S(\mathbf{p}, w)$ is *negative semidefinite*
(mathematicians sometimes restrict the term to symmetric matrices; but symmetry of $S(\mathbf{p}, w)$ is *not* implied by Walras' law, WARP and homogeneity for $L > 2$)
- Negative semidefiniteness requires that, in particular, $s_{i,i} = \partial x_i^s / \partial p_i$ is non-positive for every i (echoing that the compensated law of demand requires $\Delta p_i \Delta x_i \leq 0$);
that's why $\partial x_i / \partial p_i = \partial x_i^s / \partial p_i - \partial x_i / \partial w > 0$ (i.e., i is a Giffen good) requires $\partial x_i / \partial w < 0$
- Given that the virtual substitution effects $\partial x_i^s(\mathbf{p}, \mathbf{z}) / \partial p_k$ can be inferred from real and, at least in principle, observable price and wealth effects at (\mathbf{p}, w) , the joint hypothesis of a consumer's behavior satisfying Walras' law, homogeneity of degree zero, and WARP can be tested empirically

Remarks

- Negative semidefiniteness of $S(\mathbf{p}, w)$ is a necessary implication of WARP (given Walras' law and homogeneity), but not yet sufficient to guarantee that a differentiable demand function satisfies WARP
(sufficiency requires that $\Delta \mathbf{p} \cdot S(\mathbf{p}, w) \Delta \mathbf{p} < 0$ holds strictly if $\Delta \mathbf{p}$ is not proportional to \mathbf{p})
- A theory of consumer demand based on the assumption of homogeneity of degree zero, Walras' law, and WARP is a bit less restrictive than one based on rational preference maximization;
as we'll see in the next chapter, rational preferences force the Slutsky matrix to be *symmetric* at all (\mathbf{p}, w)



UNIVERSITÄT
BAYREUTH

Advanced Microeconomics I

Summer Term 2025

Schedule for lectures

#	Date	Topic	Chs. in MWG
1	28.04.	Introduction	
2	05.05.	Preference and choice	1.A–D
3	12.05.	Consumer choice	2.A–F
4	19.05.	Classical demand theory	3.A–E, G
5	26.05.	Aggregate demand	3.I; 4.A–D
6	02.06.	Choice under risk	6.A–D, F
7	16.06.	Static games of complete information	7.A–E; 8.A–D, F
8	23.06.	Dynamic games of complete information	9.A–B; 12. App. A
9	30.06.	Games of incomplete information	8.E; 9.C
10	07.07.	Competitive markets	10.A–G
11	14.07.	Market power	12.A–F
12	21.07.	Question session for exam (→ t.b.a.)	

4. Preference-based demand theory

- The classical approach to consumer theory tries to explain demand by rational preferences \succsim over commodity bundles
[vs. description of choices from Walrasian budget sets by $(\mathcal{B}^w, x(\cdot))$]
- We'll assume that \succsim can be represented by a utility function u , and that u is sufficiently “smooth” /differentiable
- A rational consumer's demand can be seen as the result of
 - ▶ maximizing utility under the constraint that a given budget is not blown
 - or of
 - ▶ minimizing expenditure under the constraint of a target utility level
- The latter perspective will be useful for comparing individual welfare across different price vectors (e.g., policy interventions)

4.1 Preference relations and utility

- Many qualitative properties of \succsim imply analogue properties of u :
 - \succsim is *strictly monotone* : $\Leftrightarrow \{\mathbf{y} \geq \mathbf{x} \wedge \mathbf{y} \neq \mathbf{x} \Rightarrow \mathbf{y} \succ \mathbf{x}\}$
 $\Leftrightarrow u$ is strictly increasing
 - \succsim is *locally nonsatiated* : $\Leftrightarrow \forall \mathbf{x} \in X : \forall \epsilon > 0 : \exists \mathbf{y} \in U_\epsilon(\mathbf{x}) : \mathbf{y} \succ \mathbf{x}$;
this is implied by monotonicity
 - \succsim is *convex* : \Leftrightarrow upper contour sets $\{\mathbf{y} \in X : \mathbf{y} \succsim \mathbf{x}\}$ are convex
 $\Leftrightarrow \{\mathbf{y} \succsim \mathbf{x} \wedge \mathbf{z} \succsim \mathbf{x} \Rightarrow \forall \alpha \in (0, 1) : \alpha \mathbf{y} + (1 - \alpha) \mathbf{z} \succsim \mathbf{x}\}$
 $\Leftrightarrow u$ is quasiconcave*
 - \succsim is *strictly convex* : $\Leftrightarrow \{\mathbf{y} \succsim \mathbf{x} \wedge \mathbf{z} \succsim \mathbf{x} \wedge \mathbf{y} \neq \mathbf{z} \Rightarrow \forall \alpha \in (0, 1) : \alpha \mathbf{y} + (1 - \alpha) \mathbf{z} \succ \mathbf{x}\}$
 $\Leftrightarrow u$ is strictly quasiconcave
- * : \Leftrightarrow upper level sets $\{\mathbf{x} \in X : u(\mathbf{x}) \geq a\}$ are convex for all $a \in R$
 $\Leftrightarrow \forall \mathbf{x} \neq \mathbf{y} : \forall \lambda \in (0, 1) : u(\lambda \mathbf{x} + (1 - \lambda) \mathbf{y}) \geq \min \{u(\mathbf{x}), u(\mathbf{y})\}$

4.1 Preference relations and utility

- \succsim is *homothetic* $:\Leftrightarrow \{\mathbf{x} \sim \mathbf{y} \Rightarrow \forall \alpha \geq 0 : \alpha \mathbf{x} \sim \alpha \mathbf{y}\}$
 $\Leftrightarrow \exists u : u$ is homogeneous of degree 1*
 - \succsim is *quasilinear* w.r.t. good i
 $:\Leftrightarrow \{\text{good } i \text{ is desirable}^{**} \wedge$
 $\mathbf{x} \sim \mathbf{y} \Rightarrow \forall \alpha \in \mathbf{R} : (\mathbf{x} + \alpha \mathbf{e}_i) \sim (\mathbf{y} + \alpha \mathbf{e}_i)\}$
 $\Leftrightarrow \exists u : u(\mathbf{x}) = x_i + \phi(\mathbf{x}_{-i})$
- *: $\Leftrightarrow \forall \mathbf{x} : \forall \lambda > 0 : u(\lambda \cdot \mathbf{x}) = \lambda \cdot u(\mathbf{x})$
- **: $\Leftrightarrow \forall \mathbf{x} : \forall \alpha > 0 : (\mathbf{x} + \alpha \mathbf{e}_i) \succ \mathbf{x}$

4.2 Utility maximization problem

- If all $p_i > 0$ and u is continuous, the consumer's *utility maximization problem*

$$\max_{\mathbf{x} \geq 0} u(\mathbf{x}) \quad \text{s.t.} \quad \mathbf{p} \cdot \mathbf{x} \leq w \quad (\text{UMP})$$

has a solution (\rightarrow extreme value theorem):

the consumer's (*Walrasian or Marshallian*) demand $x(\mathbf{p}, w)$

- Assume u represents locally nonsatiated preferences \succsim then $x(\mathbf{p}, w)$
 - ▶ is convex-valued if u is quasiconcave (\succsim convex)
 - ▶ is singleton-valued, i.e., a function, and continuous at all $(\mathbf{p}, w) > \mathbf{0}$ if u is strictly quasiconcave (\succsim strictly convex)
 - ▶ satisfies Walras' law and is homogeneous of degree 0
- NB: Lagrange multiplier in (UMP) is the marginal utility of wealth
- The utility value of (UMP), $v(\mathbf{p}, w) := u(x(\mathbf{p}, w))$, is the consumer's *indirect utility function*

4.3 Expenditure minimization problem

- The *expenditure minimization problem*

$$\min_{\mathbf{x} \geq 0} \mathbf{p} \cdot \mathbf{x} \quad \text{s.t.} \quad u(\mathbf{x}) \geq \hat{u} \quad (\text{EMP})$$

is related to (UMP), often called its “dual problem”

- Its cost value $e(\mathbf{p}, \hat{u})$ is the consumer's *expenditure function*
- Analogously to a firm's cost function, if u is continuous and \succsim locally nonsatiated then $e(\mathbf{p}, \hat{u})$ is strictly increasing in \hat{u} , homogeneous of degree 1 in \mathbf{p} , nondecreasing in p_i , and weakly concave in \mathbf{p}
(intuition for the latter: 1. raise expenditure linearly by sticking to the old consumption quantities at new prices; 2. lower costs by re-optimizing)
- Note that $e(\mathbf{p}, v(\mathbf{p}, w)) = w$ and $v(\mathbf{p}, e(\mathbf{p}, \hat{u})) = \hat{u}$

Hicksian Demand

- (EMP)'s solution bundle(s) constitute the *Hicksian demand* (or *Hicks compensated demand*) $h(\mathbf{p}, \hat{u})$
- For strictly convex \succsim , $h(\mathbf{p}, \hat{u})$ is a function;
it is homogeneous of degree zero in \mathbf{p} , and satisfies the *compensated law of demand*

$$(\mathbf{p}' - \mathbf{p})[h(\mathbf{p}', \hat{u}) - h(\mathbf{p}, \hat{u})] \leq 0$$

- Goods l and k are called *substitutes* if $\frac{\partial h_l(\mathbf{p}, \hat{u})}{\partial p_k} > 0$
- Goods l and k are called *complements* if $\frac{\partial h_l(\mathbf{p}, \hat{u})}{\partial p_k} < 0$

4.4 Hicksian demand and the expenditure function

- Even though Hicks compensation (keeping utility constant) and Slutsky compensation (keeping the old bundle affordable) produce different demand changes for a discrete price change, they coincide for *marginal* price changes
- In particular, the Slutsky matrix $S(\mathbf{p}, w)$ equals the Jacobian of both $x(\mathbf{p}, \mathbf{p} \cdot x(\mathbf{p}, w))$ and $h(\mathbf{p}, v(\mathbf{p}, w))$ w.r.t. \mathbf{p}
- Note that $e(\mathbf{p}, \hat{u}) = \mathbf{p} \cdot h(\mathbf{p}, \hat{u})$ implies

$$\frac{\partial e(\cdot)}{\partial p_i} = h_i(\mathbf{p}, \hat{u})$$

(where $[+\sum p_j \cdot \frac{\partial h_j}{\partial p_i}] = 0$ because (h_1^*, \dots, h_L^*) is chosen optimally, i.e., $p_j = \lambda^{-1} \cdot \frac{\partial u}{\partial x_j} \big|_{x_j=h_j(\cdot)}$, and so $[\dots]$ equals $\lambda^{-1} \cdot$ total utility change from quantity adjustment which, for constant \hat{u} , must be zero)

- So the marginal expenditure change that is required to keep utility constant after a change of p_i is just equal to current quantity consumed of good i
(this mimicks *Shepard's lemma* in the theory of production)

Symmetry of (UMP)/(EMP)-implied Slutsky matrix

- Assuming $e(\mathbf{p}, \hat{u})$ is twice continuously differentiable, we have

$$\frac{\partial^2 e(\cdot)}{\partial p_i \partial p_j} = \frac{\partial h_i(\cdot)}{\partial p_j} = \frac{\partial h_j(\cdot)}{\partial p_i}$$

or in matrix notation

$$D_p^2 e(\mathbf{p}, \hat{u}) = D_p h(\mathbf{p}, \hat{u}) = S(\mathbf{p}, e(\mathbf{p}, \hat{u}))$$

- So the Hesse matrix $D_p^2 e(\mathbf{p}, \hat{u}) = S(\mathbf{p}, e(\mathbf{p}, \hat{u}))$ is symmetric, i.e., the Slutsky matrix is *symmetric*
 - Since $e(\mathbf{p}, \hat{u})$ is concave in \mathbf{p} , $S(\mathbf{p}, w)$ must moreover be *negative semidefinite*
- ⇒ Preference-based (or utility-maximizing) demand implies negative semidefiniteness *and* symmetry of the Slutsky matrix;
hence it is more restrictive than choice-based demand satisfying Walras' law, WARP and homogeneity of degree zero

Remarks

- Revealed choice-based demand can be rationalized if it satisfies Walras' law and WARP (\Rightarrow zero-homogeneity) *and* has a symmetric substitution matrix (latter is equivalent to satisfying Houthakker's SARP instead of WARP)
- That the derivative of (EMP)'s value function is simply (EMP)'s solution vector cannot have a direct equivalent in the (UMP):
indirect utility $v(\mathbf{p}, w)$ is ordinal while $x(\mathbf{p}, w)$ is cardinal
- But there exists a close analogue, in which marginal (indirect) utility is “normalized”, known as *Roy's identity*:

$$x_i(\mathbf{p}, w) = - \frac{\frac{\partial v(\mathbf{p}, w)}{\partial p_i}}{\frac{\partial v(\mathbf{p}, w)}{\partial w}}$$

- This makes indirect utility functions convenient to work with:
demand can be computed w/o solving an optimization problem



UNIVERSITÄT
BAYREUTH

Advanced Microeconomics I

Summer Term 2025

Schedule for lectures

#	Date	Topic	Chs. in MWG
1	28.04.	Introduction	
2	05.05.	Preference and choice	1.A–D
3	12.05.	Consumer choice	2.A–F
4	19.05.	Classical demand theory	3.A–E, G
5	26.05.	Aggregate demand	3.I; 4.A–D
6	02.06.	Choice under risk	6.A–D, F
7	16.06.	Static games of complete information	7.A–E; 8.A–D, F
8	23.06.	Dynamic games of complete information	9.A–B; 12. App. A
9	30.06.	Games of incomplete information	8.E; 9.C
10	07.07.	Competitive markets	10.A–G
11	14.07.	Market power	12.A–F
12	21.07.	Question session for exam (→ t.b.a.)	

