



**MATHEMATICAL METHODS  
FOR ECONOMIC ANALYSIS**

Itzhak Rasooly  
10 December 2025

## CONTENTS

<b>0</b>	<b>The mathematisation of economics</b>	<b>1</b>
0.1	Some history . . . . .	1
0.2	Some perspectives . . . . .	2
<b>1</b>	<b>Preliminaries</b>	<b>6</b>
1.1	Sets and numbers . . . . .	6
1.2	Exponents . . . . .	7
1.3	Logarithms . . . . .	9
<b>2</b>	<b>Sequences, sums and limits</b>	<b>11</b>
2.1	Motivating examples . . . . .	11
2.2	Sequences . . . . .	11
2.3	Sums . . . . .	16
<b>3</b>	<b>Solving equations</b>	<b>19</b>
3.1	Motivating examples . . . . .	19
3.2	Solving equations . . . . .	19
3.3	Linear equations . . . . .	20
3.4	Non-linear equations . . . . .	25
<b>4</b>	<b>Functions</b>	<b>27</b>
4.1	Motivating examples . . . . .	27
4.2	Introduction . . . . .	27
4.3	Examples . . . . .	29
4.4	Key concepts . . . . .	32
<b>5</b>	<b>Derivatives</b>	<b>37</b>
5.1	Motivating examples . . . . .	37
5.2	Defining derivatives . . . . .	37
5.3	Examples . . . . .	40
5.4	Rules . . . . .	41
5.5	Applications . . . . .	43
<b>6</b>	<b>Multivariate functions</b>	<b>47</b>
6.1	Motivating examples . . . . .	47
6.2	Multivariate functions . . . . .	47
6.3	Classifying functions . . . . .	50
6.4	Derivatives . . . . .	53
<b>7</b>	<b>Optimisation</b>	<b>57</b>
7.1	Motivating examples . . . . .	57
7.2	Optimisation . . . . .	57

7.3	Three basic results . . . . .	59
7.4	An application . . . . .	62
<b>8</b>	<b>Utility maximisation</b>	<b>64</b>
8.1	Motivating examples . . . . .	64
8.2	The consumer's problem . . . . .	64
8.3	The MRS method . . . . .	66
8.4	General results . . . . .	69
8.5	Three examples . . . . .	71
<b>9</b>	<b>Cost minimisation</b>	<b>74</b>
9.1	Motivating examples . . . . .	74
9.2	The firm's problem . . . . .	74
9.3	The MRTS method . . . . .	75
9.4	General results . . . . .	79
9.5	Three examples . . . . .	80
<b>10</b>	<b>Integration</b>	<b>82</b>
10.1	Motivating examples . . . . .	82
10.2	Indefinite integrals . . . . .	82
10.3	Some integration rules . . . . .	83
10.4	Definite integrals . . . . .	84
10.5	Economic applications . . . . .	86

## 0. THE MATHEMATISATION OF ECONOMICS

### 0.1 Some history

Modern economics is sometimes said to have started with the publication of Adam Smith's *The Wealth of Nations* in 1776. Here is a well-known excerpt in which [Smith \(1776\)](#) expounds his views on what we would now call 'competition policy':

People of the same trade seldom meet together, even for merriment and diversion, but the conversation ends in a conspiracy against the public, or in some contrivance to raise prices. It is impossible indeed to prevent such meetings, by any law which either could be executed, or would be consistent with liberty and justice. But though the law cannot hinder people of the same trade from sometimes assembling together, it ought to do nothing to facilitate such assemblies; much less to render them necessary.

Perhaps even more widely cited is the following passage in which Smith explains how self-interested behaviour can encourage individuals to try to satisfy the wants of others:

It is not from the benevolence of the butcher, the brewer, or the baker that we expect our dinner, but from their regard to their own interest. We address ourselves, not to their humanity, but to their self-love, and never talk to them of our own necessities, but of their advantages.

Importantly for our purposes, Smith's book did not contain very much mathematics or formal symbolism. Instead, Smith (and his contemporaries) used what we might call a 'literary' style.

For our next excerpt, we will consider Jevons's 1871 publication *The Theory of Political Economy*. In contrast to Smith's work, [Jevons \(1871\)](#) does contain substantial mathematical content. However, the style is informal and the book contains no rigorous theorems. Here is an extract from a more mathematical portion of the book:

In attempting to represent these conditions of labour with accuracy, we shall find that there are no less than four quantities concerned; let us denote them as follows:—

$t$  = time, or duration of labour ;

$l$  = amount of labour, as meaning the aggregate balance of pain accompanying it, irrespective of the produce ;

$x$  = amount of commodity produced ;

$v$  = total utility of that commodity.

Jumping forward in time once again, let us now consider economics in the middle of the twentieth century. By this time, economists had begun to adopt far more formal mathematical frameworks, especially in some theoretical subfields. In particular, they

had started (i) rigorously proving theorems based on clearly stated assumptions, and (ii) making use of substantially more complex branches of mathematics. (See [Debreu, 1991](#) for some more semi-formal evidence on the historical pace and timing of this ‘mathematisation’). Our final extract, from [Arrow and Debreu \(1954\)](#), illustrates the level of mathematical sophistication that economics had reached by the early 1950s:

2.1. Let there be  $v$  subsets of  $\mathbb{R}^l$ ,  $\mathfrak{A}_\iota$  ( $\iota = 1, \dots, v$ ). Let  $\mathfrak{A} = \mathfrak{A}_1 \times \mathfrak{A}_2 \times \dots \times \mathfrak{A}_v$ , i.e.  $\mathfrak{A}$  is the set of ordered  $v$ -tuples  $a = (a_1, \dots, a_v)$ , where  $a_\iota \in \mathfrak{A}_\iota$  for  $\iota = 1, \dots, v$ . For each  $\iota$ , suppose there is a real function  $f_\iota$  defined over  $\mathfrak{A}$ . Let  $\bar{\mathfrak{A}}_\iota = \mathfrak{A}_1 \times \mathfrak{A}_2 \times \dots \times \mathfrak{A}_{\iota-1} \times \mathfrak{A}_{\iota+1} \times \dots \times \mathfrak{A}_v$ , i.e., the set of ordered  $(v-1)$ -tuples  $\bar{a}_\iota = (a_1, \dots, a_{\iota-1}, a_{\iota+1}, \dots, a_v)$ , where  $a_{\iota'} \in \mathfrak{A}_{\iota'}$  for each  $\iota' \neq \iota$ . Let  $A_\iota(\bar{a}_\iota)$  be a function defining for each point  $\bar{a}_\iota \in \bar{\mathfrak{A}}_\iota$  a subset of  $\mathfrak{A}_\iota$ . Then the sequence  $[\mathfrak{A}_1, \dots, \mathfrak{A}_v, f_1, \dots, f_v, A_1(\bar{a}_1), \dots, A_v(\bar{a}_v)]$  will be termed an *abstract economy*.

Economics remains highly mathematical today. This might not be clear from taking undergraduate economics courses, which sometimes downplay the mathematisation of the discipline. However:

- Economics education is highly mathematical at the *graduate* level (see, for example, the classic graduate textbook [Mas-Colell et al., 1995](#)).
- Economic theory remains highly mathematical (see, for example, any recent issue of *Econometrica* or the *Journal of Economic Theory*).
- More ‘applied’ economics makes heavy use of *data* and analyses these data using mathematical methods.

This short history raises at least two questions. First, *why* has this happened? Second, is it a *good thing*? While we won’t have much to say about the first question, we will say something about the second. Unsurprisingly, there are a variety of views on the benefits (and costs) of a highly mathematical approach to economics.

## 0.2 Some perspectives

In what follows, I survey some especially influential arguments in favour of a mathematical approach to economics.

**Making assumptions precise.** A classic justification for using formal, mathematical models is that they encourage researchers to make their assumptions explicit. For example, [Rodrik \(2015\)](#) claims that:

. . . math ensures that the elements of a model—the assumptions, behavioral mechanisms, and main results—are stated clearly and are transparent. Once a model is stated in mathematical form, what it says or does is obvious to all who can read it. This clarity is of great value and is not adequately appreciated. We still have endless debates today about what Karl Marx, John Maynard Keynes, or Joseph Schumpeter really meant. Even though all three are giants of the economics

profession, they formulated their models largely (but not exclusively) in verbal form. By contrast, no ink has ever been spilled over what Paul Samuelson, Joe Stiglitz, or Ken Arrow had in mind when they developed the theories that won them their Nobel.

*Commentary.* I think there is something to this: you cannot rigorously present a mathematical model without explicitly stating all of your assumptions. Having said that, even if one *does* state one's assumptions symbolically, the connection between these symbols and the phenomena that they are supposed to represent may remain obscure. (See [Romer, 2015](#) for elaboration.) In my view, which I will not defend, many classic macroeconomic variables (e.g., 'aggregate demand', 'capital', 'labour') fall into this category. Although Romer would probably not accept these examples, he provides some less controversial examples that you may find more persuasive.

**Ensuring the validity of arguments.** A second justification for using mathematical models is that they allow one to check whether one's conclusions logically follow from one's premises. (They also allow one to check whether one's assumptions are logically consistent.) To quote again from [Rodrik \(2015\)](#):

. . . [mathematics] ensures the internal consistency of a model—simply put, that the conclusions follow from the assumptions. This is a mundane but indispensable contribution. Some arguments are simple enough that they can be self-evident. Others require greater care, especially in light of cognitive biases that draw us toward results we want to see.

*Commentary.* Again, there is definitely something to this. However, there is also room for debate about whether economists should focus so much attention on elaborate chains of deduction whose validity can only be confirmed through formal methods. One might also question the value of logical validity when economic models rely on assumptions that are rarely (if ever) fully satisfied in reality. While it is clearly desirable to construct valid (i.e., truth-preserving) arguments on the basis of accurate assumptions, it is less obvious how much value logical validity provides if one starts from less than fully accurate premises.

**Explaining and predicting quantitative phenomena.** Historically, an important justification for a mathematical approach to economics has been that economics is concerned with explaining (or predicting) numerical quantities (e.g., exchange rates, inflation rates, interest rates, ...) As [Jevons \(1871\)](#) puts it,

To me it seems that *our science must be mathematical, simply because it deals with quantities*. Wherever the things treated are capable of being *greater or less*, there the laws and relations must be mathematical in nature.

Importantly, Jevons adds that

. . . we do not render the science less mathematical by avoiding the symbols of algebra,—we merely refuse to employ, in a very imperfect science, much need-

ing every kind of assistance, that apparatus of appropriate signs, which is found indispensable in other sciences.

Thus, the idea is that, assuming we need mathematical models to explain or predict quantitative phenomena (like prices), we may as well make use of mathematical notation and known mathematical results.

*Commentary.* In my view, any successful defence of mathematical economics is going to need an argument along these lines. The first two arguments (that formalisation makes assumptions more explicit and helps check the validity of arguments) could be applied to *any* academic discipline. Despite this, however, it seems very implausible that (e.g.) intellectual historians should start exclusively expressing themselves using mathematical models. What is needed, then, is some reason why economics (as opposed to some other disciplines) is especially amenable to a mathematical treatment. Whatever its merits, Jevons' argument attempts to provide such a reason: how persuasive this reason is is something that you will need to decide.

So far, we have discussed some *defences* of mathematical economics. However, there has also been no shortage of criticisms. I won't survey all of these, not least because some of these criticisms strike me as highly implausible. However, I will briefly mention two criticisms that strike me as worth taking seriously.

**An illusion of respectability.** One idea is that the use of complex mathematics gives economics an unearned appearance of 'scientific objectivity'. As [Levinovitz \(2016\)](#) puts it, mathematisation 'imbues economic theory with unearned empirical authority' (see also [McCloskey, 1998](#).) One reason for this, in Levinovitz's view, is that mathematisation allows economics to borrow from the prestige of highly successful (and highly mathematical) fields like theoretical physics. Another proposed reason is that mathematisation makes economic theory sufficiently difficult to comprehend that outsiders are unable to spot its flaws and weaknesses.

**Bad incentives.** Another idea is that the emphasis on mathematics encourages economists to focus on proving 'elegant' theorems at the cost of engaging with practical and socially useful policy questions (see, e.g., [Krugman, 2009](#) for remarks along these lines.) The claim here need not be that abstract proving theorem is necessarily useless: rather, one might argue that it 'crowds out' other scientific activities that would deliver a greater social or intellectual payoff.

*Commentary.* I am sympathetic to both of these critiques. Having said this, it is important to realise that neither are critiques of mathematical economics *per se*: for instance, while you *can* use an elaborate mathematical model to hide the weaknesses of your arguments, you can also be 'up front' about your model's limitations. Instead, I think that it is best to view these arguments as a warning about particular ways of doing (and 'selling') mathematical economics.

So, is it a good thing that economics has become so mathematical? I am not going

to attempt to settle that issue here, although I will say that I find ‘extreme’ views on this question to be rather implausible (see, e.g., [Rothbard, 2009](#) for a maximally ‘anti-mathematical’ view and [Lucas Jr, 2001](#) for the opposite extreme.) I will say, however, that you have a good reason to learn some mathematics almost independently of whether the discipline should (or should not) have been mathematised. For example:

- You may find the mathematics interesting.
- You will need to know some mathematics if you want to master your subsequent undergraduate courses (and *especially* if you want to study economics at the graduate level).
- You will need to know some mathematics if you want to understand or contribute to economic research (since it is so mathematical).
- You also have a more *immediate* reason to learn some mathematics, namely your EC1019 exams...

Hopefully, this provides you with some motivation to master the material in this course!



## 1. PRELIMINARIES

In the following, we discuss some topics that are not large enough to merit a week in themselves but nonetheless underpin much of the material in this course.

### 1.1 Sets and numbers

We begin with a definition.

#### Definition 1.1

A *set* is a collection of objects.

Some examples may illustrate the concept. One set is the set of the English vowels:  $\{a, e, o, u, i\}$ . Note that the *objects* or *elements* (here, the vowels) are separated by commas and collected in the ‘curly’ brackets  $\{\}$ . Another (rather arbitrary) example is the set of fruits:  $\{\text{apple, banana, pear, ...}\}$ . One can also consider sets containing numbers. For example, consider the (finite) sets  $\{1, 3\}$  and  $\{5.5, 600, 10\}$ .

**Specifying a set.** There are two ways to specify a set:

- Exhaustively list all of its elements, as in the previous example. Notice that this is possible only if the set is finite.
- Describe the *property* that all elements of the set are supposed to have, e.g.,  $\{x: x \text{ is an economist}\}$ . This reads: ‘the set of all  $x$  such that  $x$  is an economist’.

**Some important sets.** The following sets arise frequently in economic theory:

- The *empty set*  $\emptyset$ . This is simply the set that contains no elements.
- The set of *natural numbers*  $\mathbb{N} = \{1, 2, 3, \dots\}$ . This is the set of ‘counting numbers’. Note that, in contrast to some texts, we do *not* count 0 as a natural number.
- The set of *integers*  $\mathbb{Z} = \{\dots, -2, -1, 0, 1, 2, \dots\}$ . This set includes not just the natural numbers but also zero along with the negative ‘whole numbers’.
- The set of *real numbers*  $\mathbb{R}$ . This is the set of *all* the numbers of the ‘number line’, including not just the ‘whole’ numbers but also numbers like 3.4,  $1/3$  and  $\pi$ .
- The *closed interval*  $[a, b]$ , *open interval*  $(a, b)$ , and *semi-closed intervals*  $[a, b)$  and  $(a, b]$ . These are formally defined as follows:

$$[a, b] = \{x: a \leq x \leq b\}$$

$$(a, b) = \{x: a < x < b\}$$

$$[a, b) = \{x: a \leq x < b\}$$

$$(a, b] = \{x: a < x \leq b\}$$

The important point is that a *closed* interval (denoted with *square* brackets) contains its ‘boundaries’ or ‘endpoints’ whereas an *open* interval (denoted with *circular* brackets) does not. A *semi-closed* interval contains one endpoint but not the other.

As an aside, note that *all* the sets above (with the exception of  $\emptyset$ ) are *infinite* in the sense that they contain an infinite number of elements. For example,  $\mathbb{N}$  is infinite since there is no largest natural number. Indeed, for any natural number that you can think of, one can produce an even larger natural number by (for example) taking that number and adding 1.

**Some important notation.** The following notation is vital to memorise:

- $x \in X$ : element  $x$  is a member of set  $X$ .
- $x \notin X$ : element  $x$  is not a member of set  $X$ .

The following problems will test your understanding of this notation as well as the sets that we have encountered so far.

#### Problem 1.1

Determine whether the following statements are true or false:

- (1)  $-3 \in \mathbb{N}$
- (2)  $1.5 \in \mathbb{Z}$
- (3)  $0.22 \notin \mathbb{R}$
- (4)  $1 \in [0, 1]$
- (5)  $0 \in (0, 1]$

## 1.2 Exponents

For our next topic, we will discuss *exponents* (otherwise known as ‘indices’ or powers). These are simplest to define when the ‘exponent’ in question is a natural number.

#### Definition 1.2

If  $a$  is a real number (called the *base*) and  $n \in \mathbb{N}$ , then

$$a^n = \underbrace{a \times a \times a \times \cdots \times a}_{n \text{ times}}$$

To illustrate, let  $a = 2$  and  $n = 5$ . Our definition then says that:

$$2^5 = 2 \times 2 \times 2 \times 2 \times 2 = 32$$

**Generalising the definition.** While this definition covers all natural exponents, it leaves various cases unresolved. For example, it does not specify the meaning of

numbers like (for example)  $2^0$ ,  $2^{-5}$ ,  $2^{5/4}$ , or  $2^\pi$ . However, there is a natural way to extend our definition so that it covers any real exponent.

To do this, consider a positive base ( $a > 0$ ) to avoid complications. We then proceed in several steps:

- To begin, define  $a^0 = 1$ . For example,  $2^0 = 1$  and  $5^0 = 1$ . This specifies what it means to raise a number to the power of zero.
- Next, for any  $n \in \mathbb{N}$ , define  $a^{-n} = 1/a^n$ . For example,

$$5^{-3} = \frac{1}{5^3} = \frac{1}{125}$$

This specifies what it means to raise a number to a negative (integer) power.

- Next, for any  $q \in \mathbb{N}$ , define  $a^{1/q}$  to be the unique positive number such that:

$$\underbrace{a^{\frac{1}{q}} \times a^{\frac{1}{q}} \times a^{\frac{1}{q}} \times \dots \times a^{\frac{1}{q}}}_{q \text{ times}} = a$$

This is also called the  $q$ th root of  $a$ , written  $\sqrt[q]{a}$ . For example,  $1000^{1/3}$  is the unique positive number that satisfies the equation

$$1000^{1/3} \times 1000^{1/3} \times 1000^{1/3} = 1000$$

Thus,  $1000^{1/3} = 10$ ; this is the ‘third root’ of 1000.

- Next, for any  $p \in \mathbb{Z}$  and  $q \in \mathbb{N}$ , define  $a^{p/q} = (a^{1/q})^p$ . For example,

$$100^{\frac{3}{2}} = (100^{\frac{1}{2}})^3 = 10^3 = 1000$$

This finally specifies what it means to raise a number to a fractional power.

While the previous steps allow us to extend our definition to any fractional (technically, ‘rational’) power, they do not precisely cover all the *real* powers. Thus, cases like  $5^\pi$  remain unresolved. However, any real number can be approximated by some fraction  $p/q$  where  $p \in \mathbb{Z}$  and  $q \in \mathbb{N}$ . Thus, it should come as no surprise that our definition can be extended to any real exponent, although we omit the details in this course.

The following result summarises the key properties of exponents.

#### Result 1.1: Properties of exponents

Let  $a, b > 0$  and  $r, s \in \mathbb{R}$ . Then

- (1)  $a^r a^s = a^{r+s}$
- (2)  $a^r / a^s = a^{r-s}$
- (3)  $(a^r)^s = a^{rs}$
- (4)  $(ab)^r = a^r b^r$
- (5)  $(a/b)^r = a^r / b^r$

It is reasonably easy to check that these properties hold for any natural exponents  $r, s$ . (You may wish to verify this as an exercise.) However, as can be seen, they also hold for any *real* exponents  $r, s$ . This is not a coincidence. Indeed, the *reason* why we have generalised the notion of natural exponents in the way that we have done above is precisely so that our notion of exponents continues to satisfy these properties.

#### Problem 1.2: Troubling doubling

Evaluate the following without computational assistance:

- (1)  $2^3$
- (2)  $2^0$
- (3)  $2^{-3}$
- (4)  $8^{1/3}$
- (5)  $8^{-1/3}$
- (6)  $4^{3/2}$
- (7)  $4^{-3/2}$

#### Problem 1.3: Properties of exponents

Evaluate the following without computational assistance:

- (1)  $2^{1000} \times 2^{-997}$
- (2)  $2^{1000}/2^{1003}$
- (3)  $(2^{1/1000})^{1000}$
- (4)  $(16 \times 81)^{3/4}$

### 1.3 Logarithms

As usual, we begin with a definition.

