An Infinite Descent into Pure Mathematics

BY CLIVE NEWSTEAD

A free PDF copy is available on the book's website:
https://infinitedescent.xyz

Contents

Preface

Hello, and thank you for taking the time to read this quick introduction to the book! I would like to begin with an apology and a warning:

This book is still under development!

That is to say, there are some sections that are incomplete, as well as other sections which are currently much more terse than I would like them to be.

The most recent version is freely available for download from the following website:

<https://infinitedescent.xyz>

As the book is undergoing constant changes, I advise that you do not print it in its entirety—if you must print anything, then I suggest that you do it a few pages at a time, as required.

This book was designed with *inquiry* and *communication* in mind, as they are central to a good mathematical education. One of the upshots of this is that there are many exercises throughout the book, requiring a more active approach to learning, rather than passive reading. These exercises are a fundamental part of the book, and should be completed even if not required by the course instructor. Another upshot of these design principles is that solutions to exercises are not provided—a student seeking feedback on their solutions should speak to someone to get such feedback, be it another student, a teaching assistant or a course instructor.

Navigating the book

This book need not be read from front to back, so if you are using it to teach a course or for self-study, then a certain degree of customisation is possible. With that said, some of the content in the book is fundamental to further study in pure mathematics and should be included in any course serving as a first introduction to proof-based mathematics—at the very least, Chapter 0, Chapter 1, Chapter 2 and Section 3.1 should be covered, and Sections 5.1 and 6.1 are highly recommended.

Within each chapter, subsequent sections depend on the sections before them—as such if you want to learn about modular arithmetic (Section 4.3), then you should first cover division (Section 4.1) and prime numbers (Section 4.2).

The following table indicates prerequisites between sections. If a section is listed as an *essential* prerequisite, it means that concepts from that section are central; if it is listed as a *recommended* prerequisite, it means that concepts from that section are required to fully understand several of the proofs and examples; and if it is listed as a *useful* prerequisite, it means that concepts from that section appear in some examples or provide useful background for understanding the material.

Section	Essential	Recommended	Useful
1.1	0		
2.1	1.3		
3.1	1.3		
3.2	2.3		
4.1	1.3	3.1	2.1, 2.2
5.1	2.2	4.3	3.2
6.1	3.2		
6.3	5.3		
7.1	1.3		5.2
7.2	2.2	6.1	
8.1	3.3	6.1, 7.3	

Note that prerequisites are cumulative, so, for example, in order to cover the material in Section 6.2, you should first work through Sections 1.1–1.3, 2.1–2.3, 3.1, 3.2 and 6.1.

What the numbers, colours and symbols mean

Much of the material in this book is broken into enumerated items which, broadly speaking, fall into one of four categories: **results** (often followed by **proofs**), **definitions**, **examples**, **exercises** and **remarks**. These items are colour-coded as indicated and are enumerated according to their section—for example, Theorem 3.1.14 is in Section 3.1. Definitions and theorems (important results) appear in a ⬚box⬚.

You will also encounter the symbols □, ◁ and ⋆, whose meanings are as follows:

□ **End of proof.** It is standard in mathematical documents to identify when a proof has ended by drawing a small square or by writing '*Q.E.D.*' (The latter stands for *quod erat demonstrandum*, which is Latin for *what was to be shown*.)

◁ **End of item.** This is *not* a standard usage, and is included only to help you to identify when an item has finished and the main content of the book continues.

⋆ **Optional content.** Sections, exercises, results and proofs marked with this symbol can be skipped over. Usually this is because the content is very challenging, or is technical in a way that is mathematically necessary but educationally not very important.

Licence

This book is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 (CC BY-NC-SA 4.0) licence. This means you're welcome to share this book, provided that you give credit to the author, and that any copies or derivatives of this book are released under the same licence, are freely available and are not for commercial use. The full licence is available at the following link:

http://creativecommons.org/licenses/by-nc-sa/4.0/

Comments and corrections

Any feedback, be it from students, teaching assistants, instructors or any other readers, would be very much appreciated. Particularly useful are corrections of typographical errors, suggestions for alternative ways to describe concepts or prove theorems, and requests for new content (e.g. if you know of a nice example that illustrates a concept, or if there is a relevant concept you wish were included in the book).

Such feedback can be sent to me by email (`cnewstead@northwestern.edu`).

Acknowledgements

When I reflect on the time I have spent writing this book, I am overwhelmed by the number of people who have had some kind of influence on their content.

This book would never have come to exist were it not for Chad Hershock's course 38-801 *Evidence-Based Teaching in the Sciences*, which I took in Fall 2014 as a graduate student at Carnegie Mellon University. His course heavily influenced my approach to teaching, and it motivated me to write this book in the first place. Many of the pedagogical decisions I made when writing this book were informed by research that I was exposed to as a student in Chad's class.

The legendary Carnegie Mellon professor, John Mackey, has been using this book (in various forms) as course notes for 21-128 *Mathematical Concepts and Proofs* and 15-151 *Mathematical Foundations of Computer Science* since Fall 2016. His influence can be felt throughout the book: thanks to discussions with John, many proofs have been reworded, sections restructured, and explanations improved. As a result, there is some overlap between the exercises in this book and the questions on his problem sheets. I am extremely grateful for his ongoing support.

Steve Awodey, who was my doctoral thesis advisor, has for a long time been a source of inspiration for me. Many of the choices I made when choosing how to present the material in this book are grounded in my desire to do mathematics *the right way*—it was this desire that led me to study category theory, and ultimately to become Steve's PhD student. I learnt a great deal from him and I greatly appreciated his patience and flexibility in helping direct my research despite my busy teaching schedule and extracurricular interests (such as writing this book).

Perhaps unbeknownst to them, many insightful conversations with the following people have helped shape the material in this book in one way or another: Jeremy Avigad, Deb Brandon, Heather Dwyer, Thomas Forster, Will Gunther, Kate Hamilton, Jessica Harrell, Bob Harper, Brian Kell, Marsha Lovett, Ben Millwood, David Offner, Ruth Poproski, Hilary Schuldt, Gareth Taylor, Katie Walsh, Emily Weiss and Andy Zucker.

Chapter 0

Getting started

Before we can start proving things, we need to eliminate certain kinds of state-ments that we might try to prove. Consider the following statement:

This sentence is false.

Is it true or false? If you think about this for a couple of seconds then you'll get into a bit of a pickle.

Now consider the following statement:

The happiest donkey in the world.

Is it true or false? Well it's not even a sentence; it doesn't make sense to even *ask* if it's true or false!

Clearly we'll be wasting our time trying to write proofs of statements like the two listed above—we need to narrow our scope to statements that we might actually have a chance of proving (or perhaps refuting)! This motivates the following (informal) definition.

Definition 0.1
A **proposition** is a statement to which it is possible to assign a **truth value** ('true' or 'false'). If a proposition is true, a **proof** of the proposition is a logically valid argument demonstrating that it is true, which is pitched at such a level that a member of the intended audience can verify its correctness.

Thus the statements given above are not propositions because there is no possible way of assigning them a truth value. Note that, in Definition 0.1, all that matters is that it *makes sense* to say that it is true or false, regardless of whether it actually *is* true or false—the truth value of many propositions is unknown, even very simple ones.

Exercise 0.2
Think of an example of a true proposition, a false proposition, a proposition whose truth value you don't know, and a statement that is not a proposition.   ◁

Results in mathematical papers and textbooks may be referred to as *propositions*, but they may also be referred to as *theorems*, *lemmas* or *corollaries* depending on their intended usage.

- A **proposition** is an umbrella term which can be used for any result.

- A **theorem** is a key result which is particularly important.

- A **lemma** is a result which is proved for the purposes of being used in the proof of a theorem.

- A **corollary** is a result which follows from a theorem without much additional effort.

These are not precise definitions, and they are not meant to be—you could call every result a *proposition* if you wanted to—but using these words appropriately helps readers work out how to read a paper. For example, if you just want to skim a paper and find its key results, you'd look for results labelled as *theorems*.

