Course outline

Preparatory course in Microeconomics, 2024-2025

Master of Science in European Economy and Business Law

Lorenzo Bozzoli*

6 agosto 2024

This document is a tentative outline of this year's lectures. The final version and the problem sets of each lectures will be posted on the Teams channel of the course on a daily basis.

Schedule: 09/09/2023 to 13/09/2023, Monday to Friday, 09.30am-1.00 pm.

Outline and objectives: We will review the axiomatic foundations of microeconomics; the theory of consumption and production; partial equilibrium with competitive markets and the theory of monopoly.

Some extra topics that we will try to cover are: choice under uncertainty and externalities.

Topics (detail):

1. **Introduction.** Definition of microeconomics; axiomatic foundations of decision theory; understanding the utility function.
2. **Theory of consumption, I.** The budget set, the indifference curves and the marginal rate of substitution.
3. **Theory of consumption, II.** Constrained optimization with example utility functions;
4. **Theory of consumption, III.** Understanding the demand function; computing demand elasticity and the consumer surplus.
5. **Theory of consumption, extra I.** Hints to choice under uncertainty, risk-aversion and insurance.
6. **Theory of consumption, extra II.** Hints to social welfare: Pareto-optimality and equity-efficiency tradeoffs.
7. **Theory of production, 1.** The firm as a production technology; isoquants, isocosts; profit-maximization problem and optimal factor demand.
8. **Theory of production, 2.** Cost-minimization problem and the geometry of costs; the supply function; fixed and variable costs.
9. **Market equilibrium, 1.** Competitive equilibrium; efficiency of partial equilibrium; if there is time, the effect of a price cap and the deadweight loss.
10. **Market equilibrium, 2.** Monopolistic equilibrium; surplus comparison with competition; if there is

*e-mail: lor.bozzoli@gmail.com

References: Varian, H.R. (2010), *Intermediate Microeconomics: a modern approach*, 8th edition, WW Norton & Company.

The CORE team, The Economy. Available at: `https://www.core-econ.org`.

1 Introduction.

Economics and its methodology. An **economic model** is a stylized representation of human behavior in an economic context. The science of economics makes use of economic models to capture the salient characteristics of social phenomena and to address normative (*should we have a minimum wage?*), interpretative (*what is the driver of inflation?*) and predictive questions (*what will be the impact of Covid-19 on labor force participation?*) about them.

The original object of economics is the study of *economic systems*: the ways in which *goods and services* (things or activities that are in some way **scarce**, and that serve to *satisfy human needs*) are *produced, allocated and consumed.* Modern contributions tend to have a broader interpretation of the discipline, applying its methodology to any situation in which *rational people take choices* (from law, institutions and politics to history, psychology and evolutionary theory).

Because of the difficulty of reducing a great portion (and maybe all) of human behavior to tractable terms, the foundations of the economic discipline are still debated.

Mainstream (so-called *Neoclassical*) contributions, despite being very different in nature, have set up a more or less consistent standard account of human nature, sometimes interpreted as a set of *meta-axioms*:

- **Mathematical formalism**: a rigorous abstract and mathematical language should be adopted to describe human behavior.
- **Methodological individualism**: economic analysis should proceed *from individual behavior* to larger scale phenomena.
- **Instrumental rationality (optimization principle)**: decision-makers act as close as possible to their individual objectives given the alternative at their disposal.
- **Equilibrium principle**: economic systems are assumed to have a state of *equilibrium*, at which they exhibit some form of order and predictability despite the lack of explicit coordination. Describing what an equilibrium would *look like*, is more interesting than understanding *how* implicit coordination happens.

It must be noted that those axioms are often criticized as being: simplistic; far-fetched in their epistemic optimism; too similar to hard science for a qualitative subject; biased towards a neoliberal viewpoint. Nonetheless, even if you are skeptical about its foundations, you should study Neoclassical economics to understand the driver of its success and formulate the right set of objections.

Microeconomics. Microeconomics is a mathematical subject. It has the objective to provide rigorous foundations on *how rational people take decisions*, and *how their individual decisions interact* and affect socioeconomic systems.

Not surprisingly, microeconomics is the language of Neoclassical economic theory, as it provides a general, consistent and tractable formalization of concepts such as rationality, individual choice, decentralized coordination and economic efficiency.

Many familiar ideas that are central in the political and economic debate are the product of microeconomic theory. Some notable commonplace (and polarizing) statements that are derived from microeconomics considerations are:

- When price goes up, demand goes down and supply goes up, and vice versa.
- The invisible hand of the free market leads to the common good despite people acting selfishly.

- Extra-profits and rents are levelled out in the long run when markets are truly competitive.

Statement as such are difficult to qualify as true or false, as they summon centuries of unsettled scholarly debate. In this course, we will review the basic microeconomics that underlies this apparently mundane ideas, and provide some tools to start exploring the complexity of this subject.

In this process, you will acquire the basic notions to approach economic theoretical subjects such as Game Theory and Industrial Organization.

Objectives. In the **first week** of the course, we will cover the foundation of **decision theory** and the **theory of consumption**, and learn how to construct the **market demand curve** from the aggregation of individual decision problems.

In the **second week** of the course, we will explore the neoclassical **theory of the firm** as a black box technology that transforms inputs into outputs under a profit-maximization objective. We will construct the market **short-run and long-run supply curve** and learn to describe market outcomes in the **partial equilibrium framework**. We will introduce some basic notions of **welfare analysis** and **compare competitive markets to monopolistic markets**.

Some additional topics that we will try to cover, time permitting, are the basic notions of *decision theory under uncertainty*, the basic ideas of the *equity-efficiency trade-off*, and possibly some hints to the *general equilibrium* framework.

1.1 Foundations of decision theory

The elementary ingredients of a microeconomic model are the individual **decision problems** facing the decision-makers that populate the economy. The essential objects of a *decision problem*, defined and studied in the field of **decision theory**, are:

- A decision-maker, which is called an **agent** and denoted with the index i.
- A **feasible set** X of alternative choices $x \in X$.
- A **preference relation** P_i, namely an ordering among the alternatives according to the agent's objectives or tastes.
- A subset $B \subset X$ of alternatives that are *effectively* available to the agent, given some *economic* constraints in addition to physical feasibility.
- A very simple behavorial rule, used to make predictions: *the agent always selects her preferred alternatives among those available.* We call x^* the **choice function or correspondence** i.e. the mapping that associates to each subset $B \subseteq X$ the preferred choices of the agent within B, according to the ranking P_i.

This is a very general framework, that can be applied beyond markets. We now introduce some examples that do not involve a market economy: this should help you to familiarize with abstract economic thinking, a basic requirement for your future game theoretical studies. The next lectures instead will be devoted to the analysis of markets, providing the basic ingredients for applied subjects like Industrial Organization.

Feasible sets. Any set X can be the **decision space** of a microeconomic model. The existence of X relies on the assumption that the agent knows about X, is able to tell apart the elements $x \in X$ standing for any feasible alternative, and can select any $x \in X$ to take place, if no external social/institutional constraints are imposed on her behavior.

- E1. Suppose you are having breakfast and you can choose between juice (J), coffee (C) or tea (T). In this case, the feasible set is $X = \{0; J; C; T\}$, where 0 stands for no breakfast.

- E2. Suppose you can drink more than one beverage for breakfast. In this case, the feasible set is

$$X = \{0; J; C; T; (J, C); (J, T); (C, T); (C, J, T)\}.$$

- E3. Suppose you are a student-worker and you need to allocate your time awake, say $H = 15$ hours, between studying (s), working (w) and playing (p). The feasible set is then

$$X = \{(s, w, p) \in \mathbb{R}^3 : s \geq 0; w \geq 0; p \geq 0; s + w + p = 15\}.$$

Additional constraints. Some of the alternatives in X may not be available to the agent due to social and institutional factors out of her control. The set of effectively available alternatives is $B \subset X$: the agent is able to recognize the existence of such constraints and to tell apart which alternatives belong to B and which don't.

- E1. Assume that no additional constraint exists. In this case, $B = X$ and then $B = \{J; C; T\}$.

- E2. Suppose now that drinking three different things is forbidden, as a rule between you and your roommate. In this case

$$B = \{0; J; C; T; (J, C); (J, T); (C, T)\},$$

namely $B \subset X$; in particular $(J, C, T) \in X$ but $(J, C, T) \notin B$, as it physically feasible but subject to a social constraint.

- E3. Suppose you are required to attend 3 hours of lectures per day: in this case, an additional constraint is $s \geq 3$. Then:

$$B = \{(s, w, p) \in X : s \geq 3\}.$$

Preference relationship. A preference relationship P_i is a **complete ordering** over the set of feasible alternatives X. Saying that the agent i is endowed with a preference relationship P_i means that, for any pair of alternatives $(x, y) \in X$, the agent is able to say whether she prefers x, y or is indifferent.

Let us introduce some notation in this respect.

We write $x \succ_i y$ to say that i **strictly prefers** x **to** y: this means that, in ranking her available alternatives, the agent unambiguously attributes a higher position to x. We write that $x \sim_i y$ to say that i is **indifferent between** x **and** y: that is, in a ranking among the elements in X, they occupy the same position. We write $x \succcurlyeq_i y$ to say that i **weakly prefers** x **to** y: such a statement excludes the case that $y \succ_i x$, but not that $x \sim_i y$.