4.4 Individual welfare evaluation

- We can evaluate whether a consumer is better off under price vector \mathbf{p}' or \mathbf{p}'' by checking if $v(\mathbf{p}', w) - v(\mathbf{p}'', w)$ is positive or negative
- Recall that we obtain an equivalent (indirect) utility function \tilde{u} (\tilde{v}) if we apply a strictly increasing transformation to u (v);
e.g., $e(\mathbf{p}', v(\mathbf{p}, w))$ is also an indirect utility function
- It is *money metric*: it evaluates \mathbf{p} -vectors by the euro amount that the consumer would need to get (\mathbf{p}, w) -situation utility under fixed reference prices \mathbf{p}' :
 - ▶ If under \mathbf{p}' , say, 100€ would be needed to obtain utility $v(\mathbf{p}^0, w)$ while 120€ would be needed to obtain $v(\mathbf{p}^1, w)$, then welfare can, loosely speaking, be said to be 20€ higher for \mathbf{p}^1 than for \mathbf{p}^0

Compensating variation

- Suppose we want to use $e(\mathbf{p}', v(\mathbf{p}, w))$ in order to quantify the change in a consumer's welfare caused by going from \mathbf{p}^0 to \mathbf{p}^1 : what should be the reference price \mathbf{p}' ?
- One natural choice is $\mathbf{p}' = \mathbf{p}^1$, i.e., we use new prices as our reference
- The change $CV(\mathbf{p}^0, \mathbf{p}^1, w) := e(\mathbf{p}^1, v(\mathbf{p}^1, w)) - e(\mathbf{p}^1, v(\mathbf{p}^0, w)) = w - e(\mathbf{p}^1, v(\mathbf{p}^0, w))$ is known as the *compensating variation*
- It measures the welfare effect of $\mathbf{p}^0 \rightarrow \mathbf{p}^1$ on the consumer by answering the question: How much money would need to be paid to the consumer (could be extracted from her) under the less (more) favorable \mathbf{p}^1 in order for her to be indifferent to the change, i.e., to feel fully compensated in the new situation?
[a negative sign indicates: *without* compensation, the consumer is worse off ...]

Equivalent variation

- Another natural choice is $\mathbf{p}' = \mathbf{p}^0$, i.e., we use old prices as our reference
- The change $EV(\mathbf{p}^0, \mathbf{p}^1, w) := e(\mathbf{p}^0, v(\mathbf{p}^1, w)) - e(\mathbf{p}^0, v(\mathbf{p}^0, w)) = e(\mathbf{p}^0, v(\mathbf{p}^1, w)) - w$ is known as the *equivalent variation*
- It measures the welfare effect of $\mathbf{p}^0 \rightarrow \mathbf{p}^1$ on the consumer by answering the question: How much money would need to be paid to the consumer (could be extracted from him) under \mathbf{p}^0 in order for him to be indifferent to the change to a more (less) favorable \mathbf{p}^1 , i.e., what is the cash equivalent of the welfare change from a pre-change perspective?

[here a negative sign indicates: consumer is willing to pay for preventing change ...]

Consumer surplus


- If \mathbf{p}^0 and \mathbf{p}^1 differ only in the price of a normal good i then

$$CV(\mathbf{p}^0, \mathbf{p}^1, w) < \Delta CS < EV(\mathbf{p}^0, \mathbf{p}^1, w)$$

where ΔCS is the change in (Marshallian) *consumer surplus*

- CS adds up *marginal willingness to pay* for all units of good i (from 0 up to $x_i(\mathbf{p}, w)$) and subtracts the actual payment for $x_i(\mathbf{p}, w)$:
 - ▶ Denote by $p_i(x_i)$ good i 's price s.t. consumer would buy x_i units (for given \mathbf{p}_{-i} and w)
 - ▶ She'd strictly prefer to buy the last marginal unit of total x_i if $p_i < p_i(x_i)$
but is indifferent if $p_i = p_i(x_i)$
 $\implies MWTP_i(x_i) = p_i(x_i)$
- Remarks:
As $MWTP_i(x_i)$ and ΔCS relate to uncompensated demand, their interpretation is obscured by income effects and induced surplus changes for other products;
if multiple prices change, product-specific CS-changes cannot meaningfully be added

CV, EV, and consumer surplus



- If there is no wealth effect for good i (e.g., \succsim is quasilinear w.r.t some good $j \neq i$, so that any extra utility from $w \uparrow$ comes via $x_j \uparrow$), then $h_i(\mathbf{p}, u^1) = h_i(\mathbf{p}, u^0)$ and all three measures coincide

5. Aggregate demand

- Aggregate demand in an economy is readily obtained by adding individual demand $x^i(\mathbf{p}, w^i)$ across all individuals, i.e.,

$$x(\mathbf{p}, w^1, \dots, w^I) = \sum_i x^i(\mathbf{p}, w^i)$$

- Tracking vector (w^1, \dots, w^I) in, e.g., comparative static analysis is cumbersome; one is tempted to work with aggregate wealth $w = \sum_i w^i$ and to pretend that $x(\mathbf{p}, w)$ is the demand of a single 'representative' agent
- This raises questions:
 - ▶ When is it possible to work with w instead of the full wealth distribution (w^1, \dots, w^I) ?
 - ▶ Assuming that individual demands are preference-based and (\mathbf{p}, w) determines aggregate demand, are the choices $x(\mathbf{p}, w)$ compatible with existence of a single *rational* representative consumer?
 - ▶ Can the representative consumer's (money-metric) indirect utility function be used for welfare statements?

5.1 When doesn't the wealth distribution matter?

- Total demand $x(\mathbf{p}, w^1, \dots, w^I) = \sum_i x^i(\mathbf{p}, w^i)$ can be expressed as a function $x(\mathbf{p}, w)$ of total wealth $w = \sum_i w^i$ only in special cases
- Distribution independence requires that individual wealth effects exactly cancel out as we shift Δw between consumers i and j , i.e.,

$$\left. \frac{\partial x_k^i}{\partial w} \right|_{(\mathbf{p}, w^i)} = \left. \frac{\partial x_k^j}{\partial w} \right|_{(\mathbf{p}, w^j)} \quad \text{for all } k \text{ and arbitrary } i, j, w^i \text{ and } w^j$$

- This necessitates that consumers (for the relevant wealth range) have parallel straight lines as their wealth expansion paths
- That turns out to be equivalent to each \succsim_i admitting a utility representation s.t. indirect utility functions are of the *Gorman form*

$$v_i(\mathbf{p}, w^i) = a_i(\mathbf{p}) + b(\mathbf{p}) \cdot w^i$$

with *identical* wealth multiplier $b(\mathbf{p})$ for all i

Cases when $x(\mathbf{p}, w^1, \dots, w^I) = x(\mathbf{p}, w)$

- This is the case (mainly) if
 - ▶ all \succsim_i equal the *same* homothetic \succsim
(e.g., Cobb-Douglas, perfect substitutes, or complements)or
 - ▶ all \succsim_i are quasilinear w.r.t. the same good k and we only consider sufficiently big wealth levels
- But we can also, trivially, drop (w^1, \dots, w^I) and simply write $x(\mathbf{p}, w)$ if each w^i can be expressed as a function $w^i(\mathbf{p}, w)$ of \mathbf{p} and w
(e.g., because of wealth redistribution according to a fixed rule, or as an empirical ‘regularity’)

5.2 Aggregate demand $\stackrel{?}{=}$ demand of a single \succsim

- Even when $\sum_i x^i(\mathbf{p}, w^i) = x(\mathbf{p}, w)$:
that each $x^i(\mathbf{p}, w^i)$ satisfies WARP, or results from a rational \succsim_i , does *not* guarantee that $\sum_i x^i(\mathbf{p}, w^i)$ satisfies WARP, or comes from a 'representative' rational \succsim
- WARP ($\hat{=}$ compensated law of demand) doesn't 'aggregate':
a price-wealth change that is compensated for the aggregate may fail to be compensated for some individuals ...
- The stronger *uncompensated law of demand* (ULD)

$$(\mathbf{p}' - \mathbf{p}) \cdot [x^i(\mathbf{p}', w^i) - x^i(\mathbf{p}, w^i)] \geq 0$$

does aggregate when $w^i \equiv \alpha_i \cdot w$

- So, if all $x^i(\cdot)$ satisfy ULD (and hence also CLD), the $x(\cdot)$ -induced choice structure will satisfy WARP (example: all \succsim_i are homothetic)

Positive representative consumer

- We say that a *positive representative consumer* exists for a given economy if one can find a fictional individual whose optimal behavior would result in aggregate demand $x(\mathbf{p}, w^1, \dots, w^I)$ if she could spend the society's budget $w = \sum w^i$
- Existence requires that
 - ▶ distribution (w^1, \dots, w^I) doesn't matter, so that $x(\mathbf{p}, w^1, \dots, w^I) = x(\mathbf{p}, w)$ and
 - ▶ $x(\mathbf{p}, w)$ satisfies WARP
(in fact, even Houthakker's SARP)
- Note that it is also possible that aggregate demand satisfies more stringent 'consistency requirements' than individual demands do: individual violations of, say, ULD may 'average out'

5.3 Aggregate welfare evaluation

- A social planner, who evaluates different (\mathbf{p}, w) -situations for society as a whole may plausibly consider a (Bergson-Samuelson) *social welfare function* $W: \mathbf{R}^I \rightarrow \mathbf{R}$ which is defined on (indirect) utility vectors (u_1, \dots, u_I) and is non-decreasing in every u_i
- Prominent examples:
 - ▶ utilitarian welfare $W^U(u_1, \dots, u_I) = \sum_i u_i$
 - ▶ ‘Rawlsian’ welfare $W^R(u_1, \dots, u_I) = \min\{u_1, \dots, u_I\}$
- Such a social aggregation rule implicitly requires interpersonal comparability of utility or an ‘all individuals are alike’ assumption

Normative representative consumer

- To what extent can social welfare evaluation be simplified to individual welfare evaluation for the representative consumer?
- Answer depends on the considered social welfare function
- The positive representative consumer with preferences \succsim is called a *normative representative consumer relative to social welfare function* $W(\cdot)$ if the value function $W^*(\mathbf{p}, w)$ of the planner's welfare maximization problem

$$\begin{aligned} & \max_{w^1, \dots, w^I} W(v_1(\mathbf{p}, w^1), \dots, v_I(\mathbf{p}, w^I)) \\ \text{s.t. } & \sum w^I \leq w \end{aligned}$$

is an indirect utility function for \succsim , i.e., if the representative consumer's demand corresponds to the aggregate demand* which would result from *utility-maximizing individual demands after an optimal wealth redistribution*

(*: apply Roy's identity to $W^*(\mathbf{p}, w)$)

Welfare vs. normative representative consumer

- If a normative representative consumer exists, we can, in principle, say that $\mathbf{p}^0 \rightarrow \mathbf{p}^1$ is socially beneficial or detrimental by looking at $CV(\mathbf{p}^0, \mathbf{p}^1, w)$, $EV(\mathbf{p}^0, \mathbf{p}^1, w)$, or ΔCS for that consumer
- But: w 's optimal distribution (w^{1*}, \dots, w^{I*}) , which maximizes $W(\cdot)$, generally depends on \mathbf{p} ;
hence, saying “ $\mathbf{p}^0 \rightarrow \mathbf{p}^1$ is socially beneficial because $\Delta CS > 0$ ” is only warranted in the sense that there *exists* a redistribution scheme s.t. welfare is higher under \mathbf{p}^1
 (“potential welfare” $W^*(\mathbf{p}, w)$ is higher while actual welfare $W(v_1(\mathbf{p}, w^1), \dots, v^I(\mathbf{p}, w^I))$ may be lower for $\mathbf{p} = \mathbf{p}^1$ than $\mathbf{p} = \mathbf{p}^0$ if wealth is not redistributed)

Existence of a normative representative consumer

- Conditions for existence of a positive representative consumer were already very demanding
- And if a positive representative consumer happens to exist, there is no guarantee that he is also a normative one for the considered welfare function $W(\cdot)$; it is even possible that his preferences have no normative content for *any* social welfare function
- However, if all consumers have indirect utility of the Gorman form *with identical* $b(\mathbf{p})$, then the positive representative consumer also is a normative one
(the Gorman form imposes sufficient structure for $v(\mathbf{p}, w) = \sum_i a_i(\mathbf{p}) + b(\mathbf{p}) \cdot w$ to be a strictly increasing transformation of the planner's value function for any social welfare function $W(\cdot)$)



UNIVERSITÄT
BAYREUTH

Advanced Microeconomics I

Summer Term 2025

Please go online and browse to

`http://www.pingo.coactum.de`

Enter “Zugangsnummer” 200427

Schedule for lectures

#	Date	Topic	Chs. in MWG
1	28.04.	Introduction	
2	05.05.	Preference and choice	1.A–D
3	12.05.	Consumer choice	2.A–F
4	19.05.	Classical demand theory	3.A–E, G
5	26.05.	Aggregate demand	3.I; 4.A–D
6	02.06.	Choice under risk	6.A–D, F
7	16.06.	Static games of complete information	7.A–E; 8.A–D, F
8	23.06.	Dynamic games of complete information	9.A–B; 12. App. A
9	30.06.	Games of incomplete information	8.E; 9.C
10	07.07.	Competitive markets	10.A–G
11	14.07.	Market power	12.A–F
12	21.07.	Question session for exam (→ t.b.a.)	

