Lecture notes

Course: Preparatory course in Microeconomics
Teaching load: 15h
Instructor: Guillaume Pommey
Last updated: August 27, 2025

Foreword

These lecture notes are intended to serve as a support to the material taught in class. They focus on selected topics of microeconomics and do not cover the entirety of a standard course of introductory microeconomics.

Microeconomics is a branch of economic theory providing rigorous mathematical foundations on how rational people take decisions, and how their individual decisions interact, coordinate and aggregate to affect the outcomes of socioeconomic systems. Microeconomic theory provides a general, consistent and tractable formalization of concepts such as rationality, individual choice, decentralized coordination and economic efficiency. Methodologically, we follow what is known as the *neoclassical approach*, where we assume that all agents are *rational* and seek to maximize their objectives subject to some constraints.

1. Preferences

1.1. Preference relations

The first step in modeling individual behavior of an agent consists in describing (i) what are the options available to the agent, and (ii) how the agent ranks the options (e.g., first, second, third) based on their *preferences.*

Set of alternatives. We denote by X the set of possible (mutually exclusive) alternatives from which the agent must choose. For this section, we do not need to make further assumption on this set which can include a very broad rage of applications. In what follows, however, this set will essentially represents *consumption bundles*. We let $x \in X$ denote a typical element of this set.

Examples

1. The agent is offered (for free!) the choice between an apple, a pear, and an orange. Their set of alternatives in that case writes as $X := \{\text{apple}, \text{pear}, \text{orange}\}$. An element of the set $x \in X$ is one of the fruits, for instance $x = \text{pear}$.

2. Alternatives can be more complicated objects than a single fruit. Assume instead that the agent is offered to pick a fruit (apple or orange) *and* a beverage (tea or coffee). The set of alternatives now contains all the combinations of pairs fruit-beverage, that is, $X := \{$(apple, tea), (apple, coffee), (orange, tea), (orange, coffee)$\}$. An element of this set is a pair, for instance $x =$ (orange, tea)

3. Alternatives can also represents quantities. Assume that the agent is offered fruits and beverages (of any kind) but can pick at most 3 items in total. Let x_1 and x_2 denote the quantity of fruits and beverages chosen by the agent, respectively. The set of alternatives can be written as follows:

$$X := \{(x_1, x_2) \in \mathbb{N}^2 \mid x_1 + x_2 \leqslant 3\}.$$

For instance, $(0,1)$, $(1,2)$, and $(3,0)$ are elements of X as they all have at most three items in the bundle. Instead, bundles like $(3,1)$ or $(5,2)$ are not in X.

Preference relations. Once we have described what the agent *can choose* (i.e., the set X) we want to be able to describe how the agent ranks the different alternatives in the set. For that purpose, we introduce an *order* on the set of alternatives, that is, a binary relation denoted by $\succeq$. Formally, $\succeq$ allows the comparison of pairs of alternatives. Take two alternatives $x, y \in X$, we read $x \succeq y$ as "x is at least as good as y".

From the relation $\succeq$ we can derive two other relations on X: the *strict preference* relation $\succ$ and the *indifference* relation $\sim$. We read $x \succ y$ as "x is preferred to y" and $x \sim y$ as "x is indifferent to y". Formally, $\succ$ is defined by

$$x \succ y \Leftrightarrow x \succeq y \text{ but not } y \succeq x.$$

That is, x is preferred to y if and only if x is at least as good as y but y is *not* at least as good as y. The indifference relation $\sim$ is defined by

$$x \sim y \Leftrightarrow x \succeq y \text{ and } y \succeq x.$$

Important

The relation $\succeq$ is defined for a given agent an is *tied* to their set of alternatives. Later, we will have several agents $i = 1, 2, \ldots$ interacting and our theory will allow for each of them to have their *own* preferences $\succeq_i$ on a set of alternatives X.

The preference relation is said to be *rational* when it satisfies two properties: *Completeness* and *Transitivity.*

Completeness means that the agent is able to compare any two pairs of alternatives with $\succeq$.

$$\textbf{Completeness:} \ \succeq \ \text{ is complete if for all } x, y \in X, x \succeq y \text{ and/or } \ y \succeq x.$$

It means that for any two alternatives x and y, the agent must be able to tell whether they strictly prefer x to y, y to x, or if they are indifferent between x and y.

Transitivity is a subtler concept as it imposes some relation between three alternatives.

$$\textbf{Transitivity:} \ \succeq \ \text{ is transitive if for all } x, y, z \in X, x \succeq y \text{ and } \ y \succeq z, \text{ then } \ x \succeq z.$$

Example

Let $X = \{\text{apple, pear, orange}\}$. Transitivity implies that if apple is at least as good as pear, and pear is at least as good as orange, then *it must also be true* that apple is at least as good as orange. In other words, the relation between apple and pear and that between pear and orange imposes a relation between apple and orange.

Transitivity imposes some structure on the relation preference and forbids that the agent's preference *cycles.* A cycle is when for instance $x \succeq y$, $y \succeq z$, and $z \succeq x$. This would imply that $x \succeq y \succeq z \succeq x \succeq y \succeq \ldots$

Exercise 1.1.1 *Show that if $\succeq$ is complete and transitive it is also reflexive (reflexiveness means that for all $x \in X$, $x \succeq x$).*

Exercise 1.1.2 *Show that if $\succeq$ is complete and transitive then $\succ$ and $\sim$ are transitive.*

1.2. Utility

Preference relations are technically rich enough to derive all the results of this course (and beyond) as they entirely capture individual preferences and rationality. However, they can sometimes be difficult objects to manipulate or to interpret. For convenience, we usually like to represent a preference relation by the means of a *utility function.*

A utility function is a function $u : X \mapsto \mathbb{R}$ that attaches a number to each element of the set of alternatives X. We say that $u(x)$ is the *utility* or the *utility score* of alternative $x \in X$. By convention, we define the utility function such that if an alternative x is preferred to another alternative y, it must have a higher utility score.

Formally, we say that $u : X \mapsto \mathbb{R}$ represents the preference relation $\succeq$ if, for all $x, y \in X$,

$$x \succeq y \Leftrightarrow u(x) \geqslant u(y).$$

In other words, the utility function allows us to compare alternatives only by checking their respective scores.

Exercise 1.2.1 *Show that $x \sim y$ implies $u(x) = u(y)$ and that $x \succ y$ implies $u(x) > u(y)$.*

Example

Let $X = \{\text{apple, pear, orange}\}$. If we want to represent the preference relation $\succeq$ such that apple $\succeq$ pear, pear $\succeq$ orange, and apple $\succeq$ orange we can choose u such that $u(\text{apple}) = 10$, $u(\text{pear}) = 5$, and $u(\text{orange}) = 2$.

Notice however, that the utility representation is *not unique.* For the same preference relation $\succeq$ there can be an infinite number of utility functions that represents it. Assume u represents $\succeq$, then any $v = f \circ u$ where $f : \mathbb{R} \mapsto \mathbb{R}$ is a strictly increasing function also represents $\succeq$. We say that utility functions are *invariant under monotonic transformations.*

Exercise 1.2.2 *Show that $v = f \circ u$ indeed represents $\succeq$.*

Example

Let $X = \{\text{apple, pear, orange}\}$ and assume that $\succeq$ is the same as in the previous example. Define $f(x) = 2x$, then $v(x) = f(u(x)) = 2u(x)$ so that $v(\text{apple}) = 20$, $v(\text{pear}) = 10$, and $v(\text{orange}) = 4$ also represents $\succeq$ as the utility scores have the same "ranking".

Exercise 1.2.3 *Find another strictly increasing function f such that $v = f \circ u$ represents $\succeq$ of the previous example.*

The fact that the utility representation of preferences is not unique conveys an important message: Only the ranking of utility scores matters, their absolute value does not mean anything *per se.* If the utility score of x is 2 and that of y is 1, we can only conclude that $x \succeq y$ and not that x is twice as good as y. We say that the utility representation is an *ordinal* concept, as opposed to a *cardinal* concept.

Important

While the utility representation may seem less *realistic* than the preference relation as it requires the agent to score each alternative (which can prove difficult in practice), the two concepts are formally strictly equivalent. Agents may not have a utility function *for real*, but we can do *as if* they had one in order to simplify our computations, without loss of generality.
The real questionable assumption is not that of the existence of an utility function, but that of rationality. It is not always obvious that an agent's preference relation is complete or transitive in practice and many empirical inconsistencies have been document.

