ECC4650/ECC5650/BEX5650 Microeconomics Workshop
2025 edition

Filip Premik

December 12, 2025

Contents

1 Math Review

This notes provide a summary of important mathematical tools that will be used in the subsequent part of the course.

Notation. In all that follows, we will be living in the N-dimensional Euclidean space. We use $\mathbb{R}$ to denote the set of real numbers, and $\mathbb{R}_+$ the set of non-negative real numbers. Also, we will use the concept of neighborhood of a point defined as an open ball around it. With a small abuse of notation, whenever $x \in \mathbb{R}^N$ with $N > 1$, we will write $x = (x_1, x_2, ..., x_N)$.

List of contents. A function. Functions of a single and many variables. Graphical representation of a function. Parametrized functions. Properties of functions: monotonicity, convexity, continuity. Derivative of a function. Graphical interpretation of a derivative. Frequently used derivatives. Partial derivatives. Unconstrained convex optimization. Graphical interpretation of the optimum. Constrained convex optimization. Lagrangian.

1.1 Functions

We begin with a definition of a **function**. A function is a basic concept that allows us to model relationship between two (sets) of variables. A function maps elements from a set of inputs into elements from a set of permissible outputs.

Definition 1.1.1 (function)**.** *A **function** f from a set $X \subseteq \mathbb{R}^N$ to a set $Y \subseteq \mathbb{R}$ is defined as a rule that assigns to each element x in X exactly one element y in Y. This relationship is summarized by $f : X \to Y$.*

Notationally, we write:

$$y = f(x)$$

where:

- $\hookrightarrow$ *x is an element of the* domain X*,*
- $\hookrightarrow$ *y is an element of the* codomain Y*,*
- $\hookrightarrow$ *$f(x)$ is the value of the function at x.*

In our definition of a function, the domain lives in $\mathbb{R}^N$. In the simplest case, $N = 1$ and we have a function of one variable, or a **uni-dimensional** function.

Example 1.1.1 (a function of one variable)**.** *An example of a function of one variable is $f(x) = x^2$. In this case:*

- $\hookrightarrow$ *the domain $X = \mathbb{R}$,*
- $\hookrightarrow$ *the codomain $Y = \mathbb{R}_+$,*
- $\hookrightarrow$ *obviously, for each $x \in X$ there exists only one $y \in Y$ such that $y = x^2$. In other words, the function is well defined.*

If $N > 1$, we talk about functions of many variables.

Example 1.1.2 (a function of two variables)**.** *An example of a function of two variables is $f(x_1, x_2) = x_1^2 + x_2^2$. In this case:*

- $\hookrightarrow$ *the domain $X = \mathbb{R}^2$,*
- $\hookrightarrow$ *the codomain $Y = \mathbb{R}_+$,*
- $\hookrightarrow$ *obviously, for each $x = (x_1, x_2) \in X$ there exists only one $y \in Y$ such that $y = x_1^2 + x_2^2$. In other words, the function is well defined.*

Functions can also be interpreted as a set of points:

$$\{(x, y) \in X \times Y : y = f(x)\}$$

If our function is low-dimensional, meaning that $N \leq 2$ (the domain is at most two-dimensional), this formulation becomes very human-friendly. This is because it allows us to draw a figure representing our function.

Example 1.1.3 (a graph of a function of one variable)**.** *Let $f : \mathbb{R} \to \mathbb{R}$ be a function such that $f(x) = 2x + 5$. It can be represented by a set of points*

$$\{(x, y) : y = 2x + 5, x \in \mathbb{R}\}$$

Graphically:

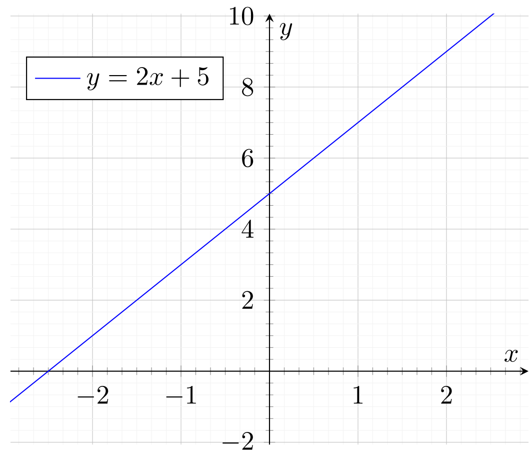

A function describes relationship between variables. The relation itself may depend on a set of parameters shaping the relationship.

Definition 1.1.2 (a parametrized family of functions)**.** *A* ***parametrized family of functions*** *indexed by a (potentially multidimensional) parameter α is a set of functions:*

$$\{f^\alpha : X \to Y \text{ such that } \alpha \in \mathbb{R}^L\} \text{ for some integer } L$$

Example 1.1.4 (a family of linear functions)**.** *An example of a parametrized family of functions is a set of all affine functions. Let $\alpha = (a, b) \in \mathbb{R}^2$. Any linear function can be written as:*

$$f^\alpha(x) = ax + b$$

Function parameters should not be confused with variables. For each function, we will interpret parameters a fixed set of numbers that describe the relation between variables in domain and codomain of the function.

Example 1.1.5. *Consider the set of linear functions $f^\alpha(x) = ax + b$. Every combination of parameters $\alpha = (a, b) \in \mathbb{R}^2$ define a separate function, for example:*

- $\hookrightarrow$ $f^\alpha(x) = 2x + 5$ *for* $\alpha = (a, b) = (2, 5)$
- $\hookrightarrow$ $f^\alpha(x) = -3x$ *for* $\alpha = (a, b) = (-3, 0)$
- $\hookrightarrow$ $f^\alpha(x) = 6$ *for* $\alpha = (a, b) = (0, 6)$

In each of these cases, we can draw a graph of a function, like in the example above.

Parametrizing functions is extremely useful in economics. We will be able to solve a parametrized model once, and the results would hold for many (typically: infinitely many) special cases defined by specific values of parameters.

1.2 Properties of functions

In this course, we will use functions to describe behavior of economic agents. We will typically be interested in functions that have certain properties. These properties will not only make our the behavior consistent with intuitions, but also will allow us to derive desired properties of economic models.

Definition 1.2.1 (monotonicity of a function)**.** *Consider a function $f : X \to Y$. We say that function f is:*

1. ***weakly increasing*** *if, for any two points x_1 and x_2 in its domain such that $x_1 > x_2$, the following holds:*

$$f(x_1) \geq f(x_2)$$

2. ***weakly decreasing*** *if, for any two points x_1 and x_2 in its domain such that $x_1 > x_2$, the following holds:*

$$f(x_1) \leq f(x_2)$$

3. ***constant*** *if, for any two points x_1 and x_2 in its domain such that $x_1 \neq x_2$, the following holds:*

$$f(x_1) = f(x_2)$$

The monotonicity is strict if we can replace $\geq$ ($\leq$) with $>$ ($<$).

Example 1.2.1 (a strictly increasing function)**.** *A function $f(x) = (x-1)^3$ is strictly increasing because:*

$$x_1 > x_2 \Leftrightarrow x_1 - 1 > x_2 - 1 \Leftrightarrow (x_1 - 1)^3 > (x_2 - 1)^3 \Leftrightarrow f(x_1) > f(x_2)$$

A function may be monotone only over a subset of its domain.

Example 1.2.2 (monotonicity on a subset of domain)**.** *A function $f(x) = x^2$ is:*

- $\hookrightarrow$ *strictly decreasing for $x \in (-\infty, 0)$, because for $x_1, x_2 < 0$ it follows that $x_1 > x_2 \Leftrightarrow x_1^2 < x_2^2 \Leftrightarrow f(x_1) < f(x_2)$*
- $\hookrightarrow$ *strictly increasing for $x \in (0, \infty)$, because for $x_1, x_2 > 0$ it follows that $x_1 > x_2 \Leftrightarrow x_1^2 > x_2^2 \Leftrightarrow f(x_1) > f(x_2)$*
- $\hookrightarrow$ *not monotone for $x \in \mathbb{R}$, because $f(-2) > f(1)$ and $f(1) < f(2)$.*

Monotonicity may depend on function parameters.

Example 1.2.3 (monotonicity and parametrization)**.** *Consider a parametrized set of linear functions $f^\alpha(x) = ax + b$ for $\alpha = (a, b) \in \mathbb{R}^2$. A function $f^\alpha(x)$ is:*

- $\hookrightarrow$ *strictly increasing for $\alpha > 0$, because for $x_1 > x_2$ it follows that $ax_1 > ax_2 \Leftrightarrow ax_1 + b > ax_2 + b \Leftrightarrow f(x_1) > f(x_2)$*
- $\hookrightarrow$ *strictly decreasing for $\alpha < 0$, because for $x_1 > x_2$ it follows that $ax_1 < ax_2 \Leftrightarrow ax_1 + b < ax_2 + b \Leftrightarrow f(x_1) < f(x_2)$*
- $\hookrightarrow$ *constant for $\alpha = 0$ because $f(x) = 0x + b = b$*

Note that only parameter a affects the monotonicity property of our function in this case. In turn, parameter b has nothing to do with it.

Monotonicity is easy to read from the graphical representation of a function.

Example 1.2.4 (monotonicity on a graph)**.**

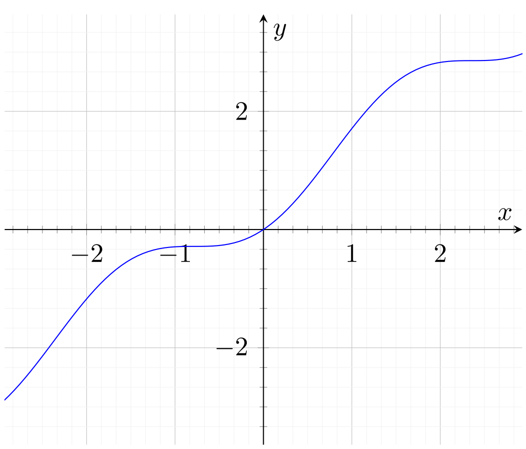

an increasing function

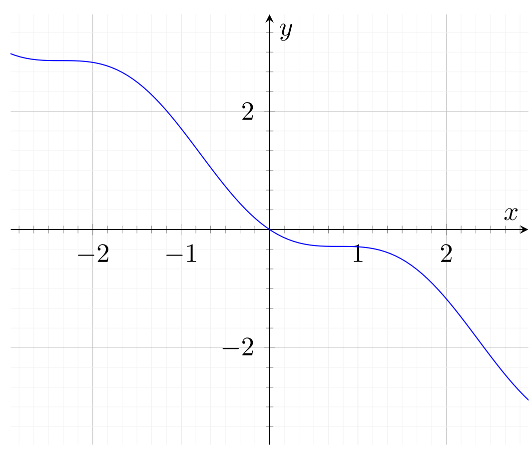

a decreasing function

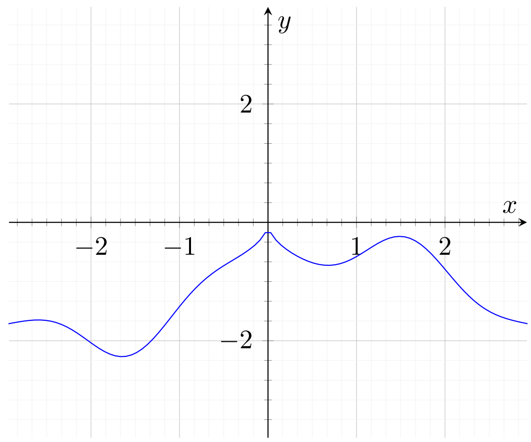

a non-monotonic function

Function monotonicity is a crucial aspect in economics. For example, we typically assume that a consumer prefers to have more units of a good than fewer units of a good. If there is a function summarizing consumer's happiness or *utility* from having a good, it would better be monotone increasing in the number of units of this good.

Another important property of a function is convexity/concavity.

Definition 1.2.2 (convex and concave functions)**.** *Consider a function $f : X \to Y$. We say that function f is:*

↪ ***convex*** *if, for all $x_1, x_2 \in X$ and for all $\lambda \in [0, 1]$, the following inequality holds:*

$$f(\lambda x_1 + (1 - \lambda)x_2) \leq \lambda f(x_1) + (1 - \lambda)f(x_2)$$

This means that the line segment between any two points on the graph of the function lies above or on the graph.

↪ ***concave*** *if, for all $x_1, x_2 \in X$ and for all $\lambda \in [0, 1]$, the following inequality holds:*

$$f(\lambda x_1 + (1 - \lambda)x_2) \geq \lambda f(x_1) + (1 - \lambda)f(x_2)$$

This means that the line segment between any two points on the graph of the function lies below or on the graph.

Furthermore, if $x_1 \neq x_2$, $\lambda \in (0, 1)$, and the above inequalities are strict, the convexity or concavity property is strict.

Example 1.2.5 (graphical interpretation of convexity)**.**

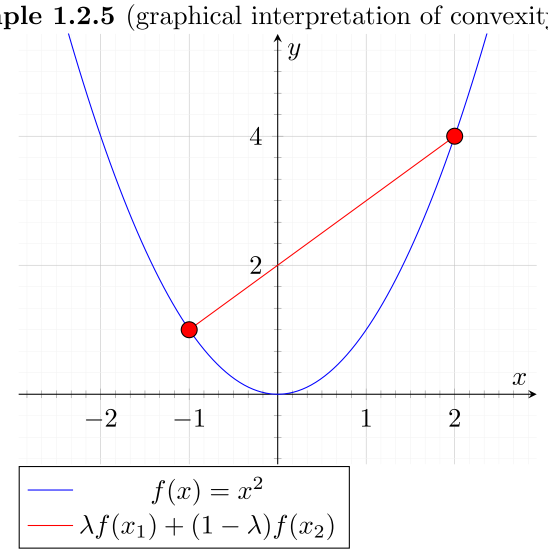

a convex function

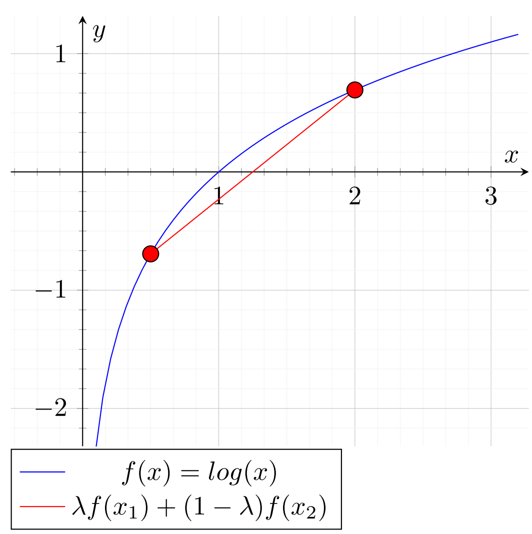

a concave function

Convexity and concavity are important properties facilitating search for maximum or minimum values of the function. We will see it clearly in the following sections, after introducing derivatives.

Now, we move to continuity of a function.

Definition 1.2.3 (continuous functions)**.** *A function $f : X \to Y$ is **continuous** at a point $c \in X$ if the following three conditions hold:*

1. *the function f is defined at c. i.e. $f(c)$ exists*
2. *the limit of $f(x)$ as x approaches c exists:*
$$\lim_{x\to c} f(x) \text{ exists}$$
3. *the limit of $f(x)$ as x approaches c is equal to the function value at c:*
$$\lim_{x\to c} f(x) = f(c)$$

*If a function f is continuous at every point in its domain, then f is said to be a **continuous function**.*

Intuitively, a function continuous if one can draw its graph without lifting your pen from the paper.

Example 1.2.6 (graphical interpretation of continuity)**.**

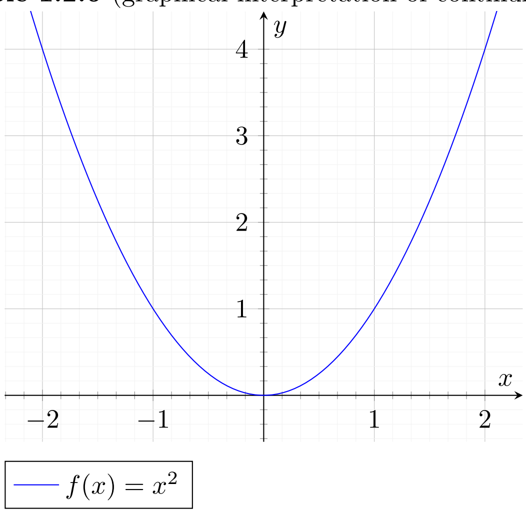

a continuous function

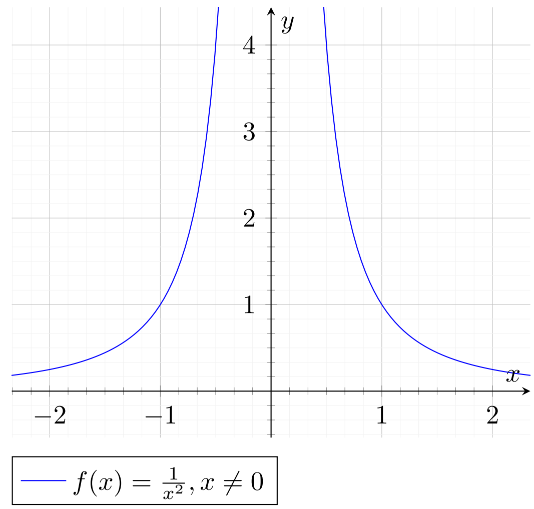

a non-continuous function

Continuity is an important property ensuring that small changes in the input of a function result only in small changes in the output – a result desired in most economic studies. Continuity is also needed for differentiability, which is a key property that we will use to establish model solutions and their properties throughout the course.

1.3 Derivatives of a Function

A derivative of a function is a measure of how the function's value changes as its input changes a little bit. We start our analysis with derivatives of functions of one variable, and then move to partial derivatives of functions of more than one variable.

Definition 1.3.1 (derivative of a function)**.** *Let $X \subseteq \mathbb{R}$. The **derivative** of a continuous function $f : X \to Y$ at a point $x \in X$ is defined as the limit of the average rate of change of the function over an interval as the interval approaches zero.*

$$f'(x) = \lim_{\Delta x\to 0} \frac{f(x + \Delta x) - f(x)}{\Delta x}$$

Definition 1.3.2 (differentiability of a function)**.** *Let $X \subseteq \mathbb{R}$. A continuous function $f : X \to Y$ is said to be **differentiable at a point** if it has a derivative at this point. A continuous function $f : X \to Y$ is said to be **differentiable** if it has a derivative at every point in its domain.*

Although the derivatives are defined as limits, for many functions analytical formulas for derivatives exist.

Proposition 1.3.1 (commonly used derivatives (I))**.** *The following table contains derivatives of commonly used functions:*

Function	**Derivative**
$f(x) = c$	$f'(x) = 0$
$f(x) = x$	$f'(x) = 1$
$f(x) = x^n$	$f'(x) = nx^{n-1}$
$f(x) = e^x$	$f'(x) = e^x$
$f(x) = \ln(x)$	$f'(x) = \frac{1}{x}$
$f(x) = \sin(x)$	$f'(x) = \cos(x)$
$f(x) = \cos(x)$	$f'(x) = -\sin(x)$

Table 1: Commonly Used Derivatives (I)

Function	**Derivative**
$\big(f(x) + g(x)\big)'$	$f'(x) + g'(x)$
$\big(f(x) - g(x)\big)'$	$f'(x) - g'(x)$
$\big(f(x) \cdot g(x)\big)'$	$f'(x) \cdot g(x) + f(x) \cdot g'(x)$
$\big(\frac{f(x)}{g(x)}\big)'$	$\frac{f'(x)g(x) - f(x)g'(x)}{g(x)^2}$
$\Big(f\big(g(x)\big)\Big)'$	$f'\big(g(x)\big) \cdot g'(x)$

Table 2: Commonly Used Derivatives (II)

In addition, there are ready-to-use formulas for derivatives of functions of functions.

Proposition 1.3.2 (commonly used derivatives (II))**.** *Let f, g, h be differentiable functions defined on $\mathbb{R}$. Then:*

It turns out that derivatives have a very useful graphical interpretation.

Proposition 1.3.3 (graphical interpretation of the derivative)**.** *Consider a function $f : X \to Y$ and point $x = c \in X$. The line tangent to the graph of function $f(x)$ at point c is given by:*

$$y - f(c) = f'(c)(x - c)$$

Therefore, the derivative $f'(x)$ provides the slope of the tangent line to the graph of the function at the point x.

Example 1.3.1 (graphical interpretation of derivative)**.** *Consider function $f(x) = \frac{1}{x}$, $x \neq 0$ and a point $c = 1$. Then:*

- $\hookrightarrow$ $f(1) = 1$
- $\hookrightarrow$ $f'(x) = (x^{-1})' = -x^{-2}$ *by the table above*
- $\hookrightarrow$ $f'(1) = -1$
- $\hookrightarrow$ *the line tangent to $f(x)$ at $x = c = 1$ is given by:* $y - 1 = (-1) \cdot (x - 1) \Leftrightarrow y = -x + 2$

Now we can draw it:

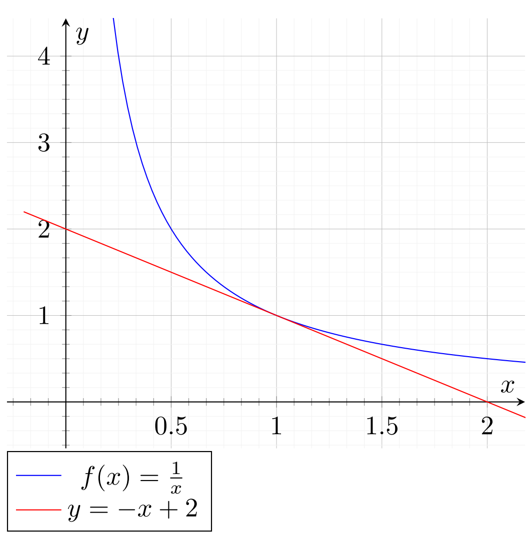

It works!

An immediate corollary from the graphical interpretation of the derivative relates to function's monotonicity.

Proposition 1.3.4 (derivative and monotonicity)**.** *Let $f : X \to Y$ be a uni-dimensional function that is differentiable on an interval $I \subseteq X$. The monotonicity of f on I can be determined by the sign of its derivative $f'(x)$ as follows:*

- ↪ *If $f'(x) \geq (>) 0$ for all x in I then function f is weakly (strictly) increasing on I.*
- ↪ *If $f'(x) \leq (<) 0$ for all x in I then function f is weakly (strictly) decreasing on I.*

Proof. One way of proving this proposition relies on the Mean Value Theorem, which states that if f is continuous on the closed interval $[a, b]$ and differentiable on the open interval (a, b), then there exists a point $c \in (a, b)$ such that:

$$f'(c) = \frac{f(b) - f(a)}{b - a}$$

Using this theorem, we can show that the sign of $f'(x)$ determines whether f is increasing or decreasing on the interval I. □

Now, we move to derivatives of functions of many variables.

Definition 1.3.3 (partial derivative of a function)**.** *Let $X \subseteq \mathbb{R}^N$. A **partial derivative** of a continuous function $f : X \to Y$ at a point $x \in X$ with respect to n-th variable is defined as the limit of the average rate of change of the function over an interval along nth dimension as it approaches zero.*

$$\frac{\partial f}{\partial x_n}(x_1, x_2, \ldots, x_n, \ldots, x_N) = \lim_{\Delta x_n \to 0} \frac{f(x_1, x_2, \ldots, x_n + \Delta x_n, \ldots, x_N) - f(x_1, x_2, \ldots, x_n, \ldots, x_N)}{\Delta x_n}$$

Partial derivatives with respect to nth variable are sometimes also denoted as $f_n'(x)$. Similar as before, the function of many variables is differentiable at x if it is differentiable at every dimension of x at x.

Note that in the above definition, every variable except the one with respect to which we take the derivative is held constant. Partial derivatives hence explain the growth rate of a function when we move only one of its dimensions at a time. This observation is crucial when it comes to calculating partial derivatives.

Example 1.3.2 (partial derivatives)**.** *Let $f(x_1, x_2) = x1^2 + x_1 x_2 + x_2^2$. Then:*

- ↪ $\frac{\partial f}{\partial x_1} = 2x_1 + x_2$
- ↪ $\frac{\partial f}{\partial x_2} = x_2 + 2x_2$

Often, the derivative function is also differentiable. This leads to the second and higher order derivatives.

Definition 1.3.4 (higher level derivatives of uni-dimensional functions)**.** *Let $X \subseteq \mathbb{R}$. The **second derivative** of a continuous function $f : X \to Y$ at a point $x \in X$ (if exists) is defined as:*

$$f''(x) = \big(f'(x)\big)'$$

*Following the same logic, we define **m-th order derivative** (if exists) recursively as:*

$$f^m(x) = \big(f^{m-1}(x)\big)'$$

Example 1.3.3 (higher order derivatives of a uni-dimensional function)**.** *Let $f(x) = x^5 + 3x^3 + 2$. Then:*

- ↪ $f'(x) = 5x^4 + 9x^2$
- ↪ $f''(x) = 20x^3 + 18x$
- ↪ $f^3(x) = 60x^2 + 18$
- ↪ $f^4(x) = 120x$

↪ $f^5(x) = 0$

↪ $f^6(x) = 0$

↪ *and so on..*

It turns out that convexity and concavity is related to the sign of second derivative. Before we introduce that, we start with a short definition.

Definition 1.3.5 (continuous differentiability). *Let $f : X \to Y$ be a uni-dimensional function. We say that f is **twice continuously differentiable** if $f'(x)$ and $f''(x)$ exist and are continuous.*

Proposition 1.3.5 (derivatives and convexity). *Let $f : X \to Y$ be a uni-dimensional function that is twice continuously differentiable on an interval $I \subseteq X$. If*

↪ *$f''(x) > 0$ for every $x \in I$, then function f is convex on I*

↪ *$f''(x) < 0$ for every $x \in I$, then function f is concave on I*

Extending the concept of higher order derivatives to functions of many variables requires paying attention to which dimensions have been used prior in differentiation. In our practice, we will only need second partial derivatives.

Definition 1.3.6 (second partial derivatives). *Let $X \subseteq \mathbb{R}^N$ and consider its two dimensions i, j (potentially equal). The second partial derivative of a continuous function $f : X \to Y$ at a point $x \in X$ (if exists) is defined as:*

$$\frac{\partial^2 f}{\partial x_i \partial x_j}(x) = \frac{\partial}{\partial x_j}\Big(\frac{\partial f}{\partial x_i}(x)\Big)$$

Example 1.3.4 (second order partial derivatives). *Let $f(x_1, x_2) = x1^2 + x_1x_2 + x_2^2$. Then:*

↪ $\frac{\partial f}{\partial x_1}(x_1, x_2) = 2x_1 + x_2$

↪ $\frac{\partial f}{\partial x_2}(x_1, x_2) = x_2 + 2x_2$

And hence:

↪ $\frac{\partial^2 f}{\partial x_1 \partial x_1}(x_1, x_2) = 2$

↪ $\frac{\partial^2 f}{\partial x_1 \partial x_2}(x_1, x_2) = 1$

↪ $\frac{\partial^2 f}{\partial x_2 \partial x_2}(x_1, x_2) = 2$

The derivatives will be a very important tool for us to find extrema: minima and/or maxima of functions.

1.4 Unconstrained Convex Optimization

In this section, we discuss extrema of functions, that is points at which the function attains the highest or the lowest value in certain neighborhood.

Definition 1.4.1 (local extrema of a function). *Let $f : X \to Y$ and consider a point c belonging to the interior of its domain. We say that function f has a:*

↪ *a **local maximum** at a point c if there exists a neighborhood $I \subseteq X$ containing c such that:*

$$f(c) \geq f(x)$$

for all $x \in I$. This means that $f(c)$ is greater than or equal to the function values at all nearby points.

↪ *a **local minimum** at a point c if there exists a neighborhood $I \subseteq X$ containing c such that:*

$$f(c) \leq f(x)$$

for all $x \in I$. This means that $f(c)$ is less than or equal to the function values at all nearby points.

Example 1.4.1 (local extrema)**.**

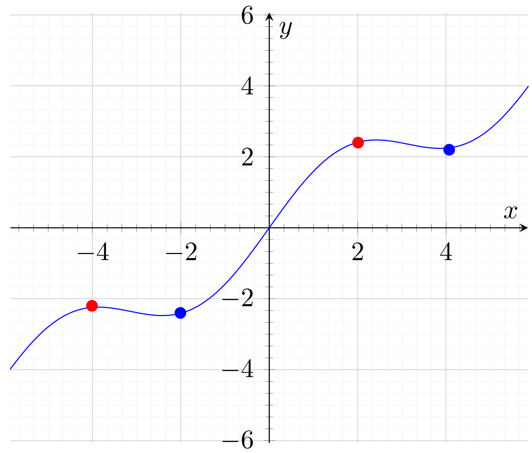

blue dots: local minima, red dots: local maxima

Finding local extrema of a function is not an easy task in general. Fortunately, it gets much easier to do it for functions that share certain properties. For us, the key will be continuity and differentiability. Note that continuity prevents any discrete jumps. Differentiability ensures that the derivative exists. It turns out that derivatives are crucial in searching for extrema of continuous functions.

Proposition 1.4.1 (local extrema and the derivative)**.** *Let $f : X \to Y$ be a continuous and differentiable. Let $c \in X$ be an interior point in X. If f attains a local extremum at c, then $f'(c) = 0$.*

Think about a local maximum (minimum would be analogous) of a uni-dimensional function. If a continuous function has a local maximum at c, it means that there is an open interval (a, b) such that $c \in (a, b)$ in which values of the function are lower or equal than $f(c)$. Since we have a local maximum in $x = c$, for any $a < x_r < c$ it has to be that $f(x_r) \leq f(c)$. Therefore, the function needs to be increasing on at least some small interval to the right of c. Analogous logic applies for points to the left of c: the function must be decreasing at least over a small interval just to the of c. Given the proposition above, it means that $f'(x)$ is positive on points just left to c and negative on points right to c. Hence, it must be equal to zero at c.

Example 1.4.2 (local extrema and the derivative)**.** *The following graph depicts a uni-dimensional function $f(x) = -x^2+4$ (blue curve) with tangent lines (red lines) at various points. We can clearly see that as we approach the local (and global) maximum, lines become less and less steep. Eventually, at the maximum, the tangent line is horizontal – in other words, its slope is equal to zero.*

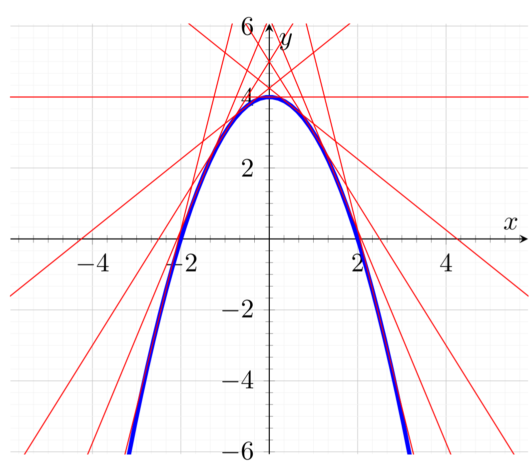

This observation brings us an important observation, saying that if there is a local extremum at c, the derivative must be zero. However, this is only a necessary condition. That means, that we cannot say that if $f'(c) = 0$ then there is a local extremum at c.

Example 1.4.3 (a zero derivative needs not imply a local extremum)**.** *Let $f(x) = 0$ and consider $c = 0$. Note that $f'(x) = 3x^2$, so $f'(0) = 0$. However, $c = 0$ is clearly not a local extremum of the function f.*

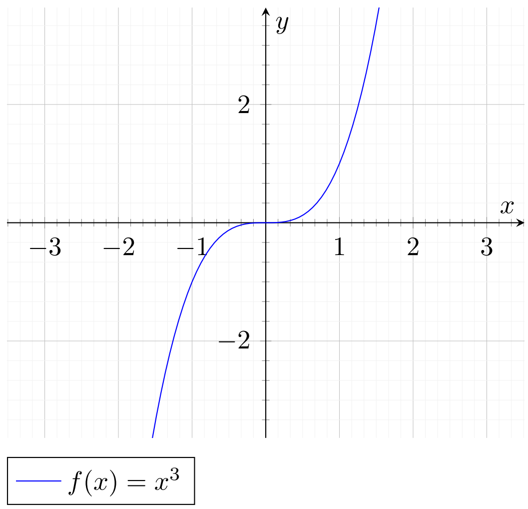

What we are missing are the sufficient conditions for a local extremum.

Proposition 1.4.2 (sufficient conditions for local extrema in one dimension)**.** *Let $f : X \to Y$ be a uni-dimensional function that is twice continuously differentiable on an interval I containing a point c. f attains:*

1. *a local maximum at c if $f'(c) = 0$ and $f''(c) < 0$ (in other words, f is concave in the neighborhood of c)*
2. *a local minimum at c if $f'(c) = 0$ and $f''(c) > 0$ (in other words, f is convex in the neighborhood of c)*

The last proposition gives us a recipe on how to find local extrema of a twice continuously differentiable function. First, we need to find all points at which the first derivative vanishes. Second, we need to verify the sign of the second derivative at each of these points.

These considerations easily extend to functions of many variables. To simplify the notation, we introduce a definition of a gradient.

Definition 1.4.2 (gradient)**.** *Let $f : X \to Y$ be a differentiable N-dimensional function. The vector of first partial derivatives:*

$$\nabla f(x) = \begin{bmatrix} \frac{\partial f}{\partial x_1} & \frac{\partial f}{\partial x_2} & \cdots & \frac{\partial f}{\partial x_N} \end{bmatrix}$$

*is called **gradient** of function f at x.*

The necessary condition states that at a local extremum every first partial derivative needs to be zero.

Proposition 1.4.3 (local extrema and partial derivatives)**.** *Let $f : X \to Y$ be a continuous and differentiable function of many variables. Let $c = (c_1, .., c_N) \in X$ be an interior point in X. If f attains a local extremum at c, then $\frac{\partial f}{\partial x_n} = 0$ for each $n \in \{1, 2, .., N\}$. Alternatively, we can express this condition using gradient:*

$$\nabla f(x) = 0$$

There is also an analogous proposition stating sufficient conditions. To introduce them, we first need to define an analog of a second derivative in the multi-dimensional world and an analog of it being positive or negative at a given point $x \in X$.

Definition 1.4.3 (hessian)**.** *Let $f : X \to Y$ be a twice-continuously differentiable N-dimensional function. Matrix of second partial derivatives:*

$$H(x) = \begin{bmatrix} \frac{\partial f^2}{\partial x_1^2} & \frac{\partial f^2}{\partial x_1 \partial x_2} & \cdots & \frac{\partial f^2}{\partial x_1 \partial x_N} \\ \frac{\partial f^2}{\partial x_2 \partial x_1} & \frac{\partial f^2}{\partial x_2^2} & \cdots & \frac{\partial f^2}{\partial x_2 \partial x_N} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial f^2}{\partial x_N \partial x_1} & \frac{\partial f^2}{\partial x_N \partial x_2} & \cdots & \frac{\partial f^2}{\partial x_N^2} \end{bmatrix}$$

*is called **hessian** of function f at x.*

Definition 1.4.4 (Hessian's definite-ness)**.** *Let $f : X \to Y$ be a twice-continuously differentiable N-dimensional function and let $H(x)$ be its hessian at x. We say that $H(x)$ is:*

1. *positive semi-definite if for every $z \in X$, it holds that $z^T H(x) z \geq 0$, where ^{T} denotes the transposition operator.*
2. *negative semi-definite if for every $z \in X$, it holds that $z^T H(x) z \leq 0$, where ^{T} denotes the transposition operator.*

If in addition we could write $>$ ($<$), we say that hessian is positive (negative) definite.

Now we are ready to provide the sufficient conditions for a local extremum of a multi-dimensional function.

Proposition 1.4.4 (sufficient conditions for local extrema in multiple-dimensions)**.** *Let $f : X \to Y$ be a multi-dimensional function that is twice continuously differentiable on a neighborhood I containing a point c. f attains:*

1. *a local maximum at c if $\nabla f(x) = 0$ and $H(x)$ is negative semi-definite*
2. *a local minimum at c if $\nabla f(x) = 0$ and $H(x)$ is positive semi-definite*

Unsurprisingly, definiteness of a hessian is related to convexity.

Proposition 1.4.5 (sufficient conditions for local extrema in multiple-dimensions)**.** *Let $f : X \to Y$ be a multi-dimensional function that is twice continuously differentiable at every point of its domain. The following holds:*

1. *$f(x)$ is convex if and only if its hessian $H(x)$ is positive semi-definite for all $x \in X$.*
2. *$f(x)$ is concave if and only if its hessian $H(x)$ is positive semi-negative for all $x \in X$.*

We see that searching for local extrema of multi-dimensional functions is analogous to searching for local extrema of uni-dimensional functions.

Lastly, we move to defining global extrema of a function.

Definition 1.4.5 (global extrema of a function)**.** *Let $f : X \to Y$ be a uni-dimensional function. We say that function f:*

- $\hookrightarrow$ *attains a* ***global maximum*** *at $c \in X$ if for every point $x \in X$, $f(c) \geq f(x)$*
- $\hookrightarrow$ *attains a* ***global minimum*** *at $c \in X$ if for every point $x \in X$, $f(c) \leq f(x)$*

1.5 Constrained Convex Optimization

In applications in economics, we typically need to consider a range of constraints on function arguments. For example, given certain budget we can afford only a subset of bundles. Therefore, we need to be able to find highest or lowest values of a function while being restricted to a subset of its domain.

We start with providing a useful existence result.

Proposition 1.5.1 (lowest and highest value of a function on a closed set)**.** *Let $f : X \to Y$ be a continuous function defined on a closed and bounded set $X \subseteq \mathbb{R}^N$. Then, f attains both its maximum and minimum values on D. In other words, there exist points $x_{max}, x_{min} \in X$ such that:*

$$f(x_{min}) \leq f(x) \leq f(x_{max}) \quad \textit{for all} \quad x \in X.$$

Example 1.5.1 (local extrema on a closed set)**.** *We consider function $f(x) = -x^2 + 4$ defined on the closed interval $x \in [-2.5, 1]$. On this interval, the function attains a global maximum at $x = 0$, and a global minimum at $x = -2.5$.*

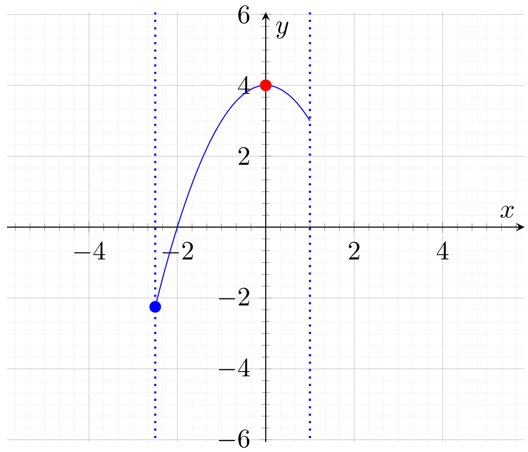

blue dots: global minima, red dots: global maxima

The example above suggests a way of looking for extrema of a function on a closed subset of domain. We need to find any interior local extrema using the general theory from the previous section, as well as check values of the function at the boundary of our closed domain.

However, sometimes the restriction of the domain is more sophisticated than just a simple interval. In general, we may want to find the maximum (minimum) value of a functions given that its arguments satisfy certain equation, say $h(x) = 0$.

Fortunately, there is a way around that.

Theorem 1.1 (Lagrange Multipliers)**.** *Let $f : X \to Y$ be a function in $\mathbb{R}^N$ that is continuously differentiable. Let also $h(x)$ be differentiable. Consider an optimization problem:*

$$\max_{x \in X} f(x), \quad \text{such that } h(x) = 0$$

*For a multiplier $\lambda \in \mathbb{R}$, define the **Lagrangian** function $\mathcal{L}$ as:*

$$\mathcal{L}(x, \lambda) = f(x) + \lambda \cdot h(x)$$

Let the optimum be x^*. Then there exists a unique λ^* such that:

$$\begin{cases} \underset{k=1,..,N}{\forall} \; \frac{\partial \mathcal{L}}{\partial x_k}(x^*, \lambda^*) = \frac{\partial f}{\partial x_k}(x^*) + \lambda^* \frac{\partial h}{\partial x_k}(x^*) = 0 \\ \frac{\partial \mathcal{L}}{\partial \lambda}(x^*, \lambda^*) = h(x^*) = 0 \end{cases}$$

The above theorem provides an algorithm of finding constrained extrema of a continuous differentiable function subject to continuous differentiable constraints. This is a workhorse tool of microeconomic theory. In majority of our classes, we will rely on this tool to solve economic models.

2 Consumer Theory

In this and the next class, we study consumer theory. The plan is to develop theory explaining the mechanism behind consumer consumption choices that would have testable implications. That means, we will develop a model that set constraints on objects that can be observed as data, so that we can use the data to verify whether these constraints hold.

Today, we start with a refreshment of mathematical tools that will be used throughout this chapter. Next, we define the basic primitive of the model, that is, what is to be consumed. Then, we proceed with a definition of agents' preferences defined over the set of admissible consumption bundles. Lastly, we move to defining an optimal choice and studying its properties. These properties imply a set of restrictions on individual behavior that can be tested using data. Today's class aims to build a foundation necessary to move to the next class in which we will bring our model to data.

We postpone studying consumer surplus (and welfare in general) until class 6.

List of contents. Consumption Set. Consumption Bundle. Preference relations and their properties: completeness, transitivity, continuity, monotonicity, local non-satiation Indifference curve in space of preferences. Examples. Utility representation of preferences: existence, (non-)uniqueness (ordinal vs cardinal orders). Indifference curve in space of utilities. Properties of indifference curves. Marginal rate of substitution. Utility: usual functional forms. Prices. Budget set. Consumer's utility maximization problem. Methods to solve the consumer's maximization problem. First order conditions. Second order conditions. Solution objects: demand function (Marshallian, uncompensated), indirect utility. Properties of demand. Roy's identity. Consumer's expenditure minimization problem. Hicksian (compensated) demand function. Expenditure function. Shaphard's Lemma. Substitution and income effects. Slutsky equation. Inferior and normal goods.

Textbook references. Espinola: Ch. 2-4; Varian: Ch. 7-9; Jehle & Reny: Ch. 1; MWG: Ch. 1-3

2.1 Mathematical Glossary

Definition 2.1.1 (a closed set)**.** *A set $C \subseteq \mathbb{R}^N$ is called **closed** if it contains all its limit points. This means that if a sequence of points $\{c_k\}$ in C converges to a point $c \in \mathbb{R}^N$, then c must also be in C.*

Definition 2.1.2 (a convex set)**.** *A set $C \subseteq \mathbb{R}^N$ is called **convex** if, for any two points $c^1, c^2 \in C$, the line segment connecting c^1 and c^2 is entirely contained within C. This means that for any $c^1, c^2 \in C$ and any $\lambda \in [0,1]$:*

$$\lambda c^1 + (1-\lambda)c^2 \in C$$

Definition 2.1.3 (contour sets)**.** *Let $f : \mathbb{R}^N \to \mathbb{R}$ and let $c \in \mathbb{R}^N$. Define:*

1. *the **contour set** of f at c to be the set of all points in $\mathbb{R}^N$ at which function f takes the value c.*

$$\{x \in \mathbb{R}^N : f(x) = c\}$$

2. *the **upper contour set** of f at c to be the set of all points in $\mathbb{R}^N$ at which function f takes values greater or equal to c.*

$$\{x \in \mathbb{R}^N : f(x) \geq c\}$$

3. *the **lower contour set** of f at c to be the set of all points in $\mathbb{R}^N$ at which function f takes values lower or equal to c.*

$$\{x \in \mathbb{R}^N : f(x) \leq c\}$$

Definition 2.1.4 (open ball)**.** *An **open ball** in $\mathbb{R}^N$ centered at point $c \in \mathbb{R}^N$ with radius $\varepsilon > 0$ is the set of all points in $\mathbb{R}^N$ whose distance from c is less than ε:*

$$B(c, \varepsilon) = \{x \in \mathbb{R}^N : \|x - c\| < \varepsilon\}$$

where $\|\cdot\|$ denotes the Euclidean norm.