Please go online and browse to

`http://www.pingo.coactum.de`

Enter “Zugangsnummer” 200427

6. Choice under risk and uncertainty

- Lecture 2 considered preferences and choice w/o specific assumptions re. the considered alternatives $\{x_1, x_2, \dots\}$; they might involve risk, uncertainty, different points in time, space, etc.
- We now specifically consider risky alternatives, i.e., options associated with known objective probability distributions over deterministic outcomes (= *lotteries*) (vs. *uncertain / ambiguous* alternatives = *prospects*)
- One may distinguish between simple lotteries $L = (\pi_1, \dots, \pi_N)$ over deterministic outcomes $Y = \{y_1, \dots, y_N\}$, and *compound lotteries* ('lotteries over lotteries')
- From a consequentialist perspective, a compound lottery can be equated with the simple lottery which it induces; we therefore focus on the set $\Delta(Y)$ of simple lotteries

6.1 Expected utility representations

- We know that if an agent has complete, transitive and continuous preferences \succsim over the space $\Delta(Y)$ of all (simple) lotteries L , then \succsim can be represented by a utility function $U: \Delta(Y) \rightarrow \mathbf{R}$
- Here, *continuity* may, e.g. be simplified to:

$$\begin{aligned} \forall L, L', L'': \{ \alpha \in [0, 1] : \alpha L \oplus (1 - \alpha)L' \succsim L'' \} & \quad \text{and} \\ \{ \alpha \in [0, 1] : L'' \succsim \alpha L \oplus (1 - \alpha)L' \} & \quad \text{are closed sets} \end{aligned}$$

- The function $U(\cdot)$ that maps each distribution L to a number may be highly complicated and unwieldy (e.g., involve a “Choquet integral” w.r.t. a “capacity” derived from L)
- However, if \succsim additionally satisfies the *von Neumann-Morgenstern independence axiom*

$$\forall L, L', L'' : \forall \alpha \in (0, 1) : L \succsim L' \Leftrightarrow \alpha L \oplus (1 - \alpha)L'' \succsim \alpha L' \oplus (1 - \alpha)L'',$$

then $U(\cdot)$ can be chosen to have a simple functional form

Von Neumann-Morgenstern expected utility

- In particular, $U(\cdot)$ can be chosen to have the *v.N.-M.-expected utility* form, that is: there exists a (*Bernoulli*) *utility function* $u: Y \rightarrow \mathbf{R}$ defined only for deterministic outcomes $y \in Y$ such that:

$$U(L) = \sum \pi_i \cdot u(y_i) = \mathbf{E}_L[u(y)] \quad [= \int u(y) dL(y)]$$

- \succsim 's Bernoulli utility function $u(\cdot)$ is unique up to an order-preserving affine transformation, i.e.,

$$\begin{aligned} & u(\cdot) \text{ can be chosen as Bernoulli utility function for } \succsim \\ \Leftrightarrow & \alpha u(\cdot) + \beta \text{ for } \alpha > 0 \text{ can also be chosen} \end{aligned}$$

- $u(\cdot)$ is a *cardinal* utility function over deterministic outcomes:
 $u(x) - u(y) > u(z) - u(w) > 0$ now has the interpretation that x is a bigger improvement on y than z is on w :
 - ▶ one could mix x with a greater probability for a bad outcome q and the agent still prefers this to y ... than one could mix z with q and retain preference over w

Remarks on independence axiom

- Requiring “independence” when “adding” lottery L'' to L and to L' makes sense since there is no (obvious) complementarity or substitutability across distinct y -outcomes
 - An agent whose \succsim violates independence may be “Dutch-booked”, i.e., some money can be extracted from her at no risk:
 - ▶ Suppose $L_1 \succ L_2$, but $\alpha L \oplus (1 - \alpha)L_1 \prec \alpha L \oplus (1 - \alpha)L_2$
 - ▶ Let her own $\alpha L \oplus (1 - \alpha)L_1$, while you own $\alpha L \oplus (1 - \alpha)L_2$
 - ▶ Trade lotteries with her, collect a fee, and wait
 - ▶ If L isn't realized, then trade L_1 for L_2 and collect another fee \Rightarrow Your position is exactly as without the trades (L with prob. α , L_2 with prob. $1 - \alpha$), but you additionally collect a fee one or two times
 - By independence, $L \sim L' \Rightarrow$ (i) $L \sim \alpha L \oplus (1 - \alpha)L'$ and
(ii) $\alpha L \oplus (1 - \alpha)L'' \sim \alpha L' \oplus (1 - \alpha)L''$
for all $\alpha \in [0, 1]$ and any L''
- \Rightarrow \succsim 's indifference curves are straight parallel lines in the probability simplex (unless the agent is indifferent between all outcomes)

Allais paradox

- Though normatively appealing, real people frequently violate the independence axiom

- This is illustrated, e.g., by the *Allais paradox*:

For $(y_1, y_2, y_3) = (2500\text{€}, 500\text{€}, 0\text{€})$ many people reveal

- ① $L_1 = (0, 1, 0) \succ L_2 = (0.1, 0.89, 0.01)$
- ② $L_3 = (0, 0.11, 0.89) \prec L_4 = (0.1, 0, 0.9)$



- If this satisfied the v.N.-M. axioms, we could choose $u(0\text{€}) = 0$, and then infer
 - ▶ from (1): $u(500\text{€}) > 0.1 \cdot u(2500\text{€}) + 0.89 \cdot u(500\text{€}) \Leftrightarrow [1 - 0.89] \cdot u(500\text{€}) > 0.1 \cdot u(2500\text{€})$
 - ▶ from (2): $0.11 \cdot u(500\text{€}) < 0.1 \cdot u(2500\text{€})$

[L_1 and L_2 lie parallel to L_3 and L_4 in the probability simplex; so 1st choice fixes 2nd one under v.N.-M. axioms: *all* indifference lines either have greater, smaller, or same slope as these two lines]

6.2 Money lotteries and risk attitudes

- Consider lotteries over interval $[a, \infty)$ of final wealth levels as described by random variable X with cumulative distribution function $F(x) = \Pr(X \leq x)$, and v.N.-M. utility function $U(\cdot)$ with increasing Bernoulli utility $u(\cdot)$ such that $\mathbf{E}_F[u(X)]$ is finite
- The agent is said to be
 - ▶ *risk neutral* \Leftrightarrow she is indifferent between lottery F and receiving $\mathbf{E}_F[X]$ for sure, i.e., $\forall F : \mathbf{E}_F[u(X)] = u(\mathbf{E}_F[X])$
 - ▶ *(strictly) risk averse* \Leftrightarrow she (strictly) prefers $\mathbf{E}_F[X]$ for sure to F
 - ▶ *(strictly) risk loving* \Leftrightarrow she (strictly) prefers F to $\mathbf{E}_F[X]$ for sure
- By *Jensen's inequality*, u is concave iff

$$\int u(x) dF(x) \leq u\left(\int x dF(x)\right)$$

- So (strict) risk aversion is equivalent to (strict) concavity of u
- It is also equivalent to the *certainty equivalent*, i.e., sure payment $c(F, u)$ that renders agent indifferent to F , being (strictly) smaller than $\mathbf{E}_F[X]$

Quantifying and comparing risk aversion

- Risk attitudes of two individuals, or the same individual at different levels of wealth x can be compared by the *Arrow-Pratt coefficient of absolute risk aversion*

$$r_A(x; u) = -\frac{u''(x)}{u'(x)}$$

- $u_2(\cdot)$ is *more risk averse than* $u_1(\cdot)$
 - $\Leftrightarrow r_A(x; u_2) \geq r_A(x; u_1)$ for all x
 - $\Leftrightarrow c(F; u_2) \leq c(F; u_1)$ for any lottery F
 - $\Leftrightarrow u_2$ is “more concave” than u_1 , i.e., there exists an increasing concave transformation $k(\cdot)$ s.t. $u_2(x) = k(u_1(x))$

Common assumptions about risk aversion

- It is often plausible to assume that $u(\cdot)$ has *decreasing absolute risk aversion* in wealth (DARA), i.e., that $r_A(x; u)$ decreases in x
- Moreover, one often assumes that $u(\cdot)$ has *nonincreasing relative risk aversion*, i.e., the *coefficient of relative risk aversion*

$$r_R(x; u) = -x \cdot \frac{u''(x)}{u'(x)}$$

is constant or decreasing (CRRA or DRRA)

- This captures the regularity that, as an individual becomes richer, a greater absolute amount is invested in risky assets (DARA), and this amount corresponds to a weakly greater share of total wealth (CRRA or DRRA)
- Remarks:

$$-r_A(x; u) \equiv \lambda \neq 0 \text{ (CARA)} \Leftrightarrow u(x) = a_1 - a_2 \cdot e^{-\gamma x} \text{ with } a_2 > 0$$

$$-r_R(x; u) \equiv \delta \quad \text{(CRRA)} \Leftrightarrow \delta = 1 : u(x) = a_1 + a_2 \cdot \ln(x)$$

$$\delta \neq 1 : u(x) = a_1 + a_2 \cdot x^{1-\delta}$$

(Partial) orderings of random variables

- Any two agents, who like higher x better, agree that lottery F_1 is better than lottery F_2 if $F_1(x) \leq F_2(x)$ for all x ,
i.e., F_1 places less probability on small realizations of X than F_2
 $\Leftrightarrow F_1$ *first-order stochastically dominates* F_2
- Any two risk averters agree that lottery F_1 is better than lottery F_2
if F_1 and F_2 have the same mean ($\hat{=}$ expected value) and F_2 can be generated from F_1
by shifting probability towards the extremes
 $\Leftrightarrow F_2$ is a *mean-preserving spread* of F_1
- F_2 being a mean-preserving spread of F_1 is a special case of:
 F_1 *second-order stochastically dominates* F_2

Please go online and browse to

`http://www.pingo.coactum.de`

Enter “Zugangsnummer” 200427

6.3 Subjective probability theory

- If agents choose between *uncertain* prospects for which *no* objective probabilities are given, their behavior may still be represented in an “as-if”-fashion as expected utility maximization for *subjective probabilities* $(\tilde{\pi}_1, \dots, \tilde{\pi}_N)$
- The key requirements for this to be possible is *Savage’s sure thing principle (STP)*: the ranking of two *prospects* P_1 and P_2 ($\hat{=}$ mappings from states of the world to, e.g, wealth) depends only on provisions for states in which P_1 and P_2 actually differ
- In particular,

$$P_1: \begin{array}{|c|c|c|} \hline s_1 & s_2 & s_3 \\ \hline x & y & z \\ \hline \end{array} \succsim P_2: \begin{array}{|c|c|c|} \hline s_1 & s_2 & s_3 \\ \hline x' & y' & z \\ \hline \end{array}$$

if and only if

$$P_3: \begin{array}{|c|c|c|} \hline s_1 & s_2 & s_3 \\ \hline x & y & z' \\ \hline \end{array} \succsim P_4: \begin{array}{|c|c|c|} \hline s_1 & s_2 & s_3 \\ \hline x' & y' & z' \\ \hline \end{array}$$

Ellsberg paradox

- Intuitively reasonable choices under uncertainty can violate subjective expected utility maximization (e.g., because the latter cannot account for *ambiguity aversion*)
- Example:
Suppose a ball is drawn from an urn with 30 red balls, and 60 white or blue balls in unknown proportion
 - Many people strictly prefer P_1 in
 P_1 : 100€ for red, 0€ otherwise
vs. P_2 : 100€ for blue, 0€ otherwise
 - ▶ And they strictly prefer P_4 in
 P_3 : 100€ for red or white, 0€ otherwise,
vs. P_4 : 100€ for blue or white, 0€ otherwise
- The first choice indicates $\tilde{\pi}_{\text{blue}} < 1/3 = \tilde{\pi}_{\text{red}}$;
the second one indicates $2/3 > 1 - \tilde{\pi}_{\text{blue}} \Leftrightarrow \tilde{\pi}_{\text{blue}} > 1/3$
[Homework: find violation of STP if $P_1 \succ P_2$ and $P_3 \prec P_4$]





UNIVERSITÄT
BAYREUTH

Advanced Microeconomics I

Summer Term 2025

Schedule for lectures

#	Date	Topic	Chs. in MWG
1	28.04.	Introduction	
2	05.05.	Preference and choice	1.A–D
3	12.05.	Consumer choice	2.A–F
4	19.05.	Classical demand theory	3.A–E, G
5	26.05.	Aggregate demand	3.I; 4.A–D
6	02.06.	Choice under risk	6.A–D, F
7	16.06.	Static games of complete information	7.A–E; 8.A–D, F
8	23.06.	Dynamic games of complete information	9.A–B; 12. App. A
9	30.06.	Games of incomplete information	8.E; 9.C
10	07.07.	Competitive markets	10.A–G
11	14.07.	Market power	12.A–F
12	21.07.	Question session for exam (\rightarrow t.b.a.)	