1.3. Examples of preferences

Preferences can be chosen so as to represent meaningful situations we are interested in studying. In the case in which alternatives are consumption bundles, preferences over these bundles describes how different goods are related to each others. We now introduce two important concepts in microeconomics: *complements* and *substitutes*. Roughly speaking, some goods are complement when it is somewhat better to consume them together and they are substitute when it is somewhat easy to replace one by the other.

Perfect complements. Assume that an agent can choose how many cups of coffee and small sugar packets they order. We further assume that our agent only likes a cup of coffee when it contains exactly one sugar packet. In that case, we say that cups of coffee and sugar packets are *perfect complements*. Finally, we assume that our agent is always happy to consumer more coffee with sugar. How can we represent their preferences?

We can let $X := \mathbb{N}^2$ where an element $(x_1, x_2) \in X$ means that the agent orders x_1 cup(s) of coffee and x_2 packet(s) of sugar. The preference relation $\succeq$ on X must translates the idea that for instance ordering 3 cups and 2 sugar packets is the same as ordering 2 cups and 2 sugar packets as the agent will not drink the cup without sugar. At the same time, the agent always like more coffee with sugar so that (2,2) must be preferred to (1,1) for instance.

Using the preference relation approach we can define $\succeq$ as follows: for $x = (x_1, x_2) \in X$ and $y = (y_1, y_2) \in X$, $x \succeq y$ if and only if $\min\{x_1, x_2\} \geqslant \min\{y_1, y_2\}$. We can describe preferences also with the following utility function: $u(x_1, x_2) = \min\{x_1, x_2\}$. This utility function is often call a *Leontief utility function.*

Exercise 1.3.1 *Assume the agent needs not one but two sugar packets in each cup of coffee to enjoy it. Find a utility representation of these preferences.*

Perfect substitutes. The agent can now choose between cups of coffee and cups of tea. We assume that the agent *equally* likes coffee and tea. Once again, we assume that the agent is always better off with more cups than less.

We let $X := \mathbb{N}^2$ where an element $(x_1,x_2) \in X$ means that the agent orders x_1 cup(s) of coffee and x_2 cup(s) of tea. To translate the idea that the agent equally likes coffee and tea, the preference relation $\succeq$ must be such that what matters is only the number of cups the agent orders, regardless of whether they contain coffee or tea. For instance, the agent must prefer $(1,2)$ to $(0,2)$ and must be indifferent between $(3,1)$ and $(2,2)$.

The preference relation in that case is defined as follows. For $x = (x_1,x_2) \in X$ and $y = (y_1,y_2) \in X$, $x \succeq y$ if and only if $x_1 + x_2 \geqslant y_1 + y_2$. A utility function representing these preferences can simply be $u(x_1,x_2) = x_1 + x_2$.

Imperfect substitutes. Most of the time, there exists some substitutability between goods but it is not perfect as in the previous case. That is, the agent may accept to substitute some quantity of good x to good y but still exhibit some strict preference for good x. In that case goods are only *imperfect substitutes.*

Assume as in the previous example that the agent can order cups of coffee and cups of tea. However, the agent has a preference for coffee over tea in the sense that one cup of coffee is *worth* two cups of tea. Tea and coffee are still substitutes as one can be used to replace the other, but it is now imperfect substitutability as replacing one unit of coffee requires twice as much units of tea.

Once again, let $X := \mathbb{N}^2$ and $(x_1,x_2) \in X$ means that the agent orders x_1 cup(s) of coffee and x_2 cup(s) of tea. The preference relation in that case must be such that for $x = (x_1,x_2) \in X$ and $y = (y_1,y_2) \in X$, $x \succeq y$ if and only if $2x_1 + x_2 \geqslant 2y_1 + y_2$. A utility function to represent these preferences can be $u(x_1,x_2) = 2x_1 + x_2$.

Notice that the utility functions $v(x_1,x_2) = x_1 + 0.5x_2$ or $w(x_1,x_2) = 100x_1 + 50x_2$ also represents these preferences. More generally, any $u(x_1,x_2) = ax_1 + bx_2$ where $a/b = 2$ represents the same preferences.

A more complex case. Let us combine our previous preferences describing complements and substitutes to build a more complex preference representation.

Assume that the agent can order cups of coffee, cups of tea, and sugar packets. For convenience let us denote by $(c,t,s) \in \mathbb{N}^3$ the vector of consumed quantities of coffee, tea, and sugar, respectively. We make the following assumptions. First, coffee and tea can be substituted but the agent likes coffee twice as much as tea. Second, the agent only likes coffee or tea when there is exactly one sugar packet poured into the cup.

The set of alternatives writes $X := \mathbb{N}^3$ and we can represents the agent's preferences with the following utility function: $u(c,t,s) = 2\min\{c,s\} + \min\{t,s\}$. We can check with examples that we correctly represented the agent's preferences. For instance, we should expect that the agent prefers $(2,0,3)$ to $(1,1,3)$ as in both cases the agent can consume two cups with sugar but the first bundle is better because both cups are coffee cups and the agent prefers coffee to tea. Indeed, $u(2,0,3) = 4 > u(1,1,3) = 3$. We should also expect that $(1,0,5)$ is indifferent to $(0,2,7)$ as in the first case the agent can only consume one cup of coffee with sugar while in the second they can consumer two cups of tea with sugar. Indeed, $u(1,0,5) = 2 = u(0,2,7) = 2$.

Substitutes and complements. Sometimes two goods can also be both substitutes and complements at the same time. This situation occurs when an agent would accept to replace some units of good x with some units of good y but also attributes some value to the joint consumption of the two goods.

A classical example is that of burger and fries. Assume that the agent cannot imagine ordering a burger without fries, nor fries without a burger. However, as long as both the burger and fries are present, the agent accepts to receive less (resp. more) fries if the burger is bigger (resp. smaller). In that case, the two goods are indeed substitutes but they are also complements.

Let the vector $(x_1, x_2) \in X = \mathbb{R}^2$ be an alternative where x_1 is the *burger size* and x_2 the *number of fries*. A simple utility function to represent the agent's preferences is for instance $u(x_1, x_2) = x_1 x_2$. We can see that there is some complementarity between the goods as for instance the bundle $(5,5)$ is preferred to $(15,1)$ as the first yields a utility of 25 while the second of 15. In words, the agent prefers a smaller burger with a decent amount of fries rather than a bigger burger with very few fries. However, it is still possible to substitute the two goods: for instance $(25,1)$ gives the same utility as $(5,5)$ so that a huge burger can at some point compensate the lack of fries.

Exercise 1.3.2 *Assume that for the agent to consider the meal to be decent, any serving with a burger size less than $a > 0$ and/or an amount of fries less than $b > 0$ is perceived as the same as the same as the zero bundle $(0,0)$. Find a utility function that represents these preferences.*

1.4. Monotonicity and convexity of preferences

On top of the rationality assumption – preferences are complete and transitive – we generally make two additional key assumptions on preferences.

Monotonicity. We already implicitly used the notion of *monotonicity* of preferences in some of the examples above. This assumption simply ensures that preferences are such that *more* is always better than *less*: three oranges is better than two oranges, which is better than one orange.

Formally, preferences are (weakly) *monotonic* if $x, y \in X = \mathbb{R}^m_+$ such that $x_i > y_i$ for all $i = 1, \ldots, m$ implies that $x \succ y$.

Important

The monotonicity assumption is arguably quite strong: While it seems reasonable to say that two oranges are better than one, it is much less reasonable that one million oranges is preferred to two oranges. If you were to receive one million oranges you would have to store them or to find a way to dispose of them, which would entail additional costs or inconveniences.

In advanced microeconomics, we usually use a weaker condition – local nonsatiation – that allows for some level of satiation of preferences.

Convexity. The convexity assumption reflects the idea that an agent has an inclination for *diversification,* that is, the agent has some taste for *variety.* Loosely speaking, the agent prefers to consume a bit of everything rather than a huge amount of a restricted set of commodities.

Formally, we say that preferences are convex if $x \succeq z$ and $y \succeq z$, then $\alpha x + (1-\alpha)y \succeq z$ for any $\alpha \in [0,1]$. In words, if we take two alternatives x and y that are both at least as good as alternative z, then any convex combination (i.e. mix) of x and y is also at least as good as z. For instance, it means that if the agent prefers 3 oranges to 2 apples and 2 pears to 2 apples, then the agents also prefers a mix of 1.5 oranges and 1 apple to 2 pears.

Importantly, preferences need not to satisfy monotonicity and convexity to be rational. We make these two additional assumptions – monotonicity and convexity – as they are seen as mostly realistic and make the consumption problem tractable even without adding much more structure.