Definition 2.1.5 (inequality relations on vectors)**.** *Let $c^1, c^2 \in \mathbb{R}^N$. We say that:*

1. *c^1 is weakly larger than c^2, i.e. $c^1 \geq c^2$, if for every element $i \in \{1, .., N\}$ it holds that $c^1_i \geq c^2_i$.*
2. *c^1 is strictly larger than c^2, i.e. $c^1 \gg c^2$, if for every element $i \in \{1, .., N\}$ it holds that $c^1_i > c^2_i$.*

Definition 2.1.6 (quasi-concave functions)**.** *A function $f : \mathbb{R}^N \to \mathbb{R}$ is called **quasi-concave** if, for any two points $c^1, c^2 \in \mathbb{R}^N$ and any $\lambda \in [0,1]$:*

$$f(\lambda c^1 + (1-\lambda)c^2) \geq \min\{f(c^1), f(c^2)\}$$

2.2 Consumption Goods

We consider a general setting in which there are N different consumption goods. We assume that each good is perfectly divisible, which means that any real quantity of a given good can be consumed. For convenience, we also assume that only nonnegative quantities can be consumed, though this assumption is not necessary for the methods we develop in this chapter to work.

We begin with defining the consumption bundle and consumption set.

Definition 2.2.1 (consumption bundle)**.** *A **consumption bundle** (shortly: a bundle) is an ordered list of non-negative real numbers, reflecting quantities of goods. We denote it by $c = (c_1, c_2, ..., c_N)$.*

Definition 2.2.2 (consumption set)**.** *The set of all possible consumption bundles is called* ***consumption set****. We denote it by C.*

By definition, every consumption bundle belongs to the consumption set: $c \in C$. In this chapter, we assume that the consumption set is a subset of $\mathbb{R}^N$: $C \subseteq \mathbb{R}^N_+$, and that it is closed, convex, and contains zero as an element. These properties will be useful to establish existence of consumer's optimal consumption choice.

In a typical application, $C = \mathbb{R}^N_+$: that is every nonnegative amount of every good belongs to the consumption set. However, you can easily imagine in which the consumption set may be bounded from above.

2.3 Preferences

In this chapter, we study behavior of a single consumer. Our consumer may *like* some bundles *better* than others. We model this phenomenon using a mathematical concept of *binary relation.*

Definition 2.3.1 (a binary relation)**.** *A binary relation R between two elements c^1, c^2 is a statement $c^1 R c^2$, which can be either true or false.*

Definition 2.3.2 (consumer's preference relation)**.** ***Consumer's preference relation*** *is a binary relation that $\succsim$ such that for any $c^1, c^2 \in C$, if $c^1 \succsim c^2$ then we say that c^1 is* as good as *c^2 for the consumer.*

Assuming that consumer preference relation is a binary relation is too general to describe a typical economic behavior. We introduce more assumptions reflecting desired properties of the consumer reference relation.

Definition 2.3.3 (a complete relation)**.** *A binary relation $\succsim$ is* ***complete*** *over a set C if for every $c^1, c^2 \in C$, either $c^1 \succsim c^2$ or $c^2 \succsim c^1$.*

Definition 2.3.4 (a transitive relation)**.** *A binary relation $\succsim$ is* ***transitive*** *over a set C if for every $c^1, c^2, c^3 \in C$, if $c^1 \succsim c^2$ and $c^2 \succsim c^3$, than it must be that $c^1 \succsim c^3$.*

Completeness and transitivity are fundamental properties of consumer's preferences. The form states that every pair of admissible consumption bundles can be compared. The latter ensures a preference order over all elements can be established by using just pair-wise comparisons. They make a preference relation *make sense*, gaining it a honorable title of rationality.

Definition 2.3.5 (rational preference relation)**.** *A preference relation $\succsim$ over a consumption set C is called* ***rational*** *if it is complete and transitive.*

Rational preferences are a backbone of modern economics.

Now, we extend the notion consumer's preference relation $\succsim$ by defining a strict preference relation and an indifference preference relation.

Definition 2.3.6 (a strict preference relation)**.** *Let $\succsim$ be a preference relation. For any $c^1, c^2 \in C$, define* ***strict preference relation*** *by:*

$$c^1 \succ c^2 \Leftrightarrow c^1 \succsim c^2 \text{ and } c^2 \not\succsim c^1$$

The strict relation describes a situation in which one good is *better* than other from the points of view of our consumer.

Definition 2.3.7 (an indifference preference relation)**.** *Let $\succsim$ be a preference relation. For any $c^1, c^2 \in C$, define* ***indifference preference relation*** *by:*

$$c^1 \sim c^2 \Leftrightarrow c^1 \succsim c^2 \text{ and } c^2 \succsim c^1$$

The strict relation describes a situation in which one good is *as good as* as the other for the consumer.

You should be able to show that rationality extends to strict and indifference consumer's preference relations.

Exercise 2.3.1. *Suppose that consumer preference relation $\succsim$ is rational. Show that:*

a) strict preference relation

b) indifference preference relation

are also rational preference relations, i.e. they are complete and transitive binary relations.

The class of rational preference relations is still too general for us to proceed with. We need to impose more assumptions on it. As we will see, these properties are not only useful in terms of desired mathematical properties tools but also seem to fit many real life situations.

We start with introducing a concept of continuity of preferences. To define it, we need to understand what are the preference contour sets.

Definition 2.3.8 (preference contour sets)**.** *Let $\succsim$ be a preference relation over C and let $\overline{c} \in C$. Define:*

1. *the **upper contour set** of $\succsim$, that is the set of elements weakly better than $\overline{c}$:*

$$\succsim (\overline{c}) = \{c \in C : c \succsim \overline{c}\}$$

2. *the **lower contour set** of $\succsim$, that is the set of elements weakly worse than $\overline{c}$:*

$$\precsim (\overline{c}) = \{c \in C : \overline{c} \succsim c\}$$

3. *the **contour set** of $\succsim$, that is the indifference curve for an element $\overline{c}$:*

$$\sim (\overline{c}) = \{c \in C : \overline{c} \sim c\}$$

Example 2.3.1 (upper contour set and indifference curve)**.** *Let $C = \mathbb{R}^2_+$. In the figure below, the dotted line presents an indifference curve for some preference relation. Consider consumption bundle $c \in C$, as marked on the figure. The shaded area north-east form c represents the upper contour set of $\succsim$ at c, that is, the set of (weakly if the indifference curve is included, strictly otherwise) preferred bundles.*

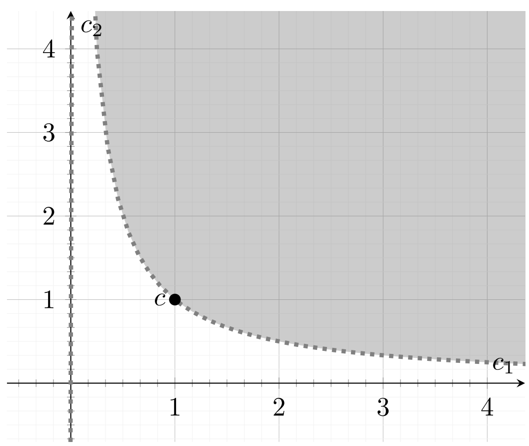

We define continuity of preference relation using the contour sets.

Definition 2.3.9 (continuity of preference relation)**.** *A preference relation $\succsim$ over a consumption set C is called **continuous** if for every consumption bundle $c \in C$, the set of weakly better elements $\succsim (c)$ and the set of weakly worse elements $\precsim (c)$ are both closed in $\mathbb{R}^N_+$.*

Continuous preferences imply that the consumer is not going to sharply reverse their preference over consumption bundles that are *close* to each other. We typically think that this is a feature of consumer behavior. In addition, continuity is a key technical assumption that will allow us to construct utility representation of preferences.

Another desired property of preferences is local non-satiation.

Definition 2.3.10 (local non-satiation of preferences)**.** *A preference relation $\succ$ over a consumption set C is **locally non-satiated** if for every consumption bundle $c \in C$ and every $\varepsilon > 0$, there exists $\overline{c} \in B(c, \varepsilon)$ such that $\overline{c} \succ c$.*

Local non-satiation implies that consumer's *desires* are unlimited – for every consumption bundle, there is another one that is located arbitrarily close and is strictly preferred. Local non-satiation is violated if there exists a (local) **bliss point** - a point for which there exists an open ball around it for which every point would be *worse* than it in terms of preference relation.

Local non-satiation doesn't say what makes another consumption bundle *better*. Intuitively, the more, the better. This concept is summarized by monotonicity of preferences.

Definition 2.3.11 ((strict) monotonicity of preferences)**.** *We say that preference relation is* ***strictly monotone*** *if for every* $c^1, c^2 \in C$*, if:*

1. *if* $c^1 \geq c^2$ *then* $c^1 \succsim c^2$
2. *if* $c^1 \gg c^2$ *then* $c^1 \succ c^2$

Local non-satiation and monotonicity are related, which you should be able to show.

Exercise 2.3.2. *Prove or provide a counterexample:*

a) *does strict monotonicity imply local non-satiation?*

b) *does local non-satiation imply strict monotonicity?*

The last desired property of preferences is convexity.

Definition 2.3.12 ((strict) convexity of preferences)**.** *A preference relation* $\succsim$ *is called* ***convex*** *if for every* $c^1, c^2 \in C$ *and every* $\lambda \in [0, 1]$*:*

$$c^1 \succsim c^2 \Rightarrow \lambda c^1 + (1 - \lambda)c^2 \succsim c^2$$

Furthermore, preference relation $\succsim$ *is called* ***strictly convex*** *if for every* $c^1, c^2 \in C$ *such that* $c^1 \neq c^2$ *and every* $\lambda \in [0, 1]$*:*

$$c^1 \succsim c^2 \Rightarrow \lambda c^1 + (1 - \lambda)c^2 \succ c^2$$

Convexity ensures the desired curvature of the indifference curves. In particular, the indifference curves will be (weakly) convex relative to the origin. That means that the amount of one good that needs to be traded for a unit of another to remain at the same indifference curve will be relatively higher (lower) if the former good is in abundance (shortage). In other words, convexity of preferences implies the principle of *diminishing marginal rate of substitution.* Convexity is also very useful from the technical point, as it will be inherited (to some extent) by the utility representation giving us the sufficient conditions for consumer's optimal behavior.

From now on, unless stated differently, we will work with preferences that are rational, continuous, strictly monotone, and strictly convex on $C \subseteq \mathbb{R}^N$.

So far, we have built a notion of consumer preferences in a very abstract form. In the next section, we link it with the utility representation of preferences, that should be more familiar to you.

2.4 Utility Representation

We start with formalizing the notion of a utility function representing a preference relation.

Definition 2.4.1 (utility function)**.** *Function* $u : C \to \mathbb{R}$ *is called a* ***utility function*** *representing preference relation* $\succsim$ *if for every* $c^1, c^2 \in C$ *the following holds:*

$$c^1 \succsim c^2 \Leftrightarrow u(c^1) \geq u(c^2)$$

It turns out that a broad class of preferences can be represented by a utility function.

Theorem 2.1 (existence of utility representation)**.** *Let* $\succsim$ *be a preference relation that is complete, transitive, continuous, and strictly monotone. Then there exists a continuous function* $u : C \to \mathbb{R}$ *that represents* $\succsim$*.*

The proof of this theorem lies beyond the scope of out class. However, if you aim at a PhD program, you should be able to prove it yourself.

A key result of the above theorem is that the resulting utility representation is continuous. That is a key property that will enable us to use calculus to seek consumer's optimal behavior. Moreover, the more properties of preference relations are also inherited by a utility function representing them.

Proposition 2.4.1 (properties of utility function)**.** *Let* $\succsim$ *be a preference relation that is complete, transitive, continuous, and strictly monotone, and* $u : C \to \mathbb{R}$ *be a utility function that represents* $\succsim$*. The following statements hold:*

1. $u(c)$ *is strictly increasing if and only if* $\succsim$ *is strictly monotone*

2. *$u(c)$ is quasi-concave if and only if $\succsim$ is convex*
3. *$u(c)$ is strictly quasi-concave if and only if $\succsim$ is strictly convex*

Exercise 2.4.1 (properties of utility function). *Prove the proposition above.*

From now on, we assume that the utility representation is also differentiable. Differentiability can also be shown to be a consequence of specific assumptions regarding preference relation but we skip this in our class.

The results above imply that utility representation exists for a large class of preference relations. However, it is not unique.

Proposition 2.4.2 ((non-) uniqueness of utility representation). *Let $u : C \to \mathbb{R}$ represent preference relation $\succsim$, and $f : \mathbb{R} \to \mathbb{R}$ be a strictly increasing uni-dimensional function. Function $v(c) = f\big(u(c)\big)$ also represents $\succsim$.*

Exercise 2.4.2 ((non-) uniqueness of utility representation). *Prove the proposition above.*

The non-uniqueness proposition brings us to an important distinction between ordinal and cardinal preferences. **Ordinal preferences** refer to a way of representing consumer preferences where only the order of preferences matters, not the specific intensity of those preferences. In turn, **cardinal preferences** assign specific numbers to each bundle, enabling statements like c^1 is twice as much preferred as c^2.

Utility representation emphasizes the ordinal nature of preferences. The absolute distance between bundles in the preference space does not really matter, as we can represent the same preference relation with a different set of numbers induced by any strictly increasing transformation of the original utility function.

However, even though the utility representation naturally describes only the ordinal aspect of preferences, we are still able to talk about substitution patterns between goods. In particular, we may ask how much of one good one needs to be supplemented in exchange for a unit of another good to keep the utility level constant. To do this, we first analyze the set of bundles which are equally preferred, i.e. are associated with the same utility.

Definition 2.4.2 (indifference curve). *Let $u : C \to \mathbb{R}$ represent consumer preferences. The indifference curve at utility value $\overline{u}$ is a set:*

$$IC(\overline{u}) = \{c \in C : u(c) = \overline{u}\}$$

Note that the graph of a utility function is $N+1$ dimensional, including all the consumption dimensions plus utility values. The indifference curve, in turn, fixes a utility value: hence it is a N dimensional object. That enables a comvienient graphical analysis of consumer's preference even with a two-dimensional consumption set.

Example 2.4.1 (indifference curve in two dimensions). *Consider $C = \mathbb{R}^2_+$. Let $u : C \to \mathbb{R}$ represent consumer's preferences. The indifference curve associated with utility level $\overline{u}$ is given by:*

$$u(c_1, c_2) = \overline{u}$$

This formulation is a convenient way to get the graph of the function. Note that we can use the equation above to derive c_2 as a function of c_1 ($\overline{u}$ will just be a parameter):

$$c_2 = f^{\overline{u}}(c_1)$$

and draw the graph in (c_1, c_2) space.

To make things less abstract, consider another example.

Example 2.4.2 (indifference curve in two dimensions (II)). *Let $u(c_1, c_2) = c_1 \cdot c_2$, and $\overline{u} = 2$. The line combining all bundles that bring utility level 2 is given by:*

$$c_2 = \frac{2}{c_1}$$

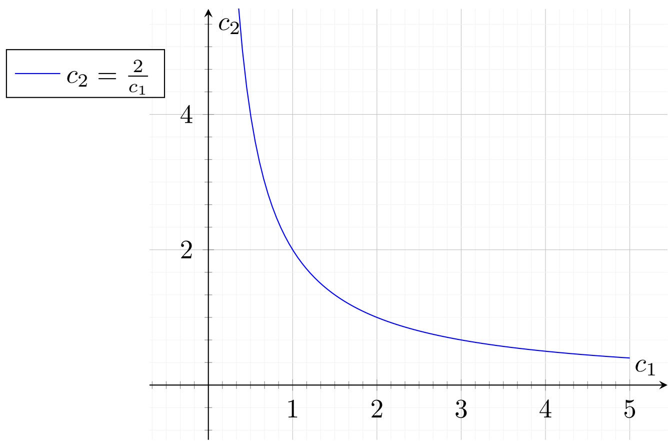

Note that bundles presenting utility higher than 2 are in the upper count set of this line, whereas bundles presenting utility level smaller than two lie below this curve.

Indifference curves have a few interesting properties. We leave them for you to prove.

Exercise 2.4.3 (properties of utility function)**.** *Let $C = R^2_+$ and consider a strictly increasing and quasi-convex utility function. Show that:*

1. *two indifference curves associated with different utility values cannot cross*
2. *the indifference curve is downward sloping*

Related to indifference curves is the notion of marginal rate of substitution.

Definition 2.4.3 (marginal rate of substitution)**.** ***Marginal rate of substitution*** *between good i and good j, $MRS_{ij}(c)$, denotes the amount of good j that needs to be changed in response to a unit change in good i in order to keep the utility level unchanged.*

How to derive MRS? Let's start to build our intuition in a two-dimensional setting.

Example 2.4.3 (MRS in two dimensions)**.** *In the previous examples, we showed that if C is two-dimensional, the indifference curve can be expressed simply as a relation between c_1 and c_2:*

$$c_2 = f^{\overline{u}}(c_1)$$

Note that the definition of MRS in two dimensions coincides with the interpretation of derivative of $f^{\overline{u}}(c_1)$:

$$MRS_{12}(c) = \frac{d}{dc_1} f^{\overline{u}}(c_1)$$

Furthermore, we can plug the $c_2 = f^{\overline{u}}(c_1)$ into the equation defining the indifference curve to obtain:

$$u\big(c_1, f^{\overline{u}}(c_1)\big) = \overline{u}$$

Taking a total derivative with respect to c_1:

$$\frac{\partial u}{\partial c_1} + \frac{\partial u}{\partial c_2} \cdot \frac{d}{dc_1} f^{\overline{u}}(c_1) = 0$$

Using the fact that $MRS_{12}(c) = \frac{d}{dc_1} f^{\overline{u}}(c_1)$, we can now obtain the equation for MRS in terms of utilities:

$$MRS_{12}(c) = -\frac{\frac{\partial u}{\partial c_1}}{\frac{\partial u}{\partial c_2}}$$

Analogously, we can derive the $MRS_{ij}(c)$ for a general N-dimensional consumption set. We leave this as an exercise for you.

In practice, economists typically begin their analysis of consumer's behavior departing straight from a definition of a utility function (as opposed to defining a binary preference relation). Below we list frequent choices for consumer's utility functions.

description	function
perfect substitutes	$u(c_1, c_2) = \alpha c_1 + \beta c_2,\ \alpha, \beta > 0$
perfect complements	$u(c_1, c_2) = \min\{\alpha c_1, \beta c_2\},\ \alpha, \beta > 0$
quasi - linear	$u(c_1, c_2) = c_1 + f(c_2)$
logarithmic	$u(c_1, c_2) = \alpha \log c_1 + \beta \log c_2,\ \alpha, \beta > 0$
CRRA	$u(c_1, c_2) = \frac{c_1^{1-\sigma}}{1-\sigma} + \beta \frac{c_2^{1-\sigma}}{1-\sigma},\ \beta > 0, \sigma > 0, \sigma \neq 1$
Cobb-Douglas	$u(c_1, c_2) = c_1^{\alpha} c_2^{\beta},\ 0 \geq \alpha + \beta \leq 1$
CES	$u(c_1, c_2) = (\alpha c_1^{\rho} + \beta c_2^{\rho})^{\frac{1}{\rho}},\ \alpha, \beta > 0$

Example 2.4.4 (frequently used utility function)**.** *Let $C = \mathbb{R}_+^2$ with a typical element $c = (c_1, c_2)$. Frequently, the following utility functions are assumed to represent consumer's utility: CRRA stands for Constant Relative Risk Aversion, and CES for Constant Elasticity of Substitution.*

You should get familiar with the functional forms of utility representation presented in the table above.

Exercise 2.4.4 (frequently used utility functions)**.** *For every class of functions presented above, derive and draw the indifference curves for some assumed values of the parameters. What can you tell about the nature of preference relation that they represent?*

2.5 Consumer's Problem

Now we know pretty well how to deal with consumer's preferences. We are ready to analyze consumer's choices.

Given our assumptions on consumer's preferences, in particular monotonicity, our consumer would consume infinite amounts of goods in absence of any further restrictions. However, economics is all about dealing with limited resources. This will apply also to our consumer. We further restrict the set of consumption bundles that are *feasible* for the consumer by exploiting the concept of affordability, by assuming that the consumer has limited budget and can only purchase goods up to a certain limit.

We formalize this by letting $p = (p_1, p_2, ..., p_N) \in R_{++}^N$ be a N-dimensional vector of strictly positive prices, and $w > 0$ be the endowment of wealth/income of the consumer. We assume that both p and w are taken as given by our consumer while deciding about consumption levels.

We begin the analysis of consumer choice with formally defining what is *feasible* to the consumer.

Definition 2.5.1 (affordable bundle)**.** *A consumption bundle $c \in C$ is called **affordable** (feasible) if its monetary value is not greater than income:*

$$p \cdot c = \sum_{i=1}^{N} p_i c_i \leq w$$

Knowing what is an affordable bundle, we may be tempted to define the set of all affordable bundles.

Definition 2.5.2 (budget set)**.** ***Budget set** contains all bundles affordable by the consumer at p and w:*

$$B(p, w) = \{c \in C : p \cdot c \leq w\}$$

The budget set has a series of useful properties. Make sure you know how to formally prove them.

Exercise 2.5.1 (properties of the budget set)**.** *Let $C \subseteq \mathbb{R}^2$ and $0 \in C$. Suppose $p \gg 0$ and $w > 0$. Show that the budget set $B(p, w)$ is:*

1. *non-empty*
2. *closed*
3. *convex*

At this stage, we are ready to analyze optimal consumption. The task for our consumer is to choose a consumption bundle from among the affordable bundles that yields highest utility.

$$\max_{c \in C} u(c)$$
$$\text{such that } p \cdot c \leq w$$

Finding a bundle that maximizes utility is the same as finding the most preferred bundle. This seemingly obvious fact is left for you to prove as an exercise.

Exercise 2.5.2 (utility maximization and preferences)**.** *Suppose there exists a solution to the consumer maximization problem as defined above.*

1. *Show that the resulting bundle is weakly preferred (in terms of $\succsim$) over any other affordable bundle.*
2. *What would we need to assume about the preference relation for it to be strictly preferred (in terms of $\succ$) over any other affordable bundle?*

The utility maximization problem is a constrained optimization problem. The only constraint, called the **budget constraint**, is an inequality constraint. This sort of mathematical programming problems can be solved using a Kuhn-Tucker method, which we skip in our class. Instead, we will note one fact about the construction of consumer's problem under standard assumptions regarding preferences that will allow us to use our well-known Lagrange multipliers approach in finding the optimal consumption bundle.

Proposition 2.5.1 (budget constraint is binding at the optimum)**.** *Suppose preferences are strictly monotone and let $c^* \in C$ solve the consumer problem. Then $p \cdot c^* = w$*

Proof. Suppose not. Then it must be that $p \cdot c^* < w$. Let $\bar{c}$ be such that $\bar{c}_i = c_i^* + \frac{w - p \cdot c^*}{p_i} \frac{1}{N}$ for $i = 1, .., N$. Now note that $\bar{c}$ is affordable:

$$p \cdot \bar{c} = \sum_{i=1}^{N} p_i \Big(c_i^* + \frac{w - p \cdot c^*}{p_i} \frac{1}{N} \Big) = \sum_{i=1}^{N} \Big(p_i c_i \Big) + N \frac{(w - p \cdot c^*)}{N} = p \cdot c^* + w - p \cdot c^* = w$$

Moreover, since $\bar{c}_i > c_i^*$ for every dimension i, it follow that $\bar{c} \gg c^*$. By monotonicity of preferences, it must be that:

$$u(\bar{c}) > u(c^*)$$

but this contradicts optimality of c^*. □

Given the result above, for strictly monotone preferences (or strictly increasing utility functions) we can rewrite the optimization problem as:

$$\max_{c \in C} u(c)$$
$$\text{such that } p \cdot c = w$$

To solve it, we first formulate the Lagrangian function:

$$\mathcal{L}(c, \lambda) = u(c) - \lambda(p \cdot c - w)$$

Second, we obtain the first order conditions by differentiating the Lagrangian with respect to all of its arguments (elements of the consumption bundle and the multiplier):

$$\begin{cases} \frac{\partial \mathcal{L}}{\partial c_i} = \frac{\partial u}{\partial c_i} - \lambda p_i = & 0 \quad , \ i = 1, ..., N \\ \frac{\partial \mathcal{L}}{\partial \lambda} = p \cdot c - w = & 0 \end{cases}$$

Before we move to further analysis of the optimum solution, let's have a break and consider two short examples building more intuition.

Example 2.5.1 (FOC and MRS)**.** *The first order conditions illustrate an intuitive property that must be held in the optimum: the MRS between any pair of consumption goods must be equal to the ratio of their prices.*

$$\frac{\frac{\partial u}{\partial c_i}}{\frac{\partial u}{\partial c_j}} \equiv MRS_{ij} = \frac{p_i}{p_j}, \ i \neq j \in \{1, ..., N\}$$

That means that in the optimum the relative price of a pair of goods in units of utility is equal to their relative price in monetary terms.

Example 2.5.2 (FOC in two dimensions)**.** *We can analyze the FOC in two-dimension to gain more intuition about the mechanics of the optimal solution. First, note that in dimensional setting, the ratio of prices is in fact the slope of the budget constraint in (c_1, c_2) space:*

$$p_1 c_1 + p_2 c_2 = w \Leftrightarrow c_2 = -\frac{p_1}{p_2} c_1 + \frac{w}{p_2}$$

In turn, the MRS is simply the slope of a line tangent to the indifference curve, because (using the notation from a previous example):

$$MRS = \frac{d}{dc_1} f^{\overline{u}}(c_1)$$

Hence, the slope of budget constraint has to be equal to the slope of indifference curve. This implies that, at the optimum, the budget line has to be tangent to an indifference curve.

Coming back to the FOC, they constitute $N+1$ (possibly non-linear) equations with $N+1$ unknowns, so there is hope we can get a solution. Typically we would express λ in terms of c_i's and thus get rid of it – we do not need to know the actual value of λ. Eventually, any $c = (c_1, .., c_N)$ that satisfies FOCs may be an optimum – a point that maximizes consumer's utility given the budget constraint. To make sure we found the optimal choice, we need to verify the second order conditions. Alternatively, we can use the following proposition:

Proposition 2.5.2 (sufficient conditions for quasi-concave utility function)**.** *Suppose $u : C \to \mathbb{R}$ is continuous and quasi-concave on C. Let $p \gg 0$ and $w > 0$. Let λ be the Lagrange multiplier associated with the budget constraint. If u is differentiable at $c \in C$, $(c, \lambda) \gg 0$ solves the FOCs, then c solves the consumer's problem at p and w.*

Finally, we arrive with (a set of) $c \in C$ that are solutions to the consumer utility maximization problem. In general, the optimal solution is a function of (p, w). We denote it by $c(p, w)$ and call the Marshallian demand:

Definition 2.5.3 (Marshallian demand)**.** *The optimal solution to consumer maximization problem is called* ***Marshallian consumer demand****:*

$$c(p, w) = \big(c_1(p, w), ..., c_N(p, w)\big)$$

Consumer's utility maximization problem is only one way in which we may analyze consumer's behavior. Another way would be to minimize expenditure while ensuring the consumer gets at least $\overline{u}$ utils. This kind of a problem is called the consumer's expenditure minimization problem:

$$\min_{c^h \in C} p \cdot c^h$$
$$\text{such that } u(c^h) \geq \overline{u}$$

It turns out that expenditure minimization is a well-defined problem.

Proposition 2.5.3. *Given our assumptions, there exists an optimal solution to the minimum expenditure problem.*

Exercise 2.5.3. *Prove the above proposition.*

Given the proposition above, we can find (a set of) $c^h \in C$ that are solutions to the consumer expenditure minimization problem. In general, the optimal solution is a function of $(p, \overline{u})$. We denote it by $c^h(p, \overline{u})$ and call it the Hicksian demand:

Definition 2.5.4 (Hicksian demand)**.** *The optimal solution to expenditure minimization problem is called* ***Hicksian consumer demand****:*

$$c^h(p, \overline{u}) = \big(c_1^h(p, \overline{u}), ..., c_N^h(p, \overline{u})\big)$$

Hicksian demand is sometimes called the **compensated** demand function, because it holds the utility level fixed – in response to change in prices it assumes a shift in income to compensate the consumer for the change in prices. In turn, the Marshallian demand is sometimes called the **uncompensated** demand function, because holding the income fixed, changes in prices will induce changes in the utility level.

If you think that Marshallian and Hicksian demands are related, you are right. It turns out that the utility maximization and expenditure minimization are *dual* problems. Although we do not define it formally yet, you can think that in some notion they are solutions to the same problem, just formulated in different ways. These differences are crucial: in the data, we cannot observe utility (the objective of utility maximization) but we can observe expenditures (the objective of expenditure minimization). Similarly, we can observe income (the second argument of Marshallian demand), but we cannot observe $\overline{u}$ (the second argument of Hicksian demand). We will make advantage of this duality to be able to bring our model to data by deriving testable restrictions. To do this, we need to study the properties of both demand functions.

2.6 Properties of Consumer's Demand

In this section, we study properties of both uncompensated and compensated demand functions.

Proposition 2.6.1 (differentiability of Marshallian demand)**.** *Let c^* be the optimal solution to the consumer maximization problem and suppose:*

1. *the utility function is twice-continuously differentiable on C,*
2. $\frac{\partial u}{\partial c_i}(c^*) > 0$ *for each i,*
3. *the Hessian of the Lagrange function $\mathcal{L}$ has a non-zero determinant at (c^*, λ^*)*

Than $c(p, y)$ is differentiable at (p, y).

The differentiability of the Marshallian demand is a key property that will allow us to define, among others, the income and substitution effects.

Second, Marshallian demand is linked to the income.

Proposition 2.6.2 (adding-up property of Marshallian demand)**.** *Let $c(p, w)$ be the Marshallian demand function. Then:*

$$\sum_{i=1}^{N} p_i c_i(p, w) = w$$

Proof. This proposition is proved above, in the section describing the utility maximization problem. □

The uncompensated demand is also homogeneous of degree zero.

Proposition 2.6.3 (homogeneity of Marshallian demand)**.** *The Marshallian demand function $c(p, w)$ is homogeneous of degree 0 in both arguments. It means that for every $t > 0$, the following holds:*

$$c(tp, tw) = c(p, w)$$

The homogeneity of degree zero of Marshallian demand has an important interpretation. It is sometimes called **no money illusion**. Note that it means that if we increase all prices and consumer's income proportionally $t > 0$ times, the demand will remain unchanged. That means that the value of money is relative. What matters are **real** terms – relative prices and real income.

Definition 2.6.1 (relative prices)**.** *Relative price p_{ij} is the number of units of good j than needs to be given up to acquire 1 unit of good i:*

$$p_{ij} = \frac{p_j}{p_i}$$

Definition 2.6.2 (real income)**.** *Real income in terms of good j is the amount of good j that can be acquired by spending the (nominal) income:*

$$w_j = \frac{w}{p_j}$$

To see that nominal terms do not matter, we use the homogeneity of Marshallian demand. Without loss of generality, we will normalize all monetary terms to the price of good 1. It follows that:

$$c(p_1, p_2, ..., p_N, w) = c\Big(p_1\frac{p_1}{p_1}, p_1\frac{p_2}{p_1}, ..., p_1\frac{p_N}{p_1}, p_1\frac{w}{p_1}\Big) \underset{\text{homogeneity of degree 0}}{=} c\Big(1, \frac{p_2}{p_1}, ..., \frac{p_N}{p_1}, \frac{w}{p_1}\Big)$$

In other words, consumer's behavior changes only with changes in $N-1$ relative prices and real income.

Marshallian demand gives rise to another useful object, the indirect utility function.

Definition 2.6.3 (indirect utility function)**.** *The optimal value of consumer utility at (p, w) is called the **indirect utility function**:*

$$v(p, w) = u\big(c(p, w)\big)$$

It turns out that in many applications in economics, in particularly empirical, it is more convenient to directly define the indirect utility function and build the model without directly relating to the original utility maximization. The indirect utility function has a few desired properties.

Proposition 2.6.4 (properties of indirect utility)**.** *Suppose $u(c)$ is continuous and strictly increasing. Then the indirect utility $v(p,w)$ is:*

1. *continuous in all of its arguments*
2. *homogeneous of degree zero in (p,w), meaning that for every $t>0$:*
$$v(tp,tw)=v(p,w)$$
3. *strictly increasing in w*
4. *decreasing in p*
5. *quasi-convex in (p,w)*
6. *satisfies **Roy's identity**. Namely, if $v(p,w)$ is differentiable at (p,w) and $\frac{\partial v}{\partial w}\neq 0$, then:*
$$c_i(p,w)=\frac{\frac{\partial v}{\partial p_i}(p,w)}{\frac{\partial v}{\partial w}(p,w)},\ i\in\{1,...,N\}$$

Now, we move to analyzing the Hicksian demand. Unsurprisingly, we get a homogeneity result also for the Hicksian demand.

Proposition 2.6.5 (homogeneity of Hicksian demand)**.** *The Hicksian demand function $c^h(p,\overline{u})$ is homogeneous of degree 0 in prices. It means that for every $t>0$, the following holds:*
$$c(tp,\overline{u})=c(p,\overline{u})$$

The Hicksian demand gives rise to the expenditure function.

Definition 2.6.4 (expenditure function)**.** *The optimal value of consumer expenditure at $(p,\overline{u})$ is called the **expenditure function**:*
$$e(p,\overline{u})=p\cdot c^h(p,\overline{u})$$

Proposition 2.6.6 (properties of expenditure function)**.** *Suppose $u(c)$ is continuous and strictly increasing. The expenditure function is:*

1. *continuous in all of its arguments,*
2. *strictly increasing and unbounded above in $\overline{u}$*
3. *increasing in p*
4. *homogeneous of degree one in p, which means that for any $t>0$:*
$$e(tp,\overline{u})=te(p,\overline{u})$$
5. *concave in p*
6. *if in addition u is strictly quasi-concave, the expenditure function satisfies **Shephard's lemma**. Namely, if $e(p,\overline{u})$ is differentiable in p at $(p,\overline{u})$ with $p\gg 0$ and:*
$$c_i^h(p,\overline{u})=\frac{\partial e}{\partial p_i}(p,\overline{u}),\ i=1;,...,N$$

With the derived properties above, we are able to describe the relation between solution objects in two sides of the consumer's problem: utility maximization and expenditure minimization. Intuitively, first fix (p,w) and let $\overline{u}=v(p,w)$. This must imply that $e(p,v(p,w))\leq w$. Second, now fix $(p,\overline{u})$ and let $w=e(p,\overline{u})$. It must be that $v(p,e(p,\overline{u}))\geq\overline{u}$. Under our assumptions, these inequalities turn into equalities.

Proposition 2.6.7 (indirect utility and expenditure function)**.** *Let $v(p, y)$ and $e(p, \overline{u})$ be the indirect utility and expenditure functions respectively for a consumer whose utility function u is continuous and strictly increasing. Then for any strictly positive price vector $p \gg 0$, income $w \geq 0$ and utility level $\overline{u}$ it follows that:*

1. $e\big(p, v(p, w)\big) = w$
2. $v\big(p, e(p, \overline{u})\big) = \overline{u}$

Proposition 2.6.8 (Marshallian and Hicksian demands)**.** *For any strictly positive price vector $p \gg 0$, income $w \geq 0$ and utility level $\overline{u}$ it follows that:*

1. $c_i(p, w) = c_i^h\big(p, v(p, w)\big)$
2. $c_i^h(p, \overline{u}) = c_i\big(p, e(p, \overline{u})\big)$

The following two graphs, borrowed from Deaton & Muellbauer (1980), summarizes the relation between objects discussed in this section.

Figure 1: from value demands to value functions

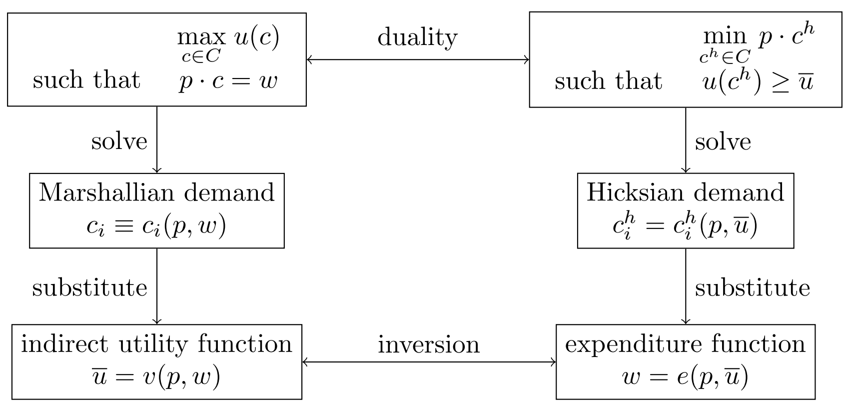

Figure 2: from value functions to demands

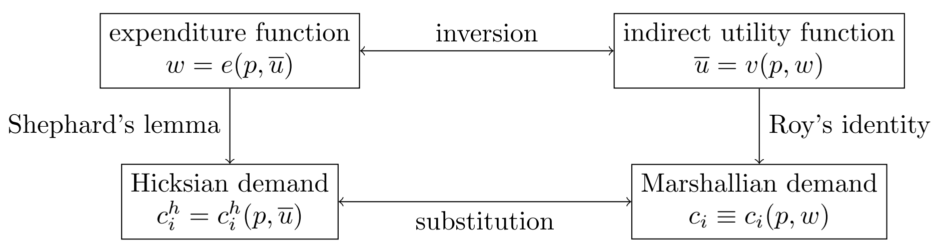

The two propositions above highlight the dual nature of both consumer's problems. They represent the same problem, just tackled from two different perspectives. Essentially, the solutions are the same. The last series of proposition provides recipe of how to get from one object to another using their properties. If you aim at a PhD program, you should be able to prove formally these statements.

2.7 Testable Implications

The following proposition lists restrictions implied by the consumer theory that can be testified empirically.

Proposition 2.7.1 (testable properties of demand)**.**

1. ***adding-up***. *The total expenditures sum up to income.*

$$\sum_{i=1}^{N} p_i c_i(p, w) = w \quad \text{and} \quad \sum_{i=1}^{N} p_i c_i^h(p, \overline{u}) = w$$

2. ***homogeneity***, *or no money illusion.*

$$c_i(tp, tw) = c_i(p, w) \quad \text{and } c_i^h(tp, \overline{u}) = c_i^h(p, \overline{u}) \text{ for any } t > 0$$

3. ***symmetry****. cross-price derivatives of Hicksian demand are symmetric.*

$$\frac{\partial c_i^h(p,\overline{u})}{\partial p_j} = \frac{\partial c_j^h(p,\overline{u})}{\partial p_i} \text{ for any } i \neq j$$

4. ***negativity****. the hessian of the expenditure function (in prices) is negative semi-definite.*

We have discussed the first two properties in detail above. The remaining to are more technical and find their sources in the properties of preferences. We shortly provide intuition behind them. The symmetry ensures consistent choices. The negativity basically means that the expenditure function is concave in prices. All of these implications of the theoretical model carry important empirical content that will serve to verify the consumer theory using data.

2.8 Consumer's Comparative Statics

So far, we have derived optimal consumer's behavior conditionally on a fixed environment given by prices p and income w. However, in real world the environment in which consumer operates can be very volatile. In this section, we analyze changes in consumer's behavior in response to (small) changes in the environment. This type of analysis is often called the **comparative statics**.

Imagine that a price of one goods increases a little bit. This affects the relative prices and real income – other goods become relatively cheaper, but the real income goes down. The consumer is likely to adjust their optimal choice. First, they may demand more of the relatively cheaper goods and less of the relatively more expensive good. Second, the drop in purchasing power may lead to an overall smaller consumption. We formalize these two driving forces below.

Definition 2.8.1 (substitution effect)**.** *A change in the demanded bundle that is driven only by changes in relative prices (holding the income fixed) is called the* ***substitution effect****.*

Definition 2.8.2 (income effect)**.** *A change in the demanded bundle that is driven only by changes in the real income is called the* ***income effect****.*

The total effect of a change in price is a sum of these two objects. Before we prove it formally, let's first develop a graphical intuition.

Example 2.8.1 (substitution and income effect in two dimensions)**.** *Consider and example with two goods. Suppose $p_1 = p_2 = 1$ and $w = 4$. That gives rise to the original budget constraint (blue line) given by $c_2 = -c_1 + 4$. Suddenly, price of good 2 increases to 2, $p_2' = 2$. The new budget constraint will be given by $c_2 = -\frac{1}{2}c_1 + 2$ (red line). The red line reflect both changes in relative prices and income. To get a budget constraint that only accounts for change in the relative prices, holding the original income fixed, we need to find a line that have the same slope as the new budget constraint but also intersects the original optimum (the green line). We call this line the* ***pivoted*** *budget constraint. The left panel below illustrates the three budget constraints.*

The left panel adds the optimal indifference curves at each budget constraint together with the optimal bundle. A move from the black point to the blue is called the substitution effect – it reflects the effect of changes in relative prices holding income fixed. The move from the blue point to the red point is called the income effect – it reflect the effect of changing income holding relative prices fixed.

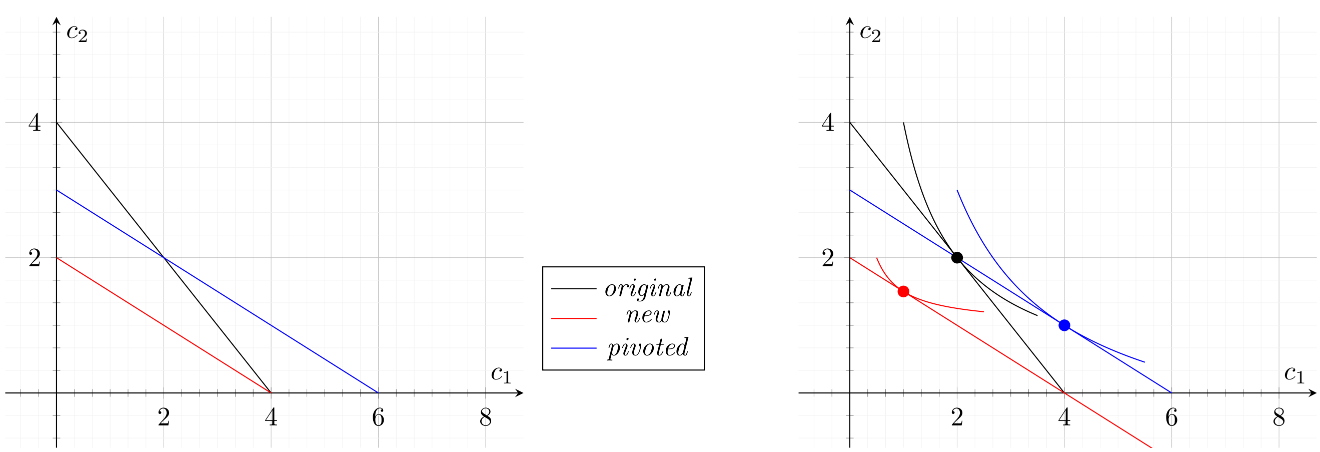

Now, we formalize the considerations above.

Proposition 2.8.1 (the Slutsky equation)**.** *Let $c(p,w)$ be the Marshallian demand function, and $\overline{u}$ be the (maximal) utility that consumer achieves at prices p and income w. Then:*

$$\underbrace{\frac{\partial c_i(p,w)}{\partial p_j}}_{\text{total effect}} = \underbrace{\frac{\partial c_i^h(p,\overline{u})}{\partial p_j}}_{\text{substitution effect}} \underbrace{- c_j(p,w)\frac{\partial c_i(p,w)}{\partial w}}_{\text{income effect}}$$

Exercise 2.8.1. *Prove the result above.*

Proposition 2.8.2. *Let $c_i^h(p,w)$ be the Hicksian demand for good $i \in \{1,...,N\}$. Then:*

$$\frac{\partial c_i^h(p,\overline{u})}{\partial p_i} \leq 0$$

Exercise 2.8.2. *Prove the result above.*

Definition 2.8.3 (normal and inferior goods)**.** *A good is **normal** if the demand for it increases with income. A good is called an **inferior** good, if the demand for it decreases with income.*

That brings us to the celebrated result of the Law of Demand, which is the last result for today.