7. Static games of complete information

- GT \equiv multiperson decision theory
- Each agent's utility possibly depends on actions of other agents; optimal decisions thus depend on individual beliefs about other agents' choices (which depend on *their* beliefs)
- GT works with models of real-life situations, called “games”; to these, it applies “solution concepts”
- GT helps to understand how decision makers interact if they are *rational* and reason *strategically*, i.e., if they pursue a well-defined objective and make optimal use of their knowledge about other decision makers
- Illustration by “70%-Beauty Contest game”:
 - ▶ Submit a number $s_i \in [0; 100]$
 - ▶ We'll compute the average $\bar{s} = 1/n \cdot \sum s_i$
 - ▶ The person(s) whose number is closest to $0.7 \cdot \bar{s}$ receives (share) the prize

Please go online and browse to

`http://www.pingo.coactum.de`

Enter “Zugangsnummer” 6595

Some distinctions

- We focus on *non-cooperative* GT: [\leftrightarrow cooperative GT; evolutionary GT; ...]
players may communicate but cannot commit to any agreed action;
order of moves and players' information is explicitly specified
- Players' information in a game can be
 - ▶ *complete*:
all know the game's structure and everybody's preferences
(though maybe not all of others' actions prior to a move)
 - ▶ *incomplete*:
at least one player lacks information, e.g., about others' preferences

Some distinctions

- A non-cooperative game can be
 - ▶ in *normal form* or *static* or *simultaneous-move*:
players choose a strategy (= a complete plan of action covering all contingencies) once and “simultaneously”
 - ▶ in *extensive form* or *dynamic* or *sequential-move*:
players act sequentially based on *perfect* or *imperfect* information about what has happened so far
- An extensive form game can be translated into normal form, and vice versa; dynamic information is often useful, but sometimes also distracting

A book (just in case you get hooked ...)

- PDF version can be downloaded for free by UBT students:
<https://link.springer.com/content/pdf/10.1007/978-3-642-31963-1.pdf>



7.1 Basic notation and preliminaries

- Notation:

- ▶ $N = \{1, 2, \dots, n\}$: set of agents or *players*
- ▶ S_i : set of (pure) *strategies* available to player i
- ▶ $s_i \in S_i$: a strategy of player i
- ▶ $\mathbf{S} \equiv S_1 \times \dots \times S_n$: strategy space of the game
- ▶ $\mathbf{s} = (s_1, \dots, s_n) \in \mathbf{S}$: a *strategy profile*
- ▶ $\mathbf{s}_{-i} = s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n$: profile of all except player i 's strategies
- ▶ $\mathbf{S}_{-i} \equiv S_1 \times \dots \times S_{i-1} \times S_{i+1} \times \dots \times S_n$
- ▶ $u_i: \mathbf{S} \rightarrow \mathbf{R}$: player i 's v.N.-M. utility or *payoff function*
- ▶ $\mathbf{u}: \mathbf{S} \rightarrow \mathbf{R}^n$ with $\mathbf{u}(\mathbf{s}) \equiv (u_1(\mathbf{s}), \dots, u_n(\mathbf{s}))$
- ▶ $\Delta(S_i)$: set of all probability distributions over S_i ($=i$'s *mixed strategies*)
- ▶ $\sigma_i \in \Delta(S_i)$: a mixed strategy of i
- ▶ $\boldsymbol{\sigma}, \boldsymbol{\sigma}_{-i}$: analogous

Normal form

- The *normal* or *strategic form* of a game is a triplet $\langle N, \mathbf{S}, \mathbf{u} \rangle$ specifying the players, their strategies and payoff functions
- The *mixed extension* of $\langle N, \mathbf{S}, \mathbf{u} \rangle$, denoted by $\langle N, \mathbf{\Sigma}, \mathbf{u} \rangle$ with $\mathbf{\Sigma} = \Delta(S_1) \times \dots \times \Delta(S_n)$, explicitly allows the use of mixed strategies, i.e., players can *independently* randomize over their pure strategies
- Remarks:
 - ▶ Pure strategies are just particular (degenerate) mixed strategies
 - ▶ Often the analysis concerns $\langle N, \mathbf{\Sigma}, \mathbf{u} \rangle$, but only $\langle N, \mathbf{S}, \mathbf{u} \rangle$ is mentioned
 - ▶ Utility on \mathbf{S} naturally extends to $\mathbf{\Sigma}$ by the assumption of v.N.-M. utilities

Complete information and common knowledge

- Unless otherwise stated, we will consider games of *complete information*, i.e., we assume that $\langle N, \mathbf{S}, \mathbf{u} \rangle$ and the rationality underlying \mathbf{u} are *common knowledge*
- Some fact x is called *common knowledge* if
 - ▶ everybody knows x ,
 - ▶ everybody knows that everybody knows x ,
 - ▶ everybody knows that everybody knows that everybody knows x ,
 - ▶ etc. ad infinitum
- We presume that with any facts x, y , and z players know all the logical implications of x, y , and z , too

7.2 Dominant strategies and rationalizability

- Question: Which predictions follow from (common knowledge of) rationality?
- Strategy $\sigma_i^* \in \Sigma_i = \Delta(S_i)$
 - ▶ *strictly dominates strategy $s'_i \in S_i$ (or s'_i is strictly dominated by σ_i^*)*
 $\Leftrightarrow \quad \forall \mathbf{s}_{-i} \in \mathbf{S}_{-i} : u_i(\sigma_i^*, \mathbf{s}_{-i}) > u_i(s'_i, \mathbf{s}_{-i}),$
i.e., σ_i^* is strictly better than s'_i no matter what (player i believes that) other players do
 - ▶ *weakly dominates s'_i (or s'_i is weakly dominated by σ_i^*)*
 $\Leftrightarrow \quad \begin{aligned} &\forall \mathbf{s}_{-i} \in \mathbf{S}_{-i} : u_i(\sigma_i^*, \mathbf{s}_{-i}) \geq u_i(s'_i, \mathbf{s}_{-i}) \\ &\wedge \quad \exists \mathbf{s}_{-i} \in \mathbf{S}_{-i} : u_i(\sigma_i^*, \mathbf{s}_{-i}) > u_i(s'_i, \mathbf{s}_{-i}) \end{aligned}$
i.e., σ_i^* is never worse than s'_i and sometimes strictly better
- s_i^* is *strictly dominant* if it strictly dominates *all* other $s'_i \in \mathbf{S}_i$
- If a strictly dominant strategy exists, rationality dictates its use
- For $n = 2$, a profile σ is consistent with common knowledge of rationality, i.e., is *rationalizable* iff all involved s_i survive *iterated elimination of strictly dominated strategies*

7.3 Nash equilibrium

- When many strategy profiles are rationalizable, more specific predictions can be obtained if players are assumed to have beliefs consistent with each other, i.e., i 's beliefs about \mathbf{s}_{-i} are correct for every $i \in N$
- NB: this is *not* implied by common knowledge of rationality and the game, but requires extra motivation!
- Strategy profile $\mathbf{s}^* = (s_1^*, \dots, s_n^*) \in S$ is a *Nash equilibrium* (NE)
 $\Leftrightarrow \quad \forall i \in N : \forall s_i \in S_i : u_i(s_i^*, \mathbf{s}_{-i}^*) \geq u_i(s_i, \mathbf{s}_{-i}^*),$
i.e., everybody plays a *best response*¹ to (his or her correct beliefs about the) strategy choices \mathbf{s}_{-i}^* of everybody else
[¹ There may be others!]

Remarks

- Mixed strategy NE: same for profiles $\sigma^* \in \Delta(S_1) \times \dots \times \Delta(S_n)$
- A strategy profile \mathbf{s}^* is a *strict Nash equilibrium* iff it is a NE and above inequality is strict, i.e., everyone has a unique best response to \mathbf{s}_{-i}^*
[NB: a game may have several strict NE]
- Why game theorists care about NE so much:
 - ▶ Though NE is not implied by rationality, it is “focal” amongst all rationalizable profiles: only a NE involves *consistent beliefs*
 - ▶ If there is any “unique predicted outcome” or a stable social convention for playing a particular game w/o external coordination, then it must be a NE
 - ▶ If players can talk prior to the game and agree on some profile \mathbf{s} *without* exogenous commitment or coordination, only NE are *self-enforcing*
 - ▶ A NE may be viewed as a “steady state” of play where an unspecified dynamic process has brought about correct expectations;
many learning dynamics or evolutionary processes converge to a NE

Mixed-strategy NE

Proposition

Consider the mixed extension of finite game $\langle N, \mathbf{S}, \mathbf{u} \rangle$:

σ^* is a NE of $\langle N, \Sigma, \mathbf{u} \rangle$

\Leftrightarrow For all $i \in N$, every pure strategy s_i played with positive probability under σ_i^* (i.e., any s_i that is in the *support* of σ_i^*) is a best response to σ_{-i}^*

Proof:

Recall that
$$u_i(\sigma) = \sum_{s_i \in S_i} \sigma_i(s_i) \cdot u_i(s_i, \sigma_{-i}).$$

“ \Rightarrow ” Assume some s_i in $\text{supp}(\sigma_i^*)$ is *no* best response to σ_{-i}^* .

Then $u_i(\sigma^*)$ can be *increased* by shifting probability from s_i to some s_i^* that is a best response. ⚡ to NE implying “ σ_i^* is a best response”

“ \Leftarrow ” Assume σ^* is *no* NE, i.e., for some i , σ_i^* is no best response to σ_{-i}^* while σ_i' is a best response to σ_{-i}^* . Hence some s_i' in $\text{supp}(\sigma_i')$ gives higher payoff against σ_{-i}^* than some s_i in $\text{supp}(\sigma_i^*)$. ⚡ to the premise above

Mixed-strategy NE

- That truly mixed NE involve indifference reduces their appeal
- Defense of mixed NE:
 - ▶ In some games, players try to be unpredictable and mixed NE has empirical support (penalty kicks, tennis serves, R-S-P game, ...)
 - ▶ In zero-sum games, σ_i^* maximizes i 's guaranteed expected payoff, i.e., is a “safe” strategy with minimal knowledge requirements
 - ▶ Probabilities σ_i in mixed NE may also be interpreted as *other* players' subjective beliefs about i 's play
 - ▶ A mixed NE may describe a large population where individuals are randomly matched and play pure strategies in the “right” population proportions
 - ▶ A mixed NE can be viewed as approximating a pure (Bayesian) NE of a game in which part of players' payoffs is private knowledge [\rightarrow *purification* of mixed NE à la Harsanyi]

Existence of NE

- Games with infinite pure strategy spaces may fail to have any NE
- Nash (1950) proved that every *finite* game has “an equilibrium point” (= mixed NE)
[using *Kakutani's fixed point theorem*]
- Nash's result can be extended to games with general, convex and compact strategy spaces or to show that symmetric finite games must have at least one symmetric NE

7.4 Equilibrium selection and refinement

- The key “problem” is usually not existence but multiplicity of NE
- What would you play in

(a)

1\2	F	H
F	7,7	0,0
H	0,0	9,9

?

(b)

1\2	F	H
F	7,7	8,0
H	0,8	9,9

?

(c)

1\2	f	h
F	3,1	0,0
H	2,2	2,2

?

(d) $S_1 = S_2 = [0, 100]$, $u_i(s_i, s_j) = s_i$ if $s_i + s_j = 100$ and 0 otherwise?