1.5. Marginal utility, indifference curves, and marginal rate of substitution

We now introduce three key important concepts in microeconomics: Marginal utility, indifference curves, and marginal rate of substitution. As we will mostly focus on the case in which the set of alternatives is the consumption set in the case of two commodities we will set $X := \mathbb{R}^2_+$.

We will further assume that preferences are *continuous*, which we will define as preferences that can be represented by a continuous utility function u. We will also generally assume that the utility function is *differentiable.*

Marginal utility. The notion of marginal utility captures the additional utility the agent derives from the increase in consumption of one good while the consumption of the other goods is held constant. In general, it is defined as the change of total utility divided by the change in number of units.

Assume that we start from the consumption bundle (x_1, x_2) and that we want to evaluate the marginal utility of consuming $(x_1 + \varepsilon, x_2)$, where $\varepsilon > 0$. The marginal utility of an ε-increase of good 1 starting at (x_1, x_2) is given by

$$\frac{u(x_1 + \varepsilon, x_2) - u(x_1, x_2)}{\varepsilon}.$$

Notice that if $\varepsilon = 1$, that is, we consider an increment of exactly one unit of good 1, the marginal utility is simply given by the difference between the utility after and before consuming the additional unit.

In general, we assume that the utility function is differentiable so that the concept of marginal utility coincides exactly with that of partial derivatives. Indeed, the above formula corresponds to the definition of the partial derivative of u with respect to x_1 when we let $\varepsilon \to 0$. From now on, we will define the marginal utility of good i at consumption bundles x as follows:

$$\text{marginal utility: } \frac{\partial}{\partial x_i} u(x).$$

Important

The marginal utility is a *local* concept. In general, we cannot compute what is the marginal utility of an increase in good 1 without knowing what the agent is currently consuming. That is, the marginal utility of an increase in good 1 is in general different when evaluated at different consumption bundles.

Example

1. Assume that $m = 2$, let $u(x_1, x_2) = x_1 + x_2$. In this case the marginal utility of good 1 is simply $\frac{\partial}{\partial x_1} u(x) = 1$. Noticeably, it does not depend on the consumption bundle we are evaluating it and is always equal to 1. The same holds for the marginal utility of good 2.

2. Let $u(x_1, x_2) = x_1 x_2$. We can compute the marginal utility of good 1 simply as $\frac{\partial}{\partial x_1} u(x) = x_2$. Contrary to the previous example, the marginal utility now entirely depends on the value of x_2. It follows that the marginal utility of good 1 at $(4, 3)$ is equal to 3 but is equal to 2 at $(1, 2)$. Hence, evaluating the marginal utility at different consumption bundles gives different values.

Marginal rate of substitution. Using the concept of marginal utility of a good, we can define a second very useful tool, namely the *marginal rate of substitution* (hereafter "MRS"). The MRS represents the rate at which a consumer is willing to trade one good for another while keeping their utility constant.

We define the MRS between good i and good j as follows:

$$MRS_{i,j}(x) := -\frac{\frac{\partial u(x)}{\partial x_i}}{\frac{\partial u(x)}{\partial x_j}},$$

that is, the MRS between good i and good j is simply the ratio of the marginal utility of good i and that of good j.

To better understand the MRS, it is useful to explicitly construct it in the case $m = 2$. We want to know, starting from (x_1, x_2), how many units of good 2 the agent is willing to give up to obtain one additional unit of good 1, while keeping their utility constant. Mathematically, we can simply take the total derivative of $u(x_1, x_2) = \overline{u}$, that is:

$$du(x_1, x_2) = \frac{\partial u(x)}{\partial x_1} dx_1 + \frac{\partial u(x)}{\partial x_2} dx_2 = 0.$$

Rearranging the above equation we can get that

$$\frac{dx_2}{dx_1} = -\frac{\frac{\partial u(x)}{\partial x_i}}{\frac{\partial u(x)}{\partial x_j}}.$$

Notice that setting $dx_1 = 1$, the above equation tells us what should be the change in good 2 (dx_2) that keeps the utility level constant when the agent consumes one more unit of good 1 ($dx_1 = 1$).

Important

Like the marginal utility, the MRS is a *local* concept and depends on the consumption bundle at which it is evaluated.

Example

1. Assume that $m = 2$, let $u(x_1, x_2) = x_1 + x_2$. The MRS between good 1 and 2 is given by $MRS_{1,2}(x) = -1$. In this particular case, it is constant and simply says that the agent needs to receive one unit of good 1 to compensate for the loss of one unit of good 2 in order to keep their utility constant.

2. Let $u(x_1, x_2) = x_1 x_2$. In that case $MRS_{1,2}(x) = -x_2/x_1$ and therefore the value of the MRS depends on where we evaluate it. For instance, $MRS_{1,2}(4,2) = -1/2$ and $MRS_{1,2}(6,2) = -1/3$.

Indifference curves. A central tool of microeconomic theory is that of *indifference curves.* An indifference curve is a level curve (or level set) of a utility function: namely, a set that contains all the elements of X that give the same utility to the agent.

Formally, we start by defining the indifference set of utility u at utility level $\overline{u}$ by

$$L_{\overline{u}}(u) := \{x \in X \mid u(x) = \overline{u}\},$$

that is, a set that contains all the alternatives that give the agent the same utility $\overline{u} \in \mathbb{R}$. Notice that $L_{\overline{u}}(u)$ can be empty, or contain only one bundle in some cases. Naturally, if we take the collection $(L_{\overline{u}}(u))_{\overline{u}\in\mathbb{R}}$ we retrieve the set of all possible alternatives X. Finally, it is worth stressing that if $L_{\overline{u}}(u) \cap L_{\overline{u}'}(u) = \varnothing$ for any $\overline{u} \neq \overline{u}'$, that is, each consumption bundle belongs to only one indifference set.

In the case of two commodities, $X := \mathbb{R}^2_+$ with continuous and monotonic preferences, we can easily represent the indifference curves graphically. Take utility level $\overline{u}$, we look for the consumption bundles such that $u(x_1, x_2) = \overline{u}$. From the monotonicity assumption we can always invert this function so as to get a well-defined function $x_2 = \varphi_{\overline{u}}(x_1)$. This function gives the relationship between x_1 and x_2 in order to keep the utility level constant at $\overline{u}$.

At this stage, it is more instructive to take examples of utility function and derive their indifference curves.

□ *Cobb-Douglas.* Let $u(x_1,x_2) = x_1x_2$. For a utility level $\bar{u}$, we simply solve $x_1x_2 = \bar{u}$ for x_2 and immediately obtain $x_2 = \frac{\bar{u}}{x_1}$ (i.e. we have $\varphi_{\bar{u}}(z) = \bar{u}/z$).

□ *Perfect substitutes.* Let $u(x_1,x_2) = x_1 + x_2$. For a utility level $\bar{u}$, we simply solve $x_1 + x_2 = \bar{u}$ for x_2 and immediately obtain $x_2 = \bar{u} - x_1$ (i.e. we have $\varphi_{\bar{u}}(z) = \bar{u} - z$).

□ *Perfect complements.* Let $u(x_1,x_2) = \min\{x_1,x_2\}$. In that case we have to be a bit more careful when deriving the indifference curves as this utility function is only *weakly* monotonic in each good. Indeed notice that even if we fix a x_1 and $\bar{u}$, the equation $\min\{x_1,x_2\} = \bar{u}$ does not always return a unique value for x_2. If $x_1 > \bar{u}$, then $x_2 = \bar{u}$ but if $x_1 = \bar{u}$ then any $x_2 \in [\bar{u}, +\infty)$ satisfies the equation.

The indifference curves for each of the three examples above are represented in Figure 1. Along each curve, the utility level is constant. Indifference curves that are *higher* correspond to higher levels of utility.

MRS and indifference curves. Importantly, we can connect the MRS with the indifference curves. At any point (x_1,x_2), the slope of the indifference curve is equal to the MRS at this point.

Take the Cobb-Douglas case $u(x_1,x_2) = x_1x_2$. At utility level $\bar{u}$, the slope of the indifference curve is given by $\frac{\partial}{\partial x_1}(\frac{\bar{u}}{x_1}) = -\frac{\bar{u}}{x_1^2}$. Notice that at point (x_1,x_2), the indifference curve yields that $x_2 = \frac{\bar{u}}{x_1}$, so that we can rewrite the slope at this point as $-\frac{\bar{u}}{x_1^2} = -\frac{\bar{u}}{x_1}\frac{x_2}{\bar{u}} = \frac{-x_2}{x_1} = MRS_{1,2}$.