It is not much good trying to prove results if we don't have anything to prove results about. With this in mind, we will now introduce the *number sets* and prove some results about them in the context of four topics, namely: division of integers, number bases, rational and irrational numbers, and polynomials. These topics will provide context for the rest of the material in Chapters 1 and 2.

We will not go into very much depth in this chapter. Rather, think of this as a warm-up exercise—a quick, light introduction, with more proofs to be provided in the rest of the book.

Number sets

Later in this chapter, and then in much more detail in Section 2.1, we will encounter the notion of a *set*; a set can be thought of as being a collection of objects. This seemingly simple notion is fundamental to mathematics, and is so involved that we will not treat sets formally in this book. For now, the following definition will suffice.

> **Definition 0.3** (to be revised in Definition 2.1.1)
> A **set** is a collection of objects. The objects in the set are called **elements** of the set. If X is a set and x is an object, then we write $x \in X$ (LaTeX code: `x \in X`) to denote the assertion that x is an element of X.

The sets of concern to us first and foremost are the *number sets*—that is, sets whose elements are particular types of *number*. At this introductory level, many details will be temporarily swept under the rug; we will work at a level of precision which is appropriate for our current stage, but still allows us to develop a reasonable amount of intuition.

In order to define the number sets, we will need three things: an infinite line, a fixed point on this line, and a fixed unit of length.

So here we go. Here is an infinite line:

The arrows indicate that it is supposed to extend in both directions without end. The points on the line will represent numbers (specifically, *real numbers*, a misleading term that will be defined in Definition 0.26). Now let's fix a point on this line, and label it '0':

This point can be thought of as representing the number zero; it is the point against which all other numbers will be measured. Finally, let's fix a unit of length:

This unit of length will be used, amongst other things, to compare the extent to which the other numbers differ from zero.

Definition 0.4
The above infinite line, together with its fixed zero point and fixed unit length, constitute the (**real**) **number line**.

We will use the number line to construct five sets of numbers of interest to us:

- The set $\mathbb{N}$ of *natural numbers*—Definition 0.5;

- The set $\mathbb{Z}$ of *integers*—Definition 0.11;

- The set $\mathbb{Q}$ of *rational numbers*—Definition 0.25;

- The set $\mathbb{R}$ of *real numbers*—Definition 0.26; and

- The set $\mathbb{C}$ of *complex numbers*—Definition 0.32.

Each of these sets has a different character and is used for different purposes, as we will see both later in this chapter and throughout this book.

Natural numbers ($\mathbb{N}$)

The *natural numbers* are the numbers used for counting—they are the answers to questions of the form 'how many'—for example, I have *three* uncles, *one* dog and *zero* cats.

Counting is a skill humans have had for a very long time; we know this because there is evidence of people using tally marks tens of thousands of years ago. Tally marks provide one method of counting small numbers: starting with nothing, proceed through the objects you want to count one by one, and make a mark for every object. When you are finished, there will be as many marks as there are objects. We are taught from a young age to count with our fingers; this is another instance of making tally marks, where now instead of making a mark we raise a finger.

Making a tally mark represents an *increment* in quantity—that is, adding one. On our number line, we can represent an increment in quantity by moving to the right by the unit length. Then the distance from zero we have moved, which is equal to the number of times we moved right by the unit length, is therefore equal to the number of objects being counted.

Definition 0.5

The **natural numbers** are represented by the points on the number line which can be obtained by starting at 0 and moving right by the unit length any number of times:

In more familiar terms, they are the *non-negative whole numbers*. We write $\mathbb{N}$ (LaTeX code: `\mathbb{N}`) for the set of all natural numbers; thus, the notation '$n \in \mathbb{N}$' means that n is a natural number.

The natural numbers have very important and interesting mathematical structure, and are central to the material in Chapter 3. A more precise characterisation of the natural numbers will be provided in Section 3.1, and a mathematical construction of the set of natural numbers can be found in Section B.1 (see Construction B.2.5). Central to these more precise characterisations will be the notions of 'zero' and of 'adding one'—just like making tally marks.

Aside

Some authors define the natural numbers to be the *positive* whole numbers, thus excluding zero. We take 0 to be a natural number since our main use of the natural

numbers will be for counting finite sets, and a set with nothing in it is certainly finite! That said, as with any mathematical definition, the choice about whether $0 \in \mathbb{N}$ or $0 \notin \mathbb{N}$ is a matter of taste or convenience, and is merely a convention—it is not something that can be proved or refuted.   ◁

Number bases

Writing numbers down is something that may seem easy to you now, but it likely took you several years as a child to truly understand what was going on. Historically, there have been many different systems for representing numbers symbolically, called *numeral systems*. First came the most primitive of all, tally marks, appearing in the Stone Age and still being used for some purposes today. Thousands of years and hundreds of numeral systems later, there is one dominant numeral system, understood throughout the world: the **Hindu–Arabic numeral system**. This numeral system consists of ten symbols, called *digits*. It is a *positional* numeral system, meaning that the position of a symbol in a string determines its numerical value.

In English, the *Arabic numerals* are used as the ten digits:

$$0 \quad 1 \quad 2 \quad 3 \quad 4 \quad 5 \quad 6 \quad 7 \quad 8 \quad 9$$

The right-most digit in a string is in the units place, and the value of each digit increases by a factor of ten moving to the left. For example, when we write '2812', the left-most '2' represents the number two thousand, whereas the last '2' represents the number two.

The fact that there are ten digits, and that the numeral system is based on powers of ten, is a biological accident corresponding with the fact that most humans have ten fingers. For many purposes, this is inconvenient. For example, ten does not have many positive divisors (only four)—this has implications for the ease of performing arithmetic; a system based on the number twelve, which has six positive divisors, might be more convenient. Another example is in computing and digital electronics, where it is more convenient to work in a *binary* system, with just two digits, which represent 'off' and 'on' (or 'low voltage' and 'high voltage'), respectively; arithmetic can then be performed directly using sequences of *logic gates* in an electrical circuit.

It is therefore worthwhile to have some understanding of positional numeral systems based on numbers other than ten. The mathematical abstraction we make leads to the definition of *base-b expansion*.

> **Definition 0.6**
> Let $b > 1$. The **base-b expansion** of a natural number n is the[a] string $d_r d_{r-1} \ldots d_0$ such that
>
> - $n = d_r \cdot b^r + d_{r-1} \cdot b^{r-1} + \cdots + d_0 \cdot b^0$;
>
> - $0 \leqslant d_i < b$ for each i; and
>
> - If $n > 0$ then $d_r \neq 0$—the base-b expansion of zero is 0 in all bases b.
>
> Certain number bases have names; for instance, the base-2, 3, 8, 10 and 16 expansions are respectively called *binary*, *ternary*, *octal*, *decimal* and *hexadecimal*.
>
> ---
> [a]The use of the word 'the' is troublesome here, since it assumes that every natural number has only one base-b expansion. This fact actually requires proof—see Theorem 4.3.56.

Example 0.7

Consider the number 1023. Its decimal (base-10) expansion is 1023, since

$$1023 = 1 \cdot 10^3 + 0 \cdot 10^2 + 2 \cdot 10^1 + 3 \cdot 10^0$$

Its binary (base-2) expansion is 1111111111, since

$$1023 = 1 \cdot 2^9 + 1 \cdot 2^8 + 1 \cdot 2^7 + 1 \cdot 2^6 + 1 \cdot 2^5 + 1 \cdot 2^4 + 1 \cdot 2^3 + 1 \cdot 2^2 + 1 \cdot 2^1 + 1 \cdot 2^0$$

We can express numbers in base-36 by using the ten usual digits 0 through 9 and the twenty-six letters A through Z; for instance, A represents 10, M represents 22 and Z represents 35. The base-36 expansion of 1023 is SF, since

$$1023 = 28 \cdot 36^1 + 15 \cdot 36^0 = S \cdot 36^1 + F \cdot 36^0$$

$\triangleleft$

Exercise 0.8

Find the binary, ternary, octal, decimal, hexadecimal and base-36 expansions of the number 21127, using the letters A–F as additional digits for the hexadecimal expansion and the letters A–Z as additional digits for the base-36 expansion. $\triangleleft$

We sometimes wish to specify a natural number in terms of its base-b expansion; we have some notation for this.