- E1. Say that juice is better than coffee, and coffee is as good as latte. Namely, $J \succ_i C$, $C \sim_i L$. Also, $J \succcurlyeq_i C$ is a true statement.

- E2. Say that you strictly prefer coffee with juice than coffee or juice alone, but that coffee and latte together are worse than coffee or latte alone. Then, $(J, C) \succ_i C$ and $(J, C) \succ_i C$ are true. Also, $C \succ_i (C, L)$ and $L \succ_i (C, L)$.

- E3. Say that you prefer to allocate equal hours to each of the three activity, rather than doing the same activity all day long. Then, for example, we can say that $(s = 5, w = 5, p = 5) \succ_i (s = 15, p = 0, w = 0)$.

Axiomatic definition of rationality. Not any preference ordering is considered rational. Say for example that, in the example E1, the agent strictly prefers latte over juice. That would be puzzling, since preferences would be circular: $J \succ_i C$ and $C \sim_i L$, but $L \succ_i J$.

With circular preferences, the agent's optimal choice from a feasible set $\{C; J; L\}$ is ambiguous.

A basic notion of **individual rationality**, ruling out unpredictable behavior, is captured by the following three axioms:

- **Completeness**: the agent is able to compare any couple of alternatives. In symbols, if $x \in X$ and $y \in X$, either $x \succcurlyeq_i y$ or $y \succcurlyeq_i x$.
- **Reflexivity**: for any $x \in X$, $x \succcurlyeq_i x$, namely, any alternative is at least as good as itself.
- **Transitivity**: Say that $x \in X$, $y \in X$ and $z \in X$. If $x \succcurlyeq_i y$ and $y \succcurlyeq_i z$, then $x \succcurlyeq_i z$: *cycles* of preferences are ruled out.

Optimal choice. Given a set of alternatives $B \subset X$ at her effective disposal, it is **rational** for the agent to select any alternative that is **optimal**, namely any alternative $x \in B$ such that there is *no other* $y \in B$ *strictly preferred to* x:

$$x^*(B) = \{x \in B : x \succcurlyeq_i y \quad \forall\, y \in B\}.$$

In other words the expression $x^*(B)$ denotes the subset of elements in B that are not strictly worse than anything else available in B itself.

Exercise.

- E1. What is $x^*(B)$, according to the information at your disposal?
- E2. What is your guess for $x^*(B)$?

1.2 The utility function

Given an agent i and a feasible set X, a **utility function** is a function $U_i : X \to \mathbb{R}$ that attaches a utility score, which can be any real number, to any feasible item in X.

We say that the function U_i *represents* the preference ordering P_i if the following relationships holds for any (x, y) in X:

- $U_i(x) > U_i(y)$ if and only if $x \succ_i y$.
- $U_i(x) = U_i(y)$ if and only if $x \sim_i y$

Then, if we rank the alternatives $x \in X$ according to the value assumed by $U_i(x)$, we obtain the same ordering as the one implied by P_i.

Each U_i is totally identified by the associated ordering among choices; different utility functions U_i and V_i keeping the same ordering over X are equivalent: they induce an identical *choice* $x^*(B)$ for any possible set B of available alternatives.

Utility functions are thus **invariant over monotonic transformation**: any algebraic manipulation $U_i \to V_i$ that leaves the ordering of $V_i(x)$ over X identical to those of U_i implies an identical description of the agent's behavior. This is true for example when, for some $\alpha \in \mathbb{R}, \beta > 0$:

$$V_i(x) = \alpha + \beta U_i(x) \quad \forall x \in X,$$

namely for linear transformations, but also when

$$V_i(x) = 6 + 12 \cdot [U_i(x)]^3$$

For this reason, a utility function is said to be ***ordinal* rather than *cardinal***: the specific numbers attached to each decision have no meaning, what matters is the *order* relation that they describe among feasible options.

Note that, having defined U_i from P_i, the following definition of $x^*(B)$ can be put forward for each $B \subset X$:

$$x^*(B) = \text{argmax}_{x \in B}\, U_i(x)$$

Exercise.

- Propose a utility function to represent P_i in the examples E1 and E2.
- What is $x^*(B)$ if $U(s, w, p) = p\,w\,s$? What is $x^*(B)$ if $U(s, w, p) = 3\,p\,w\,s + 2$?

2 Theory of consumption, I.

Today we introduce a much more structured class of decisions problem, which posits the existence of **commodities** and **markets**: the Neoclassical **theory of consumption**.

The theory of consumption studies how a rational agent optimally allocates her income to purchase various commodities in a market system. The framework is obtained by adding specific assumptions to the basic model of *decision theory* studied yesterday.

2.1 Feasible set.

In the baseline model of the theory of consumption, it is assumed that a finite number M of different goods are available. Goods are indexed with $m = 1, \dots, M$.

So, for example, if there only two goods, say *apples* and *oranges*, $M = 2$, where $m = 1$ is the index of apples, and $m = 2$ of oranges.

The consumer's choice consists in combining quantities of different goods under a set of economic constraints. A *feasible* level of consumption of each good is represented by a positive real number. Goods are assumed to be perfectly separable, meaning for example that you can consume a grand total 1.5, π or $\sqrt{2}$ apples.

Denote by $x_m \in \mathbb{R}_+$ a feasible level of consumption of the good m. A vector $x = (x_1, \dots, x_M)$ is called a **bundle** of goods: it specifies a level of consumption for each existing good in the market. A *bundle* $x \in X$ is a possible basket of commodities that the agent can buy from the market; any $x \in X$ fully describes a possible consumption plan of the agent.

The set of all feasible consumption bundles is $X = \mathbb{R}^M_+$, the set of all real vectors with M non-negative components.

In case $M = 2$, a consumption bundle is $(x_1, x_2) \in X$: if you consume 1 apple and 3 oranges, your consumption is described by the vector $(1, 3) \in \mathbb{R}^2_+$, where $x_1 = 1$ is your consumption of apples and $x_2 = 3$ is your consumption of oranges.

The case with $M = 2$ is easy to visualize: in this case, any *feasible consumption bundle* can be represented as a **point on the positive orthant** of the Cartesian plane (x_1, x_2).

2.2 The budget set

In the theory of consumption, B is called the **budget set** of the agent. It contains all the consumption bundles that the agent **can afford**, given the purchasing power of her income.

In the standard problem, the budget set has a simple (indeed, **linear**) mathematical structure.

Each good m is assumed to have a unique and constant **unit price**, represented by a positive real number $p_m \in \mathbb{R}_+$. Call $p = (p_1, \ldots, p_M) \in \mathbb{R}^M_+$ the **price vector** faced by the consumer, and denote the consumer's income with the letter Y.

Given Y and given $p = (p_1, \ldots, p_M)$, a bundle $x = (x_1, \ldots, x_M) \in X$ is affordable if and only if:

$$p_1x_1 + p_2x_2 + \cdots + p_Mx_M \leq Y$$

Or, shortly, $\sum_{m \in M} p_mx_m \leq Y$. This inequality is called the **budget constraint** of the agent. Its fulfillment means that x costs less than the total income Y at the agent's disposal, given the market prices p.

Using the notation of the previous lecture, we say that $x \in B$ if $x \in \mathbb{R}^M_+$ and $\sum_{m \in M} p_mx_m \leq Y$, and call B the **budget set**. More formally,

$$B = \left\{ x : x_m \geq 0 \ \forall m \in M; \sum_{m \in M} p_mx_m \leq Y \right\}.$$

Consumption bundles outside of the budget set are *physically possible*, but the agent is prevented from choosing them because of an important *institutional constraint*: the existence of **private property**, which is characteristic of a *market economy.*

Consider now the case in which $M = 2$ and focus on the Cartesian plane (x_1, x_2).

There, a bundle is affordable if $p_1x_1 + p_2x_2 \leq Y$. This can be rewritten as $x_2 \leq \frac{Y - p1x_1}{p_2}$, corresponding to the portion of the positive orthant located *south-west* of the the **budget line** $x_2 = \frac{Y - p_1x_1}{p_2}$.

The budget line represents the **frontier** of the affordable bundles for the agent: if a bundle lies *above* the frontier, it is *not* affordable, if it lies *on the frontier or below*, it *is* affordable. Also, if an affordable bundle is *not* on the frontier, there is some other affordable bundle that contains *more* of all goods; instead, starting from a bundle that lies *on* the frontier, I *cannot* affordably increase consumption of everything: if I want to consume *more* of $h \in M$ I need to consume *less* of at least another good $j \neq h$.

The shape of the budget line is thus economically insightful. It has intercept $\frac{Y}{p_m}$ on each x_m-axis, and negative slope $-\frac{p_1}{p_2}$. The negativity of the slope has a simple interpretation: if I start from a bundle that lies *on the budget line*, exhausting all of my income, and I want to consume more of one good, I need to consume less of the other. The quantity of good $m = 2$ that I need to give up in order to purchase one more unit of $m = 1$ is found by solving:

$$p_1(x_1 + 1) + p_2(x_2 + \Delta x_2) = Y$$

meaning that

$$-\Delta x_2 = \frac{p_1(x_1 + 1) + p_2x_2 - Y}{p_2}.$$

Substituting out $p_1x_1 + p_2x_2 = Y$, you obtain $\Delta x_2 = -\frac{p_1}{p_2}$. The relative price $\frac{p_1}{p_2}$ expresses the amount of good $m = 2$ that I need to sacrifice in order to purchase another unit of good $m = 1$.