The reason why the slope and the MRS coincide is simple. Moving along an indifference curve amounts to substitute one good to another while keeping the utility constant which is precisely what the MRS is measuring in the first place: how many unit of good 2 should I give up to get one more unit of good 1 and be indifferent about it.

2. Consumer choice

We now turn to develop a theory of how agents make their consumption choice. From now on, we will refer to the set of alternatives X as the *consumption set* and define it as follows:

$$X := \{(x_1,\dots,x_m) \in \mathbb{R}^m \mid x_1 \geqslant 0,\dots,x_m \geqslant 0\},$$

or equivalently $X := \mathbb{R}^m_+$. This set contains all the *consumption bundles* of $m > 0$ goods. Notice that we impose that the quantity consumer of each good must be nonnegative. In this environment, a preference relation $\succeq$ or a utility function describes how an agent ranks the consumption bundles of their consumption set.

Example

Assume that $m = 2$, then the set of consumption bundles is $X = \mathbb{R}^2_+$ and an element $(x_1,x_2) \in X$ corresponds to the agent consuming x_1 unit of good 1 and x_2 units of good 2.

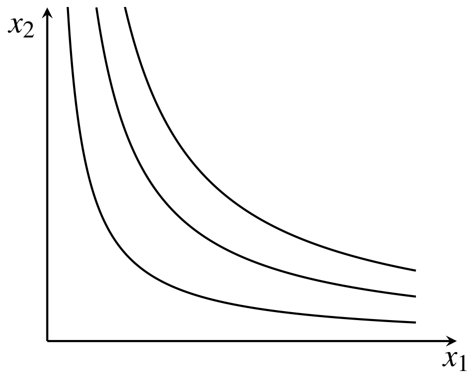

(a) Case: $u(x_1, x_2) = x_1 x_2$.

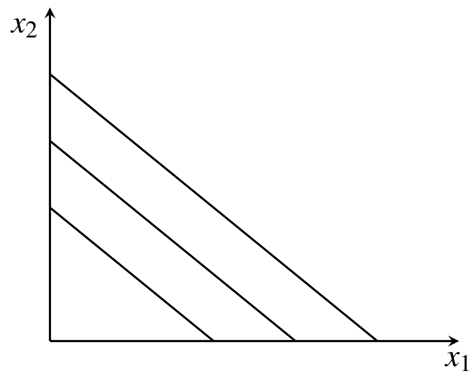

(b) Case: $u(x_1, x_2) = x_1 + x_2$.

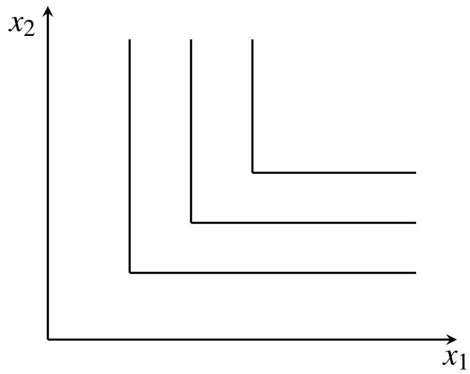

(c) Case: $u(x_1, x_2) = \min\{x_1 + x_2\}$.

Figure 1: Indifference curves.

Notice that so far, we have described how the agent ranks alternatives in their set of alternatives X as if all alternatives where *feasible*. That is, as if the agent could always freely pick any $x \in X$ as they wish. In this *unconstrained* scenario, determining what is the agent's optimal choice is rather trivial: The agent's optimal choice simply consists in choosing an alternative $x^* \in X$ such that $x^* \succeq y$ for all $y \in X$. Using the utility representation, this statement is equivalent to say that the agent's optimal choice is the alternative $x^* \in X$ such that $u(x^*) \geqslant u(y)$ for all $y \in X$, that is, $x^* \in \arg\max_x u(x)$.

In practice, agents are subject to constraints which may prevent them to pick some of the elements in the set of alternatives. As our focus is consumer theory, our agents will typically face the following two types of constraints.

Resource constraints. There is naturally a finite amount of resources available to consume meaning that no agent can consume more units of a good than it exists in the economy. While important in general, we will mostly ignore this type of constraints in this course as the second type of constraints will prevail.

Budget constraints. In our theory of consumption, agents have to *buy* goods in order to consume them. Therefore, the consumption bundles available to an agent will ultimately depend on the price of each good as well as the agent's wealth/income/budget.

2.1. Budget set, prices, and wealth

We assume that each good $i = 1,\dots,m$ can be bought be the agent at unit price $p_i \in \mathbb{R}_+$. We will sometimes refer to $p := (p_1,\dots,p_m)$ as the price vector. For instance, if the agent wants to consumer x_i units of good i with price p_i they will have to spend $p_i x_i$ euros.

Additionally, an agent has a limited amount of money to spend. We denote this amount by $w \in \mathbb{R}_+$ and we refer to it as the agent's *wealth*.

Important

It is assumed that both prices and income are *given parameters*. Prices are beyond the influence of the agents, that is, none of the agents can affect prices through their behavior. We say that agents are *price-takers*. A similar assumption applies to the agents' incomes.

Budget set. For a given price vector p and a given wealth w, we define the budget set B as the set of *affordable* consumption bundles. Formally, the budget set is defined by:

$$B := \left\{ (x_1,\dots,x_m) \in X \;\middle|\; \sum_{i=1}^{m} p_i x_i \leqslant w \right\}.$$

That is, the budget set B is a subset of the consumption set X that contains only consumption bundles that cost at most as much as the total wealth of the agent. Figure 2 represents the set of affordable bundles in the case of two goods.

Example

Assume that $m = 2$, $p_1 = 2$, $p_2 = 4$ and $w = 10$. Then, any consumption bundle (x_1, x_2) such that $2x_1 + 4x_2 \leqslant 10$ belongs to the budget set B. For instance, $(1,1)$ and $(5,0)$ respectively cost 6 and 10 euros and therefore belong to the budget set. On the contrary, $(3,2)$ and $(1,5)$ do not belong to the budget set as they cost 14 and 22, respectively.

If, for some reason, the price of good 2 were to change to $p_2 = 2$, then the bundle $(3,2)$ would become affordable and therefore belong to the budget set.

The agent's optimal choice must therefore lie in the budget set B. That is, the agent must choose their preferred consumption bundle among the ones that are affordable according to prices and their wealth.

Price and wealth changes. It is important to stress that the budget set is defined for a given price vector p and level of wealth w. Sometimes we may want to write it as $B_{p,w}$ so as to emphasize this dependency. Here we briefly explore how the budget set varies with prices and wealth in the case $m = 2$.

First, it is useful to detail how we proceeded to obtain Figure 2. Recall that the budget set is characterized by the consumption bundles such that $p_1x_1 + p_2x_2 \leqslant w$. As we want to represent it in

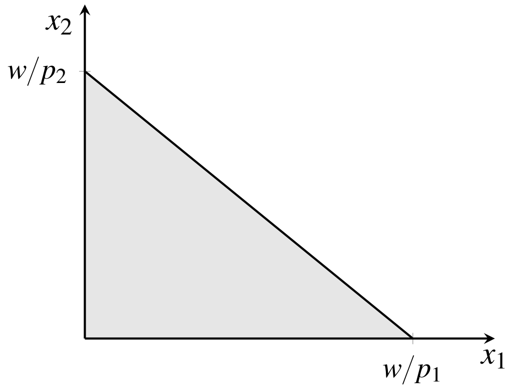

Figure 2: Representation of the budget set (gray area) in the case $m = 2$.

the (x_1, x_2) plane, we need to rearrange this expression so as to solve for x_2. That is, $x_2 \leqslant \frac{w}{p_2} - \frac{p_1}{p_2}x_1$. In practice, we like to work with the equation $x_2 = \frac{w}{p_2} - \frac{p_1}{p_2}x_1$ that characterizes the consumption bundles that are on the boundary of the budget set B, that is, the consumption bundles such the agent spends all their wealth.

To understand how changes in price or wealth affect the budget set, it is instructive to further interpret the equation $x_2 = \frac{w}{p_2} - \frac{p_1}{p_2}x_1$. Notice first that if the agent does not consume any good 1, $x_1 = 0$, they can consume *at most* $x_2 = \frac{w}{p_2}$ units of good 2. Similarly, if they consume zero unit of good 2, $x_2 = 0$, they can consume *at most* $x_1 = \frac{w}{p_1}$ units of good 1.