Notation 0.9

Let $b > 1$. If the numbers $d_0, d_1, \ldots, d_r$ are base-b digits (in the sense of Definition 0.6), then we write

$$d_r d_{r-1} \ldots d_{0\,(b)} = d_r \cdot b^r + d_{r-1} \cdot b^{r-1} + \cdots + d_0 \cdot b^0$$

for the natural number whose base-b expansion is $d_r d_{r-1} \ldots d_0$. If there is no subscript (b) and a base is not specified explicitly, the expansion will be assumed to be in base-10.

Example 0.10

Using our new notation, we have

$$1023 = 1111111111_{(2)} = 1101220_{(3)} = 1777_{(8)} = 1023_{(10)} = 3FF_{(16)} = SF_{(36)}$$

$$\triangleleft$$

Integers ($\mathbb{Z}$)

The *integers* can be used for measuring the difference between two instances of counting. For example, suppose I have five apples and five bananas. Another person, also holding apples and bananas, wishes to trade. After our exchange, I have seven apples and only one banana. Thus I have two more apples and four fewer bananas.

Since an increment in quantity can be represented by moving to the right on the number line by the unit length, a *decrement* in quantity can therefore be represented by moving to the *left* by the unit length. Doing so gives rise to the integers.

Definition 0.11

The **integers** are represented by the points on the number line which can be obtained by starting at 0 and moving in either direction by the unit length any number of times:

We write $\mathbb{Z}$ (LATEX code: `\mathbb{Z}`) for the set of all integers; thus, the notation '$n \in \mathbb{Z}$' means that n is an integer.

The integers have such a fascinating structure that a whole chapter of this book is devoted to them—see Chapter 4. This is to do with the fact that, although you can add, subtract and multiply two integers and obtain another integer, the same is not true of division. This 'bad behaviour' of division is what makes the integers interesting. We will now see some basic results about division.

Division of integers

The motivation we will soon give for the definition of the rational numbers (Definition 0.25) is that the result of dividing one integer by another integer is

not necessarily another integer. However, the result is *sometimes* another integer; for example, I can divide six apples between three people, and each person will receive an integral number of apples. This makes division interesting: how can we measure the failure of one integer's divisibility by another? How can we deduce when one integer is divisible by another? What is the structure of the set of integers when viewed through the lens of division? This motivates Definition 0.12.

Definition 0.12 (to be repeated in Definition 4.1.4)
Let $a, b \in \mathbb{Z}$. We say b **divides** a if $a = qb$ for some integer q. Other ways of saying that b divides a are: b is a *divisor* of a, b is a *factor* of a, or a is a *multiple* of b.

Example 0.13
The integer 12 is divisible by 1, 2, 3, 4, 6 and 12, since

$$12 = 12 \cdot 1 = 6 \cdot 2 = 4 \cdot 3 = 3 \cdot 4 = 2 \cdot 6 = 1 \cdot 12$$

It is also divisible by the negatives of all of those numbers; for example, 12 is divisible by -3 since $12 = (-4) \cdot (-3)$. ◁

Exercise 0.14
Prove that 1 divides every integer, and that every integer divides 0. ◁

Using Definition 0.12, we can prove some general basic facts about divisibility.

Proposition 0.15
Let $a, b, c \in \mathbb{Z}$. If c divides b and b divides a, then c divides a.

Proof
Suppose that c divides b and b divides a. By Definition 0.12, it follows that

$$b = qc \quad \text{and} \quad a = rb$$

for some integers q and r. Using the first equation, we may substitute qc for b in the second equation, to obtain

$$a = r(qc)$$

But $r(qc) = (rq)c$, and rq is an integer, so it follows from Definition 0.12 that c divides a. □

Exercise 0.16
Let $a, b, d \in \mathbb{Z}$. Suppose that d divides a and d divides b. Given integers u and v, prove that d divides $au + bv$. ◁

Some familiar concepts, such as evenness and oddness, can be characterised in terms of divisibility.

> **Definition 0.17**
> An integer n is **even** if it is divisible by 2; otherwise, n is **odd**.

It is not just interesting to know when one integer *does* divide another; however, proving that one integer *doesn't* divide another is much harder. Indeed, to prove that an integer b does not divide an integer a, we must prove that $a \neq qb$ for *any* integer q at all. We will look at methods for doing this in Chapter 1; these methods use the following extremely important result, which will underlie all of Chapter 4.

> **Theorem 0.18** (Division theorem, to be repeated in Theorem 4.1.1)
> Let $a, b \in \mathbb{Z}$ with $b \neq 0$. There is exactly one way to write
> $$a = qb + r$$
> such that q and r are integers, and $0 \leqslant r < b$ (if $b > 0$) or $0 \leqslant r < -b$ (if $b < 0$).

The number q in Theorem 0.18 is called the **quotient** of a when divided by b, and the number r is called the **remainder**.

Example 0.19
The number 12 leaves a remainder of 2 when divided by 5, since $12 = 2 \cdot 5 + 2$. ◁

Here's a slightly more involved example.

Proposition 0.20
Suppose an integer a leaves a remainder of r when divided by an integer b, and that $r > 0$. Then $-a$ leaves a remainder of $b - r$ when divided by b.

Proof
Suppose a leaves a remainder of r when divided by b. Then
$$a = qb + r$$
for some integer q. A bit of algebra yields
$$-a = -qb - r = -qb - r + (b - b) = -(q+1)b + (b-r)$$
Since $0 < r < b$, we have $0 < b - r < b$. Hence $-(q+1)$ is the quotient of $-a$ when divided by b, and $b - r$ is the remainder. □

Exercise 0.21
Prove that if an integer a leaves a remainder of r when divided by an integer b, then a leaves a remainder of r when divided by $-b$. ◁

We will finish this part on division of integers by connecting it with the material on number bases—we can use the division theorem (Theorem 0.18) to find the base-b expansion of a given natural number. It is based on the following observation: the natural number n whose base-b expansion is $d_r d_{r-1} \cdots d_1 d_0$ is equal to

$$d_0 + b(d_1 + b(d_2 + \cdots + b(d_{r-1} + bd_r) \cdots))$$

Moreover, $0 \leqslant d_i < b$ for all i. In particular n leaves a remainder of d_0 when divided by b. Hence

$$\frac{n - d_0}{b} = d_1 + d_2 b + \cdots + d_r b^{r-1}$$

The base-b expansion of $\frac{n-d_0}{b}$ is therefore

$$d_r d_{r-1} \cdots d_1$$

In other words, the remainder of n when divided by b is the last base-b digit of n, and then subtracting this number from n and dividing the result by b truncates the final digit. Repeating this process gives us d_1, and then d_2, and so on, until we end up with 0.

This suggests the following algorithm for computing the base-b expansion of a number n:

- **Step 1.** Let d_0 be the remainder when n is divided by b, and let $n_0 = \frac{n-d_0}{b}$ be the quotient. Fix $i = 0$.

- **Step 2.** Suppose n_i and d_i have been defined. If $n_i = 0$, then proceed to Step 3. Otherwise, define d_{i+1} to be the remainder when n_i is divided by b, and define $n_{i+1} = \frac{n_i - d_{i+1}}{b}$. Increment i, and repeat Step 2.

- **Step 3.** The base-b expansion of n, is

$$d_i d_{i-1} \cdots d_0$$

Example 0.22
We compute the base-17 expansion of 15213, using the letters A–G to represent the numbers 10 through 16.

- $15213 = 894 \cdot 17 + 15$, so $d_0 = 15 = $ F and $n_0 = 894$.

- $894 = 52 \cdot 17 + 10$, so $d_1 = 10 = $ A and $n_1 = 52$.

- $52 = 3 \cdot 17 + 1$, so $d_2 = 1$ and $n_2 = 3$.

- $3 = 0 \cdot 17 + 3$, so $d_3 = 3$ and $n_3 = 0$.

- The base-17 expansion of 15213 is therefore 31AF.