Regarding the intercepts, they can be seen as expressions of the agent's *real* income, expressed in terms of each commodity: in fact, they correspond to the agent's level of consumption when she allocates all of her income to a single good.

Say now that the consumer receives extra income from Y to $Y' > Y$. In this case, the budget line moves to a higher parallel line: the consumer can afford a strictly larger region of bundles, but the cost of one good in terms of another is unchanged. Both intercepts are higher because the agent is richer according to any *real* unit of measure.

Say instead that p_1 moves to $p_1' > p_1$. In this case, the intercept on the x_1-axis decreases, while that on the x_2-axis is unchanged: this implies that the agent is effectively poorer, and can afford a strictly smaller regions of bundles after the price change. Also, the slope of the line changes: it becomes steeper because $m = 1$ becomes relatively more expensive in terms of good $m = 2$.

Say now that $p' = kp$, namely that all prices are multiplied by a constant $k > 1$. The new budget constraint is

$$kp_1x_1 + kp_2x_2 = Y$$

which is equal to

$$p_1 x_1 + p_2 x_2 = \frac{Y}{k}.$$

This is as if the consumer's income shrinks from Y to $Y' = \frac{Y}{k}$, and prices are unchanged. This has the interpretation that **only relative prices are relevant**: just like cardinal utility scores, nominal prices have no economic meaning.

To elaborate on the last point, start from an arbitrary budget line $p_1 x_1 + p_2 x_2 = Y$ and divide everything by p_1. You obtain:

$$x_1 + \frac{p_2}{p_1} x_2 = \frac{Y}{p_1}$$

which means that, by redefining $p_1' = 1$, $p_2' = \frac{p_2}{p_1}$ and $Y' = \frac{Y}{p_1}$, you get a new budget line

$$x_1 + p_2' x_2 = Y'$$

identical to the original one. One way to put this, is to say that one nominal price is a **redundant** parameter: a simpler budget expression can always be obtained by re-parametrizing $p_1 = 1$ and expressing all *nominal variables* as *real variables* measured in terms of $m = 1$. The good chosen as a unit of measure is called the **numeraire**.

Exercises.

- Say that $M = 4$, $Y = 10$, $x = (3, 3, 3, 2)$. If $p = (1, 1, 1, 1)$, does $x \in B$? What if $p = (0.5, 0.5, 1, 1)$?
- Consider the examples with $M = 2$ and say that oranges are twice are expensive as apples. What is the slope of the budget line? What are the intercepts if $p_1 = 1$ and $Y = 6$?
- In the latter example, say $p_1 = p_2 = 1.5$. Write the old and the new budget lines and comment the changes.

Further discussion. The budget set is a simple mathematical object that hides a lot of assumptions.

First, it postulates that M **markets exist**, one for each good, and that the consumer takes all purchasing decision at the same time across different markets.

Second, it assumes that each good m has some uniquely defined *unit price* p_m. This means that:

- Unit prices do not depend on quantities (**linear pricing**: no quantity discounts exist).
- The consumer can demand any quantity of the good that she can afford.
- The consumer can buy any infinitesimal quantity of any good at a uniform price (there are producers willing to supply an atom of an apple).
- **Law of one price**: you cannot purchase the same good at different prices from different providers.

Finally, it is assumed that the consumer has **no power to affect prices**: she is a **price-taker**; prices are *parametric* in her optimization problem. There is **no bargaining** going on between consumers and producers.

2.3 Utility and marginal utility

As we have seen in the first lecture, any preference relation P_i can be represented by means of a utility function U_i which respects the basic axioms of rationality. Additional assumptions can be introduced in a market setup, in order to provide P_i and U_i with a set of *desirable* characteristics.

Well-behaved utility functions. A utility function U_i is said to be **well-behaved** when, in addition to rationality, it respects:

- **Monotonicity**. If for each $m \in M$, $x_m \geq x'_m$, then $U_i(x) \geq U_i(x')$. Moreover, if $x_m > x'_m$ for each $m \in M$, then $U(x) > U(x')$.

- **Convexity.** Take two bundles x and y such that $U_i(x) = U_i(y)$. Suppose a third bundle z is such that, for some $\lambda \in (0,1)$ and all $m \in M$, $\lambda x_m + (1-\lambda) y_m = z_m$: then, $U(z) \geq U(x)$ and $U(z) \geq U(y)$.

The first requirement states that the agent is *never worse off* when **adding more of one good**, and she is always made *strictly better off* when adding **more of all goods**. When the first property holds with strict inequality, preferences are said to be **strictly monotonic**.

The second requirement states that, if you *mix in constant proportion two different bundles that have the same value*, you obtain a third bundle that is **weakly better than both** of them. If the property holds with *strict inequality*, you say that preferences are **strictly convex**. To understand why this is reasonable, think of example E3 in the first lecture: it is better to dedicate the same time to three activities than to do the same thing all day long. In general, convexity assumes that *variety* in consumption is pleasant or, equivalently, that the *more* you consume something, the *less* you value additional consumption of it with respect to other things.

Preferences need **not** to satisfy monotonicity and convexity to be rational. These are properties that are seen as realistic and that also make the consumption problem tractable even without adding much more structure.

Indifference curves. An **indifference curve** is a **level curve** of the utility function: namely, a *set* that contains *all the elements of* X that have the **same utility** for the consumer.

For a fixed level of utility U^0, an indifference curve is identified with the expression $U_i(x) = U^0$.

Take the case with $M = 2$. Indifference curves of well-behaved functions have the following properties:

- Indifference curves **do not cross**. If indifference curves U^0 and U^1 cross at $x = x'$, the bundle x' is strictly preferred to itself which violates reflexivity.

- Higher indifference curves than U^0 are wholly located north-east of U^0 **and** all points north-east of U^0 are better than U^0 (**exercise**: prove it using monotonicity).

- Indifference curves are weakly **decreasing**. Indeed, moving from bundle $x = (x_1, x_2)$ to $x' = (x'_1, x'_2)$, with $x'_1 > x_1$, because of monotonicity you need to diminish $x'_2 < x_2$ in order to stay indifferent.

- Indifference curves are **weakly convex**: any linear combination of indifferent bundles is weakly preferred to both.

Relaxing convexity and monotonicity. If U_i is not monotonic, it may show **satiation**: namely, there can be a critical amount $\bar{x}_m$ of a good such that, over this threshold, there is no additional value in consuming a higher x_m. If there is a specific bundle $\bar{x} = (\bar{x}_1, \bar{x}_2)$ that maximizes utility in the whole set X, including all the bundles that are more expensive, we say that $\bar{x}$ is a **bliss point**. This is for example the case when indifferent curves form concentric circles, whose center is $\bar{x}$.

If one good is actually a *bad* (has negative marginal utility: you dislike to consume more of it), indifference curves are increasing. You can show this as an exercise. Similarly, if one good is a *neutral*, indifference curves are parallel lines to one axis.

If indifference curves are (strictly) concave, mixtures of indifferent bundles are (strictly) less preferred than the original ones. This is the case when you don't want to mix different goods, for example ice cream and anchovies, or when *the more you consume something, the more you value additional consumption.*

Marginal utility. A central concept in microeconomics is the **marginal utility** of a good, namely the *additional value* that the agent attributes to an increase in consumption of a good $m \in M$, starting from the bundle $x = (x_1, \ldots, x_M)$.

Think again to our example with apples and oranges. The **theory of value** of Neoclassical economics only allows to define the value of an apple for the agent with reference to *how many apples and oranges are already in her basket*; that is to say, *the value of one more apple given that I already have one apple and three oranges.*

To see how this applies in more general terms, consider a utility function $U_i : \mathbb{R}^M \to \mathbb{R}$ that describes the preference of agent i over M goods.

Since goods are perfectly separable, starting from any $x \in X$, it is possible to compute the increment of the agent's utility for any variation Δx_m of consumption of a single good $m \in M$. This is

$$\Delta U_i(x, \Delta x_m) = U_i(x_1, \ldots, x_m + \Delta x_m, \ldots, x_M) - U_i(x).$$

The **marginal utility of** m ***evaluated at*** x coincides with the limit of $\Delta U_i(x, \Delta x_m)$ for an arbitrarily small variation Δx_m.

Under the due assumption about the differentiability of U_i, the marginal $MU_m(x)$ utility of good m at x can be defined as follows:

$$MU_m(x) = \frac{\partial U_i}{\partial x_m}(x_1, \ldots, x_M)$$

In the example with apples and oranges, this is expressed as $MU_1(1,3) = \frac{\partial U_i}{\partial x_m}(1,3)$.

If preferences are **monotonic**, the following holds:

$$MU_m(x) \geq 0 \quad \forall m \in M$$

restating the fact that it *never hurts* to consume *more* of a good.

Convexity of preferences is implied by the following property, the **law of diminishing marginal utility**:

$$\frac{\partial MU_m(x)}{\partial x_m} \leq 0 \quad \forall m \in M$$

which is a popular assumption in utility theory stating that, if you consume more of good m, the marginal utility of it does **not** increase. In other words (with **strict** diminishing marginal utility), the more apples you consume, the less you want to consume additional apples (or, in the **weak** case, your appetite for apples does *not increase* along with your apple consumption).

Example of utility functions.