Example

Assume $p_1 = 2$, $p_2 = 5$, and $w = 20$. If the agent consumes only good 1, they can buy at most $w/p_1 = 10$ units of it. If instead the agent consumes only good 2, they can buy at most $w/p_2 = 4$ units of it.

The slope of the budget set's boundary, $-\frac{p_1}{p_2}$, has important economic interpretations. The fact that is it negative first indicates that if the agent wants to consume more of one of the good, they have to consume less of the other good. How much one good must be exchanged with the other depends on the *price ratio* $\frac{p_1}{p_2}$. Indeed, if the two goods have the same price $p_1 = p_2$ then the price ratio is one, meaning that consuming one more unit of good 1 requires to consume exactly one less unit of good 2 to keep the spending constant. Instead, if for instance $p_1 = 0.5p_2$, consuming one more unit of good 1 requires to consume two less units of good 2 as good 1 is twice as expensive as good 1.

Important

Looking at how goods should be exchanged to keep the spending constant has a similar "flavor" to the concept of "substitutability" between goods that we have seen before. The price ratio plays an similar role as that of the preferences in the sense that it indicates how easy/hard it is to exchange one good to another.

Be careful, though, despite their similarities the two concepts are distinct. One is about intrinsic preferences – how much the agent likes one good with respect to another – while the other is about economic constraints – how much the agent has to deprive themselves of the consumption of one good to be able to consume more of another without spending more money.

When solving for the agent's consumption problem we will see how the two concepts interact in an intuitive way.

Using the equation $x_2 = \frac{w}{p_2} - \frac{p_1}{p_2}x_1$, we can study how a change in price or wealth affects the budget set. For clarity, let us explicit denote by $B_{p,w}$ the budget set associated with prices p and wealth w. Consider a change in the price of good 1, $\tilde{p}_1 > p_1$. What changes is the slope of the budget line, it becomes steeper as $-\frac{\tilde{p}_1}{p_2} < -\frac{p_1}{p_2}$ and the new budget set $B_{\tilde{p},w}$ shrinks as represented in Figure 3. Instead, consider a change in wealth $\tilde{w} > w$, then the budget line shifts downwards as represented in Figure 4.

In general, a price increase in one or several goods or a decrease in wealth makes the budget set *smaller*, in the sense that $B_{\tilde{p},\tilde{w}} \subseteq B_{p,w}$ if either $\tilde{w} \geqslant w$ or $\tilde{p}_i \geqslant p_i$ for some $i = 1, \dots, m$. On the contrary, if some prices decrease or the agent's wealth increases the budget set *expands*. If some price go up and other go down, how the budget set changes is more complex as some consumption bundles may become affordable while others may become too costly and there is no clear-cut way to describe the change in general.

Relative prices. An important thing to notice is that our theory suggests that only *relative prices* are relevant to determine how easy it is to exchange one good to another and satisfy the budget constraint. Too see that, assume $m = 2$ and start with prices p_1, p_2 and wealth w. Further assume that both prices are multiplied by a constant $\alpha > 0$, such that the new budget constraint writes $\alpha p_1 x_1 + \alpha p_2 x_2 \leqslant w$. This inequality is equivalent to $p_1 x_1 + p_2 x_2 \leqslant \frac{w}{\alpha}$. Hence, the case in which both prices increase in the same proportion is strictly equivalent to the case in which prices are unchanged but the agent's wealth has decreased to $\tilde{w} = w/\alpha$. In other words, the overall price increase has effectively made the agent *poorer* but as the price ratio $\frac{p_1}{p_2} = \frac{\alpha p_1}{\alpha p_2}$ remained the same, it did not change the relationship between the two goods with respect to their relative costs.

Homogeneity. Another interesting property of the budget set is what happens when all prices and wealth are multiplied by the same constant $\alpha \in R_+$, so that $\tilde{p}_1 = \alpha p_1$, $\tilde{p}_2 = \alpha p_2$ and $\tilde{w} = \alpha w$. The budget set $B_{\tilde{p},\tilde{w}}$ is characterized by $\alpha p_1 + \alpha p_2 \leqslant \alpha w$ which is strictly equivalent to $p_1 x_1 + p_2 x_2 \leqslant w$,

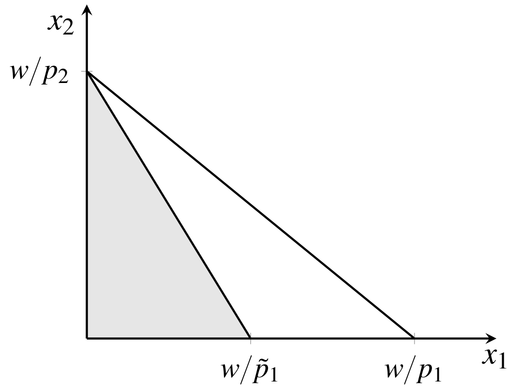

Figure 3: Budget set change when $\tilde{p}_1 > p_1$.

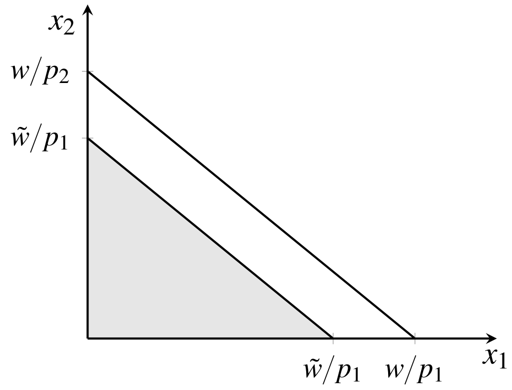

Figure 4: Budget set change when $\tilde{w} > w$.

that is, to $B_{p,w}$ the initial budget set. This shows that the *nominal* value of prices and wealth are irrelevant – if all prices and wealth are multiplied by 2, 3 or 10, the budget set remains unchanged and the exact same consumption bundles are available to the agent.

2.2. *The consumer's program*

Now that we have introduced how to model preferences and the budget set of the agent, we can turn to the agent's optimal choice. We say that a consumption bundle $x^* \in X$ is an *optimal choice* for the agent if (i) this bundle is feasible, that is, it belongs to the budget set $x^* \in B$, and (ii) there is no other feasible bundle that is strictly preferred to x^*. In other words, among all the affordable consumption bundles the agent's optimal choice consists in choosing one that is *better* than all others.

Optimal choice. More formally, we can define an optimal choice as an $x^* \in B$ such that there exists no $x \in B$ satisfying $x \succeq x^*$. In utility terms, we can write that $x^* \in B$ is such that $u(x^*) \geqslant u(x)$ for all $x \in B$. We can also write it in a more compact way, that is, $x^* \in \arg\max_{x \in B} u(x)$. All three formulations are strictly equivalent.

Example

Assume $m = 2$, $p_1 = 1$, $p_2 = 2$, and $w = 2$. Preferences are such that $u(x_1, x_2) = x_1 + x_2$ and $X := \{(0,0),(0,1),(1,0),(1,1)\}$. First notice that $B = X \backslash \{(1,1)\}$ as $(1,1)$ costs $p_1 + p_2 = 3 > w = 2$. Clearly, the two affordable consumption bundles $(0,1)$ and $(1,0)$ gives the agent a utility of 1 while $(0,0)$ gives 0 utility. Hence both $(0,1)$ and $(1,0)$ could be the agent's optimal choice as they are both affordable and *better* than all other affordable bundles. Without any doubt, $(1,1)$ is the *best* bundle as it yields a utility of 2 but unfortunately for the agent, it is not affordable.

"Intuitive approach". The above example is arguably simple to solve thanks to the limited number of consumption bundles we have to compare to find the agent's optimal choice. When we face a *larger* consumption set such as $X := \mathbb{R}^2_+$, the problem may become much harder to solve. We attempt to solve a particular problem using an "intuitive" approach.

Take the following example: $m = 2$, $p_1 = 1$, $p_2 = 1$, and $w = 4$. Let $X := \mathbb{R}^2_+$ and $u(x_1,x_2) = x_1x_2$. We cannot conceivably compute the utility of each consumption bundle as they is an infinite number of bundles, even when restricting ourselves to the budget set. Let us therefore try to reduce the complexity of the problem.