#### Definition 1.3

The *logarithm*  $\log_b x$  for a base  $b > 0$  (with  $b \neq 1$ ) and a number  $x > 0$  is the unique  $y$  such that  $b^y = x$ . That is,

$$\log_b x = y \iff b^y = x$$

This definition can be recast in still another (although almost equivalent) way. When confronted with the expression  $\log_b x$ , we ask ourselves: ‘ $b$  to the power of *what* is equal to  $x$ ?’ The answer to this question is precisely  $\log_b x$ . For example, when confronted with the expression  $\log_{10}(100)$ , we ask ourselves: ‘10 to the power of *what* is equal to 100?’ Since  $10^2 = 100$ , the answer is 2; and so  $\log_{10}(100) = 2$ .

The next result collects some important properties of logarithms. The first two results follow immediately from the definition of a logarithm. The subsequent three results

are less immediate and require a bit more work to establish. Nonetheless, it is vital that you memorise all of the rules below.

### Result 1.2: Properties of logarithms

Let  $b > 0$ ,  $b \neq 1$ , and  $x, y > 0$ . Then

- (1)  $\log_b(1) = 0$
- (2)  $\log_b(b) = 1$
- (3)  $\log_b(xy) = \log_b(x) + \log_b(y)$
- (4)  $\log_b(x/y) = \log_b(x) - \log_b(y)$
- (5)  $\log_b(x^r) = r \log_b(x)$  for any  $r \in \mathbb{R}$

We conclude this section with some problems. The first set of problems below tests your understanding of the concept of logarithms. The next set gives you the opportunity to apply the logarithm rules that we have just outlined. When solving the problems, make sure not to use a calculator — I promise that you don't need one!

### Problem 1.4: Basic logarithms

Evaluate the following without computational assistance:

- (1)  $\log_2(8)$
- (2)  $\log_2(1)$
- (3)  $\log_2(1/8)$
- (4)  $\log_3(9)$
- (5)  $\log_{10}(1000)$

### Problem 1.5: Properties of logarithms

Simplify the following expressions:

- (1)  $\log_{57}(1)$
- (2)  $\log_{57}(57)$
- (3)  $\log_5(25x)$
- (4)  $\log_5(100) - \log_5(4)$
- (5)  $\log_3(27^{5/2})$

## 2. SEQUENCES, SUMS AND LIMITS

### 2.1 Motivating examples

By the end of this lecture, you should be able to solve the following problems:

#### Problem 2.1: Family matters

Your eccentric relative offers you one of two options:

- Option 1: You will receive £100,000 in 20 years time.
- Option 2: Your relative will invest £40,000 in a savings account that pays 5% annual interest and give you the accumulated savings in 20 years time.

Which option would you choose?

#### Problem 2.2: Escalating spending

A company plans to increase its advertising spending. In the first week, it will spend £1,000. Each subsequent week, it plans to increase its spending by 1%. If the company follows this plan for a full year, how much will it spend on advertising in total?

### 2.2 Sequences

We begin with a (slightly simplified) definition.

#### Definition 2.1

A *sequence* is a list of numbers.

To illustrate, consider the sequences  $(1, 3)$ ,  $(8, 4, 8)$ , or  $(10, -1, 0, 5)$ . While there is nothing special about these sequences, one should note that they are *finite* in the sense that they have a finite number of elements. However, we can also consider *infinite* sequences that continue without end. For example, we can consider the sequence formed by listing all the natural numbers in ascending order:  $(1, 2, 3, \dots)$ .

Talking about sequences becomes easier once we introduce some notation:

- If a sequence has finitely many terms, say  $N$ , it can be written as  $(a_1, a_2, \dots, a_N)$ . Here,  $a_1$  is the first term,  $a_2$  is the second term,  $\dots$ , and  $a_N$  is the final term. More compactly, one can write the sequence as  $(a_n)_{n=1}^N$ .
- If a sequence is infinite (i.e. continues without end), it can be written as  $(a_1, a_2, \dots)$ . Here, the ' $\dots$ ' indicates that the sequence continues indefinitely. More compactly, one can write this as  $(a_n)_{n=1}^\infty$ .

- If we don't want to specify whether a sequence is finite or infinite, we will simply write  $(a_n)$ .

△ In the case of a finite sequence, do carefully distinguish between  $n$  (which plays the role of a 'running index' or 'dummy variable') and  $N$  (the final index in the sequence).

There are different ways in which one can define a sequence:

- In some instances, a sequence follows a particular pattern and can be *specified using a formula*. For example, one may specify that  $a_n = 2n$  (for every natural number  $n$ ), which then gives rise to the sequence  $(2, 4, 6, \dots)$ .
- One can also *define sequences recursively*. That is, one can define a sequence by specifying
  - The first term of the sequence  $a_1$  (this corresponds to the 'initial condition' in many economic models).
  - A rule that specifies how every subsequent term can be computed using the term just before it; this rule is sometimes called a *recurrence relation*.

For example, one can specify that  $a_1 = 0$  (the 'initial condition') and that  $a_n = a_{n-1} + 1$  for  $n \geq 2$  (the 'recurrence relation'). This then recursively defines the sequence  $(0, 1, 2, 3, \dots)$ .

Some examples may clarify.

#### Problem 2.3: Notation

How many terms does the sequence  $(a_n)_{n=1}^7$  have?

#### Problem 2.4: A formula

Consider the (infinite) sequence defined by the formula  $a_n = 1/n$  for every natural number  $n$ . What is the fourth term  $a_4$ ?

#### Problem 2.5: A recursion

Consider the sequence defined by  $a_1 = 1$  and the rule that  $a_n = a_{n-1} + 3$  for every natural number  $n \geq 2$ . What is the third term  $a_3$ ?

#### Problem 2.6: A philosophical investigation

Find the next term in the sequence 1, 5, 11, 19, 29, ...?

Although an infinite sequence never ends, it may get closer and closer to a particular value. If so, this value is called the sequence's *limit*. A bit more precisely:

### Definition 2.2

Suppose that, for every notion of ‘close’ that one specifies, one can find that a point in a sequence  $(a_n)_{n=1}^{\infty}$  such that, from that point onwards, all terms in the sequence are ‘close’ to  $L$ . Then  $L$  is the *limit* of the sequence.

To explain this even formally, suppose that a sequence has a limit  $L$ . This means that, for any  $\epsilon > 0$  that you pick, there is a point in the sequence such that, from that point onwards, all terms in the sequence are within an  $\epsilon$  of  $L$ . This is close to the formal definition of a limit and is included in case it further illuminates the concept. However, you don’t need to know it for your exams.

In such a scenario, i.e. when  $L$  is the limit of the sequence  $(a_n)_{n=1}^{\infty}$ , one can also write:

- $(a_n)_{n=1}^{\infty}$  converges to  $L$ .
- $a_n \rightarrow L$  as  $n \rightarrow \infty$ .
- $\lim_{n \rightarrow \infty} a_n = L$ .

These statements all mean the same thing. More important than notation, however, is the observation that, while some sequences have limits, other sequences do not. For example, a sequence may increase without bound, decrease without bound, or ‘cycle’ between different values without ever ‘settling down’. In such circumstances, one can loudly declare that ‘the limit does not exist’. (This point will be readily appreciated by those who have seen the exciting finale of the 2004 classic ‘Mean Girls’.)

Again, some examples can help illustrate.

### Problem 2.7: Limits

Consider the infinite sequences defined by

- (1)  $a_n = 1/n$
- (2)  $a_n = (n + 2n^2)/n^2$
- (3)  $a_n = 3$
- (4)  $a_n = 2n$
- (5)  $a_n = (-1)^n$

In every case, write out the first few terms of the sequence, determine whether the sequence has a limit, and (if so) find this limit.

Before ending our discussion of sequences, it is worth defining two ‘types’ of sequence that frequently arise in economic models. We start by considering *arithmetic sequences*, which are sequences where each new term is obtained by adding a constant number to the preceding term. More formally,



### Definition 2.3

A sequence  $(a_n)$  is *arithmetic* when there exists a ‘common difference’  $d$  such that  $a_n = a_{n-1} + d$  for all  $n \geq 2$ .

To illustrate, suppose that  $a_1 = 10$  and  $d = 2$ . To get the second term, we simply add 2 to the initial value; thus,  $a_2 = 12$ . To get the third term, we add 2 again; thus,  $a_3 = 14$ . This process continues recursively, yielding the sequence  $(10, 12, 14, 16, \dots)$ .

Another common type of sequence within economics is the *geometric* sequence. This is simply a sequence where each new term is obtained by multiplying the preceding term by a ‘common ratio’. More formally,

### Definition 2.4

A sequence  $(a_n)$  is *geometric* when there exists a ‘common ratio’  $r$  such that  $a_n = r a_{n-1}$  for all  $n \geq 2$ .

To illustrate, suppose that  $a_1 = 1$  and  $r = 2$ . To get the second term, we multiply the first term by  $r$  (i.e., we double it); this yields  $a_2 = 2$ . To get the third term, we double the second term, yielding  $a_3 = 4$ . Continuing in this fashion, we get the sequence  $(1, 2, 4, 8, 16, \dots)$ .

The next set of problems tests your understanding of arithmetic and geometric sequences.

### Problem 2.8: Checking sequences

Determine whether each of the following sequences is arithmetic, geometric, neither, or both:

- (1)  $(-1, -5, -9)$
- (2)  $(2, 6, 18, 54)$
- (3)  $(17, 17)$
- (4)  $(-2, 4, -10)$

We have defined both arithmetic and geometric sequences using a rule that computes each successive term in the sequence from the previous term (i.e., a recurrence relation). For example, we defined arithmetic sequences using the rule  $a_n = a_{n-1} + d$ . Interestingly, however, these recurrence relations turn out to be logically equivalent to rather simple formulae. Specifically, one can show that:

### Result 2.1: The arithmetic formula

A sequence  $(a_n)$  is arithmetic with common difference  $d$  if and only if every term in the sequence satisfies the formula  $a_n = a_1 + (n - 1)d$ .

### Result 2.2: The geometric formula

A sequence  $(a_n)$  is geometric with common ratio  $r$  if and only if every term in the sequence satisfies the formula  $a_n = a_1 r^{n-1}$ .

To prove these results, one can use ‘proof by induction’. However, these results can also be grasped more intuitively. For example, take the case of a geometric sequence. By the recurrence relation,  $a_2 = a_1 r$ ; i.e., to get  $a_2$ , we multiply the initial term by  $r$  once. Likewise,  $a_3 = a_2 r = a_1 r^2$ ; i.e., to get  $a_3$ , we multiply the initial term by  $r$  twice. Continuing in this fashion, one ‘sees’ that, to obtain the  $n$ th term  $a_n$ , one multiplies the initial term  $a_1$  by  $r$  in total  $(n - 1)$  times.

### Problem 2.9: The arithmetic formula

Consider the arithmetic sequence defined by  $a_1 = 4$  and  $d = 0.5$ . Find a formula for the  $n$ th term of this sequence.

### Problem 2.10: The geometric formula

Consider the geometric sequence defined by  $a_1 = 0.5$  and  $r = 3$ . Find a formula for the  $n$ th term of this sequence.

**Convergence.** As can be seen from the previous formulae, an arithmetic sequence almost never converges to a fixed value. If  $d > 0$ , it increases without bound; whereas if  $d < 0$ , it decreases without bound. (Technically, the sequence does have a limit, namely  $a_1$ , in the special case of  $d = 0$ ). In contrast, a geometric sequence *can* converge. As it turns out, this depends almost entirely on its common ratio:

### Result 2.3: Geometric convergence

A geometric sequence  $(a_n)_{n=1}^{\infty}$  with initial value  $a_1 \neq 0$  and common ratio  $r$  converges if and only if  $r \in (-1, 1]$ .

⚠ Notice that the sequence does *not* converge when  $r = -1$ . Indeed, in such a case, the sequence oscillates between  $a_1$  and  $-a_1$  without ever ‘settling down’.

The next problem tests your understanding of the geometric convergence result.

### Problem 2.11: Geometric convergence

Consider the geometric sequence given by  $a_1 = 10$  and  $r = -0.5$ . Write out the first few terms of the sequence. Does the sequence converge?

## 2.3 Sums

Having considered sequences, we now consider *sums* of sequences. These are defined in exactly the way in which you would expect.

### Definition 2.5: Sum of a sequence

Given a sequence  $(a_n)_{n=1}^N$ , the *sum of the sequence* is

$$\sum_{n=1}^N a_n = a_1 + a_2 + \dots + a_N$$

The *summation operator*  $\Sigma$  provides us with a way to describe sums compactly. Here is an example to check that you understand what it means.

### Problem 2.12: The summation operator

Evaluate  $\sum_{n=1}^3 (2n + 1)$ .

As it turns out, there is a remarkably simple way to add up the terms in an arithmetic sequence. Following our notation, the first term of the sequence is  $a_1$  and the last term is  $a_N$ . Since the sequence increases (or decreases) at a constant rate, this means that the ‘average term’ is  $(a_1 + a_N)/2$ . Meanwhile, the total number of terms is  $N$ . From this, one can more or less deduce the following.

### Result 2.4: The sum of an arithmetic sequence

If  $(a_n)_{n=1}^N$  is an arithmetic sequence with common difference  $d$ , then

$$\sum_{n=1}^N a_n = N \left( \frac{a_1 + a_N}{2} \right)$$

Here is a problem to check that you can understand the formula.

### Problem 2.13: The young Gauss

Without computational assistance, add up all the integers from 1 to 100.

What about summing a geometric sequence? It turns out that there is a simple formula for that too. If there is time, we may derive the formula in the lectures.

**Result 2.5: The sum of a geometric sequence**

If  $(a_n)_{n=1}^N$  is a geometric sequence with common ratio  $r \neq 1$ , then

$$\sum_{n=1}^N a_n = \frac{a_1(1 - r^N)}{1 - r}$$

The next problem illustrates the formula.

**Problem 2.14: Saving up**

Rachel decides to put £1,000 in her cash ISA at the start of every year. Her savings earn 5% annual interest, paid on the last day of each year. How much money will she have in her ISA in 20 years (just before she makes her 21st deposit)?

The formula above describes the sum of a finite geometric sequence. But what about an infinite geometric sequence? To put the question a bit more rigorously, what happens when we consider the sum of the first  $N$  terms of an infinite geometric sequence and let  $N \rightarrow \infty$ ?

The formula says that, if  $(a_n)_{n=1}^\infty$  is a geometric sequence (with  $r \neq 1$ ), then the sum of the first  $N$  terms is given by

$$\sum_{n=1}^N a_n = \frac{a_1(1 - r^N)}{1 - r}$$

If  $|r| < 1$  (say,  $r = 0.5$ ), then  $r^N \rightarrow 0$  as  $N \rightarrow \infty$ . (For example, if  $r = 0.5$ , the sequence is  $0.5, 0.25, 0.125, \dots$ ) Thus, the formula converges to

$$\lim_{N \rightarrow \infty} \sum_{n=1}^N a_n = \frac{a_1(1 - 0)}{1 - r} = \frac{a_1}{1 - r}$$

In summary, then, we see that:

**Result 2.6: The sum of an infinite geometric sequence**

If  $(a_n)_{n=1}^\infty$  is a geometric sequence with common ratio  $|r| < 1$ , then

$$\lim_{N \rightarrow \infty} \sum_{n=1}^N a_n = \frac{a_1}{1 - r}$$

The next ‘problem’ illustrates this result rather nicely.

Problem 2.15:  $1 + 1/2 + 1/4 + \dots$

Explain this joke: *An infinite number of mathematicians walk into a bar. The first one orders a pint, the second orders a half pint, and the third orders a quarter pint. Frustrated, the bartender pours two pints and slams them on the table. 'If you're going to drink', he tells his guests, 'then you should know your limits!'*

### 3. SOLVING EQUATIONS

#### 3.1 Motivating examples

##### Problem 3.1: Research

A researcher spends all of her salary on instant coffee (specifically, Nescafé Azera) and milk. She has calculated that each pot of Azera requires three pints of milk. A pot of Azera costs £7, and a pint of milk costs £1. If her monthly salary is £100, how many pots of Azera will she buy?

##### Problem 3.2: Event planning

A group of students are organising a public lecture. They plan to spend £10 on snacks for each of their guests; their invited speaker (an economist) has also demanded a £75 ‘show-up fee’. The students expect that, if they charge a ticket price of £ $p$ , with  $p \in \{0, 1, \dots, 30\}$ , then they will get  $30 - p$  attendees. At what ticket price(s) should they expect to break even?

#### 3.2 Solving equations

Solving equations plays a large role in economic research. But what does it mean to ‘solve’ an equation? For example, what does it mean to ‘solve’ the equation

$$2^x + \log_2(x) + x = 0?$$

We will say that:

##### Definition 3.1: Solving an equation

To *solve* an equation means to identify all the possible values (if any) that make the equation true.

Thus, ‘solving’ the equation  $2^x + \log_2(x) + x = 0$  means identifying *all* the values of  $x$ , if there are any, that make the equation hold true. Note that, in general, there may be several values of  $x$  that satisfy the equation (‘multiple solutions’), exactly one such value of  $x$  (a ‘unique solution’), or no such value of  $x$  (‘no solution’). The following example emphasises these points.

### Problem 3.3: Uniqueness and existence

Solve the following equations

- (1)  $x^2 = 0$
- (2)  $x^2 + 1 = 0$
- (3)  $x^2 - 1 = 0$

There are several ways to go about solving an equation. In particular, one can:

- Make educated guesses and check if they satisfy the statement in question.
- Guess and check possible values in a more structured and systematic way (*numerical* methods).
- Re-write (or ‘manipulate’) the equation into a simpler form (an example of *analytic* methods).

Although analytic methods will be emphasised in this course, numerical methods can be of great value when solving more complex equations (or ‘systems’ of equations).

### 3.3 Linear equations

We will start by considering perhaps the simplest type of equation: linear equations. In such equations, the unknown variables (e.g.,  $x$ ) are not raised to any powers greater than 1: thus, there are no  $x^2$ ,  $x^3$ , or  $x^4$  terms. The next definition spells out the meaning of a linear equation (in a single variable) with a little more precision.

#### Definition 3.2

A *linear equation in one variable* is an equation that can be written in the form  $ax = b$  where  $a \neq 0$ .

For example, consider the equations  $2x = 4$ ,  $x = -9$ , or  $-3x = 7.5$ . As the definition indicates, we also include equations that, although written in the form  $ax = b$ , can be *re-written* in this way. For example, the equation  $7x + x + 7 = 2x$  is a linear equation in one variable since it is logically equivalent to the equation  $6x = -7$ .

The next result is an obvious implication of the assertion that  $ax = b$  under the restriction that  $a \neq 0$ .

#### Result 3.1

A linear equation in one variable has exactly one solution, namely

$$x = \frac{b}{a}$$

Although obvious, this result does hide one bit of nuance. The second example in the following problem attempts to draw this nuance out.

#### Problem 3.4

Solve the equations

- (1)  $2x = 4$
- (2)  $x + 1 = x$

Somewhat more interesting is the case of *two* linear equations in (at most) *two* variables. The next definition gives this a more precise definition.

#### Definition 3.3

A *system of two linear equations in two variables* is a set of equations that can be written in the form

$$a_1x + b_1y = c_1$$

$$a_2x + b_2y = c_2$$

where  $(a_1, b_1) \neq (0, 0)$  and  $(a_2, b_2) \neq (0, 0)$ .

To illustrate, consider the system of equations

$$2x + 4y = 10$$

$$-x + 10y = -2$$

Note that, as before, we allow for systems that, although not written in this exact form, can be *rewritten* so that they appear in the form  $a_1x + b_1y = c_1$ ,  $a_2x + b_2y = c_2$ .

Given any such system, there exists a simple procedure that can be used to (i) determine the number of solutions that the system has, (ii) compute the solutions (if there are any). The following result describes this procedure in detail.

#### Result 3.2

Consider a system of two linear equations in two variables and define

$$\Delta = a_1b_2 - a_2b_1, \quad \Delta_x = c_1b_2 - c_2b_1, \quad \Delta_y = a_1c_2 - a_2c_1$$


Then:

- If  $\Delta \neq 0$ , the system has exactly one solution, namely  $x = \Delta_x/\Delta$  and  $y = \Delta_y/\Delta$ .
- If  $\Delta = 0$  and  $\Delta_x = \Delta_y = 0$ , the system is satisfied by any  $(x, y)$  pair that satisfies either of the equations individually and has infinitely many solutions.
- If  $\Delta = 0$  and either  $\Delta_x \neq 0$  or  $\Delta_y \neq 0$ , the system has no solution.




The basic idea of this result can also be described graphically. It turns out that any such system of two linear equations can be visualised as two straight lines in  $(x, y)$  space. The first equation says that any solution must lie on the first line. The second equation says that any solution must lie on the second line. There are then three possibilities (see also the figure below):


- One possibility is that the lines have different slopes. If so, the lines will cross exactly once; this is the unique solution to the system.
- Another possibility is that the lines have the same slope and are in fact the *very same line*. In this case, the system is solved by *any*  $(x, y)$  pair that lies on the line; thus, the system has an infinite number of solutions.
- A final possibility is that the lines have the same slope but are different lines. In that case, no point can lie on both lines simultaneously, so the system has no solution.



One solution



Infinitely many solutions



No solution

The next problem gives you the opportunity to solve a system using this procedure.

### Problem 3.5

Consider the linear equations  $x + 2y = 3$ ,  $2x + y = 3$ . Compute  $\Delta$ ,  $\Delta_x$ ,  $\Delta_y$  and thus solve the system.