Proposition 2.8.3. *Suppose the price of good i decreases. If:*

1. *good i is normal, then demand for i will increase.*
2. *good i is inferior, then demand for i will decrease.*

3 Demand Systems

In the previous class, we derived a coherent theory of consumer behavior based on a low-level and somewhat abstract assumptions on the structure of preference relations. Today, we will look at how to bring the theory to the data.

The key object of interest while bringing consumer theory to the data is the demand function describing the optimal consumption choices. We would typically concentrate on Marshallian demand – optimal consumption as a function of prices and income, because its arguments are typically observed in the data. In turn, the Hicksian demand depends on a utility level which no one has ever seen.

The link between theory and empirics are so called testable implications – restrictions that theory places on model parameters that must hold if the data is generated according to the theory. We will study a range of approaches of how economists did and do bring theory and data together.

List of contents. Log-log demand model. Income elasticity of demand. Price elasticity of demand. Adding-up constraint. Homogeneity constraint. Symmetry constraint. Expenditure shares. Stone's demand system. Compensated elasticities. Price indices. Stone's price index. Linear expenditure system. LES price index. Almost ideal demand system. Engel curves. Necessity and luxury goods. PIGLOG class of preferences. AIDS price index. Subutility. Separability. Hicksian and functional separability. Weak separability of a function. Functional separability. Conditional demand function. Two-stage budgeting. Homothetic preferences. Aggregate demand function. Gorman theorem.

Literature References. REFERENCESREFERENCES

[1] Stone, Richard (1954). Linear expenditure systems and demand analysis: an application to the pattern of British

[2] Theil, H. (1965). The information approach to demand analysis. Econometrica, 33, 67-87 demand. The Economic Journal, 64(255), 511-527.

[3] Deaton, Angus and John Muellbauer (1980) An almost ideal demand system, The American Economic Review 70.3 (1980): 312-326.

[4] Deaton, Angus and John Muellbauer (1980) Economics and consumer behavior. Cambridge university press. Chapters 1-3 (demand systems), 5 (separability and aggregation).

[5] Varian, Hal (1992). Microeconomic analysis (Vol. 3). New York: Norton. Chapter: 9 (separability and aggregation).

3.1 Mathematical Glossary

Definition 3.1.1 (elasticity). *Let $f : X \to Y$ be a multi-dimensional function. An* ***elasticity*** *of f with respect to x_i, $i \in \{1, ..., N\}$, is given by:*

$$s_i^f(x) = \frac{\partial \log f}{\partial \log x_i} = \frac{\partial f}{\partial x_i}\frac{x_i}{f}$$

The elasticity describes how the value of a function would change in relative terms if we move x_i by 1%, holding everything else constant.

3.2 The Log-Log Model

Early empirical models of consumer demand were oriented almost completely to the data. The economists would specify a simple (that means linear) equation linking quantities with prices and income and estimate their parameters without thinking of how to link such a representation to the data.

Nevertheless, the early models still carried some theoretical content linking them to the consumer theory. Before we show this, we introduce a concept of elasticities of demand.

Definition 3.2.1 (income elasticity of demand). ***Income elasticity of demand*** *is defined as:*

$$e_i = \frac{\partial \log c_i(p, w)}{\partial \log w}$$

Definition 3.2.2 (price elasticity of demand). ***Price elasticity of demand*** *is defined as:*

$$e_{ij} = \frac{\partial \log c_i(p, w)}{\partial \log p_j}$$

Elasticities are one of most favorite objects among (particularly empirical) economists. They summarize the relationship between two variables in a way that is independent from their units of measurement, greatly facilitating the interpretation and making the numbers (more) comparable across studies.

Recall that the Marshallian demand is a function of prices and income $c(p, w)$. No wonder, the standard choice of specification for a theory-free empirical model of demand links c, p and w in a log-linear form. This functional form implies that parameters of the model can directly be interpreted as elasticities. Moreover, linearity makes the equation easy to estimate which was important in the early days.

Definition 3.2.3 (log-log demand model). *The* ***log-log*** *(double logarithmic)* ***demand model*** *links specifies log-linear relationship between quantities and prices.*

$$\log c_i = \alpha_i + e_i \log w + \sum_{k=1}^{N} e_{ik} \log p_k,\ i \in \{1, ..., N\}$$

It turns out that with minimal assumptions on the environment, we still may derive some theoretical restrictions on parameters of the log-log model. The restrictions are derived from the budget constraint.

Proposition 3.2.1. *Suppose that a demand function $c = c(p, w)$ exists, and that the budget constraint holds with equality:*

$$\sum_{i=1}^{N} p_i c_i = w$$

Then, the following properties hold:

- ***adding-up****.*

$$\sum_{i=1}^{N} p_i \frac{\partial c_i}{\partial w} = 1 \quad \text{and} \quad \sum_{i=1}^{N} p_i \frac{\partial c_i}{\partial p_j} + c_j = 0,\ j \in \{1, .., N\}$$

- ***homogeneity****.*

$$\sum_{i=1}^{N} p_i \frac{\partial c_j}{\partial p_i} + w \frac{\partial c_j}{\partial w} = 0$$

Proof.

↪ **adding-up**. The first result follows from differentiating the budget constraint with respect to income. Analogously, the second follows from differentiating it with respect to p_j and re-arranging terms.

↪ **homogeneity**. Try it yourself!

□

The result above derives the adding-up and homogeneity restrictions in terms of derivatives of the Marshallian demand function $c_i(p, w)$. We need to re-express these restrictions in terms of elasticities to be able to link them with the log-log demand model.

The elasticity formulation of the restrictions requires us to first define the expenditure shares.

Definition 3.2.4 (expenditure share)**.** *Good i's **expenditure share** as the ratio of total expenses on good i divided over the budget:*

$$y_i = \frac{p_i c_i}{w},\ i = \{1, ..., N\}$$

Exercise 3.2.1. *Show that $y_i \in [0, 1]$ for every $i \in \{1, .., N\}$.*

Now, we are ready to formulate the restrictions in terms of demand elasticities.

Proposition 3.2.2. *Suppose that a demand function $c = c(p, w)$ exists, and that the budget constraint holds with equality:*

$$\sum_{i=1}^{N} p_i c_i = w$$

Then the following properties hold:

↪ ***adding-up**.*

$$\sum_{i=1}^{N} y_i e_i = 1 \quad \text{and} \quad \sum_{i=1}^{N} y_i e_{ij} + y_j = 0,\ j \in \{1, .., N\}$$

↪ ***homogeneity**.*

$$\sum_{i=1}^{N}\sum_{j=1}^{N} e_{ij} + e_i = 0,\ i \in \{1, .., N\}$$

Proof. In the proof we use the results from the previous proposition, re-expressing them in terms of elasticities and expenditure shares.

↪ **adding-up**. The first result:

$$1 = \sum_{i=1}^{N} p_i \frac{\partial c_i}{\partial w} = \sum_{i=1}^{N} p_i \frac{\partial c_i}{\partial w}\frac{w}{c_i}\frac{c_i}{w} = \sum_{i=1}^{N} \frac{p_i c_i}{w} e_i = \sum_{i=1}^{N} y_i e_i$$

The second result:

$$\sum_{i=1}^{N} p_i \frac{\partial c_i}{\partial p_j} + c_j = 0 \Leftrightarrow \sum_{i=1}^{N} p_i \frac{\partial c_i}{\partial p_j}\frac{p_j}{c_i}\frac{c_i}{p_j} + c_j = 0 \Leftrightarrow \sum_{i=1}^{N} \frac{p_i c_i}{p_j} e_{ij} + c_j = 0 \Leftrightarrow \sum_{i=1}^{N} \frac{p_i c_i}{w} e_{ij} + \frac{p_j}{w} c_j = 0 \Leftrightarrow \sum_{i=1}^{N} y_i e_{ij} + y_j = 0$$

↪ **homogeneity**. Try it yourself!

□

The adding-up and homogeneity restrictions imply specific relations between elasticities and expenditure shares. The expenditure share are typically observed. The elasticities are functions of parameters that can be estimated. Thus, we can confirm or reject these theoretical restrictions by using a statistical test.

However, in the general case, only one of these properties is directly testable.

Proposition 3.2.3 (testable implications in log-log model)**.** *In the log-log model, one may test:*

1. *the homogeneity property.*
2. *the adding-up property only if $e_i = 1$ for all i*

3.3 The Stone's Model

The first economist who actively tried to link the theoretical foundations with empirical modeling was Richard Stone. He was not yet interested in testing whether the data is generated by a process consistent with the consumer theory. Instead, his strategy was derive a functional form for his estimable demand system to make sure it satisfies the restrictions implied by the theory.

Before we move to describing the Stone's model, we need to introduce the concept of compensated price elasticity of demand.

Definition 3.3.1 (compensated price elasticity of demand)**.** *The* ***compensated price elasticity of demand*** *is defined as:*

$$e^*_{ij} = \frac{\partial \log c^h_i(p, \overline{u})}{\partial \log p_j}$$

It turns out that the (uncompensated) demand elasticities and compensated demand elasticities are related.

Proposition 3.3.1. *For any $i, j \in \{1, .., N\}$:*

$$e_{ij} = e^*_{ij} - e_i y_j$$

Proof. We use two properties of demand functions derived last week. First, the Shephard's lemma implies that:

$$\frac{\partial e(p, \overline{u})}{\partial p_i} = c^h_i(p, \overline{u})$$

Second, recall the relation between the uncompensated and compensated demand functions:

$$c_i\big(p, e(p, \overline{u})\big) = c^h_i(p, \overline{u})$$

Differentiating the latter with respect to p_j:

$$\frac{\partial c_i}{\partial p_j} + \frac{\partial e(p, \overline{u})}{\partial p_j}\frac{\partial c_j}{\partial w} = \frac{\partial c^h_i}{\partial p_j} \underset{\text{substitute for the former}}{\Leftrightarrow} \frac{\partial c_i}{\partial p_j} + c^h_j(p, \overline{u})\frac{\partial c_j}{\partial w} = \frac{\partial c^h_i}{\partial p_j}$$

Again, using the fact that $c^h_j(p, \overline{u}) = c_j(p, w)$, and switching to elasticities, and cleaning the result:

$$\frac{\partial c_i}{\partial p_j}\frac{p_j}{c_i}\frac{c_i}{p_j} + c_j\frac{\partial c_j}{\partial w}\frac{c_j}{w}\frac{w}{c_j} = \frac{\partial c^h_i}{\partial p_j}\frac{p_j}{c^h_i}\frac{c^h_i}{p_j} \Leftrightarrow e_{ij}\frac{c_i}{p_j} + e_i\frac{c_i c_j}{w} = e^*_{ij}\frac{c_i}{p_j} \Leftrightarrow e_{ij}\frac{c_i}{p_j} + e_i\frac{c_i c_j}{w}\frac{p_j}{p_j} = e^*_{ij}\frac{c_i}{p_j} \Leftrightarrow e_{ij}\frac{c_i}{p_j} + e_i y_j\frac{c_i}{p_j} = e^*_{ij}\frac{c_i}{p_j} \Leftrightarrow e_{ij} + e_i y_j = e^*_{ij}$$

The result then follows. □

Now, we are back to the Stone's model. The analysis departs from the log-log demand model:

$$\log c_i = \alpha_i + e_i \log w + \sum_{k=1}^{N} e_{ik} \log p_k, \; i \in \{1, ..., N\}$$

The Stone's model is still derived from the log-log function form. As we learned in the previous section, we can only impose the homogeneity assumption.

In the first step, we substitute e_{ik}s using the result above:

$$\begin{aligned}
\log c_i &= \alpha_i + e_i \log w + \sum_{k=1}^{N} e_{ik} \log p_k \\
&= \alpha_i + e_i \log w + \sum_{k=1}^{N} (e^*_{ik} - e_i y_j) \log p_k \\
&= \alpha_i + e_i\big(\log w - \sum_{k=1}^{N} y_k \log p_k\big) + \sum_{k=1}^{N} e^*_{ik} \log p_k
\end{aligned}$$

Our demand model now depends on real expenditures and compensated prices. It is as if we moved from a Marshallian formulation of demand model to the Hicksian formulation, at least approximately.

To proceed, we need to define a price index – an index that measures the overall level of prices.

Definition 3.3.2 (Stone's price index)**.** *The Stone's price index is given by:*

$$\log P = \sum_{k=1}^{N} y_k \log p_k$$

Stone's formulation for a price index – the average log price weighted by expenditure shares – is one of the most intuitive. Note that if we translate this index to consider the overall level of prices (not their logarithms), we would get a geometric average.

Now, substitute for the price index in the log-log model:

$$\log c_i = \alpha_i + e_i\big(\log \frac{w}{P}\big) + \sum_{k=1}^{N} e^*_{ik} \log p_k$$

The last missing element is a non-homogeneous (of degree zero) term including $\log p_k$. We deal with that starting with a small result concerning the compensated elasticities.

Proposition 3.3.2. *The homogeneity restriction implies that* $\sum_{k=1}^{N} e^*_i k = 0$.

Proof. The homogeneity implies that $\sum_{k=1}^{N} e_{ik} + e_i = 0$. Using the previous result, it implies $\sum_{k=1}^{N}(e^*_{ik} - e_i y_k) + e_i = 0$. Furthermore:

$$0 = \sum_{k=1}^{N}(e^*_{ik} - e_i y_k) + e_i = \sum_{k=1}^{N} e^*_{ik} + e_i \sum_{k=1}^{N} y_k + e_i = \sum_{k=1}^{N} e^*_{ik} + e_i - e_i = \sum_{k=1}^{N} e^*_{ik}$$

□

Now, since $\sum_{k=1}^{N} e^*_{ik} = 0$, then $\sum_{k=1}^{N} e^*_{ik} \log P = 0$, which brings us to the estimating equation for the Stone's model.

Definition 3.3.3 (Stone's demand model)**.** *The* ***Stone's demand model*** *is given by the following equation:*

$$\log c_i = \alpha_i + e_i\big(\log \frac{w}{P}\big) + \sum_{k \in K} e^*_{ik} \log \frac{p_k}{P},\ i \in \{1, ..., N\}$$

Note that the summation goes over $k \in K$ which is a set of close substitutes for food i. This is a way in which Stone aimed to decrease the number of parameters, easily justifiable on the ground of theory. In his work, he aimed to estimate a demand system of $N = 48$ food categories using 19 yearly observations. We can immediately understand why he really had to count the number of parameters in his model.

Even though Stone attempted to accommodate the theory straight into the functional form, we can still say something about how well the estimates fit the theory. For example, we should expect negative signs of own price elasticities of demand, and positive signs of cross price elasticities of demand. Moreover, we are also able to say something about Slutsky symmetry using the estimated elasticities.

3.4 Linear Expenditure System

The **linear expenditure system** (LES), derived also by Richard Stone, is the first demand system entirely based on the theory. Similarly to the Stone's mode above, it focuses on accommodating the theoretical implications of consumer theory into the estimable demand specification. However, as usually, there some implications with additional empirical content can be derived.

Conceptually, we depart as usually from the Marshallian demand system: $c(p, w)$. Consider a general form of linear demand:

$$p_i c_i = \beta_i w + \sum_{j=1}^{N} \beta_{ij} p_j,\ i \in \{1, ..., N\}$$

First, we want to know what the consumer theory implies in relation to parameters of this equation.

Proposition 3.4.1. *Suppose a demand function is linear, and satisfies adding-up, homogeneity, and symmetry. Then it must be given by:*

$$p_i c_i = p_i \gamma_i + \beta_i (w - \sum_{k=1}^{N} \gamma_k p_k)$$

where $\sum_{k=1}^{N} \beta_k = 1$.

You should skip the proof of this proposition.

It turns out that the theory does not leave us much choice regarding the choice of specification.

In the last class, we derived a series of connections between different objects in consumer theory. This is going to be useful here. The following proposition says that the demand function pins the functional forms of the utility, the indirect utility, and the expenditure function. As a result, we can characterize the whole decision environment just by specifying the demand function.

Proposition 3.4.2. *Suppose consumer demand function is given by:*

$$p_i c_i = p_i \gamma_i + \beta_i (w - \sum_{k=1}^{N} \gamma_k p_k)$$

Then:

1. ***expenditure function*** *is given by*
$$e(p, \overline{u}) = \sum_{k=1}^{N} p_k \gamma_k + \overline{u} \prod_{k=1}^{N} p_k^{\beta_k}$$
2. ***indirect utility function*** *is given by:*
$$v(p, w) = \frac{w - \sum_{k=1}^{N} p_k \gamma_k}{\prod_{k=1}^{N} p_k^{\beta_k}}$$
3. ***the utility function*** *is given by:*
$$u(c) = \prod_{k=1}^{N} (c_k - \gamma_k)^{\beta_k}$$

Proof. Try yourself to prove this result. □

The functional form of the utility function suggests interpretation of γ's. If $\gamma_i \geq 0$, then we could perceive it as a minimum level of consumption of good i that is necessary to sustain the consumer, in other words, a **subsistence** quantity of good i.

Furthermore, the demand functional form suggests that the consumer would first acquire the subsistence level of goods – $p_i \gamma_i$ – and then the spend the remaining income $w - \sum p_i \gamma_i$ (also called a supernumerary expenditure) in fixed proportions across goods.

The cost function tells the same story, being a sum of a fixed-cost element (allowing no substitution) $\sum p_i \gamma_i$ and a term increasing utility at a constant marginal price $P = \prod_{k=1}^{N} p_k^{\beta_k}$. We interpret P as a price index measuring the marginal cost of living.

Definition 3.4.1 (LES price index)**.** *The LES price index is a geometric average of prices:*

$$P = \prod_{k=1}^{N} p_k^{\beta_k}$$

where $\sum \beta_k = 1$

Proposition 3.4.3. *The LES expenditure function* $e(p, \overline{u}) = \sum_{k=1}^{N} p_k \gamma_k + \overline{u} \prod_{k=1}^{N} p_k^{\beta_k}$ *is concave if* $\beta_k \geq 0$ *for all* $k \in \{1, .., N\}$, $w \geq \sum_{k=1}^{N} p_k \gamma_k$ *and* $c_k \geq \gamma_k$ *for every* $k \in \{1, .., N\}$.

The above proposition gives a set of theoretical restrictions on the model. Recall that last week we proved that the expenditure function in the consumer theory model is concave. If the estimated parameters of the LES violate the conditions stated above, it means that the data were not obtained by observing agents maximizing their utility subject to the budget constraint.

LES is restrictive in the sense that it imposes a functional form on the demand function. This particular functional from excludes analysis of complementary or inferior goods (they would require negative βs, violating concavity of expenditure function). On the other hand, it presents a parsimonious empirical model with only $2N-1$ independent parameters to estimate.

3.5 Almost Ideal Demand System

The Almost Ideal Demand System (AIDS) by Deaton and Muellbauer (1980) was a significant milestone for the literature of estimable demand systems. It still remains an important tool to study consumer behavior on the markets.

AIDS has been design to test the implications of consumer theory, instead of simply imposing them on the functional form of demand as the demand systems described above. AIDS specification is rich – that means it has many parameters. Consumer theory places restrictions over these parameters. A strategy to testify the theory would be to estimate a general form of the model and statistically test for whether these restrictions hold.

At this point it is worth noticing that AIDS was not the first to have a flexible specification or to allow for direct tests of theoretical implications. The pioneering work includes so called **Rotterdam** demand system and **translog** demand system. However, AIDS is the most coherent and appealing of all these three and this is why it sort of won the race of demand systems (in the space of goods).

We introduce the AIDS using so called Engel curves.

Definition 3.5.1 (Engel curves)**.** *An **Engel curve** is formed by pairs (w, y_i) describing relation between income and expenditure for a given good.*

Engel curves allow us to distinguish necessity goods from luxury goods.

Definition 3.5.2 (necessity good)**.** *Good i is called a **necessity** good if the slope of its Engel curve goes upwards at a decreasing rate.*

Definition 3.5.3 (luxury good)**.** *Good i is called a **luxury** good if the slope of its Engel curve goes upwards at an increasing rate.*

Necessity goods are goods which constitute larger expenditure share among poorer consumers, opposite holds for the luxuries.

The simplest, seemingly a-theoretical empirical model of an Engel curve is given by the following:

$$y_i = \alpha_i + \beta_i \log w,\ i \in \{1, .., N\}$$

Proposition 3.5.1 (adding-up restriction in the simplest Engel curve model)**.** *The adding-up restriction in the simplest Engel curve model is given by:*

$$\sum_{i=1}^{N} \alpha_i = 1 \quad \text{and} \quad \sum_{i=1}^{N} \beta_i = 0$$

Proof. Sum equations of the simplest Engels curve model over goods:

$$\sum_{i=1}^{N} y_i = \sum_{i=1}^{N} \alpha_i + \sum_{i=1}^{N} \beta_i \log w \Leftrightarrow 1 = \sum_{i=1}^{N} \alpha_i + \log w \sum_{i=1}^{N} \beta_i$$

since expenditure share sum up to 1. It must that $\sum_{i=1}^{N} \alpha_i = 1$ and $\sum_{i=1}^{N} \beta_i = 0$. □

The simplest Engel curve model is rich enough to allow us to differentiate between necessity and luxury goods.

Proposition 3.5.2 (necessity and luxury goods in the simplest Engel curve model)**.** *The following hold:*

1. *if $\beta_i < 0$ then good i is a necessity good*

2. *if $\beta_i > 0$ then good i is a luxury good*

The simplicity of simplest Engel curve model may seem inappropriate for modern applications. However, it turns out it generalizes to a broad set of rich models. To see this, we introduce the PIGLOG class of expenditure functions.

Definition 3.5.4 (PIGLOG expenditures)**.** *An expenditure function belongs to the **PIGLOG** class if it can be expressed as:*

$$\log e(p, \overline{u}) = a(p) + ub(p)$$

where $a(b)$ and $b(p)$ are functions of prices.

Proposition 3.5.3. *PIGLOG expenditure functions imply a demand system that can be expressed in the form of the simplest Engel curve model.*

The AIDS model assumes a PIGLOG expenditure function in the following form:

Definition 3.5.5 (AIDS specification)**.** *In AIDS:*

$$\begin{cases} a(p) & = \alpha_0 + \sum_{k=1}^{N} \alpha_k \log p_k + \frac{1}{2} \sum_{k=1}^{N} \sum_{l=1}^{N} \gamma_{kl}^* \log p_k \log p_l \\ b(p) & = \beta_0 \prod_{k=1}^{N} p_k^{\beta_k} \end{cases}$$

The next proposition gives us the AIDS estimable equation.

Proposition 3.5.4. *The AIDS expenditure function induces Engel curves of the following form:*

$$y_i = \alpha_i + \sum_{j=1}^{N} \gamma_{ij} \log p_j + \beta_i \log \frac{w}{P}$$

where the price index is given by:

$$\log P = \alpha_0 + \sum_{k=1}^{N} \alpha_k \log p_k + \frac{1}{2} \sum_{k} \sum_{l} \gamma_{kl} \log p_k \log p_l$$

and $\gamma_{ij} = \frac{1}{2}(\gamma_{ij}^* + \gamma_{ji}^*)$

Proof. First note that:

$$\frac{\partial \log e(p, \overline{u})}{\partial \log p_i} = \underbrace{\frac{\partial e}{\partial p_i}}_{c_i \text{ by Shephard's lemma}} \underbrace{\frac{p_i}{e}}_{\frac{p_i}{w}} = \frac{p_i c_i}{w_i} = y_i$$

Therefore, the Engel curve is just the price elasticity of the expenditure function. Calculating the elasticity using AIDS expenditure functions yields the result. □

As mentioned in the beginning, the AIDS model generates a series of testable implications of consumer theory. The following proposition summarizes these restrictions.

Proposition 3.5.5 (AIDS testable conditions)**.** *The consumer theory generates the following restrictions on AIDS's parameters:*

1. ***adding-up**.*

$$\sum_{k=1}^{N} \alpha_k = 1 \quad \text{and} \quad \sum_{k=1}^{N} \beta_k = 0 \quad \text{and} \quad \sum_{k=1}^{N} \gamma_{kl} = 0$$

2. ***homogeneity**.*

$$\sum_{k=1}^{N} \gamma_{jk} = 0, \ j \in \{1, .., N\}$$

3. ***symmetry**.*

$$\gamma_{jk} = \gamma_{kj}, \ k \neq j \in \{1, .., N\}$$

The beauty of this proposition is that all the restrictions refer directly to the parameters of the model. We acknowledge that the negativity (hessian) condition can also be checked, however it would require more work.

The AIDS price index is more complicated the the previously introduced price indices, and perhaps less intuitive.

Definition 3.5.6 (AIDS price index)**.** *The AIDS price index is given by:*

$$\log P = \alpha_0 + \sum_{k=1}^{N} \alpha_k \log p_k + \frac{1}{2} \sum_k \sum_l \gamma_{kl} \log p_k \log p_l$$

Note that the only source of non-linearity in the AIDS model comes through the price index. Its complicated nature of may make estimation difficulty, particularly if prices are collinear. A solution would be to replace the AIDS price index with another one that you think may be a good approximation, e.g. Stone's index.

Researchers find AIDS good because of its generality. It relies on a flexible specification that can be seen as a first-order approximation to any demand system. Moreover, instead of imposing the consumer theory as it predecessors, it accommodates it as a special case, allowing for statistical testing.

However, as the name indicates, the model is only *almost ideal.* The flexible specification intrinsically requires many parameters to be estimated. This can be troublesome in practice. In the next section, we exploit theory-rooted ways of decreasing the number of parameters in a demand system.

3.6 Aggregation Across Goods

The part of any demand system that generates the largest number of parameters is the dependence of quantity of good i on prices of every other good. A concept of **separability** is helpful in reducing the dimension of a demand system.

To understand the topic better, suppose that the consumption bundle can be partitioned into two parts $(c, d) = \big((c_1, .., c_n), (d_1, ..., d_m$ For example, we could interpret c-goods as food, and d-goods as entertainment. Analogously, we define also the respective partition of prices $(p, q) = \big((p_1, .., p_n,), (q_1, .., q_m)\big)$.

In this setting, the standard consumer maximization problem takes the form of:

$$\max_{c,d} u(c, d)$$
$$\text{such that } p \cdot c + q \cdot d = w$$

We are looking for conditions under which we could study the problem of choosing the optimal quantities of c-goods without having to worry about how the demand is allocated across d-goods.

Formally, we want to ask when there exist:

1. a scalar quantity index $C = g(c)$
2. a scalar price index $P = f(p)$

such that the optimization problem:

$$\max_{C,d} U(C, d)$$
$$\text{such that } PC + q \cdot d = w$$

yields *the same* result as the standard problem. The same means here:

$$C(P, q, w) \equiv C(f(p), q, w) = g\big(c(p, q, w)\big)$$

In other words, we need to be able to obtain the same C regardless of whether we:

1. first aggregate prices P and maximize $U(C, d)$
2. first maximize $u(c, d)$ and then aggregate the result $C = g(c)$

It turns out there are two distinct way in which we can ensure this separability property.

The first is **Hicksian Separability**. Suppose the relative prices within the group of c-goods cannot change, that is, $p = tp^0$ for some fixed reference price vector p^0 and $t > 0$. Define the indices as:

1. a scalar price index $P = t$
2. a scalar quantity index $C = p^0 \cdot c$

Proposition 3.6.1. *Constant relative prices within a group of goods implies separability.*

Proof. With constant relative prices, the indirect utility has the following form:

$$V(P, q, w) = \left\{ \max_{c,d} u(c, d) \text{ such that } Pp^0 \cdot c + q \cdot d = w \right\}$$

An application of Roy's identity:

$$C(P, q, w) = -\frac{\partial V/\partial P}{\partial V/\partial w} = p^0 \cdot c$$

yields the result. To see this, note that we get the same result if we first aggregate prices and then maximize $U(C, d)$ as we get if we maximize $u(c, d)$ and then aggregate quantities. □

The second type of separability is the **Functional Separability**. The functional separability is derived from an assumption on individual's preferences.

Definition 3.6.1 (functional separability of preferences)**.** *We say that preferences are* ***functionally separable*** *if for any four consumption bundles c, c', d, d' in the consumption set, the following holds:*

$$(c, d) \succ (c', d) \Leftrightarrow (c, d') \succ (c', d')$$

The functional separability is a type of an independence assumption between a group of c-goods and d-goods.

Definition 3.6.2 (weakly separable utility)**.** *If the utility function can be written in a form*

$$u(c, d) = U\big(v(c), d\big)$$

where $U(v, d)$ is increasing in v, then we call it ***weakly separable****. $v(c)$ is called a subutility.*

There is a connection between separable preferences and utility functions.

Proposition 3.6.2. *If the preference relation $\succsim$ is functionally separable and locally non-satiated, then a utility function representing $\succsim$ is weakly separable.*

The proof of this claim is beyond the scope of our class.

Proposition 3.6.3. *Let $c(p, q, w)$ and $d(p, q, w)$ be the optimal demands for c-gooda and d-goods respectively, and $w_c = p \cdot c(p, q, w)$ be the optimal expenditure on c-goods. If the overall utility $u(c, d)$ if weakly separable in (c, d), then the demand for c-goods can be found through:*

$$\max_c v(c)$$
$$\text{such that } p \cdot c = w_c$$

Proof. Suppose not. Then $c(p, q, w)$ doesn't solve the above problem. That means there exists c' that satisfies $p \cdot c' = w_c$ and yields strictly higher subutility. However, that would imply that $u(c', d) > u(c, d)$ which contradicts the definition of demand function. □

The last proposition provides a useful result stating that as long as we know w_c we can solve for optimal demand structure within c-goods. d-goods affect it only through impacting the expenditure w_c.

Definition 3.6.3 (conditional demand function)**.** *A consumption bundle $c(p, w_c)$ that solves the subutility maximization problem is called the* ***conditional demand*** *function.*

A related concept is **two-stage budgeting**. In two stage-budgeting, the consumer first decides about how much of their expenditure would be spent on broader categories, e.g. c- and d- goods. In the second step, the consumer chooses the exact quantities within the categories.

Before we proceed, we need to introduce the concept of homothetic preferences.

Definition 3.6.4 (homothetic utility)**.** *Let $u : C \to \mathbb{R}$ be a utility function representing some preferences $\succsim$. We say that u is homothetic if it is a strictly increasing transformation of a function that is homogeneous of degree 1.*

The homothetic preferences have implications on expenditure.

Proposition 3.6.4. *Suppose the utility function representing consumer preference is homothetic. Then:*

$$e(p, \overline{u}) \equiv \overline{u}h(p), \text{ for some function exclusively of prices } h(p)$$

The property above will be key to determine when two-stage budgeting is consistent with consumer theory. To see this, let $e(p, v)$ be expenditure function for the subutility v.

Consider the overall utility maximization problem:

$$\max_{v,d} U(v, d)$$
$$\text{such that } e(p, v) + q \cdot d = w$$

The expenditure function tells us how much we need to spend on c goods to achieve the subutility v. If preferences are homothetic, $e(p, v) \equiv ve(p)$. Then, we can choose the price index to be $P = e(p)$ and the quantity index $C = v(c)$. That makes the budget constraint linear, so the dual problem would yield the same solution. So homothetic preferences are necessary for two-stage budgeting.

Separability allows us to analyze the demand sequentially, within nodes defined by partitioned consumption set. This may lead to significant reduction in dimensionality of the demand system.

3.7 Aggregation Across Consumers

Lastly, we consider another type of aggregation – aggregation across consumers. The main question to ask is when can we aggregate individual decisions such that they work as if the aggregate consumption was chosen by a representative consumer.

Definition 3.7.1 (aggregate demand function)**.** *Suppose there are I consumers in the economy. The* ***aggregate demand function*** *is the sum of individual demands of all agents in the economy. With a abuse of notation, we write:*

$$C(p, w^1, ..., w^I) = \sum_{i=1}^{I} c(p, w^i)$$

Notably, the aggregate demand function depends on prices and the distribution of incomes of all agents.

In general, consumer theory does not produce restrictions on aggregate behavior, so we cannot say much about the aggregate demand in the general case. However, there is a special case, in which we can.

Theorem 3.1 (Gorman aggregation)**.** *The indirect utility function for every consumer is of the Gorman form:*

$$v_i(p, w^i) = a_i(p) + b(p)w^i$$

if and only if there is a representative consumer with indirect utility:

$$V(p, W) = \sum_{i=1}^{I} a_i(p) + b(p)W = A(p) + B(p)W$$

with $W = \sum_{i=1}^{I} w^i$

The Gorman theorem states that the set of preferences for which the exact consumer aggregation is possible is pretty restrictive. This is why the modern macro-economics departed from the representative agent models, focusing more on the heterogeneous agents models.

4 Producer Theory

The last two meetings taught us how the consumer makes their optimal consumption decisions. However, we did not take any stand on where the consumption goods come from. In this and the next class, we study producer theory. Producer theory describes how and why some goods are transformed in other goods. The output of such a transformation may be the consumption good that our consumer was so vividly interested last week.

This class lies foundation for producer theory. We develop the production environment and discuss various associated assumptions. Then we move to analyzing firm's decision. We carefully study cost minimization and profit maximization, deriving the optimal solutions and their properties that can be later on verified with data. We finish with a powerful result – Topkis theorem – allowing us to study properties of optimal solutions even in very complicated and abstract models. All of these is a necessary step before we dive into how the theory is used together with data to answer the most relevant economic questions, which will happen next week.

List of contents. Net output. Production plan. Production possibility set. Feasible technologies. Restricted production possibility set. Input requirement set. Regular, monotone, and convex input requirement sets. An isoquant. Diminishing marginal returns in production. Production function. Marginal product. Marginal rate of technical substitution. Elasticity of substitution. Returns to scale: constant, increasing, decreasing. Frequently assumed production function functional forms. (Total) cost function. Conditional input demand function. Average cost function. Marginal cost function. Minimal efficient scale of production. Fixed and variable inputs. Short-run cost function. Short-run conditional input demand function. Profit maximization problem. Profit function. (Unconditional) input demand function. Supply function. Lattices and lattice operations. Supermodular functions. Functions with non-decreasing and non-increasing differences. Topkis theorem.

Textbook references. Espinola: Ch. 7-8; Varian: Ch. 1-6; Jehle & Reny: Ch. 3; MWG: Ch. 5

4.1 Mathematical Glossary

As usually, we begin with a set of mathematical tools that will be used in today's class.

Definition 4.1.1 (homogeneous function)**.** *We say that a multi-dimensional function $f : X \to Y$ is* ***homogeneous*** *of degree $\alpha > 0$ if for every $t > 0$:*

$$f(tx) = t^\alpha f(x)$$

Definition 4.1.2 (homothetic function)**.** *We say that a multi-dimensional function $f : X \to \mathbb{R}$ is* ***homothetic*** *if it can be expressed as a strictly increasing transformation of a homogeneous of degree one function. Formally, $f(x)$ is homothetic if it can be expressed as a composition of function $g : X \to \mathbb{R}$ that is homogeneous of degree one and $\phi : \mathbb{R} \to \mathbb{R}$ that is strictly increasing:*

$$f(x) = \phi\big(g(x)\big)$$

Definition 4.1.3. *The* ***image*** *of a function $F : X \to Y$ is the set of all output values that the function can produce:*

$$Im(F) = \{F(x) : x \in X\}$$

where X is the domain and Y is the codomain.

Definition 4.1.4 (inverse of a function)**.** *Let $f : X \to Y$ be a strictly increasing uni-dimensional function. Its* ***inverse*** *is a function $f^{-1} : Y \to X$ assigning to every point of codomain ($y \in Y$) its respective point in the domain ($x \in X$):*

$$f^{-1}(y) = \{x : f(x) = y\}$$

Proposition 4.1.1. *The inverse of a strictly increasing function is a strictly increasing function.*

Proposition 4.1.2. *Suppose that F is a homothetic function so that it can be written as $F(x) = f(g(x))$, where f is strictly increasing, and g is homogeneous of degree one. If $Im(F) = R_+$ then $f^{-1}(y) > 0$ for all $y > 0$.*

4.2 Production Technology

Suppose there are M goods in the economy. Producer theory explains how a subset of goods is transformed into another subset of goods. Intuitively, production transforms a set of inputs into a set of outputs. Nevertheless, many goods can be outputs in one technology and inputs in another. Our notation needs to be general enough to account for this fact.

Definition 4.2.1 (net output)**.** *Consider a technology that uses y_i^{inp} units of good i as inputs and produces y_i^{out} units of good i. The **net output** y_i is the difference between the net outputs and net inputs:*

$$y_i = y_i^{\text{out}} - y_i^{\text{inp}},\ i \in \{1, ..., M\}$$

We will focus our attention on net outputs.

Definition 4.2.2 (production plan)**.** *A **production plan** is a M-dimensional vector y listing the quantities of net production outputs.*

Definition 4.2.3 (production possibility set)**.** *The set of all production plans, $Y \subset R^M$, is called the **production possibility set**.*

The set of production possibility fully describes **feasible** technology. We may be interested in placing restrictions on what is technologically feasible.

Example 4.2.1 (production inputs fixed in the short term)**.** *We often assume that a subset of production factors cannot be adjusted in the short term. Suppose in the short term only $\overline{y}_i$ is available for some good $i \in \{1, ..., M\}$. The resulting **restricted** production possibility set is given by:*

$$Y(\overline{y}_i) = \{y \in Y : y_i = \overline{y}_i\}$$

This way of defining technology is very general. In our class, we will only analyze production settings in which there are N production inputs (i.e. negative net outputs using the general language) and a single output (i.e. positive net output using the general language). In this setting $M = N + 1$. For sake of notational convenience, we adjust a little bit the notation compared to the general framework. From now on, we denote the set of inputs by $X \subseteq \mathbb{R}_+^N$ with a typical element $x \in X$, expressing the inputs in terms of non-negative values (which in terms of the general model would be written as $(-x, y) \in Y$).

We may be interested in production plans that allow us to produce a certain level of output.

Definition 4.2.4 (input requirement set)**.** *The **input requirement** set is a restriction over the production possibility set containing all production plans which include sufficient quantities of inputs to produce (single) output level $\overline{y} > 0$:*

$$IR(\overline{y}) = \{x \in X : (-x, \overline{y}) \in Y\}$$

The input requirement set is fairly general object. You can think of it in terms of describing feasibility.

Definition 4.2.5 (isoquant)**.** *The **isoquant** of production technology is a set of inputs enabling the maximal production level $\overline{y} > 0$:*

$$IQ(\overline{y}) = \{x \in X : x \in IR(\overline{y} \text{ and } x \notin IR(y') \text{ for } y' > y\}$$

The interpretation of an isoquant goes beyond feasibility, as it relates to the highest feasible production level.

At this stage, our definition of technology is very broad, including many unintuitive production technologies that we may want to discard. Therefore, the question is: what else do we want to assume on the production process?

Definition 4.2.6 (regularity)**.** *Production is **regular** if for every $\overline{y}$ the set $IR(\overline{y})$ is closed and nonempty.*

The regularity means that there is a way to produce every positive quantity of the output. Closedness, implying roughly that $V(\overline{y})$ includes its own boundary, is a technical assumption enabling well defined optima in subsequent analysis.

Definition 4.2.7 (monotonicity)**.** *Production is **monotone** if for every production level $\overline{y}$:*

$$x \in IR(\overline{y}) \text{ and } x' \geq x \Rightarrow x' \in IR(\overline{y})$$

Monotonicity implies that if you can produce $\overline{y}$ units of output with certain levels of inputs, then you can produce the same $\overline{y}$ with higher levels of these inputs.

Definition 4.2.8 (convexity)**.** *Production is **convex** if for every two elements of the input requirement set $x, x' \in IR(\overline{y})$ and $\lambda \in [0, 1]$, the following holds:*

$$\lambda x + (1 - \lambda)x' \in IR(\overline{y})$$

Convexity is useful because it enables us to use the tools of calculus in search of the optimal firm's decisions. It also has a desired interpretation in terms of **diminishing marginal returns** which we describe in details later.

In our setting, in which there is a single production output, it is convenient to describe the production technology using the concept of production function.

Definition 4.2.9 (production function)**.** *Function $f : X \to \mathbb{R}_+$ is called a **production function** if it describes the highest quantity of production output y that can be produced using vector if production inputs x:*

$$y = f(x)$$

The production functions is essentially the boundary of the production possibilities set. Not surprisingly, the properties of production functions are related to the properties of sets constructed from the production possibilities set.

Proposition 4.2.1. *$IR(\overline{y})$ is a convex set if and only if the production function is quasi-concave.*

Proof. $IR(\overline{y}) = \{x : f(x) \geq y\}$ is the upper contour set of $f(x)$. A function is quasi-concave if and only if it has a convex upper contour set. □

Therefore, it does not really matter whether we make certain assumptions on the technology expressed through production possibility sets or via a production function. The latter seems more intuitive so we stick to it in our class. From now on, we will impose the following a set of assumptions on production functions in our analysis.

Definition 4.2.10 (assumptions on production function)**.** *In what follows, we assume that the production function $f : X \to \mathbb{R}_+$ is:*

1. *continuous*
2. *strictly increasing*
3. *strictly quasi-concave*
4. $f(0) = 0$ *(possibility of inaction)*

There are many analogies between assumptions on production functions in producer theory and assumptions on utility functions in consumer theory. They both have technical importance, enabling us to use familiar tools in calculus to analyze optimal choices. They also carry deeper meanings, allowing us to describe the observed economic processes with language of mathematics.

However, in the contrast to preferences in consumer theory, the production technology is not derived from a set of more abstract axioms. Economics typically assumes that production is a sort of a black box in which for certain level of inputs we receive certain level of outputs without getting into details of how it actually works. This approach may be good because it allows us to apply the same framework to study different production processes – across many industries, on micro or macro level, and even in pretty intangible markets such as education or skill formation. This approach may be bad if the complexities of a given production process are not well represented by the general production function approach.

Having defined the production function, we define a series of objects analogous to what we defined using the utility function while studying the consumer theory.

Definition 4.2.11 (marginal product)**.** *The **marginal product** of input i is the rate at which the production changes if one increases i by a unit, holding other inputs fixed:*

$$MP_i(x) = \frac{\partial f}{\partial x_i}(x),\ i \in \{1, .., N\}$$

Concavity of production function implies the law of diminishing marginal product stating that marginal product of each input (eventually) decreases as we increase the quantity of that input in use. This is fairly common and intuitive assumption. Think about the relative increment of work that can be done if you hire an additional employee when initially there is only one, compared to an analogous increment in a situation in which you already have a thousand employees.

Next, we define the analog of an indifference curve.

Definition 4.2.12 (isoquant of production function)**.** *The **isoquant of production function** is a set of inputs that yield the same maximal production level $\overline{y}$:*

$$IS(\overline{y}) = \{x \in X : \overline{y} = f(x)\}$$

While analyzing production, we are often interesting in studying the substitution patterns across inputs.

Definition 4.2.13 (marginal rate of technical substitution)**.** *The **marginal rate of technical substitution** (MRTS) tells us how many units of good i we need to change if we change the amount of good j by a unit in the production plan such that we keep the same level of maximal production (or equivalently, we stay at the same isoquant):*

$$MRTS_{ij}(x) = \frac{\partial f/\partial x_i}{\partial f/\partial x_j}$$

MRTS is analogous to MRS in the consumer theory, and inherits the same properties. For example, in two dimensional input set, MRTS is the slope of the isoquant at a given point x.

Another measure of substitution is the elasticity of substitution.

Definition 4.2.14 (elasticity of substitution)**.** *The **elasticity of substitution** in a production function measures how easily one input can be substituted for another while maintaining the same level of output. Formally, it is defined as the percentage change in the ratio of two inputs divided by the percentage change in their marginal rate of technical substitution (MRTS).*

$$\sigma_{ij}(x) = \frac{\partial \frac{x_j}{x_i}}{\partial MRTS_{ij}} \frac{MRTS_{ij}}{\frac{x_j}{x_i}}$$

You can think of the elasticity of substitution in terms of describing how flexible a production process is in switching between different inputs. For example, if the elasticity of substitution is high, it means that the inputs can be easily substituted for each other without significantly affecting production. Conversely, a low elasticity indicates that the inputs are not easily interchangeable.

An important feature of production functions are returns to scale.