- **Quadratic utility.** With a single good, a quadratic utility function is:

$$U_i(x) = x - \frac{x^2}{2}.$$

- **Cobb-Douglas utility.** Cobb-Douglas utility functions are standard in consumption theory. With $m = 2$ they look like:

$$U_i(x_1, x_2) = x_1^\alpha \, x_2^\beta.$$

- **Quasi-linear utility.** A quasi-linear utility function is linear in one good only:

$$U_i(x_1, x_2) = u(x_1) + x_2$$

  where $u'(x_1) > 0$ and $u''(x_1) < 0$. Typically:

$$U_i(x_1, x_2) = \sqrt{x_1} + x_2$$

or

$$U_i(x_1, x_2) = \log(x_1) + x_2$$

- **Perfect substitutes.** This is the case of a linear utility function:

$$U_i(x_1, x_2) = \alpha x_1 + \beta x_2.$$

- **Léontief utility function (perfect complements).** In this case,

$$U_1(x_1, x_2) = \min\{\alpha x_1; \beta x_2\}.$$

Then:

Exercise.

- Draw the indifference curves of Cobb-Douglas, perfect substitutes and perfect complements, and comment their shape.
- Show that:
  - Quadratic preferences are not monotonic.
  - Cobb-Douglas preferences are strictly convex.
  - Perfect substitutes are not stricly convex.
  - Perfect complements are not strictly monotonic.

3 Theory of consumption, II.

3.1 Constrained optimization: graphical

In the case with $M = 2$, graphical analysis provides insights into the individual optimization problem when U_i is well-behaved.

Given a budget set B and a well-behaved utility function U_i:

- If there is a bundle such that the **indifference curve is tangent to the budget line**, that bundle is an **optimal** choice.
- If there is **no such bundle**, the optimal choice coincides with a **corner** of the budget line.

Moreover, if U_i is strictly convex, the optimal choice $x^*(B)$ is unique.

We visualize the argument in two steps.

Walras' Law. The **Walras' Law** states that, given a well-behaved utility function U_i and a budget set B, any optimal bundle $x \in x^*(B)$ **must belong to the budget line**: namely, the *total expenditure* $\sum_m p_m x_m$ of $x \in x^*(B)$ must be *equal to the level of income* Y.

Indeed, if the agent's income is not exhausted by x, namely $\sum_m p_m x_m < Y$, there exists *another affordable bundle* x' that contains *strictly more of all goods*, i.e. $x'_m > x_m$ for all $m \in M$, and then $U_i(x') > U_i(x)$ because of **monotonicity**.

Tangency condition. Suppose now by contradiction that there is an optimal $x \in x^*(B)$ that is not an extreme of B, and is not such that the indifference curve passing through x is tangent to the budget line.

Graphically, this has the implication that a region of the space (x_1, x_2) lies north-east of the indifference curve and south-west of the budget line. Any bundle y on the south-west boundary of the region (the part coinciding with the indifference curve $U_i(x)$) is indeed not optimal, because it does not exhaust the agent's income. But, since $U_i(y) = U_i(x)$, also x cannot be optimal.

This proves that interior bundles *not* verifying tangency are **never** optimal. To prove the other way round, suppose that an interior bundle x verifies the tangency condition and is not optimal. You can see graphically that, because of the convexity of the indifference curve $U_i(x)$, this cannot apply: any point lying north-east of $U_i(x)$ would be **unaffordable**.

However, it may be that no point on the budget line verifies the tangency condition. In this case, the optimum needs to be an extreme point of the budget line: as an exercise, show that with a geometric argument.

Exercise.

- Represent graphically the optimal choice for each utility function that we have studied.

3.2 The Marginal Rate of Substitution

As we have reviewed yesterday, Neoclassical economists think of the *value of a good* as a *marginal quantity*, a *local* measure, a *limit* expression that only makes sense when *evaluated at a certain point.*

The marginal utility $MU_m(x)$ of a good is not a univocal measure of its additional value at x. The value of $MU_m(x)$ depends in fact on the specific functional form of U_i, but many equivalent functional forms exist given the same underlying preference ordering P_i.

A much more *objective* (and, as we shall see, *observable*) definition for the marginal value of a good is the *ratio* of its marginal utility to that of a different commodity. Evaluated at a given point x, this ratio is called the **Marginal Rate of Substitution** (MRS_{hj}) of good $h \in M$ with respect to $j \in M$. This corresponds to:

$$MRS_{hj}(x) = \frac{MU_h(x)}{MU_j(x)}.$$

In this formulation, the marginal value of h at x is expressed in terms of that of j. The good j becomes a *unit of measure* to express *how worth* h is (locally).

Alternatively, this means that the MRS_{hj} can be interpreted as the **maximum real price in terms of** h at which the consumer wants to buy an additional marginal quantity of j.

Consider for example a bundle of one apple and three oranges. If $MRS_{12}(1, 3) = 2$, an additional *atom* of an apple is *twice as worth* as an additional atom of an orange. Therefore, it the agent's income is exhausted at $x = (1, 3)$, the consumer is willing to exchange a *positive variation of apples with a negative variation of oranges* whenever their relative price $\frac{p_1}{p_2}$ is less than two.

MRS and Indifference Curves. Consider the case of $M = 2$ and take an arbitrary indifference curve with equation $U_i(x_1, x_2) = U^0$.

The $MRS_{12}(x)$ has an important interpretation when $M = 2$: it corresponds to the slope of the indifference curve passing through $x = (x_1, x_2)$. To see why, note that adding a small quantity dx_1 of one good provides an additional utility of $MU_1(x)dx_1$, meaning that, in order to keep the agent indifferent with respect to (x_1, x_2),

$$MU_1(x)dx_1 + MU_2(x)dx_2 = 0$$

must hold. This can be rearranged into:

$$\frac{dx_2}{dx_1} = -\frac{MU_1(x)}{MU_2(x)}.$$

You can think of the $dx_j = MRS_{hj}(x)dx_h$ as the **maximum** quantity of good j that the agent is **willing to give away** in order to obtain dx_h more units of good h, when dx_h is very small.

This has an interesting implication.

Note in fact that *all the utility functions U_i mapping the **same system of preferences** P_i generate the **same indifference curves***. Therefore, under a monotonic transformation of U_i we obtain a map of indifference curves with the same slope for each $x \in X$, meaning that, under U_i and its monotonic transformation, the **Marginal Rate of Substitution** $MRS_{hj}(x)$ is the same at every $x \in X$, for any $h, j \in M$. As a consequence, **the MRS is invariant under monotonic transformations** and constitutes a measure of the marginal value of a good which is independent from the cardinal metrics adopted to represent U_i.

Exercises.

- Compute the $MRS_{12}(x)$ of a generic bundle x with Cobb-Douglas utility, quasilinear utility, perfect complements and perfect substitutes.

3.3 Constrained optimization: analytical

When $M = 2$, the consumer's decision can be formulated as an **unconstrained** optimization problem by substituting away x_2 in the objective function:

$$\max_{x_1 \in \mathbb{R}_+} \quad U_i\left(x_1, \frac{Y - p_1 x_1}{p_2}\right) \tag{1}$$

This can be done because the Walras' Law guarantees that the budget constraint is always optimally binding.

By equating to zero the first derivative of the objective function, it can be found that a sufficient condition for (x_1, x_2) to be optimal is to solve the following system:

$$\begin{cases} MU_1(x_1, x_2) & = \frac{p_1}{p_2} MU_2(x_1, x_2) \\ p_1 x_1 + p_2 x_2 & = Y \end{cases}$$

In order to give an economic interpretation of the first condition, start by noting that any bundle on the budget line such that $x'_1 > x_1$ would in fact verify $MRS_{12}(x') \leq MRS_{12}(x)$ (due to $MRS_{12}(x) = \frac{p_1}{p_2}$ and to $MRS_{12}(x)$ being decreasing in x_1). Then, two cases can be distinguished:

- $MRS_{12}(x') = MRS_{12}(x) = \frac{p_1}{p_2}$. Then, x' verifies the sufficient condition and is also optimal.
- $MRS_{12}(x') < \frac{p_1}{p_2} = MRS_{12}(x)$. Thus the $MRS_{12}(x')$ is lower than the relative price, and the agent prefers to curb consumption of $m = 1$ with respect to x'_1: x' cannot be optimal.

A similar reasoning applies to bundles such that $x' < x_1$: the agent wants to consume more of good $m = 1$ than x'_1 when the inequality is strict, otherwise x' is also optimal.

If there exists no feasible x that verifies the system of sufficient conditions characterized above, the optimal bundle has either $x_1 = 0$ or $x_2 = 0$: this is called a **a corner solution**, namely a solution such that the feasibility constraint $x_m \geq 0$ binds for some $m \in M$. To interpret this, note that, it the sufficient conditions doesn't hold for any feasible bundle, all interior points of the budget line cannot be optimal: indeed, for any such x' the MRS would be either higher or lower than relative prices, and the agent would a tension to consume more or less of $m = 1$ than x'_1. Once the feasibility constraint is hit, the agent cannot push x_1 lower (or higher), and an optimum bundle can be identified.