The previous result provides a complete solution to any two equation, two variable linear system. However, we will now describe some *additional* methods of solving the system. This might seem rather pointless – don't we know how to solve such systems already? These new methods are useful, however, since they (and their variants) extend systems with more equations and more variables.

We begin by describing the substitution method. This is perhaps easiest to explain in the two equation, two variable case:

- First, solve one of the equations for one variable. For concreteness, let's suppose that we solve one of the equations for  $y$ , yielding an equation of the form  $y = \dots$

- Next, substitute this expression for  $y$  into the other equation. This *eliminates*  $y$  from the other equation and yields a single equation solely in  $x$ .
- One then solves the resulting equation for  $x$ .
- Finally, one plugs the solution for  $x$  back into either original equation to determine the solution for  $y$ .

More generally, the substitution method works by ‘isolating’ a variable (say  $y$ ) and then plugging this expression into the rest of the system, thus eliminating all mention of  $y$  from the rest of the system. One then *repeats* the process, each time eliminating a new variable from the system. Eventually, one ends up with an equation in a single variable, which one then solves. Using this solution, one then recursively ‘backs out’ the values of the other variables.

The following problem is an opportunity to practice this procedure.

#### Problem 3.6: The substitution method

Use the substitution method to solve the system

$$x + y = 2, \quad x - y = 0$$

Next, we describe the *elimination method*. Here is how it works in the two equation, two variable case:

- First, choose one of the variables to eliminate. For concreteness, let’s suppose that we want to eliminate  $y$ .
- Then multiply one of the equations by a suitable constant so that the coefficients on that variable ( $y$ ) are equal across the two equations.
- Subtract one equation from the other to eliminate the variable, yielding a single equation in one variable ( $x$ ).
- Solve the resulting equation for that variable ( $x$ ).
- Finally, substitute back to determine the value of the other variable ( $y$ ).

The following problem asks you to put this method into practice.

#### Problem 3.7: The elimination method

Use the elimination method to solve the system

$$x + y = 2, \quad x - y = 0$$

As a review, you may now wish to solve the following systems using *all* of the methods that we have discussed so far. That is, you can use the general theory (i.e. compute  $\Delta$ ,

$\Delta_x, \Delta_y$ ), the substitution method, and the elimination method. However, the really important point is that you can solve the following systems using *some* method.

### Problem 3.8: A review

Solve the following systems of equations:

- (1)  $x + y = 5, \quad x - y = 1$
- (2)  $y = 1 - x, \quad x + y = 0$
- (3)  $2x = 6 - 2y, \quad x + y = 3$

To wrap up our discussion of linear systems, we briefly consider the general case of  $m$  equations and  $n$  unknowns.

### Definition 3.4

A *system of linear equations* is a set of equations that can be written in the form

$$a_{11}x_1 + a_{12}x_2 + \dots + a_{1n}x_n = c_1$$

$$a_{21}x_1 + a_{22}x_2 + \dots + a_{2n}x_n = c_2$$

...

$$a_{m1}x_1 + a_{m2}x_2 + \dots + a_{mn}x_n = c_m$$

In general, what can we say about the solution to such a system? In fact, general linear systems turn out to have a very simple structure.

### Result 3.3: Linear systems (solutions)

A system of linear equations has either

- No solution.
- Exactly one solution.
- Infinitely many solutions.

The intuition for this result is similar to the intuition in the two variable, two equation case: a pair of lines will either cross exactly once, never (if the two lines are different but parallel), or at infinitely many points (if the two lines are the same).

Although this result tells us something about the solutions of linear systems, it does not tell us how to calculate them. Fortunately, this is in some sense a ‘solved problem’, although the solution to this problem goes beyond the scope of this course. Specifically:

### Result 3.4: Linear systems (algorithms)

There are known algorithms that are guaranteed to solve any system of linear equations in a finite number of steps.

It should be noted that these algorithms are not only guaranteed to solve a system eventually but are also rather ‘fast’. For example, although a system with 1000 equations and 1000 unknowns may seem quite forbidding, these algorithms can solve such systems in well under a second when implemented on a modern laptop.

## 3.4 Non-linear equations

So far, we have just discussed *linear* systems of equations. In such systems, any unknowns (e.g., ‘ $x$ ’) only appear to the power of 1. One can wonder, however, what happens when unknowns are raised to higher powers (e.g.,  $x^2$ ,  $x^3$ , ...). In such cases, the resulting equations become *non-linear*.

The simplest non-linear equation arises when an unknown is raised to the power of 2; such an equation is termed ‘quadratic’. More formally,

### Definition 3.5

A *quadratic equation in one variable* is an equation that can be written in the form  $ax^2 + bx + c = 0$  where  $a \neq 0$ .

To illustrate, consider the equation  $4x^2 + 2x + 3 = 0$ ; in this example,  $a = 4$ ,  $b = 2$ , and  $c = 3$ . One could also consider an equation like  $7x^2 = 4$  since this equation can be rewritten in the quadratic form (with  $a = 7$ ,  $b = 0$ , and  $c = -4$ ).

Much like linear systems, quadratic equations in a single variable are very much a ‘solved problem’. Specifically:

### Result 3.5: The quadratic formula

Consider a quadratic equation in one variable and define  $\Delta = b^2 - 4ac$ . Then:

- If  $\Delta < 0$ , the equation has no solution.
- If  $\Delta = 0$ , the equation has exactly one solution, namely

$$x = -\frac{b}{2a}$$

- If  $\Delta > 0$ , the equation has exactly two solutions, namely

$$x = \frac{-b \pm \sqrt{\Delta}}{2a}$$

As indicated in the result above, computing  $\Delta$  when faced with a quadratic equation allows one to determine the number of (real) solutions that the equation has. The next example encourages you to put this procedure into practice.

**Problem 3.9: The quadratic formula**

Solve the following quadratic equations:

(1)  $x^2 + 2x + 5 = 0$

(2)  $x^2 + 2x + 1 = 0$

(3)  $x^2 + 2x - 3 = 0$

As we have seen, quadratic equations present no special difficulties despite their non-linearity. Unfortunately, however, solving non-linear equations (or systems of non-linear equations) is not always this easy. In such situations, one may want to reach for *numerical* methods to deliver approximate solutions. For example, although the equation  $x^5 - x - 1 = 0$  cannot be solved using analytic methods, a numerical analysis suggests the approximate solution  $x \approx 1.1673$ . Moreover, using an analytical argument, one can establish that this solution is unique.

## 4. FUNCTIONS

### 4.1 Motivating examples

#### Problem 4.1: Average costs

A company produces a quantity of a good  $q \geq 0$  at the total cost of  $500 + 40q - 4q^2 + 0.2q^3$ . Find the average cost of producing  $q > 0$  units. What happens to the average cost when  $q \rightarrow 0$  and  $q \rightarrow \infty$ ?

#### Problem 4.2: The Laffer curve

Arthur believes that, if the government taxes a share  $t \in [0, 1]$  of their citizens' income, the amount of tax revenue they will raise will be proportional to  $R(t) = t(1 - t)$ . (i) Compute  $R(0)$  and  $R(1)$ . (ii) Sketch  $R(t)$  from  $t = 0$  to  $t = 1$ . Are Arthur's beliefs reasonable?

### 4.2 Introduction

Economists are often interested in understanding the relationship between two different variables. For example, they may want to understand the relationship between the money supply and the price level, the relationship between the level of taxation and economic performance, or the relationship between interest rates and inflation. *Functions* are useful since they provide a precise way to describe and thus study these relationships.

We begin by defining the notion of a function.

#### Definition 4.1


A *function*  $f: D \rightarrow C$  is a rule that associates every element of  $D$  with exactly one element in  $C$ .

In the definition above, the set  $D$  is called the *domain* and the set  $C$  is called the *codomain*. One can think of  $D$  as the 'set of inputs' and  $C$  as a 'set of possible outputs'. Thus, a function can be thought of taking a particular 'input' (in  $D$ ) and associating it with (or 'mapping it to') a particular 'output' (in  $C$ ).

It is important to appreciate that, although every 'output' *must* belong to the codomain  $C$ , there is no requirement that every element in  $C$  is an 'output'. In other words,  $C$  can contain elements that are never associated with elements in the domain  $D$ . If one *does* want to focus on the *actual* outputs of the function, as opposed to the 'potential outputs', one should instead consider the function's *range*.<sup>1</sup>

<sup>1</sup> More formally, a function's range is the set  $\{y: y = f(x) \text{ for some } x \in D\}$ .

While these comments may seem abstract, they can perhaps be illuminated by the figure below. In the figure, one sees the domain  $D$  represented by the black circle. This is the ‘set of inputs’ and contains 4 elements, which are represented by the small ‘dots’. One can also see the codomain  $C$ , which is represented by the red circle. This is the ‘set of possible outputs’ and contains 5 elements, again represented by the small ‘dots’. Finally, one sees a function  $f: D \rightarrow C$  that maps each input in  $D$  to an output in  $C$ . This mapping is represented using arrows.



Several points should be emphasised:

- First, every element in  $D$  is mapped to some output. This is ‘mandatory’ for any function defined with  $D$  as its domain.
- Second, every element in  $D$  is mapped to *exactly one* output. Again, this is an integral part of the definition of a function.
- Finally, in this example, there are two elements in  $C$  that are not associated with any element in  $D$ . (These are the two ‘isolated dots’.) This is unproblematic since the codomain is just a set of possible outputs, not the set of actual outputs. In contrast, those two elements would *not* belong to the function’s range.

At this point, it may be useful to introduce some further notation. Throughout this lecture, we will write  $x$  to describe a typical element in the domain of  $f$  and  $f(x)$  to describe its corresponding output. We can then view a function as a map

$$x \mapsto f(x)$$

Before getting to some examples, we make a final comment about dimensionality. The previous definition is very general and allows for all kinds of domains  $D$ . For example, the elements of  $D$  could be pairs of numbers, triples of numbers, etc. In this lecture, however, we will keep things simple by considering functions that take a single variable as their input. Such functions are called *univariate* (‘one’ ‘variable’) and have one-dimensional domains: for example, the domain might be  $\mathbb{N}$ ,  $\mathbb{Z}$ , or  $\mathbb{R}$ .

### 4.3 Examples

We now consider some functions that arise frequently in economics. In all the examples that follow, the function's domain is either the full real line ( $\mathbb{R}$ ) or its strictly positive subset, i.e.  $(0, \infty)$ . While the examples that follow are important to understand, it should be emphasised that they certainly do not exhaust the large universe of possible functions that one could study.


We begin with perhaps the simplest type of function, namely the 'monomial'. This is defined as follows.

#### Definition 4.2

A function  $f: \mathbb{R} \rightarrow \mathbb{R}$  is *monomial* if it can be written in the form  $f(x) = ax^k$  for some constant  $a \in \mathbb{R}$  and power  $k \in \mathbb{N}$ .

For example, one could consider the function  $f(x) = x$  (letting  $a = 1$  and  $k = 1$ ). This is called the identity function and simply returns whatever number is plugged into it: for example, the number 4 returns 4 as its output. To take a second example, one could consider the function  $f(x) = 3x^2$  (here,  $a = 3$  and  $k = 2$ ). This function is *quadratic* since it has two as its highest power.

The following figure illustrates monomial functions in the case where  $a = 2$  and  $k = 1$  (that is,  $f(x) = 2x$ ). To plot the function, one can start by computing the 'y-intercept', i.e.  $f(0) = 2 \times 0 = 0$ . One could then compute some further points, e.g.  $f(1) = 2$ ,  $f(2) = 4$ , and  $f(3) = 6$ . At this point, one can hopefully convince oneself that this function is *linear* and can be represented using a straight line. Indeed, this is a general property of functions of the form  $f(x) = ax$  where  $a \neq 0$ .



We now generalise the notion of a monomial function.




### Definition 4.3

A function  $f: \mathbb{R} \rightarrow \mathbb{R}$  is *polynomial* if it can be written as the sum of monomial functions. That is,

$$f(x) = a_0 + a_1x + a_2x^2 + \dots + a_nx^n$$

To understand this, suppose that all  $a_i$  coefficients are zero except for  $a_0$ . One then obtains the *constant function*  $f(x) = a_0$ : graphically, this can be viewed as a horizontal line. Next, suppose that all  $a_i$  coefficients are zero except for  $a_0$  and  $a_1$ . One then obtains the *linear function*  $f(x) = a_0 + a_1x$ : graphically, this can be viewed as a straight line with ‘ $y$ -intercept’  $a_0$  and slope  $a_1$ . Next, suppose that all  $a_i$  coefficients are zero except for (say)  $a_0$  and  $a_2$ . One then obtains the *quadratic function*  $f(x) = a_0 + a_2x^2$ . More generally, a polynomial is the sum of an arbitrary collection of monomials. The highest power that appears in a polynomial is called the polynomial’s *degree*.

The next figure illustrates this notion by considering the case  $f(x) = 1 + x + x^2$ . In this example,  $a_0 = a_1 = a_2 = 1$ ; all other  $a_i$  coefficients are zero. The resulting function is quadratic and has the quadratic function’s characteristic shape. One can begin verifying the figure’s correctness by computing the ‘ $y$ -intercept’  $f(0) = 1 + 0 = 1$ .



Next, we consider exponential functions.

### Definition 4.4

A function  $f: \mathbb{R} \rightarrow (0, \infty)$  is *exponential* if it can be written in the form


$$f(x) = ab^x,$$

where  $a > 0$ ,  $b > 0$ , and  $b \neq 1$ .

As can be seen, these functions are *not* polynomial since the  $x$  appears as the function’s

*exponent*, not as the function's base. For example, the function  $f(x) = 2^x$  is exponential (but not polynomial) since 2 is raised to the power of  $x$ . In contrast, the function  $f(x) = x^2$  is polynomial (but not exponential).

There are several important points to note about this family of functions. First, the function can describe either *exponential growth* (when  $b > 1$ ) or *exponential decay* (when  $0 < b < 1$ ). Second, the coefficient  $a$  is simply a scaling factor that vertically stretches the function's graph. Finally, if we set the base  $b$  equal to  $e \approx 2.71828$ , we obtain what is known as *the* exponential function. This function plays an important role in economics and is plotted below.



As our final example, we consider a class of functions that are (in a very literal sense) the ‘mirror image’ of exponential functions. These are logarithmic functions and are defined below.


#### Definition 4.5

A function  $f: (0, \infty) \rightarrow \mathbb{R}$  is *logarithmic* if it can be written in the form

$$f(x) = a \log_b(x),$$

where  $a \neq 0$ ,  $b > 0$ , and  $b \neq 1$ .

For any  $x > 0$ , this function simply returns the logarithm of  $x$  given a base  $b > 0$  (with  $b \neq 1$ ). Again, several points should be noted. First, the function increases when  $b > 1$  but decreases when  $0 < b < 1$ . Second, the coefficient  $a$  is again a scaling factor that stretches the function's graph. Finally, when  $b = e \approx 2.71828$ , we obtain the *natural logarithm*, written  $f(x) = a \ln(x)$ . This function is plotted in the next figure.



The problems that follow test your understanding of the material that you have encountered so far. In particular, they test your understanding of the abstract definition of a function in addition to your ability to sketch the graph of particular functions.

#### Problem 4.3: A finite domain

Let  $D = \{0, 4, 5\}$  and  $C = \{0, 1, 2, 3, 6\}$ . Consider the function  $f: D \rightarrow C$  that sets  $f(0) = 1$ ,  $f(4) = 1$ , and  $f(5) = 0$ . How many elements are in the function's domain? How many elements are in the function's range?

#### Problem 4.4: A linear function

Consider the function  $f(x) = 2 - 3x$ . (i) Compute  $f(-2)$ ,  $f(0)$ , and  $f(2)$ . (ii) Sketch the graph of the function.

#### Problem 4.5: The natural logarithm

Consider the function  $f(x) = 3\ln(x)$ . (i) Compute  $f(1)$  and  $f(e)$ . (ii) Sketch the graph of the function.

## 4.4 Key concepts

We now discuss some core topics in the theory of functions.

**Composing functions.** Consider a function  $f$  that maps  $x$  to  $f(x)$ . Although  $f(x)$  is the 'output' of this function, one can also use this as the 'input' for another function  $g$ . If one does this, one ends up with  $g(f(x))$ : the so-called 'composition' of the two functions. More formally:

#### Definition 4.6

Let  $f: D \rightarrow C$  and  $g: C \rightarrow E$  be functions. The *composition* of  $g$  and  $f$ , denoted  $g \circ f: D \rightarrow E$ , is the function that returns  $g(f(x))$  given any element  $x$  in  $D$ .

Of course, one can continue composing functions for as long as one likes. For example, one can compute

$$x \mapsto f(x) \mapsto g(f(x)) \mapsto h(g(f(x))) \mapsto \dots$$

The following example checks your ability to compose two functions, each defined with finite domains.

#### Problem 4.6: Composing functions with a finite domain

Let  $D = \{1, 2, 3\}$ ,  $C = \{a, b\}$ , and  $E = \{x, y\}$ . Define  $f: D \rightarrow C$  with

$$f(1) = a, f(2) = b, f(3) = b,$$

and  $g: C \rightarrow E$  with

$$g(a) = y, g(b) = x.$$

Compute  $(g \circ f)(1)$ ,  $(g \circ f)(2)$ , and  $(g \circ f)(3)$ .

**Function inverses.** Some functions have a corresponding ‘inverse function’ that ‘undoes’ their effect. For example, the function  $f(x) = 2x$  has the inverse of  $g(x) = x/2$ : if one multiplies a number by 2 (applying  $f$ ), but then divides the resulting number by 2 (applying  $g$ ), one ends up exactly where one started. More formally:

#### Definition 4.7

Let  $f$  be a function with domain  $D$  and range  $R$ . If  $g: R \rightarrow D$  is the *inverse* of  $f$ , then  $g(f(x)) = x$  for all  $x \in D$ .

Importantly, not all functions have inverses. Indeed, a function has an inverse if and only if different elements in  $D$  are mapped to different elements in  $C$ . To put this slightly differently, a function is invertible if and only if no two ‘inputs’ generate the same ‘output’. We call such functions *one-to-one*.

The next problem tests your understanding of invertibility.

#### Problem 4.7: Invertibility

Let  $D = \{1, 2, 3\}$  and  $C = \{a, b, c\}$ . Define  $f: D \rightarrow C$  with

$$f(1) = a, f(2) = b, f(3) = b$$

Does  $f$  have an inverse? Would this change if  $f(3) = c$ ?

**Function limits.** The limit of a function as  $x$  becomes large is defined in a very similar way to the limit of a sequence. Specifically:

- If  $f(x)$  becomes arbitrarily close to  $L$  as  $x \rightarrow \infty$ ,  $\lim_{x \rightarrow \infty} f(x) = L$ .
- If  $f(x)$  grows without bound as  $x \rightarrow \infty$ ,  $\lim_{x \rightarrow \infty} f(x) = \infty$ .
- If  $f(x)$  decreases without bound as  $x \rightarrow \infty$ ,  $\lim_{x \rightarrow \infty} f(x) = -\infty$ .

(Limits in the  $x \rightarrow -\infty$  case are handled exactly analogously.)

One can also define the limit of a function as  $x$  approaches a finite value  $a$ , which is either contained in the domain  $D$  or is ‘just outside’ of  $D$ . (For example, if the domain is  $(0, 1)$ , the value  $a$  could be any value in  $[0, 1]$ .) Unfortunately, one cannot define such limits in exactly the same way in which one defines the limits of a sequence. Roughly speaking, however, one can say that:

- If  $f(x)$  becomes arbitrarily close to  $L$  as  $x \rightarrow a$ ,  $\lim_{x \rightarrow a} f(x) = L$ .
- If  $f(x)$  grows without bound as  $x \rightarrow a$ ,  $\lim_{x \rightarrow a} f(x) = \infty$ .
- If  $f(x)$  decreases without bound as  $x \rightarrow a$ ,  $\lim_{x \rightarrow a} f(x) = -\infty$ .

The following examples will give the opportunity to calculate some limits. Notice that, in the final example, the value  $a$  is (just) outside of the function’s domain.

#### Problem 4.8: Function limits

Let  $f(x) = 1/x$  with domain  $D = (0, \infty)$ . Compute:

- (1)  $\lim_{x \rightarrow \infty} f(x)$
- (2)  $\lim_{x \rightarrow 2} f(x)$
- (3)  $\lim_{x \rightarrow 0} f(x)$

**Continuity.** An important class of functions are the *continuous* functions. Although these are not easy to define formally, they can be roughly characterised as follows.


#### Definition 4.8

A function  $f: \mathbb{R} \rightarrow \mathbb{R}$  is *continuous* if one can draw the function without lifting one’s pen from the page.

To illustrate, consider the function

$$f(x) = \begin{cases} 1 & \text{if } x < 0 \\ 2 & \text{if } x \geq 0 \end{cases}$$

This function is defined ‘piecewise’, which just means that it is defined differently in different parts of its domain. As can be seen below, this function is *not* continuous since it has a ‘jump’ at zero (see the illustration below). For this reason, one could not draw the function without lifting one’s pen from the page. In contrast, all the other functions that we have encountered so far (e.g., polynomials, exponentials) *are* continuous.



Here is a quick problem to check your understanding of continuity.

**Problem 4.9: Continuity**

Define  $f: \mathbb{R} \rightarrow \mathbb{R}$  with

$$f(x) = \begin{cases} x^2, & \text{if } x \geq 0 \\ 0 & \text{if } x < 0 \end{cases}$$

Is  $f$  continuous?