First, the budget set is such that only bundles satisfying $x_1 + x_2 \leqslant 4$ have to be taken into account. Second, notice that we can restrict our search to bundles such that the agent spends all their wealth, that is, $x_1 + x_2 = 4$. Indeed, assume that we start from $(\tilde{x}_1, \tilde{x}_2)$ such that $\tilde{x}_1 + \tilde{x}_2 < 4$. This bundle yields utility $\tilde{x}_1\tilde{x}_2$ and is always worse than bundle $(\tilde{x}_1 + \varepsilon, x_2)$ where $\varepsilon > 0$. If we set $\varepsilon = 4 - (\tilde{x}_1 + \tilde{x}_2) > 0$, bundle $(\tilde{x}_1 + \varepsilon, x_2)$ is strictly preferred and is such that the agent spends all their wealth. Hence, we know that in any case $x_2 = 4 - x_1$ and we can reduce the problem to finding a $x_1 \leqslant 4$ such that $u(x_1, 4 - x_1) = x_1(4 - x_1)$ is maximum.

Exercise 2.2.1 *Find the unique optimal choice of the consumer in the above example.*

Constrained optimization problem. We now introduce the standard methodology to solve for the agent's optimal choice. Not only it provides a systematic solution to the agent's problem but it also uncovers some general characterization results whose economic interpretation is of importance.

In this part we assume that the utility function is differentiable and concave, that $p >> 0$, i.e. all prices are strictly positive, and that $w > 0$. The standard maximization program that we aim at solving writes as follows:

$$\begin{aligned} \max_{x \in X} \quad & u(x) \\ \text{s.t.} \quad & \sum_{i=1}^{m} p_i x_i \leqslant w. \end{aligned}$$

We can establish the Lagrangian of this problem as $\mathscr{L}(x,\lambda) := u(x) + \lambda(w - \sum_{i=1}^m p_i x_i)$, where $\lambda \geqslant 0$ is the Lagrange multiplier associated with the budget constraint. The KKT first-order condition necessary for a maximum point are:

$$\frac{\partial \mathscr{L}(x,\lambda)}{\partial x_i} = 0, \text{ for all } i = 1, \ldots, m,$$

together with the complementary slackness condition $\lambda(w - \sum_{i=1}^m p_i x_i) = 0$.

Important

The first-order conditions are also sufficient in our case as the utility function is concave and the constraint set is convex.

The first-order conditions explicitly write as:

$$\frac{\partial u(x)}{\partial x_i} = \lambda p_i, \text{ for all } i = 1, \ldots, m.$$

From our assumption that u is monotonic in each x_i, we know that the marginal utility must alwways be strictly positive, that is $\frac{\partial u(x)}{\partial x_i} > 0$. This immediately implies that $\lambda p_i > 0$ and therefore that $\lambda > 0$. Hence, from the complementary slackness condition we must have $w - \sum_{i=1}^{m} p_i x_i = 0$, that is, the agent must spend all their wealth. This is our first important finding.

Now, take the first-order condition for good i and that for good j, that is, $\frac{\partial u(x)}{\partial x_i} = \lambda p_i$ and $\frac{\partial u(x)}{\partial x_j} = \lambda p_j$. Simply divide the first equation by the second and we immediately obtain that

$$\frac{\frac{\partial u(x)}{\partial x_i}}{\frac{\partial u(x)}{\partial x_j}} = \frac{p_i}{p_j}, \text{ for all } i, j.$$

On the left-hand side we recognize the marginal rate of substitution between i and j, $MRS_{i,j}(x)$. On the right-hand side we have the price ratio between good i and good j. Hence, this equation tells us that at the optimal consumption bundle, the MRS of any two goods i and j must equal their price ratio.

Recall that we previously observed that while the MRS is providing us with a measure of substitution between two goods with respect to preferences, the price ratio was instead providing a measure of substitution in terms of maintaining and equal budget. What the above equation is telling us is that the agent is balancing out these two effects at their optimal consumption choice.

Example

Assume $u(x_1, x_2) = x_1 x_2$ where $X = \mathbb{R}^2_+$.

Substitution method. We first provide an alternative method to solve the agent's problem and then check that it corresponds to our general result. We want to find $\max_{x \in X} u(x)$ s.t. $p_1 x_1 + p_2 x_2 \leqslant w$. We have already shown before that the agent must spend all their wealth, therefore we must have that $p_1 x_1 + p_2 x_2 = w \Leftrightarrow x_2 = \frac{w}{p_2} - \frac{p_1}{p_2} x_1$. We can therefore substitute this value of x_2 into the objective function and simply solve:

$$\max_{x_1 \in \mathbb{R}_+} u(x_1, \frac{w}{p_2} - \frac{p_1}{p_2} x_1) = x_1 (\frac{w}{p_2} - \frac{p_1}{p_2} x_1).$$

The first-order condition of this unconstrained problem immediately writes as:

$$\frac{w}{p_2} - \frac{p_1}{p_2} x_1 - x_1 \frac{p_1}{p_2} = 0.$$

Solving for x_1 yields that $x_1^* = \frac{w}{2p_1}$. Finally, plugging x_1^* into the budget constraint we obtain that $x_2^* = \frac{w}{2p_2}$.

Lagrangian method. Relying on our general result, we only have to compute the marginal utility in 1 and 2. That is, $\frac{\partial u(x)}{\partial x_1} = x_2$ and $\frac{\partial u(x)}{\partial x_i} = x_1$. We must therefore have that the ratio of MRS is equal to the price ratio:

$$\frac{x_2}{x_1} = \frac{p_1}{p_2}.$$

We can solve this equation for x_2 to obtain that $x_2 = \frac{p_1}{p_2} x_1$. Given that the budget constraint is satisfied with an equality we also have that $x_2 = \frac{w}{p_2} - \frac{p_1}{p_2} x_1$. Hence we must have that $\frac{p_1}{p_2} x_1 = \frac{w}{p_2} - \frac{p_1}{p_2} x_1 \Leftrightarrow x_1^* = \frac{w}{2p_1}$. Finally $x_2^* = \frac{p_1}{p_2} x_1^* = \frac{w}{2p_2}$.

We now turn to define some standard concepts in microeconomics.

Marshallian demand function. The Marshallian demand function of a good i is simply the quantity the agent is willing to buy given prices and their wealth level. It is therefore simply the solution to the agent's utility maximization program.

Formally, the vector of Marshallian demands is defined by $x(p, w) \in \arg\max_{x \in B(p,w)} u(x)$, where $x(p, w) = (x_1(p, w), \ldots, x_m(p, w))$. We explicitly denote each demand $x_i(p, w)$ as a function of price vector and wealth.

Example

In the previous example, when $u(x_1,x_2) = x_1x_2$, the Marshallian demand for good i is $x_i(p,w) = \frac{w}{2p_i}$.

Indirect utility function. The indirect utility function, usually denoted by $v(p,w)$, is the maximal utility level the agent can reach for a given price vector and level of wealth. Mathematically speaking, the indirect utility function is the *value function* of the agent's utility maximization problem:

$$v(p,w) := \begin{cases} \max_{x\in X} & u(x) \\ \text{s.t.} & \sum_{i=1}^m p_i x_i \leqslant w. \end{cases}$$

Another equivalent way to define the indirect utility function is via Marshallian demands:

$$v(p,w) := u(x_1(p,w),\dots,x_m(p,w)).$$

In words, the indirect utility is already taking into account the agent's optimization behavior and only reflects the best the agent can do given a price vector and a level of wealth. Notice that it is indeed a function that depends solely on prices and wealth and not on the quantities of good.

Example

Once again assume that $u(x_1,x_2) = x_1x_2$. We saw that the Marshallian demand for good i is $x_i(p,w) = \frac{w}{2p_i}$.
Hence, the indirect utility function is defined by

$$\begin{aligned} v(p,w) := u(x_1(p,w),x_2(p,w)) &= u\left(\frac{w}{2p_1},\frac{w}{2p_2}\right) \\ &= \frac{w}{2p_1}\frac{w}{2p_2} = \frac{w^2}{4p_1p_2}. \end{aligned}$$

Therefore, we can easily deduce how the agent's utility changes with price and wealth as long as the agent chooses opitmally.

3. The firm: Theory of production

We now turn to the supply side of the economy, that is, the study of the process of how goods are produced. In this part we will focus on the problem of *the firm*, the entity whose activity is to transform inputs – or factors of production – into output, i.e. quantity of some good.

3.1. Technology

The firm uses inputs $z := (z_1, \ldots, z_n) \in \mathbb{R}^n_+$ to produce a quantity $q \in \mathbb{R}_+$ of single output. The technology, or production function, f describes how each combination of inputs x transforms into a quantity of output q, that is, $q = f(z_1, \ldots, z_n)$.