Definition 4.2.15 (returns to scale)**.** *Production technology described by a production function $f : X \to \mathbb{R}_+$ exhibits:*

1. ***constant returns to scale** (CRS) if for every $t > 0$, $f(tx) = tf(x)$.*
2. ***increasing returns to scale** (IRS) if for every $t > 1$, $f(tx) > tf(x)$.*
3. ***decreasing returns to scale** (DRS) if for every $t > 1$, $f(tx) < tf(x)$.*

Returns to scale describe how the production reacts to a proportional increase in every input. The returns are constant, if a proportional increase in output follows. The returns are decreasing, if the output increases less than proportionally. These two situations are most often expected to hold in the real world. If the increase in output is more than proportional, it may create incentives to unboundedly increasing the inputs which is not something observed in the data.

Exercise 4.2.1. *We say that the input requirement set satisfies constant returns to scale if $x \in IR(\overline{y}) \Rightarrow tx \in IR(t\overline{y})$ for all $\overline{y} > 0$ and $t > 0$. Show that the input requirement set satisfies constant returns to scale if and only if the production function exhibits CRS.*

The concept of returns to scale is essentially a description of the degree of homogeneity of a production function. By definition, a production function exhibits constant returns to scale if and only if it is homogeneous of degree one. Increasing (decreasing) returns to scale mean the degree of homogeneity if above (below) one.

Homogeneity (and thus returns to scale) has direct implications on other properties of a function.

Proposition 4.2.2. *Let $f : X \to \mathbb{R}_+$ be a production function. If f is homogeneous of degree $\alpha \in (0, 1]$ then $f(x)$ is a concave function of x.*

Proof. Suppose first that $\alpha = 1$, so our production function is homogeneous of degree one. Take any $x^1, x^2 \gg 0$ and let $y^1 = f(x^1)$ and $y^2 = f(x^2)$. Since f is strictly increasing and $f(0) = 0$, it follows that $y^1 > 0$ and $y^2 > 0$.

Since f is homogeneous of degree one:

$$y^1 = f(x^1) \Rightarrow y^1 = y^1 f\Big(\frac{x^1}{y^1}\Big) \Rightarrow f\Big(\frac{x^1}{y^1}\Big) = 1$$

The same follows for the second plan, i.e. $f\Big(\frac{x^2}{y^2}\Big) = 1$.

Since f is strictly quasi-concave:

$$f\Big(\lambda\frac{x^1}{y^1}+(1-\lambda)\frac{x^2}{y^2}\Big)\geq \min\{f\Big(\frac{x^1}{y^1}\Big),f\Big(\frac{x^2}{y^2}\Big)\}=1 \text{ for any } \lambda\in(0,1)$$

In particular it has to hold for $\lambda=\frac{y^1}{y^1+y^2}$ and thus $1-\lambda=\frac{y^2}{y^1+y^2}$:

$$f\Big(\frac{y^1}{y^1+y^2}\frac{x^1}{y^1}+\frac{y^2}{y^1+y^2}\frac{x^2}{y^2}\Big)=f\Big(\frac{x^1}{y^1+y^2}+\frac{x^2}{y^1+y^2}\Big)\geq 1$$

Since f is homogeneous of degree one:

$$f\Big(\frac{x^1}{y^1+y^2}+\frac{x^2}{y^1+y^2}\Big)\geq 1\Rightarrow\frac{1}{y^1+y^2}f(x^1+x^2)\geq 1\Rightarrow f(x^1+x^2)\geq y^1+y^2=f(x^1)+f(x^2)$$

where the last equality follows by definition of y^1 and y^2. This holds for any $x^1,x^2\gg 0$ but invoking continuity of f we can extend this for every $x^1,x^2\geq 0$.

Lastly, for any $x^1,x^2\geq 0$ and $\lambda\in[0,1]$, we get that (by homogeneity)

$$\begin{cases} f(\lambda x^1)=\lambda f(x^1)\\ f((1-\lambda)x^2)=(1-\lambda)f(x^2)\end{cases}$$

Using the last three results:

$$f\big(\lambda x^1+(1-\lambda)x^2\big)\geq f(\lambda x^1)+f\big((1-\lambda)x^2\big)=\lambda f(x^1)+(1-\lambda)f(x^2)$$

which yields the result for $\alpha=1$. For $\alpha<1$, first note that $f^{\frac{1}{\alpha}}$ is homogeneous of degree one, hence the result above holds: $f^{\frac{1}{\alpha}}$ is concave. But then $f=\Big(f^{\frac{1}{\alpha}}\Big)^\alpha$ is also concave, as $\alpha<1$.

□

This proposition ensures existence of a well defined optimum in the firm's problem if the production function exhibits constant of decreasing returns to scale. It singles out production functions with increasing returns to scale, for which we have already built some suspicion.

Although useful, requiring homogeneity is pretty restrictive. It turns out that we can generalize this concept and still maintain a lot of useful properties of a function. The generalization concerns homothetic functions, which are strictly increasing transformations of functions homogeneous of degree one. Homothetic functions are a rich class of functions, hence useful for our analysis as they can accommodate many shapes of production technology as special cases.

This brings us to the choice of functional form for the production function.

Example 4.2.2 (frequently used production functions)**.** *Let $X=\mathbb{R}^2_+$ with a typical element $x=(x_1,x_2)$. Frequently, the following utility functions are assumed to describe the production technology:*

description	**function**
perfect substitutes	$f(x_1,x_2)=\alpha x_1+\beta x_2,\ \alpha,\beta>0$
perfect complements (Leontief)	$f(x_1,x_2)=\min\{\alpha x_1,\beta x_2\},\ \alpha,\beta>0$
Cobb-Douglas	$f(x_1,x_2)=x_1^\alpha x_2^\beta,\ 0\geq\alpha+\beta\leq 1$
CES	$f(x_1,x_2)=(\alpha x_1^\rho+\beta x_2^\rho)^{\frac{1}{\rho}},\ \alpha,\beta>0$

CES stands for for Constant Elasticity of Substitution.

Exercise 4.2.2. *Show that the CES exhibits constant elasticity of substitution.*

Exercise 4.2.3. *Which of the above production functions are homogeneous of degree one? Which of them are homothetic?*

4.3 Cost Functions

What is the producer's problem? It may depend on the environment in which they operate. The optimal production level may vary with the consumer demand for firm's products as well as the composition of firm's competitors and their actions. Before we analyze firm's profit maximization which will require us to take a stand on such factors, we study cost functions. The cost functions tell us what is the minimum level of cost that enables producing y units of output.

The cost is essentially a monetary term. Hence, we need to introduce money to our model. Suppose input prices are characterized by a strictly positive vector $w \gg 0$, which is takes as given in our analysis.

Definition 4.3.1 (cost function)**.** *The* ***cost function*** *is a solution to the cost minimization problem at a given production output level y:*

$$c(w, y) = \min_{x \in X} w \cdot x$$
$$\text{such that } f(x) \geq y$$

Thus, if $x(p, w)$ solves the minimization, then $c(w, y) = w \cdot x(w, y)$.

Definition 4.3.2 (conditional input demand functions)**.** *The optimal solution to the cost minimization problem is called the* ***conditional input demand function*** *and is denoted by $x(w, y)$.*

The word *conditional* in the conditional input demand functions indicates that we condition on a particular output level y. So far, y is arbitrary and does not need to be a profit maximizer (or any-maximizer).

Proposition 4.3.1 (properties of the cost function)**.** *The cost function $c(w, y)$ satisfies:*

1. $c(w, 0) = 0$,
2. *is continuous,*
3. *for every $w \gg 0$, is strictly increasing and unbounded above in y,*
4. *is increasing in w,*
5. *is homogeneous of degree one in w,*
6. *is concave in w,*
7. *if $c(w, y)$ is differentiable then the cost function satisfies Shephard's lemma:*

$$\frac{\partial c(w, y)}{\partial w_i} = x_i(w, y)$$

Exercise 4.3.1. *Prove the above proposition.*

More important properties can be derived from homotheticy of production function.

Proposition 4.3.2. *Suppose the production function satisfies assumptions of this chapter and is in addition homothetic. Then:*

1. *the cost function is multiplicatively separable in w and y:*

$$c(w, y) = h(y)c(w, 1)$$

   where $h(y)$ is a strictly increasing function solely depending on production level, and $c(w, 1)$ is the minimal cost of producing one unit of output.

2. *the conditional input demand function is multiplicatively separable in w and y:*

$$x(w, y) = h(y)x(w, 1)$$

   where $h(y)$ is a strictly increasing function solely depending on production level, and $x(w, 1)$ is the conditional input demand associated with producing one unit of output.

Proof. Let f be a homothetic production function. That means, we can write it down as:

$$f(y) = \phi\big(g(y)\big)$$

where $\phi(\cdot)$ is a strictly increasing function, and $g(\cdot)$ is homogeneous of degree one. For simplicity, assume that when varying y, f attains any weakly positive value (the image of f is equal to $\mathbb{R}_+$). We know (see the math glossary) that in such a case $\phi^{-1}(y) > 0$ for all $y > 0$. Hence, fix $y > 0$ and consider $t = \frac{\phi^{-1}(1)}{\phi^{-1}(y)}$. It must be that $t > 0$ by previous observation. Now:

$$\phi\big(g(x)\big) \geq y \Leftrightarrow g(x) \geq \phi^{-1}(y)$$

Moreover, using the fact that g is homogenous of degree 1, the above holds if and only if:

$$g(tx) = tg(x) \geq t\phi^{-1}(y) = \frac{\phi^{-1}(1)}{\phi^{-1}(y)}\phi^{-1}(1)$$

Lastly, putting ϕ on both sides, this is equivalent to:

$$\phi\big(g(tx)\big) \geq \phi\big(\phi^{-1}(1)\big) = 1$$

So the cost function is defined through:

$$\begin{aligned}
c(w, y) &= \min_{x \in \mathbb{R}^N_+} \; w \cdot x \text{ such that } \phi\big(g(x)\big) \geq y \\
&= \min_{x \in \mathbb{R}^N_+} \; w \cdot x \text{ such that } \phi\big(tg(x)\big) \geq 1 \\
&= \frac{1}{t} \min_{x \in \mathbb{R}^N_+} \; w \cdot tx \text{ such that } \phi\big(tg(x)\big) \geq 1 \\
&= \frac{1}{t} \min_{z \in \mathbb{R}^N_+} \; w \cdot z \text{ such that } \phi\big(tg(z)\big) \geq 1 \\
&= \frac{\phi^{-1}(y)}{\phi^{-1}(1)} c(w, 1) \\
&= h(y)c(w, 1)
\end{aligned}$$

Since ϕ is strictly increasing, so is ϕ^{-1} hence the result holds for every $y > 0$. For $y = 0$, note that $c(w, 0) = 0$ because $f(0) = 0$ by assumption. Moreover, $g(0) = 0$ by homogeneity of first degree. The second part from Shephard's lemma for the cost function. $\square$

The multiplicative separability implies additive separability in logs:

$$\log c(w, y) = \log h(y) + \log c(w, 1)$$

and is a fantastic result from the estimation point of view. It drastically reduces the amount of variation in the data necessary to recover the underlying cost function. Intuitively, instead of requiring to observe co-movement in (w, y), it is enough for us to focus on variation in w and y separately to be able to recover the cost function. Moreover, (log)-linearity enables use of the standard, fast and simple empirical tools in estimation.

Not surprisingly, the result extends to the homogeneous functions, and by that – to functions exhibiting constant returns to scale. You are invited to prove these two simple propositions.

Exercise 4.3.2. *Prove the following statement: Suppose the production function satisfies assumptions of this chapter and is in addition homogeneous of degree $\alpha > 0$. Then:*

1. $c(w, y) = y^{\frac{1}{\alpha}} c(w, 1)$
2. $x(w, y) = y^{\frac{1}{\alpha}} x(w, 1)$

Exercise 4.3.3. *Show that if f exhibits constant returns to scale then*

1. $c(w, y) = yc(w, 1)$
2. $x(w, y) = yx(w, 1)$

Based on the cost function, also known as **total cost function**, we can build more related objects that allow us to study the optimal behavior in more detail. In a few following results we focus on variation in costs at induced by the output. For notational simplicity, we are going to drop the dependence on w (which is held fixed and taken as given anyways). Thus, $c(w,y) \equiv c(y)$.

Definition 4.3.3 (average cost function)**.** *The* ***average cost function*** *is the total cost divided by the quantity of output:*

$$ac(y) = \frac{c(y)}{y}$$

Definition 4.3.4 (marginal cost function)**.** *The* ***marginal cost function*** *describes the additional cost incurred by producing one more unit of output. Formally, it defined as the derivative of the total cost function with respect to the quantity of output:*

$$mc(y) = \frac{\partial c(y)}{\partial y}$$

Both the average cost and marginal cost functions are measures of cost associated with production and depend explicitly on y. ac provides an overall measure of the cost efficiency of production. It explains how the total cost is spread over the units produced. In turn, mc provides a measure of the incremental cost of production. ac and mc provide a useful characterization of (long-term) optimal costs.

Definition 4.3.5 (minimal efficient scale)**.** *The* ***minimal efficient scale*** *refers to the production level at which the average cost is minimized (usually: in the long term).*

The minimal efficient scale represents the point at which the firm fully exploits economies of scale, and any further increase in output does not result in a significant decrease in average cost.

Proposition 4.3.3. *At the minimal efficient scale, $mc(y) = ac(y)$.*

Proof. Fix w and write $c(w,y) \equiv c(y)$. Differentiate the definition of ac with respect to y:

$$ac'(y) = \frac{c'(y)y - c(y)}{y^2} = \frac{mc(y) - ac(y)}{y}$$

If $mc(y) > ac(y)$ then$ac'(y) > 0$ so $ac(y)$ is increasing. If $mc(y) < ac(y)$ then$ac'(y) < 0$ so $ac(y)$ is decreasing. So it must be that $ac(y)$ reaches its minimum when $mc(y) = ac(y)$, and thus $ac'(y) = 0$. □

It may be that some inputs are fixed (cannot be adjusted) in the short run. Consider a manufacturing firm that uses both labor and capital to produce goods. In the short term, the amount of capital (e.g., machinery, equipment, factory buildings) is fixed because it takes time to acquire, install, or modify these assets. Therefore, the firm cannot easily adjust its capital stock in response to changes in production needs. To account for that we distinguish short and long term cost functions.

Before we move to analyze differences between these two, we need to introduce more notation. Define a partition of the space of inputs $x = (x^v, x^f)$ where x^v are inputs that can be adjusted even in the short term – the **variable** inputs, and x^f are inputs that cannot be adjusted in the short term – the **fixed inputs**. Suppose that the fixed inputs assume values $x_i^f = \overline{x}_i^f$ for every i in the set of fixed inputs. Define also an analogous partition of the input price vector $w = (w^v, w^f)$.

Definition 4.3.6 (short-run cost minimization)**.** *The* ***short-run cost function*** *$c^{SR}(w, y, \overline{x}^f)$ represents the minimum cost of producing a given level of output y when fixed inputs take value $\overline{x}^f$:*

$$\begin{aligned} c^{SR}(w, y, \overline{x}^f) = \min_{x \in X} \; & w \cdot x \; \left(= w^v \cdot x^v + w^f \cdot x^f \right) \\ \text{such that } & f(x) \geq y \\ & x_i^f = \overline{x}_i^f, \quad \text{for every fixed input } x_i^f \end{aligned}$$

Let $x(w, y, \overline{x}^f)$ be the ***short-run*** *(constrained)* ***input demand function****. Then:*

$$c^{SR}(w, y, \overline{x}^f) = \underbrace{w^v \cdot x^v(w, y, \overline{x}^f)}_{\text{total variable cost}} + \underbrace{w^f \cdot \overline{x}^f}_{\text{total fixed cost}}$$

In the long term, all of the inputs are freely adjustable so the long term cost function coincides with the definition of the total cost function discussed above.

It turns out that the short-term and long-term cost function are related to each other in a rather fascinating way.

Proposition 4.3.4. *The long-run total cost curve is the lower envelope of the entire family of short-run total cost curves.*

Proof. Let the the optimal quantity of fixed inputs, i.e. one that maximizes the short-term cost function, be $x^f(y)$. Then:

$$c(w,y) = c^{SR}(w,y,\overline{x}^f)$$

Note that since we chose the fixed inputs to maximize the short term cost functions, it must be that

$$\frac{\partial c^{SR}(w,y,\overline{x}^f)}{\partial x_i^f} = 0$$

Now, differentiate the equation linking the long-term cost function with the short-term cost function:

$$\begin{aligned}\frac{dc(w,y)}{dy} &= \frac{\partial c^{SR}(w,y,\overline{x}^f)}{\partial y} + \sum_{i\in\text{fixed inputs}} \underbrace{\frac{\partial c^{SR}(w,y,\overline{x}^f)}{\partial x_i^f}}_{=0} \frac{\partial x_i^f}{\partial y} \\ &= \frac{\partial c^{SR}(w,y,\overline{x}^f)}{\partial y}\end{aligned}$$

Now we are ready to collect three facts:

1. Since the constrained (short-term) cost minimization includes in general more restrictions, it cannot yield lower value than the total cost:
$$c(w,y) \leq c^{SR}(w,y,\overline{x}^f)$$
so the graph of short-term cost will lie above the graph of the long-term costs in the (y,c) space.
2. For each y, every short-run cost function (depending on various values of fixed inputs) and the long-run cost would be equal to each other at at least one point.
3. The slope of the short-run cost function is equal to the slope of the long-run cost function in the (y,c) space at their intersection point.

The last two observations imply that the graph of the short-run cost function has to be tangent to the graph of the long-run cost function at their intersection point. This, together with the first observation, yields the result. □

Example 4.3.1. *The following graph, i.e. figure 3.6 borrowed from the Jehle and Reny's textbook, visualizes points made in the proposition above:*

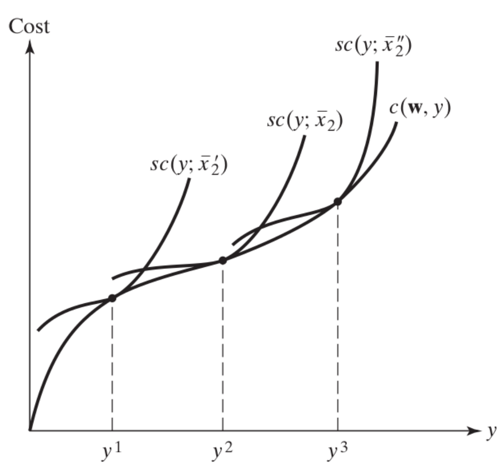

Figure 3.6. Long-run total cost is the envelope of short-run total cost.

Our last two propositions in this section provide results with important empirical content.

Proposition 4.3.5. *Let $c : R^N_{++} \times \mathbb{R}_+ \to \mathbb{R}_+$ be a total cost function satisfying all the properties of proposition 4.3.1. Then the function $f : X \to \mathbb{R}_+$ defined by:*

$$f(x) = \max_{y \geq 0} \left\{ y \text{ such that } w \cdot x \geq c(w, y) \text{ for every } w \gg 0 \right\}$$

is increasing, unbounded above and quasi-concave. Moreover, the cost function generated by f is c.

This proposition (which you should not attempt to prove) establishes a sense of equivalence between the cost function and the production function. We already know that we can derive a cost function from a production function. Now, we also know we can go the other way around: given a cost function, we can reconstruct the production function. This is very important empirically, because it requires us to observe only one from the firm's cost and production aspects to be able to reconstruct the whole environment.

But, do we believe so? Well, it can be formally checked. The producer theory generates a range of testable restrictions on the firm's optimal behavior (which are very similar to those generated by the consumer theory). Specifically, we focus on the properties of the conditional input demand.

Proposition 4.3.6. *A continuously differentiable function $x(w, y) : R^N_{++} \times \mathbb{R}_+ \to \mathbb{R}_+$ is the conditional input demand function generated by some strictly increasing, quasi-concave production function if and only if it is:*

1. *homogeneous of degree zero in w,*
2. *its substitution matrix defined by $\frac{\partial x_i(w,y)}{\partial w_j}$ is symmetric and negative definite,*
3. *$w \cdot x(w, y)$ is strictly increasing in y.*

4.4 Production on Competitive Markets

What motivates firm to produce is its profit from selling the output. As we already agreed, the functional form of the profit function depends on the type of a market in which the firm operates. Today, we consider the simplest form of the market structure – perfect competition. In perfectly competitive markets, firms take as given the price of the final good (their output). It is set by the market. In particular, firm's choice of y does not affect the market price. This a consequence of the assumption that there are *many* firms on the competitive market and each firm has infinitesimal market share. This also implies that any *reasonable* level of the output y chosen by each firm would sell – so we do not need to put restrictions on the sellable quantities, or bring into the picture consumer's demand.

Now, we are ready to define firm's profit maximization problem in the competitive market.

Definition 4.4.1 (firm's profit maximization)**.** *The competitive firm's* ***profit maximization problem*** *is defined as:*

$$\max_{x \in X} pf(x) - w \cdot x$$

This is an unconstrained optimization. The assumptions that we have already made on the production function ensure a well defined interior maximum. By first order conditions, the maximum satisfies:

$$\frac{\partial f}{\partial x_i} = \frac{w_i}{p}$$

i.e. the marginal product of input i has to be equal to the its price (in real terms – value of output).

Likewise, using the first order conditions, it is easy to show that the MRTS is equal to the ratio of input prices in the optimum:

$$\frac{\partial f / \partial x_i}{\partial f / \partial x_j} = \frac{w_i}{w_j},\ i, j \in \{1, ..., N\}$$

The profit maximization problem gives rise to the profit and input demand functions.

Definition 4.4.2 (profit function)**.** *The firm's* ***profit function*** *$\pi(p, w)$ is defined as the difference between total revenue and total cost:*

$$\pi(p, w) = \max_{x \in X} pf(x) - w \cdot x$$

Definition 4.4.3 (input demand function). *The firm's **input demand function** $x(p, w)$ is the vector of inputs that maximizes firm's profits:*

$$x(p, w) = \arg\max_{x \in X} pf(x) - w \cdot x$$

Note that this is the **unconditional** input demand – it doesn't depend on the production level y, instead assuming implicitly its optimal value.

Definition 4.4.4 (supply function). *The firm's **supply function** $y(p, w)$ is the level of production maximizing firm's profits at (p, w):*

$$y(p, w) = f\big(x(p, w)\big)$$

It turns out you can frame the same problem in two ways. First, you can choose the input levels to maximize profits. Second, you can reformulate the same problem in a two-step way in which in the first step you find the least costly set of inputs to produce at least y units of output, and in the second step use it to maximize firm's profits just in terms of y:

$$\max_{y \geq 0} py - c(w, y)$$

We omit the proof of this statement.

The profit, input demand, and supply functions have a range of important properties that can be used to test the producer theory.

Proposition 4.4.1 (properties of profit function). *Suppose the production function f satisfies assumptions of this chapter, $p \geq 0$, and $w \geq 0$. Then the profit function $\pi(p, w)$ (if well-defined) is continuous and:*

1. *increasing in p*
2. *decreasing in w*
3. *homogeneous of degree one in (p, w)*
4. *convex in (p, w)*
5. *differentiable in $(p, w) \gg 0$*
6. *if f is strictly concave, then Hotelling's lemma holds:*

$$\frac{\partial \pi(p, w)}{\partial p} = y(p, w) \text{ and } \frac{\partial \pi(p, w)}{\partial w_i} = x_i(p, w)$$

Proposition 4.4.2 (properties of input demand and supply functions). *Suppose that the production function f satisfies assumptions of this chapter and is strictly concave. Then for all $p > 0$, and $w \gg 0$ (if well-defined) :*

1. *homogeneity of degree zero. For every $t > 0$:*

$$y(tp, tw) = y(p, w) \text{ and } x_i(tp, tw) = x_i(p, w), \ in \in \{1, ..., N\}$$

2. *own price effects:*

$$\frac{\partial y(p, w)}{\partial p} \geq 0 \text{ and } \frac{\partial x_i(p, w)}{\partial w_i} \leq 0, \ in \in \{1, ..., N\}$$

Exercise 4.4.1. *prove this proposition.*

4.5 Monotone Comparative Statics*

This section introduces a very useful theoretical result allowing us to learn how the optimal solution of a problem depends on parametrization of the objective. We start with the necessary math concepts and notation. They may look difficult. Set your mindset algorithmically in a way to rather learn how to use the result than to think of why it works.

Definition 4.5.1 (lattice operations). *For two vectors $x, y \in \mathbb{R}^N$, define:*

1. *the **infimum** as $x \wedge y = \big\{ \min\{x_1, y_1\}, \min\{x_2, y_2\}, ..., \min\{x_N, y_N\}, \big\}$*

2. *the* ***supremum*** *as* $x \vee y = \Big\{ \max\{x_1, y_1\}, \max\{x_2, y_2\}, ..., \max\{x_N, y_N\}, \Big\}$

Definition 4.5.2 (lattice)**.** *A set* $X \subset \mathbb{R}^N$ *is a lattice if for every* $x, y \in X$ *both* $x \wedge y \in X$ *and* $x \vee y \in X$*.*

Proposition 4.5.1. *The following sets are lattices:*

1. *a closed interval* $[a, b] \subset \mathbb{R}^N$
2. $\mathbb{R}^N_+$
3. $\mathbb{R}^N$

Definition 4.5.3 (supermodular function)**.** *A multi-dimensional function* $f : X \to \mathbb{R}$*, where* X *is a lattice, is called* ***supermodular*** *if for every* $x, y \in X$*:*

$$f(x \vee y) - f(x) \geq f(y) - f(x \wedge y)$$

Definition 4.5.4 (non-decreasing (non-increasing) differences of a function)**.** *We say that a multi-dimensional function* $f : X \times T \to \mathbb{R}$ *has* ***non-decreasing (non-increasing) differences*** *in* (x, t) *if for every* $x, x' \in X$ *such that* $x \geq x$ *and every* $t, t' \in T$ *such that* $t' \geq t$*:*

$$f(x', t') - f(x, t') \geq (\leq) f(x', t) - f(x, t)$$

Proposition 4.5.2. *Let* $f : X \times T \to \mathbb{R}$ *be twice continuously differentiable. Then* f *has non-decreasing differences in* (x, t) *if and only if:*

$$\frac{\partial^2 f}{\partial x_i \partial t_k}(x, t) \geq 0$$

for every i, k *and* (x, t)*.*

With all the definitions above, we are ready to deliver the celebrated Topkis theorem. Consider a function $f : X \times T \to \mathbb{R}$, where $X = \mathbb{R}^N$ or $X = \mathbb{R}^N_+$ and $T \subset \mathbb{R}^M$ (for some $M \geq 1$). Let $S \subset X$. Consider the following maximization problem:

$$\max_x f(x, t)$$
$$\text{such that } x \in S$$

Let $\varphi^*(t) = \arg\max_{x \in S} f(x, t)$ be the set of solutions to the problem above.

Theorem 4.1 (Topkis)**.** *If:*

1. S *is a lattice,*
2. $f(x, t)$ *is supermodular in* x *for every* t*,*
3. $f(x, t)$ *has non-decreasing (non-increasing) differences in* (x, t)

Then $\varphi^*(t)$ *is non-decreasing (non-increasing) in* t*.*

Example 4.5.1. *We will use Topkis theorem to show that the input demand* $x(p, w)$ *in a non-decreasing function of output price* p *in the competitive firm's profit maximization problem. First, formulate the maximization problem:*

$$x(p, w) = \arg\max_{x \in \mathbb{R}^N_+} pf(x) - w \cdot x$$

From now on, we focus on p*, holding fixed* w*, so for notational simplicity write* $x(p) \equiv x(p, w)$*.*

Suppose that the production function is supermodular.

1. *we need to show that the set of feasible solutions is a lattice. Here it holds because* $x \in \mathbb{R}^N_+$ *and* $\mathbb{R}^N_+$ *is a lattice.*
2. *the objective function is* $F(x, p) = pf(x) - w \cdot x$*. We need to show it is supermodular in* x *for every* p*. Take* $x, x' \geq 0$*. Since the production function is supermodular:*

   $$f(x \vee x') - f(x) \geq f(x') - f(x \wedge x') \Leftrightarrow pf(x \vee x') - pf(x) \geq pf(x') - pf(x \wedge x')$$

   Now, note that

   $$x \vee x' - x = \Big\{ \max\{x_i, x'_i\} \Big\} - x \geq x' - x \geq x' - \Big\{ \min\{x_i, x'_i\} \Big\} = x' - x \wedge x'$$

so it follows that $w \cdot (x \vee x' - x) \geq w \cdot (x' - x \wedge x')$. *Plugging back to the previous equation, we get that:*

$$f(x \vee x') - w \cdot (x \vee x') - \big(f(x) - w\big) \cdot x \geq f(x') - w \cdot x' - \big(f(x \wedge x') - w \cdot (x \wedge x')\big) \Leftrightarrow F(x \vee x', p) - F(x,p) \geq F(x',p) - F(x \wedge x', p)$$

so F *is supermodular in* x *for every* p.

3. *Now we need to show that* $F(x,t)$ *has non-decreasing differences in* (x,t). *Take* $x' \geq x$ *and* $p' \geq p$. *In fact:*

$$\begin{aligned} F(x',p') - F(x,p') &= p' f(x') - w \cdot x' - \big(p' f(x) - w \cdot x\big) \\ &= p'\big(f(x') - f(x)\big) - w \cdot (x' - x) \underbrace{\geq}_{f \text{ increasing and } x' \geq x} p\big(f(x') - f(x)\big) - w \cdot (x' - x) \\ &= F(x',p) - F(x,p) \end{aligned}$$

All three conditions of Topkis theorem hold. Therefore, by Topkis theorem, the input demand function is a non-decreasing function of output price p.

Topkis theorem is extremely powerful result allowing us to study properties of solutions in potentially very complicated theoretical models.

5 Productivity and Production

Today, we link the producer theory to data. We are particularly interested in the so called productivity term, that permits firms with the same level of input produce different level of outputs − a phenomenon widely observed in the data. We analyze how the existence of such a productivity term affects firm's problem and what are its implications for the economists willing to identify the parameters of production function. We spend some time discussing different measurement of output that leads to potentially different properties of the retrieved productivity term. We finish with discussing possible applications, including studying reallocation or markups.

List of contents. Production function. Inputs into production technology. Capital inputs: physical capital, human capital, financial capital, natural capital, social capital. Variable inputs: labor, materials, energy. Intermediate inputs. Hick-neutral productivity. Total factor productivity. Identification challenges related to production functions. Olley-Pakes two-stage identification argument: investment as a proxy for productivity. Control function. Lumpy investment. Levinsohn-Petrin: materials as a proxy for productivity. Measurement approaches: value-added, gross output, physical output. Physical productivity (TFPQ), revenue productivity (TFPR). Applications: reallocation, markups.

Literature References.

[1] Richard Ericson and Ariel Pakes (1995). Markov-perfect industry dynamics: A framework for empirical work. The Review of Economic Studies, 62(1), 53-82.

[2] G. Steven Olley and Ariel Pakes (1996). The Dynamics of Productivity in the Telecommunications Equipment Industry. Econometrica, 1263-1297.

[3] James Levinsohn and Amil Petrin (2003). Estimating production functions using inputs to control for unobservables. The Review of Economic Studies, 70(2), 317-341.

[4] Lucia Foster, John Haltiwanger, and Chad Syverson (2008). Reallocation, firm turnover, and efficiency: Selection on productivity or profitability?. American Economic Review, 98(1), 394-425.

[5] Daniel A. Ackerberg, Kevin Caves, and Garth Frazer (2015). Identification properties of recent production function estimators. Econometrica, 83(6), 2411-2451.

[6] Jan De Loecker and Frederic Warzynski (2012). Markups and firm-level export status. American Economic Review, 102(6), 2437-2471.

5.1 Production Inputs in Economics Framework

In the previous class, we studied firm's technology of converting a bundle of goods (called inputs) X into another good Y (called output) and explained through a production function F.

Economists typically focus on two types of inputs within the production framework. The first type is capital K, which represents durable resources in the economy. We usually assume a firm possesses a certain stock of capital that is crucial for production. Capital is inherently persistent; however, it depreciates over time and requires ongoing investments to maintain its stock. In particular, capital stock cannot be adjusted fully in the short term. Economists recognize various kinds of capital:

- ↪ **physical capital** including tangible assets like machinery, buildings, and infrastructure needed for production of goods and services,
- ↪ **human capital** referring to the skills, knowledge, and experience possessed by individuals, which can be enhanced through education and training,
- ↪ **financial capital** encompassing funds available for investment, including money, stocks, bonds, and other financial instruments,
- ↪ **natural capital** containing natural resources such as land, minerals, and ecosystems that contribute to economic production,
- ↪ **social capital** involving networks, relationships, and social norms that facilitate cooperation and economic transactions.

The second type of inputs are so called freely adjustable inputs. The most important representative of this category would be labor L, referring to employee count of the firm. In fact, on different markets labor may be more or less freely adjustable, nevertheless it is typically assumed that within a period (usually a (financial) year) the firms is free to adjust their labor input. Another free inputs is materials M which describes the raw materials and intermediate goods used in the production process.

The second type of inputs is referred to as freely adjustable inputs. The most significant representative of this category is labor L, which usually refers to the number of employees in a firm. While the factual degree of labor adjustability can vary across different markets, it is generally assumed that firms can adjust their labor input freely within a specified period (usually a financial year). Another example of freely adjustable input are materials M, which encompass the raw materials and other intermediate goods used in the production process.

Tody, we center our analysis around a simple two-input specification of a production function:

$$Y = F(K, L) \tag{1}$$

5.2 The Missing Piece – Productivity

The formulation of the production function given by equation 2 implies that every firm holding the same input levels K and L will necessarily produce Y units of output. However, this is not something we typically observe in the data. Quite contrary, there is substantial variation in production outputs generated by seemingly similar firms (in terms of K and L). This dispersion is systematic, meaning that some firms consistently generate more output from the same inputs than others.

A fundamental question arises, what is missing in out framework? Economists address the above data patterns by introducing **productivity** term, Ω:

$$Y = \Omega F(K, L) \tag{2}$$

We assume that each firm has its own Ω that scales the default production levels determined by the input levels K and L and the functional form of production F. The Ω term is referred to as **Hicks-neutral** productivity (neutral because it scales the production level independently of the input levels) or **total factor productivity** (TFP, as it measures the overall efficiency of inputs), or shortly – productivity.

The TFP term is central to the interest of economists across many fields, including macroeconomics, growth, trade, industrial organization, policy evaluation, innovation, technological change, and many others. Today, we will study the firm's productivity term in detail, departing from how to identify it from data and finishing at presenting a plethora of applications utilizing the concept of TFP.

To set the stage for our analysis, we will assume that the production function follows a Cobb-Douglas functional form. This assumption is not strictly necessary, but it simplifies our notation and allows us to communicate the main points more clearly. The output level Y_{it} of firm i at period t can be expressed as:

$$Y_{it} = \Omega_{it} K_{it}^{\beta_k} L_{it}^{\beta_\ell} \exp\{\varepsilon_{it}\}$$

where the additional term ε_{it} measures the impact of completely random events – in particular, independent from anything else in the model.

It is convenient to operate in logarithmic terms rather than levels of variables. Following the practice in the literature, define $a = \log A$ for $A \in \{Y, K, L, \Omega\}$. Our production function can then be rewritten as:

$$y_{it} = \beta_k k_{it} + \beta_\ell \ell_{it} + \omega_{it} + \varepsilon_{it} \tag{3}$$

Despite being linear, equation 3 is notoriously difficult to identify. This is because we, as economists and econometricians, cannot observe the productivity factor ω_{it}. However, there is no reason to assume the firm does not observe it. If the firm knows its ω_{it} at the time inputs are decided, the optimal and chosen inputs are functions of ω_{it}.

To see more clearly why this is a problem, consider a firm experiencing a high realization of ω_{it}. High productivity firms are likely to produce more simply due to the scaling effect of TFP. This, in turn, increases the marginal gain from production, creating incentives to hire additional inputs. With more inputs, there will be even more output. Consequently, the perceived productivity of these inputs may be overestimated if we fail to account for ω_{it} while studying the production function, leading to an **overestimation** of the parameters β_k and β_ℓ.

In the following sections, we will establish a framework to identify the parameters of the production function, accounting for the unobserved productivity term. Given the nature of the problem—two additively separable terms in the same equation—identification can only be achieved by relying on the theoretical implications of a model that describes firm behavior.

Before we proceed, two remarks are to be made. First, TFP is defined in relation to the specific functional form of the production function and thus does not have its own units. Generally, we study relative values of TFP across firms, industries, or over time. Second, despite its complex interpretation, ω_{it} is essentially a residual in the production function. We can never lose this aspect out of the picture, as it has significant consequences for our understanding of the production process – more so than many economists would wish to acknowledge.

5.3 Theoretical Background*

The modern analysis of productivity is based on the industry dynamics model developed by Ericson and Pakes (1995). Here, we draft a simplified model just to present the main idea. The full specification as well as solving the model is beyond the scope of our course.

Assume that firms operate in a perfectly competitive market. Within each period, labor is freely adjustable, while capital is fixed being a result of past investments. Therefore, the within-period profit function can be expressed as:

$$\pi(K_{it}, \Omega_{it}) = \max_{L} \left\{ \Omega_{it} K_{it}^{\beta_k} L^{\beta_\ell} - r_{it} K_{it} - w_{it} L_{it} \right\}$$

where r_{it} and w_{it} are rental rates of capital and labor respectively. Assuming competitive input markets, these are taken as given.

Every period, a fraction $\gamma \in (0, 1)$ of the capital depreciates. Capital can be replenished through investment I_{it}:

$$K_{it+1} = (1 - \gamma) K_{it} + I_{it}$$

Furthermore, we assume that the productivity for the next period depends only on the productivity of the current period, along with a random shock η_{it} that is unrelated to anything else in the model:

$$\Omega_{it+1} = \Psi(\Omega_{it}) \exp\{\eta_{it}\}$$

This is the celebrated **Markov process** formulation, in which the value of a variable today depends only on its value yesterday adjusted possibly by an unpredictable shock.

The long-term profit optimization problem of the firm is quite complicated. One way to think about it is that at any given point in time $t = \tau$, the firm selects the optimal investment path contingent upon all possible future productivity realizations:

$$\max_{\{I_{it}\}_{t=\tau}^{\infty}} \sum_{t=\tau}^{\infty} \delta^{t-\tau} \Big(\pi(K_{it}, \Omega_{it}) - C(I_{it}) \Big)$$

where $\delta \in (0, 1)$ is a discount factor that indicates how much the firm values future profits today and $C(I_{it})$ is the cost of investment.

A key insight from this framework is that the optimal solution for investment, denoted as $I_{it} \equiv I(K_{it}, \Omega_{it})$, exists and is strictly increasing with respect to the unobserved productivity Ω_{it} (as long as $I_{it} > 0$).

5.4 Identification of Production Function Parameters – The Control Function Argument

In this section, we examine all variables expressed in natural logarithms.

As shown in the previous section, the theory implies that the optimal investment is a strictly increasing function of the current productivity:

$$i_{it} = i(k_{it}, \omega_{it})$$

This means that we can invert this function to express productivity as a function of capital and investment – two quantities observed in the data!

$$\omega_{it} = h(k_{it}, i_{it})$$

where function $h(\cdot) \equiv i^{-1}(\cdot)$ is the inverse investment with respect to the productivity term. Although we don't know the functional form of $h(k_{it}, i_{it})$, we can approximate it with great accuracy, typically through a polynomial expansion (which you don't need to know for the purpose of succeeding in our class). The function h is sometimes referred to as a **control function**, as it effectively controls for the unobserved component ω_{it}.

The inversion of the optimal investment is key for the identification argument. Note that now we can substitute it back to the production function (equation 3):

$$y_{it} = \beta_k k_{it} + \beta_\ell \ell_{it} + \underbrace{h(k_{it}, i_{it})}_{\omega_{it}} + \varepsilon_{it} \tag{4}$$

Take a closer look at this equation. On the right-hand side, every variable is observed, except for a purely random shock (which is not really a concern). The only potential challenge is that we do not know what the function h is. This is why, at this stage, we cannot identify the parameter β_k separately from $h(k_{it}, i_{it})$. Intuitively, how can we differentiate between $\beta_k k_{it}$ and $h(k_{it}, i_{it})$ without knowing what h is?

A different argument can be made regarding β_ℓ. Intuitively, $\beta_\ell \ell_{it}$ is the only place where ℓ_{it} appears in equation 4. Even though it may be correlated with the unobserved ω_{it}, this is no longer an issue for us, as we *control* for it by incorporating function $h(\cdot)$ into the equation. Thus, by regressing output levels on labor input along with a chosen polynomial in capital and investment, we can recover β_ℓ:

$$y_{it} = \beta_\ell \ell_{it} + \underbrace{\phi(k_{it}, i_{it})}_{\beta_k k_{it} + h(k_{it}, i_{it})} + \varepsilon_{it} \tag{5}$$

where $\omega_{it} = h(k_{it}, i_{it}) = \phi(k_{it}, i_{it}) - \beta_k k_{it}$. As a result of this **first stage**, we identify β_ℓ and $\phi(k_{it}, i_{it})$ (which we denote simply by ϕ_{it} for notational convenience).

The final step is to identify β_k. To achieve this, we utilize the law of motion of the unobserved productivity. In log terms, we express it as:

$$\omega_{it} = \psi(\omega_{it-1}) + \eta_{it}$$

Since $\omega_{it-1} = h(k_{it-1}, i_{it-1}) = \phi_{it-1} - \beta_k k_{it-1}$, we can rewrite it as follows:

$$\begin{aligned}
\underbrace{y_{it} - \beta_\ell \ell_{it}}_{\text{identified at the } 1^{st} \text{ stage}} &= \beta_k k_{it} + \omega_{it} + \varepsilon_{it} \\
&= \beta_k k_{it} + \psi(\omega_{it-1}) + \underbrace{\eta_{it} + \varepsilon_{it}}_{\equiv \tilde{\varepsilon}_{it}} \\
&= \beta_k k_{it} + \psi(\phi_{it-1} - \beta_k k_{it-1}) + \tilde{\varepsilon}_{it}
\end{aligned}$$

Now, let's take a closer look at the right-hand side. The capital variable k_{it} and its lag k_{it-1} are observed. The term ϕ_{it-1} is identified in the first stage. The term $\tilde{\varepsilon}_{it}$ is a sum of two independent random disturbances, so it remains a random disturbance (well-behaved—unrelated to anything else in the model). The only unknown is the function ψ, which describes the law of motion of the productivity term. However, this can be approximated, again, by a polynomial expansion in ϕ_{it-1} and $\beta_k k_{it-1}$. Intuitively, this **second stage** boils down to a sort of regression in which k_{it} shows up only as a linear term (albeit other terms enter possibly in a non-linear fashion), completing our identification argument.

Before we proceed, note that we need both stages to build our identification argument. This is because ℓ_{it} depends on ω_{it} (as a solution to the within-period profit maximization), and thus on η_{it}, which in turn enters $\tilde{\varepsilon}_{it}$—the random disturbance in the second stage equation. Consequently, we could not identify β_ℓ in the second stage.

This section demonstrates that the theoretical model produces restrictions that allow us retrieve its unobserved primitives—here, the production function parameters and implied productivity realizations—from data. This is because we actually showed that is a one-to-one mapping between the realizations of exogenous shocks in the model—hence the data—and model parameters. This is what the economists understand under the term model **identification**. Take a moment and think whether you understand this link.

Even though showing how to actually estimate the parameters and which statistical theories allow us to proceed with inference is beyond the scope of our course, the identification arguments likely provides a set of excellent hints on how to do so.

Having identified β_k and β_ℓ enables us to comfortably assume that we already know their values. This also means that we can calculate ω_{it} (up to a random disturbance ε that by construction averages to zero).

5.5 Extensions*

Before discussing how to use the concept of productivity to study various economic phenomena, it's important to recognize that an entire body of literature has emerged based on the insights of Olley and Pakes (1996).

The first significant refinement to the work of Olley and Pakes comes from Levinsohn and Petrin (2002). They observe that many firms exhibit zero investment in practice, which contradicts assumptions necessary for the identification argument introduced above. Additionally, investment is often **lumpy**, meaning that firms tend to make large and infrequent investments rather than smaller, regular ones. This indicates that investment does not really respond to productivity shocks occurring in each period, complicating its use as a **proxy** for unobserved productivity shocks. To address this issue, Levinsohn and Petrin propose a similar approach in which investment variable is replaced with materials—typically admitting non-zero values in empirical applications.