- A large literature discusses *selection criteria* for NE or proposes *refinements of NE* (e.g. “trembling-hand perfect NE”) that reflect specific plausibility considerations



UNIVERSITÄT
BAYREUTH

Advanced Microeconomics I

Summer Term 2025

Schedule for lectures


#	Date	Topic	Chs. in MWG
1	28.04.	Introduction	
2	05.05.	Preference and choice	1.A–D
3	12.05.	Consumer choice	2.A–F
4	19.05.	Classical demand theory	3.A–E, G
5	26.05.	Aggregate demand	3.I; 4.A–D
6	02.06.	Choice under risk	6.A–D, F
7	16.06.	Static games of complete information	7.A–E; 8.A–D, F
8	23.06.	Dynamic games of complete information	9.A–B; 12. App. A
9	30.06.	Games of incomplete information	8.E; 9.C
10	07.07.	Competitive markets	10.A–G
11	14.07.	Market power	12.A–F
12	21.07.	Question session for exam (→ t.b.a.)	

8.1 Game tree

- A *dynamic* or *sequential-move* or *extensive (form)* game adds to the information provided in static games an explicit description of
 - ▶ the timing of players' actions
 - ▶ the information about play so far on which these actions can be conditioned
- We keep the assumption of complete information, i.e., the game (incl. all preferences) is common knowledge

8.1 Game tree



- Central to the modeling of dynamic games is the concept of a game tree, e.g.



- A tree is a particular type of *directed graph*, with *nodes* (or vertices) and *edges*, each connecting two nodes

Game tree

- Formally, a *tree* is defined by
 - a set of nodes N
 - a transitive and asymmetric (i.e., $a \prec b \Rightarrow \neg(b \prec a)$) precedence relation \prec satisfying the *arborescence properties*:
 - ★ there is a unique *initial node* $n^0 \in N$ without predecessor
 - ★ if n and n' precede n'' , then either $n \prec n'$ or $n' \prec n$
(in particular, every node except n^0 has a unique direct predecessor)
- For example,



are no trees

Game tree

- Nodes without successors are called *terminal nodes*; all non-terminal nodes are called *decision nodes*
- Given N and \prec with decision nodes D , a function

$$\iota: D \rightarrow N \cup \{\text{Nature}\}$$

specifies for every decision node which player has to move

- The additional player “Nature” is a trick to model chance moves (if needed)
- For $n \in D$, $A(n)$ denotes the set of *actions* available to player $\iota(n)$
- Each $a \in A(n)$ leads to a different direct successor n' of n as defined by a function

$$\alpha(n): A(n) \rightarrow \text{Succ}(n)$$


[i.e., each non-initial node n' is reached from a unique n by a unique action $a \in A(n)$]

8.2 Information sets

- The player $\iota(n)$ who has to move at n may not know that the game is currently exactly at n , e.g., because moves of other players are imperfectly observed
- This is reflected by a partition \mathbf{P} of D into *information sets* $\{n^0\}, P^2, \dots, P^k \in \mathbf{P}$ that capture what players know when moving

Information sets

- Example:



- Here:

- ▶ 1st-moving player 1 (always) knows the entire “empty history”
- ▶ Player 2 knows 1's choice when making his first choice
- ▶ Player 1 does not know whether 2 played *up* or *down*;
neither does 2 know if 1 played *u* or *d* when making his second choice

Information partition

- The *information partition* \mathbf{P} of D into information sets must satisfy the following conditions:
 - ▶ the same player $\iota(n)$ and action set $A(n)$ are assigned to all $n \in P^j$
(so we may simply write $\iota(P^j)$ and $A(P^j)$)
 - ▶ if $n \in P^j$, then no successor of n is also contained in P^j
- Player $\iota(P^j)$ called to select an action $a \in A(P^j)$ at a node in P^j knows that moves leading to $P^i \neq P^j$ were *not* played, but doesn't know which move(s) led into P^j if that's non-singleton

8.3 Extensive game

- Formally, the collection $\langle N, N, \prec, \iota, \{A(n)\}_{n \in N}, \{\alpha(n)\}_{n \in N}, \mathbf{P} \rangle$ defines an *extensive game form*.

An extensive game form together with

- ▶ v.N.-M. utilities u_i over all (lotteries over) terminal nodes for all $i \in N$
- ▶ a probability distribution $\rho(n)$ on $A(n)$ for each n at which Nature “moves”

define an *extensive (form) game* Γ .

- Remarks:
 - ▶ The definition of a “game form” may include $\rho(n)$, too
 - ▶ Above 9-tuple (or 10-tuple in MWG) is rarely written down: usually, Γ is “defined” by a diagram or verbal description
 - ▶ We assume that players have *perfect recall*, i.e., do not forget what they learned at some stage (\rightarrow restricts possible partitions \mathbf{P})
 - ▶ If all information sets are singletons then we speak of a *game of perfect information*, otherwise of a *game of imperfect information*

Strategies in extensive games

- In extensive games, *actions* (at some information set) need to be clearly distinguished from *strategies*; strategies are complete plans that prescribe an action *for every contingency* that calls a player to move
- Denoting the set of information sets P such that $\iota(P) = i$ by \mathbf{P}_i , a (*pure*) *strategy* of player i in an extensive game is a function

$$s_i: \mathbf{P}_i \rightarrow \bigcup_{P \in \mathbf{P}_i} A(P)$$

which maps each of i 's information sets $P \in \mathbf{P}_i$ to a feasible action $s_i(P) \in A(P)$

- *Histories of play* often substitute for information sets in the description of strategies
- A player may randomize either over his pure strategies (\rightarrow *mixed strategy*) or independently over feasible actions at each information set (\rightarrow *behavior strategy*)

8.4 Backward induction

- Extensive games of perfect information can be solved by *backward induction* if there is a “last period”, i.e., if every possible history is finite:
 - ▶ One determines optimal choices for the respective last-moving players in all next-to-terminal nodes
 - ▶ One replaces these decision nodes by the selected terminal nodes (or marks the pertinent edges appropriately), and then repeats the exercise until the initial node is reached
- Every finite game of perfect information has a solution to backward induction; for “generic” games – i.e., if no two payoffs are the same – the solution is unique

8.5 Subgame perfect equilibrium

- The idea of players behaving rationally (and others anticipating this) throughout the entire game (= *sequential rationality*) can also be applied to games of imperfect information or without “last period” ...
- A *subgame* Γ_n of an extensive game Γ is an extensive game starting in a singleton information set $\{n\}$ (of Γ), containing exactly all successors of n as its other nodes, not cutting through any of Γ 's information sets and inheriting payoffs, information sets, etc. from Γ
- A strategy profile \mathbf{s}^* of Γ is a *subgame perfect equilibrium* (SPE) iff \mathbf{s}^* induces a NE in every subgame of Γ
- In games with finitely many stages, SPE can be found by a generalization of backward induction: determine a NE in all “final” subgames, replace these by the respective payoffs, and repeat until the initial node is reached

One-deviation principle

- Consider a game of perfect information or one where at each stage players move simultaneously and afterwards observe all actions:
 - ▶ Obviously, \mathbf{s}^* is a SPE *only if* no player i has a strategy s'_i that differs from s_i^* in *just one* information set $P \in \mathbf{P}_i$ and does strictly better than s_i^* conditional on P being reached
 - ▶ The reverse is also true and known as the
- One-deviation principle:
 \mathbf{s}^* is a SPE if no player i has a strategy s'_i that differs from s_i^* in *just one* information set $P \in \mathbf{P}_i$ and does strictly better than s_i^* conditional on P being reached

8.6 Finitely repeated games

- Suppose that extensive game Γ^T consists of $T < \infty$ iterations of *exactly* the same normal form game $\langle N, \mathbf{S}, \mathbf{u} \rangle$ and players try to maximize their undiscounted sum of payoffs
- Knowing the NE of Γ , what can we say about SPE of Γ^T ?
- If stage game Γ has a *unique* NE \mathbf{s}^* then T -fold play of \mathbf{s}^* independently of the current history is Γ^T 's *unique* SPE
- If \mathbf{s}^* is *any* NE of stage game Γ , then T -fold play of \mathbf{s}^* independently of the current history is a SPE of Γ^T
- In case of multiple stage game NE, there may also exist other SPE which are history-dependent and involve play of a stage game NE only in an “end-game” phase

Infinitely repeated games

- Let Γ^∞ denote the *infinite* repetition of normal form game $\Gamma = \langle N, \mathbf{s}, \mathbf{u} \rangle$ in which players maximize their discounted sum of payoffs
(with common discount factor $\delta \in (0, 1)$)
- A payoff vector \mathbf{x} is called *strictly individually rational* iff for every player i , x_i strictly exceeds i 's *minmax payoff* M_i in Γ ,
i.e., the lowest payoff that players $-i$ can impose as punishment on a player i who correctly anticipates σ_{-i} and best-responds to it
- *Nash Folk Theorem / Perfect Folk Theorem:*
Let \mathbf{x} be feasible and strictly individually rational. Then, for δ sufficiently close to 1, there exists a *NE / SPE* of Γ^∞ with average payoff $\cong \mathbf{x}$.
(for games with $n > 2$ players, an additional technical condition related to reward opportunities has to be satisfied for the Perfect Folk Theorem)



UNIVERSITÄT
BAYREUTH

Advanced Microeconomics I

Summer Term 2025

Schedule for lectures

#	Date	Topic	Chs. in MWG
1	28.04.	Introduction	
2	05.05.	Preference and choice	1.A–D
3	12.05.	Consumer choice	2.A–F
4	19.05.	Classical demand theory	3.A–E, G
5	26.05.	Aggregate demand	3.I; 4.A–D
6	02.06.	Choice under risk	6.A–D, F
7	16.06.	Static games of complete information	7.A–E; 8.A–D, F
8	23.06.	Dynamic games of complete information	9.A–B; 12. App. A
9	30.06.	Games of incomplete information	8.E; 9.C
10	07.07.	Competitive markets	10.A–G
11	14.07.	Market power	12.A–F
12	21.07.	Question session for exam (→ 12.8.25)	

9. Games of incomplete information

- So far, we assumed that players have complete information about the game; in particular, every player knows
 - ▶ every other player's preferences (associated with their rationality)
 - ▶ every other player's strategy space
 - ▶ every other player's information partition
- What use are NE or SPE, which rest on correct beliefs about others' behavior in the game, when there is *incomplete information* on one of the above aspects, i.e., about which game is played?

9.1 Harsanyi's transformation

- John C. Harsanyi (1967/68) proposed a powerful framework for analyzing games of incomplete information
 - ① Introduce different *types* of each player:
 - ★ A particular type θ_i of player i is identified with a particular preference, strategy space and information partition
 - ★ Each player i knows his or her own type θ_i but possibly not that of other players
 - ② Introduce *Nature* as an additional player:
 - ★ *Nature* moves first and assigns each player i his or her type $\theta_i \in \Theta_i$
 - ★ Nature's move is a random draw from an exogenous and commonly known joint probability distribution ρ on $\Theta \equiv \Theta_1 \times \cdots \times \Theta_n$
 - ★ Each player i rationally updates the common prior ρ after learning θ_i
- Thus, a game of *incomplete information* is transformed into an (extensive) game with *complete (but imperfect) information*

Example

- Suppose a potential entrant and the incumbent monopolist simultaneously decide whether to enter and whether to boost capacity, respectively
- Cost of a capacity increase is *high* or *low*, and private information of the incumbent
- Profits are