A quick verification gives

$$31\text{AF}_{(17)} = 3 \cdot 17^3 + 1 \cdot 17^2 + 10 \cdot 17 + 15 = 15213$$

as desired.   ◁

Exercise 0.23
Find the base-17 expansion of $408\,735\,787$ and the base-36 expansion of $1\,442\,151\,747$.   ◁

Exercise 0.24
The video-sharing website *YouTube* assigns to each video a unique identifier, which is a string of 11 characters from the set

$$\{\text{A}, \text{B}, \ldots, \text{Z}, \text{a}, \text{b}, \ldots, \text{z}, 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, \text{-}, \_\}$$

This string is actually a natural number expressed in base-64, where the characters in the above set represent the numbers 0 through 63, in the same order—thus C represents 2, c represents 28, 3 represents 55, and _ represents 63. According to this schema, find the natural number whose base-64 expansion is dQw4w9WgXcQ, and find the base-64 expansion of the natural number $7\,159\,047\,702\,620\,056\,984$.   ◁

Rational numbers ($\mathbb{Q}$)

Bored of eating apples and bananas, I buy a pizza which is divided into eight slices. A friend and I decide to share the pizza. I don't have much of an appetite, so I eat three slices and my friend eats five. Unfortunately, we cannot represent the proportion of the pizza each of us has eaten using natural numbers or integers. However, we're not far off: we can count the number of equal parts the pizza was split into, and of those parts, we can count how many we had. On the number line, this could be represented by splitting the unit line segment from 0 to 1 into eight equal pieces, and proceeding from there. This kind of procedure gives rise to the *rational numbers*.

Definition 0.25

The **rational numbers** are represented by the points at the number line which can be obtained by dividing any of the unit line segments between integers into an equal number of parts.

$$-5 \quad -4 \quad -3 \quad -2 \quad -1 \quad 0 \quad 1 \quad 2 \quad 3 \quad 4 \quad 5$$

The rational numbers are those of the form $\frac{a}{b}$, where $a, b \in \mathbb{Z}$ and $b \neq 0$. We write $\mathbb{Q}$ (LaTeX code: `\mathbb{Q}`) for the set of all rational numbers; thus, the notation '$q \in \mathbb{Q}$' means that q is a rational number.

The rational numbers are a very important example of a type of algebraic structure known as a *field*—they are particularly central to algebraic number theory and algebraic geometry.

Real numbers ($\mathbb{R}$)

Quantity and change can be measured in the abstract using *real numbers*.

Definition 0.26

The **real numbers** are the points on the number line. We write $\mathbb{R}$ (LaTeX code: `\mathbb{R}`) for the set of all real numbers; thus, the notation '$a \in \mathbb{R}$' means that a is a real number.

The real numbers are central to real analysis, a branch of mathematics introduced in Chapter 7. They turn the rationals into a *continuum* by 'filling in the gaps'—specifically, they have the property of *completeness*, meaning that if a quantity can be approximated with arbitrary precision by real numbers, then that quantity is itself a real number.

We can define the basic arithmetic operations (addition, subtraction, multiplication and division) on the real numbers, and a notion of ordering of the real numbers, in terms of the infinite number line.

- **Ordering.** A real number a is less than a real number b, written $a < b$, if a lies to the left of b on the number line. The usual conventions for the symbols $\leqslant$ (LaTeX code: `\le`), $>$ and $\geqslant$ (LaTeX code: `\ge`) apply, for instance '$a \leqslant b$' means that either $a < b$ or $a = b$.

- **Addition.** Suppose we want to add a real number a to a real number b. To do this, we *translate* a by b units to the right—if $b < 0$ then this amounts to translating a by an equivalent number of units to the left. Concretely, take two copies of the number line, one above the other, with the same choice of unit length; move the 0 of the lower number line beneath the point a of the upper number line. Then $a + b$ is the point on the upper number line lying above the point b of the lower number line.

  Here is an illustration of the fact that $(-3) + 5 = 2$:

- **Multiplication.** This one is fun. Suppose we want to multiply a real number a by a real number b. To do this, we *scale* the number line, and perhaps *reflect* it. Concretely, take two copies of the number line, one above the other; align the 0 points on both number lines, and stretch the lower number line evenly until the point 1 on the lower number line is below the point a on the upper number line—note that if $a < 0$ then the number line must be reflected in order for this to happen. Then $a \cdot b$ is the point on the upper number line lying above b on the lower number line.

  Here is an illustration of the fact that $5 \cdot 4 = 20$.

  and here is an illustration of the fact that $(-5) \cdot 4 = -20$:

Exercise 0.27

Interpret the operations of subtraction and division as geometric transformations of the real number line. ◁

We will take for granted the arithmetic properties of the real numbers in this chapter, waiting until Section 7.1 to sink our teeth into the details. For example,

we will take for granted the basic properties of rational numbers, for instance

$$\frac{a}{b} + \frac{c}{d} = \frac{ad + bc}{bd} \quad \text{and} \quad \frac{a}{b} \cdot \frac{c}{d} = \frac{ac}{bd}$$

Rational and irrational numbers

Before we can talk about irrational numbers, we should say what they are.

Definition 0.28
An **irrational number** is a real number that is not rational.

Unlike $\mathbb{N}, \mathbb{Z}, \mathbb{Q}, \mathbb{R}, \mathbb{C}$, there is no standard single letter expressing the irrational numbers. However, by the end of Section 2.1, we will be able to write the set of irrational numbers as $\mathbb{R} \setminus \mathbb{Q}$.

Note in particular that 'irrational' does not simply mean 'not rational'—that would imply that all complex numbers which are not real are irrational—rather, the term 'irrational' means 'real and not rational'.

Proving that a real number is *irrational* is not particularly easy. We will get our foot in the door by allowing ourselves to assume the following result, which is restated and proved in Proposition 3.1.48.

Proposition 0.29
The real number $\sqrt{2}$ is irrational. □

We can use the fact that $\sqrt{2}$ is irrational to prove some facts about the relationship between rational numbers and irrational numbers.

Proposition 0.30
Let a and b be irrational numbers. It is possible that ab be rational.

Proof
Let $a = b = \sqrt{2}$. Then a and b are irrational, and $ab = 2 = \frac{2}{1}$, which is rational. □

Exercise 0.31
Let r be a rational number and let a be an irrational number. Prove that it is possible that ra be rational, and it is possible that ra be irrational. ◁

Complex numbers ($\mathbb{C}$)

We have seen that multiplication by real numbers corresponds with scaling and reflection of the number line—scaling alone when the multiplicand is positive, and scaling with reflection when it is negative. We could alternatively interpret this reflection as a *rotation* by half a turn, since the effect on the number line is the same. You might then wonder what happens if we rotate by arbitrary angles, rather than only half turns.

What we end up with is a *plane* of numbers, not merely a line—see Figure 1. Moreover, it happens that the rules that we expect arithmetic operations to satisfy still hold—addition corresponds with translation, and multiplication corresponds with scaling and rotation. This resulting number set is that of the *complex numbers*.

> **Definition 0.32**
> The **complex numbers** are those obtained by the non-negative real numbers upon rotation by any angle about the point 0. We write $\mathbb{C}$ (LaTeX code: `\mathbb{C}`) for the set of all complex numbers; thus, the notation '$z \in \mathbb{C}$' means that z is a complex number.

There is a particularly important complex number, i, which is the point in the complex plane exactly one unit above 0—this is illustrated in Figure 1. Multiplication by i has the effect of rotating the plane by a quarter turn anticlockwise. In particular, we have $i^2 = i \cdot i = -1$; the complex numbers have the astonishing property that square roots of *all* complex numbers exist (including all the real numbers).

In fact, every complex number can be written in the form $a + bi$, where $a, b \in \mathbb{R}$; this number corresponds with the point on the complex plane obtained by moving a units to the right and b units up, reversing directions as usual if a or b is negative. Arithmetic on the complex numbers works just as with the real numbers; in particular, using the fact that $i^2 = -1$, we obtain

$$(a+bi)+(c+di) = (a+c)+(b+d)i \quad \text{and} \quad (a+bi)\cdot(c+di) = (ac-bd)+(ad+b$$

We will discuss complex numbers further in the portion of this chapter on polynomials below.