The **constrained optimization problem** with $M > 2$ goods is as follows:

$$\begin{aligned} \max_{x \in X} \quad & U_i(x) \\ \text{s.t.:} \quad & \sum_m p_m x_m \leq Y \\ & x_m \geq 0 \qquad \forall m \in M \end{aligned}$$

If a corner solution is excluded (which can be ensured by introducing additional assumptions on U_i, not covered here) the program has the following FOCs with Lagrange multiplier λ:

$$\begin{aligned} MU_m(x) &= \lambda p_m \qquad \forall\, m \in M \\ \sum_m p_m x_m &= Y \end{aligned}$$

Which, by substituting away λ, can be restated as

$$MRS_{hj}(x) = \frac{p_h}{p_j} \qquad \forall (h,j) \in M$$

Therefore, the equivalence requirement between MRSs and relative prices extends to any pair of goods in the general framework with $M > 2$.

Exercises.

- Suppose $U(x_1, x_2) = 2\sqrt{x_1 x_2}$, $p_1 = 1$, $p_2 = 1$. Which kind of utility function is this? Solve the maximization problem analytically. What happens as $p_1 = k$ varies? What if Y changes?
- Suppose $U(x_1, x_2) = \log(x_1) + x_2$, $p_1 = k$ and $p_2 = 1$. Find the optimal choice as a function of k and Y.

3.4 Comparative statics.

An important domain of microeconomic analysis is **comparative statics**: the study of *how optimal choices respond to the variation of exogenous parameters.*

An interesting problem of comparative statics in the theory of consumption is how changes in **prices** and **income** affect consumption decisions.

Change of income. Start from the case in which p is taken as fixed and Y is changed to $Y' > Y$. If, when income *increases*, the optimal consumption of all $m \in M$ *increases*, all goods are said to be **normal goods**; otherwise, if the optimal consumption of x_m *decreases* for some $m \in M$, then m is said to be an *inferior good.* An inferior good is thus a good that is consumed *less* when people get *richer.*

The effect of a change of income on optimal consumption is called the **income effect**: if the *income effect* is **negative**, then the good is **inferior**. You can think of inferior goods as goods that you want to substitute out with better alternatives once that you have provided for subsistence: for example, low-quality food and clothing.

Note that the income effect is *local*: a good can be inferior or not, according to the level of income or the price vector.

Change of own price. Say now that p changes and in particular, everything else being equal, $p'_m < p_m$ for some $m \in M$ (all the definitions are valid in the opposite directions if $p'_m > p_m$, with no substantial difference).

If the optimal x_m increases when $p'_m < p_m$, m is said to be an **ordinary good**; otherwise, it is said to be a **Giffen good** (an equivalent definition is that x_m increases when p_m increases).

The existence of Giffen goods may seem puzzling, as it defies the common interpretation of demand being *decreasing in prices*. This intuitive logic translates in the fact that for any good, there exists a **substitution effect** going on the opposite direction of the price change: in simple terms, as a price of a good goes up, the consumer is **always** pushed to diminish consumption of the more expensive good, with the intent of buying more of the alternatives that have become relatively cheaper. However, it must be noted that, when a price changes, the agent becomes also **richer** or *poorer* in *real* terms. Then, the *income effect* kicks in, which can have the *opposite sign* if the good is inferior. A Giffen good is a good *so inferior* that the *income effect* overcomes the *substitution effect*.

Income and substitution effect: graphical For the case of $M = 2$, it is possible to disentangle the income and substitution effects graphically.

For an arbitrary income Y and price vector p such that x is the optimal choice of the agent, consider a new price vector p' such that $p'_1 > p_1$ and $p' = p_2$ and trace the new budget line. Since the consumer is poorer in real terms, her affordable consumption possibilities are reduced after the price variation, and the new budget line is entirely south-west of the old one.

The new optimal choice, say x', lies at the point of tangency between the new budget line and a new indifference curve. The shift from x to x' compounds *income and substitution effects*: it combines the effect of $m = 1$ being relatively more expensive and of the agent being poorer after the price variation.

Call now $\tilde{x}$ **the point of the *original indifference curve* that is tangent to a line parallel to the *new budget line***. The *difference* between x_m and $\tilde{x}_m$ following a price change from p to p' is the **substitution effect**.

The substitution effect has always the ***opposite*** *sign of the variation of the own price*, namely:

$$(p'_m - p_m)(\tilde{x}_m - x_m) < 0.$$

It can be noted that $\tilde{x}$ is obtained by providing the agent, after the price change, with **additional** income up to the point in which the budget line is tangent with the original indifference curve: namely, with as much income as it is needed for her to be **as well off as before the price change**. This serves to **compensate** the loss of welfare due to the variation of the agent's real income, and to *isolate* the *impact of substitution* between the two goods as their relative affordability varies.

The **income effect** is defined instead as $\tilde{x}_m - x'_m$, which is the *part of variation* in consumption of m that is **not explained** by the change of relative prices.

Change of other prices. Consider again a price change of $m \in M$ such that $p'_m < p_m$ but focus on the effect on optimal consumption of another good $k \in M$. If x_k increases, the two goods are said to be **complementary goods**, otherwise they are **substitute goods**.

Complementary goods are goods that are **consumed better together**: for example, coffee and sugar. If coffee is *more affordable* after a price variation, and the agent consumes more of it, consumption of sugar rises even if the price of sugar is the same, because the availability of more coffee makes sugar more valuable.

Substitute goods satisfy to some extent the same needs: think for example of coffee and tea. When the price of coffee increases, you can be reasonably expected to adjust your consumption habit by substituting away some coffee with tea.

When goods are only valuable when consumed together, they are said to be **perfect complements**: for example, left shoes and right shoes. This is the case when the utility function has the form

$$U_i(x_1, x_2) = \min\{x_1; x_2\}$$

in which the agent only cares about the minimum of the two quantities. If for example $x_2 > x_1$, the excess consumption $(x_2 - x_1)$ of good $m = 2$ has therefore no value for the agent (the third right shoe is useless if you only have two left shoes).

When two goods satisfy exactly the same needs, they are said to be **perfect substitutes**. This is the case of the utility function

$$U_i(x) = \sum_m x_m$$

in which, since the goods are perfectly fungible, you only care about the *sum* of quantities you consume, no matter the *composition* of your basket.

4 Theory of consumption, III.

In this lecture, we describe the **individual and aggregate demand functions** for a good, define their **elasticities** with respect to various parameters, derive them from the individual consumer's maximization problem and, finally, define the **consumer's surplus**.

4.1 The demand function

Consider a market in which a good $m \in M$ is traded. If, for example, the individual i has a *linear* demand function of the form

$$D^i_m(p_m) = \begin{cases} 5 - 3p_m & \text{if } 5 - 3p_m \geq 0 \\ 0 & \text{if } 5 - 3p_m < 0 \end{cases}$$

it means that, if the price of good m is $p_m = 1$, the agent demands

$$x_m = D^i_m(1) = 5 - 3 \cdot 1 = 2$$

units of the good m.

More in general, the **individual demand function** $D^i_m(p_m)$ expresses the amount of a good that an agent i is willing to consume if its unit price is p_m.

If there are $N > 1$ agents populating a market, the **aggregate demand function** $D(p_m) = \sum_i D^i_m(p_m)$ expresses the cumulative amount of good m that all the agents in the economy are going to consume at the unit price p_m. For example, if there are $N = 4$ agents with the linear demand expressed before,

$$D(p_m) = 20 - 12p_m$$

and, if $p_m = 1$, $D(1) = 8$ is the aggregate market demand for the good m.

Sometimes, during your economic studies, you will find a demand function $D^i_m(p_m)$ represented in terms of its **inverse demand function** $P^i_m(x_m) = D^{i\,-1}_m(x_m)$: for each quantity x_m, $P^i_m(x_m)$ expresses the *maximum* price at which the agent is willing to purchase at least x_m units of the good m.

In our linear example,

$$P^i_m(x_m) = \begin{cases} \frac{5-x_m}{3} & \text{if } \frac{5-x_m}{3} \geq 0 \\ 0 & \text{if } \frac{5-x_m}{3} < 0 \end{cases}$$

Keep in mind the following: when you aggregate demand functions graphically, you **sum quantities vertically** in the most natural way. However, most of the times, demand is represented graphically in its *inverse* form $P_m(x_m)$: when this happens, remember to **aggregate quantities horizontally**, in order to obtain a symmetric graph to the one of $D(p_m)$ with respect to the (x_1, p)-bisector.

Slope and elasticity. The **slope** of the demand curve, $\frac{\partial D^i_m(p_m)}{\partial p_m}$ is the *local* measure of the **absolute variation** of demand for good m with respect to its own price. Common sense suggests that the slope of the demand function should be *negative*: as we reviewed yesterday, this is not the case when the good exhibits Giffen behavior.

When the demand function is linear, its *slope* is constant: in our example, $\frac{\partial D^i_m(p_m)}{\partial p_m} = -3$ for each value of p_m such that the demand is positive.

Since the most natural reason to compute or estimate a demand curve is to understand *what would be the reaction of the consumer to a price change*, looking at the slope of $D^i_m(p_m)$ seems to provide a valid answer. Perhaps not intuitively, this is rarely the case.

One reason is that the slope of a demand function is totally dependent on the unit of measure that you adopt to express prices and quantities. A much more interesting object is the **(own price-)elasticity** of the demand function: the local measure of the **relative (percentage) change** of quantities demanded with respect to a relative change in prices.