Incumbent\Entrant	<i>enter</i>	<i>stay out</i>
<i>invest</i>	0,-1	2,0
<i>don't invest</i>	2,1	3,0

in case incumbent has *high* costs

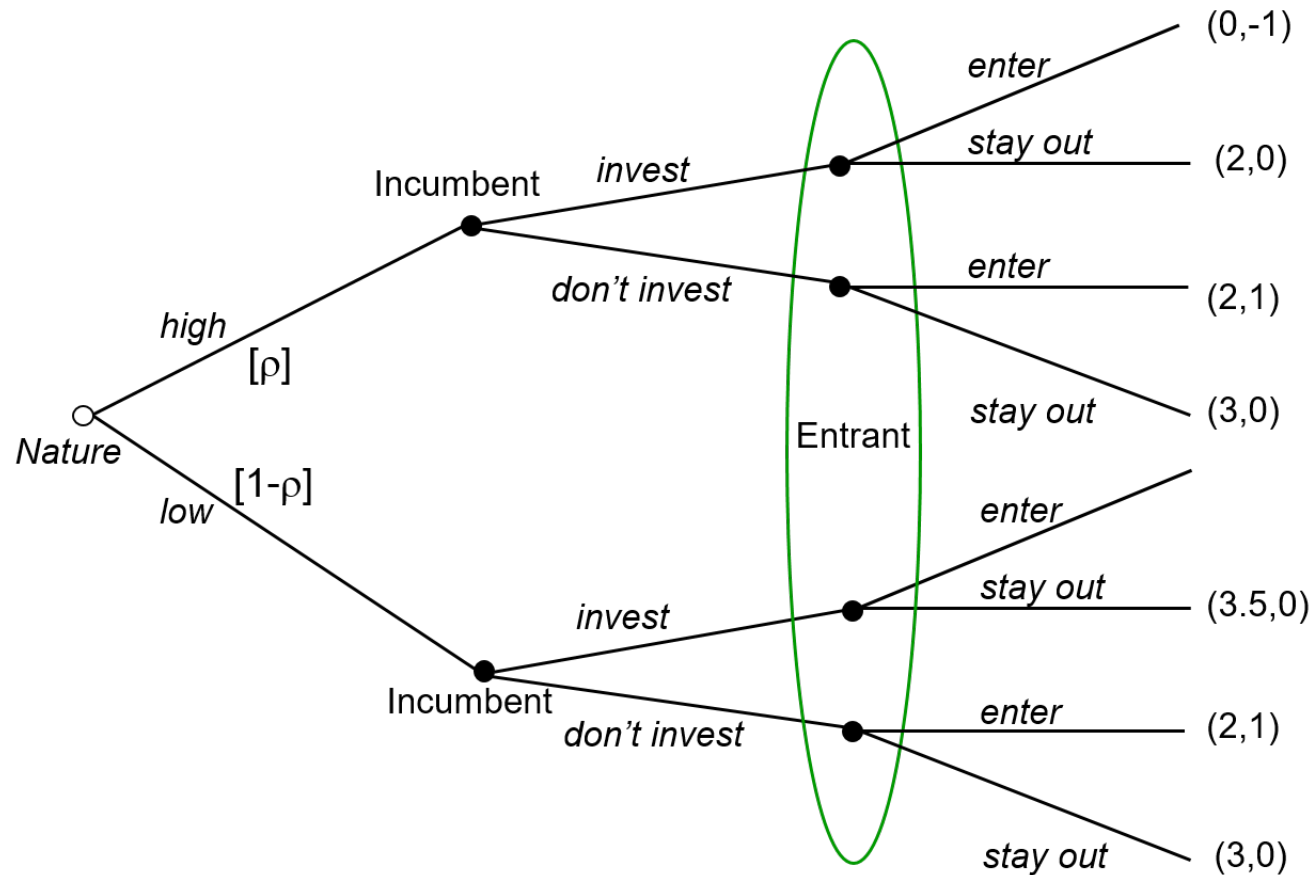
and

Incumbent\Entrant	<i>enter</i>	<i>stay out</i>
<i>invest</i>	1.5,-1	3.5,0
<i>don't invest</i>	2,1	3,0

in case incumbent has *low* costs

Example

- *Nature* 'selects' high costs with probability ρ , i.e., we obtain:



9.2 Bayesian games

- A *Bayesian (normal form) game* is a collection $\langle N, \Theta, \rho, \mathbf{A}, \mathbf{u} \rangle$ where
 - ▶ type space $\Theta \equiv \Theta_1 \times \dots \times \Theta_n$ specifies all possible types of players $i \in N$
 - ▶ actual types are drawn from joint probability distribution ρ on Θ
 - ▶ players' (pure) strategy sets S_i are *implicitly* defined as the set of all functions $s_i: \Theta_i \rightarrow A_i$ which map every possible type θ_i of player i to an action $s_i(\theta_i) \in A_i$ (elements of A_i are strategies in the original game of incomplete information)
 - ▶ u_i is defined on $\mathbf{A} \times \Theta_i$
 - We assume that $\langle N, \Theta, \rho, \mathbf{A}, \mathbf{u} \rangle$ is common knowledge
- ⇒ Rational players update the prior ρ using *Bayes' rule*:

$$Pr(A | B) = \frac{Pr(A \cap B)}{Pr(B)}$$

Best responses

- Comparing two actions $a_i, a'_i \in A_i$, player i with type θ_i will (in equilibrium: correctly) anticipate some *strategy* profile \mathbf{s}_{-i} but – in the spirit of players having incomplete information – must treat other players' *types* and hence actions as random variables
- So player i 's type θ_i compares

$$\mathbf{E}u_i(a_i, \mathbf{s}_{-i}, \theta_i) \equiv \sum_{\boldsymbol{\theta}_{-i} \in \Theta_{-i}} \rho(\boldsymbol{\theta}_{-i} | \theta_i) \cdot u_i(a_i, \mathbf{s}_{-i}(\boldsymbol{\theta}_{-i}), \theta_i)$$

to $\mathbf{E}u_i(a'_i, \mathbf{s}_{-i}, \theta_i)$

- If players use mixed strategies, then $u_i(a'_i, \mathbf{s}_{-i}(\boldsymbol{\theta}_{-i}), \theta_i)$ is simply replaced by expected payoff $u_i(a'_i, \boldsymbol{\sigma}_{-i}(\boldsymbol{\theta}_{-i}), \theta_i)$
- Strategy s_i^* of player i (in a Bayesian game) is a *best response* to \mathbf{s}_{-i} iff it specifies an optimal action $s_i^*(\theta_i) \in A_i$ for each type θ_i that player i might happen to be, i.e.,

$$\forall \theta_i \in \Theta_i : \forall a'_i \in A_i : \mathbf{E}u_i(s_i^*(\theta_i), \mathbf{s}_{-i}, \theta_i) \geq \mathbf{E}u_i(a'_i, \mathbf{s}_{-i}, \theta_i)$$

9.3 Bayesian Nash equilibrium

- A *Bayesian Nash equilibrium (BNE)* of the game $\langle N, \Theta, \rho, \mathbf{A}, \mathbf{u} \rangle$ is a strategy profile $\mathbf{s}^* = (s_1^*, \dots, s_n^*)$ such that for each player $i \in N$ the strategy s_i^* is a best response to \mathbf{s}_{-i}^* , i.e.,

$$\forall \theta_i \in \Theta_i : a_i = s_i^*(\theta_i) \in A_i \text{ maximizes } \mathbf{E}u_i(a_i, \mathbf{s}_{-i}, \theta_i)$$

(with expectation \mathbf{E} based on $\rho(\boldsymbol{\theta}_{-i} | \theta_i)$)

- A mixed-strategy BNE $\boldsymbol{\sigma}^*$ is defined analogously
- As in games of complete information, mixed strategy σ_i^* is a best response to $\boldsymbol{\sigma}_{-i}$ iff each action a_i played with a probability $\sigma_i(\theta_i)(a_i) > 0$ maximizes $\mathbf{E}u_i(a_i, \boldsymbol{\sigma}_{-i}, \theta_i)$

Example

- Again consider

$1_h/1_l \backslash \backslash 2$	<i>enter</i>	<i>stay out</i>
<i>invest</i>	0/1.5,-1	2/3.5,0
<i>don't invest</i>	2/2,1	3/3,0


with a specific probability $\rho \in [0, 1]$ for firm 1 having high costs

- Suppose $\rho = 0.5$; then
 - ▶ $\sigma^* = ((1_h \mapsto \text{don't invest}, 1_l \mapsto \text{don't invest}); \text{enter})$
and
 - ▶ every $\sigma^{**} = ((1_h \mapsto \text{don't invest}, 1_l \mapsto \text{invest}); (q, 1 - q))$ with $q \in [0, 1/2]$
- are BNE
(q refers to probability of *enter*)

9.4 Dynamic games of incomplete information

- Two complications arise when we apply the Harsanyi transformation to an extensive game of incomplete information:
 - ① If θ_i is private information, $-i$'s information sets are never singletons
 \Rightarrow there are no proper subgames started by $-i$'s moves

Example




9.4 Dynamic games of incomplete information

- Two complications arise when we apply the Harsanyi transformation to an extensive game of incomplete information:
 - ① If θ_i is private information, $-i$'s information sets are never singletons
 - \Rightarrow there are no proper subgames started by $-i$'s moves
 - \Rightarrow *subgame perfection* does not restrict $-i$'s moves off the NE path
 - \Rightarrow sequentially irrational behavior can survive (e.g., empty threats)
 - ② While $-i$'s beliefs about θ_i should be updated after any of i 's moves a_i^t , *Bayes' rule* only defines the conditional probability $\rho(\theta_i | a_i^t, \theta_{-i})$ after moves a_i^t which have positive probability under strategy profile σ^*

Example

- Consider the following game of complete but imperfect information (not even involving a move by *Nature*):



- $(ENTER, accommodate)$ and $(stay out, fight)$ are NE and SPE [because the game itself is its only subgame]
 - For the incumbent, *fight* is strictly dominated conditional on \bigcirc being reached; this makes SPE $(stay out, fight)$ rather implausible
- ⇒ We need a better formalization of sequential rationality than SPE

Strategies *and* beliefs

- More refined equilibrium concepts require optimal behavior in every “continuation game” starting in some information set, rather than only proper subgames
- For player i to be able to identify an *optimal* action in an arbitrary information set P^j at which she has the move she must
 - ▶ anticipate a particular (mixed) strategy σ_{-i} played by other players
 - ▶ have conditional beliefs $\mu_i(\cdot | P^j)$ about which decision node $n \in P^j$ she is in (= a probability distribution μ_i on P^j) given that P^j was reached
- The beliefs held by any player i and the equilibrium strategy profile σ^* depend on each other:
 - ▶ each player i 's strategy σ_i must maximize expected utility given μ_i
 - ▶ each belief μ_i must be consistent with prior ρ and the anticipated strategy σ_{-i}

9.5 Perfect Bayesian equilibrium

- A (weak) *Perfect Bayesian (Nash) equilibrium (PBE)* of the game $\Gamma = \langle N, \Theta, N, \prec, \iota, \{A(n)\}_{n \in N}, \{\alpha(n)\}_{n \in N}, \mathbf{P}, \{\rho(n)\}, \mathbf{u} \rangle$ is a combination (σ^*, μ^*) of a strategy profile $\sigma^* = (\sigma_1^*, \dots, \sigma_n^*)$ and a system of beliefs $\mu^* = (\mu_1^*, \dots, \mu_n^*)$ such that for each player $i \in N$

- ① strategy σ_i^* is “*sequentially rational*” in the sense that it prescribes a best response to σ_{-i}^* in any information set $P^j \in \mathbf{P}_i$ given the beliefs described by μ_i^* , i.e.,

$$\forall \theta_i \in \Theta_i : \forall P^j \in \mathbf{P}_i : \sigma_i^*(\theta_i) \in \Delta(A(P^j)) \text{ maximizes } \mathbf{E}u_i(\cdot, \sigma_{-i}^*, \theta_i | P^j)$$


[expectation \mathbf{E} is based on μ_i^* , and i chooses $\sigma_i^*(\theta_i)(a_i) > 0$ only if a_i max'es $\mathbf{E}u_i(\cdot)$]

- ② beliefs described by μ_i^* are *consistent with σ^** ,
i.e., they are derived from σ^* and Bayes' rule where that can be applied
[namely, in all information sets which have positive probability under σ^*]

- A combination of a strategy profile and a system of beliefs, (σ, μ) , is also called an *assessment* \rightarrow a PBE is a “sequentially rational” and consistent assessment

Example

- Consider




- When the incumbent's information set is reached, sequential rationality requires *accommodate* for *any* belief $(\mu, 1 - \mu)$ about the true history
 - Anticipating $\sigma_2^* = \text{accommodate}$, rationality requires $\sigma_1^* = \text{ENTER}$
 - Anticipating σ_1^* , incumbent must believe that *ENTER* was played with probability 1
- $\Rightarrow (\sigma^*, \mu^*)$ with $\sigma^* = (\text{ENTER}, \text{accommodate})$ and $\mu^* = 1$ is the unique PBE

Remarks

- If players use completely mixed strategies in a PBE, every information set is reached with positive probability and the system of beliefs is well-defined by Bayes' rule everywhere
- Otherwise, there is no restriction on conditional beliefs in information sets reached only after a deviation, i.e., the respective player i who has the move is free to interpret j 's deviation as, e.g., a fully informative indication of any particular action or type θ_j , or as not revealing any information, or ...