Examples

1. A firm is constituted by two workers, A_1 and A_2, who respectively work z_1 and z_2 hours. Each unit of the good produced by the firm requires exactly one hour of work of either A_1 or A_2 – they are each equally skilled in manufacturing the good.
Hence, we can define the firm's production function as $f(z) = z_1 + z_2$. One hour of each worker can be perfectly substituted by one hour of the other worker. The two inputs are perfect substitutes.

2. A carrier is providing transportation services to other firms. The unit of production is each truck the firm sends on the road. For each unit the carrier needs a truck and one driver operating the truck. A truck without a driver or a driver without a truck does not produce anything.
Hence, we can define $f(z) = \min\{z_1, z_2\}$, where for instance z_1 is the amount of trucks and z_2 is the amount of drivers. In that case, the two inputs are perfect complements.

Technical rate of transformation. We can now define the **technical rate of transformation** (hereafter TRT) which describes the rate at which one input can be substituted for another input while keeping the total output constant. Conceptually, it is very similar to the MRS that we have defined in the consumer theory part.

For simplicity assume that a firm produces a good using only two inputs, that is, its technology writes as $f(z_1, z_2)$. Taking the total derivative of this function and equating to zero (we want to keep output constant) we obtain:

$$df = \frac{\partial f}{\partial z_1} dz_1 + \frac{\partial f}{\partial z_2} dz_2 = 0.$$

We can rearrange this equation so that

$$\frac{dz_2}{dz_1} = -\frac{\frac{\partial f}{\partial z_1}}{\frac{\partial f}{\partial z_2}}.$$

If we set $dz_1 = 1$, the above equation tells us by how much input 2 should decrease when we use one more unit of input 1 and want to keep total output constant. We can define the TRT as follows:

$$TRT_{z_1, z_2} = -\frac{\frac{\partial f}{\partial z_1}}{\frac{\partial f}{\partial z_2}}.$$

Short and long-run. In general, we make a distinction between the short-run and the long-run regarding the choice of inputs. Some factors of production are fixed in the short-run as it may be difficult to adjust time in a short period of time. This is for instance the case with capital (or land) as it may be difficult to quickly increase the size of an entire factory. Other inputs may be more easily adjusted in the short run such as labor or energy consumption. In the long-run, however, we allow all factor of production to vary.

Hence, we sometimes consider the short-run technology of a firm to be a function $f(z_1, \bar{z}_2)$ where one of the input, say input 2, is considered to be fixed and only input 1 can vary.

Returns to scale. In the long-run, when all inputs are allowed to vary we are often interested in analyzing the effect on output of scaling all inputs up by a constant factor $t > 1$.

- *Constant returns to scale (CRS).* A production function exhibits CRS if for $t > 1$ we have that

$$f(tx_1, \dots, tx_n) = tf(x_1, \dots, x_n).$$

- *Increasing returns to scale (IRS).* A production function exhibits IRS if for $t > 1$ we have that

$$f(tx_1, \dots, tx_n) > tf(x_1, \dots, x_n).$$

- *Decreasing returns to scale (DRS).* A production function exhibits DRS if for $t > 1$ we have that

$$f(tx_1, \dots, tx_n) < tf(x_1, \dots, x_n).$$

In words, CRS means that doubling the inputs doubles the output while IRS (resp. DRS) means that doubling more (resp. less) than doubles the output.

Example

The technology $f(x_1, x_2) = x_1 + x_2$ exhibits CRS as $f(tx_1, tx_2) = tx_1 + tx_2 = tf(x_1, x_2)$.

The technology $f(x_1, x_2) = x_1 x_2$ exhibits IRS as $f(tx_1, tx_2) = tx_1 tx_2 = t^2 f(x_1, x_2) > tf(x_1, x_2)$.

The technology $f(x_1, x_2) = x_1^{1/2} + x_2^{1/2}$ exhibits DRS as $f(tx_1, tx_2) = (tx_1)^{1/2} + (tx_2)^{1/2} < tf(x_1, x_2)$.

3.2. Cost minimization

The first problem faced by the firm is that of cost minimization. Each factor of production z_i has a unit cost $w_i > 0$ that the firm has to pay (capital, salaries, etc). For each level of output q, the cost minimization problem is as follows.

$$\begin{aligned} \min_{z \in \mathbb{R}^n_+} \quad & \sum_{i=1}^{n} w_i z_i \\ \text{s.t.} \quad & f(z) = q. \end{aligned}$$

In words, the firm must find the combination of inputs z that allows to produce a quantity q of the output at the minimal possible cost.

The value function of this problem is the total cost function and is denoted by $C(q,w)$. This function describes what is the minimal total cost the firm must incur to produce an amount q when the input cost vector is w. To compute the total cost function we have to solve the firm's cost minimization problem as it is the function defined by $C(q,w) = \sum_{i=1}^{n} w_i z_i^*$ where $z^* = (z_1^*, \ldots, z_n^*)$ is a cost-minimizing choice of inputs.

Let $\lambda \in \mathbb{R}$ be the Lagrange multiplier associated with the constraint $f(z) = q$. The first-order conditions of the Lagrangian of this problem write as

$$w_i + \lambda \frac{\partial f}{\partial z_i}(z) = 0, \text{ for all } i = 1, \ldots, n.$$

Given that we assumed that $w_i > 0$ it follows that $\lambda > 0$. We can therefore take any two of these equations, say $w_i + \lambda \frac{\partial f}{\partial z_i}(z) = 0$ and $w_j + \lambda \frac{\partial f}{\partial z_j}(z) = 0$ and rearrange to obtain:

$$\frac{w_i}{w_j} = \frac{\frac{\partial f}{\partial z_i}(z)}{\frac{\partial f}{\partial z_j}(z)}.$$

This condition states that for any two inputs z_i and z_j, cost minimization implies that the price ratio of input i and j must equate their TRT.

Example

Assume $f(z_1, z_2) = z_1 z_2$. The Lagrangian associated with this cost minimization problem writes as $\mathscr{L} = w_1 z_1 + w_2 z_2 + \lambda(z_1 z_2 - q)$. First-order conditions write:

$$w_1 + \lambda z_2 = 0,$$
$$w_2 + \lambda z_1 = 0.$$

Hence we must have that $w_1/w_2 = z_2/z_1$. Using the technology constraint we have that $z_1 \frac{w_1}{w_2} z_1 = q$ so that $z_1^* = \sqrt{\frac{w_2}{w_1} q}$ and $z_2^* = \sqrt{\frac{w_1}{w_2} q}$.
Finally, we can compute the total cost function as $C(q,w) = w_1 z_1^* + w_2 z_2^*$, that is,

$$C(q,w) = w_1 \sqrt{\frac{w_2}{w_1} q} + w_2 \sqrt{\frac{w_1}{w_2} q}$$

3.3. Profit maximization

Once the firm knows how to produce each quantity level at the minimal possible cost, we are interested in how much quantity it actually wants to produce. Given the price of the good p on the market (the firm is a price-taker), the firm maximizes its revenue minus its cost, that is:

$$\max_{q \in \mathbb{R}_+} \quad pq - C(q,w).$$

The solution to this elementary unconstrained maximization problem is immediately given by its first-order condition:

$$p = \frac{\partial}{\partial q} C(q^*, w).$$

where q^* denotes the firm's optimal production choice given the price p and input prices w.

This equation is probably one of the most 'famous' equations in economics: it states a price-taking competitive firm chooses to produces a quantity q such that 'price equals marginal cost'.

We can therefore easily establish the firm's profit $\pi(p,w)$ as the profit resulting from profit-maximization as:

$$\begin{aligned}\pi(p,w) := & \; pq^* - C(q^*, w) \\ = & \; \frac{\partial C(q^*, w)}{\partial q} q^* - C(q^*, w).\end{aligned}$$

This expression is important as it describes the firm's equilibrium profit only in terms of 'costs'.

We can further rearrange this expression to obtain that

$$\pi(p,w) = q^* \left[\frac{\partial C(q^*,w)}{\partial q} - \frac{C(q^*,w)}{q^*} \right],$$

where $\frac{C(q^*,w)}{q^*}$ is the *average cost* function evaluated at q^*.

Hence we can see that whether the firm is making negative, null or positive profit depends on the difference between its marginal cost and its average cost evaluated at q^*. We will see why this relationship is important in the next section.