**Increasing and decreasing functions.** As our final topic, we discuss increasing and decreasing functions. Strictly increasing functions are those that slope upwards (see, for example, the previous graph of  $e^x$ ). Weakly increasing functions can never slope downwards (and may, or may not, slope upwards). Similarly, strictly decreasing functions slope downwards, and decreasing functions never slope upwards. These points are precisified by the definition below.

#### Definition 4.9

A function  $f: \mathbb{R} \rightarrow \mathbb{R}$  is *increasing* if

$$x_1 < x_2 \text{ implies that } f(x_1) \leq f(x_2)$$

and *strictly increasing* if

$$x_1 < x_2 \text{ implies that } f(x_1) < f(x_2).$$

Decreasing functions are defined analogously.

If a function is either strictly increasing or strictly decreasing, then it is called *strictly monotone*. Such functions are necessarily ‘one-to-one’: if they have two different inputs, then these inputs must produce different outputs. For this reason, any strictly monotone function is invertible.

The final problem encourages you to check whether a particular function is increasing. While this can be done using the definitions given above, it can also be accomplished using the methods that we will develop in the next lecture.

#### Problem 4.10: Increasing functions

Let  $f: (0, \infty) \rightarrow \mathbb{R}$  with

$$f(x) = 2 - \frac{1}{x}$$

Is  $f$  weakly increasing? Is it strictly increasing?

## 5. DERIVATIVES

### 5.1 Motivating examples

#### Problem 5.1: Returns to scale

If a firm hires  $L \geq 0$  employees, they will be able to produce  $L^2$  units of their product. Does the firm experience increasing or decreasing returns to labour?

#### Problem 5.2: Price sensitivity

At the current price it charges, the price elasticity of demand for a firm's product is -0.5. If the firm slightly reduces its price, will its total revenue increase or decrease?

### 5.2 Defining derivatives

Rational choice involves 'thinking at the margin'. This means that, when people make optimal choices, they often need to consider questions like:

- Should I buy a little more rice?
- Should I hire a few more employees?
- Should the government slightly increase taxes?

Derivatives are useful (in part) because they provide a way to formalise such kinds of marginal thinking.

So, what is a derivative? Roughly speaking, the *derivative* of a function tells us how quickly the function changes when its 'input' changes very slightly. In other words, it is a measure of the *sensitivity* of the function. As we will see, the derivative of a function can also be viewed as the *slope* of the function at a particular point.

At this stage, it is helpful to introduce the formal definition of a derivative.

#### Definition 5.1

Consider a function  $f: \mathbb{R} \rightarrow \mathbb{R}$ . The *derivative* of  $f$  at  $x_0$  is the limit

$$f'(x_0) = \lim_{h \rightarrow 0} \frac{f(x_0 + h) - f(x_0)}{h}$$


if this limit exists. If the limit exists,  $f$  is called *differentiable*.

To understand this, consider an arbitrary point  $x_0$  and consider changing this to some point  $x_0 + h$ . If one does this, the *change* in  $x$ , often denoted  $\Delta x$ , is simply  $(x_0 + h) - x_0 =$



$h$ : this is the denominator in the formula above. Moreover, the change in the function output is  $\Delta f = f(x_0 + h) - f(x_0)$ : this is the numerator in the formula above. Thus, a derivative is nothing more than the limit of the ratio  $\Delta f / \Delta x$  as  $h \rightarrow 0$  (i.e., as the change in  $x$  becomes infinitesimally small).

The notion of a derivative can also be explained graphically. In the left figure, one can see the horizontal distance  $h$  (the difference between  $x_0 + h$  and  $x_0$ ): this is the denominator of the ratio. One can also see the vertical distance  $f(x_0 + h) - f(x_0)$ : this is the numerator of the ratio. The ratio itself is simply the slope of the dotted line connecting the points  $(x_0, f(x_0))$  and  $(x_0 + h, f(x_0 + h))$ . As  $h \rightarrow 0$  (see the right figure), the point  $x_0 + h$  moves to the left until the two points ‘collapse’ into a single point. From this, one sees that *the derivative at  $x_0$  is the slope of the function at  $x_0$* .



The following example checks your grasp of the definition of a derivative.

#### Problem 5.3: Computing a derivative

Let  $f(x) = x^2$  for all  $x \in \mathbb{R}$ . Using the definition of a derivative, compute  $f'(0)$ .

**The  $h = 1$  approximation.** We have defined derivatives as the limit

$$f'(x_0) = \lim_{h \rightarrow 0} \frac{f(x_0 + h) - f(x_0)}{h}$$

In some applications, however,  $h = 1$  is ‘sufficiently small’ to well approximate this limit. In these cases,

$$f'(x_0) \approx \frac{f(x_0 + 1) - f(x_0)}{1} = f(x_0 + 1) - f(x_0)$$

For this reason, derivatives can (very roughly) be viewed as *the effect of increasing  $x$  by one unit on  $f(x)$* . This observation is an important one since such approximations pervade economics and economic theory. To provide some examples:

- Suppose that a firm produces output  $Y$  by hiring ‘labour’  $L$ . (In less dated language,  $L$  could be viewed as the firm’s number of employees.) One can then define the *marginal product of labour* as the amount of extra output that an extra employee will generate, i.e.  $Y(L + 1) - Y(L)$ . Alternatively, one could define the marginal product as the *derivative* of the ‘production function’  $Y$  with respect to  $L$ , i.e.  $Y'(L)$ . In certain cases (namely, when  $h = 1$  provides a ‘good approximation’ to  $h \rightarrow 0$ ), these two definitions will almost coincide.
- Suppose that, in order to produce a certain level of output  $Y$ , the firm must incur a total cost of  $C$ . This defines a ‘cost function’  $C(Y)$ . One can then define the *marginal cost of output* as the additional cost involved in producing one more unit of output, i.e.  $C(Y + 1) - C(Y)$ . Alternatively, one can define the marginal cost of output as the derivative of the cost function, i.e.  $C'(Y)$ . Again, these two definitions may give reasonably similar results in practical cases.
- To provide a final example, suppose an individual will obtain a level of ‘utility’  $u$  given the amount of money they have  $m$ . This defines a *utility function*  $u(m)$ . One can then define the *marginal utility of money* as the amount of extra utility that the individual gets if they obtain an additional £1, i.e.  $u(m + 1) - u(m)$ . Alternatively, one could define the marginal utility of money as  $u'(m)$ . As usual, one might hope that, in practice,  $u(m + 1) - u(m) \approx u'(m)$ .

The next example allows you to explore the extent to which these two different ways of defining ‘marginal effects’ match up in a particular case.

#### Problem 5.4: The marginal product of labour

Suppose that, if a firm hires  $L \geq 0$  employees, it produces  $Y = L^2$  units of output.

- (i) How much does its output increase if it increases its number of employees from 0 to 1? (ii) How does this compare to  $Y'(0)$ ?

**Notation.** We conclude this section with some remarks on notation. So far, we have written  $f'(x_0)$  to denote the derivative of  $f$  at the point  $x_0$ . However, there are other ways to denote derivatives. Another option would be to write

$$\left. \frac{d}{dx} f(x) \right|_{x=x_0}$$

which simply means: ‘differentiate  $f(x)$ , and then evaluate the result at  $x_0$ ’. If we let  $y = f(x)$  denote the output of the function, one can also write the derivative as

$$\left. \frac{dy}{dx} \right|_{x=x_0}$$

This last notational choice is quite intuitive since, as discussed, the derivative is just the change in the output  $\Delta y$  divided by an infinitesimal change in the input  $\Delta x$ .

So far, we have talked about the derivative of a function at a particular point  $x_0$ . However, since we can compute this at *any* point  $x_0$ , this defines a new *function*  $f'(x)$ . Note that this function returns the derivative of  $f$  at *any* particular point  $x$ . Similarly to before, one can also denote this function by

$$\frac{d}{dx}f(x)$$

Alternatively, letting  $y = f(x)$ , one can write

$$\frac{dy}{dx}$$

Again, this last bit of notation provides a nod to the definition of a derivative, i.e. the ratio  $\Delta y/\Delta x$  as  $\Delta x \rightarrow 0$ .

### 5.3 Examples

We now consider some examples of functions that are especially easy to differentiate. As our first example, we consider *power functions*. You may remember that, if  $f(x) = x^k$  for some integer  $k$ ,  $f$  is called a *monomial function*. The notion of a power function generalises this by allowing the exponent  $k$  to be any real number. For example,  $f(x) = x^2$ ,  $f(x) = x^{2.5}$  and  $f(x) = x^\pi$  are all valid power functions. The next result shows how to compute the derivative of such functions.

#### Result 5.1: Differentiating a power function

If  $f(x) = x^k$  for some real number  $k$ , then  $f'(x) = kx^{k-1}$ .

To illustrate, consider the function  $f(x) = x^3$ . Upon differentiation, this becomes  $f'(x) = 3x^2$ : thus, one ‘takes the power down’ and then ‘subtracts one from the power’. Likewise,  $x^4$  becomes  $4x^3$ ,  $x^5$  becomes  $5x^4$ , and  $x^{5.5}$  becomes  $5.5x^{4.5}$ .

There are two special cases of this result that should perhaps be emphasised. First, suppose that  $k = 0$  and so  $f(x) = x^k = x^0 = 1$ . The result then says that  $f'(x) = 0 \times x^{-1} = 0$ . Thus, if  $f(x) = 1$ ,  $f'(x) = 0$ . This illustrates the general rule that *constants vanish if differentiated*. Second, suppose that  $k = 1$  and so  $f(x) = x^k = x^1 = x$ . The result then says that  $f'(x) = 1 \times x^0 = 1$ . Thus, if  $f(x) = x$ ,  $f'(x) = 1$ . This illustrates the general rule that *linear functions have a constant slope*.

Next, we consider *exponential functions*. (Hopefully, you remember these from the previous lecture on functions!) It turns out that these are also easy to differentiate:

### Result 5.2: Differentiating an exponential function

If  $f(x) = b^x$  for a constant base  $b > 0$ ,  $b \neq 1$ , then

$$f'(x) = b^x \ln(b)$$

An important special case arises when the base  $b$  is equal to the number  $e \approx 2.718$ , i.e. when the function in question is  $f(x) = b^x = e^x$ . The result then says that  $f'(x) = b^x \ln(b) = e^x \ln(e) = e^x$ . Thus, this result implies that the derivative of the exponential function  $e^x$  is just  $e^x$ : differentiation does not change the function.

Finally, we consider *logarithmic functions*. (Again, these should also be familiar.) Such functions can be differentiated as follows:

### Result 5.3: Differentiating a logarithmic function

Let  $f(x) = \log_b(x)$  with  $b > 0$ ,  $b \neq 1$ , and  $x > 0$ . Then

$$f'(x) = \frac{1}{x \ln(b)}$$

Similarly to before, an important special case arises when  $b = e \approx 2.718$ , i.e. when the function in question is  $f(x) = \log_e(x) = \ln(x)$ . The result then says that  $f'(x) = 1/(x \ln(e)) = 1/x$ . Thus, the derivative of the natural logarithm  $\ln(x)$  is simply  $1/x$ .

We conclude this section with a warning. In all the previous examples, our functions have well-defined derivatives. However, some functions *cannot* be differentiated, at least at certain points, since the relevant limit (as  $h \rightarrow 0$ ) is not well-defined. Examples included functions with a ‘sharp corner’ like  $f(x) = |x|$  or functions with ‘jumps’ (i.e. discontinuities). Although this course will generally focus on differentiable functions, you should remember that not all functions have this property.

## 5.4 Rules

We now provide a brief overview of certain rules that govern differentiation. These rules are useful since they allow us to *generalise* the results that we have outlined so far and thus obtain the derivatives of a much larger set of functions.

**Linearity.** Our first rule combines two results. The first result is that the derivative of *sum* of functions is the sum of the derivatives of the individual functions. In other words, if  $h(x) = f(x) + g(x)$ , then  $h'(x) = f'(x) + g'(x)$ . The second result is that the derivative of a function multiplied by some constant is that constant multiplied by the derivative of the function. In other words, if  $h(x) = af(x)$  for some  $a \in \mathbb{R}$ , then  $h'(x) = af'(x)$ . If one puts these two results together, one obtains the following:

#### Result 5.4: Linearity of differentiation

Let  $f$  and  $g$  be differentiable functions, and let  $a, b \in \mathbb{R}$ . Then

$$\frac{d}{dx}(af(x) + bg(x)) = af'(x) + bg'(x)$$

The next problem allows you to put this result to work.

#### Problem 5.5: Differentiating a polynomial

Let  $f(x) = 3x + 4x^2$ . What is the derivative  $f'(x)$ ?

**The product rule.** Our second rule tells us how to differentiate the *product* of two functions.

#### Result 5.5: The product rule

If  $f$  and  $g$  are differentiable functions, then

$$\frac{d}{dx}(f(x)g(x)) = f'(x)g(x) + f(x)g'(x)$$

Notice that the derivative of the product  $f(x)g(x)$  is *not* (in general) equal to the product of the derivatives,  $f'(x)g'(x)$ . Instead, it is equal to the derivative of  $f$  multiplied by  $g$  plus the derivative of  $g$  multiplied by  $f$ . The next problem encourages you to make use of this fact.

#### Problem 5.6: Applying the product rule

Let  $f(x) = x^2$  and  $g(x) = \ln(x)$  for all  $x > 0$ . Use the product rule to compute the derivative of

$$h(x) = f(x)g(x) = x^2 \ln(x)$$

**The quotient rule.** Our ‘third’ rule is essentially the product rule once again but expressed in a rather different way. This time, we want to differentiate a ratio (or ‘quotient’)  $f(x)/g(x)$ . Notice that, as before, the derivative of this ratio is absolutely *not* the ratio of the two derivatives:

### Result 5.6: The quotient rule

If  $f$  and  $g$  are differentiable and  $g(x) \neq 0$ , then

$$\frac{d}{dx} \left( \frac{f(x)}{g(x)} \right) = \frac{f'(x)g(x) - f(x)g'(x)}{[g(x)]^2}$$

### Problem 5.7: Applying the quotient rule

Let  $f(x) = x^2 + 1$  and  $g(x) = x$  for all  $x > 0$ . Use the quotient rule to compute the derivative of

$$h(x) = \frac{f(x)}{g(x)} = \frac{x^2 + 1}{x}$$

**The chain rule.** Our final rule tells us how to differentiate *composite* functions, i.e. functions obtained by substituting one function into another. Specifically, the rule relates the derivative of a composite function  $f \circ g$  to the derivatives of the ‘underlying’ functions  $f$  and  $g$ :

### Result 5.7: The chain rule

If  $f$  and  $g$  are differentiable and  $h(x) = f(g(x))$ , then

$$h'(x) = f'(g(x)) \cdot g'(x)$$

### Problem 5.8: Applying the chain rule

Let  $f(x) = e^x$  and  $g(x) = x^3$ . Use the chain rule to compute the derivative of the composite function

$$h(x) = f(g(x)) = e^{x^3}$$

## 5.5 Applications

Derivatives are extremely useful and have a number of important applications. In this section, we discuss three of these applications. It should be emphasised, however, that these are just a few of the many uses of derivatives in economics and economic theory. For example, the theory of optimisation (which we will develop in the subsequent lectures) rests to a large extent on derivatives.

**Monotonicity.** In the previous lecture, we defined what it means for a function to be strictly increasing or strictly decreasing. In practice, however, the definition that we gave can be rather difficult to check. Fortunately, it is *also* possible to check whether a function is strictly increasing (or strictly decreasing) using derivatives:

### Result 5.8: Derivatives and monotonicity

Let  $f$  be a differentiable function on  $\mathbb{R}$ .

- If  $f'(x) > 0$  for all  $x$  in some interval, then  $f$  is *strictly increasing* on that interval.
- If  $f'(x) < 0$  for all  $x$  in some interval, then  $f$  is *strictly decreasing* on that interval.
- If  $f'(x) = 0$  for all  $x$  in some interval, then  $f$  is *constant* on that interval.

This result should not come as a surprise. For example, if  $f'(x) > 0$  for all  $x$  in some interval, then the function  $f$  has a positive slope over that interval. Obviously, then,  $f$  is strictly increasing (over that same interval). Although not very surprising, this result can be useful, as indicated by the subsequent example.

### Problem 5.9: Increasing functions

Let  $f: (0, \infty) \rightarrow \mathbb{R}$  with

$$f(x) = 2 - \frac{1}{x}$$

Is  $f$  weakly increasing? Is it strictly increasing?

**Concavity and convexity.** Concavity and convexity are important notions in economics. For an increasing function, concavity corresponds to *decreasing marginal returns*: think, for example, of a factory that can produce more if it hires more employees, but finds that each new employee adds less to their production than the last. Likewise, convexity corresponds to the notion of *increasing marginal returns*: in this case, each additional input creates *more* output than the previous one.

Typically, one says that a real function is *strictly concave* on an interval  $I$  if, for all distinct  $x_1, x_2 \in I$  and all  $t \in (0, 1)$ ,

$$f(tx_1 + (1 - t)x_2) > tf(x_1) + (1 - t)f(x_2)$$

(If the inequality is reversed, the function is *strictly convex*.) However, these definitions are somewhat complex and can be difficult to check. Fortunately, one does not necessarily need to use these definitions; instead, one can use derivatives.


Specifically, we start with a function  $f$  and compute its derivative  $f'$ . We then differentiate (again!) to obtain the second derivative  $f''$ . Intuitively, this is just the rate of change of  $f'$ . When this can be done (namely, when a function is *twice differentiable*), this provides us with exactly the information we need to check concavity and convexity:

### Result 5.9: Concavity and the second derivative


Let  $f$  be twice differentiable.

- If  $f''(x) < 0$  for all  $x$  in an interval, then  $f$  is strictly concave on that interval.
- If  $f''(x) > 0$  for all  $x$  in an interval, then  $f$  is strictly convex on that interval.

The figures below illustrate this result. The figure on the left shows a concave function that exhibits *diminishing marginal returns*. As one can see, the function starts off quite steep (with a large and positive slope) but gradually flattens, leading its derivative  $f'$  to fall. For this reason, its second derivative  $f''$  is negative: increasing  $x$  makes the slope (i.e.  $f'$ ) fall. The figure on the right shows a convex function that exhibits *increasing marginal returns*. Here, the function starts off quite flat (with a small slope  $f'$ ) but eventually becomes steep, leading its derivative  $f'$  to increase. For this reason, its second derivative  $f''$  is now positive: increasing  $x$  makes the slope (i.e.  $f'$ ) rise.



(a) A concave function ( $f'' < 0$ )



(b) A convex function ( $f'' > 0$ )

### Problem 5.10: Concavity

Use the concavity test to establish that the function  $f(x) = x^{1/2}$  is indeed strictly concave over  $(0, \infty)$ .

**Elasticities.** In economics, we often want to know how *sensitive* one variable is to changes in another. For example, we want to know how much the demand for a good changes when its price changes; this is called the good's *elasticity of demand*. Roughly speaking, the *elasticity* of  $y$  with respect to  $x$  can be viewed as the percentage change in  $y$  divided by the percentage change in  $x$ . For example, if  $x$  increases by 1%, and  $y$  correspondingly increases by 3%, then the elasticity of  $y$  with respect to  $x$  is 3.

Although this is one way of defining elasticities, it is not the standard way in mi-



economic theory. Instead, one typically defines elasticities in terms of *derivatives*:

**Definition 5.2**

Let  $y = f(x)$  with  $x > 0$  and  $y > 0$ . The *elasticity of  $y$  with respect to  $x$*  is

$$\epsilon_{y,x} = \frac{dy}{dx} \times \frac{x}{y}$$

One then says that:

- If  $|\epsilon_{y,x}| > 1$ ,  $y$  is *elastic*: it responds more than proportionally to  $x$ .
- If  $|\epsilon_{y,x}| < 1$ ,  $y$  is *inelastic*: it responds less than proportionally to  $x$ .
- If  $|\epsilon_{y,x}| = 1$ ,  $y$  changes proportionally with  $x$ .

How does this all connect with the notion of a ratio of percentage changes? To understand this, recall that the ‘percentage change’ between two numbers is just the difference between the numbers divided by the ‘original value’. Thus,

$$\frac{\% \Delta y}{\% \Delta x} = \frac{\left(\frac{\Delta y}{y}\right)}{\left(\frac{\Delta x}{x}\right)} = \frac{\Delta y}{y} \times \frac{x}{\Delta x} = \frac{\Delta y}{\Delta x} \times \frac{x}{y}$$

But this is very reminiscent of our definition of an elasticity! Indeed, as  $\Delta x \rightarrow 0$  (i.e., we consider an infinitesimally small change in  $x$ ), this *exactly* converges to  $\epsilon_{y,x}$  as defined above. From this, we deduce that *the elasticity of  $y$  with respect to  $x$  is the limit (as  $\Delta x \rightarrow 0$ ) of the percentage change in  $y$  divided by the percentage change in  $x$ .*

Our final problem tests your understanding of elasticities. In particular, it encourages you to compare the derivatives based definition with the ratio  $\% \Delta y / \% \Delta x$ .

**Problem 5.11: Elasticities**

If a firm sets a price  $p > 0$ , it will sell  $100/p$  units of its product. The firm initially sets  $p = 1$ . (i) Compute the product’s elasticity of demand. (ii) What will be the percentage change in demand if the firm increases its price by 10%?