In other words, the slope of a demand function says for example: *if price goes up by* 1*, demand goes down by* -3, which is an empty statement if you do not specify 1 euro/dollar/pound/million of euros, and -3 gallons/kilograms/shoes/hundreds of cars, or what else. Moreover, the variation captured by the local slope of a demand function is absolute and is not informative per se on the proportional size of the effect.

Instead, the elasticity of a demand function may say: *if price goes* 10% *up, demand goes* 5% *down.* That is a much more robust statement, because it is independent of the unit of measure, and because it provides an idea of the magnitude of the effect irrespective of its scale.

For a sizeable price variation Δp, the elasticity of demand is:

$$\bar{\epsilon}^i_m(p_m, \Delta p) = \frac{D^i_m(p_m + \Delta p) - D^i_m(p_m)}{\Delta p} \cdot \frac{p_m}{D^i_m(p_m)}$$

As $\Delta p \to 0$, the limit of the expression above is the *local* measure of elasticity evaluated at p_m:

$$\epsilon^i_m(p_m) = \frac{\partial D^i_m(p_m)}{\partial p_m} \cdot \frac{p_m}{D^i_m(p_m)} = \frac{\partial \log(D^i_m(p_m))}{\partial \log(p_m)}$$

The expression above can be used for example to compute the elasticity of the linear demand function $D^i_m(p_m) = 5 - 3p_m$:

$$\epsilon^m_i(p_m) = -3 \times \frac{5 - 3p_m}{p_m} = \frac{-15 + 9p_m}{p_m}$$

Interestingly, a linear demand function has a constant slope, but it doesn't have a constant elasticity. The magnitude $|\epsilon^m_i(p_m)|$ of the own price-elasticity is decreasing in p_m, meaning that the more expensive the good, the higher the reactivity of the consumer to price changes.

When $\epsilon^i_m(p_m) = 0$ and $\epsilon^i_m(p_m) = -\infty$, two extreme cases are identified:

- **Infinitely elastic demand.** When $\epsilon^i_m(p'_m) = -\infty$, demand is *infinitely elastic*: this is for example the case when consumption drops *from a positive quantity to zero* for *any price* $p_m > p'_m$, no matter how small the increase. This is the case when there is a perfect substitute good available on the market whose unit price is exactly p'_m (you can prove it as an exercise).

- **Inelastic demand.** When $\epsilon^i_m(p_m) = 0$, demand **does not change** at all when price varies: the consumer wants to consume a fixed amount of the good, in a way that is unresponsive to the price. This is the case for things that you really need to consume in a specific quantity: you can imagine an (at least locally) inelastic demand for addictive goods (cigarettes), for subsistence goods, housing or accommodation, or for life-saving drugs.

Exercises.

- Draw a linear demand function and try to explain what is the meaning of the two intercepts.
- What is the elasticity of a linear demand curve $D^i_m(p_m) = \alpha - \beta p$ at the intercepts? When is is it equal to one?
- Consider the demand function $D^i_m(p_m) = \frac{1}{p}$ and prove that it has constant elasticity equal to -1. What does it mean?

Income and cross-price elasticity. We have drawn and characterized the demand function $D^i_m(p_m)$ of the agent i for good m. Implicit behind the expression $D^i_m(p_m)$ lies the assumption that, as price p_m changes, all the other relevant factors in the decision problem of the consumer stand still. This is the so-called *ceteris paribus* assumption: namely, $D^i_m(p_m)$ expresses the response of the individual demand for good m to own-price changes at *parity of other conditions*.

Two important *other conditions* that may affect the demand for a good are cross-price variations (variation of the price of other goods) and income variations.

First, in fact, we can define the elasticity of demand for m with respect to the price of other goods, rather than the own price p_m. Denote the (local) **cross-price-elasticity** by

$$\epsilon^i_{mh}(p_h) = \frac{\partial D^i_m(p_h)}{\partial p_h} \frac{p_h}{D^i_m(p_h)}$$

where p_m is an *exogenous parameter* of $D^i_m(p_h)$, in the sense that it is kept fixed before and after the variation of p_h.

If h is a *substitute* of m, $\epsilon^i_{mh}(p_h) > 0$, and the opposite holds if the goods are *complements.*

Second, the elasticity of demand can be computed with respect to the income of the consumer. We can denote the local **income-elasticity** of m by

$$\epsilon^i_{mY}(Y) = \frac{\partial D^i_m(Y)}{\partial p_h} \frac{p_h}{D^i_m(Y)}$$

whose sign depends on the income effect. A good $m \in M$ whose income elasticity is *greater than 1* is said to be a **luxury good**: this is a good whose relative consumption increases in a higher proportion with respect to income variations, meaning that, as your income goes up, you devote a larger share of your earnings to the consumption of m (and vice versa). If $\epsilon^i_{mY} < 1$, the good is said to be a **subsistence good**.

4.2 Derivation of the demand function

The demand curve, expressed in the form $D^i(p_m)$, is the simplest and most popular model of consumption behavior on a market. Its computation or estimation provides a simple and intuitive guidance for policy-making of investors, managers, entrepreneurs, politicians, et cetera.

Yet, the demand curve is a very complex object. First, it is an elusive entity as you can **never observe** the whole of it in the real world: you can only observe equilibrium prices and quantities, and their variations over time. Second, as we see in this section, each point of the demand function is the result of a complicated optimization program that happens *in the back of the mind of the consumer*, taking into account the relative price of all the goods available on the market and real income in terms of every good, for every price vector that may show up in hypothetical terms.

Despite this complexity, microeconomics provides some instruments to construct a fairly tractable model of market demand. Starting from the infinite hypothetical maximization problems of each individual, we can provide a useful representation of the demand curve that can account for different observed phenomena (income effects, subsistence versus luxury, complementarity versus substitution...) in a mathematically tractable way, also providing a framework for *normative* considerations that we will see at the end of the lecture (the theory of the **consumer surplus**).

There are many caveats and many simplifications in this process, and many objections that can be made; nevertheless, the power and the elegance of this framework partly explains why the Neoclassical paradigm has been so successful despite its obvious shortcomings.

We will address for simplicity the case in which $M = 2$ and preferences are *strictly convex*, meaning that, for each budget set B, the optimal bundle $x^*(B)$ of the consumer is uniquely identified.

Assume that good $m = 2$ is the numeraire, meaning that we keep $p_2 = 1$ with no loss of generality (because only relative prices matter). Denote as Y the real income of the agent expressed in terms of $m = 2$, and as $p = \frac{p_1}{p_2} = p_1$ the relative price which, in our unit of measure, coincides with the price of $m = 1$.

In the optimization problem of the consumer, (p, Y) are *exogenous* parameters, namely quantities that are *relevant* for the consumer's decision-making, but the consumer cannot change. Since each (p, Y) identifies a budget constraint B, it is the same to say $x^*(p, Y)$ and $x^*(B)$.

Remember that $x^*(p, Y)$ is a vector, that has components $x^*(p, Y) = (x_1^*(p, Y), x_2^*(p, Y))$. At this point, analytically, we easily obtain our demand curve $D_m^i(p_m)$ for m: it coincides with the $m = 1$ component $x_1^*(p, Y)$ of the optimal bundle, computed for each value of p keeping Y fixed:

$$D_1^i(p) = x_1^*(p, Y) \qquad \text{for a given } Y.$$

Geometrically, for $p = p'$, we can identify $x^*(p', Y)$ as the point of tangency between the budget line and the highest affordable indifference curve. The x_1-coordinate of $x^*(p', Y)$ identifies $D_1^i(p')$. In this formulation of the demand function, Y remains implicit and *hidden* in the shape of D_1^i. As p' varies to p'', the budget line pivots around its intercept on the x_2 axis (which remains constant to $x_2 = Y$) and a new point of tangency $x^*(p'', Y)$ can be identified.

It can be noted that, by letting p take all admissible values, it is possible to trace out a locus on the positive (x_1, x_2) orthant by connecting all the optimizal bundles $x^*(p, Y)$: this is called the **price-offer curve**.

Then, for any p (any slope of the budget constraint), a point $x^*(p, Y)$ of the price-offer curve is identified: the demand curve $D_1^i(p)$ of $m = 1$ translates all the points of the price-offer curve on the (p, x_1)-plane, meaning that, for each $x^*(p, Y)$ on the price offer curve, it returns the point $\big(p, D_1^i(p) = x_1^*(p, Y)\big)$ of the demand function.

Changes *on* and *of* the demand function. Say now that the price of a good changes from p' to p''. The individual demand of the good changes from $D_1^i(p')$ to $D_1^i(p'')$. A variation has occurred *on* the demand function, moving the description of the consumer's behavior from one point of it to another.

Say instead that income changes from Y to Y' but $p = p'$ stands still. This means that the *shape* of D_1^i changes totally, as it reproduces a *new* price-offer curve that is obtained by shifting the original budget line to a parallel line, and traced out by pivoting the new budget line around its new x_2-intercept.

In this case, the demand for a good does not change because price is changed (a shift **on** the demand curve), but it changes because income has changed even if price is unaltered (shift **of** the demand curve). This difference is extremely important: for example, practitioners trying to estimate a demand function need to disentagle one effect from the other, as they observe demand shifts taking place over time.

Suppose for example $m = 1$ is a *normal good* for each value of Y and p. In this case, an increase in income will *shift the whole demand function* north-east, meaning that the consumer is willing to **purchase more of the good at any price** after the change.