More modern refinements raise potential problems with the original identification argument. Ackerberg, Caves, and Frazer (ACF, 2015) note that if the variable input—here: labor—is chosen in the per-period profit maximization function as a function only of k_{it} and ω_{it}: $\ell_{it} \equiv \ell(k_{it}, \omega_{it})$ then its coefficient in the production function β_ℓ is not identified in the first stage. To see this, plug $\ell(k_{it}, \omega_{it})$ back to the equation 4 and note that every element on the right hand side is now a function of k_{it} and ω_{it}. There is no variation left in ℓ_{it} to identify β_ℓ. So either we are comfortable in assuming some additional variation in a specific application (input-specific prices, measurement error in variable input, optimization error in variable input, variation in firm's decision timing and availability of various pieces of information there), or the first stage only identifies function

$$\phi(k_{it}, i_{it}) = \beta_\ell \ell(k_{it}, i_{it}) + \beta_k k_{it} + h(k_{it}, i_{it})$$

in equation 5. If the latter is the case, the second stage needs to be updated to break the link of ℓ_{it} and $\tilde{\varepsilon}_{it}$. This can be done by instrumenting ℓ_{it} with its lag ℓ_{it-1}. See ACF for more details if you are interested.

5.6 Output Measurement and Productivity Measures

Taking the production functions to data is a tricky thing. A lot of how to do it depends on what data is available, requiring the pursuit of different approaches.

Typically, empirical analyses of production functions rely on accounting data derived from firms' mandatory financial reporting. This usually means that output is expressed in monetary terms as the total value of sales of goods and services produced by the company. Researchers often pursue two main approaches:

↪ **value-added** specification which includes only capital and labor as inputs, excluding in particular materials, energy inputs, and other analogous intermediate inputs. The value-added method was adopted by Olley and Pakes, who defined output as total sales minus materials expenditure (i.e., the accounting value added). This method, which can seem somewhat a-theoretical, is now referred to as the **restricted profit** approach. While this may lack a strong theoretical justification, a relatively simple fix is available. If we are willing to assume that the production function is Leontief of the following form:

$$Y = \min\{\Omega F(K, L), \gamma M\}$$

for some parameter γ, the resulting framework would allow us to abstract from materials in the functional form of production technology. This approach is sometimes called the **structural value added** approach.

↪ **gross-output** specification which includes intermediate inputs such as materials, energy, or fuel in the equation of production function.

There are subtle differences between these approaches, which we do not detail here. The choice of a specific specification largely depends on data availability and the assumptions made about the production process.

Nonetheless, it is important to note that the origins of production theory pertain to quantities, or **physical output**. Unfortunately, physical output data are most often unavailable to researchers, which is why this specification is rarely observed in empirical studies.

It turns out that choosing a specifications from among the above is not without loss of generality. Foster, Haltiwanger, and Syverson (FHS, 2008) used unique data where both physical output and (deflated) sales were observed to explore differences in productivity measures induced by different output specifications. They defined two key objects:

↪ **physical productivity** (TFPQ) based on physical quantities of the output,
↪ **revenue productivity** (TFPR) based on producer revenues.

TFPQ represents the producer's *true* technical efficiency level. In turn, FHS show theoretically that producers can have high TFPR levels because they are efficient, but this can also be driven by high producer-specific demand.

Empirically, TFPQ and TFPR showed correlation in the FHS study, though imperfect. Moreover, physical TFPQ turned out to be inversely correlated with prices, while TFPR was positively correlated with prices. This indicates that TFPR might be higher (indicating greater revenue productivity) either due to increased output (holding inputs constant) or increased prices. Distinguishing between these two factors is vital from the perspective of welfare analysis. Unfortunately, in 99% of cases, we can only rely on TFPR due to data limitations. This reliance may obscure the effects of consumer demand. In other words, if we do not assume perfectly competitive markets, we may run into issues. To accurately separate the effects of demand from productivity in the TFPR measure, it may be essential to incorporate a fully specified demand model into the analysis.

5.7 Application (I)—Reallocation

Productivity measures can be used to analyze the allocation of resources across firms.

Olley-Pakes study examines how resources are reallocated among firms as a result of market deregulation. To see how this is done, let ω_{it} represent plant-level productivity and s_{it} denote the share of firm i in total output on the industry, national, or other level of aggregation. Define the aggregate productivity in period t as:

$$\omega_t = \sum_{i=1}^{N_t} s_{it}\omega_{it}$$

Next, we decompose the productivity and shares into the (unweighted) averages and individual deviations from these averages:

$$\omega_{it} = \overline{\omega}_t + \Delta\omega_{it} \quad \text{and} \quad s_{it} = \overline{s}_t + \Delta s_{it}$$

As an exercise, you can prove that the deviations must sum to zero across all firms, i.e., $\sum_{i=1}^{N_t} \Delta\omega_{it} = 0$ and $\sum_{i=1}^{N_t} \Delta s_{it} = 0$.

Using this observation, it follows that:

$$
\begin{aligned}
\omega_t =& \sum_{i=1}^{N_t} (\bar{s}_t + \Delta s_{it})(\bar{\omega}_t + \Delta\omega_{it}) \\
=& \sum_{i=1}^{N_t} \bar{s}_t \bar{\omega}_t + \sum_{i=1}^{N_t} \bar{s}_t \Delta\omega_{it} + \sum_{i=1}^{N_t} \Delta s_{it} \bar{\omega}_t + \sum_{i=1}^{N_t} \Delta s_{it} \Delta\omega_{it} \\
=& N_t \bar{s}_t \bar{\omega}_t + \bar{s}_t \underbrace{\sum_{i=1}^{N_t} \Delta\omega_{it}}_{=0} + \bar{\omega}_t \underbrace{\sum_{i=1}^{N_t} \Delta s_{it}}_{=0} + \sum_{i=1}^{N_t} \Delta s_{it} \Delta\omega_{it} \\
=& N_t \bar{s}_t \bar{\omega}_t + \sum_{i=1}^{N_t} \Delta s_{it} \Delta\omega_{it} \\
=& \bar{\omega}_t + \sum_{i=1}^{N_t} \Delta s_{it} \Delta\omega_{it}
\end{aligned}
$$

because $\frac{\bar{s}_t}{N_t} = 1$.

This means we can decompose the weighted average productivity into the unweighted average and the covariance between productivity and output. A larger covariance indicates that a greater share of output is (re-)allocated to more productive firms, which in turn raises overall industry productivity.

The table below illustrates what Olley and Pakes found in their study of deregulation in telecommunication equipment industry.

TABLE XI

DECOMPOSITION OF PRODUCTIVITY[a]
(EQUATION (16))

Year	p_t	$\bar{p}_t$	$\Sigma_i \Delta s_{it} \Delta p_{it}$	$\rho(p_t, k_t)$
1974	1.00	0.90	0.01	−0.07
1975	0.72	0.66	0.06	−0.11
1976	0.77	0.69	0.07	−0.12
1977	0.75	0.72	0.03	−0.09
1978	0.92	0.80	0.12	−0.05
1979	0.95	0.84	0.12	−0.05
1980	1.12	0.84	0.28	−0.02
1981	1.11	0.76	0.35	0.02
1982	1.08	0.77	0.31	−0.01
1983	0.84	0.76	0.08	−0.07
1984	0.90	0.83	0.07	−0.09
1985	0.99	0.72	0.26	0.02
1986	0.92	0.72	0.20	0.03
1987	0.97	0.66	0.32	0.10

They show that (cite) *unweighted average productivity has not changed much since 1975* (deregulation allowing more entry), *but there has been a reallocation of output from less productive to more productive plants. This reallocation of output, and not an increase in average productivity, is behind the increase in productivity at the industry level. Moreover the allocation of output seems to have improved dramatically following the certification and registration program, and then again following divestiture* (break-up of the downstream monopolist 1984).

5.8 Application (II)—Mark-ups

Another fundamental application of production function framework is to study mark-ups. It turns out that if we believe firms are cost-minimizers, we can identify the firm-level markups. We show it by following the derivations by De Loecker

& Warzynski (2012), acknowledging that the original idea dates back to Hall (1988).

Suppose the production function is given by $F(X_{it}^1, X_{it}^2, .., X_{it}^V, K_{it}, \Omega_{it})$, where X_{it}^v are freely variable inputs—inputs that can be fully adjusted in the short term. K_{it} represents capital stock (a dynamic input), and Ω_{it} denotes productivity. We assume that $F(\cdot)$ is twice continuously differentiable.

Consider firm's cost minimization problem in this setting:

$$\min_{\{X_{it}^v\}_{v=1}^V, K_{it}} \sum_{v=1}^N P_{it}^v X_{it}^v + r_{it} K_{it}$$
$$\text{such that } F(X_{it}^1, X_{it}^2, .., X_{it}^V, K_{it}, \Omega_{it}) = \overline{Y}$$

for some fixed production level $\overline{Y}$.

We proceed with standard tools to solve this problem. The Lagrangian is given by:

$$\mathcal{L}(X_{it}^1, X_{it}^2, .., X_{it}^V, K_{it}) = \sum_{v=1}^N P_{it}^v X_{it}^v + r_{it} K_{it} + \lambda_{it}\big(F(X_{it}^1, X_{it}^2, .., X_{it}^V, K_{it}, \Omega_{it}) - \overline{Y}\big)$$

A first order condition with respect to a variable input implies:

$$\frac{\partial \mathcal{L}_{it}}{\partial X_{it}^v} = P_{it}^v - \lambda_{it} \frac{\partial F(\cdot)}{\partial X_{it}^v} = 0$$

A few algebraic steps derives the final expression of interest:

$$\begin{aligned} P_{it}^v = \lambda_{it} \frac{\partial F(\cdot)}{\partial X_{it}^v} &\Leftrightarrow \frac{1}{\lambda_{it}} P_{it}^v = \frac{\partial F(\cdot)}{\partial X_{it}^v} \\ &\Leftrightarrow \frac{1}{\lambda_{it}} P_{it}^v \frac{X_{it}^v}{F_{it}} = \frac{\partial F(\cdot)}{\partial X_{it}^v} \frac{X_{it}^v}{F_{it}} \\ &\Leftrightarrow \frac{\partial F(\cdot)}{\partial X_{it}^v} \frac{X_{it}^v}{F_{it}} = \frac{P_{it}}{\lambda_{it}} \frac{P_{it}^v X_{it}^v}{P_{it} F_{it}} \end{aligned}$$

First, note that that λ_{it} describes the marginal cost of production, as $\lambda_{it} = \frac{\partial \mathcal{L}_{it}}{\partial F_{it}}$. Second, economics we usually define markups in economics to describe the margin between the price of final good and the marginal costs. Hence, it is natural to define the markup $\mu_{it} = \frac{P_{it}}{\lambda_{it}}$ as simply the ratio of output price and marginal cost.

Now, let's stare more closely at the last equation in the light of these considerations:

$$\underbrace{\frac{\partial F(\cdot)}{\partial X_{it}^v} \frac{X_{it}^v}{F_{it}}}_{\text{output elasticity of variable input}} = \underbrace{\frac{P_{it}}{\lambda_{it}}}_{\mu_{it} \equiv \text{markup}} \overbrace{\frac{P_{it}^v X_{it}^v}{P_{it} F_{it}}}^{\text{input share in total sales}}$$

This equation binds three objects: the markup, variable input's share in total sales, and output elasticity of the variable input. The first is what we want to identify. The second is typically observed in the data. What about the third? Well, this object is derived from a production function. Once we have identify the parameters of production function, this object can is also identified.

Note that in every step above, we condition on the choice of dynamic inputs, such as K_{it}, but these do not alter any derivation. Markups are recovered solely from the conditional cost function (i.e. conditional on dynamic input values). This is a great news, because working with dynamic inputs is typically much, much harder.

The last question would be why bother? The analysis of markups can deliver myriads of valuable insights into behavior of firms, and hence consumer welfare, and equilibrium outcomes. In the tutorial next week, we will examine two papers that utilize this method to (i) explore markup creation among exporters and non-exporters, and (ii) investigate the rise of market power in the long term as measured by firms' ability to charge higher markups.

5.9 What Else Do People Do With Productivity?

We finish this note by borrowing a very nice literature review from Kohei Kawaguchi at HKUST Business School to show you the scope of applications of today's topic in various fields of economics.

↪ **Industrial Organization**.

- Olley & Pakes (1996): *How much did the deregulation in the U.S. telecommunication industry, in particular the divestiture of AT&T in 1984, spurred the productivity growth of the incumbent, facilitated entries, and increased the aggregate productivity?*
- Doraszelski & Jaumandreu (2013): *What is the role of R&D in determining the differences in productivity across firms and the evolution of firm-level productivity over time?*

↪ **Development**.

- Hsieh & Klenow (2009): *How large is the misallocation of inputs across manufacturing firms in China and India compared to the U.S? How will the aggregate productivity of China and India change if the degree of misallocation is reduced to the U.S. level?*
- Gennaioli, La Porta, Lopez-de-Silanes, & Shleifer (2013): *What are the determinants of regional growth? Do geographic, institutional, cultural, and human capital factors explain the difference across regions?*

↪ **Trade**.

- Haskel, Pereira, & Slaughter (2007) *Are there spillovers from FDI to domestic firms?*
- De Loecker (2011): *Does the removal of trade barriers induces efficiency gain for producers?*

↪ **Management**.

- Bloom & Van Reenen (2007): *How do management practices affect the firm productivity?*
- Braguinsky, Ohyama, Okazaki, & Syverson (2015): *How do changes in ownership affect the productivity and profitability of firms?*

↪ **Education**.

- Cunha, Heckman, & Schennach (2010): *How do childhood and schooling interventions "produce" the cognitive and non-cognitive skills of children?*

6 General Equilibrium in Exchange Economies

The theme of today's class is a good news that the competitive markets are efficient. We explore one of the most important result of modern economics—the theory of general equilibrium, for which Kenneth Arrow and Gerard Debreu received Nobel Prizes (and virtually any other of the most prestigious awards in the profession). The basic concern of the theory of general equilibrium is how goods are allocated among the optimizing economic agents. We dive into the mechanics of how the celebrated *invisible market hand* of Adam Smith clears the markets populated by selfish agents focusing solely on their optimal choices in an efficient way. We also discuss normative contents of the general equilibrium theory.

It is my belief that every economist should understand the nature of competitive markets. To make it easier for you, we work with the simplest general equilibrium environment: exchange economies in which consumers can trade their initial endowments. We start with defining what the exchange economies are and considering when trade may occur. Next, we define equilibrium and show its existence. Lastly, we move to the most important results of today's class: the (fundamental) welfare theorems.

This meeting is the most theoretically rigorous among all our classes. This time all the proofs are provided in the notes. Take some time to read through them and make sure you understand the logic driving the results.

List of contents. Exchange economy. Initial endowment. Allocation. Feasible allocation. Gains from trade. Pareto-efficiency. Competitive Equilibrium. Individual demand function. (Aggregate) excess demand function. Walras' law. Free goods. Desirable goods. Price normalizations. Numeraire price system. Simplex price system. Existence of competitive equilibrium. First welfare theorem. Competitive equilibrium with transfers. Second welfare theorem.

Textbook references. Espinola: Ch. 9; Varian: Ch. 17; Jehle & Reny: Ch. 5; MWG: Ch. 15-17

6.1 Mathematical Glossary

Definition 6.1.1 (unit simplex)**.** *A $N-1$ dimensional **unit simplex** is a set of N dimensional vectors whose elements are non-negative and sum to one:*

$$S^{N-1} = \left\{ p \in \mathbb{R}_+^N : \sum_{j=1}^{N} p_j = 1 \right\}$$

Theorem 6.1 (Brouwer's fixed point theorem)**.** *Suppose $f : S^{N-1} \to S^{N-1}$ is a continuous function from a unit simplex into itself. Then, there exists $p \in S^{N-1}$ such that:*

$$p = f(p)$$

Definition 6.1.2 (singleton)**.** *A set containing only one element is called **singleton**.*

Proposition 6.1.1 (properties of singletons)**.** *Singletons are non-empty and convex sets.*

Definition 6.1.3 (strict monotonicity of preferences)**.** *We say that preference relation is **strictly monotone** (or $u(\cdot)$ is strictly increasing) if for every $c, c' \in C$ if $c' \geq c$ and $c' \neq c$ then $u(c') > u(c)$.*

Theorem 6.2 (Separating hyperplane theorem)**.** *Suppose A and B are two non-empty, disjoint, convex sets in $\mathbb{R}^N$. Then there exists vector $0 \neq p \in R^N$ such that for every $x \in A$ and every $y \in B$:*

$$p \cdot x \geq p \cdot y$$

6.2 Exchange Economies

We consider a setting with N goods, and define the consumption set to be $C = \mathbb{R}_+^N$. A typical bundle in the consumption set takes a form $c = (c_1, c_2, ..., c_N)$. There are I consumers indexed by $i = \{1, ..., I\}$. We index consumers by upper script and goods by lower script. Every consumer has a preference relation $\succsim^i$ defined on the consumption set. We assume that these preference relations are *well-behaved*, i.e. satisfy the assumptions we made in the consumer theory class (i.e. utility is continuous, strictly increasing, strictly quasi-concave).

Every consumer is **endowed** with an initial vector of good quantities $e^i = (e_1^i, e_2^i, \ldots, e_N^i) \in C$. We assume that the **aggregate endowment** $e = \sum_{i=1}^{I} e^i \gg 0$. In today's class, we are studying how consumers can trade goods from their initial endowment to achieve a more preferred bundle – an **exchange economy**.

Definition 6.2.1 (exchange economy)**.** *A collection $\mathcal{E} = (\succsim^i, e^i)_{i=1}^I$ defines a **exchange economy**.*

Definition 6.2.2 (allocation)**.** *Vector (of vectors) $c = (c^1, c^2, ..., c^I)$, where $c^i = (c_1^i, c_2^i, \ldots, c_N^i) \in C$, including a consumption bundle for every consumer in the economy is called **allocation**.*

Example 6.2.1 (allocation)**.** *The initial endowment in the economy:*

$$e = (e^1, ..., e^N)$$

where $e^i = (e_1^i, e_2^i, \ldots, e_N^i) \in C$ for every $i \in \{1, .., I\}$, in an allocation.

Definition 6.2.3 (feasible allocation)**.** *An allocation c is called **feasible** if the total (across consumers) quantity of every good does not exceed the total (across consumers) quantity of initial endowment for this good:*

$$\sum_{i=1}^{I} c_j^i \leq \sum_{i=1}^{I} e_j^i, \ j \in \{1, ..., N\}$$

Definition 6.2.4 (set of all feasible allocations)**.** *Collect all feasible allocations in the economy in the set $FA(e)$:*

$$FA(e) = \left\{ c : \sum_{i=1}^{I} c_j^i \leq \sum_{i=1}^{I} e_j^i, \ j \in \{1, ..., N\} \right\}$$

6.3 Trade

In exchange economies, consumers endowed with some initial values can trade among themselves to achieve a more preferred bundle. The following example summarizes **gains from trade** available to consumers.

Example 6.3.1 (exchange with two goods – gains from trade). *Consider an economy with two goods $j \in \{1, 2\}$ and two agents $i \in \{A, B\}$. Agents' preferences are represented by utility functions $u(c^i) = \sqrt{c_1^i c_2^i}$ and their initial endowments are given by $e^A = (1, 0)$ and $e^B = (0, 1)$.*

In this economy, gains from trade are particularly visible. The preferences favor variety, whereas endowments contain positive quantities of only one good. If agents were just to consume their endowments, they would receive the lowest possible stream of utility normalized to zero. Each one of them would be willing to give up some amount of good in their possession to receive some amount of the other good.

When the trade will and will not occur? A useful concept to study gains from trade is **Pareto-efficiency**.

Definition 6.3.1 (Pareto-efficient allocation). *A feasible allocation $c \in FA(e)$ is called **Pareto-efficient** if there is no other feasible allocation $\tilde{c} \in FA(e)$ such that $\tilde{c}^i \succsim^i c^i$ for every consumer i, and $\tilde{c}^j \succ^j c^j$ for at least one consumer j.*

In particular, if the initial endowment is a Pareto-efficient allocation, there is no gain from trade is possible. No one can get better off without some other agent being worse off. Therefore, no voluntary transaction can occur.

This property of Pareto-efficient allocations place them as a good candidate for the result of trade. However, not every Pareto-efficient allocation can result from trade. Consider allocation $c = (1, 1), (0, 0)$ in which all goods are consumed by agent A. It is clearly Pareto-efficient—you cannot improve the situation of agent B without deteriorating the situation of agent B. However, no voluntary trade would lead from the initial endowment to this allocation.

In the class, we will learn **Edgeworth box** — a convenient way of studying exchange economies with two agents and two goods. Using Edgeworth box, we will shed more light on which Pareto-efficient allocations can be a result of trade.

At the end of this section, note that the notion of Pareto-efficiency is related to efficiency, but not any sense of (more or less subjective) fairness. We will come back to this thought later on in this class.

Another aspect of trade is the exchange rate between a pair of goods, or **prices**. How many units of good 1 does agent A need to give up to receive a unit of good 2 from agent B? It turns out that price-setting mechanism plays a fundamental role in general equilibrium theory.

6.4 Competitive Equilibrium

In exchange economies, consumers endowed with initial quantities of goods can trade among themselves to achieve a more preferred allocation. We assume that the trade is **decentralized** — every consumer cares only about their utility facing their affordability constraint. We consider **competitive equilibria**, in which consumers take prices as given. The role of prices is to equilibrate demand and supply of goods. This is the mechanism that Adam Smith called the invisible market hand. Every consumer only cares about their utility and their affordability and their joint decisions imply prices that equilibrate the market on the aggregate level.

Definition 6.4.1 (competitive equilibrium). *A **competitive** (Walrasian) **equilibrium** of an exchange economy $\mathcal{E}$ is a set of:*

1. *allocation $c = (c^1, ..., c^I)$*
2. *prices $p \in R^N_{++}$*

such that:

1. *given prices, every consumer makes optimal choices:*

$$c^i = \arg\max_{x \in C} u^i(x) \text{ such that } p \cdot x = p \cdot e^i$$

2. *allocation is feasible:*

$$\sum_{i=1}^{I} c_j^i \leq \sum_{i=1}^{I} e_j^i$$

A few remarks. First, assuming strictly increasing preferences, we can show that the inequality above holds with equality. Second, consumer's income (or wealth) now is defined through initial endowment and prices. Consumers can spend only as much as their initial endowment is worth in market prices. Third, our equilibrium concept brings some notion of stability—in the sense that no one wants to trade anymore given the equilibrium prices (because every consumer maximizes their utility given prices and their initial endowment). You are right by connecting this sense of stability with Pareto-efficiency. In the next section, we will emphasize this connection.

Before we get there, however, we need to ask if a competitive equilibrium exists; and if the answer is yes, under which conditions it does exist. A key object allowing us to make progress in this area are individual demand functions.

Definition 6.4.2 (individual demand)**.** *The **individual demand** of a consumer $i \in \{1, .., I\}$ is a solution to the utility maximization problem:*

$$c^i(p, p \cdot e^i) = \arg\max_{x \in C} u^i(x) \text{ such that } p \cdot x = p \cdot e^i$$

Proposition 6.4.1 (individual demand functions)**.** *Under the assumption of well-behaved preferences, individual demands: exist, are unique, and are continuous in (strictly positive) prices.*

This is nothing new. Note that strictly positive prices imply that the budget set is compact: closed and bounded. Therefore, a demand function exists. It is also unique (single-valued) because of strict quasi-concavity of the utility function. Continuity comes from strictly positive prices and continuity of preferences.

In the next step, we define the excess demand function which will allow us to prove existence of the competitive equilibrium.

Definition 6.4.3 (excess demand function)**.** *The aggregate **excess demand function for good** $j \in \{1, ..., N\}$ is defined as:*

$$z_j(p) = \sum_{i=1}^{N} c^i(p, p \cdot e^i) - \sum_{i=1}^{N} e^i$$

*The aggregate **excess demand function** is a vector-valued function collecting all the aggregate excess demands for each good:*

$$z(p) = \big(z_1(p), z_2(p), ..., z_N(p)\big)$$

If $z_j(p) > 0$, then the consumers demand jointly more quantity of good j than exists on the market in the form of endowments, i.e. we observe excess demand. If $z_j(p) < 0$, then the consumers demand less than the total quantity existing on the market, leading to an excess market supply of that good. Intuitively, in equilibrium we would expect that the demanded quantities are equal to supply quantities. Therefore, in search for equilibrium we will be looking for a vector of prices p^* that zeroes the excess demand function.

Proposition 6.4.2 (properties of aggregate excess demand)**.** *Given assumptions of this chapter, for all $p \gg 0$, the aggregate demand function $z(p)$:*

1. *is continuous in p,*
2. *is homogeneous of degree zero in p: $z(tp) = z(p)$ for every $t > 0$,*
3. *satisfies Walras' law:*

$$p \cdot z(p) = 0$$

Proof. 1. continuity follows from continuity of the individual demand functions. Note that the demand is continuous only throughout the range of strictly positive prices. When price of a good is zero, it brings a few technical complications that we skip in out class.

2. homogeneity also follows from homogeneity of the individual demand functions. To see this, note that $c^i(p, p \cdot e^i)$ solves the consumer problem at prices p and income $p \cdot e^i$. That means, $p \cdot c^i = p \cdot e^i$. Multiplying both sides by $t > 0$, we get $tp \cdot c^i = tp \cdot e^i$. Note that this is still the same constraint! So if $c^i(tp, tp \cdot e^i)$ solves the problem at the rescaled budget constraint, it has to be the same demand function as the one that solves the original consumer problem:

$$c^i(p, p \cdot e^i) = c^i(tp, tp \cdot e^i), \ t > 0$$

In the last step:

$$z^j(tp) = \sum_{i=1}^{N} c^i(tp, tp \cdot e^i) - \sum_{i=1}^{N} e^i = \sum_{i=1}^{N} c^i(p, p \cdot e^i) - \sum_{i=1}^{N} e^i = z_j(p), \ j \in \{1, .., N\}, \ t > 0$$

3. Walras' law comes from the strict monotonicity of preferences that imply that the budget constraint holds in equality:

$$
\begin{aligned}
p \cdot c^i(p, p \cdot e^i) = p \cdot e^i \Leftrightarrow & \sum_{j=1}^{N} p_j \left(c^i(p, p \cdot e^i) - e^i_j\right) = 0 \\
\underbrace{\Leftrightarrow}_{\text{sum over individuals}} \quad & \sum_{i=1}^{I}\sum_{j=1}^{N} p_j \left(c^i(p, p \cdot e^i) - e^i_j\right) = 0 \\
\Leftrightarrow & \sum_{j=1}^{N}\sum_{i=1}^{I} p_j \left(c^i(p, p \cdot e^i) - e^i_j\right) = 0 \\
\Leftrightarrow & \sum_{j=1}^{N} p_j \sum_{i=1}^{I} \left(c^i(p, p \cdot e^i) - e^i_j\right) = 0 \\
\Leftrightarrow & \sum_{j=1}^{N} p_j z_j(p) = 0
\end{aligned}
$$

□

The Walras' law states that the value of excess demand is always zero in the economy (at positive prices). Therefore, if there is excess demand of one good, then there also must be excess supply of another good. We can see this even more clearly on market with two goods.

Example 6.4.1 (Walras law in the two-good economy)**.** *In two good economy, Walras' law implies that:*

$$p_1 z_1(p_1, p_2) + p_2 z_2(p_1, p_2) = 0$$

Then, for example, the excess demand for good 1 implies excess supply for good 2. However, equilibrium on market 1, i.e. $z_1(p) = 0$*, implies that* $z_2(p) = 0$*, hence equilibrium on the market for good 2!*

These considerations can be generalized for N-good economy. It turns out that it only matters what happens on $N - 1$ market, in the sense that the situation on N^{th} market is then implied.

Proposition 6.4.3 (market clearing)**.** *Suppose the demand equals supply on* $N-1$ *markets and* $p_N > 0$*. Then the demand must equal supply also on the* N^{th} *market.*

Exercise 6.4.1. *Prove the proposition above. Hint: show that* $z_N(p) \neq 0$ *would violate the Walras' law.*

In the earlier paragraphs, we suggested that while looking for an equilibrium, we will focus on prices that zero the excess demand. This is a sufficient approach if we can rule out zero prices. For completeness, we consider what happens when prices are zero.

Proposition 6.4.4 (free goods)**.** *If good* j *is in excess supply at Walrasian equilibrium, then it must be a free good:*

$$z_j(p) < 0 \Rightarrow p_j = 0$$

Proof. Since p is a Walrasian equilibrium, it has to be that $z_j(p) \leq 0$. You can show this by deriving the excess demand functions under the assumption that the budget constraint holds with inequality $\leq$. Since prices are non-negative, $p \cdot z(p) \leq 0$. If $z_j(p) < 0$ and $p_j > 0$, we would get that $p \cdot z(p) < 0$, which contradicts Walras' law. □

Definition 6.4.4 (desirability)**.** *Good* j *is* ***desirable*** *if*

$$p_j = 0 \Rightarrow z_j(p) > 0$$

Proposition 6.4.5 (equality of demand and supply)**.** *If all goods are desirable and* p^* *is the price vector in a Walrasian equilibrium, then* $z(p^*) = 0$*.*

Proof. Suppose not. Then there is a good for which the excess demand is negative: $z_j(p^*) < 0$. By proposition 6.4.4, it must be that $p^*_j = 0$. Since all goods are desirable, $z_j(p^*) > 0$, which is a contradiction. □

Recall from the consumer theory that the homogeneity property of individual demand function implies that nominal price levels do not matter for the optimal choice. In other words, if we rescale all the prices and income by the same positive constant, the optimal choice remains the same. In the exchange economies, the income is defined through prices (and initial endowment) – it is enough to rescale all prices by the same factor to get the same demand.

This property implies that prices are not unique in the equilibrium. Proposition 6.4.3 shows the source of non-uniqueness: we have an excess degree of freedom. You can think of it as a system of N equations with $N-1$ unknowns – one has to be redundant. To fix this **indeterminacy** issue, we need a **normalization**.

Proposition 6.4.6 (price normalizations)**.** *The following price systems are without loss of generality, i.e. generate the same consumer choices as the default price vector:*

1. ***numeraire** system:* $\Big\{p \in \mathbb{R}^N_+ : p_1 = 1\Big\}$.
2. ***simplex** system:* $\Big\{p \in \mathbb{R}^N_+ : \sum_{j=1}^N p_j = 1\Big\}$.

Proof. Let $c^i(p, p \cdot e^i)$ be an individual demand function at prices $p \geq 0$.

1. **numeraire** system. Let $p_1 > 0$. Using homogeneity:

$$c^i(p, p\cdot e^i) = c^i(p_1, p_2, ..., p_N, \sum_{j=1}^N p_j e^i_j) = c^i(\frac{p_1}{p_1}, \frac{p_2}{p_1}, ..., \frac{p_N}{p_1}, \sum_{j=1}^N \frac{p_j}{p_1} e^i_j) = c^i(1, \hat{p}_2, ..., \hat{p}_N, \hat{p}\cdot e^i) = c^i(\hat{p}, \hat{p}\cdot e^i)$$

   where $\hat{p}_1 = 1$ and $\hat{p}_j = \frac{p_j}{p_1}$ for $j > 1$.

2. **simplex** system. Using homogeneity:

$$\begin{aligned} c^i(p, p\cdot e^i) &= c^i(p_1, p_2, ..., p_N, \sum_{j=1}^N p_j e^i_j) = c^i(\frac{p_1}{\sum_{j=1}^N p_j}, \frac{p_2}{\sum_{j=1}^N p_j}, ..., \frac{p_N}{\sum_{j=1}^N p_j}, \sum_{j=1}^N \frac{p_j}{\sum_{j=1}^N p_j} e^i_j) = c^i(\hat{p}_1, \hat{p}_2, ..., \hat{p}_N, \hat{p}\cdot e^i) \\ &= c^i(\hat{p}, \hat{p}\cdot e^i) \end{aligned}$$

   where $\hat{p}_j = \frac{p_j}{\sum_{j=1}^N p_j}$ for $j \in \{1, ..., N\}$.

□

Proposition 6.4.6 implies that the choice of specific price normalization is irrelevant from the perspective of analyzing the demand. However, some normalizations may be more convenient for working with our model. For example, we would typically use a numeraire system while solving for equilibria. In turn, we will prove existence of equilibrium using the simplex formulation.

Theorem 6.3 (existence of competitive equilibrium)**.** *Suppose $z : S^{N-1} \to \mathbb{R}^N$ is a continuous function that satisfies Walras' law. Then there exists $p^* \in S^{N-1}$ such that $z(p^*) \leq 0$.*

Proof. Define the following mapping:

$$g_j(p) = \frac{p_j + \max\{0, z_j(p)\}}{1 + \sum_{k=1}^N \max\{0, z_k(p)\}},\ j \in \{1, ..., N\}$$

Since z_j's and max are continuous functions, g_j is continuous for each j. Moreover, $g_j(p) \geq 0$ and:

$$\sum_{j=1}^N g_j(p) = \sum_{j=1}^N \frac{p_j + \max\{0, z_j(p)\}}{1 + \sum_{k=1}^N \max\{0, z_k(p)\}} = \frac{\sum_{j=1}^N p_j + \max\{0, z_j(p)\}}{1 + \sum_{k=1}^N \max\{0, z_k(p)\}} = \frac{1 + \sum_{j=1}^N \max\{0, z_j(p)\}}{1 + \sum_{k=1}^N \max\{0, z_k(p)\}} = 1$$

This follows from the fact that since we consider the simplex price normalization, the prices sum up to one.

By Brouwer's fixed point theorem, there exists a fixed point of this mapping: $p^* = g(p^*)$:

$$p^*_j = \frac{p^*_j + \max\{0, z_j(p^*)\}}{1 + \sum_{k=1}^N \max\{0, z_k(p^*)\}},\ j \in \{1, ..., N\}$$

It is left to show that p^* is a Walrasian equilibrium. Using the equation above, we get:

$$p_j^*\big(1+\sum_{k=1}^{N}\max\{0,z_k(p^*)\}\big)=p_j^*+\max\{0,z_j(p^*)\}\underset{\text{re-arrange}}{\Leftrightarrow} p_j^*\sum_{k=1}^{N}\max\{0,z_k(p^*)\}=\max\{0,z_j(p^*)\}$$

$$\underset{\text{multiply by} z_j(p^*)}{\Leftrightarrow} z_j(p^*)p_j^*\sum_{k=1}^{N}\max\{0,z_k(p^*)\}=z_j(p^*)\max\{0,z_j(p^*)\}$$

$$\underset{\text{sum over goods}}{\Leftrightarrow}\sum_{j=1}^{N}z_j(p^*)p_j^*\sum_{k=1}^{N}\max\{0,z_k(p^*)\}=\sum_{j=1}^{N}z_j(p^*)\max\{0,z_j(p^*)\}$$

$$\Leftrightarrow\Big(\sum_{k=1}^{N}\max\{0,z_k(p^*)\}\Big)\Big(\sum_{j=1}^{N}z_j(p^*)p_j^*\Big)=\sum_{j=1}^{N}z_j(p^*)\max\{0,z_j(p^*)\}$$

By Walras law, $\sum_{j=1}^{N}z_j(p^*)p_j^*=0$, so it must be that:

$$\sum_{j=1}^{N}z_j(p^*)\max\{0,z_j(p^*)\}=0$$

Note that each of the terms above are non-negative because they either equal to zero or $z_j(p^*)^2$. Hence, if any of there terms was strictly positive, the equality would not hold. Therefore, every term has to be equal to zero. Therefore, it must be either $z_j(p)<0$ or $z_j(p)=0$—hence $z_j(p^*)\leq 0$ for every $j\in\{1,..,N\}$. □

The proof above is pretty general, as it allows as a possibility that $z_j(p^*)<0$. In our applications, every good will be desirable, so by proposition 6.4.5 implies that the equilibrium price vector p^* that zeroes the excess demand exists.

Theorem 6.3 says that a vector of equilibrium prices p^* exists. By proposition 6.4.1, we know that the equilibrium allocation $c^*=(c^{1*},...,c^{I*})$, such that $c^{1*}\equiv c^i(p^*,p^*\cdot e^i)$ for $i\in\{1,...,N\}$ also exists! This completes our analysis of existence of Walrasian equilibrium.

Focus on how general the existence theorem is. It only needs the excess demand to be continuous and satisfy Walras' law. There is just one problem there: the excess demand is not likely to be continuous around the boundary, that is, where some prices are equal to zero. This challenge is no longer an issue if the aggregate excess demand function is unbounded from above as some (but not all) prices approach zero [1]. It can be proved that this is implied by a non-zero aggregate endowment of each good and our assumptions on the utility functions.

6.5 Welfare Theorems

What is the connection between competitive equilibrium and Pareto-efficiency? Now, we finally can discuss the two main results of today's class: the celebrated **welfare theorems**. The first one says that competitive equilibria are efficient. In other words, the market mechanism itself guarantees allocation in which there is no room for improvement for anyone without hurting others.

Theorem 6.4 (1st welfare theorem for exchange economies)**.** *If (p,c) is a Walrasian equilibrium, then c is Pareto-efficient.*

Proof. Suppose not: then (p,c) is a Walrasian equilibrium but c is not Pareto-efficient. It means that there exists another feasible allocation $\overline{c}$ such that is $\overline{c}^i$ weakly preferred than c^i by all consumers and strictly preferred by at least one of them.

We begin the proof with two useful claims.

1. **Claim 1.1**: $u(\overline{c}^i)>u(c^i)$ implies that $p\cdot\overline{c}^i>p\cdot c^i$. Suppose not. Then $u(\overline{c}^i)>u(c^i)$ but $p\cdot\overline{c}^i\leq p\cdot c^i$. Since c^i is the optimal choice, it means that it must be affordable. But $\overline{c}^i$ is weakly cheaper, hence also affordable for a consumer at prices p. But this contradicts c^i being the optimal choice.

[1] formally it means that if $\{p^m\}$ is a sequence of price vectors in $\mathbb{R}^N_{++}$ indexed by m and converging to $\overline{p}\neq 0$ and $\overline{p}_j=0$ for some good j, then for some good k with $\overline{p}_k=0$ the associated sequence of excess demands in the market for good k $\{z_k(p^m)\}$ is unbounded from above.

2. **Claim 1.2**: $u(\overline{c}^i) \geq u(c^i)$ implies that $p \cdot \overline{c}^i \geq p \cdot c^i$. Suppose not. Then $u(\overline{c}^i) \geq u(c^i)$ but $p \cdot \overline{c}^i < p \cdot c^i$. Therefore, $\overline{c}^i$ is affordable. If $u(\overline{c}^i) > u(c^i)$, than it contradicts optimality of c^i for the same reason as it does in claim 1. Suppose now that $u(\overline{c}^i) = u(c^i)$. Consider another allocation $\hat{c}^i$ such that $\hat{c}^i_j = \overline{c}^i_j + \frac{p \cdot c^i - p \cdot \overline{c}^i}{p_j N}$ for every good j (note this requires all prices non-negative. It would be straightforward to account for zero prices – we skip it for notational convenience). Since $p \cdot c^i - p \cdot \overline{c}^i > 0$, it follows that $\hat{c}^i \gg \overline{c}^i$. By strict monotonicity of preferences, it must be strictly preferred than c^i: $u(\hat{c}^i) > u(\overline{c}^i) = u(c^i)$. Moreover:

$$p \cdot \hat{c}^i = \sum_{j=1}^{N} p_j \left(\overline{c}^i_j + \frac{p \cdot c^i - p \cdot \overline{c}^i}{p_j N} \right) = p \cdot \overline{c}^i + p \cdot c^i - p \cdot \overline{c}^i = p \cdot c^i$$

so $\hat{c}^i$ is also affordable. But this contradicts optimality of c^i.

With the two claims above, we note that as c is not Pareto-efficient, the following holds:

$$\begin{cases} u(\overline{c}^i) \geq u(c^i) & \text{for every } i \in \{1,..,N\} \\ u(\overline{c}^i) > u(c^i) & \text{for some } i \in \{1,..,N\} \end{cases} \underset{\text{claims 1 and 2}}{\Rightarrow} \begin{cases} p \cdot \overline{c}^i \geq p \cdot c^i & \text{for every } i \in \{1,..,N\} \\ p \cdot \overline{c}^i > p \cdot c^i & \text{for some } i \in \{1,..,N\} \end{cases}$$

Summing across consumers the two equations on the right:

$$\sum_{i=1}^{N} p \cdot \overline{c}^i > \sum_{i=1}^{N} p \cdot c^i = \sum_{i=1}^{N} p \cdot e^i$$

Where the last equation follows from strictly increasing preferences. Hence:

$$\sum_{i=1}^{N} p \cdot (\overline{c}^i - e^i) > 0 \Rightarrow p \cdot \left(\sum_{i=1}^{N} \overline{c}^i - e^i \right) > 0$$

Since $p \geq 0$, it must be that there is a good j for which $p_j > 0$ and $\sum_{i=1}^{N} \overline{c}^i_j > \sum_{i=1}^{N} e^i_j$. But this contradicts feasibility of $\overline{c}$. □

The second welfare theorem goes the other way around. It tells us that for every Pareto-efficient allocation, we can redistribute the initial endowments and find a vector of prices that would support this allocation as a competitive equilibrium. We begin with introducing a concept of competitive equilibrium with transfers.

Definition 6.5.1 (competitive equilibrium with transfers)**.** *An allocation c^* is supported in* ***competitive equilibrium with transfers*** *if there exists a price vector p^* and a set of lump-sum transfers $T = (T^1, T^2, ..., T^I)$ that are budget-balanced budget-balanced, i.e. $\sum_{i=1}^{I} T_i = 0$, such that:*

1. *given prices, every consumer make optimal choices:*

$$c^{*i} = \arg\max_{x \in C} u^i(x) \text{ such that } p^* \cdot x = p^* \cdot e^i + T^i$$

2. *allocation is feasible:*

$$\sum_{i=1}^{I} c^{*i}_j \leq \sum_{i=1}^{I} e^i_j$$

Finally, we are ready to state and prove the second welfare theorem.

Theorem 6.5 (2^{nd} welfare theorem for exchange economies)**.** *Suppose c^* is a Pareto-efficient allocation in which every consumer holds a positive amount of each good. Under our assumptions on preferences (utility functions continuous, strictly increasing, strictly quasi-concave), c^* can be supported in competitive equilibrium with transfers.*

Proof. Define a set of bundles that consumer i strictly prefers over their bundle c^{*i} in allocation c^*:

$$P^i = \left\{ c^i \in \mathbb{R}^N_+ : u(c^i) > u(c^{*i}) \right\}$$

This set is non-empty (since utilities are strictly increasing) and convex (since utilities are quasi-concave).

Define also the sum of P^i sets:

$$P = \sum_{i=1}^{I} P^i = \Big\{ \sum_{i=1}^{N} c^i : c^i \in P^i \text{ for every } i \in \{1, ..., I\} \Big\}$$

Set P is the set of bundles that could be redistributed among our N consumers to make each of them better off. As a sum of non-empty and convex sets, it is non-empty and convex.

Consider a bundle summing over the initial endowment bundles: $\omega = \sum_{i=1}^{N} e^i$.

Claim 2.1: $\omega \notin P$. Suppose not: $\omega \in P$. Then it can be written as a sum of bundles $\omega = \sum_{i=1}^{I} x^i$ such that $x^i \in P^i$, so $u(x^i) > u(x^{*i})$ for every i. At the same time, as a redistributed sum of endowments, it is feasible. Hence, $x = (x^1, ..., x^I)$ Pareto-dominates c^*, which contradicts c^* being Pareto-efficient. So it must be that $\omega \notin P$. □

Now, consider two sets: $\{\omega\}$ (a singleton) and P. They are disjoint, as $\omega \notin P$. They both are non-empty and convex. Then, by the separating hyperplane theorem, there exists a non-zero vector p^* such that:

$$p^* \cdot z \geq p^* \cdot \omega \text{ for all } z \in P$$

Now, we need to show that p^* can be interpreted as a price vector and that it supports c^* as a competitive equilibrium with transfers.