'Direct' profit-maximization problem. In practice, when we do not need an explicit formula for the total cost function we often skip the cost minimization problem and directly solve the firm's maximization problem. In that case, we have to write the firm's profit-maximization problem as follows:

$$\max_{(z_1,z_2)\in\mathbb{R}^2_+} \quad pf(z_1,z_2) - (w_1 z_1 + w_2 z_2).$$

This problem is formally equivalent to the one we have seen above but does not require to have an explicit formula for the total cost function.

The first-order conditions write:

$$p\frac{\partial f(z_1,z_2)}{\partial z_1} = w_1,$$
$$p\frac{\partial f(z_1,z_2)}{\partial z_2} = w_2.$$

Notice that dividing the first equation by the second immediately gives $w_1/w_2 = \frac{\partial f}{\partial z_1} / \frac{\partial f}{\partial z_2}$, that is, exactly the condition characterizing the cost-minimizing inputs.

Example

Consider $f(z_1,z_2) = z_1^{\alpha} z_2^{\beta}$, where $\alpha+\beta<1$. Then the first-order conditions associated with the 'direct' profit-maximization problem write $\alpha p z_1^{\alpha-1} z_2^{\beta} = w_1$ and $\beta p z_1^{\alpha} z_2^{\beta-1} = w_2$. Taking the ratio of these two equations immediately gives that $\frac{\alpha}{\beta}\frac{z_2}{z_1} = \frac{w_1}{w_2}$ so that for instance $z_2 = \frac{w_1}{w_2}\frac{\beta}{\alpha} z_1$. Plugging this last equation into the first first-order conditions gives that $\alpha p z_1^{\alpha-1} (\frac{w_1}{w_2}\frac{\beta}{\alpha} z_1)^{\beta} = w_1$. Solving for z_1 gives that $z_1^* = [\frac{\alpha p}{w_1} (\frac{w_1}{w_2}\frac{\beta}{\alpha})^{\beta}]^{\frac{1}{1-\alpha-\beta}}$. Similarly, solving for z_2 we obtain that $z_2^* = [\frac{\beta p}{w_2} (\frac{w_2}{w_1}\frac{\alpha}{\beta})^{\alpha}]^{\frac{1}{1-\alpha-\beta}}$.

4. Market structures

4.1. Perfect competition

A cornerstone of standard microeconomic theory is the paradigm of **perfect competition**. This market structure arises in idealized environment in which the following conditions must hold:

- Large number of sellers and buyers.
- Rational sellers and buyers.
- Homogeneous products.
- Perfect information.

The first condition ensures that there are enough sellers and buyers in the market so that none of them can significantly influence the outcome. We implicitly used this condition when assuming that the consumer and firms were taking price as given. The second condition is simply that we assume that firms maximize profits and consumer maximize utility. The third condition excludes the case in which firms sell similar but different enough products so that they can gain market power (e.g. smartphone producers). In other words, the homogeneous products condition ensures that firms produce goods that are similar enough so that consumers are indifferent where they buy them. The fourth condition assumes that all prices can be perfectly observed, that firms know their costs, that consumers know their utilities and that there is no asymmetric information between sellers and buyers.

There are other conditions such as the ’no externality’ condition and ’no barrier to entry’ condition. The former is important but beyond the scope of this course. The latter states that there is no significant cost for firms to enter the market (such as a very large fixed cost to establish a line of production).

Long-run entry. The no barrier to entry condition is usually used to assume that in the long-run firms will enter as long as there are profit opportunities in the market. In other words, it means that in the long-run all the firms in the market will make zero profit. Combining this observation with our previous result on the firm’s equilibrium profit we can immediately deduce than in the long-run perfect competition implies that each firm’s marginal cost equates its average cost of production, that is, $\frac{\partial C(q^*,w)}{\partial q} - \frac{C(q^*,w)}{q^*}$.

Overall, what characterizes perfect competition on a specific market is usually the fact that (i) price is equal to marginal cost and (ii) firms make zero profit. Of course, the previous conditions and this result have to be understood as an *idealized* market structure. They serve as a benchmark to assess what is the best outcome that could be attained and as a reminder that any violation of one or more of the conditions would result in imperfect competition.

4.2. Monopoly

When there is only one firm producing a given good, we say that this firm is a monopoly. There are several reasons why a firm may be in a monopoly position. It can be because there is a *natural monopoly* due to extremely high fixed costs and other barriers to entry (e.g. a water plant facility or the railway industry). In other cases the monopoly position may be due to network effects or barriers to entry such as patents.

This case is important both as it provides intuition about how a firm makes uses of market power and as a benchmark in more advanced microeconomics (notably in the case of Industrial Organization). Let us first assume that the monopolist incurs a total cost of $C(q)$ when producing an amount q of the good. We further assume that C is differentiable, increasing and convex.

Demand an inverse demand. As the monopoly is the only firm in the market there is a one-to-one correspondence between the quantity produced by the monopoly and its price. This relationship can be captured in two equivalent ways: (i) through the demand function $D(p)$, or (ii) through the inverse demand function $P(q)$. The demand function returns the aggregate demand for each given price p while the inverse demand function returns the price corresponding to selling a quantity q of the good. As its name indicates, the inverse demand function is formally the inverse function of the demand function, i.e., $P = D^{-1}$.

We make the following assumptions: D is twice differentiable and increasing in p. This immediately implies that P is also twice differentiable and increasing in q.

Example

A very common easy demand function is the linear demand function defined by $D(p) = a - bp$ where $a,b \in \mathbb{R}_+$. It is easy to find the associated inverse demand function by solving $q = a - bP(q) \Rightarrow P(q) = \frac{1}{b}[a-q]$.

The monopoly's problem. Given that there is a one-to-one correspondence between quantity and price, the monopoly's problem can be expressed either in terms of choosing a quantity or in terms of choosing a price. When the monopolist chooses the price, its problem writes as follows:

$$\max_{p \in \mathbb{R}_+} \quad pD(p) - C(D(p)).$$

When instead the monopolist chooses the quantity, the problem writes:

$$\max_{q \in \mathbb{R}_+} \quad P(q)q - C(q).$$

Important

The two programs above are rigorously equivalent from a mathematical viewpoint. Indeed, we have that for any p the corresponding demand is $q = D(p)$ so that $P(q) = p$ as $P = D^{-1}$. Hence for any couple (p,q) satisfying $q = D(p)$ we immediately have that $pD(p) - C(D(p)) = P(q)q - C(q)$.

Be careful though, this equivalence between 'choosing price' and 'choosing quantity' holds only for the monopoly. You will see later that choosing price or quantity can drastically change the outcome in oligopoly theory.

Let us now derive the monopoly's solution by focusing on the second formulation, namely,

$$\max_{q \in \mathbb{R}_+} \quad P(q)q - C(q).$$

We can simply take the first-order condition and rearrange it to obtain that:

$$P'(q)q + P(q) = C'(q).$$

The left-hand side is the marginal revenue and the right-hand side is the marginal cost. We can further interpret the left-hand by analyzing the two effects at play on the marginal revenue. When the monopoly increases q it creates two effects opposite effects on revenue: (i) a negative effect due to a decrease in price and (ii) a positive effect due to an increase in overall quantity sold. The term $P'(q)q$ captures the marginal loss of revenue (i) due to a decrease in price on all units sold. The second term $P(q)$ captures the additional revenue (ii) obtained by selling one more unit.

We can further rearrange the first-order condition and get that:

$$P(q)\left[1 + \frac{1}{\varepsilon(q)}\right] = C'(q),$$

where $\varepsilon(q) = \frac{P(q)}{qP'(q)}$ is the elasticity of the inverse demand function capturing the sensitiveness of price to changes in quantity. In the extreme case in which the elasticity is infinite for all q, which happens when the demand function is perfectly elastic to price, the monopoly chooses 'price equals marginal cost'. In all other cases the monopoly chooses a price that is always greater than the marginal cost. This contrasts with what we saw previously, namely that competitive firms produce such that price equals marginal cost. This means that the monopoly has *market power*.

Example

Assume the demand function is $D(p) = 4 - 0.5p$ and the total cost is $C(q) = 2q^2$. It immediately follows that the inverse demand function is $P(q) = 8 - 2q$. Let us solve $\max_q P(q)q - C(q)$, that is, $\max_q (8-2q)q - 2q^2$. The first-order condition writes:

$$8 - 2q - 2q - 4q = 0 \Leftrightarrow q^m = 1.$$

Hence we can find that the monopoly price is $p^m := P(q^m) = 8 - 2 * 1 = 6$. Notice that the marginal cost, $C'(q) = 4q$, evaluated at q^m is $C'(q^m) = 4$ is lower than the monopoly price.