Problematic Example

- Consider



- $((\emptyset \mapsto \text{Out}, \text{In} \mapsto A), f), \mu = 1)$ is a PBE:
 - Anticipating that 1 will stay out, Bayes' rule doesn't restrict 2's beliefs for the zero-probability event that 2 has to make a move; 2 *may* think that 1 made another "mistake", so that $\mu = 1$
 - Based on $\mu = 1$, *fight* is indeed optimal for 2
 - If 1 anticipates that 2 would *fight*, it is best to choose *Out* and to *Accommodate* after involuntary entry
- This implausible beliefs-based PBE isn't even a SPE: (A, f) is no NE of the subgame following *In*

9.6 Sequential equilibrium

- Kreps and Wilson (1982) proposed to avoid complete arbitrariness of beliefs in information sets reached with probability zero by requiring existence of some fully mixed strategy profiles – which reach every information set with positive probability – that “justify” the beliefs in (σ^*, μ^*)
- A *sequential equilibrium (SE)* of the (mixed extension of) game Γ is an assessment (σ^*, μ^*)
 - ① which constitutes a perfect Bayesian equilibrium
 - ② for which a sequence $\{\sigma^k\}_{k=1,2,\dots}$ of completely mixed strategy profiles with $\sigma^k \rightarrow \sigma^*$ exists such that the sequence of beliefs implied by σ^k and Bayes’ rule, $\{\mu^k\}_{k=1,2,\dots}$, converges to μ^*

Remarks

- Every finite game has at least one SE; any SE is a PBE, but the reverse is not true
- In games in which only players' types are private information but all actions are observed, PBE and SE coincide
 - ▶ if each player has at most two possible types or
 - ▶ if the game has only two periods (e.g., simple signaling games)
- NB: The sequence $\{\sigma^k\}_{k=1,2,\dots}$ need not consist of equilibria; requiring that the (σ^k, μ^k) form PBEs in “perturbed games” Γ^k that require positive probability for all s_i leads to *(trembling-hand) perfect equilibria (PE)* in extensive games, which are a “refinement” of SE introduced by Selten (1975)



UNIVERSITÄT
BAYREUTH

Advanced Microeconomics I

Summer Term 2025

Schedule for lectures

#	Date	Topic	Chs. in MWG
1	28.04.	Introduction	
2	05.05.	Preference and choice	1.A–D
3	12.05.	Consumer choice	2.A–F
4	19.05.	Classical demand theory	3.A–E, G
5	26.05.	Aggregate demand	3.I; 4.A–D
6	02.06.	Choice under risk	6.A–D, F
7	16.06.	Static games of complete information	7.A–E; 8.A–D, F
8	23.06.	Dynamic games of complete information	9.A–B; 12. App. A
9	30.06.	Games of incomplete information	8.E; 9.C
10	07.07.	Competitive markets	10.A–G
11	14.07.	Market power	12.A–F
12	21.07.	Question session for exam (→ 12.8.25)	

10. Competitive markets

- In a *perfectly competitive economy*, every relevant good is traded, voluntarily and without transaction costs, by agents with no market power nor information asymmetries
- A *general competitive equilibrium* is an allocation and a price vector s.t.
 - ① all firms' production and factor demand plans maximize their respective profits,
 - ② all consumers' consumption and factor supply plans maximize their respective utility,
 - ③ these plans match, i.e., all markets clear
- Properties of competitive equilibria have fundamental importance:
 - ▶ Do market allocations satisfy “minimal quality standards” from a collective point of view?
 - ▶ How do competitive market interaction and social objectives relate?

Two requirements for market outcomes

- A first minimal requirement is that the allocations brought about by the market are *Pareto efficient*
- NB: Pareto efficiency doesn't involve any equitability concerns
- So, a second ambition is that specific normatively desired allocations somehow can be brought about by the market, too ...
- These issues are addressed for the economy as a whole by *general equilibrium theory*; we here restrict attention to a single market which constitutes a small part of the overall economy, i.e., *partial equilibrium*

10.1 Partial equilibrium competitive analysis

- Generally, a consumer's welfare depends on the optimal use of all her endowments (time, talents, goods, ...), and thus on *all* prices in the economy
 - We study a good k on which consumers spend only a small part of their budgets
- Then it is reasonable to ignore wealth effects and “general equilibrium effects”, e.g., of a tax on this good on the price of other goods, labor supply, wages, etc.

Partial equilibrium competitive analysis

- Fixed prices for all other goods and no wealth effects can most easily be captured by assuming quasilinear utility

$$u_i(x_i, m_i) = \phi_i(x_i) + m_i$$

for sufficiently rich consumers $i = 1, \dots, I$, where m_i captures i 's expenditure on “other goods” (treated as a composite *numeraire good*) and $\phi'_i > 0$ with $\phi''_i < 0$

- The price of the numeraire is usually normalized to equal 1;
the considered good k has price p

Optimization by firms

- Assuming that consumers have no initial endowment of good k , all consumption has to be produced by profit-maximizing firms
- Firm j 's transformation of numeraire into good k is captured by cost function $c_j(q_j)$; with $c'_j > 0$ and $c''_j \geq 0$, the necessary and sufficient condition for a solution to

$$\max_{q_j \geq 0} p^* \cdot q_j - c_j(q_j)$$

is

$$(I) \quad p^* \leq c'_j(q_j^*), \quad \text{with equality if } q_j^* > 0$$

Optimization by consumers

- Consumer i chooses consumption (x_i, m_i) to solve

$$\begin{aligned} \max_{x_i, m_i \geq 0} \quad & \phi_i(x_i) + m_i \\ \text{s.t.} \quad & m_i + p^* \cdot x_i \leq \omega_{mi} + \sum \theta_{ij} \cdot (p^* \cdot q_j - c_j(q_j)) \end{aligned}$$

[ω_{mi} is i 's endowment of the numeraire good, θ_{ij} is i 's share of firm j 's profits]

- Monotonicity of preferences implies that the budget is exhausted, and

$$\max_{x_i \geq 0} \phi_i(x_i) + \left[\omega_{mi} + \sum \theta_{ij} \cdot (p^* \cdot q_j - c_j(q_j)) \right] - p^* \cdot x_i$$

calls for

$$(II) \quad \phi'_i(x_i^*) \leq p^*, \text{ with equality if } x_i^* > 0$$

- x_i^* is unique if we assume that $\phi''_i(\cdot) < 0$

Competitive equilibrium

- Conditions

- (I) for all firms $j = 1, \dots, J$
- (II) for all consumers $i = 1, \dots, I$, and
- (III) $\sum x_i^* = \sum q_i^*$

define a *competitive equilibrium (CE)*

- For quasilinear preferences, sufficiently rich consumers' shares θ_{ij} in firms and initial numeraire endowments play no role in their optimal consumption and production decisions, (I) and (II), hence for p^*
- *Market supply* and *demand* for the good are defined by (I) and (II) for arbitrary p
- The inverse of the supply function, $q^{-1}(\cdot)$, can be viewed as the *industry marginal cost function* $C'(\cdot)$ [with each next unit produced by the most efficient firm]
- Inverse $P(x) = x^{-1}(x)$ of the demand function equals the *marginal social benefit* of the next unit of the good *if* quantity x is distributed efficiently amongst consumers

10.2 Fundamental Welfare Theorems

- For any given consumption and production plans, \mathbf{x} and \mathbf{q} , and (sufficient) total endowments $\omega_m = \sum_i \omega_{mi}$ of the numeraire, any utility vector in set

$$\left\{ (u_1, \dots, u_I) \in \mathbf{R}^I \mid \sum u_i \leq \sum \phi_i(x_i) + \omega_m - \sum c_j(q_j) \right\}$$

could be realized by appropriate transfers of the numeraire in this quasilinear case
[because the numeraire has the same constant marginal utility for everyone]

- For given \mathbf{x} and \mathbf{q} , the RHS above is a constant, so the boundary of this utility possibility set is a hyperplane with normal vector $(1, 1, \dots, 1)$; variations of \mathbf{x} and \mathbf{q} imply parallel shifts of it

Pareto optimal plans

- Plans \mathbf{x}^* and \mathbf{q}^* are Pareto-optimal iff they maximize the RHS, i.e., they solve

$$\begin{aligned} \max_{\mathbf{x}, \mathbf{q} \geq 0} \quad & \sum \phi_i(x_i) + \omega_m - \sum c_j(q_j) \\ \text{s.t.} \quad & \sum x_i - \sum q_j = 0 \end{aligned}$$

- Given our convexity assumptions ($c_j'' \geq 0, \phi_i'' \leq 0$), the maximization of the Lagrangean

$$\mathcal{L}(x_1, \dots, x_I, q_1, \dots, q_J, \lambda) = \sum \phi_i(x_i) - \sum c_j(q_j) - \lambda \cdot (\sum x_i - \sum q_j)$$

yields the necessary and sufficient conditions ($j = 1, \dots, J; i = 1, \dots, I$):

- (i) $-c_j'(q_j^*) + \lambda \leq 0 \Leftrightarrow \lambda \leq c_j'(q_j^*)$, with equality for $q_j^* > 0$
- (ii) $\phi_i'(x_i^*) - \lambda \leq 0 \Leftrightarrow \phi_i'(x_i^*) \leq \lambda$, with equality for $x_i^* > 0$
- (iii) $\sum x_i^* = \sum q_j^*$

- These correspond exactly to the conditions (I)–(III) that characterize a competitive equilibrium, with λ replacing p^*

First Fundamental Welfare Theorem

- Hence, if price p^* and allocation $(x_1^*, \dots, x_I^*, q_1^*, \dots, q_J^*)$ constitute a CE, then this allocation is Pareto optimal
- This result is also known as the *First Fundamental Theorem of Welfare Economics*
- Good k 's price p^* in a CE exactly reflects the good's marginal social value (in units of the numeraire), i.e., the “shadow price” of the resource constraint:
 - ▶ each firm, in its resp. profit maximization, equates own marginal production cost to the marginal social value of its output
 - ▶ each consumer consumes up to the point where own marginal utility equals marginal cost of production (in units of the numeraire)
- The theorem vindicates Adam Smith's “invisible hand” for perfectly competitive markets, and holds more generally than considered here

Remarks

- Imperfections such as market power, asymmetric information, or market incompleteness often lead to less cheerful conclusions ...
- Nothing is said yet about actual *existence* of a CE, or how it might be reached (if at all) by a dynamic adaptation or *tâtonnement* process with decentralized information ...
- In the quasilinear case, CE price p^* and individually consumed and produced quantities of good k do not depend on the distribution of total endowment ω_m
[NB: except for corner solutions, in which some agents are too poor to consume both good k and the numeraire]

Second Fundamental Welfare Theorem

- So, ignoring corner solutions, changing the initial distribution $(\omega_{m^1}, \dots, \omega_{m^I})$ changes individual consumption of the numeraire but not $(x_1^*, \dots, x_I^*, q_1^*, \dots, q_J^*)$:
one moves within the Pareto efficient hyperplane
- For *any* Pareto optimal levels of utility (u_1^*, \dots, u_I^*) , there are transfers (T_1, \dots, T_I) of the numeraire good with $\sum T_i = 0$ such that a CE reached from the redistributed endowments $(\omega_{m^1} + T_1, \dots, \omega_{m^I} + T_I)$ yields exactly the utilities (u_1^*, \dots, u_I^*)
- This result is also known as the *Second Fundamental Theorem of Welfare Economics*
- Hence, pursuing a particular distributional goal does *not* conflict with having competitive markets: one can achieve the goal by appropriate endowment transfers and then “let the market work”
- This result generalizes, too, but not as much as the First Theorem
[in particular, preferences and technology should rather be convex]

10.3 Welfare Analysis in Partial Equilibrium

- What “yardstick” can we use for comparing different allocations (esp. Pareto-incomparable ones)?
- The value of $\sum \phi_i(x_i) - \sum c_j(q_j)$ in the maximization problem which characterizes Pareto efficient allocations is known as the *(Marshallian) aggregate surplus*
- It is an indicator of social welfare under *any* (increasing) social welfare function $W(u_1, \dots, u_I)$ in the quasilinear case:
 - ▶ greater surplus implies a larger utility possibility set
 - ▶ the planner can select a utility vector with a greater (maximized) W -value through appropriate endowment transfers
- Aggregate surplus can be derived very simply from market demand and supply functions; it is thus a convenient tool and used in many applications

Aggregate surplus and CE

- Start from (possibly non-CE) total consumption and production $x = \sum x_i = \sum q_j = q$
- Increases by $(\Delta x_1, \dots, \Delta x_I)$ and $(\Delta q_1, \dots, \Delta q_J)$ s.t. $\sum \Delta x_i = \sum \Delta q_j \equiv \Delta x > 0$
change surplus by

$$\Delta S \approx \sum \phi'_i(x_i) \cdot \Delta x_i - \sum c'_j(q_j) \cdot \Delta q_j$$

- For given x , the planner maximizes surplus by allocating consumption and production s.t. $\phi'_i(x_i) = P(x)$ and $c'_j(q_j) = C'(x)$ for all i, j
- Then $\Delta S \approx [P(x) - C'(x)] \cdot \Delta x$ or $dS/dx = P(x) - C'(x)$ for marginal changes
- So aggregate surplus under an optimal distribution of output x is

$$S(x) = S(0) + \int_0^x [P(s) - C'(s)] ds$$

$S(0)$ reflects fixed costs; $S(x) - S(0)$ is the area between demand and supply curves

- $S(x)$ increases up to x^* s.t. $P(x^*) = C'(x^*)$, i.e., the CE level
- ⇒ Surplus is maximal in the undistorted laissez-faire CE
[but given one distortion, adding another *may raise* surplus...]