## 6. MULTIVARIATE FUNCTIONS

### 6.1 Motivating examples

#### Problem 6.1: Production isoquants

If a firm hires  $K \geq 0$  units of capital and  $L \geq 0$  units of labour, it can generate  $\min(K, 2L)$  units of output. Sketch all the  $(K, L)$  combinations that generate 8 units of output.

#### Problem 6.2: Returns to scale

If a firm hires  $K \geq 0$  units of capital and  $L \geq 0$  units of labour, it can generate  $KL^2$  units of output. What happens to the firm's output if it scales up its inputs by a common factor?

### 6.2 Multivariate functions

In the course so far, we have studied *univariate* functions. In economics, however, outcomes of interest typically depend on *many* variables. For example, an economy's rate of unemployment might depend on government policies (e.g., minimum wage rules), labour market institutions (e.g., unionisation), demographic trends (e.g., the number of working versus retired individuals), and many other factors besides. For this reason, we now turn our attention to *multivariate* functions, which are simply functions whose output depends on two or more variables. From a more formal point of view, one can say that a function is multivariate if it has a *multidimensional domain*.

**Multivariate domains.** What are some multivariate domains that arise? To keep things simple in the lecture, we will generally consider domains that are *two-dimensional*. Important examples of such domains include:

- $\mathbb{R}^2 = \{(x, y) : x \in \mathbb{R}, y \in \mathbb{R}\}$ . (Translation: 'the set of all  $(x, y)$  pairs such that both  $x$  and  $y$  are real numbers'.)
- $\mathbb{R}_{\geq 0}^2 = \{(x, y) : (x, y) \in \mathbb{R}^2, x \geq 0, y \geq 0\}$ . (Translation: 'the set of all  $(x, y)$  pairs such that both  $x$  and  $y$  are non-negative real numbers'.)
- The *closed unit square*  $[0, 1]^2 = \{(x, y) : x \in [0, 1], y \in [0, 1]\}$ . (Translation: 'the set of all  $(x, y)$  pairs such that both  $x$  and  $y$  are real numbers between 0 and 1'.)

More generally, given a set  $A$  and a set  $B$ , one can form the new set  $A \times B = \{(a, b) : a \in A, b \in B\}$ . In this context, the symbol ' $\times$ ' is called the *Cartesian product*. For example, if  $A = \{1, 2\}$  and  $B = \{3, 5\}$ , the Cartesian product of the sets is  $A \times B = \{(1, 3), (1, 5), (2, 3), (2, 5)\}$ . When referring to the Cartesian product of a set with itself  $A \times A$ , one can also write  $A^2$ : this explains the 'squared' notation used above.

The following example allows you to test your understanding of two-dimensional spaces.

**Problem 6.3: Multidimensional domains**

Determine whether the following statements are true or false:

- (1)  $0.22 \in \mathbb{R}^2$
- (2)  $(-4.5, 2.1) \in \mathbb{R}_{\geq 0}^2$
- (3)  $(0.5, 0.7) \in [0, 1]^2$
- (4)  $(0.5, 0.5) \in [0, 1] \times [2, 3]$

**Multivariate functions.** Multidimensional functions are common in economics. For example, a firm's *production function* describes how its level of output  $Y$  depends on the amount of labour  $L$  and capital  $K$  it hires. (In theory, this can be generalised to any number of production inputs; however, we will stick to the case of two inputs to keep things simple.) Since  $K \geq 0$  and  $L \geq 0$ , the domain is taken to be  $\mathbb{R}_{\geq 0}^2$ . Popular examples of production functions in applied models include:

- The 'Cobb-Douglas' production function  $Y = AK^\alpha L^{1-\alpha}$ . Here,  $A > 0$  is a scaling parameter and  $\alpha \in (0, 1)$  controls the relative marginal contribution of each factor to the total output.
- The linear production function  $Y = A(\alpha K + (1 - \alpha)L)$ . Again  $A > 0$  is a scaling parameter and  $\alpha \in (0, 1)$  controls the relative importance of the two factors.
- The 'Leontief' production function  $Y = A \min\{\alpha K, (1 - \alpha)L\}$ . (Note that the min function simply returns its smallest argument: for example,  $\min\{2, 4\} = 2$  and  $\min\{2, 2\} = 2$ .) As will become clear, this function describes a situation where the two factors  $K$  and  $L$  must be used in fixed proportions.

Multivariate functions are also used in consumer theory. Specifically, an individual's *utility function* describes how their level of utility  $u$  depends on the amount of good  $x$  and good  $y$  that they consume. Since consumers cannot buy negative quantities,  $x \geq 0$  and  $y \geq 0$ : thus, the domain is again  $\mathbb{R}_{\geq 0}^2$ . Just as in the case of production functions, some popular specifications for utility functions include:

- The 'Cobb-Douglas' utility function  $u = x^\alpha y^{1-\alpha}$ .
- The linear utility function  $u(x, y) = \alpha x + (1 - \alpha)y$ .
- The 'Leontief' utility function  $u = \min\{\alpha x, (1 - \alpha)y\}$ .

The parameter  $\alpha \in (0, 1)$  in the previous utility functions again controls the relative marginal contributions of  $x$  and  $y$  to the consumer's utility. As before, one can also multiply the utilities by a scaling factor  $A > 0$ . However, this turns out to matter less than in the production case. In fact, there is an important sense in which the utility function  $Au$  is the 'same' as the utility function  $u$  (without the scaling factor).

#### Problem 6.4: A production function

Suppose that a firm produces  $Y(K, L) = K^{1/2}L^{1/2}$  units of output given that they hire  $L \geq 0$  employees and  $K \geq 0$  units of capital.

- (1) Compute  $Y(0, 0)$ ,  $Y(1, 0)$ ,  $Y(0, 1)$  and  $Y(1, 1)$ .
- (2) What kind of production function does the firm have?

**Level curves.** One way to visualise a two-dimensional function is to draw its *level curves*. Intuitively, this shows *different inputs* that generate the *same* output. For example, a hiker's map might show different locations (each described by a latitude and longitude) that correspond to the same vertical height. These curves (or 'contours') are thus an example of level curves. More formally, we say that:

#### Definition 6.1

Let  $f: \mathbb{R}^2 \rightarrow \mathbb{R}$ . The *level curve* of  $f$  at level  $c$  is the set of  $(x, y)$  pairs in the domain such that  $f(x, y) = c$ .

Some level curves arise so frequently in economics that they have acquired special names. For example:

- A consumer's *indifference curve* is the set of  $(x, y)$  bundles that provide her with some fixed level of utility. For example, one can draw the set of bundles that give a consumer 5 units of utility, the set of bundles that give a consumer 10 units of utility, etc.
- A firm's *production isoquant* is the set of  $(K, L)$  inputs that generate some fixed level of output. For example, one can draw the set of inputs that generate 5 units of output, the set of bundles that generate 10 units of output, etc.
- A firm's *isocost line* is the set of  $(K, L)$  inputs that cost the same amount. Again, one can draw different isocost lines corresponding to different total costs: for example, one isocost line is the set of  $(K, L)$  input choices that all cost £100.

(To help remember the meaning of 'isoquant' and 'isocost', recall that the prefix 'iso' derives from the Ancient Greek word 'isos', meaning 'equal' or 'the same'. Think isosceles, isotope, isobar...)

The next problem allows you to draw some indifference curves. While this may be the first time you are drawing indifference curves, I promise that it won't be the last!

#### Problem 6.5: Level curves

If Anna consumes  $a \geq 0$  apples and  $b \geq 0$  bananas, she obtains the utility level  $u(a, b) = a^{1/2}b^{1/2}$ . Plot her indifference curves at the levels  $u = 1$ ,  $u = 2$  and  $u = 3$ .

## 6.3 Classifying functions

In the lecture on univariate functions, we introduced some important function classifications. Some of these classifications extend to the multivariate case in an obvious way. For example, a multivariate function  $\mathbb{R}^2 \rightarrow \mathbb{R}$  is:

- *Strictly increasing* if it is strictly increasing in each of its components. That is,  $x' > x \implies f(x', y) > f(x, y)$  for any fixed  $y$ ; and  $y' > y \implies f(x, y') > f(x, y)$  for any fixed  $x$ .
- *Strictly decreasing* if it is strictly decreasing in each of its components. That is,  $x' > x \implies f(x', y) < f(x, y)$  for any fixed  $y$ ; and  $y' > y \implies f(x, y') < f(x, y)$  for any fixed  $x$ .

**Homogeneity.** We now introduce two classifications that we had not discussed previously. First, we introduce the idea of a homogeneous function.

### Definition 6.2

A function  $f: \mathbb{R}^2 \rightarrow \mathbb{R}$  is homogeneous of degree  $n$  if, for all  $\lambda > 0$  and all  $(x, y)$  in the domain,

$$f(\lambda x, \lambda y) = \lambda^n f(x, y)$$

To understand this, imagine changing the inputs of a function by a common factor  $\lambda > 0$ ; in other words, the inputs start at  $(x, y)$  and become  $(\lambda x, \lambda y)$ . The definition says that, if the function is homogeneous, the output will change from  $f(x, y)$  to  $\lambda^n f(x, y)$ . That is, changing the inputs by a factor  $\lambda$  will change the output by a factor  $\lambda^n$ .

Why is this classification useful? Let's say that an increasing function  $f: \mathbb{R}_{\geq 0}^2 \rightarrow \mathbb{R}_{\geq 0}$  exhibits

- *Decreasing returns to scale* if increasing the inputs by a factor  $\lambda > 1$  leads to a less than proportionate increase in output. That is,  $f(\lambda x, \lambda y) < \lambda f(x, y)$  for all  $\lambda > 1$  and for all  $(x, y)$ .
- *Constant returns to scale* if increasing the inputs by a factor  $\lambda > 1$  leads to a proportionate increase in output. That is,  $f(\lambda x, \lambda y) = \lambda f(x, y)$  for all  $\lambda > 1$  and for all  $(x, y)$ .
- *Increasing returns to scale* if increasing the inputs by a factor  $\lambda > 1$  leads to a more than proportionate increase in output. That is,  $f(\lambda x, \lambda y) > \lambda f(x, y)$  for all  $\lambda > 1$  and for all  $(x, y)$ .

Intuitively, one would expect a connection between homogeneity and returns to scale. For example, suppose that  $\lambda = 2$ ; so we are doubling the inputs of a function. If the function is homogeneous of degree  $n$  with  $n > 1$ , the output will increase by a factor  $2^n > 2$ . Thus, the output will increase more than proportionately ('increasing

returns to scale'). Meanwhile, if the function is homogeneous of degree  $n$  with  $n < 1$ , the output will increase by a factor  $2^n < 2$ . Thus, the output will increase less than proportionately ('decreasing returns to scale'). Using this kind of reasoning, one sees that:

**Result 6.1: Homogeneity and returns to scale**

Let  $f: \mathbb{R}_{\geq 0}^2 \rightarrow \mathbb{R}_{\geq 0}$  be an increasing function that is homogeneous of degree  $n > 0$ . Then:

- If  $n < 1$ ,  $f$  exhibits decreasing returns to scale.
- If  $n = 1$ ,  $f$  exhibits constant returns to scale.
- If  $n > 1$ ,  $f$  exhibits increasing returns to scale.

The next problem lets you put this result into practice.

**Problem 6.6: Homogeneity of a production function**

Let  $Y(K, L) = K^{0.3}L^{0.7}$  for all  $K \geq 0$  and  $L \geq 0$ . Show that this function is homogeneous. Does it exhibit decreasing, constant, or increasing returns to scale?

**Quasi-concave functions.** We now introduce the notion of a quasi-concave function.

**Definition 6.3**

Let  $f: \mathbb{R}^2 \rightarrow \mathbb{R}$ . The function  $f$  is *weakly quasi-concave* if, for all  $t \in (0, 1)$  and for all  $(x_1, y_1) \neq (x_2, y_2)$ ,

$$f(tx_1 + (1-t)x_2, ty_1 + (1-t)y_2) \geq \min \{f(x_1, y_1), f(x_2, y_2)\}$$

The function  $f$  is *strictly quasi-concave* if the inequality holds strictly.

This says that a function is quasi-concave if, for any number  $t \in (0, 1)$ , and for any distinct inputs  $(x_1, y_1)$  and  $(x_2, y_2)$ , a certain inequality holds. What is the meaning of this inequality? The right hand side is just the *minimum* of  $f(x_1, y_1)$  and  $f(x_2, y_2)$ : thus, it is either equal to  $f(x_1, y_1)$ , or equal to  $f(x_2, y_2)$ , depending on which is smaller. To understand the left hand side, notice that  $tx_1 + (1-t)x_2$  is just a weighted average of  $x_1$  and  $x_2$ . Likewise,  $ty_1 + (1-t)y_2$  is just a weighted average of  $y_1$  and  $y_2$ . Thus, the left hand side is the value of the function when evaluated at a weighted average of the two inputs  $(x_1, y_1)$  and  $(x_2, y_2)$ . Therefore, *a function is quasi-concave if the function evaluated at an average of the inputs is larger than the minimum of the function evaluated at either input.*

Why is this an important notion in economics? To understand this, consider a consumer with utility function  $u$  and consider two consumption bundles,  $(x_1, y_1)$  and  $(x_2, y_2)$ . If

the function is strictly quasi-concave, then the inequality

$$u(tx_1 + (1-t)x_2, ty_1 + (1-t)y_2) > \min \{u(x_1, y_1), u(x_2, y_2)\}$$


holds for all  $t \in (0, 1)$  and distinct bundles  $(x_1, y_1)$  and  $(x_2, y_2)$  in the domain of  $u$ . (So far, this is just a relabelling of  $f$  as  $u$ .) Now let's suppose that the consumer is indifferent between the bundles, i.e.  $u(x_1, y_1) = u(x_2, y_2)$ . If this is true,  $\min \{u(x_1, y_1), u(x_2, y_2)\} = u(x_1, y_1) = u(x_2, y_2)$ . It follows that

$$u(tx_1 + (1-t)x_2, ty_1 + (1-t)y_2) > u(x_1, y_1) = u(x_2, y_2)$$

Thus, strict quasi-concavity implies that, *if a consumer is indifferent between two bundles, they must prefer any 'average' of the bundles to either bundle*. For example, suppose that a consumer is indifferent between a) one tonne of coffee and one gallon of milk, and b) one tonne of tea and one gallon of milk. If the consumer's preferences are quasi-concave, they will strictly prefer any average of these bundles to either of the bundles. For example, they would prefer half a tonne of coffee and half a tonne of tea (along with the gallon of milk) to either option a) or option b).

These ideas can also be developed graphically. To do this, consider again a consumer who obtains utility  $u(x, y)$  given that they consume  $x$  units of one good and  $y$  units of the other. What do their indifference curves look like? First, if the consumer always wants more of both goods, their indifference curves must be *downward sloping*. The reason is that, if they gain more of one good, the only way they can remain indifferent is if they lose some of the other good; thus, the benefit and the cost can cancel out. Second, if the individual's utility function is continuous, their indifference curves must be continuous as well. Thus, the consumer must have some kind of continuous, downward-sloping indifference curve.

Suppose we now add the assumption that the consumer's utility function is strictly quasi-concave. This means that, if a consumer is indifferent between two bundles, they must prefer any 'average' of the bundles to either bundle. Consider then two bundles  $A$  and  $B$  that lie on the same indifference curve (see the figure below). The possible 'averages' of these bundles are given by the dotted line connecting  $A$  and  $B$ : for example,  $C$  is one such average. If the consumer's utility function is strictly quasi-concave, we know that the consumer must prefer a bundle like  $C$  to either  $A$  or  $B$ . This implies that the consumer's indifference curves must be *convex*: that way, all the bundles on the dotted line will be preferred to either  $A$  or  $B$ .



△ Sometimes, it is said that, if a consumer has strictly quasi-concave utility, then they prefer ‘average bundles’ to ‘extreme bundles’. However, this is complete nonsense (or, at the very least, wildly imprecise). A slightly better statement would be that, if the consumer is indifferent between two ‘extreme’ bundles, they prefer an average of the bundles to either extreme.

## 6.4 Derivatives

**Partial derivatives.** Just as one can differentiate univariate functions, one can differentiate multivariate functions. To do this, one uses partial differentiation: this tells us how the function changes when we slightly increase one variable, holding the other variables fixed. In other words, we treat the other variables like constants: thus, ‘partial differentiation’ is essentially just regular differentiation in the presence of fixed parameters.

For the curious, here is a formal definition of a partial derivative. It should be emphasised, however, that there is nothing really new here: partial differentiation is essentially just the same as differentiation in the univariate case.

### Definition 6.4

Let  $f : \mathbb{R}^2 \rightarrow \mathbb{R}$  and consider a point  $(x_0, y_0)$  in  $f$ ’s domain. At this point, the *partial derivative* of  $f$  with respect to  $x$  is

$$f_x(x_0, y_0) = \lim_{h \rightarrow 0} \frac{f(x_0 + h, y_0) - f(x_0, y_0)}{h}$$

assuming that this limit exists. The partial derivative with respect to  $y$  is defined analogously.



### Problem 6.7: Partial differentiation

Let  $f(x, y) = 3xy^2$  for all  $(x, y) \in \mathbb{R}^2$ . Compute the partial derivatives  $f_x$  and  $f_y$ .

So far, we have explained how to differentiate a multivariate function once. But what happens if we do this twice? If we differentiate with respect to the *same* variables twice, we obtain the *second* derivatives  $f_{xx}$  and  $f_{yy}$ . If we differentiate with respect to  $x$  and then with respect to  $y$ , we obtain the cross derivative  $f_{xy}$ . Similarly, if we differentiate with respect to  $y$  and then with respect to  $x$ , we obtain the cross derivative  $f_{yx}$ . As usual, these derivatives can tell us something about rates of change. For example,  $f_{xy}$  tells us how quickly the first derivative  $f_x$  changes when we slightly increase  $y$ .

An interesting fact is that, under a relatively mild condition, the two cross partial derivatives will match up. This result is sometimes known as ‘Schwarz’s theorem’:

### Result 6.2: Schwarz’s theorem

Let  $f: \mathbb{R}^2 \rightarrow \mathbb{R}$ . If the function  $f$  has continuous second-order partial derivatives, then for all  $(x, y)$ ,

$$f_{xy}(x, y) = f_{yx}(x, y)$$

The next example allows you to compute some partial derivatives as well as to check whether the cross-partials do indeed match up in a particular case.

### Problem 6.8: Higher derivatives

Let  $f(x, y) = 3xy^2$  for all  $(x, y) \in \mathbb{R}^2$ . Compute the second derivatives  $f_{xx}$ ,  $f_{yy}$ ,  $f_{xy}$  and  $f_{yx}$ . Does Schwarz’s theorem apply to this example?

**The chain rule.** Partial differentiation tells us what happens when we increase one variable, holding the others fixed. But what if all the variables change at the same time? To investigate this, consider a two-dimensional function  $x, y \mapsto f(x, y)$  and suppose that both inputs  $x$  and  $y$  are functions of  $t$ . When  $t$  changes, both  $x$  and  $y$  may change. A generalisation of the chain rule reveals how this alters  $f(x, y)$ .

### Result 6.3: The chain rule in two dimensions

If  $x, y$  are differentiable functions of  $t$ ,

$$\frac{d}{dt}f(x(t), y(t)) = f_x(x(t), y(t))x'(t) + f_y(x(t), y(t))y'(t)$$

What is the intuition behind this result? One reason why changing  $t$  can affect  $f(x, y)$  is via changing  $x$ . This channel is captured by the effect of  $t$  on  $x$  ( $x'(t)$ ) multiplied by the effect of  $x$  on  $f$ : this is the first term in the formula. The second term in the formula captures the effect of  $t$  on  $f(x, y)$  via  $y$  in an analogous way.

### Problem 6.9: Applying the chain rule

Let  $Y(K, L) = KL$  with  $K(t) = e^{t/2}$  and  $L(t) = e^t$  for all  $t \geq 0$ . Compute  $\frac{d}{dt}Y(K(t), L(t))$  at  $t = 0$ .

**Implicit differentiation.** Sometimes, one variable is an implicit function of another, even if this functional dependence is not spelled out. For example, the equation  $2x + 3y = 0$  implicitly defines  $x$  as a function of  $y$ : if  $y$  changes,  $x$  must also change if this equation is to continue to ‘balance’. The implicit function theorem tells us *how quickly*  $x$  must change in response to a small change in  $y$ .

### Result 6.4: The implicit function theorem

If  $g(x, y) = 0$  defines  $y$  implicitly as a function of  $x$  near a point with  $g_y \neq 0$ , then

$$\frac{dy}{dx} = -\frac{g_x}{g_y}$$

### Problem 6.10: Applying the implicit function theorem

Consider the relation  $2x + 3y = 0$ . After writing  $y$  as an explicit function of  $x$ , compute  $dy/dx$ . Verify your answer using the implicit function theorem.

In the previous example, the theorem was rather useless: we could simply ‘solve’ the equation for  $y$  and then differentiate directly! In other examples, however, the theorem has a lot of bite. For example, consider the relation  $x = ye^y$ . You won’t be able to explicitly solve  $y$  for  $x$ . However, using the implicit function theorem, you should be able to verify that  $dy/dx = 1$  at  $(x, y) = (0, 0)$ .