If $M > 2$, another important factor that can trigger changes **of** the demand curve is the *change of the prices of other goods*. For example, with $M = 3$, the optimal bundle is expressed by $x^*(p_1, p_3, Y)$, as it also responds to variations of the relative price p_3 of good $m = 3$ in terms of the numeraire $m = 2$. Again, the dependence of $D_1^i(p_1)$ on p_3 is implicit in the shape of the function. The variation of p_3 can trigger shifts of $D_1^i(p_1)$ in ambiguous directions, according to the goods $m = 1$ and $m = 3$ being substitutes or complements.

Exercise

- Suppose that $M = 3$, and that $m = 1$ and $m = 3$ are complements. If p_3 increases, how will $D_1^i(p_1)$ shift?
- Compute the demand function for Cobb-Douglas, perfect complements and perfect substitutes, and represent graphically the shift of the demand function for a positive change of income.

4.3 The surplus of the consumer

We have seen that, by inverting the demand function $D_m^i(p_m)$ for a good, one obtains its *inverse demand function* $P_m^i(x_m)$ which expresses *the maximum unit price that the consumer is available to pay in order to purchase x_m units.*

This seems to provide an intuitive (but wrong) solution to the problem of **value** of economics: namely, inverse demand seems to give an *objective* definition of *what is the value of a good for a consumer.*

Supppse the inverse demand for m is $P_m^i(x_m) = \frac{5-x_m}{3}$ if $x_m \leq 5$, from our running example. We are tempted to interpret the **(maximum) expenditure function**, or **revenue function**

$$R_m^i(x_m) = P_m^i(x_m) \cdot x_m$$

as a measure of the **total value** of x_m units of $m \in M$ in terms of money, of whatever is the numeraire. This seems the case, as $R_m^i(x_m)$ returns the largest amount of money (or of the numeraire) that the consumer is willing to give up in order to purchase x_m units of the good. The implication *appears* to be that, by taking the difference between $R_m^i(x'_m) - R_m^i(x_m)$, I obtain a monetary comparison of the consumer's utility in two different situations in which either x'_m or x^m units of the good are purchased.

In theory, this is wrong. In practice, this is a reasonable approximation that that can be used to provide practitioners with policy advice (and caveats).

To see why the variation of the expenditure function $R_m^i(x_m)$ is not a monetary measure of variation in the agent's welfare when consuming different x_m, you need to recall that, when prices change, there are **income effects** entering the picture. This means that, as you take the difference between $R_m^i(x_m)$ and $R_m^i(x'_m)$, you are also capturing the fact that, following the price change that shifted her consumption from x_m to x'_m, the agent has to recalibrate expenditure according to her new *real* income: her increase/decrease in expenditure for a good m is not only due to the variation $x' - x_m$ of quantity that she consumes; it also reflects that the agent has got richer or poorer after the price change.

If you want to compute a monetary measure of welfare variations after a price variation, you need a way to quantify what we called the **substitution effect** in the previous lecture. We do not go further in this direction in this introductory course; if you are interested, Varian provides nice instructions in Chapter 14 on how to do that. For the rest of our lectures, we content ourselves of using the $R_m^i(x_m)$ as an approximation of what is called the **consumer surplus**.

Quasi-linear utility. To conclude on this topic, we need to consider one of the most beloved functional forms by applied economists: **quasilinear utility**.

If $M = 2$, under the condition that $u(x_1)$ is an increasing and concave function of x_1, and the normalization that $u(0) = 0$, quasilinear preferences are represented by the utility function

$$U_i(x_1, x_2) = u(x_1) + x_2,$$

and, setting x_2 to be the numeraire, calling $p = p_1$ and substituting out the budget line rewritten as

$$x_2(x_1) = Y - px_1$$

within $U_i(x_1, x_2)$, we obtain that the consumer maximizes

$$\max_{x_1 \leq \frac{Y}{p_1}} = u(x_1) + Y - px_1.$$

Under the simplifying assumption that $u'(0) = \infty$ (that is, you really do not want to starve out of $m = 1$), the problem is solved by taking the derivative of the objective function and checking whether it attains a zero for a value of x_1 respecting the budget constraint.

Consider then the (unconstrained) FOC

$$u'(x_1) - p = 0$$

Define implicitly $x_1(p)$ as the value of x_1 solving the equation above. If $x_1(p) \leq \frac{Y}{p}$, then the optimal bundle is

$$x^*(p, Y) = (x_1(p), Y - px_1(p))$$

Otherwise, a corner solution is optimal and the entire income of the agent is optimally allocated to good $m = 1$

$$x^*(p, Y) = \left(\frac{Y}{p}, 0\right)$$

The inverse demand function is thus very easy to pin down with quasilinear utility. It is simply:

$$P_1^i(x_1) = u'(x_1)$$

Note that now, in computing the inverse demand function $P_1^i(x_1)$, **income is irrelevant** except when a corner solution is obtained. This is because a quasi-linear utility function has **parallel indifference curves**, meaning that there is absolutely **no income effect** going on among interior solutions (you can verify it graphically as an exercise, using the *comparative statics* techniques from last lecture).

This has a nice implication: if you interpret your numeraire as money (where *money* is intended as a **composite good**, standing for *cumulative expenditure on anything else*), and if your consumer has quasi-linear preferences, you can effectively measure the monetary variation of her welfare due to the shift of consumption of $m = 1$ from some x_1 to another quantity x_1'. This is done by taking the integral of the inverse demand function over $[x_1, x_1']$:

$$\Delta u(x_1, x_1') = \int_{x_1}^{x_1'} P_m^i(x)dx$$

Which is an obvious equivalence since $P_1^i(x_1) = u'(x_1)$.

The expression $\Delta u(x_1, x_1')$ is called the **variation of the *gross consumer surplus* from x_1 to x_1'**. By noting that $\Delta u(0, x_1') = u(x_1')$, one obtains a monetary measure of the total welfare that the consumer obtains from the consumption of x_1' units of $m = 1$.

Consider finally the following expression:

$$CS(x_1) = u(x_1) - px_1$$

This is the **(net) consumer surplus**. It expresses the **additional welfare** that the consumer obtains from the transaction when buying x_1 units at the unit monetary price p. We obtain that by subtracting, from the monetary value $u(x_1)$ of consumption, the amount of money px_1 that the consumer needs to pay.

If you consider a graphical representation of inverse demand $P_m^i(x_1)$ on the (x_1, p)-plane, you can easily visualize $u(x_1)$ to be the area underlying $P_m^i(x_1)$ from $P_m^i(0)$ to $P_m^i(x_1)$, and $CS(x_1)$ to be the difference between that area and the underlying rectangle $\{(0, 0); (x_1, 0); (0, p); (p, x_1)$.

Exercises.

- Consider the quasilinear quadratic function $U_i(x_1) = x_1 - \frac{x_1^2}{2} + x_2$ and show that it implies a linear inverse demand function by fixing an arbitrary level of Y. Does U_i respect the monotonicity (non-satiation) assumption? How is that related to the linearity of $P_1^i(x_1)$?

5 Theory of consumption, extra topics.

5.1 Interpersonal comparisons

We have presented a framework to understand and represent the preferences of a single agent over a choice set X. In this context, we have introduced a *subjective* theory of value that allows to compare the welfare of a single agent in different situations.

Many normative questions in economics require interpersonal welfare comparisons: as that there is no objective metrics for utility, microeconomics alone can only provide a few answers on this topic.

Pareto-efficiency. For the rest of the lecture, we will refer to a choice $x \in X$ in multiagent settings as an **allocation**. An important criterion that normative economics adopts to rank different allocations $x \in X$ is the **Pareto criterion**.

The *Pareto criterion* is an **incomplete ordering** that, starting from subjective utilities as we have defined in the previous classes, allows to unambiguously rank *some, but not all, of the allocations*, according to their *social* desirability.

Suppose that the set of all feasible allocations in an economy is X and denote the N agents populating the economy as $i = 1, 2, \dots, N$, each endowed with a utility function $U_i : X \to \mathbb{R}$. A choice $x \in X$ is said to Pareto-dominate $x' \in X$ if:

- For all the agents in the economy, $U_i(x) \geq U_i(x')$.
- For at least one agent i, $U_i(x) > U_i(x')$.

Namely, x makes **no one worse off** with respect to x', and makes **at least one agent strictly better off**. If *all the agent* are **strictly better off** in x, then x **strictly Pareto-dominates** x'.

An allocation $x \in X$ such that *no other* $x' \in X$ ***Pareto-dominates*** x is said to be **Pareto-efficient**.

A nice way to understand Pareto-dominance is the following: suppose that x' is the status quo in the economy, and that for some reason there is the possibility to change from x' to x. Then, if x Pareto-dominates x', no one would object to the change, and at least some agent would strictly benefit from the change.

It is reasonable to say that a benevolent planner with dictatorial power would impose a Pareto-efficient allocation: we don't know which one. In fact, since the Pareto criterion is an *incomplete* ranking, there may be multiple Pareto-efficient allocations, as we will review in the next examples.

The Pareto criterion is thus fairly weak; in particular, it does not take any **equity** consideration into account. Putting it informally, if we compare an allocation $x' \in X$, in which all the agents are doing quite well, to another $x \in X$ in which an agent is doing just slightly better and everyone else is starving, the Pareto-criterion rankes x' and x equally.