UNIVERSITÄT
BAYREUTH

Advanced Microeconomics I

Summer Term 2025

Please mail your questions for Lecture #12/Q&A to vw14@uni-bayreuth.de before 18.7.25, 8:30h, with "Subject: Adv. Micro 1 Q&A"

Schedule for lectures

#	Date	Topic	Chs. in MWG
1	28.04.	Introduction	
2	05.05.	Preference and choice	1.A–D
3	12.05.	Consumer choice	2.A–F
4	19.05.	Classical demand theory	3.A–E, G
5	26.05.	Aggregate demand	3.I; 4.A–D
6	02.06.	Choice under risk	6.A–D, F
7	16.06.	Static games of complete information	7.A–E; 8.A–D, F
8	23.06.	Dynamic games of complete information	9.A–B; 12. App. A
9	30.06.	Games of incomplete information	8.E; 9.C
10	07.07.	Competitive markets	10.A–G
11	14.07.	Market power	12.A–F
12	21.07.	Question session for exam (→ 12.8.25)	

11. Market power

- Price-taking behavior is implausible if there are only a few producers (or consumers)
- Several “workhorse” models of *industrial organization* capture the performance differences that market power can cause

11.1 Monopoly

- A first benchmark is an uncontested *monopolist* who can
 - ▶ produce quantity x of a good at cost $C(x)$, and
 - ▶ sell it at a constant unit price p to consumers, whose demand is described by demand function $D(p)$
- The monopolist maximizes $\Pi(p) = p \cdot D(p) - C(D(p))$
- The necessary condition for an interior profit maximum is

$$\begin{aligned} [p - C'(D(p))] \cdot D'(p) &= -D(p) \\ \Leftrightarrow \frac{[p - C'(D(p))]}{p} &= \frac{-D(p)}{[D'(p) \cdot p]} = \frac{1}{|\varepsilon|} \end{aligned}$$

\Rightarrow In the monopolist's profit maximum, the price-cost margin $[p^m - C']/p^m$ (also known as *Lerner index*) equals the inverse of the (absolute) price elasticity $|\varepsilon| = -D'(p^m) \cdot p^m / D(p^m)$

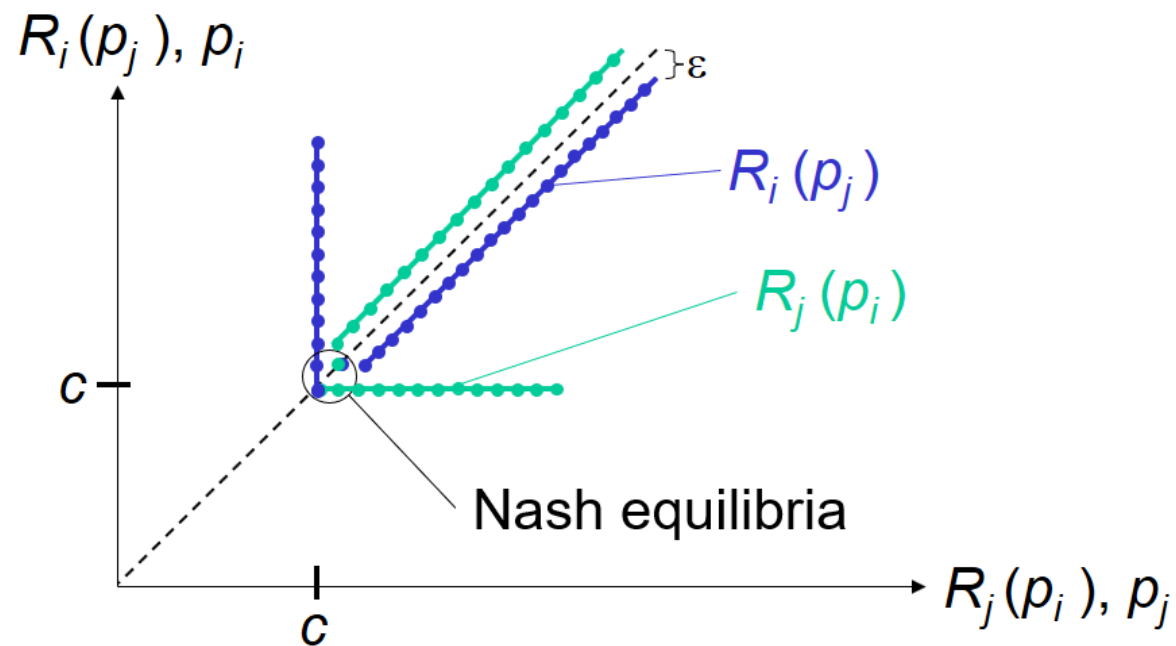
Deadweight loss of monopoly

- Except for perfectly elastic demand ($|\varepsilon| = \infty$), $p^m > C'(D(p^m))$ and quantity $x^m = D(p^m)$ is smaller than x^* in the CE
- A quantity $x < x^*$ results in an inefficient allocation and entails a *deadweight (welfare) loss*: surplus which could be generated by further trade is left unrealized
 - ▶ Having sold $D(p^m)$ units at price p^m , the monopolist would gain from selling *additional* units at any price p with $C'(D(p)) < p < p^m$
 - ▶ All consumers with willingness to pay v satisfying $p < v < p^m$ would gain from buying these additional units

11.2 Bertrand competition

- The *Bertrand duopoly model* considers
 - ▶ two firms $i \in \{1, 2\}$ that simultaneously announce their respective price p_i for a homogenous good, which – in the baseline case – can be produced at identical constant marginal cost c without capacity constraints, and
 - ▶ consumers that buy only at the cheaper firm if $p_1 \neq p_2$, and otherwise split demand $D(p)$ with $D'(p) < 0$ equally between firms 1 and 2
- If prices are discrete (e.g., multiples of a currency unit ε), firm i 's best responses $R_i(\cdot)$ to prices $p_j \in \{c, c + \varepsilon, \dots, p^m - \varepsilon, p^m\}$ are
 - ▶ $R_i(p_j) = \{p_j - \varepsilon\}$ for $p_j > c + \varepsilon$
 - ▶ $R_i(p_j) = \{c + \varepsilon\}$ for $p_j = c + \varepsilon$
 - ▶ $R_i(p_j) = \{p_i : p_i \geq c\}$ for $p_j = c$

NE in the discrete Bertrand game



- The discrete Bertrand game has two NE:
 - ▶ $(p_1^*, p_2^*) = (c, c)$
 - ▶ $(p_1^{**}, p_2^{**}) = (c + \varepsilon, c + \varepsilon)$

Bertrand paradox

- If p_1 and p_2 may be chosen from $[0, \infty)$, then $(p_1^b, p_2^b) = (c, c)$ is the unique NE

→ The “*Bertrand Paradox*”:

Price competition between two symmetric firms with CRS results in the same market outcome as perfect competition, namely $p^* = c$

- Asymmetric case:
 - ▶ If firm j has a *non-drastic* cost advantage over its competitors, it supplies the entire market at price $p_j^b = \min_{k \neq j} c_k$ (or ε below)
 - ▶ For a *drastic* advantage, it chooses $p_j^b = p_j^m < \min_{k \neq j} c_k$
- Even symmetric firms can avoid the paradox
 - ▶ if technology commits them not to undercut their rival for some $p > c$ (e.g., for capacity constraints)
 - ▶ if they differentiate their products, i.e., make them imperfect substitutes
 - ▶ if they collude

11.3 Edgeworth competition

- Consider price competition with *exogenous capacities* $q_1, q_2 < D(c)$, i.e., a single firm cannot serve the whole market at $p^* = c$
- If firm i 's capacity q_i is already exhausted for $p_i = p_j$, it will not undercut firm j
- If capacities q_1 and q_2 are “small” (namely, $\leq x_i^c$ given by Cournot NE), equilibrium prices $p_1^e = p_2^e = p^e$ are defined by $D(p^e) = q_1 + q_2$:
 - ▶ Unilateral undercutting of p^e is unprofitable because the firm's capacity is already exhausted
 - ▶ A unilateral increase of p^e (i.e., selling below capacity) is unprofitable if outputs are “small” and profit margins high already

11.4 Cournot competition

- The *Cournot duopoly model* considers two firms $j \in \{1, 2\}$ that
 - ▶ simultaneously produce a respective output x_j of a homogenous good at cost $C_j(x_j)$, and
 - ▶ sell at market clearing price $p = P(x_1 + x_2)$, i.e., such that $D(p) = x_1 + x_2$
- The Cournot game can be interpreted as the reduced form of a two-stage extensive game in which
 - ▶ first, firms invest in capacities x_j , incurring costs $C_j(x_j)$ for this
 - ▶ second, they engage in Edgeworth competition with fixed capacities $q_j = x_j$ and zero costs of production
- We assume that no firm has a *drastic cost advantage*: costs are sufficiently similar that both firms want to produce in equilibrium

Reaction function of firm i

- Firm i maximizes

$$\Pi_i(x_i, x_j) = P(x_i + x_j) \cdot x_i - C_i(x_i)$$


- Best response $x_i^* = R_i(x_j)$ to the anticipated competitor output x_j is defined by

$$P(x_i + x_j) + P'(x_i + x_j) \cdot x_i = C'_i(x_i)$$

- For $x_j = 0$, i should behave like a monopolist, i.e., $R_i(0) = x_i^m$
- If the competitor already produces the CE quantity $x_j = x^*$, it is optimal not to produce, i.e., $R_i(x^*) = 0$
- Under standard assumptions – i.e., $P''(x) \leq 0$ and $C''_i(x_i) \geq 0$ – the reaction function $R_i(x_j)$ is strictly decreasing on $(0, x^*)$
- This means firms' quantity decisions are *strategic substitutes*:
firm i reacts to a larger output x_j with a reduction of own output x_i

Nash equilibrium in the Cournot game

Symmetric case:



n -firm Cournot game

- With $x_\Sigma = \sum x_i$, the NE $\mathbf{x}^c = (x_1^c, \dots, x_n^c)$ of Cournot competition between n firms is characterized by:

$$P(x_\Sigma) + P'(x_\Sigma) \cdot x_i = C'_i(x_i) \text{ for } i \in \{1, \dots, n\}$$

or, expressed in market shares $s_i = x_i/x_\Sigma$ and with $p^c = P(x_\Sigma^c)$,

$$\frac{[p^c - C'_i(x_i^c)]}{p^c} = \frac{s_i}{|\varepsilon(p^c)|}$$

- So, in the Cournot NE, the Lerner index (\approx profitability, market power) of firm i is proportional to its *market share* s_i ;
unequal market shares derive from *technology* differences
- For symmetric firms: $s_i = 1/n$ and mark-up ratio $1/[n \cdot |\varepsilon(p^c)|]$ approaches 0 as $n \rightarrow \infty$

11.5 Product differentiation

- The *Hotelling duopoly model* describes “horizontal” product differentiation with
 - ▶ a continuum of consumers who want to buy at most one unit of a differentiated good regarding which they have uniformly distributed ideal points in a *one-dimensional product space* $X = [0, 1]$, and
 - ▶ firms 1 and 2 who are – in the baseline case – located at the extremes of X , and simultaneously announce prices p_1 and p_2 for the good produced at constant marginal cost c
- Let consumers suffer from quadratic disutility of distance, tx^2 or $t(1 - x)^2$ for $t > 0$, and each have sufficiently high valuation for one unit of the good
⇒ each one buys from the firm for which price plus “transportation cost” is minimal

Hotelling model with fixed locations

- The consumer at location x buys from 1 if $p_1 + tx^2 \leq p_2 + t(1 - x)^2$, otherwise from 2
- ⇒ Firm 1 faces demand $D_1(p_1, p_2) = (p_2 - p_1 + t)/2t$,
firm 2 faces $D_2(p_1, p_2) = 1 - D_1(p_1, p_2)$
- Maximization of $\Pi_i(p_1, p_2) = (p_i - c) \cdot D_i(p_1, p_2)$ yields reaction functions
 $R_i(p_j) = 1/2 \cdot (p_j + c + t)$
(NB: firms' prices are strategic complements)
- ⇒ Nash equilibrium: $p_1^* = p_2^* = c + t$
- Profits $\Pi_i(p_1^*, p_2^*) = t/2$ are positive;
they increase in differentiation parameter $t > 0$

11.6 Collusive behavior

- *Collusion* refers to anti-competitive coordination of firms' prices, quantities, etc. in markets where cartel agreements cannot be enforced in court
- Firms' always have an interest in full coordination: they *could duplicate* the non-cooperative outcome; not doing so reveals that they strictly increase profits ...
- Such coordination is, however, not self-enforcing if firms interact only once, or over a definite time-horizon
- If firms interact repeatedly over an in(de)finite time horizon, collusion can be supported by strategies that involve credible punishment of free-riding deviators (provided that a deviator's forgone long-term collusion rents are important enough relative to short-term gains from deviation → #8: *Folk Theorems*)

Collusion in the symmetric CRS Bertrand oligopoly

- In the symmetric Bertrand n -firm oligopoly with CRS, a firm's per-period profit is

$$\Pi^* \approx 0 \quad \text{if all firms compete,}$$

$$\Pi^k = \frac{\Pi^m}{n} \quad \text{if all firms collude, and}$$

$$\Pi^d \approx \Pi^m \quad \text{if the firm deviates}$$

- Collusion can be realized by an SPE in *Nash reversion strategies* iff firms discount future profits by a factor δ that is no smaller than the *critical discount factor*

$$\delta_{crit}^b = \frac{\Pi^d - \Pi^k}{\Pi^d - \Pi^*} = \frac{(n-1)}{n}$$

- The critical discount factor increases in n , and converges to 1