**Applications.** Implicit differentiation has many applications in economics. For example, consider again indifference curves, i.e. equations of the form

$$u(x, y) = \bar{u}$$

Suppose we want to know the slope of the indifference curve. Using the implicit function theorem, the slope is

$$-\frac{u_x}{u_y}$$

In other words, the slope is  $\partial u / \partial x$  (the ‘marginal utility of  $x$ ’) divided by  $\partial u / \partial y$  (the ‘marginal utility of  $y$ ’). This ratio plays an important role in consumer theory and is called the *marginal rate of substitution*.

As a second example, consider again production isoquants, i.e. equations of the form

$$F(K, L) = \bar{q}$$

Suppose we want to know the slope of the isoquant (with  $L$  on the  $x$ -axis). Using the implicit function theorem, the slope is

$$-\frac{F_L}{F_K}$$

In other words, the slope is  $\partial F/\partial L$  (the ‘marginal product of labour’) divided by  $\partial F/\partial K$  (the ‘marginal product of capital’). This expression is called the *marginal rate of technical substitution*.

**Problem 6.11: The slope of an isoquant**

Consider the production function  $f(K, L) = K^{0.8}L^{0.2}$ . What is the slope of the isoquant associated with  $f(K, L) = 10$ ?

## 7. OPTIMISATION

### 7.1 Motivating examples

#### Problem 7.1: Optimal napping

A student is deciding how many hours  $h \in [0, 24]$  to spend sleeping in a given day. While asleep, they obtain a happiness level of 0. While awake, they obtain an average happiness level of  $2h$ . How much sleep maximises the total amount of happiness they experience during the day?

#### Problem 7.2: Revenue maximisation

If a firm charges a price  $p \in [0, 6]$ , they will sell  $6 - p$  units of their product. What price maximises their total revenue?

### 7.2 Optimisation

In this lecture, we study optimisation. Informally, this means getting the most that one can given the constraints that one faces. Unsurprisingly, the theory of optimisation plays a large role in economics. For example, the standard theory assumes that

- Firms make whatever choices maximise their profits.
- Consumers purchase the products that maximise their utility.

One reason why the theory of optimisation is useful is that it allows one to deduce what these maximisation assumptions imply about observable behaviour. However, that is not the only reason: the theory can also play a more normative role. For example, suppose a government is trying to determine what rate of taxation it should set, how high a minimum wage to impose, or what level of unemployment benefits to pay. Alternatively, suppose that an individual is trying to decide how many hours to spend working, how many hours to spend sleeping, or how best to invest their money. Although the theory of optimisation can never by itself ‘solve’ such problems, it can bring out their logical structure and thus guide us towards better decisions.

**Terminology.** We now introduce some more precise terminology. A *maximisation problem* involves choosing some variable  $x$  to maximise some function  $f$  subject to the constraint that  $x \in S$ . More succinctly, the problem is to

$$\max_{x \in S} f(x)$$

The function to be maximised  $f$  is called the *objective function*: for example, it might be the amount of profit that a company makes. The variable  $x$  to be chosen is called the *choice variable*: for example, it might be the quantity  $q$  of its product that the firm

chooses to produce. The *constraint set*  $S$  is the set of feasible options that the decision maker has available to them. For example, if the choice variable is the quantity  $q$  that a firm produces, then the constraint set might be  $\mathbb{R}_{\geq 0}$ : this reflects the fact that the firm cannot choose a negative quantity.

What does it mean to ‘solve’ a maximisation problem? Intuitively, this involves:

- Identifying a feasible ‘maximiser’ that makes the objective function as large as it can be.
- Calculating the value of the objective function at this maximiser; that is, the ‘maximum’ of the function.

More formally, we say that:


### Definition 7.1

A point  $x^* \in S$  solves the problem

$$\max_{x \in S} f(x)$$

if  $f(x^*) \geq f(x)$  for all  $x \in S$ . In this case,  $x^*$  is a *maximiser* of  $f$  on  $S$  and  $f(x^*)$  is the *maximum* (or *maximum value*) of  $f$  on  $S$ .

Two points about this definition should be emphasised. First, as it turns out, some maximisation problems have *no solution* since the function  $f$  has no maximiser on the constraint set  $S$ .<sup>2</sup> The details of this issue, however, go beyond the scope of this course. Second, and rather more obviously, a function  $f$  can have *many* maximisers on  $S$ . For example, in the figure below,  $x'$  and  $x''$  *both* maximise the objective function since the function attains the same (maximum) value at these two points. Notice, however, that even when a function has multiple maximisers, it still has just one maximum value.



<sup>2</sup> For example, there is no solution to the problem

$$\max_{x \in (0,1)} x$$

So far, we have just discussed maximisation problems. However, one can analogously define the *minimisation problem*  $\min_{x \in S} f(x)$ . In fact, all our results can be translated into results about minimisation problems since the value  $x \in S$  that minimises  $f$  is the *same* as the value  $x \in S$  that maximises  $-f$ . However, when we engage in this translation, we need to be careful to multiply by  $-1$  at the appropriate times. In particular, this means that the concavity assumptions that we will make in the context of maximisation problems become convexity assumptions in the context of minimisation problems.

The following problem tests your understanding of the terminology that we have just introduced.

**Problem 7.3: Maxima and minima**

Let  $D = \{1, 3, 5\}$  and define the function  $f: D \rightarrow \mathbb{R}$  with  $f(1) = 7$ ,  $f(3) = 2$  and  $f(5) = 2$ . Identify the maximisers and minimisers of  $f$  on  $D$ . What are  $f$ 's maximum and minimum values?


### 7.3 Three basic results

We now state three basic but important results in the theory of optimisation. We start in as simple a way as possible: we will optimise a function  $f$  that depends on a single variable  $x$ . We will also assume that  $x$  can take any value in  $\mathbb{R}$ : in this sense,  $x$  is ‘unconstrained’. The following result provides a simple recipe for solving such maximisation problems when the objective function is strictly concave.

**Result 7.1: Unconstrained maximisation in one variable**

Let  $f: \mathbb{R} \rightarrow \mathbb{R}$  be strictly concave and differentiable. If  $f'(x^*) = 0$ , then  $x^*$  is the unique maximiser of  $f$  on  $\mathbb{R}$ .

To understand this, recall that  $f'(x)$  can be viewed as the slope of the objective function  $f$  at a point  $x$ . If  $f'(x^*) = 0$ , then the slope of the function is zero at the point  $x^*$  (see the figure below). Moreover, if the objective function is strictly concave, it has a strictly decreasing derivative: thus,  $f'(x) > 0$  when  $x < x^*$  and  $f'(x) < 0$  when  $x > x^*$ . Thus, the function initially increases (when  $x < x^*$ ) before eventually decreasing (when  $x > x^*$ ). From this, one concludes that  $x^*$  is the ‘unique’ (i.e. only) maximiser of  $f$ .



The result just given applies when there is a solution to the equation  $f'(x^*) = 0$ . This equation is called the *first-order condition* since it involves the first derivative of the objective function  $f$ . This phrase is worth memorising since it underpins all of the results in this lecture: in every result, we identify conditions under which any solution to the first-order condition must be the unique maximiser.

The following problem allows you to use the result to identify the unique maximiser of an objective function. To apply the result, you need to check that the objective function is strictly concave (sometimes, this is called the ‘second-order condition’): make sure not to neglect this very crucial step.

#### Problem 7.4: Unconstrained optimisation

Let  $f(x) = 10x - x^2$  for all  $x \in \mathbb{R}$ . Identify the maximiser of  $f$  on  $\mathbb{R}$  as well as the value attained by  $f$  at this maximiser.

The result just presented gives us a simple way to solve (some) maximisation problems. For the reasons discussed earlier, it can also be translated into a result about minimisation problems. (Specifically, it implies that, if  $f'(x^*) = 0$  where  $f$  is differentiable and strictly *convex*, then  $x^*$  is the unique minimiser of  $f$  on  $\mathbb{R}$ .) Despite being useful, however, the result is somewhat limited since it only applies to maximisation in a *single* variable  $x$ : in contrast, many situations in economics (e.g., a consumer choosing which bundle of goods to buy) involve optimisation in several variables.

To this end, we now consider a maximisation problem in two variables (i.e., a ‘two-dimensional’ problem). We restrict ourselves to two variables to keep the exposition as simple as possible. However, once one sees how to handle the two-dimensional case, it becomes obvious how to handle an arbitrarily large number of dimensions.

### Result 7.2: Maximisation in two variables

Let  $f: \mathbb{R}^2 \rightarrow \mathbb{R}$  be differentiable and strictly concave. If  $f_x(x^*, y^*) = 0$  and  $f_y(x^*, y^*) = 0$ , then  $(x^*, y^*)$  is the unique maximiser of  $f$  on  $\mathbb{R}^2$ .

As can be seen, this result extends the previous result in a very natural way. Previously, we computed the derivative of  $f$  and set this derivative equal to zero: this was the first-order condition. Now, we compute the derivative of  $f$  with respect to *both* variables and set *both* derivatives equal to zero: this is the new first-order condition. Under the assumption that  $f$  is strictly concave, any solution to this pair of equations must again be the unique maximiser of the objective function.<sup>3</sup>

### Problem 7.5: Bivariate maximisation

Let  $f(x, y) = 4x + 4y - x^2 - y^2$ . Identify the maximiser of  $f$  on  $\mathbb{R}^2$  and the associated maximum value. (You can assume that  $f$  is strictly concave.)

While useful to build one's intuition, neither of these results is very useful in economics since the choice variables are unconstrained. In practice, however, constraints are everywhere. For example:

- A consumer cannot spend more than their income and cannot purchase a negative quantity of a product.
- Likewise, even if we assume that a firm can produce any quantity of a good it likes, we surely will not assume that it can produce negative quantities of a good.

In both of these examples, the constraint set is an interval: for example, in the second example, the firm is restricted to choosing a quantity  $q$  on the interval  $q \geq 0$ . To tackle such problems, one needs to know how to solve maximisation problems where the choice variable  $x$  is restricted to some interval  $[a, b]$ . The next result reveals how to do this, at least in a relatively simple case.

### Result 7.3: Maximisation on an interval

Let  $f: [a, b] \rightarrow \mathbb{R}$  be strictly concave and differentiable. Then

- If  $f$  has a critical point  $x^* \in (a, b)$ ,  $x^*$  is the unique maximiser of  $f$  on  $[a, b]$ .
- Otherwise, the unique maximiser is either  $a$  or  $b$ .

The result reveals that maximisation problems on an interval can be solved in almost the same way as unconstrained maximisation problems. As before, one starts by checking that the objective function is strictly concave so that the result is applicable: a

<sup>3</sup>For the purposes of this course, you do not need to know how to establish the concavity of a two-dimensional function  $f$ . For the curious, however, we note that a sufficient condition is that the function's second partial derivatives satisfy the inequalities  $f_{xx} < 0$ ,  $f_{yy} < 0$  and  $f_{xx}f_{yy} > f_{xy}^2$ .



sufficient condition is that  $f''(x) < 0$  for all  $x$  in the domain. Next, one writes down the first-order condition  $f'(x) = 0$  and checks if this equation has a solution in  $(a, b)$ . If it does, then we have identified the unique maximiser of the objective function (namely, the critical point). Otherwise, we need to evaluate the function at the boundaries  $a$  and  $b$ . If  $f(a) > f(b)$ , then the maximiser is  $a$ ; if  $f(a) < f(b)$ , then the maximiser is  $b$ .

The next problem gives you the opportunity to put this algorithm into practice. As ever, do first check that the function is strictly concave so that our result can be used to solve the problem.

**Problem 7.6: Maximisation on an interval**

Let  $f(x) = 1 - x^2$  for all  $x \in [1, 2]$ . Determine the maximiser of  $f$  on  $[1, 2]$  and the value that  $f$  attains at this point.

## 7.4 An application

To emphasise the usefulness of the theory, we now provide an economic application. Specifically, consider a firm that chooses its price  $p \geq 0$  to maximise its profits. Given that the firm chooses a price  $p$ , it sells  $q(p)$  units. For simplicity, assume that the demand curve  $q$  is linear in the firm's price:

$$q(p) = \begin{cases} \alpha - \beta p, & \text{if } p \leq \frac{\alpha}{\beta} \\ 0 & \text{if } p > \frac{\alpha}{\beta} \end{cases}$$

where  $\alpha > 0$  and  $\beta > 0$ . In other words, demand decreases at a constant rate as the firm increases its price, until demand hits zero at the price  $p = \alpha/\beta$ .

If the firm sells  $q$  units, it must pay a total cost  $C(q)$ . For simplicity, assume constant marginal costs:  $C'(q) = cq$ , for some  $c > 0$ . For reasons that will become clearer, we will also assume that  $c < \alpha/\beta$ , so the marginal cost is 'not too high'. The main reason for this assumption is to avoid triviality: without the assumption, the firm would choose not to sell anything at all.

Given that the firm charges a price  $p \in [0, \alpha/\beta]$ , its profits are

$$\pi(p) = \underbrace{p q(p)}_{\text{revenue}} - \underbrace{c q(p)}_{\text{costs}}$$

Since  $q(p) = \alpha - \beta p$ , we can write profits purely in terms of the price  $p$ , yielding

$$\pi(p) = p(\alpha - \beta p) - c(\alpha - \beta p) = \alpha p - \beta p^2 - \alpha c + c\beta p$$

Thus, maximising profits turns out to be a maximisation problem in a single variable  $p$  subject to the interval constraint  $p \in [0, \alpha/\beta]$ . (Of course, one could allow for prices

above  $\alpha/\beta$ . However, since such prices lead to zero sales and thus zero profits, allowing for such prices would not change the analysis.)

To solve the problem, we consider the first-order condition

$$\pi'(p) = \alpha + \beta c - 2\beta p = 0$$

Rearranging, we identify the critical point

$$p = \frac{\alpha + \beta c}{2\beta}$$

In light of Result 7.3, we need to check if this critical point lies in the interval  $(0, \alpha/\beta)$ . Since  $\alpha$ ,  $\beta$ , and  $c$  are all positive, the critical point is clearly positive. Moreover, under our assumption that  $c < \alpha/\beta$ , one can verify that

$$\frac{\alpha + \beta c}{2\beta} < \frac{\alpha}{\beta}$$

Thus, the critical point does indeed lie in  $(0, \alpha/\beta)$ .

Finally, we need to check that our objective function  $\pi$  is strictly concave in  $p$ . To do this, we compute its second derivative  $\pi''(p) = -2\beta < 0$ . Since the second derivative is negative, the function is strictly concave as required. This implies that

$$p^* = \frac{\alpha + \beta c}{2\beta} = \frac{\alpha}{2\beta} + \frac{c}{2}$$

is indeed the unique maximiser of  $\pi$ .

Two important lessons follow from this formula:

- (1) When the firm's marginal costs  $c$  increase, the term  $c/2$  increases and so the price rises.
- (2) When consumers become less sensitive to price increases (lower  $\beta$ ), the term  $\alpha/2\beta$  increases and so the price rises.

In other words, the firm will choose to charge more if it faces higher (marginal) costs or if it can get away with higher prices without scaring off its customers. Since both of these results seem very intuitive, one might suspect that they do not depend on the specific functional forms that we have assumed. Indeed, these insights persist in more general models of pricing.

Our final problem gives you the opportunity to repeat the previous analysis in a special case. After doing this, you may find it helpful to re-read the more general analysis.

#### Problem 7.7: Profit maximisation

Suppose that, if a firm charges a price  $p \in [0, 5]$ , it will sell  $q = 5 - p$  units of its product. Suppose that the firm has a constant marginal cost of 1. What price maximises the firm's profits?

## 8. UTILITY MAXIMISATION

### 8.1 Motivating examples

#### Problem 8.1: Perfect substitutes

If a consumer buys  $x \geq 0$  packs of ibuprofen and  $y \geq 0$  packs of Nurofen, their utility is  $u(x, y) = x + y$ . A pack of ibuprofen costs £0.50 while a pack of Nurofen costs £2.30. If the consumer has £5 to spend, how many packs of ibuprofen and Nurofen will they buy?

#### Problem 8.2: Optimal teaching

If an academic spends  $t \geq 0$  hours on teaching and  $r \geq 0$  hours on research, their utility is  $u(t, r) = t^{1/4}r^{3/4}$ . In total, they plan to spend 100 hours working over the course of a year. How many hours should they spend on their research?

### 8.2 The consumer's problem

In this lecture, we discuss a model of consumer behaviour that plays an important role in economic theory. In the model, it is assumed that consumers solve a certain *optimisation problem*. Specifically, it is assumed that they solve the problem

$$\max_{(x,y) \in \mathbb{R}_{\geq 0}^2} u(x, y)$$

subject to the constraint

$$xp_X + yp_Y \leq m$$

where  $u(x, y)$  is the utility that the consumer attains given that they consume  $x \geq 0$  units of good X and  $y \geq 0$  units of good Y,  $p_X > 0$  and  $p_Y > 0$  are the prices of the two goods, and  $m > 0$  is the consumer's income.

The various ingredients of this problem can be recast using the language of optimisation that we introduced earlier:

- The consumer's utility function  $u$  is the *objective function*: that is, the function to be optimised.
- The bundle of goods that the consumer purchases  $(x, y)$  is the *choice variable*.
- The set of affordable bundles, i.e.  $\{(x, y) : x \geq 0, y \geq 0, xp_X + yp_Y \leq m\}$ , is the *constraint set*. The non-negativity constraints  $x \geq 0$  and  $y \geq 0$  arise since the consumer cannot buy negative quantities. The budget constraint arises since the value of the consumer's spending  $xp_X + yp_Y$  cannot exceed her income  $m$ .

The *solution* to the consumer's problem is the *optimal demands*  $(x^*, y^*)$  that maximise the consumer's utility given the constraints that they face. Of course,  $(x^*, y^*)$  may depend on the prices  $p_X$  and  $p_Y$  as well as the consumer's income. (For example, the consumer may choose to buy less of good X when its price  $p_X$  rises.) Assuming that the optimal demands are unique, one can thus speak of the optimal demand *functions*  $x^*(p_X, p_Y, m)$  and  $y^*(p_X, p_Y, m)$ . Sometimes, these are called *Marshallian demands*.

Before determining how a consumer would solve this problem, we pause to discuss the model's realism. Notice that:

- The model assumes that the consumer can only buy up to two different goods. Obviously, this is unrealistic: in reality, there is a huge number of different products that a typical consumer could in theory choose to purchase. Fortunately, however, the model can easily be extended to allow for an arbitrary number of different products: we assume two goods here purely for expositional reasons.
- The model assumes that the consumer maximises a 'utility function'. Taken literally, this seems implausible: most people do not make purchasing decisions on the basis of elaborate mathematical calculations. However, it turns out that, under fairly mild conditions, a consumer can be described *as if* they are maximising some utility function.<sup>4</sup> Thus, although not literally realistic, one may be able to view the assumption of utility maximisation as a convenient fiction.
- The model assumes that the consumer pays a constant price for each unit they buy (thus, they spend  $x p_X$  on X and  $y p_Y$  on Y). This is called *linear pricing*. In some situations, this is implausible. For example, in the supermarket, a consumer may find that they need to pay  $p_X$  for one unit of a good but only  $0.5 p_X$  for a second unit ('buy one and get one half price'). To take a different kind of example, a wealthy person who owns most of the stock of a large company would quickly discover that prices are non-linear if they attempted to rapidly sell all of their stock holdings (this is called 'price impact'). In many other situations, however, prices are indeed linear: thus, our model still has wide applicability.
- The model assumes the consumer can buy any real quantities  $x \geq 0$  and  $y \geq 0$ . This means, for example, that the model assumes that a consumer could buy 2.3 (or even  $\sqrt{2}$ ) apples. Of course, this is unrealistic. However, this may be a reasonable approximation in situations where the quantities that the consumer chooses to purchase are reasonably large.

Given these points, one could conclude that, although the model does not adequately describe all purchasing situations, it does provide a reasonable approximation of some purchasing situations.

---

<sup>4</sup> Sufficient conditions are provided by [Debreu \(1954\)](#). A more general version of the 'as if' argument presented here is provided by [Friedman \(1953\)](#).

### 8.3 The MRS method

We now describe a method (or ‘algorithm’) for solving these kinds of optimisation problems. The method is very simple and only involves three steps.

**Step 1.** The first step is to write down what is called the *MRS condition*:

$$\frac{u_x}{u_y} = \frac{p_X}{p_Y}$$

To understand this, recall that  $u_x$  is the partial derivative of  $u$  with respect to  $x$ ; that is, the marginal utility of X. Likewise,  $u_y$  is the partial derivative of  $u$  with respect to  $y$ ; that is, the marginal utility of Y. Meanwhile,  $p_X/p_Y$  is just the ratio of the prices of the two goods. Thus, the condition says that the ratio of the marginal utilities (i.e., the marginal rate of substitution or MRS) should be equal to the ratio of the prices.

**Step 2.** The second step is to write down the budget line equation

$$xp_X + yp_Y = m$$

This just says that the total amount of money that the consumer spends  $xp_X + yp_Y$  should be equal to the consumer’s income  $m$ .

**Step 3.** The third and final step is to solve the two equations simultaneously. Assuming that the method works—which, unfortunately, is not always guaranteed—the solution to the MRS condition and the budget line equation is the utility maximising bundle.