Claim 2.2. $p^* \geq 0$. Let $1^j = (0, 0, ..., 0, 1, 0.., 0)$ be j^{th} unit vector, that is having 1 as j^{th} coordinate and zeroes otherwise. Consider a vector:

$$\sum_{i=1}^{I} c^{*i} + 1^j \equiv (\sum_{i=1}^{I} c_1^{*i}, \sum_{i=1}^{I} c_2^{*i}, ..., \sum_{i=1}^{I} c_N^{*i}) + 1^j$$

It has a unit more of good j than an analogous vector constructed from our Pareto-efficient allocation c^*. By strict monotonicity of preferences, the additional unit can be redistributed across consumers making everybody strictly better off, so this vector is an element of P. Since p^* separates $\{\omega\}$ and P:

$$\begin{aligned} p^* \cdot \Big(\sum_{i=1}^{I} c^{*i} + 1^j\Big) \geq p^* \cdot \omega &\Leftrightarrow p^* \cdot \Big(\sum_{i=1}^{I} c^{*i} + 1^j - \omega\Big) \geq 0 \\ &\Leftrightarrow p^* \cdot \Big(\sum_{i=1}^{I} c^{*i} + 1^j - \sum_{i=1}^{I} c^{*i}\Big) \geq 0 \\ &\Leftrightarrow p^* \cdot 1^j \geq 0 \\ &\Leftrightarrow p_j^* \geq 0 \end{aligned}$$

where the second equivalence comes from that fact that $\sum_{i=1}^{I} c^{*i} = \omega$ (strictly increasing preferences imply that every Pareto-efficient allocation has to use all the resources), and the fourth uses the fact that 1^j is the j^{th} unit vector. Since this works for any $j \in \{1, .., N\}$, we just proved that every coordinate from the separating hyperplane p^* is non-negative, hence it is a perfectly legitimate price vector. □

In the next step, we move our attention to the transfers. Define:

$$T^i = p^* \cdot c^{*i} - p^* \cdot e^i$$

Claim 2.3. Transfers T^i are balanced. Note that:

$$\begin{aligned} \sum_{i=1}^{I} T_i =& \sum_{i=1}^{I} p^* \cdot c^{*i} - p^* \cdot e^i \\ =& p^* \cdot \sum_{i=1}^{I} c^{*i} - p^* \cdot \sum_{i=1}^{I} e^i \\ =& p^* \cdot \omega - p^* \cdot \omega \\ =& 0 \end{aligned}$$

because $\sum_{i=1}^{I} c^{*i} = \omega = \sum_{i=1}^{I} e^i$. □

What is left to show is that c^* is a competitive equilibrium allocation, i.e. is feasible and maximizes consumers' utility. The first part follows trivially from the fact that c^* is Pareto-efficient, hence feasible. The second part requires more work. In the following claims, $\overline{c} = (\overline{c}^1, ..., \overline{c}^I)$ denote another allocation.

Claim 2.4. If $u^i(\overline{c}^i) > u^i(c^{*i})$ then $p \cdot \overline{c}^i \geq p \cdot c^{*i}$ for every consumer $i \in \{1, ..., N\}$. Suppose then $u^i(\overline{c}^i) > u^i(c^{*i})$. Then, by continuity of utility function, there exists $\varepsilon > 0$ sufficiently small such that $u^i\big(\overline{c}^i(1-\varepsilon)\big) > u^i(c^{*i})$. Define another allocation x such that:

1. $x^i = (1-\varepsilon)\overline{c}^i$ — so that $u(x^i) > u(c^{*i})$
2. $x^j = c^{*j} + \frac{\varepsilon}{I-1}\overline{c}^i$ for $j \neq i$ — so that $u(x^j) > u(c^{*j})$ (by strict monotonicity)

The new allocation x is essentially a redistribution of $\varepsilon\overline{c}^i$ across the remaining consumers. Note that every consumer i strictly prefers x^i over c^{*i}, hence $x \in P$. Then:

$$\begin{aligned}
p^* \cdot (\overline{c}^i - c^{*i}) =& p \cdot \Big(\overline{c}^i + \sum_{j\neq i} c^{*j} - \sum_{j=1}^{N} c^{*j}\Big) \\
=& p^* \cdot \Big((1-\varepsilon)\overline{c}^i + \varepsilon\overline{c}^i + \sum_{j\neq i} c^{*j} - \sum_{j=1}^{N} c^{*j}\Big) \\
=& p^* \cdot \Big((1-\varepsilon)\overline{c}^i + \sum_{j\neq i} \big(c^{*j} + \frac{\varepsilon\overline{c}^i}{I-1}\big) - \sum_{j=1}^{N} c^{*j}\Big) \\
=& p^* \cdot \Big(x^i + \sum_{j\neq i} x^j - \sum_{j=1}^{N} c^{*j}\Big) \\
=& p^* \cdot \Big(\sum_{j=1}^{N} x^j - \sum_{j=1}^{N} c^{*j}\Big) \\
=& p^* \cdot \Big(\sum_{j=1}^{N} x^j - \sum_{j=1}^{N} e^j\Big) && (6) \\
=& p^* \cdot \Big(x - \omega\Big) && (7) \\
\geq& 0 && (8)
\end{aligned}$$

where (6) follows because since c^* is Pareto-efficient then $\sum_{i=1}^N c^{*i} = \sum_{i=1}^N e^i$, (7) follows because $\sum_{i=1}^N e^i = \omega$, and (8) follows because of the separating hyperplane argument (as $x \in P$). So we conclude that if $u^i(\overline{c}^i) > u^i(c^{*i})$ then $p \cdot \overline{c}^i \geq p \cdot c^{*i}$ for every consumer $i \in \{1, ..., N\}$. □

Claim 2.5. If $u^i(\overline{c}^i) > u^i(c^{*i})$ then $p \cdot \overline{c}^i > p \cdot c^{*i}$. Suppose not. Then $u^i(\overline{c}^i) > u^i(c^{*i})$ and $p \cdot \overline{c}^i \leq p \cdot c^{*i}$. By previous claim, we can rule out strict inequality, so it has to be that $p \cdot \overline{c}^i = p \cdot c^{*i}$. Note that $c^{*i} \gg 0$ (by assumption—every consumer holds a positive amount of every good) and $p^* \geq 0$ (by claim above), it follows that $p^* \cdot c^{*i} > 0$ and thus $p \cdot \overline{c}^i > 0$.

By continuity, there exists $0 < \delta < 1$ small enough such that $u^i\big((1-\delta)\overline{c}^i\big) > u^i(c^{*i})$. By the previous claim: $p^* \cdot (1-\delta)\overline{c}^i \geq p^* \cdot c^{*i}$.

Combining the two results together:

$$p^* \cdot (1-\delta)\overline{c}^i \geq p^* \cdot \overline{c}^i \Leftrightarrow (1-\delta)p^* \cdot \overline{c}^i \geq p^* \cdot \overline{c}^i \underset{p^* \cdot \overline{c}^i > 0}{\Leftrightarrow} 1 - \delta \geq 1 \Leftrightarrow \delta \leq 0$$

which is a contradiction. So it must be that if $u^i(\overline{c}^i) > u^i(c^{*i})$ then $p \cdot \overline{c}^i > p \cdot c^{*i}$. □

That concludes the proof of 2WFT. Note that if at least one consumer strictly prefers a bundle $\overline{c}^i \neq c^{*i}$ over c^{*i}, then $p^* \cdot \sum_{i=1}^I \overline{c}^i > p^* \cdot \sum_{i=1}^I c^{*i} = p^* \cdot \sum_{i=1}^I e^i$. This in turn implies $\sum_{i=1}^I \overline{c}^i_j > \sum_{i=1}^I e^i_j$ for at least one good j, so $\overline{c}$ is not a feasible allocation. □

Recall our previous considerations that not every Pareto-efficient allocation can be a result of trade given the initial endowment, because some agents would not voluntarily engage in trade that deteriorates their position. In turn, the second welfare theorem says that any Pareto-efficient allocation can be supported as a competitive equilibrium with transfers. The transfers are key here. By redistributing the initial endowment across consumers (potentially worsening positions of some of them), we can always find a price vector that will support the resulting Pareto-efficient allocation as a competitive equilibrium.

6.6 Remarks

The first welfare theorem is a positive result. It says that competitive markets are efficient—no resources are wasted. It does not say anything about *fairness*. For example, an allocation in which one consumer has all the goods in the economy is Pareto-efficient, but it may contradict some desired definitions of ethical optimality. It is the second welfare theorem that produces some hints on **welfare** analysis. It separates efficiency from redistribution. Depending on a definition of a *fair* allocation, a social planner is able to support it in a competitive equilibrium provided redistributive transfers that alter consumers' wealth can be enforced on them. Note that neither of the theorems says anything about how much a consumer *should* consume.

Welfare theorem emphasize the dual role of prices. They are **allocative**, reflecting (relative) scarcity of goods. They are also **redistributive**, deciding about how much of goods a consumer can purchase and hence consume.

Today we studied the welfare theorems in the simplest setting of exchange economies. However, these results easily extend for many other settings, including competitive markets with production, competitive markets with time dynamics and forward looking consumers, competitive markets with overlapping generations, competitive markets with uncertainty, and many more.

The welfare theorems are very powerful results. They uncover very appealing properties of competitive markets. Not only they are efficient—the invisible market hand equilibrated the demand and supply in a decentralized setting in which consumers only care about themselves, but also leaves some space for equity engineering.

However, the excitement you may feel about these results should be tamed. All of these results rely on *competitiveness* of the market. Any departure from perfectly competitive markets, for example through asymmetric information, product differentiation, externalities, or existence of oligopoly/monopoly, would invalidate them. Nevertheless, they constitute a very useful benchmark which we can always use to assess how much we lose if we depart from the competitive markets.

7 Measuring Welfare

Today, we are moving from the positive (prediction of behavior) to the normative (assessment of behavior) economic analysis. We are interested in assessing the market outcomes from the social point of view. Specifically, we are looking for meaningful ways to determine whether an outcome is better or worse than another outcome. This is very important piece of economics. It enables us to define and quantify *losses* associated with departures from the competitive markets. It allows us to formulate assessments regarding government interventions and hence produce relevant economic advisory. It help us evaluate the value of new goods on the market, or non-market goods.

The core of today's class is how to measure consumer utility. Our analysis departs from the ideal definition of consumer welfare in desired units, which is typically impossible to retrieve from data. We proceed with looking for imperfect but more available alternatives and study how much we lose while using them to infer about consumer welfare.

We define many important concepts whose significance and applications reach way beyond the consumer theory. Nevertheless, we stick to consumer theory as language of describing these concepts for simplicity of exposition.

List of contents. Partial equilibrium. Compensating variation. Equivalent variation. Consumer surplus. Quasi-linear preferences. Quantity indices: ideal, Laspeyres, Paasche. Revealed preference argument. Minimum income function. Money-metric utility. Standard of living. Cost of living indices: shopping basket, ideal, Laspeyres, Paasche. Local and global estimates of demand. Deep primitives of a model. Sufficient statistics. Deadweight loss. Harberger triangle. General equilibrium. Sufficient statistics approach in taxation.

Textbook references. Espinola: Ch. 5; Varian: Ch. 10; Jehle & Reny: Ch. 4.3; MWG: Ch. 3I, Deaton & Muellbauer: Ch. 7

Literature References.

[1] Harberger, A. C. (1964). The measurement of waste. The American Economic Review, 54(3), 58-76.

[2] Chetty, R. (2009). Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. Annu. Rev. Econ., 1(1), 451-488.

7.1 Mathematical Glossary

Definition 7.1.1 (area under curve)**.** *Area under curve given by a function $f : \mathbb{R}^N \to \mathbb{R}^N$ between points a and b is given by the integral:*

$$\int_a^b f(x)\mathrm{d}x = F(b) - F(a)$$

where $F(x) = \int f(x)\mathrm{d}x$.

Theorem 7.1 (Riemann integral)**.** *Consider function $f : \mathbb{R}^N \to \mathbb{R}^N$ and interval $[a, b]$. Divide the interval into n subintervals of equal with $\Delta x = \frac{b-a}{n}$. Choose a sample point within each interval $x_i^* \in [x_{i-1}, x_i]$. The area under f on $[a, b]$ is the limit of:*

$$\int_a^b f(x)\mathrm{d}x = \lim_{n\to\infty} \sum_{i=1}^{n} f(x_i^*)\Delta x$$

Definition 7.1.2. *Suppose function $f : \mathbb{R}^N \to \mathbb{R}$ depends on N arguments, out of which every one is a function of t:*

$$f(x_1, ..., x_N) \equiv f\Big(x_1(t), .., x_N(t)\Big)$$

The derivative of f with respect to t is given by

$$\frac{\mathrm{d}f}{\mathrm{d}t} = \sum_{i=1}^{N} \frac{\mathrm{d}f}{\mathrm{d}x_i}\frac{\mathrm{d}x_i}{\mathrm{d}t}$$

Theorem 7.2 (envelope theorem)**.** *Consider a function that is defined as a result to an optimization problem:*

$$V(\theta) = \max_x f(x, \theta)$$

Under appropriate regularity conditions (continuity, differentiability, ...):

$$\frac{\mathrm{d}V}{\mathrm{d}\theta} = \frac{\mathrm{d}f}{\mathrm{d}\theta}$$

7.2 Compensating and Equivalent Variation

We begin with the most accurate definitions of welfare. Our discussion today will focus on partial equilibrium markets, which we define below.

Definition 7.2.1 (partial equilibrium)**.** *Analysis of equilibrium at a single market assuming that other markets remain unaffected is called* ***partial equilibrium*** *analysis.*

In partial equilibrium analysis, the equilibrium price and quantity are determined by the intersection of the supply and demand curves for that particular market.

Example 7.2.1 (partial equilibrium)**.** *Suppose the government imposes a tax on apples. In partial equilibrium, we analyze how this tax affects the supply and demand for apples, the equilibrium price, and the quantity of apples sold. We do not consider how the tax on apples might affect the market for oranges, or any other market.*

With this background, we begin to think about welfare. We build some intuitions by deriving the desired objects from the analysis of what happens to consumer if price of one good increases.

To introduce the notation, consider a consumer facing an increase in price of one of the goods, say, good i. Denote price of this good by p_i, and $N - 1$ dimensional vector of the remaining prices by $p_{-i} = (p_1, p_2, ..., p_{i-1}, p_{i+1}, ..., p_N)$. We want to focus exclusively on changes in price of good i, and assume there are no changes on markets for other goods. With this assumption in mind, write the indirect utility function as:

$$v(p_i, p_{-i}, w) \equiv v(p_i, w)$$

Note that we do not want to measure these changes in units of utility, or so called, **utils**. This is because i) the utility representation of preference is ordinal, so the absolute values of utility do not carry any meaning. Recall that any strictly increasing transformation of utility function represents the same preferences; and ii) even if the utility representation was cardinal, utils are not observed and hence are not convenient for analysis.

Instead, we want to define welfare measures in monetary terms.

Definition 7.2.2 (compensating variation)**.** *Suppose price of good i changes from p_i to p_i'. **Compensating variation** (CV) is the change in consumer income necessary to keep the utility level unchanged when a price of a good changes:*

$$v(p_i', w + CV) = v(p_i, w)$$

Compensating variation is a measure of welfare expressed in monetary terms. By its construction, is of main interest for policy advisory.

Exercise 7.2.1. *Show that under our assumptions, if $p_i' > p_i$ then $CV > 0$.*

The next proposition highlights the fact that compensating variation is a valid measure of welfare, as a function of Hicksian demand (i.e. holding the utility level fixed).

Proposition 7.2.1. *The compensating variation CV can be expressed as:*

$$CV = \int_{p_i}^{p_i'} c_i^h(p, v)\mathrm{d}p$$

where $c_i^h(p_i, v) \equiv c(p_i, p_{-i}, v)$ is the Hicksian demand evaluated at the initial value of indirect utility $v \equiv v(p_i, w)$.

Proof. Since $v(p_i', w + CV) = v(p_i, w)$, it must be that:

$$e\big(p_i', v(p_i, w)\big) = e\big(p_i', v(p_i, w + CV)\big) = w + CV$$

Moreover, since $w = e\big(p_i, v(p_i, w)\big)$, we can plug it to the equation above. Denote $v \equiv v(p_i, w)$. Then:

$$CV = e(p_i', v) - w = e(p_i', v) - e(p_i, v)$$

Lastly, using Shephard's lemma:

$$\begin{aligned} CV =& e(p_i', w) - e(p_i, w) \\ =& \frac{\partial e(p, v)}{\partial p}\Big|_{p=p_i'} - \frac{\partial e(p, v)}{\partial p}\Big|_{p=p_i} \\ =& \int_{p_i}^{p_i'} \frac{\partial e(p, v)}{\partial p}\mathrm{d}p \\ =& \int_{p_i}^{p_i'} c^h(p, v)\mathrm{d}p \end{aligned}$$

$\square$

Therefore, the compensating variation is simply the area under the Hicksian demand curve between points p_i and p_i'.

A related measure is the equivalent variation.

Definition 7.2.3 (equivalent variation)**.** *Suppose price of good i changes from p_i to p_i'. **Equivalent variation** (EV) describes how much of their income the consumer would like to sacrifice in order to avoid price change from p_i to p_i':*

$$v(p_i, w - EV) = v(p_i', w)$$

Proposition 7.2.2. *Let $v' \equiv v(p_i', w)$. The equivalent variation EV can be expressed as:*

$$EV = \int_{p_i}^{p_i'} c_i^h(p, v')\mathrm{d}p$$

Exercise 7.2.2. *Prove the above proposition.*

The difference between the compensating variation CV and equivalent variation EV is rather subtle. CV uses the new price as the base and asks what income change would be necessary to compensate the consumer for the price change. EV takes the current price as the base and asks what income change at the current prices would be equivalent to the proposed change in terms of its impact on utility.

The magnitudes of these two will typically vary because the value of a unit of currency depends on the relative prices. In turn, their sign will remain the same because they measure the same utility differences, just using a different (indirect) utility function.

Depending on what we want to measure, one can be better than the other. If we are looking for a compensation scheme under new prices, CV is the way to go. Otherwise, if we are after a reasonable measure of **willingness to pay** for being in a given state, EV would be a more natural choice. This is because *i*) EV measures the income change at current prices—it is easier for decision makers to judge the value of a *dollar* at current prices than a some hypothetical prices; and *ii*) if we are comparing a variety of different scenarios (i.e. different prices after a change), it is natural to seek a common base prices for comparisons—in this case, the current prices.

To further understand the differences between CV and EV, consider the following example.

Example 7.2.2 (graphical interpretation of CV and EV)**.** *Consider a consumer's problem with two goods: c_1 and c_2. Without loss, normalize the price of good two to one: $p_2 = 1$, treating it as a numeraire. Suppose that both goods are normal and that the price of good one increases from p_1 to $p_1' > p_1$. The first figure illustrates this change. Initially, the budget constraint has a lower slope, and is tangent to the (optimal) indifference curve at point E. The increase in price of good one rotates the budget constraint around its intersection with c_2 axis. The slope increases, and the new optimum is in E'.*

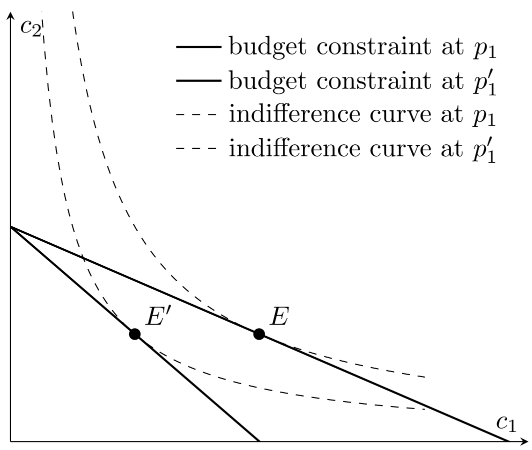

Now we can consider the compensating and equivalent variation. Since c_2 is a numeraire, changes in c_2 are equivalent to changes in monetary values. The ***compensating variation*** *denotes the monetary value of c_2 (since its numeraire—also the quantity) that needs to be given to the consumer at the* ***new prices*** *to bring them back to the* ***original indifference curve****. This amount is denoted in blue in the left panel below.*

The ***equivalent variation*** *denotes the monetary value of c_2 (since its numeraire—also the quantity) that needs to be taken away from the consumer at the* ***original prices*** *to put the consumer on the* ***new indifference curve****. This amount is denoted in red in the right panel below.*

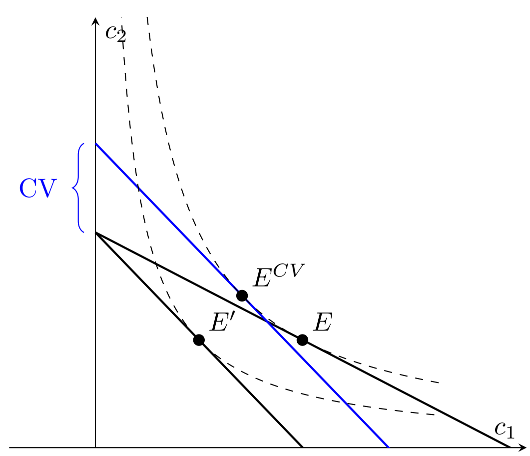

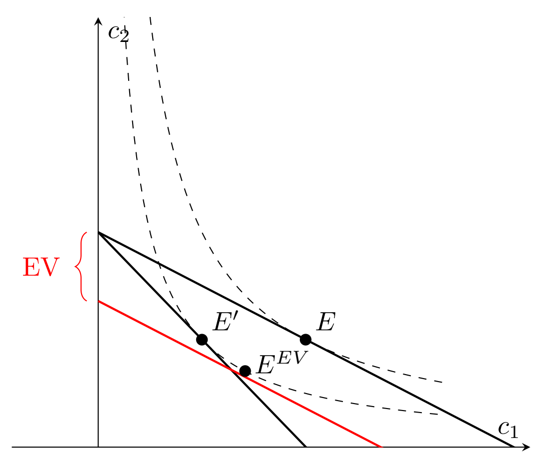

Unfortunately, compensating (and equivalent) variation is as accessible for us as the Hicksian demand is. In the next section we are looking for another measure that would (at least) approximate compensating variation and be dependent on Marshallian demand which we can observe in the data.

7.3 Consumer Surplus

The compensating and equivalent variation are measures of consumer utility—hence welfare—expressed in monetary terms. If Hicksian demand was easily observable, they would have great applicability in economics. Unfortunately, their dependence on a fixed value of indirect utility makes them less usable.

To make progress in searching for empirically good measures of welfare, we turn our attention to the Marshallian demand and associated concept of **consumer surplus**.

Definition 7.3.1 (consumer surplus)**.** *Consumer surplus is the difference between what agents are willing to pay for a product and what they actually pay.*

To build intuition consider the following the example.

Example 7.3.1. *The* ***demand curve*** *describes how much consumers are willing to pay for different quantities of a good. Each point on this curve indicates the maximum price a consumer is willing to pay for one more unit of the good. In turn, the* ***market price*** *is the actual price at which the good is sold. This is usually lower than what some consumers are willing to pay, which creates a surplus. The* ***consumer surplus*** *is the total benefit consumers get from paying less than what they are willing to pay. It's the area between the demand curve and the market price line, up to the quantity purchased.*

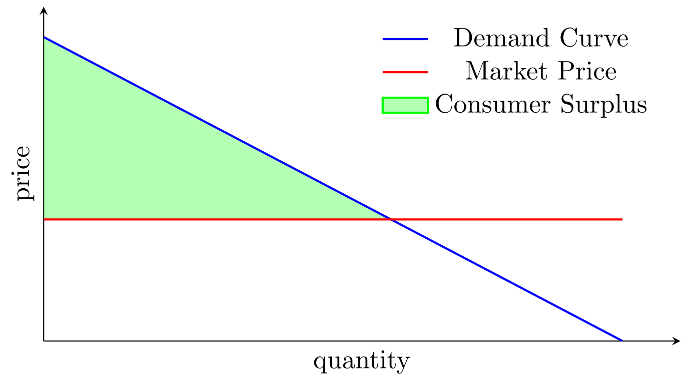

Definition 7.3.2 (consumer surplus in the quantity space)**.** *Consumer surplus is the difference between what customers are willing to pay for a product and what they actually pay.*

$$CS = \int_0^c \Big(p^D(x) - p_{\text{market}}\Big)\mathrm{d}x$$

where $p^D(x)$ is the inverse demand function evaluated at quantity bundle x, and c is the equilibrium quantity on the market.

We can also look at the same problem from a different perspective.

Example 7.3.2. *Consider the same plot but in a version in which the axes are reversed.*

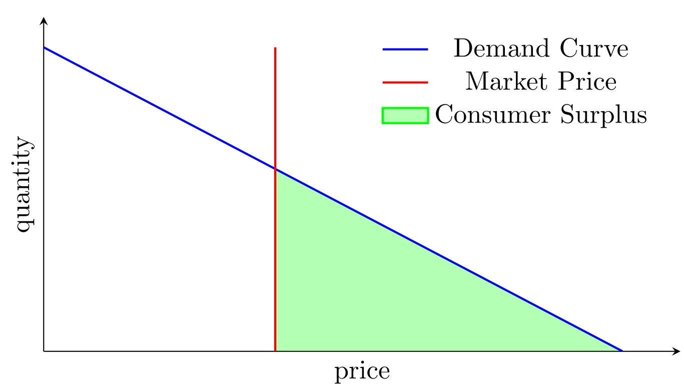

In this case, we want to integrate the demand over all prices larger than the market price.

Definition 7.3.3 (consumer surplus in the price space)**.** *Consumer surplus is the difference between what customers are willing to pay for a product and what they actually pay. Suppose* $c_i(p_i, w) \equiv c_i(p_i, p_{-i}, w)$ *is the Marshallian demand for good* i *at prices* (p_i, p_{-i}) *and income* w*. The* ***consumer surplus*** *is defined as:*

$$CS = \int_{p_{\text{market}}}^{\infty} c(p, w)\mathrm{d}p$$

The above formulation of consumer surplus allows us to consider surplus increments related to changes in price.

Definition 7.3.4 (change in consumer surplus)**.** *Change in consumer surplus as a result of price change is:*

$$\Delta CS \equiv CS(p_i', w) - CS(p_i, w) = \int_{p_i}^{p_i'} c_i(p, w)\mathrm{d}p$$

With this definition, we can see the relation between consumer surplus ΔCS and compensating variation CV. Both are obtained from integrating a measure of demand over the range of price change. However, ΔCS holds income fixed, whereas CV holds the utility level fixed. Clearly, the correct measure of welfare—the monetary value of utility—is CV. Consumer surplus is more of a measure of gains from trade because it holds income fixed and allows for (better) utilities.

We would typically want to know CV but can only calculate ΔCS. Is ΔCS a good approximation of CV?

Proposition 7.3.1. *For normal goods:*

$$EV \leq \Delta CS \leq CV$$

Proof. Consider an increase in price of good i from p_i to p_i'. We want to show that:

$$\int_{p_i}^{p_i'} c_i^h(p, v')\mathrm{d}p \leq \int_{p_i}^{p_i'} c_i(p, w)\mathrm{d}p \leq \int_{p_i}^{p_i'} c_i^h(p, v)\mathrm{d}p$$

Using propositions discussed in the consumer theory chapter, it follows that at the endpoints of integration we get:

$$\begin{cases} c_i(p_i', w) = c_i^h(p_i', u') \\ c_i(p_i, w) = c_i^h(p_i, u) \end{cases}$$

Moreover, by Slutski equation:

$$\underbrace{\frac{\partial c_i(p, w)}{\partial p_i}}_{\text{total effect}} = \underbrace{\frac{\partial c_i^h(p, u)}{\partial p_i}}_{\text{substitution effect}} \underbrace{- c_i(p, w)\frac{\partial c_i(p, w)}{\partial w}}_{\text{income effect}}$$

Note that since the Marshallian demand is weakly increasing in income w (normal good), it follows that:

$$\begin{cases} \frac{\partial c_i(p', w)}{\partial p'} \leq \frac{\partial c_i^h(p', u')}{\partial p'} \\ \frac{\partial c_i(p, w)}{\partial p} \leq \frac{\partial c_i^h(p, u)}{\partial p} \end{cases}$$

Now, since at $p = p_i$ we have $c_i(p, w) = c_i^h(p, u)$ and $\frac{\partial c_i(p,w)}{\partial p} \leq \frac{\partial c_i^h(p,u)}{\partial p}$, then for $p > p_i$ it must be that $c_i(p_i, w) \leq c_i^h(p_i, u)$. Moreover, since at $p = p_i'$ we have $c_i(p_i', w) = c_i^h(p_i', u')$ and $\frac{\partial c_i(p,w)}{\partial p} \leq \frac{\partial c_i^h(p,u)}{\partial p}$, for $p < p_i'$ it must be that $c_i(p, w) \geq c_i^h(p, u')$.

The result then follows. □

The mechanics between the proposition above can be understood better using a graphical example.

Example 7.3.3 (bounds on CS—a graphical example)**.** *Consider Hicksian demand curve as a set of points in which we could change prices and income at the same time to remain at the same indifference curve, hence holding utility level fixed. Suppose we only analyze changes in prices of good* i*. Let* u *be the original utility level, and* $u' < u$ *the utility level at the new optimum resulting from increase in price of good* i*. The Hicksian demand curves as defined above can be illustrated graphically in* (c_i, p_i) *space.*

Using the definitions of EV *and* CV*, we quickly note that they are equal to the areas as depicted at the left and center panel below, respectively:*

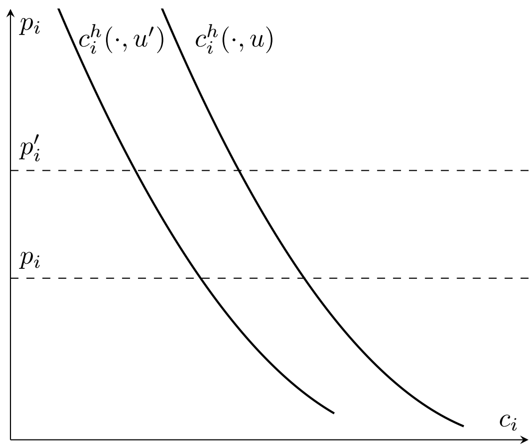

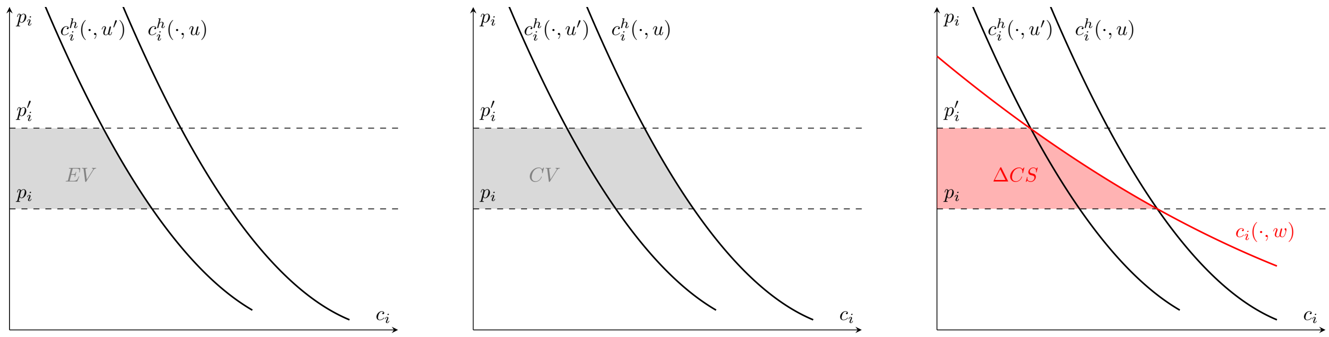

Now, we add the Marshallian demand curve (defined analogously) to the picture. Note that it needs to cross the Hicksian demand associated with u' at p_i', and the Hicksian demand associated with u at p_i. We also expect it to be decreasing in prices. By the definition of ΔCS, it is equal to the red area marked at the right panel above.

Therefore, the (change) in consumer surplus is composed of the equivalent variation plus a part of compensating variation. As we show in detail below, these differences are related to income effects that arise when prices change.

The proposition above gives a very promising result bounding the consumer surplus between compensating and equivalent variation. In the next step, we would like to learn when this approximation is good. The first step is to introduce quasi-linear preferences.

Definition 7.3.5 (quasi-linear preferences)**.** *We say that preferences are **quasi-linear** (with respect to good 1) when the utility function that represents them can be expressed as a strictly increasing transformation of a quasi-linear function:*

$$u(c) = c_1 + \sum_{i=2}^{N} u_i(c_i)$$

where $u_i(\cdot)$ are uni-dimensional, concave functions.

Proposition 7.3.2. *Suppose c is an interior solution to the consumer utility maximization problem with quasi-linear utility function. Then the demand for good $i \geq 2$ does not depend on income:*

$$c_i(p, w) = c_i(p_i), i \geq 2$$

Proof. Normalize $p_1 = 1$ to be numeraire. Write down the Lagrangian:

$$\mathcal{L}(c, \lambda) = c_1 + \sum_{i=2}^{N} u_i(c_i) - \lambda(p \cdot c - w)$$

The first order conditions are:

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial c_1} =& 1 - \lambda p_1 = 0 \\ \frac{\partial \mathcal{L}}{\partial c_i} =& u_i'(c_i) - \lambda p_i = 0, \ i \geq 2 \\ \frac{\partial \mathcal{L}}{\partial \lambda} =& p \cdot c - w = 0\end{aligned}$$

From the first FOC, we learn that $\lambda = 1$. Plugging it into the second FOC, we get:

$$u_i'(c_i) = p_i$$

If u_i is strictly quasi-concave, the marginal utility is strictly monotone. Therefore, we can invert the equation above to get the (optimal) consumption in terms of (exclusively!) price of this good. Therefore, we conclude that:

$$c_i(p, w) = c_i(p_i), i \geq 2$$

□

Example 7.3.4 (quasi-linear utility in two dimensions)**.** *Consider a simple example of quasi-linear utility with two goods. The consumer faces the following utility maximization problem:*

$$\max_{c \geq 0} c_1 + u(c_2) \text{ such that } c_1 + p_2 c_2 = w$$

Note that we normalized the price of good one to one, thus treat it as a numeraire. Assume interior solution to the problem above, meaning that in the optimum the consumer uses strictly positive amounts of both goods.

Think of starting with $w = 0$ and adding slowly additional units of income. If the marginal utility of consuming $c_2 = w$ is smaller than 1, the consumer would only choose positive c_1, hence we wouldn't be in the interior solution. Hence, it must be that for small values of w, $u'(w) > 1$. Given well behaved preferences, the marginal utility is going to monotonically decrease. Eventually, the consumer will choose c_2^* such that $u'(c_2^*) = 1$. This will cost the consumer $p_2 c_2^*$ monetary units. The remaining income will be spent on good one.

Obviously, c_2^* does not depend on w but c_1 does. In general, only the optimal quantity of the numeraire good is going to depend on income in quasi-linear world.

Proposition 7.3.3. *Suppose consumer's preferences are quasi-linear. Then:*

$$EV = \Delta CS = CV$$

Proof. The proof follows the same steps as the proof of bounds for ΔCS above. The only difference is that with quasi-linear preferences, the income effects in Slutsky's equation are zero:

$$c_i(p, w) \underbrace{\frac{\partial c_i(p, w)}{\partial w}}_{=0} = 0$$

As a result, the Marshallian demand is equal to the Hicksian demands at the endpoints of integration and their derivatives are the same for every $p \in (p_i, p_i')$. Thus it must be that the integrals are equal to each other, which yields the desired result. □

The results above show us that ΔCS diverges from CV (and EV) if the demand depends on income. This is due to the so called income effects of a price change. If the income effect are not very large, for example because the change from p_i to p_i' is relatively small, then ΔCS should provide a good approximation to the total amount of income that consumers are willing to give up for the price decrease.

7.4 Quantity Indices

The approaches discussed above aimed at quantifying the absolute value of welfare effects expressed in monetary terms. As we have seen, this is not an easy task. Sometimes we may be tempted to approximate welfare using simpler methods, though potentially on the cost of lower accuracy of results.

This brings us to methods that assess changes in welfare without assigning them an absolute-value numbers. We can use the **index number theory** to compare two vectors (for example: prices or quantities) by using just a scalar[2]. In this section, we consider quantity indices, focusing on comparing expenditures needed to attain different level of utility at the same prices.

Definition 7.4.1 (ideal welfare index)**.** *The **ideal welfare index** is defined as:*

$$U(u', u, p) = \frac{e(p, u')}{e(p, u)}$$

How to interpret u and u'? Following Deaton & Muellbauer (1980), we could think of u' as utility of \$100k-a-year family and u as utility of \$80k-a-year family, in the sense that $e(p, u)$ describe the minimum expenditures necessary to achieve a given *standard of living* given prices p.

The ideal welfare index is even harder to measure, as it depends on two unobserved utility levels. Despite this obstacle, we still can make some progress.

Definition 7.4.2 (Laspeyres quantity index)**.** *The **Laspeyres** quantity index is defined as:*

$$Q^{\text{Laspeyres}} = \frac{p \cdot c'}{p \cdot c}$$

Definition 7.4.3 (Paasche quantity index)**.** *The **Paasche** quantity index is defined as:*

$$Q^{\text{Paasche}} = \frac{p' \cdot c'}{p' \cdot c}$$

Under some circustance, the Laspeyers and Paasche indices may allows us to rank the consumption baskets. If $Q^{\text{Laspeyres}} \leq 1$, that is $p \cdot c' \leq p \cdot c$, we can rank c as at least as good as c'. By the **revealed preference argument**, c' was affordable at prices p but not chosen. Analogously, if $Q^{\text{Paasche}} \geq 1$, so $p' \cdot c' \geq p' \cdot c$, by the **revealed preference argument** we can infer that c' is at least as good as c because c was affordable at prices p' yet not chosen.

7.5 Cost of Living Indices

In this section, we want to quantify changes in price levels, thus inflation. This is a difficult task, because changes in prices induce both substitution and income effects that confound the analysis.

We focus on cost of living indices as price indices relying on some **standard of living** as a reference. In the simplest example, we consider a single, pre-defined consumption bundle $\overline{c}$ as the reference level.

Definition 7.5.1 (shopping basket price index)**.** *Consider a bundle $\overline{c}$. The **shopping basket price index** is defined as:*

$$P(p', p, \overline{c}) = \frac{p' \cdot \overline{c}}{p \cdot \overline{c}}$$

Unfortunately, defining a standard of living through a single consumption bundle is pretty restrictive. We would rather think of an indifference curve as a good reference to a standard of living.

Definition 7.5.2 (ideal cost of living index)**.** *Fix a utility level $\overline{u}$. The **ideal** (true) **cost of living index** is defined as:*

$$P(p', p, \overline{u}) = \frac{e(p', \overline{u})}{e(p, \overline{u})}$$

where $e(\cdot, \cdot)$ is the expenditure function.

[2] Note that the problem is trivial when the vectors are proportional, e.g. $p' = 5p$. However, usually they are not proportional, so we need to develop methods allowing for meaningful comparisons

The ideal cost of living measures exactly what we are after—changes in expenditures required to maintain the same level of utility. It is analogous to the compensating variation in the sense that it holds utility level fixed at u. Therefore, for the same reasons as compensating variation, it is not very useful in practice.

Therefore, we consider price indices that may approximate the ideal cost of living measure.

Definition 7.5.3 (Laspeyres price index)**.** *The **Laspeyres** price index is defined as:*

$$P^{\text{Laspeyres}} = \frac{p' \cdot c}{p \cdot c}$$

Definition 7.5.4 (Paasche price index)**.** *The **Paasche** price index is defined as:*

$$P^{\text{Paasche}} = \frac{p' \cdot c'}{p \cdot c'}$$

The Laspeyres and Paasche indices are related to the ideal index.

Proposition 7.5.1. *The following two facts hold:*

1.

$$P^{\text{Laspeyres}} = P(p', p, c) = \frac{p' \cdot c}{p \cdot c} \geq \frac{e(p', \overline{u})}{e(p, \overline{u})} = P(p', p, \overline{u})$$

2.

$$P^{\text{Paasche}} = P(p', p, c') = \frac{p' \cdot c'}{p \cdot c'} \leq \frac{e(p', \overline{u}')}{e(p, \overline{u}')} = P(p', p, \overline{u}')$$

Exercise 7.5.1. *Prove the proposition above. Hint: use the properties of expenditure function.*

The proposition above says that the Laspeyres index sets an upper bound for the ideal index at $\overline{u}$ and the Paasche index sets a lower bound for the ideal index at $\overline{u}'$. This is a pretty general statement. In general, the Laspeyres index can be both higher or lower than the Paasche index, so the ideal index does not need to lie between these two.

In general, it is difficult to think of *the* ideal index. Since it depends on $\overline{u}$, changes in prices will affect the cost of living differently for different individuals if their total level of expenditures differ. $\overline{u}$ vanishes when preferences are homothetic, in which case the expenditure function can be written as $c(p, u) = ub(p)$ for some function $b(p)$. Then, the ideal index takes form $\frac{b(p')}{b(p)}$ which doesn't depend in u. Homotheticity is sufficient and necessary condition for existence of *the* unique price index (the price index). It is also the only case in which the ideal index lies between the Paasche and Laspeyres indices. However, homotheticity is unlikely to hold, and we need to allow for the effects of welfare on cost-of-living indices.

7.6 Sufficient Statistics

To bring the welfare analysis to data we typically need to estimate demand at each possible price (and perhaps also supply). Only then we can integrate the area under the demand curve and hence calculate compensating variation or consumer surplus. Getting **local estimates** (in the neighborhood of equilibrium) of demand (and supply) is difficult due to simultaneity bias—we only observe the equilibrium outcomes, but not demand and supply separately. Getting **global estimates** of demand (and supply) is perhaps even more difficult, as identification of the **deep primitives** of the model (for example, parameters of utility function, ...) typically requires either extremely rich data (rarely available) and/or strong functional form assumptions.

Given these obstacles, researchers were looking for other ways of assessing welfare effects of policies. A clever solution is to consider the so called **sufficient statistics**. The idea is to use the theoretical model to derive formulas describing consequences of policies of interest that are functions of high-level elasticities rather than deep primitives of the model. Such elasticities are typically much easier to identify in the existing data.

The sufficient statistics approach essentially combines two paradigms in empirical economics. It is derived from theory and hence allow us to make precise statements about welfare. In turn, depending only on high level elasticities, it allows us to estimate the welfare effects of policies using reduced form methods, and hence achieve transparent and credible identification.

The mechanics of sufficient statistics approach is summarized at figure 3.

We illustrate usefulness of sufficient statistics approach with a simple example of **general equilibrium** model that dates back to Harberger (1964).

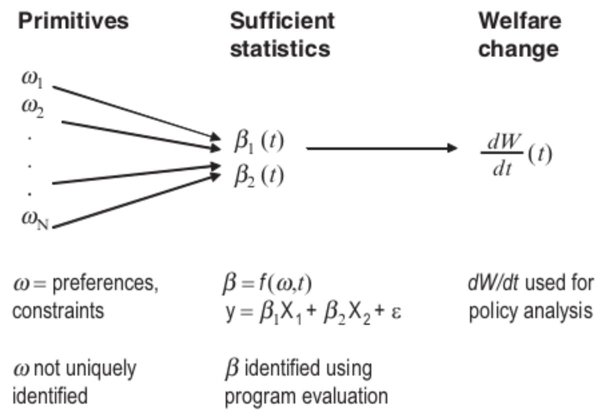

Figure 1

The sufficient-statistic approach. Here we consider a policy instrument t that affects social welfare $W(t)$. The structural approach maps the primitives (ω) directly to the effects of the policy on welfare ($\frac{dW}{dt}$). The sufficient-statistic approach leaves ω unidentified and instead identifies a smaller set of high-level parameters (β) using program-evaluation methods, e.g., via a regression of an outcome y on exogenous variables X. The β vector is sufficient for welfare analysis in that any vector ω consistent with β implies the same value of $\frac{dW}{dt}$. Identifying β does not identify ω because there are multiple ω vectors consistent with a single β vector.