Figure 1: Illustration of the complex plane, with some points labelled.

Polynomials

The integers, rational numbers, real numbers and complex numbers are all examples of *rings*, which means that they come equipped with nicely behaving notions of addition, subtraction and multiplication.

Definition 0.33

Let A be one $\mathbb{Z}$, $\mathbb{Q}$, $\mathbb{R}$ or $\mathbb{C}$. A (**univariate**) **polynomial over** A in the **indeterminate** x is an expression of the form

$$a_0 + a_1 x + \cdots + a_n x^n$$

where $n \in \mathbb{N}$ and each $a_k \in A$. The numbers a_k are called the **coefficients** of the polynomial. If not all coefficients are zero, the largest value of k for which $a_k \neq 0$ is called the **degree** of the polynomial. By convention, the degree of the polynomial 0 is $-\infty$.

Polynomials of degree 1, 2, 3, 4 and 5 are respectively called *linear*, *quadratic*, *cubic*, *quartic* and *quintic* polynomials.

Example 0.34

The following expressions are all polynomials:

$$3 \qquad 2x - 1 \qquad (3+i)x^2 - x$$

Their degrees are 0, 1 and 2, respectively. The first two are polynomials over $\mathbb{Z}$, and the third is a polynomial over $\mathbb{C}$.   ◁

Exercise 0.35

Write down a polynomial of degree 4 over $\mathbb{R}$ which is not a polynomial over $\mathbb{Q}$.   ◁

Notation 0.36

Instead of writing out the coefficients of a polynomial each time, we may write something like $p(x)$ or $q(x)$. The '(x)' indicates that x is the indeterminate of the polynomial. If α is a number[a] and $p(x)$ is a polynomial in indeterminate x, we write $p(\alpha)$ for the result of **substituting** α for x in the expression $p(x)$.

Note that, if A is any of the sets $\mathbb{N}$, $\mathbb{Z}$, $\mathbb{Q}$, $\mathbb{R}$ or $\mathbb{C}$, and $p(x)$ is a polynomial over A, then $p(\alpha) \in A$ for all $\alpha \in A$.

---

[a]When dealing with polynomials, we will typically reserve the letter x for the indeterminate variable, and use the Greek letters α, β, γ (LaTeX code: `\alpha, \beta, \gamma`) for numbers to be substituted into a polynomial.

Example 0.37

Let $p(x) = x^3 - 3x^2 + 3x - 1$. Then $p(x)$ is a polynomial over $\mathbb{Z}$ with indeterminate x. For any integer α, the value $p(\alpha)$ will also be an integer. For example

$$p(0) = 0^3 - 3 \cdot 0^2 + 3 \cdot 0 - 1 = -1 \quad \text{and} \quad p(3) = 3^3 - 3 \cdot 3^2 + 3 \cdot 3 - 1 = 8$$

◁

Definition 0.38

Let $p(x)$ be a polynomial. A **root** of $p(x)$ is a complex number α such that $p(\alpha) = 0$.

The *quadratic formula* (Theorem 1.1.31) tells us that the roots of the polynomial $x^2 + ax + b$, where $a, b \in \mathbb{C}$, are precisely the complex numbers

$$\frac{-a + \sqrt{a^2 - 4b}}{2} \quad \text{and} \quad \frac{-a - \sqrt{a^2 - 4b}}{2}$$

Note our avoidance of the symbol '$\pm$', which is commonly found in discussions of quadratic polynomials. The symbol '$\pm$' is dangerous because it may suppress the word 'and' or the word 'or', depending on context—this kind of ambiguity is not something that we will want to deal with when discussing the logical structure of a proposition in Chapter 1!

Example 0.39

Let $p(x) = x^2 - 2x + 5$. The quadratic formula tells us that the roots of p are

$$\frac{2 + \sqrt{4 - 4 \cdot 5}}{2} = 1 + \sqrt{-4} = 1 + 2i \quad \text{and} \quad \frac{2 - \sqrt{4 - 4 \cdot 5}}{2} = 1 - \sqrt{-4} = 1 - 2i$$

The numbers $1 + 2i$ and $1 - 2i$ are related in that their real parts are equal and their imaginary parts differ only by a sign. Exercise 0.40 generalises this observation.

◁

Exercise 0.40

Let $\alpha = a + bi$ be a complex number, where $a, b \in \mathbb{R}$. Prove that α is the root of a quadratic polynomial over $\mathbb{R}$, and find the other root of this polynomial. ◁

The following exercise proves the well-known result which classifies the number of real roots of a polynomial over $\mathbb{R}$ in terms of its coefficients.

Exercise 0.41

Let $a, b \in \mathbb{C}$ and let $p(x) = x^2 + ax + b$. The value $\Delta = a^2 - 4b$ is called the

discriminant of p. Prove that p has two roots if $\Delta \neq 0$ and one root if $\Delta = 0$. Moreover, if $a, b \in \mathbb{R}$, prove that p has no real roots if $\Delta < 0$, one real root if $\Delta = 0$, and two real roots if $\Delta > 0$. $\triangleleft$

Example 0.42

Consider the polynomial $x^2 - 2x + 5$. Its discriminant is equal to $(-2)^2 - 4 \cdot 5 = -16$, which is negative. Exercise 0.41 tells us that it has two roots, neither of which are real. This was verified by Example 0.39, where we found that the roots of $x^2 - 2x + 5$ are $1 + 2i$ and $1 - 2i$.

Now consider the polynomial $x^2 - 2x - 3$. Its discriminant is equal to $(-2)^2 - 4 \cdot (-3) = 16$, which is positive. Exercise 0.41 tells us that it has two roots, both of which are real; and indeed

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

so the roots of $x^2 - 2x - 3$ are -1 and 3. $\triangleleft$

Chapter 1

Logical structure

The goal of this chapter is to develop a methodical way of breaking up a proposition into smaller components and seeing how these components fit together—this is called the *logical structure* of a proposition. The logical structure of a proposition is very informative: it tells us what we need to do in order to prove it, what we need to write in order to communicate our proof, and how to explore the consequences of the proposition after it has been proved.

Sections 1.1 and 1.2 are dedicated to developing a system of *symbolic logic* for reasoning about propositions. We will be able to represent a proposition using a string of variables and symbols, and this expression will guide how we can prove the proposition and explore its consequences. In Section 1.3 we will develop techniques for manipulating these logical expressions algebraically—this, in turn, will yield new proof techniques (some have fancy names like 'proof by contraposition', but some do not).

Exploring how the logical structure of a proposition informs the structure and wording of its proof is the content of Appendix A.2.

Propositional logic

Every mathematical proof is written in the context of certain *assumptions* being made, and certain *goals* to be achieved.

- **Assumptions** are the propositions which are known to be true, or which we are assuming to be true for the purposes of proving something. They include theorems that have already been proved, prior knowledge which is assumed of the reader, and assumptions which are explicitly made using words like 'suppose' or 'assume'.

- **Goals** are the propositions we are trying to prove in order to complete the proof of a result, or perhaps just a step in the proof.

With every phrase we write, our assumptions and goals change. This is perhaps best illustrated by example. In Example 1.1.1 below, we will examine the proof of Proposition 0.15 in detail, so that we can see how the words we wrote affected the assumptions and goals at each stage in the proof. We will indicate our assumptions and goals at any given stage using tables—the assumptions listed will only be those assumptions which are made explicitly; prior knowledge and previously proved theorems will be left implicit.

Example 1.1.1
The statement of Proposition 0.15 was as follows:

Let $a, b, c \in \mathbb{Z}$. If c divides b and b divides a, then c divides a.

The set-up of the proposition instantly gives us our initial assumptions and goals:

Assumptions	Goals
$a, b, c \in \mathbb{Z}$	If c divides b and b divides a, then c divides a

We will now proceed through the proof, line by line, to see what effect the words we wrote had on the assumptions and goals.