Here is a problem to check that you understand the mechanics of the method.

#### Problem 8.3: The MRS method

Use the MRS method to find the affordable bundle  $(x^*, y^*)$  that maximises a consumer’s utility  $u(x, y) = \sqrt{xy}$  given that  $p_X = 2$ ,  $p_Y = 2$  and  $m = 20$ .

Why does the MRS method (sometimes) work? As a first observation, notice that, assuming that the consumer wants to consume as much as possible, they must spend all of their money. Thus, given the inequality  $xp_X + yp_Y \leq m$ , they will choose an optimal bundle on the boundary  $xp_X + yp_Y = m$ . This is Step 2 of the method.

Although there is no great mystery in Step 2, there is a real question about why the optimal bundle ought to satisfy the MRS condition (Step 1). In the following, we provide two different arguments as to why Step 1 ought to hold. While the first argument is verbal, the second argument uses graphical methods.

**A verbal argument.** To understand the MRS condition, it is helpful to rewrite it in a slightly different form. Using elementary algebra, one can rewrite the condition as

$$\frac{u_x}{u_y} = \frac{p_X}{p_Y} \iff \frac{u_x}{p_X} = \frac{u_y}{p_Y}$$

Since this second condition is logically equivalent to the first, we can justify the first condition by finding a justification for the second.

Imagine now that the consumer decides to spend an extra £1 on good X. If good X costs £2, they would be able to purchase an additional 1/2 units; more generally, given a price  $p_X > 0$ , they will be able to purchase an additional  $1/p_X$  units. How much utility does this give them? Each unit that they buy increases their utility by roughly the marginal utility of X, i.e.  $u_x$ . Thus, an additional  $1/p_X$  units increases their utility by roughly  $1/p_X \times u_x = u_x/p_X$  units. It follows that  $u_x/p_X$  is (roughly) the extra utility the consumer gets from spending an extra £1 on good X.

Similarly, one could imagine that the consumer decides to spend an extra £1 on good Y. Using this extra £1, they would be able to purchase an additional  $1/p_Y$  units of Y. In total, these units would increase their utility by around  $u_y \times (1/p_Y) = u_y/p_Y$  units. Thus,  $u_y/p_Y$  is roughly the extra utility that the consumer would gain from spending an extra £1 on good Y.

These preliminaries out of the way, we are (finally!) ready to make the argument. Suppose that, contrary to the MRS condition,

$$\frac{u_x}{p_X} > \frac{u_y}{p_Y}$$

at the optimal bundle. This means that spending an extra £1 on good X will give the consumer *more* utility than spending an extra £1 on good Y. But then the consumer was not choosing optimally! Indeed, they would have done better if they had spent an extra £1 on X, thus gaining  $u_x/p_X$ , and spent £1 less on Y, thus losing only  $u_y/p_Y < u_x/p_X$ . (This argument assumes that the consumer *could* spend less on Y and thus  $y^* > 0$ ; without this restriction, the argument would break down.)

Similarly, one can argue that, if

$$\frac{u_x}{p_X} < \frac{u_y}{p_Y}$$


then the consumer should spend an extra £1 on Y and spend £1 less on X. After all, the benefits from spending £1 more on Y (namely,  $u_y/p_Y$ ) would exceed the cost of spending £1 less on X (namely,  $u_x/p_X$ ). Again, this argument assumes that the consumer *could* spend less on X if they wanted to, which requires that  $x^* > 0$ .

We have argued that neither of these inequalities could be true at the optimal bundle. This then suggests that, if the consumer is choosing optimally, *and* they buy positive quantities of both goods, then


$$\frac{u_x}{p_X} = \frac{u_y}{p_Y}$$

But this is just the MRS condition (i.e., Step 1 of our method). Thus, we have found a reason why this condition could plausibly arise at the optimal bundle: it reflects the idea that the consumer could not do better by slightly increasing their spending on one good and slightly decreasing their spending on the other.


**A graphical argument.** In addition to arguing for the condition verbally, one can also develop a graphical argument for why the MRS condition ought to hold at the optimal bundle. This argument starts with the observation that, given that the consumer has an increasing utility function, any optimal bundle must lie on the boundary of their budget set (i.e., on their budget line). This budget line is depicted in red below.




For any bundle that the consumer chooses, we can draw the consumer's indifference curve 'through' this bundle. A **first** possibility is that the indifference curve through this bundle is **steeper** than the budget line: for example, consider bundle  $A$  in the figure below. However, this bundle would not be optimal. For example, in the illustration below, the consumer would do strictly better if they instead chose bundle  $B$ , which is both affordable and is on the 'better side' on the consumer's indifference curve.



We should also consider a **second** possibility: perhaps the consumer will choose a bundle like  $A'$  at which point their indifference curve is **flatter** than the budget line. However, such a bundle also cannot be optimal. For example, in the illustration, the consumer would strictly prefer  $B'$  over  $A'$  since  $B'$  is affordable and is on the 'better side' of the indifference curve through  $A'$ .



We have argued that the indifference curve at the optimal bundle cannot be flatter *or* steeper than the budget line. This implies that it must have the *same* slope as the budget line: in other words, the indifference curve must be *tangent* to the budget line. This is illustrated by the figure below.



To complete the argument, recall that, by the implicit function theorem, the slope of the indifference curve is  $-u_x/u_y$ . (You may remember this from the lecture on multivariate functions.) In addition, the slope of the budget line is  $-p_X/p_Y$ . Thus, at the optimal bundle,

$$-\frac{u_x}{u_y} = -\frac{p_X}{p_Y} \iff \frac{u_x}{u_y} = \frac{p_X}{p_Y}$$

But this is again just the MRS condition! Thus, we have found a second argument for why the ratio of marginal utilities ought to equal the ratio of the prices.

## 8.4 General results

In some cases, the MRS method allows one to calculate the utility maximising bundle. In other cases, however, it fails: in theory, the method need not yield any solution, and can even yield a solution that does not maximise the consumer's utility. Ideally, one



would want to develop a theoretical understanding on when exactly the method can be relied on to deliver correct results. We turn to this task in the present section.

The first result shows that the MRS condition arises under relatively mild assumptions. Specifically, we just need to assume that the consumer's utility function is increasing (so they spend all of their money), that it is differentiable (so we can consider small changes in the consumer's purchasing behaviour), and that  $x^* > 0$  and  $y^* > 0$  (so the consumer can always slightly reduce spending on one of the goods if this would make them better off). More formally, we can say that:

#### Result 8.1: Necessary conditions

Suppose that the utility function  $u: \mathbb{R}_{\geq 0}^2 \rightarrow \mathbb{R}$  is increasing<sup>5</sup> and differentiable. Then if an optimal bundle  $(x^*, y^*)$  is an interior solution (i.e.  $x^* > 0$  and  $y^* > 0$ ), it must satisfy the budget line and MRS conditions.

While this result allows us to see where our algorithm comes from, the result only provides a 'necessary condition'. More specifically, the result says that, *if* we have an optimal bundle  $(x^*, y^*)$ , then (given some background conditions) the bundle  $(x^*, y^*)$  must satisfy the budget line and MRS conditions. In some cases, however, other (suboptimal!) bundles may satisfy these conditions as well. Thus, the fact that a bundle satisfies the conditions does *not* automatically imply that it is optimal.

One solution to this issue is to focus on the case where just one bundle satisfies the conditions. We already know (from the previous result) that the optimal bundle will satisfy the conditions. If nothing else satisfies the conditions, we can then turn our statement around and infer that the bundle that satisfies the conditions is optimal. The next result makes this idea more precise.

#### Result 8.2: Sufficient conditions I

Suppose that the utility function  $u: \mathbb{R}_{\geq 0}^2 \rightarrow \mathbb{R}$  is increasing and differentiable. Assume that any optimal bundle is an interior solution. Finally, assume that the budget line and MRS conditions have exactly one solution, namely  $(x', y')$ . Then  $(x', y')$  is the unique utility maximising bundle.

This sufficiency result gives us a recipe for solving the consumer's problem:

- First, check that the utility function is increasing and differentiable.
- Next, try to show that any optimal bundle must be an interior solution, i.e. that  $x^* > 0$  and  $y^* > 0$ . (To do this, one needs to show that any bundle with  $x = 0$  or  $y = 0$  generates strictly lower utility than some alternative bundle.)

<sup>5</sup> There are various ways to cash out the notion of 'increasing' (in a way that makes this result true). One way would be to suppose that, for any  $x' > x$  and  $y' > y$ ,  $u(x', y') > u(x, y)$ .

- If one has succeeded in the previous two steps, use the MRS method. *If* one finds that exactly one bundle satisfies the two conditions, this must be the unique utility maximising bundle.

This recipe is useful and can be used to solve the consumer's problem. For a slightly different approach, one can also impose an additional assumption on the consumer's utility function that guarantees that they have at most one optimal bundle. The appropriate assumption here is *strict quasi-concavity*, which we introduced in a previous lecture. It turns out that, under this additional assumption, any (feasible) solution to the MRS and budget line conditions is the unique utility maximising bundle.

### Result 8.3: Sufficient conditions II

Suppose that the utility function  $u: \mathbb{R}_{\geq 0}^2 \rightarrow \mathbb{R}$  is increasing, differentiable, and strictly quasi-concave. Then any solution to the budget line and MRS equations must be the unique utility maximising bundle.

This result provides us with a slightly different recipe for solving the problem:

- First, check that the utility function is increasing, differentiable, *and strictly quasi-concave*.
- Then use the MRS method. *If* one finds that a feasible bundle satisfies the two conditions, this must be the unique utility maximising bundle.

For those with a good understanding of quasi-concavity, this second recipe may be the simpler one. However, it is only applicable in a smaller class of circumstances: not all utility functions are strictly quasi-concave.

## 8.5 Three examples

To conclude this lecture, we consider three examples. For our first example, suppose that  $u(x, y) = x^\alpha y^{1-\alpha}$  for some  $\alpha \in (0, 1)$ . As you will know, this is called *Cobb-Douglas* utility. Fortunately, one can show that this utility function is increasing, differentiable and strictly quasi-concave. Therefore, Result 8.3 tells us that any solution to the MRS and budget line conditions must be the unique utility maximising bundle. Moreover, it turns out that there *is* a (feasible) bundle that satisfies these conditions. Thus, the MRS method yields the optimal demands: this is illustrated by the next problem.

### Problem 8.4: Cobb-Douglas utility

Suppose that a consumer has the utility function  $u(x, y) = x^\alpha y^{1-\alpha}$  for some  $\alpha \in (0, 1)$ . Find their optimal demands given that  $p_X = 2$ ,  $p_Y = 2$  and  $m = 20$ .

For our second example, suppose that  $u(x, y) = \alpha x + \beta y$  for some  $\alpha, \beta > 0$ . As you

will know, this is called *linear* utility. Unfortunately, this utility function is *not* strictly quasi-concave, so we cannot apply Result 8.3. Also, there is no reason to expect that any optimal bundle is an interior solution, so we cannot apply Result 8.2 either. Of course, one might try to use the MRS method anyway, despite the apparent lack of theoretical justification. However, if one does this, one will find that the MRS condition almost never has a solution.

Although one cannot solve this problem using the MRS method, one can solve it using a *similar* technique. Recall that  $u_x/p_X$  approximates the extra utility that one gains from spending an additional £1 on good X. Similarly,  $u_y/p_Y$  approximates the extra utility that one gains from spending an additional £1 on good Y. This suggests that, if  $u_x/p_X > u_y/p_Y$ , one should spend all of one's money on good X. Likewise, if  $u_x/p_X < u_y/p_Y$ , one should spend all of one's money on good Y. If one combines this with the observations that  $u_x = \alpha$  and  $u_y = \beta$ , one arrives at the following result.

#### Result 8.4: Linear utility

Suppose that the consumer has the utility function  $u(x, y) = \alpha x + \beta y$  for some  $\alpha, \beta > 0$ . Then:

- If  $\alpha/p_X > \beta/p_Y$ , then the consumer will just buy X; i.e.,  $(x^*, y^*) = (m/p_X, 0)$ .
- If  $\alpha/p_X < \beta/p_Y$ , then the consumer will just buy Y; i.e.,  $(x^*, y^*) = (0, m/p_Y)$ .
- If  $\alpha/p_X = \beta/p_Y$ , then any bundle that exhausts the consumer's income is optimal.

The next example allows you to put this result into practice. However, it is sufficiently simple that one could solve it without studying any of the previous theory.

#### Problem 8.5: Linear utility

Suppose that a consumer has the utility function  $u(x, y) = 2x + 3y$ . Find their optimal demands given that  $p_X = 1$ ,  $p_Y = 1$  and  $m = 10$ .

For our third and final example, suppose that  $u(x, y) = \min(\alpha x, \beta y)$  for some  $\alpha, \beta > 0$ . As you will know, this is called *Leontief* utility. Since the  $\min()$  function is not differentiable, one cannot apply either Result 8.2 or 8.3. If one were to try to use the MRS method despite these 'warnings', one would find that the marginal rate of substitution is literally uncomputable since the partial derivatives  $u_x$  and  $u_y$  are not well-defined.

Fortunately, one can solve the problem using a simple insight: if  $(x^*, y^*)$  is optimal, then  $\alpha x^* = \beta y^*$ . To understand this, note that at this bundle, the consumer obtains the utility  $\min(\alpha x^*, \beta y^*) = \alpha x^* = \beta y^*$ . The consumer *could* deviate from this bundle, say by purchasing more of good X and less of good Y. But since the utility is the *minimum* of the inputs, buying more of X would be a waste: in fact, by decreasing

their consumption of Y, their utility would fall. This explains why the consumer should purchase the bundle at which the two inputs of the utility function match, i.e.  $\alpha x^* = \beta y^*$ . In combination with the budget line, this yields the utility maximising bundle.

#### Result 8.5: Leontief utility

Suppose that the consumer has the utility function  $u(x, y) = \min(\alpha x, \beta y)$  for some  $\alpha, \beta > 0$ . Then the optimal demands  $(x^*, y^*)$  are the unique solution to the equations

$$\alpha x^* = \beta y^*, \quad x^* p_X + y^* p_Y = m$$

Although these ideas might sound confusing in the abstract, they will become much clearer in the context of a particular example. In the following example, carefully consider all feasible purchasing decisions that the consumer can make and try to convince yourself that the ‘balanced’ bundle is the best one.

#### Problem 8.6: Leontief utility

Suppose that a consumer has the utility function  $u(x, y) = \min(2x, y)$ . Find their optimal demands given that  $p_X = 1$ ,  $p_Y = 1$  and  $m = 9$ .

## 9. COST MINIMISATION

### 9.1 Motivating examples

#### Problem 9.1: Perfect substitutes

If a firm hires  $L \geq 0$  units of labour and  $K \geq 0$  units of capital, it produces  $K + L$  units of output. Labour costs £10 per unit and capital costs £5 per unit. What is the cheapest way for the firm to produce 10 units of output?

#### Problem 9.2: Cobb-Douglas technology

If a firm hires  $L \geq 0$  units of labour and  $K \geq 0$  units of capital, it produces  $K^{1/5}L^{4/5}$  units of output. Labour costs £20 per unit and capital costs £5 per unit. How much will it cost the firm to produce 10 units of output?

### 9.2 The firm's problem

In this lecture, we consider the problem faced by a firm that wants to *minimise* the cost required to produce a given level of output. This problem is of interest since it is ‘one half’ of the problem faced by a firm when it determines how to maximise its profits. Specifically, in order to determine what level of output  $q$  maximises its profits, a firm *first* needs to determine the total cost of producing  $q$  units.

From a mathematical point of view, the problem is

$$\min_{(K,L) \in \mathbb{R}_{\geq 0}^2} C(K,L) = rK + wL$$

subject to the constraint

$$F(K,L) \geq \bar{q}$$

where  $C(K,L)$  is the total cost incurred by the firm given that they hire  $K \geq 0$  units of capital and  $L \geq 0$  units of labour,  $r > 0$  and  $w > 0$  are the prices of capital and labour, and  $\bar{q} \geq 0$  is the ‘target’ level of output that the firm wants to produce.

As usual, this problem can be described using the language of optimisation that we introduced earlier:

- The *objective function* is the firm's *cost function*  $C$ : this is the function to be optimised (in fact, minimised).
- The *choice variable*  $(K,L)$  is the amount of capital and labour that the firm chooses to hire.

- The *constraint set* is the set of non-negative input bundles that can produce at least  $\bar{q}$  units of output, i.e.

$$\{(K, L) : K \geq 0, L \geq 0, F(K, L) \geq \bar{q}\}$$

If  $(K^*, L^*)$  is the unique input bundle that minimises the firm's cost of producing a given level of output, then  $(K^*, L^*)$  are the firm's *cost-minimising demands*. Obviously,  $K^*$  and  $L^*$  depend on the 'target' level of output  $\bar{q}$ . Thus, one can speak of the *conditional factor demand functions*  $K^*(\bar{q})$  and  $L^*(\bar{q})$ . These then give rise to the *total cost function*  $C(\bar{q}) = rK^*(\bar{q}) + wL^*(\bar{q})$ . This is the (minimum) cost of producing  $\bar{q} \geq 0$  units and is a key ingredient in a firm's broader profit maximisation problem.

**Duality.** Mathematically, this is *not* the same as the 'consumer's problem' we solved earlier. Specifically:

- In the previous case, a consumer maximised a (possibly) *non-linear* utility function subject to a *linear* budget constraint.
- In the present case, a firm minimises a *linear* cost function subject to a (possibly) *non-linear* technological constraint.

Despite this difference, the *techniques* used to solve the problem are identical. Indeed, there is almost nothing in this lecture that is new: as a result, studying cost minimisation should consolidate your understanding of utility maximisation. In this lecture, we will try to emphasise the duality between the two topics by sticking as closely as possible to our previous exposition of the utility maximisation problem.

### 9.3 The MRTS method

We now describe a method for solving cost minimisation problems. This is called the MRTS ('Marginal Rate of Technical Substitution') method.

**Step 1.** The first step is to write down the *MRTS condition*

$$\frac{F_L}{F_K} = \frac{w}{r}$$

To understand this, recall that  $F_L$  is the partial derivative of output with respect to  $L$ ; that is, the marginal product of labour. Similarly,  $F_K$  is the partial derivative of output with respect to  $K$ ; that is, the marginal product of capital. Meanwhile,  $w/r$  is the ratio of the prices of the two inputs. Thus, the condition says that the ratio of the marginal products should be equal to the ratio of the prices.

**Step 2.** The second step is to write down the isoquant equation

$$F(K, L) = \bar{q}.$$

This just says that the total amount that the firm produces, i.e.  $F(K, L)$ , should be equal to the target level of output that the firm wants to produce, i.e.  $\bar{q}$ .

**Step 3.** The third and final step is to solve the equations simultaneously. Under certain assumptions, this will yield the cost-minimising inputs  $(K^*, L^*)$ .

Here is a problem to check that you understand the mechanics of the method.

**Problem 9.3: The MRTS method**

Use the MRTS method to find the cheapest way of producing 20 units of output given the input prices  $r = 2$  and  $w = 2$  and the production function  $F(K, L) = \sqrt{KL}$ .

Why does the MRTS method (sometimes) work? First, notice that, if the production function is increasing in both inputs, it is wasteful to use more inputs than necessary. Thus, if the firm wants to produce *at least*  $\bar{q}$  units, it may as well produce *exactly*  $\bar{q}$  units. This is the ‘isoquant equation’ (Step 2), i.e.

$$F(K, L) = \bar{q}$$

The real question, then, is how to justify the MRTS condition (Step 1). As before, this can be justified using either a verbal or graphical argument.

**A verbal argument.** To understand the MRTS condition, rewrite it as

$$\frac{F_L}{F_K} = \frac{w}{r} \iff \frac{F_L}{w} = \frac{F_K}{r}.$$

As before, we will attempt to justify the first condition by justifying the second (logically equivalent) condition.

Imagine now that the firm spends an extra £1 on labour. If it does this, it will be able to hire another  $1/w$  units of labour. (For example, if the price of labour is  $w = 0.5$ , spending £1 on labour buys you two units.) Each extra unit of labour increases output by around  $F_L$ . Thus, if the firm spends an extra £1 on labour, its output will increase by around  $1/w \times F_L = F_L/w$ .

Using similar reasoning, we observe that, if a firm spends an extra £1 on capital, it will be able to hire another  $1/r$  units of capital. Each unit of capital increases output by around  $F_K$ . Thus, if the firm spends an extra £1 on capital, its output will increase by around  $1/r \times F_K = F_K/r$ .

We can now make the argument. Suppose now that, contrary to the MRTS condition,

$$\frac{F_L}{w} > \frac{F_K}{r}$$

This means that spending an extra £1 on labour would increase output *more* than spending an extra £1 on capital. But then the firm was not choosing optimally! Specifically, it could produce *more* output at the *same* cost by spending £1 more on

labour and £1 less on capital. It could then *cut* its costs by using slightly less of both inputs while still producing its target level of output. Thus, the firm's original input bundle was not the cost minimising one. (This argument assumes that the firm *could* reduce its use of capital, which requires that it was not choosing  $K = 0$  originally.)