Figure 3: The sufficient statistics approach. *Source: Chetty (2009).*

Definition 7.6.1 (general equilibrium)**.** *Analysis of equilibrium jointly at a all markets is called* ***general equilibrium*** *analysis.*

Harberger is interested in measuring deadweight loss associated with imposing a commodity tax.

Definition 7.6.2 (deadweight loss)**.** *The* ***deadweight loss*** *measures economic inefficiency that occurs when the competitive equilibrium outcome is not achieved.*

The intuition is that a commodity tax introduces a negative supply shock. Producers offer the same quantity at higher prices to (at least partially) cover the increased tax liability. The good now is more expensive, limiting the purchasing power of individuals. This means the demand decreases, so as a result the equilibrium production is also going down. Lower production also means less inputs are used. In other words, the whole dynamism of an economy slows down. The efficiency (or deadweight) cost of a tax increase is equal to the loss in surplus from the transactions that fail to occur because of the tax.

This situation is illustrated on figure 4 using simple economy with linear demand and supply.

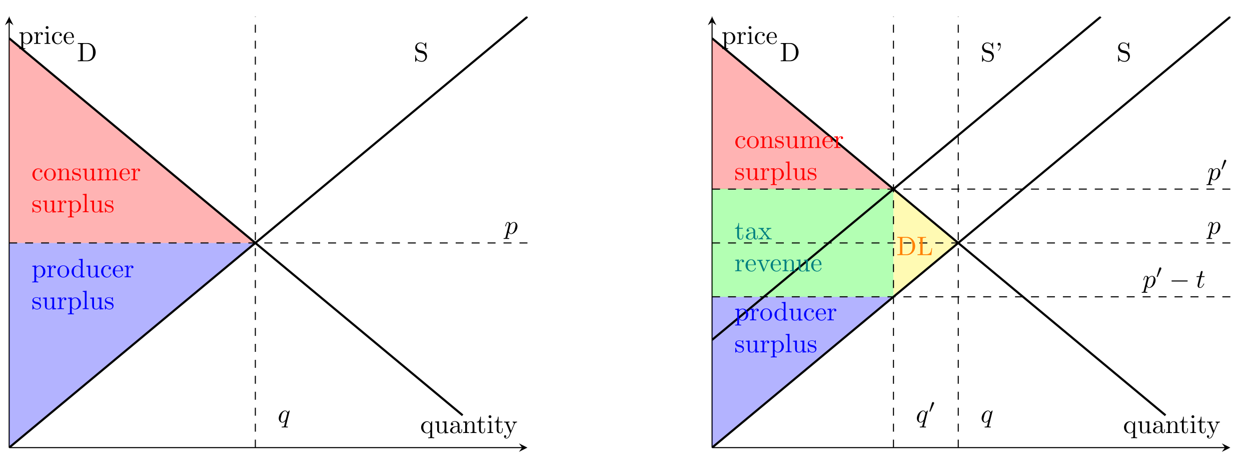

Figure 4: Harberger triangle (yellow on the right panel) – deadweight loss from taxation.

Imposition of tax t—a negative supply shock—shifts the supply curve upwards. The equilibrium quantity decreases ($q \to q'$), while equilibrium price increases ($p \to p'$). The total surplus in the economy now is divided among three agents: consumer, government, and producer. However, a part of the initial surplus is lost after imposition of the tax—it is no longer captured by any agent. This surplus is marked yellow on the right panel 4, also known as the **Harberger triangle** summarizing the deadweight loss from taxation.

Quantifying the amount of deadweight loss is of central interest for any policy maker who would like to predict consequences of their actions. We now move to discussing how to identify it from the data. We follow derivations by Chetty (2009) and analyze a simplified general equilibrium model with production and quasi-linear preferences (to abstract from income effects).

Suppose a representative consumer[3] is endowed with $\overline{L}$ units of the numeraire good ℓ (with price normalized to one). We can interpret this good as labor—firms convert it into J consumption goods $x = (x_1, .., x_J)$. Producing x_j units of good j requires $c_j(x_j)$ units of labor, so the total cost of production is $c(x) = \sum_{j=1}^{J} c_j(x_j)$. Assume that producer market is competitive and there is a representative producer. Let $p = (p_1, ..., p_J)$ be a vector of equilibrium prices and consider imposing an ad-valorem tax t on good $j = 1$.

The consumer takes prices as given and solves:

$$\max_{x,\ell} u(x_1, ..., x_J) + \ell$$
$$\text{such that } p \cdot x + tx_1 + \ell = \overline{L}$$

where $u(\cdot)$ is, by assumption, strictly quasi-concave. Suppose the solution to the consumer's problem, i.e. the Marshallian demand, is given by $x^D(p)$.

The firm takes prices as given and solves:

$$\max_{x} p \cdot x - c(x)$$

Suppose the solution to the firm's problem, i.e. the factor demand, is given by $x^S(p)$.

In equilibrium, prices equate demand and supply. The market clearing condition is:

$$x^D(p) = x^S(p)$$

From now on, consider $p(t)$ as market-clearing prices as a function of tax. To measure the inefficiency, we take the surplus from a tax-distorted allocation and then return the tax revenue to a consumer in a form of a lump-sum rebate. With quasi-linear utility, the consumer would just allocate the lump sum transfer into the numeraire good.

The social welfare is the sum of consumer utility (money-metric given quasi-linearity), producer profits, and tax revenue:

$$W(t) = \underbrace{\max_{x} \Big\{u(x) + \overline{L} - p(t) \cdot x - tx_1\Big\}}_{\text{consumer utility}} + \underbrace{\max_{x} \Big\{p(t) \cdot x - c(x)\Big\}}_{\text{firm profits}} + \underbrace{tx_1}_{\text{tax revenue}}$$

Since the same x maximizes the consumer and the firm objectives, we can simplify this expression further to:

$$W(t) = \max_{x} \Big\{u(x) + \overline{L} - tx_1 - c(x)\Big\} + tx_1$$

With the social surplus defined in terms of deep primitives of the model, we can do three things, First, we can estimate a J-good supply model and associated demand system to recover (in most applications—parametrized) utility and cost functions, which in turn would allow us to compute $W(t)$ directly. Second, we can fit a demand system to the data and recover the sought-after value indirect utility (what is possible given what we learned in meeting 2). In both cases, what we are after is estimating the full model with all its deep primitives. This usually requires $2J$ instruments (J for demand, and J for supply), but in general equilibrium exercises we typically expect J to be (very) large. Eventually, the third way is to look for **sufficient statistics** that would relieve us from the necessity of identifying the deep primitives of the model.

To proceed, differentiate the previous equation with respect to t using the envelope theorem:

$$\begin{aligned}
\frac{\mathrm{d}W(t)}{\mathrm{d}t} &= \frac{\mathrm{d}u(x)}{\mathrm{d}x}\frac{\mathrm{d}x}{\mathrm{d}t} - x_1 - t\frac{dx_1}{dt} - \frac{\mathrm{d}c(x)}{\mathrm{d}x}\frac{\mathrm{d}x}{\mathrm{d}t} + x_1 + t\frac{dx_1}{dt} \\
&= \Big(\frac{\mathrm{d}u(x)}{\mathrm{d}x_1} - \frac{\mathrm{d}c(x)}{\mathrm{d}x_1}\Big)\frac{\mathrm{d}x_1}{\mathrm{d}t} + \sum_{j=2}^{J}\Big(\frac{\mathrm{d}u(x)}{\mathrm{d}x_j} - \frac{\mathrm{d}c(x)}{\mathrm{d}x_j}\Big)\frac{\mathrm{d}x_j}{\mathrm{d}t}
\end{aligned}$$

[3]the model yields the same result with consumer heterogeneity (what is one of its major strengths) which we skip for notational simplicity.

Now, note that the first order conditions for consumer and producer imply that:

$$\frac{\mathrm{d}u}{\mathrm{d}x_1} = p_1 + t = \frac{\mathrm{d}c}{\mathrm{d}x_1} + t$$
$$\frac{\mathrm{d}u}{\mathrm{d}x_j} = p_j = \frac{\mathrm{d}c}{\mathrm{d}x_j},\ j \in \{2, .., J\}$$

Plugging it back to the main equation:

$$\frac{\mathrm{d}W(t)}{\mathrm{d}t} = \underbrace{\Big(\frac{\mathrm{d}u(x)}{\mathrm{d}x_1} - \frac{\mathrm{d}c(x)}{\mathrm{d}x_1}\Big)}_{=t}\frac{\mathrm{d}x_1}{\mathrm{d}t} + \underbrace{\sum_{j=2}^{J}\Big(\frac{\mathrm{d}u(x)}{\mathrm{d}x_j} - \frac{\mathrm{d}c(x)}{\mathrm{d}x_j}\Big)}_{=0}\frac{\mathrm{d}x_j}{\mathrm{d}t}$$
$$= t\frac{\mathrm{d}x_1}{\mathrm{d}t}$$

we get the main result. A sufficient statistics for welfare effects associated with the commodity tax is just the effect of the tax on the consumer demand in equilibrium!

By estimating $\frac{\mathrm{d}x_1}{\mathrm{d}t}$ for different values of t, we can calculate welfare consequences for any tax change within the support of t:

$$\Delta W = W(t_1) - W(T_0) = \int_{t_0}^{t_1} t\frac{\mathrm{d}x_1}{\mathrm{d}t}\mathrm{d}t$$

In particular, we do not need to estimate the full demand (or supply) curve.

The sufficient statistics approach naturally carries some limitations. We can only address small changes around the equilibrium—note that it is based on local derivations around the FOCs. That means we cannot use sufficient statistics for evaluating large counterfactual policies describing large changes in the tax on x_1 and any other good.

Lastly, we note that quasi-linearity of preferences brings very convenient interpretation. The surplus is also the compensating and equivalent variation. The derivations get a little bit (but not much) more complicated in the general case. See Harberger (1964) and Chetty (2009) (and references therein) for more details.

8 Strategic Interactions

Until now, we mostly focused on competitive markets. Agents take prices as given and their optimal actions don't depend on behavior of others. Now, we relax this assumption and move to analyzing strategic interactions between agents.

We begin with reviewing basic concepts of non-cooperative game theory that provides a language to analyze the strategic interactions. We use this language to talk about oligopoly markets, in which the number of competing firms is limited and so actions of one firm affect actions of other firms. We study a variety of arrangements in which firms may compete: by choosing quantity or price, with and without leadership of a dominant firm, and collusion. Although there is no a uniformly correct model—the structure of interactions depends on many factors, including nature of the production process—there certainly are *preferred* models. We show that different structures of firm interactions can typically be ranked according to their impact on welfare in the economy.

List of contents. Non-cooperative game theory. Information set. Action and actions set. Strategies. Pure strategies. Mixed strategies. Strategy profile. Payoff function. Best response. Nash equilibrium. Normal form game. Oligopoly. Inverse demand function. Demand function. Cournot model. Bertrand model. Models of leadership: quantity and price. Collusion. Welfare in oligopolistic games. Cartel.

Textbook references. Espinola: Ch. 12,14 ; Varian: Ch. 15,16; Jehle & Reny: Ch. 4, 7; MWG: Ch. 7,12.C

Literature References.

[1] Nash Jr, John F. (1950). Equilibrium points in n-person games. Proceedings of the national academy of sciences, 36(1), 48-49.

[2] Tremblay, C. H., and Tremblay, V. J. (2011). The Cournot–Bertrand model and the degree of product differentiation. Economics Letters, 111(3), 233-235.

8.1 Non-cooperative Game Theory

In this section, we introduce **non-cooperative games**, in which N agents play with each other without cooperation, i.e. they choose actions to maximize their own payoff. We study interactions of firm on the market using the language of non-cooperative game theory.

We begin with defining the primitives of a game. We allow the game to be multi-stage. In each stage a, as subset of players (could be all N agents, and could be only some of them) make their move simultaneously. In general, one player may be able to move in many stages.

Definition 8.1.1 (information set). *An **information set** I^a is a collection of decision points in a game that a player cannot distinguish between based on the information available to them at that a stage in the game.*

Definition 8.1.2 (action). *An **action** a_i is a move of a player at a given stage of the game.*

Definition 8.1.3 (action set). *Set $\mathcal{A}_i^a$ is a **set of all actions** of player i available at decision stage a.*

We need to distinguish actions from strategies.

Definition 8.1.4 (strategy: a loose definition). *A **strategy** σ_i is a complete plan of **behavior** that a player will follow throughout the game, contingent on every possible situation they may encounter.*

Definition 8.1.5 (pure strategy). *A **pure strategy** s_i is a complete plan of **actions** that a player will follow throughout the game, i.e. at each decision stage a in which the player makes their move, contingent on every possible situation they may encounter. That means, a pure strategy is a mapping from the set of information sets into the set of actions.*

Definition 8.1.6 (pure strategies set). *Set $\mathcal{S}_i$ is a **set of all pure strategies** of player i.*

Pure strategies are important type of strategies. They are very intuitive as a list of actions to make at various decision stages. Pure strategies are not the only strategies we may think of.

Definition 8.1.7 (mixed strategy). *Suppose player i is endowed with a set of pure strategies S_i. A **mixed strategy** σ_i is a probability distribution over the set of pure strategies, which for each pure strategy $s_i^k \in S_i$ assigns probability σ_i^k such that:*

$$\sigma_i^k \geq 0 \text{ for all } k, \qquad \sum_k \sigma_i^k = 1$$

Definition 8.1.8 (mixed strategies set). *Set $\Delta(\mathcal{S}_i)$ is a **set of all mixed strategies** of player i. This is essentially the set of all probability distributions defined over the set of actions.*

The above definition fits the situation in which the set of pure strategies is finite. It is straightforward to generalize it into infinite sets of pure strategies.

Exercise 8.1.1. *Show that every pure strategy is also a mixed strategy.*

Denote any set of relevant player i's strategies by Σ_i. We combine strategies of all players into a strategy profile.

Definition 8.1.9 (strategy profile). *Suppose player i has a set of all strategies Σ_i (can be pure or mixed). A **strategy profile** is a vector of possible strategies, one for each player:*

$$\sigma = (\sigma^1, ..., \sigma^N) \in \times_{i=1}^N \Sigma_i$$

Definition 8.1.10 (set of all strategy profiles). *The set of all strategy profiles is the Cartesian product of sets of individual strategies.*

$$\Sigma = \times_{i=1}^N \Sigma_i$$

So far, we have defined the environment associated with players' moves. Now, we move to defining what motivates their moves: payoffs.

Definition 8.1.11 (payoff function). *The **payoff function** $u : \times_{i=1}^N \Sigma_i \to \mathbb{R}^N$ is a vector-valued function assigning a vector of payoffs—one for each player—for every strategy profile σ:*

$$u(\sigma) = \big(u_1(\sigma), ..., u_N(\sigma)\big)$$

*where $u_i(\sigma)$ are **individual payoffs**.*

What is solution of the game? Which strategies we should expect players to choose? We adopt an **equilibrium concept** in which an equilibrium describes a situation in which no player has an incentive to deviate from their strategy given strategies played by other players.

Definition 8.1.12 (individual best response)**.** *Suppose all agents except agent i play strategies* $\sigma_{-i} = \{\sigma_j : \sigma_j \in \Sigma_j,\ j \neq i\}$. *A **Best Response** of player i to the set of strategies of other players is defined as:*

$$BR_i(\sigma_{-i}) = \{\sigma_i \in \Sigma_i : u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma_{i'}, \sigma_{-i}) \text{ for every } \sigma^{i'} \in \Sigma_i\}$$

Definition 8.1.13 (best response)**.** *The **Best Response** correspondence collects best responses of every player:*

$$BR(\sigma) = \Big\{BR_i(\sigma_{-i}) : i \in \{1, ..., N\}\Big\}$$

The definition above defines the Best Response as a correspondence. It may not be a function: for every set of strategies of other players $-i$, we can have more than one strategy of player i that is a best response to the rest.

Definition 8.1.14 (Nash Equilibrium)**.** *A strategy profile* $\sigma \in \Sigma$ *is a **Nash Equilibrium** if for every player i:*

$$\sigma_i = BR_i(\sigma_{-i})$$

Theorem 8.1 (existence of Nash Equilibrium (Nash))**.** *Suppose the set of strategies is finite. Then, every game has a mixed strategy Nash equilibrium.*

By definition, a Nash equilibrium is a fixed point of a Best Response correspondence. To prove the existence of Nash equilibrium, it suffices to apply a fixed point theorem to prove existence of a fixed point. The fact that the Best Response is in general a correspondence not a function brings a little bit more technical difficulty, which gets beyond the scope of our class. Thus we skip the proof of existence theorem here. Nonetheless, we can build some intuition about the proof, as it follows similar steps as proving existence of a fixed point in the general equilibrium existence proof. Loosely speaking, we need to show that the Best Response maps a closed and convex set into itself and is continuous and convex, so it must have a fixed point.

Nash' proof (Nash, 1950) was less than two pages long, yet essentially paved him a way to Nobel prize. The significance of his theorem was initially underappreciated, especially among mathematicians, as being barely a super-simple application of a fixed point theorem. Its power did not lie in mathematical complexity but in universality of applications. It has opened a whole new branch of mathematics studying strategic interactions between players, **game theory**, with extremely powerful result: the equilibrium exists in a very general setting!

Below we discuss a simple environment of normal form games to illustrate the concepts defined above.

Example 8.1.1 (normal form game)**.** *A **Normal Form Game** is characterized by three elements: players, actions and payoffs. There typically is one stage of the game in which all agents choose their strategies simultaneously. Normal form games are conveniently characterized in the form of a payoff-table. Table 3 presents an exemplary payoff table.*

		Player B	
		left	right
Player A	top	1, 2	0, 1
	bottom	2, 1	1, 0

Table 3: A payoff table of a normal form game.

In this example, there are two players: A and B. Player A can choose an action from set $\{top, bottom\}$*, and player B can choose an action from set* $\{left, right\}$*. Depending on the choices of both players, the matrix depicts what payoff each of them receives. A standard convention is to write payoffs of player A as the first coordinate (marked as red above), and player B's payoff as the second coordinate (blue above). For example, if player A plays top and player B plays* $left$ *then A receives 1 and B receives 2.*

Here, the set of pure strategies coincides with the set of actions for every player. This is because there is only one information set in the game, as players choose only one action and act simultaneously.

We can find a Nash equilibrium in pure strategies (if exists) directly from defining a Best Response for each player, and looking for a strategy profile in which the strategy of player A is a best response for the strategy of player B and vice versa.

An example of a mixed strategy of player A would be to play top with probability 1/2 and play bottom with probability 1/2.

For the purpose of our class, Nash theorem is pretty restrictive by requiring finite sets of strategies. We want to use game theory to study market interactions among firms, who likely have infinite choice sets: think of optimal quantities or prices.

Theorem 8.2 (existence of Nash Equilibrium (Debreu, Glicksberg, Fan))**.** *Suppose the set of all strategies Σ_i is compact and convex for every player. Moreover, assume that (i) $u_i(\sigma_i, \sigma_{-i})$ is continuous in σ_{-i}, and (ii) $u_i(\sigma_i, \sigma_{-i})$ is continuous and quasi-concave in σ^i. Then, every game has a* ***pure*** *strategy Nash equilibrium.*

This existence theorem opens up the path of studying market interactions. Before we discuss that, we finish this section with an even more general result.

Theorem 8.3 (existence of Nash Equilibrium (Glicksberg))**.** *Suppose the set of all strategies Σ_i is non-empty and compact metric space for every i, and that $u_i(\sigma_i, \sigma_{-i})$ is continuous in σ. Then, every game has a* ***mixed*** *strategy Nash equilibrium.*

8.2 Strategic Market Interactions

In absence of perfect competition, the optimal behavior of firms on the market depends on actions of their competitors. We describe this using language of non-cooperative game theory. In this section, we study **oligopoly**—a form of market structure, in which there are a few firms—players—who maximize their profits given consumer demand and actions of competitors. Importantly, all firms produce the same homogeneous good and consumers do not care which firm produced the very good they consume. We focus exclusively on pure strategies, though the set of all possible strategies of a firm is typically uncountable.

We begin with defining the notation and model primitives. There are N firms on the market indexed by $i \in \{1, .., N\}$. Denote the i's output level by y_i, and the vector of output levels by the remaining firms by y_{-i}. The market quantity is a sum of firm outputs:

$$y = \sum_{i=1}^{N} y_i$$

The production technology is described by cost functions $c_i(y_i)$, continuous and convex in y_i.

Firms face the **inverse consumer demand** $p(y)$ denoting market price p at which aggregate quantity y is demanded. We assume $p(\cdot)$ is continuous, concave, and strictly decreasing. With these assumptions, the **demand function** $y(p)$ denoting the aggregate quantity y demanded at market price p is well defined, continuous, concave and strictly decreasing. Depending on the nature of market interactions, it will be convenient for us to use the inverse demand function or demand function.

Every firm i maximizes its profits given by $py_i - c_i(y_i)$, where p is the market price of the good. Note that the assumptions made on the model primitives ensure existence of a Nash equilibrium, i.e. the optimal quantity y_i given optimal quantities of competitors y_{-i}.

There are many ways in which firms may interact with each other on a market as described above. We consider a few of them.

Definition 8.2.1 (Cournot model: simultaneous competition in quantities)**.** *In the* ***simultaneous quantity*** *(Cournot) model, firms compete by simultaneously choosing quantity. Given the (pure) strategies of competitors, $y_{-i} = (y_1, .., y_{i-1}, y_{i+1}, ..., y_N)$, firm i solves:*

$$\max_{y_i \geq 0} p(y_i, y_{-i})y_i - c_i(y_i)$$

for every $i \in \{1, ..., N\}$.

Recall $y(p)$ is the demand function, showing quantity demanded at a given price. The **individual demand** $y_i(p_i, p_{-i})$ returns quantity that i can sell if sets price at p_i while other firms set their prices at $p_{-i} = (p_1, .., p_{i-1}, p_{i+1}, ..., _N)$. Since the good is homogeneous and consumers do not care about producer identities, the demand can be written as:

$$y_i(p_i, p_{-i}) = \begin{cases} \frac{1}{M} y(p_i) & \text{if } p_i = \min\limits_{j \in \{1,..,N\}} p_j \text{ and } \#\{\arg\min\limits_{j \in \{1,..,N\}} p_j\} = M \\ 0 & \text{otherwise} \end{cases}$$

Definition 8.2.2 (Bertrand model: simultaneous competition in prices)**.** *In the* ***simultaneous price*** *(Bertrand) model, firms compete by choosing price. Given the (pure) strategies of competitiors, $p_{-i} = (p^1, .., p^{i-1}, p^{i+1}, ..., p^N)$, firm i solves:*

$$\max_{p_i \geq 0} y_i(p_i, p_{-i})p_i - c_i\big(y_i(p_i, p_{-i})\big)$$

for every $i \in \{1, ..., N\}$.

Now, we move to studying more complicated forms of strategic interactions. For notational simplicity, we only consider two firms on the market, noting that each of these settings easily accommodates more firms on the market.

The Cournot and Bertrand models assume equal position of every firm (up to their cost functions). It can be that some firms are stronger (e.g. more experiences), and some weaker. We model this by distinguishing a **leader** (WLOG: firm 1) and a **follower** (firm 2).

Definition 8.2.3 (Stackelberg model: quantity leadership)**.** *In the **quantity leadership** (Stackelberg) model, firms compete by choosing quantity. In the first stage, the dominant firm chooses their quantity. In the second stage, the follower responds with choosing their quantity.*

We write down the model in the reverse order. Taking y_1 as given, the follower chooses:

$$\max_{y_2 \geq 0} p(y_1, y_2)y_2 - c_2(y_2)$$

which gives rise to 2's best response to leader's strategy y_1:

$$y_2^* = y_2^*(y_1)$$

The leader, knowing the best response of the follower, solves:

$$\max_{y_1 \geq 0} p\big(y_1, y_2^*(y_1)\big)y_1 - c_1(y_1)$$

As an alternative, the leader may decide about market price.

Definition 8.2.4 (price leadership)**.** *In the **price leadership model**, the follower chooses their quantity given price level p set by the leader:*

$$\max_{y_2 \geq 0} py_2 - c_2(y_2)$$

which gives rise to 2's best response to leader's strategy p:

$$y_2^* = y_2^*(p)$$

*Define the **residual demand** of the leader: $y_1(p) = y(p) - y_2^*(p)$. The leader chooses price:*

$$\max_{p \geq 0} py_1(p) - c_1\big(y_1(p)\big)$$

Lastly, two firms may decide to collude, i.e. act as if they were a single entity.

Definition 8.2.5 (collusion)**.** *In the **collusion** model, the firms maximize joint profit:*

$$\max_{y_1, y_2 \geq 0} p\big(y_1 + y_2\big)(y_1 + y_2) - c_1(y_1) - c_2(y_2)$$

We have derived a few models of market interactions in oligopoly markets. They are very simplistic and refer more to the idea of market organization rather than its exact description. Which one of them we should expect in real world depends on the type of good and institutional environment. Following Tremblay & Tremblay (2011), *for example, farmers set quantities at local farmers' markets, while restaurants set prices.*

One more question arises: which one of the models is *preferred*?

8.3 Strategic Interactions and Welfare

The answer of which structure of market interaction may be preferred is very difficult. The game accommodates (possibly heterogeneous) firms and consumers, whose interests may go in opposite directions. In this section, we fix a market environment and derive the equilibrium outcomes across various models. The ultimate goal is to compare and contrasts different dimensions of *welfare* to take a stand on which market structure may or may not be *preferred.*

In what follows, there are two symmetric firms on the market, sharing the same constant marginal cost $c > 0$. The inverse market demand function is linear: $p(y) = a - by$, so the market demand function is also linear $y(p) = \frac{a-p}{b}$. We assume that $a > c$ and $b > 0$.

We are interested in the following metrics of market efficiency and/or welfare: **market price** p, **market quantity** y, consumer surplus, producer surplus, total surplus.

1. **consumer surplus** (CS) is defined as the area under demand and above the market price. Since the demand is linear, this area is defined by a triangle with surface equal to:

$$CS = \frac{1}{2}\big(p(0) - p\big)y = \frac{1}{2}(a-p)y$$

2. **producer surplus** (PS) is the sum of profits across firms:

$$PS = (p-c)y = \frac{(p-c)(a-p)}{b}$$

3. **total surplus** (TS) is the consumer and firm surplus combined:

$$TS = CS + PS$$

In studying welfare, it is useful to compare various market outcomes with the competitive benchmark. We follow this principle here. Note that with two firms the competitive market assumptions sounds a bit of a stretch assumption. We could think of it as accommodating two symmetric groups of *many* firms facing the same demand, if it makes us more comfortable.

Example 8.3.1 (competitive market)**.** *If the market is competitive, the market price equals to the marginal cost:* $p^* = c$*. Therefore, the market outcomes are:*

1. ***price****. By assumption,* $p^* = c$*.*
2. ***quantity****. Using the inverse demand,* $p^* = a - by^* = c$*, so* $y^* = \frac{a-c}{b}$*.*
3. ***consumer surplus****.* $CS = \frac{1}{2}(a-p)y = \frac{1}{2}\frac{(a-c)^2}{b}$*.*
4. ***producer surplus****.* $PS = 0$*.*
5. ***total surplus****.* $TS = \frac{1}{2}\frac{(a-c)^2}{b}$*.*

Example 8.3.2 (Cournot oligopoly)**.** *In Cournot oligopoly, both firms choose their quantities at the same time. Firm* i*,* $i \in \{1, 2\}$*, maximizes its profits given competitor's action* y_{-i}*:*

$$\max_{y_i \geq 0} p(y_i + y_{-i})y_i - c_i(y_i)$$

The first order conditions imply that:

$$\begin{aligned} \big(-b\big)y_i + \big(a - b(y_i + y_{-i}) - c\big) &= 0 \\ y_i^*(y_{-i}) &= \frac{a - by_{-i} - c}{2b} = \frac{a-c}{2b} - \frac{1}{2}y_{-i} \end{aligned}$$

Which gives rise to the best response functions:

$$y_1^*(y_2) = \frac{a-c}{2b} - \frac{1}{2}y_2, \; y_2^*(y_1) = \frac{a-c}{2b} - \frac{1}{2}y_1$$

In equilibrium:

$$\begin{aligned} y_1^* &= BR_1(y_2^*) \\ y_2^* &= BR_2(y_1^*) \end{aligned}$$

so:

$$y_1^* = \frac{a-c}{2b} - \frac{1}{2}y_2^* = \frac{a-c}{2b} - \frac{1}{2}\Big(\frac{a-c}{2b} - \frac{1}{2}y_1\Big) = \frac{a-c}{4b} + \frac{1}{4}y_1^*$$

and finally:

$$y_1^* = \frac{a-c}{4b} + \frac{1}{4}y_1^* \Rightarrow \frac{3}{4}y_1^* = \frac{a-c}{4b} \Rightarrow y_1^* = \frac{1}{3}\frac{a-c}{b}$$

Since firms are symmetric, $y_2^* = y_1^*$*, so* $y^* = \frac{2}{3}\frac{a-c}{b}$*.*

We are ready to calculate market outcomes under the Cournot duopoly:

1. ***price***. *Using the inverse demand:* $p^* = a - by = \frac{a+2c}{3}$
2. ***quantity***. *As derived above,* $y^* = \frac{2}{3}\frac{a-c}{b}$.
3. ***consumer surplus***. $CS = \frac{1}{2}(a-p)y = \frac{2}{9}\frac{(a-c)^2}{b}$.
4. ***producer surplus***. $PS = \frac{(p-c)(a-p)}{b} = \frac{2}{9}\frac{(a-c)^2}{b}$.
5. ***total surplus***. $TS = \frac{4}{9}\frac{(a-c)^2}{b}$.

Example 8.3.3 (Bertrand oligopoly)**.** *In Bertrand oligopoly, both firms choose their price at the same time. The market demand function* $y = \frac{a-p}{b}$ *denotes the quantity given prices. We need to define the division of output sold among firms. We assume even division, that is the individual demand function is given by:*

$$y_i = y(p_i, p_{-i}) = \begin{cases} \frac{a-p}{b} & \text{if} \quad p_i < p_{-i} \\ \frac{1}{2}\frac{a-p}{b} & \text{if} \quad p_i = p_{-i} \\ 0 & \text{if} \quad p_i > p_{-i} \end{cases}$$

Firm i*,* $i \in \{1,2\}$*, maximizes its profits given competitor's action* p_{-i}*:*

$$\max_{p_i \geq 0} \; (p_i - c)y_i(p_i, p_{-i})$$

We are looking for a Best Response of i *to competitor's choice of* p*:* $BR_i(p_{-i}) = p_i^*(p_{-i})$.

Now, note that no firm would prefer to set their price higher than the competitor because it yields zero profits. Suppose $p_i > c$*. Then* $-i$ *would prefer to set price slightly lower than* p_i *but still higher than* c*. In such a way,* $-i$ *can get the whole market for itself. Then, firm* i *would have incentive to decrease price slightly below* p_{-i}*, for the same reason. Such a game of alternating price cuts would continue as long as prices are above the marginal cost. Otherwise, the profits would become negative and firms would be better off to let their competitor fulfill the demand. Hence, the only Nash equilibrium of this game is to set prices equal to marginal cost. The same as in competitive equilibrium! Figure 5 illustrates this reasoning. The Best Responses of both players cross once at point* (c, c)*.*

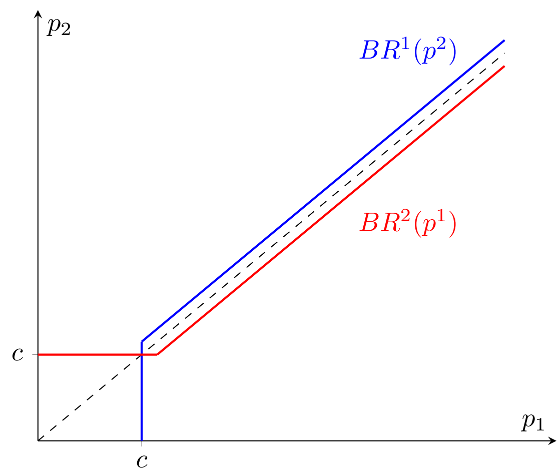

Figure 5: Best Responses in Bertrand model.

The welfare indicators for Bertrand model are the same as for the competitive market.

Bertrand equilibrium yields somewhat striking result that despite having only two firms on the market, it is possible to achieve the competitive outcomes. However, note that this requires firms to be symmetric. If one firm had a strictly smaller marginal cost than the other, they would be able to capture the whole market and set strictly positive markup. However, from a formal point of view, there would be no equilibrium in such a game.

Exercise 8.3.1. *Consider the market environment of this section. Show that if* $c_1 < c_2$*, there is no equilibrium in Bertrand game.*

Example 8.3.4 (Stackelberg oligopoly)**.** *We begin with the follower's problem. The follower takes y_1 as given and maximizes its profit:*

$$\max_{y_2 \geq 0} p(y_1 + y_2) \cdot y_2 - c_2(y_2) = \max_{y_2 \geq 0} \big(a - b(y_1 + y_2) - c\big) \cdot y_2$$

First order condition implies:

$$\big(-b\big)y_2 + \big(a - b(y_1 + y_2) - c\big) = 0$$
$$y_2^*(y_1) = \frac{a-c}{2b} - \frac{1}{2}y_1$$

*The leader **internalizes** $y_2^*(y_1) = \frac{a-c}{2b} - \frac{1}{2}y_1$ in their optimization problem:*

$$\max_{y_1 \geq 0} p(y_1 + y_2^*(y_1)) \cdot y_1 - c_1(y_1) = \max_{y_1 \geq 0} \big(a - b(y_1 + \frac{a-c}{2b} - \frac{1}{2}y_1) - c\big) \cdot y_1 =$$
$$\max_{y_1 \geq 0} \frac{1}{2}\big(a - c - by_1\big) \cdot y_1$$

First order condition implies:

$$\frac{1}{2}\big(-b\big)y_1 + \frac{1}{2}\big(a - c - by_1\big) = 0$$
$$y_1^* = \frac{a-c}{2b}$$
$$y_2^*(y_1^*) = \frac{a-c}{2b} - \frac{1}{2}y_1^* = \frac{a-c}{2b} - \frac{1}{2}\frac{a-c}{2b} = \frac{a-c}{4b}$$

Thus, the market quantity is $y^* = y_1^* + y_2^* = \frac{3}{4}\frac{a-c}{b}$. The remaining market outcomes under Stackelberg game are:

1. ***price**. Using the inverse demand: $p^* = a - by = \frac{a+3c}{4}$*
2. ***quantity**. As derived above, $y^* = \frac{3}{4}\frac{a-c}{b}$.*
3. ***consumer surplus**. $CS = \frac{1}{2}(a-p)y = \frac{9}{32}\frac{(a-c)^2}{b}$.*
4. ***producer surplus**. $PS = \frac{(p-c)(a-p)}{b} = \frac{3}{16}\frac{(a-c)^2}{b}$.*
5. ***total surplus**. $TS = \frac{15}{32}\frac{(a-c)^2}{b}$.*

The next model in the line is the price leadership. However, we skip it because it is not really well defined under our assumptions (the key is the same constant marginal cost for every firm).

Exercise 8.3.2. *Show that given assumptions of this section, the price leadership problem reduces to the Bertrand competition.*

Example 8.3.5 (collusive oligopoly)**.** *Both firms form a **cartel** and maximize joint profit. The cartel's joint maximization objective is:*

$$\max_{y_1, y_2 \geq 0} p(y_1 + y_2) \cdot (y_1 + y_2) - c(y_1 + y_2) = \max_{y_1, y_2 \geq 0} \big(a - b(y_1 + y_2) - c\big)(y_1 + y_2)$$

The cartel's objective is a function of $y_1 + y_2$, so any combination of non-negative y_1 and y_2 summing up to the equilibrium market quantity would work as an equilibrium. As firms are symmetric, we need to assume market shares. For example, assume $y_1^* = y_2^*$ noting that it does not effect our welfare calculations. First order condition implies:

$$-b(y_1 + y_2) + \big(a - b(y_1 + y_2) - c\big) = 0 \Rightarrow a - c = 2b(y_1 + y_2) \Rightarrow y_1^* + y_2^* = \frac{a-c}{2b}$$

Thus, the market quantity is $y^* = \frac{a-c}{2b}$, and the remaining outcomes are:

1. ***price**. Using the inverse demand: $p^* = a - by = \frac{a+c}{2}$*
2. ***quantity**. As derived above, $y^* = \frac{1}{2}\frac{a-c}{b}$.*

3. ***consumer surplus***. $CS = \frac{1}{2}(a-p)y = \frac{1}{8}\frac{(a-c)^2}{b}$.
4. ***producer surplus***. $PS = \frac{(p-c)(a-p)}{b} = \frac{1}{4}\frac{(a-c)^2}{b}$.
5. ***total surplus***. $TS = \frac{3}{8}\frac{(a-c)^2}{b}$.

We are ready to collect the results. It turns out that there exists a clear ordering of structures of market interactions with respect to equilibrium price and quantity.

Proposition 8.3.1. *The following holds:*

1. ***price*** *ranking.* $p^{\text{competitive}} = p^{\text{Bertrand}} < p^{\text{Stackelberg}} < p^{\text{Cournot}} < p^{\text{collusive}}$
2. ***quantity*** *ranking.* $y^{\text{competitive}} = y^{\text{Bertrand}} > y^{\text{Stackelberg}} > y^{\text{Cournot}} > y^{\text{collusive}}$

Moreover, the same ranking is preserved by the surplus!

Proposition 8.3.2. *The following holds: where all values are provided in* $\frac{(a-c)^2}{b}$ *units.*

	CS	PS	TS
competitive/Bertrand	0.5	0	0.5
Stackelberg	0.28	0.19	0.47
Cournot	0.22	0.22	0.44
collusion	0.13	0.25	0.38

Table 4: equilibrium surplus division in oligopoly models.

Not surprisingly, consumer's and producer's rankings of favorite market structures are opposite. However, note that losses to the consumer surplus are higher as we are going down table 4 than respective increases in producer surplus. This is because the ranking of equilibrium prices and quantities. Going down table 4, the quantity goes down and the price goes up—the consumer surplus gets hurt twice. In turn, for firms the losses resulting from lower quantity are more than offset by the increase in price. Therefore, from the **efficiency** point of view, the best outcome is the competitive/Bertrand equilibrium. It maximizes the dynamism of economy—in this case, ensuring the highest production as well as total surplus.

From the **social** point of view, the situation may become more complicated, as depending on how we want to weight the consumer and producer surplus, we may receive different ranking. Nevertheless, typically it is the consumer surplus which receives more weight, and hence from the social perspective we would most likely also choose the competitive/Bertrand equilibrium.

9 Antitrust

Last week, we learned that some market structures are *preferred* than others through the lenses of both efficiency and equity. Today, we focus at how the government can regulate markets—encourage desired structures and deter the undesired—to improve outcomes. We discuss the main aspects of antitrust policy. We take a deeper look at mergers, looking for guidance on what should be the conditions under which the authorities should accept or deny a merger. We conclude with some practical considerations regarding market concentration and its relevance for antitrust recommendations.

List of contents. Antitrust. Market concentration. Herfindahl-Hirschmann Index (HHI). Merger. Horizontal merger. Vertical merger. Cost synergies. CS-neutral (-increasing, -decreasing) mergers.

Literature References.

[1] Farrell, J., and Shapiro, C. (1990). Horizontal mergers: an equilibrium analysis. The American Economic Review, 107-126.

[2] Nocke, V., and Whinston, M. D. (2010). Dynamic merger review. Journal of Political Economy, 118(6), 1200-1251.

[3] Nocke, V., and Whinston, M. D. (2022). Concentration thresholds for horizontal mergers. American Economic Review, 112(6), 1915-1948.

[4] Miller, Nathan, and others (2022). On the misuse of regressions of price on the HHI in merger review. Journal of Antitrust Enforcement, 10(2), 248-259.

9.1 Mathematical Glossary

Theorem 9.1. *Under regularity conditions, the derivative of a integral with bounds is given by:*

$$\frac{\mathrm{d}}{\mathrm{d}x}\int_{g(x)}^{f(x)} h(t)\mathrm{d}t = h\big(f(x)\big)f'(x) - h\big(g(x)\big)g'(x)$$

9.2 Antitrust and Regulation

Previous sections suggest plenitude of different market outcomes depending on the nature of strategic interactions among firms. Moreover, some of these outcomes are more preferred than others. Most countries have developed institutions that aim to facilitate achieving more desired market outcomes. **Antitrust** laws are regulations that promote economic competition by limiting the market power of any particular firm. Certain public authorities, including Federal Trade Commission (FTC) in the US, or Australian Competition and Consumer Commission (ACCC), are responsible for enforcing such legislation.

The task of antitrust authorities in regulating anti-competitive firm conduct is not easy. Firms realize not only gains from anti-competitive behavior but also legal risks associated with it and pay particular attention to either conceal their wrongdoings or present them falsely as pro-competitive. In this section, we review and justify a few facts related to antitrust policy.

Fact 1. *Despite seemingly appealing properties, antitrust policy rarely focuses on enforcing price competition.*

As derived in the previous section, the efficient outcome can be implemented in Bertrand equilibrium. However, in reality it would be hard to properly define and enforce Bertrand-style of competition. Moreover, even if it was technically possible, it may still be undesired. Note that efficiency in Bertrand economy relies on the assumption of symmetric marginal costs. Once we deviate from this assumption, the market can quickly get monopolized by a firm with the lowest cost.

Fact 2. *Collusion by openly coordinating prices or quantities is legally prohibited and punishable under the law.*

Even though the most desirable outcome is rarely targeted, a different story can be told about the least desirable outcome. Collusion is particularly bad from the perspective of consumer surplus and efficiency. In most of the countries, collusion has criminal consequences. Special authorities are delegated to track and reveal collusive behavior among firms.

Fact 3. *The more firms, the better: antitrust policy focuses on limiting barriers to entry.*

Consider a Cournot model with N firms, each of them facing constant (but possibly different) marginal cost c_i. Same as before, we assume linear demand $p(y) = a - by$. Firm i's profit maximization problem is:

$$\max_{y_i \geq 0} \Big(a - b\sum_{k=1}^{N} y_k - c_i\Big)y_i$$

Denote the average marginal cost in the economy by $\overline{c} = \frac{1}{N}\sum_{i=1}^{N} c_i$.

Intuitively, with *many* firms on the market, we should be getting close to the competitive outcome. This presumption is true.

Proposition 9.2.1. *Let $y(N)$ be the Cournot equilibrium as a function of the number of firms on the market. Then:*

$$\lim_{N\to\infty} y(N) = \frac{a - \overline{c}}{b}$$

Proof. The first order condition of firm's problem implies that $a - b\sum_{k=1}^{N} y_k - c_i - by_i = 0$. Transforming this equation:

$$y_i = \frac{a - c_i}{2b} - \frac{1}{2}\sum_{j\neq i} y_j \Leftrightarrow$$

$$\frac{1}{2}y_i = \frac{a - c_i}{2b} - \frac{1}{2}\sum_{i=1}^{N} y_j \Leftrightarrow$$

$$y_i = \frac{a - c_i}{b} - y$$

where the first equivalence comes from adding $\frac{1}{2}y_i$ to both sides of the equation, and the second follows because $y = \sum_{i=1}^{N} y_i$. Now, adding up the last equations across i yields:

$$\sum_{i=1}^{N} y_i = \sum_{i=1}^{N} \left(\frac{a-c_i}{b} - y\right) \Leftrightarrow$$
$$y = \sum_{i=1}^{N} \frac{a-c_i}{b} - Ny \Leftrightarrow$$
$$y = \frac{1}{N+1}\sum_{i=1}^{N}\frac{a-c_i}{b} = \frac{1}{N+1}\frac{Na}{b} - \frac{1}{N+1}\frac{1}{b}\sum_{i=1}^{N} c_i = \frac{N}{N+1}\frac{a}{b} - \frac{1}{N+1}\frac{1}{b}N\frac{1}{N}\sum_{i=1}^{N} c_i \Leftrightarrow$$
$$y = \frac{N}{N+1}\frac{a-\bar{c}}{b}$$

Taking the appropriate limit yields the desired result. □

Fact 4. *The less concentration, the better: antitrust policy actively counteracts emergence of market leaders with market share significantly higher than others.*

When a few firms dominate the market in terms of their market share, they can gain significant market power enabling them to set higher prices, reduce output, or decrease the quality of goods and services. This is clearly an undesired outcome.