Since our goal was an expression of the form 'if... then...', it made sense to start by assuming the 'if' statement, and using that assumption to prove the 'then' statement. As such, the first thing we wrote in our proof was:

Suppose that c divides b and b divides a.

Our updated assumptions and goals are reflected in the following table.

Assumptions	Goals
$a, b, c \in \mathbb{R}$	c divides a
c divides b	
b divides a	

Our next step in the proof was to unpack the definitions of 'c divides b' and 'b divides a', giving us more to work with.

Suppose that c divides b and b divides a. By Definition 0.12, it follows that

$$b = qc \quad \text{and} \quad a = rb$$

for some integers q and r.

This introduces two new variables q, r and allows us to replace the assumptions 'c divides b' and 'b divides a' with their definitions.

Assumptions	Goals
$a, b, c, q, r \in \mathbb{Z}$	c divides a
$b = qc$	
$a = rb$	

At this point we have pretty much exhausted all of the assumptions we can make, and so our attention turns towards the goal—that is, we must prove that c divides a. At this point, it helps to 'work backwards' by unpacking the goal: what does it mean for c to divide a? Well, by Definition 0.12, we need to prove that a is equal to some integer multiplied by c—this will be reflected in the following table of assumptions and goals.

Since we are now trying to express a in terms of c, it makes sense to use the equations we have relating a with b, and b with c, to relate a with c.

Suppose that c divides b and b divides a. By Definition 0.12, it follows that

$$b = qc \quad \text{and} \quad a = rb$$

for some integers q and r. Using the first equation, we may substitute qc for b in the second equation, to obtain

$$a = r(qc)$$

We are now very close, as indicated in the following table.

Assumptions	Goals
$a,b,c,q,r \in \mathbb{Z}$	$a = [\text{some integer}] \cdot c$
$b = qc$	
$a = rb$	
$a = r(qc)$	

Our final step was to observe that the goal has at last been achieved:

Suppose that c divides b and b divides a. By Definition 0.12, it follows that

$$b = qc \quad \text{and} \quad a = rb$$

for some integers q and r. Using the first equation, we may substitute qc for b in the second equation, to obtain

$$a = r(qc)$$

But $r(qc) = (rq)c$, and rq is an integer,

Assumptions	Goals
$a,b,c,q,r \in \mathbb{Z}$	
$b = qc$	
$a = rb$	
$a = r(qc)$	
$a = (rq)c$	
$rq \in \mathbb{Z}$	

Now that there is nothing left to prove, it is helpful to reiterate that point so that the reader has some closure on the matter.

Suppose that c divides b and b divides a. By Definition 0.12, it follows that

$$b = qc \quad \text{and} \quad a = rb$$

for some integers q and r. Using the first equation, we may substitute qc for b in the second equation, to obtain

$$a = r(qc)$$

But $r(qc) = (rq)c$, and rq is an integer, so it follows from Definition 0.12 that c divides a.

◁

Symbolic logic

Consider again the proposition that we proved in Proposition 0.15 (for given integers a, b, c):

> If c divides b and b divides a, then c divides a.

The three statements 'c divides b', 'b divides a' and 'c divides a' are all propositions in their own right, despite the fact that they all appear inside a more complex proposition. We can examine the logical structure of the proposition by replacing these simpler propositions with symbols, called *propositional variables*. Writing P to represent 'c divides b', Q to represent 'b divides a' and R to represent 'c divides a', we obtain:

> If P and Q, then R.

Breaking down the proposition in this way makes it clear that a feasible *assume P and Q*, and then *derive R* from these assumptions—this is exactly what we did in the proof, which we examined in great detail in Example 1.1.1. But importantly, it suggests that the same proof strategy might work for other propositions which are also of the form 'if P and Q, then R', such as the following proposition (for a given integer n):

> If $n > 2$ and n is prime, then n is odd.

Observe that the simpler propositions are joined together to form a more complex proposition using language, namely the word 'and' and the construction 'if... then...'—we will represent these constructions symbolically using *logical operators*, which will be introduced in Definition 1.1.3.

Zooming in even more closely, we can use Definition 0.12 to observe that 'c divides b' really means '$b = qc$ for some $q \in \mathbb{Z}$'. The expression 'for some $q \in \mathbb{Z}$' introduces a new variable q, which we must deal with appropriately in our proof. Words which we attach to variables in our proofs—such as 'any', 'exists', 'all', 'some', 'unique' and 'only'—will be represented symbolically using *quantifiers*, which we will study in Section 1.2.

By breaking down a complex proposition into simpler statements which are connected together using logical operators and quantifiers, we can more precisely identify what assumptions we can make at any given stage in a proof of the proposition, and what steps are needed in order to finish the proof.

Propositional formulae

We begin our development of symbolic logic with some definitions to fix our terminology.

Definition 1.1.2

A **propositional variable** is a symbol that represents a proposition. Propositional variables may be assigned **truth values** ('true' or 'false').

We will typically use the lower-case letters p, q, r and s as our propositional variables. It is also common to use upper-case letters $P, Q, R, \ldots$, like we did earlier, or even Greek letters $\varphi, \chi, \psi, \ldots$.

We will be able to form more complex expressions representing propositions by connecting together simpler ones using *logical operators* such as $\wedge$ (which represents 'and'), $\vee$ (which represents 'or'), $\Rightarrow$ (which represents 'if... then...') and $\neg$ (which represents 'not').

The definition of the notions of *logical operator* and *propositional formula* given below is a little bit difficult to digest—it is very abstract and even appears circular—so it is best understood by considering examples of propositional formulae and instances of logical operators. Fortunately we will see plenty of these, since they are the central objects of study for the rest of this section.

Definition 1.1.3

A **propositional formula** is an expression that is either a propositional variable, or is built up from simpler propositional formulae ('subformulae') using a **logical operator**. In the latter case, the truth value of the propositional formula is determined by the truth values of the subformulae according to the rules of the logical operator.

On first sight, Definition 1.1.3 seems circular—it defines the term 'propositional formula' in terms of propositional formulae! But in fact it is not circular; it is an example of a *recursive* definition (we avoid circularity with the word 'simpler'). To illustrate, consider the following example of a propositional formula:

$$(p \wedge q) \Rightarrow r$$

This expression represents a proposition of the form 'if p and q, then r', where p, q, r are themselves propositions. It is built from the subformulae $p \wedge q$ and r using the logical operator $\Rightarrow$, and $p \wedge q$ is itself built up from the subformulae p and q using the logical operator $\wedge$.

The truth value of $(p \wedge q) \Rightarrow r$ is then determined by the truth values of the constituent propositional variables (p, q and r) according to the rules for the logical operators $\wedge$ and $\Rightarrow$.

If this all seems a bit abstract, that is because it *is* abstract, and you are forgiven if it makes no sense to you yet. From this point onwards, we will only study particular instances of logical operators, which will make it all much easier to understand.

Conjunction ('and', $\wedge$)

Conjunction is the logical operator which makes precise what we mean when we say 'and'.

Definition 1.1.4

The **conjunction** operator is the logical operator $\wedge$ (LATEX code: `\wedge`), defined according to the following rules:

- ($\wedge$I) If p is true and q is true, then $p \wedge q$ is true;

- ($\wedge$E$_1$) If $p \wedge q$ is true, then p is true;

- ($\wedge$E$_2$) If $p \wedge q$ is true, then q is true.

The expression $p \wedge q$ represents 'p and q'.

It is not always obvious when conjunction is being used; sometimes it sneaks in without the word 'and' ever being mentioned! Be on the look-out for occasions like this, such as in the following exercise.

Example 1.1.5
We can express the proposition '7 is a prime factor of 28' in the form $p \wedge q$, by letting p represent the proposition '7 is prime' and letting q represent the proposition '7 divides 28'. ◁

Exercise 1.1.6
Express the proposition 'Clive is a mathematician who lives in Pittsburgh' in the form $p \wedge q$, for propositions p and q. ◁

The rules in Definition 1.1.4 are examples of *rules of inference*—they tell us how to deduce (or 'infer') the truth of one propositional formula from the truth of other propositional formulae. In particular, rules of inference never directly tell

us when a proposition is *false*—in order to prove something is false, we will prove its *negation* is true (see Definition 1.1.37).