Using similar logic, one can argue that, if

$$\frac{F_L}{w} < \frac{F_K}{r}$$


then the firm would gain more output by spending an extra £1 on capital than it would by spending an extra £1 on labour. But, in that case, the firm would have done better by shifting some of its spending from labour to capital. (Again, this assumes that such a shift is possible, since  $L > 0$  originally). Thus, this inequality is *also* not consistent with cost minimisation.

Putting these observations together, we see that, if the firm is choosing optimally *and* uses positive quantities of both inputs, then

$$\frac{F_L}{w} = \frac{F_K}{r}$$

But this is just the MRTS condition. In other words, we have explained why, at the cost minimising bundle, the firm should not be able to increase its output by slightly shifting spending from one input to another.


**A graphical argument.** As in the previous lecture, it is also possible to justify the condition using a graphical argument. The argument starts with the observation that, since the firm does not want to incur needless costs, it will choose some input on the production isoquant corresponding to its target output level  $\bar{q}$ . This is illustrated by the isoquant below and describes the 'menu' from which the firm must choose.




For any input bundle that the firm chooses, we can draw the isocost line through this bundle: this represents all the  $(K, L)$  choices that have the same cost as the original bundle. A **first** possibility is that the firm chooses a bundle whose corresponding




isocost line is **steeper** than the isoquant. For example, consider the input bundle  $A$  which is depicted (along with its isocost line) below. However, this choice could not be optimal. For example, in the figure, the firm could produce the same level of output at a lower cost by choosing instead a bundle like  $B$ , which lies on the 'cheaper' side of  $A$ 's isocost line.



We now consider a **second** possibility: perhaps the firm will choose a bundle  $A'$  whose isocost line is **flatter** than the slope of the isoquant (see the figure below). However, this input bundle could also not be optimal. Specifically, the firm would do better by instead choosing a bundle like  $B'$  which produces the same level of output (since it lies on the same isoquant) but lies on the 'cheaper' side of  $A'$ 's isocost line.



We have argued that the isocost line going through the optimal input bundle cannot be flatter *or* steeper than the isoquant. It thus follows that the isocost line must have the same slope as the isoquant: in other words, they must be tangential. The figure below illustrates such a situation.



To conclude the argument, recall that the slope of the isoquant is  $-F_L/F_K$ . (This was derived in the lecture on multivariate functions.) Meanwhile, the slope of the isocost line is  $-w/r$ . Thus, at the optimal bundle

$$-\frac{F_L}{F_K} = -\frac{w}{r} \iff \frac{F_L}{F_K} = \frac{w}{r},$$

which is exactly the MRTS condition. Thus, we have found a second argument for why the ratio of marginal products ought to equal the ratio of prices.

## 9.4 General results

While the MRTS method sometimes identifies the cost-minimising bundle, sometimes it does not. Ideally, one would obtain a good theoretical understanding of when exactly the method can be expected to work. Fortunately, the MRTS method turns out to work under exactly the same conditions as the MRS method. Thus, this section will again provide some useful ‘revision’ of the ideas from the last lecture.

Our first result says that the MRTS condition arises for *any* production function that is increasing, differentiable and gives rise to an ‘interior solution’. The interiority assumption should not be neglected: as we explained earlier, there is no reason to expect the condition to hold if the firm has chosen just to use one of the two inputs.

### Result 9.1: Necessary conditions

Suppose that the production function  $F: \mathbb{R}_{\geq 0}^2 \rightarrow \mathbb{R}$  is increasing and differentiable. Then if a cost-minimising bundle is an interior solution (i.e.,  $K^* > 0$  and  $L^* > 0$ ), it must satisfy the isoquant and MRTS equations.

As usual, this result just provides a necessary condition: it says that (given appropriate background assumptions), any optimal bundle must satisfy the MRTS condition. However, other (suboptimal!) bundles might satisfy the condition as well. Fortunately, one does not need to worry about this if one determines that there is just one bundle

that satisfies the MRTS (and isoquant) equations: in such a case, the bundle that one has calculated *must* be the optimal bundle. The next result makes this more precise.

#### Result 9.2: Sufficient conditions I

Suppose that the production function  $F: \mathbb{R}_{\geq 0}^2 \rightarrow \mathbb{R}$  is increasing and differentiable. Assume that any cost-minimising bundle is an interior solution. Finally, assume that there is exactly one solution to the isoquant and MRTS equations, namely  $(K', L')$ . Then  $(K', L')$  is the unique cost-minimising bundle.

Alternatively, one could instead try to establish that the production function is strictly quasi-concave. Intuitively, this means that the isoquants generated by the production function will have the decreasing and convex shapes illustrated in the figures above. If one is able to establish this, and does find a solution to the MRTS and isoquant equations, then one has successfully computed the cost-minimising input bundle.

#### Result 9.3: Sufficient conditions II

Suppose that the production function  $F: \mathbb{R}_{\geq 0}^2 \rightarrow \mathbb{R}$  is increasing, differentiable, and strictly quasi-concave. Then any solution to the isoquant and MRTS equations must be the unique cost-minimising bundle.

### 9.5 Three examples

We conclude with three examples. The first example is the *Cobb-Douglas* production function  $F(K, L) = K^\alpha L^{1-\alpha}$ ,  $\alpha \in (0, 1)$ . You may recall that the MRS method is guaranteed to work when applied to a Cobb-Douglas *utility* function. For the exact same reasons, the MRTS method is guaranteed to work when applied to a Cobb-Douglas *production* function.

Our second example is the *linear* production function  $F(K, L) = aK + bL$ ,  $a, b > 0$ . You may recall that the MRS method generally fails in the context of linear utility functions. For exactly the same reasons, the MRTS method generally fails in the context of linear production functions. However, the cost-minimising bundle can be identified using an ‘MRTS-like’ strategy. Specifically, one starts by computing  $F_L/w = b/w$ : this is the extra output one gets from spending £1 on labour. Then, one computes  $F_K/r = a/r$ : this is the extra output one gets from spending £1 on capital. One then determines which of these values is larger: to minimise costs, one will exclusively produce output using this single factor. The next result spells this out more formally.

#### Result 9.4: Linear technology

Suppose that a firm must produce  $\bar{q} > 0$  units of output and has the production function  $F(K, L) = aK + bL$  for some  $a, b > 0$ . Then:

- If  $a/r > b/w$ , the firm uses only capital.
- If  $a/r < b/w$ , the firm uses only labour.
- If  $a/r = b/w$ , then any  $(K, L)$  bundle that produces  $\bar{q}$  units of output is optimal.

#### Problem 9.4: Linear technology

Suppose that a firm has the production function  $F(K, L) = 2K + 3L$ . Find its cheapest way of producing 10 units of output given that  $r = 1$  and  $w = 4$ .

As our third and final example, we consider the *Leontief* production function  $F(K, L) = \min(aK, bL)$ , where  $a, b > 0$ . Since the MRS method cannot be used to optimise Leontief utility functions, it should come as no surprise that the MRTS method cannot be used to optimise Leontief production functions. However, one can again solve the problem by equating the two arguments of the function. In this case, this means setting  $aK = bL$ . As the next result makes clear, the bundle that satisfies this ‘balance condition’ along with the requirement that  $F(K, L) = \bar{q}$  is the cost-minimising bundle.

#### Result 9.5: Leontief technology

Suppose that a firm must produce  $\bar{q} > 0$  units of output and has the production function  $F(K, L) = \min(aK, bL)$  for some  $a, b > 0$ . Then the cost-minimising input demands  $(K^*, L^*)$  are the unique solution to the equations

$$aK^* = bL^*, \quad \min(aK^*, bL^*) = \bar{q}$$

#### Problem 9.5: Leontief technology

Suppose that a firm has the production function  $F(K, L) = \min(3K, 2L)$ . Find its cheapest way of producing 18 units of output given that  $r = 2$  and  $w = 3$ .

## 10. INTEGRATION

### 10.1 Motivating examples

#### Problem 10.1: From marginal to total costs

If a firm produces  $q \geq 0$  units of output, its marginal cost of production is

$$MC(q) = 2q + 10.$$

Can we recover the *total* cost of producing  $q$  units from this marginal cost function?

#### Problem 10.2: Consumer surplus

The inverse demand curve for a good is

$$p(q) = 20 - q \quad \text{for } 0 \leq q \leq 20.$$

If the good is sold at a price  $p^* = 10$ , how much surplus do consumers obtain?

### 10.2 Indefinite integrals

Previously, we learned how to go from a function  $f$  to its derivative  $f'$ . This process is called *differentiation*. The topic of this week, *integration*, is the operation that reverses this process: roughly speaking, it allows us to move from the derivative  $f'$  back to the function  $f$ .

To develop these ideas, we first introduce the concept of an *antiderivative*.

#### Definition 10.1

Let  $f: \mathbb{R} \rightarrow \mathbb{R}$ . A function  $F$  is called an *antiderivative* of  $f$  if, for all  $x \in \mathbb{R}$ ,

$$F'(x) = f(x)$$

In other words, to say that  $F$  is an antiderivative of  $f$  is to say that differentiating  $F$  yields  $f$ . To illustrate, recall that differentiating  $x^2$  yields  $2x$ . In other words,  $2x$  is the *derivative* of  $x^2$ . Turning things around, this means that  $x^2$  is an *antiderivative* of  $2x$ .

In the previous example, we were careful to say that  $x^2$  is *an* antiderivative of  $f$ , not *the* antiderivative. Indeed,  $2x$  has many antiderivatives: for example, the functions  $x^2 + 3$ ,  $x^2 + 17$  and  $x^2 - 417$  all differentiate to  $2x$ . More generally, if  $F$  is an antiderivative of  $f$ , then  $F + c$  must be an antiderivative of  $f$  as well, where  $c$  is an arbitrary constant: the reason is that the constant vanishes upon differentiation. We call the family of

antiderivatives of a given function that function's *indefinite integral*.

### Definition 10.2

The *indefinite integral* of a function  $f$ , written  $\int f(x) dx$ , describes the family of antiderivatives of  $f$ . That is, if  $F$  is an antiderivative of  $f$ , then

$$\int f(x) dx = F(x) + c$$

Here are two examples to test your understanding of indefinite integrals. To solve the problems, ask yourself which functions differentiate to  $f(x) = 2$  and  $g(x) = x$ .

### Problem 10.3: Two simple integrals

For any  $x \in \mathbb{R}$ , let  $f(x) = 2$  and  $g(x) = x$ . Compute the indefinite integrals  $\int f(x) dx$  and  $\int g(x) dx$ .

## 10.3 Some integration rules

We previously stated various rules that allow one to go from a function  $f$  to its derivative  $f'$ . For example, we said that the function  $f(x) = x^k$  differentiates to  $f'(x) = kx^{k-1}$ . Since integration is the reverse process, we can reverse each of these rules and restate them as rules of integration.<sup>6</sup>

Our first rule says that, just as differentiation is a linear operator, integration is a linear operator. Similarly to before, this rule blends two ideas. The first idea is that the integral of a sum is the sum of integrals, i.e.  $\int f(x) + g(x) dx = \int f(x) dx + \int g(x) dx$ . The second idea is that the integral of a function multiplied by a constant  $a$  is just  $a$  times the integral of the function:  $\int af(x) dx = a \int f(x) dx$ . Putting these two ideas together, we obtain the following.

### Result 10.1: Linearity of integration

Let  $f$  and  $g$  be functions, and let  $a, b \in \mathbb{R}$ . Then

$$\int af(x) + bg(x) dx = a \int f(x) dx + b \int g(x) dx$$

Our second rule reverses our rule for differentiating power functions, i.e. that the function  $f(x) = x^k$  has derivative  $f'(x) = kx^{k-1}$ . Indeed, you should confirm the validity of the next rule by differentiating the right hand side and confirming that it is indeed an antiderivative of  $x^k$ .

<sup>6</sup> To simplify the exposition, we slightly specialise our rules for logarithms and exponents: instead of working with arbitrary bases, we restrict ourselves to base  $e$ .

### Result 10.2: Integrating a power function

If  $k$  is a real number with  $k \neq -1$ , then

$$\int x^k dx = \frac{x^{k+1}}{k+1} + c$$

Our third rule reverses our rule for differentiating exponentials. You may recall that, if  $f(x) = e^x$ , then  $f'(x) = e^x$ . In other words, nothing happens if you differentiate  $e^x$ . For this reason, nothing happens if you integrate  $e^x$ : one just needs to remember to add the constant of integration  $c$ .

### Result 10.3: Integrating the exponential function

For any  $x \in \mathbb{R}$ ,

$$\int e^x dx = e^x + c$$

Our final rule reverses our rule for differentiating the natural logarithm. You may recall that, if one differentiates the natural logarithm  $\ln(x)$ , one obtains  $1/x$ . For this reason, an antiderivative of  $1/x$  is the natural logarithm  $\ln(x)$ . To describe the family of antiderivatives, one again just needs to add the constant of integration.

### Result 10.4: Integrating $1/x$

For any  $x > 0$ ,

$$\int \frac{1}{x} dx = \ln(x) + c$$

Here are some problems to test your understanding of the rules.


### Problem 10.4: Indefinite integrals

Compute the following integrals:

- (i)  $\int 12x^5 + 2x dx$
- (ii)  $\int 4e^x + x dx$
- (iii)  $\int \frac{2}{x} + e^x dx$  (assume  $x > 0$ )

## 10.4 Definite integrals

So far, we have discussed *antiderivatives*. We now turn to an apparently unconnected problem: how to compute the *area under a curve (and above the horizontal axis)*. For example, how would you compute the shaded area that is depicted below?



This area is an example of a *definite* integral. More precisely, we say that:


**Definition 10.3**

The *definite integral* of  $f$  from  $a$  to  $b$ ,

$$\int_a^b f(x) dx,$$

is the signed area under the graph of  $f$  between  $a$  and  $b$ .

Careful readers will note that, when defining the definite integral, we have slipped from talking about the ‘area under a curve’ to a *signed area*. What’s the difference? If a function always takes non-negative values, there is no difference: the signed area is simply the area under the curve (and above the horizontal axis). However, if the function can take negative values, one first computes the area under the curve when the function takes positive values and then *subtracts* the area above the curve whenever the function takes negative values. For example, in the figure below, the signed area from  $a$  to  $b$  is the area below the curve (shaded in pink) *minus* the area above the curve (shaded in grey).



Suppose then that one wants to compute a signed area. In some simple cases, one can do this using elementary geometry: for example, one can compute the area under a horizontal line using the formula for the area of a rectangle. In other cases, however,



it is far less clear how to proceed: for example, imagine trying to calculate the area under a complex and highly non-linear function. The next result provides us with an amazingly simple way of solving such problems.

**Result 10.5: Fundamental theorem of calculus (simplified)**

Let  $f$  be a function and let  $F$  be an antiderivative of  $f$ . Then

$$\int_a^b f(x) dx = F(b) - F(a)$$

This result reveals a deep link between definite and indefinite integrals. Specifically, the result tells us that one can compute the signed area under the curve, i.e.  $\int_a^b f(x) dx$ , by evaluating an antiderivative  $F$  at the bounds  $a$  and  $b$ . Thus, all one needs to do to calculate the signed area is to find an antiderivative of the function. The computational value of this result is hard to overestimate.

Although this result is very useful, it might also seem quite mysterious: what does the signed area under a curve have to do with an antiderivative  $F$ ? The key intuition comes from considering the area under a curve from  $a$  to some variable upper bound  $x$ . As  $x$  slightly increases, the area under the curve increases by  $f(x)$ : after all, this is the ‘extra slice’ of area that we are adding on. Thus, the rate of change (or derivative) of the area is the function  $f$ . This then means that the area is an antiderivative of the function. Thus, the notions of areas and antiderivatives are closely connected.

The next problem gives you the opportunity to put this result into practice. When applying the result, it is sometimes helpful to introduce the notation  $[F(x)]_a^b = F(b) - F(a)$ . Given this notation, the result says that

$$\int_a^b f(x) dx = [F(x)]_a^b$$

**Problem 10.5: Area under a straight line**

Consider the function  $f(x) = x$ .

- (i) Use the fundamental theorem of calculus to compute

$$\int_0^2 x dx$$

- (ii) Confirm your answer using geometric methods.

## 10.5 Economic applications

We conclude this week with two economic applications.

**Marginal costs.** As our first application, we consider whether we can recover the *total* costs incurred by a firm from its *marginal* costs of production. Recall that the marginal cost  $MC$  can be defined as the derivative of the total cost  $C$  with respect to quantity: that is,  $MC(q) = C'(q)$ . Since the marginal cost function is the *derivative* of the total cost function, it follows that the total cost function is an *antiderivative* of the marginal cost function. Thus, one can compute total costs by integrating marginal costs.

To make these results more precise, we use the fundamental theorem of calculus. By the theorem,

$$\int_0^q MC(t)dt = C(q) - C(0)$$

Therefore,

$$C(q) = C(0) + \int_0^q MC(t) dt$$

This says that the total cost  $C(q)$  is the cost of producing nothing, i.e.  $C(0)$ , plus the integral of all the marginal costs from 0 to  $q$ . We can interpret  $C(0)$  as a *fixed cost*: it is a cost that must be paid no matter how much the firm produces (think, e.g., of the rental payments that a firm must pay for its office). The result thus says that, if we know a firm's fixed cost  $C(0)$ , and we know its marginal costs  $MC(t)$ , we can infer its total cost of production.

Here is a quick problem to test your understanding of this application.

#### Problem 10.6: Recovering $C(q)$

A firm has the marginal cost function

- $MC(q) = 2q + 10$
- (i) Compute  $\int_0^q MC(t) dt$ .
  - (ii) Suppose the fixed cost is  $C(0) = 50$ . Use (i) to find the total cost function  $C(q)$ .
  - (iii) What is the total cost of producing 10 units?


**Consumer surplus.** For our second application, we consider the problem of determining how much ‘value’ a group of consumers derive from a particular product. Assume that the product is produced by a firm, which charges a price  $p \geq 0$  and thus sells  $q(p)$  units. The function  $q$  is called the *demand curve*; assuming that it is decreasing, we can invert it to obtain the inverse demand curve  $p$ . Intuitively, the inverse demand curve tells us that, given that consumers want to buy  $q \geq 0$  units, the firm must be charging a price of  $p(q)$ .

We now make use of an amazing fact from economic theory: under certain assumptions, *the inverse demand curve at a quantity  $q$  can be viewed as the maximum that a consumer would pay for the  $q$ th unit*. In other words, the inverse demand curve is a ‘willingness

to pay curve' in disguise. Given this fact, it follows that the *surplus* that the consumer who buys the  $q$ th unit obtains from their purchase is  $p(q) - p^*$ . (For example, if the consumer is willing to pay up to £10 for the unit, and they are charged a price of £5, then their surplus is  $p(q) - p^* = £10 - £5 = £5$ .) Taken together, the surplus that the consumers obtain is therefore

$$CS = \int_0^{q^*} p(q) - p^* dq$$

At a conceptual level, this tells us that consumer surplus can be computed by finding an antiderivative of the inverse demand curve and evaluating this antiderivative at the bounds 0 and  $q^*$ . Geometrically, this is the area between the demand curve and the horizontal line at price  $p^*$ , from 0 to  $q^*$ ; see the figure below for an illustration.



Our final problem allows you to compute the consumer's surplus in a simple example. Although the example is stylised, the logic underlying the computation continues to inform research and policy evaluation within economics (see, e.g., [Cohen et al., 2016](#)).

#### Problem 10.7: Consumer surplus with linear demand

The inverse demand curve for a good is

$$p(q) = 20 - q \quad \text{for } 0 \leq q \leq 20.$$

Suppose that the firm charges the price  $p^* = 10$  and so the consumers purchase  $q^* = 10$  units.

- (i) Write consumer surplus as a definite integral.
- (ii) Compute this integral.
- (iii) Check your answer using basic geometry.

## REFERENCES

- Arrow, K. J. and Debreu, G. (1954). Existence of an equilibrium for a competitive economy. Econometrica, 22(3):265–290.
- Cohen, P., Hahn, R., Hall, J., Levitt, S., and Metcalfe, R. (2016). Using big data to estimate consumer surplus: The case of Uber. Technical report, National Bureau of Economic Research.
- Debreu, G. (1954). Representation of a preference ordering by a numerical function. Decision processes, 3:159–165.
- Debreu, G. (1991). The mathematization of economic theory. The American Economic Review, 81(1):1–7.
- Friedman, M. (1953). The methodology of positive economics.
- Jevons, W. (1871). The Theory of Political Economy. Macmillan and co.
- Krugman, P. (2009). How did economists get it so wrong? The New York Times, 2(09):2009.
- Levinovitz, A. J. (2016). The new astrology: By fetishising mathematical models, economists turned economics into a highly paid pseudoscience. Aeon.
- Lucas Jr, R. E. (2001). Professional memoir. Lecture given in the Nobel Economists Lecture Series at Trinity University, San Antonio, Texas.
- Mas-Colell, A., Whinston, M. D., and Green, J. R. (1995). Microeconomic theory, volume 1. Oxford University Press New York.
- McCloskey, D. N. (1998). The rhetoric of economics. Univ of Wisconsin Press.
- Rodrik, D. (2015). Economics rules: The rights and wrongs of the dismal science. WW Norton & Company.
- Romer, P. M. (2015). Mathiness in the theory of economic growth. American Economic Review, 105(5):89–93.
- Rothbard, M. N. (2009). A note on mathematical economics. Mises Daily.
- Smith, A. (1776). The Wealth of Nations. W. Strahan and T. Cadell, London.