The good things about the Pareto criterion are that it is objective (its definition is free from ethical/philosophical speculations), and totally relies on *subjective* utility. The main alternative approach, that takes into account **equity concerns**, is to define a **social welfare function** $W : X \to \mathbb{R}$ that explicitly states the objectives of the planner: of course, the choice of W reflects the philosophical and political view of the scholar. Note that, according to some of the main social welfare functions, **Pareto-inefficient allocations may dominate Pareto-efficient ones**: this contrast between the Pareto criterion and other criteria to rank social choices may give rise to an **equity-efficiency trade-off**.

Example. Suppose $i = 1, 2$ and $\{a; b; c\} \in X$.

Also,

$$U_1(a) = 10; \quad U_1(b) = 21; \quad U_1(c) = 20$$

$$U_1(a) = 10; \quad U_1(b) = -10; \quad U_1(c) = 20$$

You can see that, in this case, c strictly Pareto dominates a, while b is *Pareto-equivalent* with respect to both a and c. Therefore, c and b are Pareto-efficient, a is Pareto-inefficient. This holds despite the fact that b makes $i = 1$ just a little better off than c and has a huge negative impact on $i = 2$.

Exercise. Suppose $i = 1, 2$ are dividing a cake of size $S = 100$ and call x_i the slice of agent i. An allocation $x \in X$ is any (x_1, x_2) such that $x_1 + x_2 \leq 1$. What is the set of Pareto-efficient allocations? Is $(0, 100)$ a Pareto-efficient allocation?

Edgeworth's box (hints). Consider now the situation in which two agents exist, say $i = 1, 2$, and two goods $m = A, B$. Denote by x_{im} a level of consumption of good m by agent i.

Suppose that there exists one unit of each good. This means that for each $m \in M$ the following **feasibility constraint** exists:

$$x_{1m} + x_{2m} \leq 1.$$

Suppose also that each consumer has (strictly) well-behaved preferences for any $(x_{iA}, x_{iB}) \in X_i$ represented by $U_i : X_i \to \mathbb{R}$.

Now, you are asked to play the part of a benevolent central planner that can decide to assign any feasible distribution of each good m to the agents.

The Pareto criterion allows us to identify a restricted set of allocations to choose from: indeed, if we are free to select any feasible allocation with no further constraint, there is no reasonable motivation (i.e. no *monotonic* social welfare function) to select any Pareto-dominated allocation.

We are going to prove in class, graphically, that all the allocations such that the indifference curves of the two agents have the same slope are Pareto-efficient. The set of all these allocations is called the **contract curve** of the economy.

5.2 Uncertainty

Expected utility. Suppose that the agent, when taking a choice $x \in X$, is uncertain about her payoff outcome. In particular, the value of each x depend from something that is stochastic and that she cannot control.

For example, $X = \{Yes; No\}$ can be the choice whether to buy or not an umbrella, when the agent is uncertain about tomorrow's weather.

A choice x taken under conditions of uncertainty is called a **lottery**. We say that the *realized* utility of a lottery x depends on the **state of the world** being $\omega \in \Omega$: any ω is a full description of all those random, relevant things that the agent does not know when taking a choice, and that will be realized only later, when the choice is irreversible. In the example, $\Omega = \{Rainy; Sunny\}$ is the set of possible states of the weather tomorrow; the set Ω of all the possible realizations of ω is called the **state space**.

For any choice $x \in X$, the agent's **ex-post utility** is denoted by

$$u : X \times \Omega \to \mathbb{R},$$

where $u(X, \omega)$ is the utility of choice x provided that ω is the realization of the underlying uncertainty.

In our example, if an umbrella costs 10 but getting wet from rain has utility -20, the agent's ex-post utility function is:

$$u(Y;R) = -10; \quad u(Y;S) = -10; \quad u(N;R) = -20; \quad u(N;S) = 0$$

The standard theory of **choice under uncertainty** assumes that the agent is endowed with an assessment of the probability that each ω will happen: denote as $p_\omega \in (0,1)$ the probability that the agent allocates to the state ω: this is her **prior probability** for state ω. Clearly, it must hold that $\sum_\omega p_\omega = 1$; we denote by

$$p = (p_1, \ldots, p_\omega, \ldots, p_{|\Omega|})$$

the agent's prior distribution. In terms of notation, we write $p \in \Delta(\Omega)$ where $\Delta(\Omega)$ is the space of all probability distributions over the state space Ω.

For example, if it is equally likely that it rains or not, $p_S = p_R = 0.5$.

It is usually assumed then that the agent's evaluation of a choice x can be expressed in a convenient way, as a **Von Neumann-Morgenstern (VNM) expected utility** function. The (VNM) expected value $U(x)$ of a lottery $x \in X$ is defined as the **expected value of the agent's ex-post utility** when taking $x \in X$. Then

$$U(x) = \sum_\omega p_\omega u(x, \omega)$$

which is the *weighted sum of ex-post payoffs in each ω according to their prior probability p_ω*.

In the example,

$$U(Y) = 0.5 \cdot (-10) + 0.5 \cdot (-10) = -10$$

$$U(N) = 0.5 \cdot (-20) + 0.5 \cdot 0 = -10$$

The *standard behavioral rule* assumed in microeconomics is that, when proposed with various lotteries, the agent selects the one **maximizing her expected utility**. In the example with $p_S = p_R = 0.5$, the agent is indifferent between buying or not the umbrella.

Exercise.

- What is the optimal choice of the agent in our example, if $p_S = 0.3$? What if $p_S = 0.7$?

Example: a monetary lottery. Consider the following situation: an agent has an initial amount of money L and must decide whether to take or not a bet, i.e. a **lottery** over monetary outcomes denoted by:

$$x = \begin{cases} L + \alpha & \text{w.p. } p \in (0,1) \\ L - \beta & \text{w.p. } 1 - p \end{cases}.$$

The bet returns a positive payment $\alpha > 0$ with probability $p \in (0,1)$ and a negative payment $-\beta$ with probability $(1-p)$. The **expected value** of the lottery x is therefore $L + p\alpha - (1-p)\beta$.

The lottery is said to be **fair** if

$$p\,\alpha - (1-p)\beta = 0$$

meaning that, in expectation, the final money holdings of the agent are as worth as her initial holdings.

Common sense suggests that the agent is *not indifferent* between an *uncertain* lottery and a *certain* payment having the *same expected value*: apart from expected value, **risk** is an important dimension of the agent's preferences.

It is reasonable to think that the agent, in front of an uncertain monetary outcome, is willing to give away some money (in expectation) in order to reduce her risk: this is the motive for trade in **insurance markets**. The aversion of the agent against taking risks is unsurprisingly called **risk-aversion**: an agent is said to be *risk-averse* when, provided two lotteries with the same expectation and different distribution of payments, the agent prefers the one that is *less uncertain*.

In the rest of the lecture, we show that the agent is risk-averse if (and only if) her preferences for money exhibits *decreasing marginal utility*.

Bernouilli utility function and risk-aversion. The agent's **utility for money** is represented by some continuous and differentiable real function $u : \mathbb{R} \to \mathbb{R}$, that is called a **Bernouilli utility function** and, for monetary lotteries, coincides with the agent's *ex-post utility function*.

The utility of **not taking** a bet is $U_0 = u(L)$, which is the *certain* value of L monetary units.

The utility of accepting the lottery x is represented by the following *Von Neumann-Morgenstern (VNM) expected utility function*:

$$U(x) = p\,u(L + \alpha) + (1 - p)u(L - \beta)$$

that represents the **expected value of the lottery** x. Note that the expected value of a lottery is **different** from the *value of its expectation*, that is

$$u\big(L + p\alpha + (1 - p)\beta\big).$$

One way to understand *risk-aversion* is exactly to note that, if $p \in (0, 1)$, the **expected value of x is less than the value of the expectation of x**. This can be shown to be true whenever u is **strictly concave**, as a consequence of the *Jensen's Inequality.* This implies that **the agent is risk-averse when money has decreasing marginal utility**.

Another way to visualize risk-aversion is to define, for any lottery x, the (monetary) value $L^*(x)$ as follows:

$$u(L^*(x)) = U(x)$$

so that, for any lottery x, $L^*(x)$ represents the *certain* amount of money that makes the agent *indifferent between taking the lottery x or the certain payment $L^*(x)$*. This called the **certainty equivalent** of x.

It can be shown that, if x is **non-degenerate** (meaning that it does not put all the probability on a certain payment), **whenever u is strictly concave, the certainty equivalent $L^*(x)$ is lower than the expected value of x**.

This means that the agent that faces a risk x with expected value $L + p\alpha - (1 - p)\beta$ is willing to pay a price of

$$\rho(x) = L + p\alpha - (1 - p)\beta - L^*(x)$$

in order to avoid the lottery, namely to hedge against the riskiness that x entails. The value $\rho(x)$ is called the **risk premium** of lottery x.

When u is **linear** the agent is **risk-neutral**: in this case, $L^*(x) = L + p\alpha - (1 - p)\beta$ for any p and insurance is **not a motive for trade**: the agent would not give up any amount of her expected income in order to reduce risk. Another way to put that, is to say that the agent is **indifferent among all lotteries having the same expectation**.

In class, we will explore (analytically and graphically) the mathematical intuitions behind all these results.

Exercise.

- Consider the set of all lotteries having expected value $L = 1$ and assume $u''0$. Show analytically that the agent's favorite lottery is degenerate on one payment.