Rules of inference tell us how to use the logical structure of propositions in proofs:

- The rule (∧I) is an *introduction rule*, meaning that it tells us how to *prove a goal* of the form $p \wedge q$—indeed, if we want to prove that $p \wedge q$ is true, (∧I) tells us that it suffices to prove that p is true and prove that q is true.

- The rules (∧E₁) and (∧E₂) are *elimination rules*, meaning that they tell us how to *use an assumption* of the form $p \wedge q$—indeed, if we are assuming that $p \wedge q$ is true, we are then free to use the fact that p is true and the fact that q is true.

Each logical operator will come equipped with some introduction and/or elimination rules, which tell us how to prove goals or use assumptions which include the logical operator in question. It is in this way that the logical structure of a proposition informs *proof strategies*, like the following:

Strategy 1.1.7 (Proving conjunctions)
A proof of the proposition $p \wedge q$ can be obtained by tying together two proofs, one being a proof that p is true and one being a proof that q is true.   ◁

Example 1.1.8
Suppose we are required to prove that 7 is a prime factor of 28. In Example 1.1.5 we expressed '7 is a prime factor of 28' as the conjunction of the propositions '7 is prime' and '7 divides 28', and so Strategy 1.1.7 breaks down the proof into two steps: first prove that 7 is prime, and then prove that 7 divides 28.   ◁

Much like Strategy 1.1.7 was informed by the introduction rule for ∧, the elimination rules inform how we may make use of an assumption involving a conjunction.

Strategy 1.1.9 (Assuming conjunctions)
If an assumption in a proof has the form $p \wedge q$, then we may assume p and assume q in the proof.   ◁

Example 1.1.10
Suppose that, somewhere in the process of proving a proposition, we arrive at the fact that 7 is a prime factor of 28. Strategy 1.1.9 then allows us to use the separate facts that 7 is prime and that 7 divides 28.   ◁

Strategies 1.1.7 and 1.1.9 seem almost *obvious*. To an extent they are obvious, and that is why we are stating them first. But the real reason we are going through the process of precisely defining logical operators, their introduction and elimination rules, and the corresponding proof strategies, is that when you are in the middle of the proof of a complicated result, it is all too easy to lose track of what you have already proved and what remains to be proved. Keeping track of the assumptions and goals in a proof, and understanding what must be done in order to complete the proof, is a difficult task.

To avoid drawing this process out too long, we need a compact way of expressing rules of inference that allows us to simply read off corresponding proof strategies. We *could* use tables of assumptions and goals like in Example 1.1.1, but this quickly becomes clunky—indeed, even the very simple conjunction introduction rule ($\wedge$I) doesn't look very nice in this format:

Assumptions	Goals		Assumptions	Goals
$\vdots$	$p \wedge q$	$\rightsquigarrow$	$\vdots$	p
$\vdots$			$\vdots$	q

Instead, we will represent rules of inference in the style of *natural deduction*. In this style, we write the *premises* $p_1, p_2, \ldots, p_k$ of a rule above a line, with a single *conclusion* q below the line, representing the assertion that the truth of a proposition q follows from the truth of (all of) the premises $p_1, p_2, \ldots, p_k$.

$$\frac{p_1 \qquad p_2 \qquad \cdots \qquad p_k}{q}$$

For instance, the introduction and elimination rules for conjunction can be expressed concisely follows:

$$\frac{p \qquad q}{p \wedge q}\ (\wedge\text{I}) \qquad \frac{p \wedge q}{p}\ (\wedge\text{E}_1) \qquad \frac{p \wedge q}{q}\ (\wedge\text{E}_2)$$

In addition to its clean and compact nature, this way of writing rules of inference is useful because we can combine them into *proof trees* in order to see how to prove more complicated propositions. For example, consider the following proof tree, which combines two instances of the conjunction introduction rule.

$$\frac{\dfrac{p \qquad q}{p \wedge q} \qquad r}{(p \wedge q) \wedge r}$$

From this proof tree, we obtain a strategy for proving a proposition of the form $(p \wedge q) \wedge r$. Namely, first prove p and prove q, to conclude $p \wedge q$; and then prove r, to conclude $(p \wedge q) \wedge r$. This illustrates that the logical structure of a proposition informs how we may structure a proof of the proposition.

Exercise 1.1.11

Write a proof tree whose conclusion is the propositional formula $(p \wedge q) \wedge (r \wedge s)$, where p, q, r, s are propositional variables. Express '2 is an even prime number and 3 is an odd prime number' in the form $(p \wedge q) \wedge (r \wedge s)$, for appropriate propositions p, q, r and s, and describe how your proof tree suggests what a proof might look like. ◁

Disjunction ('or', $\vee$)

> **Definition 1.1.12**
> The **disjunction** operator is the logical operator $\vee$ (LaTeX code: `\vee`), defined according to the following rules:
>
> - ($\vee I_1$) If p is true, then $p \vee q$ is true;
>
> - ($\vee I_2$) If q is true, then $p \vee q$ is true;
>
> - ($\vee E$) If $p \vee q$ is true, and if r can be derived from p and from q, then r is true.
>
> The expression $p \vee q$ represents 'p or q'.

The introduction and elimination rules for disjunction are represented diagramatically as follows.

$$\frac{p}{p \vee q} \, (\vee I_1) \qquad \frac{q}{p \vee q} \, (\vee I_2) \qquad \frac{p \vee q \quad \overset{[p]}{\underset{r}{\wr}} \quad \overset{[q]}{\underset{r}{\wr}}}{r} \, (\vee E)$$

We will discuss what the notation $[p] \rightsquigarrow r$ and $[q] \rightsquigarrow r$ means momentarily. First, we zoom in on how the disjunction introduction rules inform proofs of propositions of the form 'p or q'.

> **Strategy 1.1.13** (Proving disjunctions)
> In order to prove a proposition of the form $p \vee q$, it suffices to prove just one of p or q. ◁

Example 1.1.14

Suppose we want prove that 8192 is not divisible by 3. We know by the division theorem (Theorem 0.18) that an integer is not divisible by 3 if and only if it leaves a remainder of 1 or 2 when divided by 3, and so it suffices to prove the following:

$$\begin{array}{ccc} \text{8192 leaves a remainder of 1} & \vee & \text{8192 leaves a remainder of 2} \\ \text{when divided by 3} & & \text{when divided by 3} \end{array}$$

A quick computation reveals that $8192 = 2730 \times 3 + 2$, so that 8192 leaves a remainder of 2 when divided by 3. By Strategy 1.1.13, the proof is now complete, since the full disjunction follows by ($\vee I_2$). ◁

Example 1.1.15

Let p, q, r, s be propositional variables. The propositional formula $(p \vee q) \wedge (r \vee s)$ represents 'p or q, and r or s'. What follows are two examples of truth trees for this propositional formula.

$$\cfrac{\cfrac{p}{p \vee q}\ (\vee I_1) \quad \cfrac{r}{r \vee s}\ (\vee I_1)}{(p \vee q) \wedge (r \vee s)}\ (\wedge I) \qquad \cfrac{\cfrac{q}{p \vee q}\ (\vee I_2) \quad \cfrac{s}{r \vee s}\ (\vee I_2)}{(p \vee q) \wedge (r \vee s)}\ (\wedge I)$$

The proof tree on the left suggests the following proof strategy for $(p \vee q) \wedge (r \vee s)$. First prove p, and deduce $p \vee q$; then prove r, and deduce $r \vee s$; and finally deduce $(p \vee q) \wedge (r \vee s)$. The proof tree on the right suggests a different strategy, where $p \vee q$ is deduced by proving q instead of p, and $r \vee s$ is deduced by proving s instead of r.