Definition 9.2.1 (market concentration)**.** ***Market concentration*** *refers to the extent to which a small number of firms dominate total sales, production, or capacity in a particular market. It reflects the degree of competition within a market.*

Market concentration is often measured by the so called HHI index.

Definition 9.2.2 (Herfindahl-Hirschman index of market concentration)**.** *The* ***Herfindahl-Hirschman*** *index of market concentration (HHI) is the sum of squared market shares:*

$$HHI = \sum_{i=1}^{N} s_i^2$$

where $s_i = \frac{y_i}{y}$

Exercise 9.2.1. *Show that* $HHI \in (0,1]$.

HHI is one of the most widely used concentration indicators, but not the only one. We can also consider, for example, the total share of M largest firms, or the ratio of 90th to 10th percentile of market shares.

Antitrust' care for market concentration brings us to one of the most important aspects of its activity: mergers.

Fact 5. *Antitrust oversees mergers, allowing or disabling their execution.*

Mergers gain particular interest, among others, because the joint profit of two competitors may be higher if both raise the price, but neither would do so unilaterally because otherwise they would just lose sales to the competitor.

9.3 Reviewing Mergers

Definition 9.3.1 (mergers)**.** *A* ***merger*** *is a business strategy where two or more companies combine to form a single entity.*

Definition 9.3.2 (types of mergers)**.** *A merger can be* ***horizontal*** *if it occurs between companies operating in the same industry and at the same stage of production competing against each other, or* ***vertical*** *if it involves companies at different stages of the same production process.*

In our class, we focus on horizontal mergers.

Horizontal mergers are directly related to market concentration, as the shares of (at least) two firms would add and work as share of a single firm, inevitably leading to increase in market concentration. Therefore, candidate firms for a merger need to convince an antitrust authority about other benefits that the merger would bring on the cost of increased concentration. A typical gain comes from potential cost **synergies**, which occur when the cost structure of merged firms allows for more efficient production compared to the sum of pre-merger, separate productions of firms that intend to merge. Such cost synergies may result in increase in total welfare and consumer surplus as a result of a merger.

The post-merger cost synergies are difficult to measure pre-merger. This is why typical guidelines for antitrust policy still suggest concentrating on measures of concentration in merger proposal reviews. Both levels of the initial HHI and predicted change in industry concentration are taken into account. Effectively, guidelines rely on thresholds in terms of HHI and ΔHHI for which the merger is either accepted, closely reviewed, or rejected (for further dispute in a court of law). Mergers that do not increase significantly the market concentration as well as keeping it low after mergers are most likely to succeed.

So the merger review traditionally relies on pre-merger HHI index. Modern research shows this is a significant short-coming, because a merger will have equilibrium effects, for (at least) two reasons. First, using only the pre-merger concentrations, if the post-merger share of the new firm must be the sum of pre-merger shares of its members, say $s_1 + s_2$, so the increase in HHI is given as $(s_1 + s_2)^2 - (s_1^2 + s_2^2) = 2s_1 s_2$. But this relies on an assumption that all other firms maintain their pre-merger shares, and hence also the output. If the merger is profitable for merging firms, there is no reason to prevent it from happening. The problem lies in the fact that this formulation does not take into account potential changes in post-merger market shares among non-merging companies in response to the new market conditions. Second (and even more serious), merger evaluations based only on pre-merger shares assume that there is a universal causal relation between market concentration and market performance, or welfare. This does not have to the case.

Proposition 9.3.1 (Farrell & Shapiro (1990))**.** *In Cournot model (defined as in the previous section), welfare rises with a small change in firms' outputs if and only if:*

$$\frac{\mathrm{d}y}{y} + \frac{1}{2}\frac{\mathrm{d}H}{H} > 0$$

where $H = \sum_{i=1}^{N} \left(\frac{y_i}{y}\right)^2$ *is the HHI index.*

Proof. Define welfare as a sum of consumer surplus and firm surplus:

$$W = CS + PS = \int_0^y p(x)\mathrm{d}x - py + \sum_{i=1}^{N} py_i - c_i y_i$$

and consider a small change in firms' output levels dy_i:

$$\mathrm{d}W = \sum_{i=1}^{N} p\mathrm{d}y_i - c_i \mathrm{d}y_i = \sum_{i=1}^{N} (p - c_i)\mathrm{d}y_i$$

In turn, the first order condition implies that $p - c_i = -p'(y)y_i$, so:

$$\mathrm{d}W = -p'(y)\sum_{i=1}^{N} y_i \mathrm{d}y_i$$

Now, note that:

$$\begin{aligned}\sum_{i=1}^{N} y_i \mathrm{d}y_i =& \frac{1}{2}\mathrm{d}\Big[\sum_{i=1}^{N} y_i^2\Big] = \frac{1}{2}\mathrm{d}\Big[y^2 \sum_{i=1}^{N}\Big(\frac{y_i}{y}\Big)^2\Big] = \frac{1}{2}\mathrm{d}\Big[y^2 H\Big] \\ =& \frac{1}{2}\Big[2yH\mathrm{d}y + y^2\mathrm{d}H\Big] = y^2 H\Big[\frac{\mathrm{d}y}{y} + \frac{1}{2}\frac{\mathrm{d}H}{H}\Big]\end{aligned}$$

Plugging back to the equation above:

$$\mathrm{d}W = -p'(y)\sum_{i=1}^{N} y_i \mathrm{d}y_i = -p'(y)y^2 H\Big[\frac{\mathrm{d}y}{y} + \frac{1}{2}\frac{\mathrm{d}H}{H}\Big]$$

$-p'(y) > 0$ by assumption. By construction $y > 0$ and $H > 0$. The result follows. □

Intuitively, increases in output ($\mathrm{d}y > 0$) tend to raise welfare because in Cournot model the price exceeds marginal costs in equilibrium. Suppose a firm with a large market share increases its output. In this case, H, y, and W will all rise.

Moreover, the equation above indicates that even $\mathrm{d}y < 0$ and $\mathrm{d}H > 0$ may be associated with $\mathrm{d}W > 0$. This is because the larger firms have smaller marginal costs in Cournot oligopoly. Reallocating production from small inefficient firms towards larger and more efficient firms may both increase concentration and welfare.

The reliance of antitrust guidelines on HHI is largely a-theoretical, somewhat arbitrary, and has significant path-dependence. Modern **regulation economics** is looking for a set of better tools to distinguish for whether to allow or disallow mergers. We want these tools to be deeply rooted in the economic theory to be capable of accounting for equilibrium effects. The remainder of this section is inspired by Nocke & Whinston (2010) and Nocke & Whinston (2022). Derivations are based on the general Cournot model with heterogeneous constant marginal costs, as introduced in the previous section. Suppose a subset of M from among N firms wants to merge.

In what follows, we will assume that the parameters of the model are such that all considered firms are active, that is, produce non-zero quantities of good and generate strictly positive profits. It is straightforward to generalize the results to situations involving non-active firms.

Proposition 9.3.2. *A merger increases consumer surplus if and only if it induces an increase in industry output.*

Proof. Recall from the previous class that the consumer surplus at market quantity y is given by:

$$CS(y) = \int_0^y \big(p(x) - p(y)\big)\mathrm{d}x$$

Therefore:

$$\frac{\mathrm{d}}{\mathrm{d}y}CS(y) = y(-p'(y)) - 0 \cdot p'(0) = -yp'(y) > 0$$

since p is downward sloping. Therefore, the consumer surplus increases if and only if the aggregate quantity increases. The result follows. □

Definition 9.3.3 (CS-neutral merger)**.** *We say that a merger is* ***CS-neutral*** *if consumer surplus following the merger is the same as before.*

Proposition 9.3.3. *A merger is CS-neutral if and only if it does not lead to changes in aggregate output.*

Proof. This is a direct corollary of proposition 9.3.2. □

Proposition 9.3.4. *Suppose a merger is CS-neutral. Then:*

1. *it causes no changes in the output of any non-merging firm or in the total output of the merging firms,*
2. *the merged firm's post-merger margin equals the sum of the merging firms' pre-merger margins:*

$$P(Y) - \bar{c}_M = \sum_{i \in M} \big(p(y) - c_i\big)$$

3. *the merger is profitable: it strictly increases the joint profit of merging firms.*

Proof.

1. Since the merger is CS-neutral, by proposition 9.3.3 the aggregate output level remains unchanged. The first order condition for a non-merging firm remains the same as it was before merger:

$$p(y) + p'(y)y_i = c_i,\ i \neq M$$

Note that for every level of y there is a unique y_i given firm's i marginal cost c_i. Therefore, it has to remain the same pre- and post- merger. Since this is the case, the total output of merging firms must remain the same as well.

2. The post-merger optimality condition is:

$$p(y) + p'(y)\sum_{i \in M} y_i - \bar{c}_M = 0$$

Now, sum the for pre-merger first order conditions across merging firms:

$$\sum_{i \in M} \big(p(y) + y_i p'(y) - c_i\big) = 0 \Rightarrow \sum_{i \in M} \big(p(y) - c_i\big) + p'(y)\sum_{i \in M} y_i = 0$$

Combing the two equations above yields:

$$p(y) - \bar{c}_M = \sum_{i \in M} \big(p(y) - c_i\big)$$

3. The merged firm's profit after merger:

$$\begin{aligned}
p(y) - \bar{c}_M \underset{\text{prop. 9.3.4, point 2}}{=} & \Big(\sum_{j\in M} p(y) - c_j\Big) \sum_{i\in M} y_i \\
&= \sum_{j\in M} \Big(p(y) - c_j\Big) \sum_{i\in M} y_i \\
&= \sum_{j\in M} \sum_{i\in M} \Big(p(y) - c_j\Big) y_i \\
&= \sum_{i\in M} (p(y) - c_i) y_i + \sum_{j\neq i} \sum_{i\in M} \Big(p(y) - c_j\Big) y_i \\
&> \sum_{i\in M} (p(y) - c_i) y_i
\end{aligned}$$

which is the sum of pre-merger profits among M firms.

□

Proposition 9.3.5. *If a merger is CS-neutral, it must reduce marginal cost below the marginal cost of the most efficient merger partner.*

Proof. Without loss, assume that firm 1 has the lowest pre-merger marginal cost: $\min\limits_{k\in M} \{c_k\} = c_1$. By point 2 of proposition 9.3.4, if the merger is CS-neutral then:

$$p(y) - \bar{c}_M = \sum_{i\in M} \big(p(y) - c_i\big)$$

This is equivalent to the following statement:

$$c_1 - \bar{c}_M = \sum_{i>1} \big(p(y) - c_i\big) > 0$$

So it must be that $\min\limits_{k\in M} \{c_k\} = c_1 > \bar{c}_M$. □

Proposition 9.3.6. *Let $\hat{c}_M(Y) = p(y) - \sum_{i\in M} \big(p(y) - c_i\big)$. The merger is :*

1. *CS-neutral if $\bar{c}_M = \hat{c}_M(y)$*
2. *CS-increasing if $\bar{c}_M < \hat{c}_M(y)$*
3. *CS-decreasing if $\bar{c}_M > \hat{c}_M(y)$*

Proof. Use proposition 9.3.4 and the fact that the post-merger output $\bar{y}$ is decreasing in merged firm's marginal cost. □

Proposition 9.3.6 defines the cost synergy levels that are necessary for merger to be CS-nondecreasing, and hence acceptable for an agency concerned with maximizing consumer surplus. The merger should be approved if $\bar{c}_M \leq \hat{c}_M(y)$ and rejected otherwise. The threshold is strictly increasing in the total output y. Hence, the merger policy in larger industries will be more liberal (in the absolute terms).

The cost synergy thresholds in proposition 9.3.6 are defined in terms of market price and cost structure within the set of M candidates for merger. It turns out we can go further than that and express the necessary synergies in terms of the change in HHI index (but not levels!). For simplicity, from now on suppose that only two firms want to merge: i and j.

Proposition 9.3.7. *Suppose that the pre-merger marginal costs of merging firms are equal to each other: $c_i = c_j \equiv c_M$ and that the post-merger marginal cost is $\bar{c}_M$. A merger is CS-neutral if and only if:*

$$\frac{c_M - \bar{c}_M}{c_M} = \frac{\sqrt{\frac{\Delta H}{2}}}{\epsilon - \sqrt{\frac{\Delta H}{2}}}$$

where ϵ is the pre-merger elasticity of demand and H is the pre-merger HHI.

Proof. We begin the proof with an observation that since both firms have the same pre-merger marginal costs, they also have the same pre-merger market shares: $s_i = s_j = \frac{s_M}{2}$.

The pre-merger elasticity of demand is given by:

$$\epsilon = -\frac{p(y)}{yp'(y)}$$

Now, we use the first order condition $p(y) - c_i + p'(y)y_i = 0$ to derive firm (pre-merger) market share:

$$s_i = \frac{y_i}{y} = -\frac{p(y) - c_i}{p'(y)y}$$

It turns out that the pre-merger marginal cost can be related to the demand elasticity and market share:

$$\begin{aligned} p(y)\Big(1 - \frac{s_M}{2\epsilon}\Big) =& p(y)\Big(1 - \frac{-2\frac{p(y)-c_i}{p'(y)y}}{-2\frac{p(y)}{yp'(y)}}\Big) = p(y)\Big(1 - \frac{p(y) - c_i}{p(y)}\Big) = p(y) - p(y) + c_i \\ =& c_M \end{aligned}$$

By propositions 9.3.4, the CS-neutral merger requires:

$$p(y) - \overline{c}_M = \big[p(y) - c_i\big] + \big[p(y) - c_j\big]$$

Using the fact that $c_i = c_j = c_M$:

$$\begin{aligned} p(y) - \overline{c}_M = \big[p(y) - c_i\big] + \big[p(y) - c_j\big] \Leftrightarrow& c_M - \overline{c}_M = p(y) - c_M \\ \Leftrightarrow& \frac{c_M - \overline{c}_M}{c_M} = \frac{p(y) - c_M}{c_M} \end{aligned}$$

Now, plug in the equation for c_M in terms of elasticity and share:

$$\frac{c_M - \overline{c}_M}{c_M} = \frac{p(y) - c_M}{c_M} = \frac{p(y) - p(y)\Big(1 - \frac{s_M}{2\epsilon}\Big)}{p(y)\Big(1 - \frac{s_M}{2\epsilon}\Big)} = \frac{\frac{s_M}{2\epsilon}}{1 - \frac{s_M}{2\epsilon}}$$

Since the merger is CS-neutral, shares of remaining firms remain the same. Hence:

$$\Delta H = \overline{H} - H = (s_i + s_j)^2 - (s_i^2 + s_j^2) = 2s_i s_j = \frac{s_M^2}{2}$$

so $\sqrt{2\Delta H} = s_M$ and:

$$\frac{s_M}{2\epsilon} = \frac{\sqrt{2\Delta H}}{2\epsilon} = \frac{\sqrt{\frac{\Delta H}{2}}}{\epsilon} \quad \text{and} \quad 1 - \frac{s_M}{2\epsilon} = \frac{\epsilon - \sqrt{\frac{\Delta H}{2}}}{\epsilon}$$

which brings us to the final conclusion:

$$\frac{c_M - \overline{c}_M}{c_M} = \frac{\sqrt{\frac{\Delta H}{2}}}{\epsilon - \sqrt{\frac{\Delta H}{2}}}$$

□

Let's think about the significance of proposition 9.3.7. For a given elasticity of demand, the required cost efficiencies are an increasing function of the change in HHI, and are completely independent from the level of HHI! This suggests that relying on HHI levels while deciding about whether or not to allow a merger is incorrect. Moreover, the model-based approach allows us to define very precise and well-justified thresholds for synergy efficiencies that would justify a merger. Nocke & Whinston (2022) conclude that the existing thresholds are *lax* compared to the thresholds implied by their model. Also, see their paper to study a more general proposition extending the proposition 9.3.7 to the case of asymmetric merging firms.

9.4 Concentration and Empirical Merger Review

Market concentration remains at the center of interest in antitrust. For example, for a merger review, someone may be tempted to verify empirically how market outcomes would change in a response to increase in concentration as an evidence in favor or against the merger. Data may come, for example, from markets for similar goods, or markets for the same good but placed in different time or at another geographical location. Since HHI is a leading indicator of market concentration[4], what we have in mind is a regression of a market outcome, typically: price, on HHI:

$$p_i = \beta_0 + \beta_1 HHI_i + \varepsilon_i$$

To see why such a regression is problematic, consider an example provided by Miller et al. (2022). There are two firms (1 and 2) that sell the same homogeneous product on three distinct (e.g. geographically) markets (A, B, C). The firms interact according to the Cournot model (quantity competition). We assume that the demand at each location is given by $p(y) = 10 - y$, and that the firms face constant marginal costs. However, the costs may vary across geographic locations. Figure 6 summarizes Miller's assumptions and equilibrium outcomes at every market. The last column presents the result had the two firms merged.

Table 1. The price–HHI relationship with cost variation

	Region A	Region B	Region C	Monopoly
Firm 1's costs	3.00	4.00	1.83	3.00
Firm 2's costs	3.00	3.00	3.00	
Firm 1's share (%)	50	38	62	100
Firm 2's share (%)	50	62	38	
Market price	5.33	5.67	4.94	6.50
Quantity	4.67	4.33	5.06	3.50
HHI	5000	5288	5288	10,000

Figure 6: source: Miller et al. (2022)

In every market, a merger, hence monopolization, raises price. However, the correlation between price and concentration is ambiguous, depending on which regions we compare (table 5)! This is despite the fact that the equilibrium outcomes have been obtained directly from a theoretical model that does not accommodate any randomness, and so we cannot blame sampling error or unobserved confounding factors. Different patterns of correlation between price and concentration are induced purely by differences in firms' marginal costs.

	A v. B	B v. C	A v. C
Δ price	-	+	-
Δ concentration	-	0	-
correlation	+	0	-

Table 5: equilibrium surplus division in oligopoly models.

With more data points, we could estimate the regression model as suggested above, and it would suffer the same issue. This is because there is no **causal relation** between price and concentration. As Miller et al. (2022) point out, price and HHI are both equilibrium outcomes, jointly determined in equilibrium by the primitives of the model. They are not causally related. This is a very peculiar problem. Note that it is not an **identification** problem, in which we just struggle to retrieve a well define object from an imperfect data environment (like well known problems of omitted variable bias, endogeneity, or simultaneity). The problem is that there is nothing to identify! Therefore, the regressions of price on HHI are not recommended for merger review. Fortunately, we have many other great tools (purely theoretical or structural) that allow us to gain accurate predictions about the expected post-merger price levels.

[4] the same argument extends to other measures of concentration.

This result should be of no surprise for us, given the theoretical results from the previous section that suggest the level of concentration does not have an impact on the direction of changes in CS in response to merger. However, the pre-post merger increment of HHI does! And this brings us to another example suggested by Miller et al. (2022), illustrating that HHI can be useful in assessing mergers. We keep the same demand, fix marginal cost at 3 for all firms, and consider 2-firm mergers on markets with varying number of existing firms. Figure 7 presents the results.

Table 2. Merger price effects and the implied HHI change

	Number of pre-merger competitors				
	Two	Three	Four	Five	Six
Percentage price increase	21.88	12.28	7.95	5.60	4.17
Percentage quantity decrease	25.00	11.11	6.25	4.00	2.78
Implied change in the HHI	5000	2222	1250	800	556

Figure 7: source: Miller et al. (2022)

There clearly is a positive correlation between the merger price effect and the implied change in HHI, and it may be retrieved from data.

10 Incomplete Information

In our last meeting, we study strategic interactions when players are unsure about some aspects of the game. This incomplete information framework requires players to choose actions depending on their beliefs regarding the unobserved pieces of the game. We introduce Bayesian games as language to study such situations and consider one of the most classical example of Bayesian games: auctions.

List of contents. Beliefs. Bayesian Game. Individual types. Pure strategy in Bayesian games. Payoffs in Bayesian games. Common prior. Private values. Independent and interdependent private values. Independent and correlated types. Bayes' rule. Expected payoffs. Bayesian Nash Equilibrium. Weakly dominant strategies. Ex-ante, interim, ex-post stages of a Bayesian Game. Ex-ante optimality. Interim optimality. Ex-post equilibrium. Auctions. Forward and reverse auctions. Bidders. Bidder valuations. Auctioneer. Independent private values auctions. Affiliated values auctions. Symmetric bidders. Bid. First-price sealed bid auctions. Second-price sealed-bid auctions. Dutch auctions. English auctions. Bidding strategy. Symmetric strategies. Strictly monotone strategies. Bayesian Nash equilibrium in auctions. Revenue equivalence. Auction efficiency.

Textbook references. Espinola: Ch. 15; Varian: Ch. 15; Jehle & Reny: Ch. 9

Literature References.

[1] Krishna, Vijay (2009). Auction theory. Academic press. Chapters 1-3.

10.1 Mathematical Glossary

Theorem 10.1 (derivative of the inverse function)**.** *Consider a strictly monotone unidimensional function $f : X \to Y$ and its inverse $f^{-1} : Y \to X$. The derivative of the inverse is given by:*

$$\left[f^{-1}(y)\right]' = \frac{1}{f'\big(f^{-1}(y)\big)}$$

10.2 Bayesian Games

The last time we studied game theory, we assumed that every piece of the game is known to every player, including payoffs and agents' preferences. In many setting, it may not be the case. Agents will typically try to hide some information in order to get advantage in the game. How their competitors should react? What would be best responses in a world with incomplete information?

Intuitively, if agent i does not have full knowledge about a certain element of the game, he would formulate his best responses in terms of what he *believes* is the missing (from his perspective) part. In a response, other players would formulate their optimal strategies in terms of not only their beliefs about the uncertain primitives of the game but also taking into account what they *believe* the beliefs of the agent i are (as they will impact his strategy choices). Then, player i should take into account beliefs of his competitors about his own beliefs. And so on.

This looks like a hopeless circle of **hierarchical beliefs**. However, it turns out we can re-express the problem in a very tractable way under relatively weak assumptions. We skip the details here, refer to Harsanyi (1967-1968) and Mertens & Zamir (1985) for more details. We focus on the result, which is the **Bayesian Games** formulation of games of incomplete information.

Consider a game with N players. For simplicity, we assume every player has exactly one decision node and that all decisions are made simultaneously (it can be easily extended to sequential games with more than one decision node per player). Every player i has a set of available actions A_i with a typical action a_i.

Definition 10.2.1 (individual type)**.** *Every agent has a type $\theta_i \in \Theta_i$ where Θ_i the set of all possible individual types for agent i.*

We introduce more notation regarding agents' types. Let the profile of types be denoted by $\theta = (\theta_1, ..., \theta_N)$ and the set of types of i's competitors be $\theta_{-i} = (\theta_1, ..., \theta_{i-1}, \theta_{i+1}, ..., \theta_N)$.

The incomplete information is related to individual types. We assume that θ_i is observed only by agent i. This affects the definition of (pure) strategies in a Bayesian Game.

Definition 10.2.2 (pure strategy)**.** *A **pure strategy** of player i is a function mapping their type space into the set of actions:*

$$s_i : \Theta_i \to A_i$$

we denote $s_i(\theta_i) \in A_i$.

Payoffs in Bayesian Games depend on actions chosen by individuals as well as their types.

Definition 10.2.3 (individual payoffs)**.** ***Individual payoffs** are characterized by a utility function mapping action profiles and type profiles into the real line:*

$$u_i : \times_{i=1}^{N} A_i \times_{i=1}^{N} \Theta_i \to \mathbb{R}$$

with a typical payoff denoted by $u_i(a_1, ..., a_N, \theta_1, ..., \theta_N) \equiv u_i(a, \theta)$.

The last element to define is the source of incomplete information and related agents' beliefs. Conceptually, we think of agents' types as random realizations drawn from a joint distribution of all possible agent type profiles.

Definition 10.2.4 (distribution of types)**.** *Every type profile θ is a drawn from a probability distribution described by a cumulative distribution $F : \times_{i=1}^{N} \Theta_i \to [0, 1]$.*

This distribution of agents' types is often referred to as the **common prior**, because we assume that everybody knows its functional form before making decisions. The common prior describes agent's beliefs about types of other agents and we will typically assume that these beliefs are correct.

The above formulation of the common prior if very general and assumes that types can depend on each other in unrestricted ways. In practice, we often impose some independence assumptions do decrease the dimensionality of the problem. For example, if i's payoffs only depend on i's type (but not on θ_j for $j \neq i$), the we deal with a **private value** model, as opposed to an **interdependent values** model otherwise. If for every i θ_i is drawn independently from any other θ_j such that $j \neq i$, we talk about a model with **independent types**. Otherwise, it would be a model with **correlated types**.

Having defined the primitive of Bayesian Game, it is time to think how the game proceeds. We conceptualize the timing in the following way. In the first step, **nature** moves by drawing a realization of type profile θ. Every player i learns their own θ_i but does not know the realizations of θ_j for $j \neq i$. However, when types are not independent, θ_i may carry some information about possible realizations of θ_{-i}. In the second step, agents update their beliefs regarding the types of other players to $F(\theta_{-i}|\theta_i)$. Third, players choose their actions (if the actions are taken sequentially, we may need to re-update the beliefs) Fourth, payoffs realize.

We use the **Bayes' rule** to update the beliefs. For discrete distributions:

$$p_i(\theta_{-i}|\theta_i) = \frac{p(\theta_i, \theta_{-i})}{\sum\limits_{\theta'_{-i} \in \Theta_{-i}} p(\theta_i, \theta'_{-i})}$$

Now, we move to describing how agents make their decisions. They choose their actions only knowing their type, without knowing types of their competitors. Given that everyone knows the common prior, it is straightforwards to define agents' expected payoffs.

Definition 10.2.5 (expected payoffs)**.** *Agents'* ***expected payoffs*** *are defined as an expected value (with respect to the distribution of θ_{-i}) of their individual payoff function:*

$$\mathbb{E}_{\theta_{-i}}\Big[u_i\big(s_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i, \theta_{-i}\big)\Big|\theta_i\Big] = \begin{cases} \sum_{\theta-i \in \Theta_{-i}} u_i\big(s_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i, \theta_{-i}\big)p_i(\theta_{-i}|\theta_i) & \text{if distribution of } \theta_{-i} \text{ is discrete} \\ \int_{\theta-i \in \Theta_{-i}} u_i\big(s_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i, \theta_{-i}\big)\mathrm{d}F_{-i}(\theta_{-i}|\theta_i) & \text{if distribution of } \theta_{-i} \text{ is continuous} \end{cases}$$

With the expected payoffs formulated, we are ready to define a Nash equilibrium of a Bayesian Game.

Definition 10.2.6 (Bayesian Nash equilibrium)**.** *A strategy profile $\Big(s_1(\theta_i), ..., s_N(\theta_N)\Big)$ is a* ***Bayesian Nash Equilibrium*** *(BNE) if :*

$$\mathbb{E}_{\theta_{-i}}\Big[u_i\big(s_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i, \theta_{-i}\big)\Big|\theta_i\Big] \geq \mathbb{E}_{\theta_{-i}}\Big[u_i\big(s'_i(\theta_i), s_{-i}(\theta_{-i}), \theta_i, \theta_{-i}\big)\Big|\theta_i\Big]$$

for every $i \in \{1, ..., N\}$, $\theta_i \in \Theta_i$, and $s'_i(\theta_i) \in A_i$.

Theorem 10.2. *Every finite game of incomplete information possesses at least one Bayesian-Nash equilibrium.*

How to find Nash equilibria of Bayesian games? Helpful may be looking at some specific kinds of strategies. A way of making a Bayesian game more tractable is to note that for each profile o types θ it generates a full information game. If a strategy profile is a Nash equilibrium in each such game, it has to be a Nash equilibrium of the Bayesian game.

Definition 10.2.7 (weakly dominant strategies)**.** *Strategy $s_i(\theta_i)$ is* ***weakly dominant*** *if it yields weakly higher payoffs than any other strategy $s'_i(\theta_i)$:*

$$u_i(s_i(\theta_i), a_{-i}, \theta) \geq u_i(s'_i(\theta_i), a_{-i}, \theta)$$

for every $a_{-i} \in A_{-i}$ and θ.

Definition 10.2.8 (undominated strategies)**.** *A strategy $s_i(\theta_i)$ is* ***undominated*** *if there exists no other strategy $s'_i(\theta_i) \in A_i$ that weakly dominates it.*

Dominant strategies are optimal for an agent in every state of the world, i.e. for every draw of θ. If every player has a dominant strategy, we can talk about equilibrium.

Proposition 10.2.1 (equilibrium in dominant strategies)**.** *Let $s(\theta) = \{s_i(\theta_i)\}_{i,\theta_i}$ be a profile of weakly dominant strategies. Then $s(\theta)$ is a* ***Nash equilibrium in dominant strategies*** *of the Bayesian game.*

We defined the Bayesian game in three stages: the **ex-ante** stage, in which agents only know the common prior, the **interim** stage in which agents only know their type, and the **ex-post** stage in which agents learn everything (and receive their payoffs). Even though the decisions are made in the interim stage, it is instructive to consider their optimality from the point of view of the information available at every stage.

Proposition 10.2.2. *Let $s(\theta)$ be a Bayesian Nash equilibrium. Then it is ex-ante optimal:*

$$\mathbb{E}_\theta\Big[u_i\big(s_i(\theta_i), s_{-i}(\theta_{-i}), \theta\big)\Big] \geq \mathbb{E}_\theta\Big[u_i\big(s'_i(\theta_i), s_{-i}(\theta_{-i}), \theta\big)\Big]$$

for every i and $s'_i(\theta)$.

It also works the other way around.

Proposition 10.2.3. *Suppose strategy profile $s(\theta)$ is ex-ante optimal. Then it is interim-optimal almost surely (except possibly a set of points θ such that $p(\theta) = 0$).*

Definition 10.2.9 (ex-post equilibrium)**.** *A Nash-equilibrium strategy profile $s(\theta)$ is an* ***ex-post equilibrium*** *if:*

$$u_i(s(\theta),\theta) \geq u_i(a_i, s_{-i}(\theta_{-i}),\theta)$$

for every action $a_i \in A_i$.

The ex-post equilibria have a useful property of not depending on the distribution of agent types.

Exercise 10.2.1. *Consider two possibly different common priors given by distributions F and G. If strategy profile $s(\cdot)$ is an ex-post equilibrium under F, then it is also an ex-post equilibrium under G.*

The ex-post equilibrium is a Nash equilibrium that is optimal at every realization of θ.

We conclude with a ranking of the different kinds of Bayesian games equilibria described above.

Proposition 10.2.4. *The following holds:*
the set of dominant strategy equilibra $\subseteq$ the set of ex-post equilibria $\subseteq$ the set of BNE

10.3 Auctions

A lot of effort in economics has been made to study mechanisms [5] of allocating objects (a good, a service, a contract, ...) among a group of agents. In our class, we focus on a task of transferring exactly one indivisible object between two agents. However, we typically assume that there is more than just two agents interested in a transaction.

Such a task can be resolved by an **auction**. An auction consists of an **auctioneer** who either sells (**forward** or selling **auction**) or buys (**reverse** or procurement **auction**) an object, and N **bidders** who are interested in buying (selling auction) or selling (procurement auction) the object, and naturally the object itself. The auctioneer solicit bids—offers of purchase or sale—from the bidders and chooses one according to a pre-specified rule, known to every agent.

In a standard setting, we assume that the object creates some value to every agent. However, every agent observes only their own **valuation**, but do not know precisely what are valuations of other agents. For this reason, auctions are conveniently expressed as Bayesian games.

To see how it works, we consider a very simple example of a forward auction[6] . In a selling auction, an **auctioneer** seeks to sell a good to any of N bidders. Bidders value the object differently. We express it by assuming that every bidder $i \in \{1,...,N\}$ receives a draw of a realization of their **valuation** (shortly: value) v_i. The joint distribution of bidders' valuations is given by $\mathcal{F}$. Using the Bayesian games language, this is the common prior. Different assumptions regarding $\mathcal{F}$ place us in a little bit different auction frameworks.

10.4 Auction Paradigms

Definition 10.4.1 (independent private value auctions)**.** *For every bidder i, let $F_i : [0,\overline{v}_i] \to [0,1]$ be the distribution of valuations in which we normalize the lower extremity of the domain to zero. The auction lies within the* ***independent private value*** *(IPV) paradigm if bidders' values are drawn independently from each other from their respective distributions:*

$$v_i \sim F_i(\cdot), i \in \{1,...,N\}$$

Definition 10.4.2 (common value auctions)**.** *In a* ***common value auction****, the value of the good is the same for every bidder, say v. However, it is not revealed and the bidders have different information about its possible realization. An example of a common value auction is an* ***affiliated values auction****, in which bidders' valuations are given by:*

$$v_i = \mathbb{E}[v] + \epsilon_i$$

where ϵ_i is a individual private shock to the estimate of the common value.

In what follows, we remain in the independent private value framework. We additionally assume that bidders are **symmetric**.

Definition 10.4.3 (symmetric bidders)**.** *Bidders are said to be* ***symmetric*** *if for every $i \in \{1,...,N\}$, c_i is drawn from the same distribution:*

$$F(\cdot) : [0,\overline{v}] \to [0,1]$$

Symmetry greatly simplifies the analysis and is useful for us to introduce the frameworks. In reality, we often have reasons that various bidders may differ in their cost distributions.

[5] we skip the formal definition of a mechanism in our class. Take Arthur's micro theory class to learn more about mechanisms

[6] there is a sense of analogy between forward and reverse auctions, but it is more subtle than you might expect. See Huh & Roundy, 2004.

10.5 Auction Formats

The bidders know their values. Now, it is time to choose a bid. Optimal bids depend on the **auction format**.

Definition 10.5.1 (auction format). ***Auction format*** *specifies the algorithm of choosing the winner from the set of bidders and the contract price.*

There are **many** auction formats used in practice. Here, we focus on four traditional formats. Denote by b_i a **bid**—price offer—submitted by bidder i, and by p_i the price that the auctioneer will pay to the winner.

Definition 10.5.2 (first-price sealed-bid auction). *In a* ***first-price sealed-bid auction****, bidders submit bids simultaneously. The contract is allocated to a bidder that submitted the highest bid:*

$$b_i \text{ wins } \Leftrightarrow b_i = \max_{k\in\{1,..,N\}} b_k$$

where ties are broken randomly. The winner receives the good in exchange for their bid:

$$p_i = b_i$$

Sometimes, economists would also use the term **high-price auction** to describe this format. We prefer to call it the first-price auction, because this nomenclature naturally extends also to the forward auctions, in which the lowest bid wins.

Definition 10.5.3 (second-price sealed-bid auction). *In a* ***second-price sealed-bid auction****, bidders submit bids simultaneously. The contract is allocated to a bidder that submitted the lowest bid:*

$$b_i \text{ wins } \Leftrightarrow b_i = \min_{k\in\{1,..,N\}} b_k$$

where ties are broken randomly. The winner receives the good in exchange for the price set at the second highest bid:

$$p_i = \max_{k\neq i} b_k$$

These two auction formats are simple in the sense that every bidder has only one decision node and all decisions are made simultaneously. We can easily find an auction format that deviates from this principle.

Definition 10.5.4 (Dutch auction). *In a* ***Dutch auction****, the auctioneer begins with a very high price and gradually decreases it. The winner is the first bidder to raise their hand, and the value of the contract is set to the current price.*

Definition 10.5.5 (English auction). *In an* ***English auction****, the auctioneer begins with a very low price and gradually increases it. At every price, each bidder signals whether they are interested in delivering the contract. If not, such bidder is excluded from the bidding. The auction concludes when only one bidder remains willing to provide the contract at the current price, and this price becomes the value of the contract.*

It turns out that these seemingly very different auction formats have a lot in common.

Proposition 10.5.1. *The following hold:*

1. *Dutch auctions are equivalent to first-price auctions.*
2. *English auctions are equivalent to second-price auctions.*

10.6 Optimal Bidding

In a typical auction game, the set of actions is the set of bids to submit. Without any further restrictions (e.g. reserve price), we can think of them as the set of all positive real numbers:

Definition 10.6.1 (action sets). *For every bidder, the set of all available actions is $\mathbb{R}_{++}$.*

Definition 10.6.2 (bidding strategy). *A* ***bidding strategy*** *is a function that maps the realization of individual private value v_i into the set of all bids:*

$$\beta_i : [0, \overline{v}] \rightarrow [0, 1]$$

Definition 10.6.3 (Bayesian Nash equilibrium)**.** *A Bayesian Nash equilibrium in the auction game is a strategy profile $\beta(v) \equiv \{\beta_i(v_i)\}_i$ such that for every i, the expected profit achieved at $\beta_i(v_i)$ is weakly higher than the expected profit achieved at any other action $a_i \in A_i$, given the strategies of other bidders.*

The above definition of BNE is very general, and hence it may be difficult to fully characterize the set of BNEs. In order to proceed, we need to restrict our attention to a reasonable subset of strategies. First, it should not be controversial to assume that the optimal bid should be increasing in the realization of the private value.

Definition 10.6.4 (strictly increasing bidding strategy)**.** *We say that the bidding strategy is* ***strictly increasing*** *if $\beta_i(v_i)$ is strictly increasing in v_i.*

Second, every bidder in our game has the same distribution of project completion costs. It is natural to assume that every bidder would choose the same bid if facing the same value of their private value.

Definition 10.6.5 (symmetric strategies)**.** *We say that the bidding strategy is symmetric if for every $i \neq j$ and every $v \in [0, \overline{v}]$:*

$$\beta_i(v) = \beta_j(v)$$

Now, we are ready to look for a symmetric Bayesian Nash equilibrium in strictly increasing strategies. For notational simplicity, assume that ties do not happen.

We begin with studying the first-price sealed-bid format. If player i wins with bid b_i, i.e. b_i is the highest bid, their payoff is equal to $(v_i - b_i)$. Otherwise, the payoff of a loser is normalized to zero. Bidder i chooses bid b_i to maximize expected profit:

$$\max_{b_i \geq 0} (v_i - b_i) \cdot \text{ Prob } \Big[b_i > b_j \text{ for all } j \neq i \Big]$$

The key object is the probability of winning, or that b_i turns out to be the highest bid. In principle, it depends on the (optimal) bidding strategy of every other bidder. We can make some progress by using the assumed properties of the bidding functions.

By strict monotonicity, we can invert the optimal bidding to obtain the **inverse bidding function** indicating underlying valuation v at every values of the submitted bid b. If $b_i = \beta_i(v_i)$ then:

$$v_i = \beta_i^{-1}(b_i)$$

Coming back to the probability of winning of bidder i:

$$\begin{aligned}
\text{Prob } \Big[b_i > b_j \text{ for all } j \neq i \Big] &= \text{ Prob } \Big[\beta_j^{-1}(b_i) > \beta_j^{-1}(b_j) \text{ for all } j \neq i \Big] \\
&= \prod_{j \neq i} \text{ Prob } \Big[\beta_j^{-1}(b_i) > \beta_j^{-1}(b_j) \Big] \\
&= \prod_{j \neq i} \text{ Prob } \Big[\beta_j^{-1}(b_i) > v_j \Big] \\
&= \prod_{j \neq i} F_j \Big[\beta_j^{-1}(b_i) \Big] \\
&= \prod_{j \neq i} F \Big[\beta^{-1}(b_i) \Big]^{N-1}
\end{aligned}$$

where the second equality follows from assuming valuations are drawn independently across bidders, and the last follows from the assumption of symmetric value distributions. Note that although $\beta^{-1}(b_i) = v_i$, winning probability is a function of the bid. So while looking for an optimal bidding strategy, we need to take into account this dependence.

Given a realization of v_i, bidder i chooses bid b_i to maximize:

$$\max_{b_i \geq 0} (v_i - b_i) F \Big[\beta^{-1}(b_i) \Big]^{N-1}$$

The first order condition implies:

$$-F\Big[\beta^{-1}(b_i)\Big]^{N-1} + (v_i - b_i)(N-1)F\Big[\beta^{-1}(b_i)\Big]^{N-2} f\big(\beta^{-1}(b_i)\big)\frac{1}{\beta'\big(\beta^{-1}(b_i)\big)} = 0$$

Simplifying:

$$F\Big[\beta^{-1}(b_i)\Big]^{N-1}\beta'\big(\beta^{-1}(b_i)\big) + (N-1)b_i F\big(\beta^{-1}(b_i)\big)^{N-2} f\big(\beta^{-1}(b_i)\big) = (N-1)v_i F\big(\beta^{-1}(b_i)\big)^{N-2} f\big(\beta^{-1}(b_i)\big)$$

Now, by definition of the inverse bidding function, i.e. $\beta^{-1}(b_i) = v_i$, we can write:

$$\underbrace{F\big(c_i\big)^{N-1}\beta'\big(v_i\big) + (N-1)\beta(v_i)F\big(v_i\big)^{N-2} f\big(v_i\big)}_{=\left(F\left(v_i\right)^{N-1}\beta(v_i)\right)'} = (N-1)v_i F\big(v_i\big)^{N-2} f\big(v_i\big)$$

So:

$$F\big(v_i\big)^{N-1}\beta(v_i) = \int_0^{v_i}(N-1)v\Big(1-F\big(v\big)\Big)^{N-2} f\big(v\big)\mathrm{d}v$$

We take a closer look at the right-hand side expression. Integrating by parts:

$$\begin{aligned}\int_0^{v_i} v(N-1)F\big(v\big)^{N-2} f\big(v\big)\mathrm{d}v &= \int_0^{v_i} v\Big[F\big(v\big)^{N-1}\Big]'\mathrm{d}v \\ &= \Big[F\big(v\big)^{N-1}v\Big]_0^{v_i} - \int_0^{v_i} F\big(v\big)^{N-1}\mathrm{d}v \\ &= F\big(v_i\big)^{N-1}v_i - \int_0^{v_i} F\big(v\big)^{N-1}\mathrm{d}v\end{aligned}$$

Plugging it back to the FOC:

$$F\big(v_i\big)^{N-1}\beta(v_i) = F\big(v_i\big)^{N-1}v_i - \int_0^{v_i} F\big(v\big)^{N-1}\mathrm{d}v$$

so finally:

$$\beta(v_i) = v_i - \frac{\int_0^{v_i} F\big(v\big)^{N-1}\mathrm{d}v}{F\big(v_i\big)^{N-1}}$$

The following proposition summarizes our derivations.

Proposition 10.6.1. *The forward first-price sealed-bid auction has a unique equilibrium in strictly increasing symmetric strategies given by:*

$$\beta(v_i) = v_i - \frac{\int_0^{v_i} F\big(v\big)^{N-1}\mathrm{d}v}{F\big(v_i\big)^{N-1}}$$

The optimal bidding consists of two terms: the realization of the private bidder valuation and the so called **bid-shading** term, decreasing the bid below the valuation.

Now, we move to second price auctions.

Proposition 10.6.2. *In the forward second-price sealed-bid auction, it is a weakly dominant strategy to bid the private valuation:*

$$\beta(v_i) = v_i$$

Proof. WLOG consider bidder 1 and suppose that the highest competing bid is given by:

$$p_1 = \max_{j \neq 1} b_j$$

By bidding v_1, bidder 1 wins if $v_1 > p_1$ and loses otherwise. Suppose for a contradiction that bidder 1 bids $w \neq v_1$. Suppose $w < v_1$. If $v_1 > w > p_1$, bidder 1 still wins and their profit remains $v_1 = p_1$. If $p_1 > v_1 > w$, bidder 1 still loses. If $v_1 > p_1 > w$, then bidder loses if posting w and wins (with a positive payoff) if posting v_1. Therefore, bidding w instead v_1 can never increase the profit but may actually decrease it. Similar argument applies for $w > v_1$. So $\beta(v) = v$ is a weakly dominant strategy. □

Since every bidder has a weakly dominant strategy, the second-price auction has a Bayesian Nash equilibrium in dominant strategies. This time, it is very simple!

Which of the two auction formats would be preferred by the buyer?

Theorem 10.3 (revenue equivalence)**.** *With independently and identically distributed private values, the expected revenue in a first-price auction is the same as the expected revenue in a second-price auction.*

Which of the two auction formats would be preferred by the social planner?

Theorem 10.4. *With independently and identically distributed private values, both first-price and second-price auction are efficient: the object goes to the bidder with the highest valuation.*