Selecting which (if any) of these to use in a proof might depend on what we are trying to prove. For example, for a fixed natural number n, let p represent 'n is even', let q represent 'n is odd', let r represent '$n \geqslant 2$' and let s represent 'n is a perfect square'. Proving $(p \vee q) \wedge (r \vee s)$ when $n = 2$ would be most easily done using the left-hand proof tree above, since p and r are evidently true when $n = 2$. However, the second proof tree would be more appropriate for proving $(p \vee q) \wedge (r \vee s)$ when $n = 1$. ◁

Aside

If you haven't already mixed up $\wedge$ and $\vee$, you probably will soon, so here's a way of remembering which is which:

fish n chips

If you forget whether it's $\wedge$ or $\vee$ that means 'and', just write it in place of the 'n' in 'fish n chips':

$$\text{fish} \wedge \text{chips} \qquad \text{fish} \vee \text{chips}$$

Clearly the first looks more correct, so $\wedge$ means 'and'. If you don't eat fish (or chips), then worry not, as this mnemonic can be modified to accommodate a wide variety of dietary restrictions; for instance 'mac n cheese' or 'quinoa n kale'.   ◁

Recall the diagrammatic statement of the disjunction elimination rule:

$$
\frac{p \vee q \qquad \begin{array}{c}[p] \\ \wr \\ r\end{array} \qquad \begin{array}{c}[q] \\ \wr \\ r\end{array}}{r} \;(\vee\text{E})
$$

The curious notation $[p] \rightsquigarrow r$ indicates that p is a *temporary assumption*. In the part of the proof corresponding to $[p] \rightsquigarrow r$, we would assume that p is true and derive r from that assumption, and remove the assumption that p is true from that point onwards. Likewise for $[q] \rightsquigarrow r$.

The proof strategy obtained from the disjunction elimination rule is called *proof by cases*.

Strategy 1.1.16 (Assuming disjunctions—proof by cases)
If an assumption in a proof has the form $p \vee q$, then we may derive a proposition r by splitting into two cases: first, derive r from the temporary assumption that p is true, and then derive r from the assumption that q is true.   ◁

The following example illustrates how Strategies 1.1.13 and 1.1.16 can be used together in a proof.

Example 1.1.17
Let n be a positive proper factor of 4, and suppose we want to prove that n is either even or a perfect square.

- Our assumption that n is a positive proper factor of 4 can be expressed as the disjunction $n = 1 \vee n = 2$.

- Our goal is to prove the disjunction 'n is even $\vee$ n is a perfect square'.

According to Strategy 1.1.9, we split into two cases, one in which $n = 1$ and one in which $n = 2$. In each case, we must derive 'n is even $\vee$ n is a perfect square', for which it suffices by Strategy 1.1.13 to derive either that n is even or that n is a perfect square. Thus a proof might look something like this:

Since n is a positive proper factor of 4, either $n = 1$ or $n = 2$.

- **Case 1.** Suppose $n = 1$. Then since $1^2 = 1$ we have $n = 1^2$, so that n is a perfect square.
- **Case 2.** Suppose $n = 2$. Then since $2 = 2 \times 1$, we have that n is even.

Hence n is either even or a perfect square. ◁

Notice that in both Case 1 and Case 2, we did not explicitly mention that we had proved that 'n is even $\vee$ n is a perfect square', leaving that deducgion to the reader—we only mentioned it after the proofs in each case were complete. ◁

The proof of Proposition 1.1.18 below splits into *three* cases, rather than just two.

Proposition 1.1.18

Let $n \in \mathbb{Z}$. Then n^2 leaves a remainder of 0 or 1 when divided by 3.

Proof

Let $n \in \mathbb{Z}$. By the division theorem (Theorem 0.18), one of the following must be true for some $k \in \mathbb{Z}$:

$$n = 3k \quad \text{or} \quad n = 3k + 1 \quad \text{or} \quad n = 3k + 2$$

- Suppose $n = 3k$. Then

$$n^2 = (3k)^2 = 9k^2 = 3 \cdot (3k^2)$$

So n^2 leaves a remainder of 0 when divided by 3.

- Suppose $n = 3k + 1$. Then

$$n^2 = (3k + 1)^2 = 9k^2 + 6k + 1 = 3(3k^2 + 2k) + 1$$

So n^2 leaves a remainder of 1 when divided by 3.

- Suppose $n = 3k + 2$. Then

$$n^2 = (3k + 2)^2 = 9k^2 + 12k + 4 = 3(3k^2 + 4k + 1) + 1$$

So n^2 leaves a remainder of 1 when divided by 3.

In all cases, n^2 leaves a remainder of 0 or 1 when divided by 3. □

Note that in the proof of Proposition 1.1.18, unlike in Example 1.1.17, we did not explictly use the word 'case', even though we were using proof by cases. Whether or not to make your proof strategies explicit is up to you—discussion of this kind of matter can be found in Appendix A.2.

When completing the following exercises, try to keep track of exactly where you use the introduction and elimination rules that we have seen so far.

Exercise 1.1.19

Let n be an integer. Prove that n^2 leaves a remainder of 0, 1 or 4 when divided by 5. ◁

Exercise 1.1.20

Let $a, b \in \mathbb{R}$ and suppose $a^2 - 4b \neq 0$. Let α and β be the (distinct) roots of the polynomial $x^2 + ax + b$. Prove that there is a real number c such that either $\alpha - \beta = c$ or $\alpha - \beta = ci$. ◁

Implication ('if…then…', $\Rightarrow$)

Definition 1.1.21

The **implication** operator is the logical operator $\Rightarrow$ (LaTeX code: `\Rightarrow`), defined according to the following rules:

- ($\Rightarrow$I) If q can be derived from the assumption that p is true, then $p \Rightarrow q$ is true;

- ($\Rightarrow$E) If $p \Rightarrow q$ is true and p is true, then q is true.

The expression $p \Rightarrow q$ represents 'if p, then q'.

$$\frac{\begin{array}{c}[p]\\ \vdots \\ q\end{array}}{p \Rightarrow q}\ (\Rightarrow\text{I}) \qquad \frac{p \Rightarrow q \qquad p}{q}\ (\Rightarrow\text{E})$$

Strategy 1.1.22 (Proving implications)

In order to prove a proposition of the form $p \Rightarrow q$, it suffices to assume that p is true, and then derive q from that assumption. ◁

The following proposition illustrates how Strategy 1.1.22 can be used in a proof.

Proposition 1.1.23

Let x and y be real numbers. If x and $x + y$ are rational, then y is rational.

Proof

Suppose x and $x + y$ are rational. Then there exist integers a, b, c, d with $b, d \neq 0$ such that

$$x = \frac{a}{b} \quad \text{and} \quad x + y = \frac{c}{d}$$

It then follows that

$$y = (x+y) - x = \frac{c}{d} - \frac{a}{b} = \frac{bc - ad}{bd}$$

Since $bc - ad$ and bd are integers, and $bd \neq 0$, it follows that y is rational. $\quad\square$

The key phrase in the above proof was 'Suppose x and $x+y$ are rational.' This introduced the assumptions $x \in \mathbb{Q}$ and $x+y \in \mathbb{Q}$, and reduced our goal to that of deriving a proof that y is rational—this was the content of the rest of the proof.

Exercise 1.1.24
Let $p(x)$ be a polynomial over $\mathbb{C}$. Prove that if α is a root of $p(x)$, and $a \in \mathbb{C}$, then α is a root of $(x - a)p(x)$. ◁

The elimination rule for implication ($\Rightarrow$E) is more commonly known by the Latin name *modus ponens*.

Strategy 1.1.25 (Assuming implications—modus ponens)
If an assumption in a proof has the form $p \Rightarrow q$, and p is also assumed to be true, then we may also assume that q is true. ◁

Strategy 1.1.16 is frequently used to reduce a more complicated goal to a simpler one. Indeed, if we know that $p \Rightarrow q$ is true, and if p is easy to verify, then it allows us to prove q by proving p instead.

Example 1.1.26
Let $f(x) = x^2 + ax + b$ be a polynomial with $a, b \in \mathbb{R}$, and let $\Delta = a^2 - 4b$ be its discriminant. Part of Exercise 0.41 was to prove that:

   (i) If $\Delta > 0$, then f has two real roots;

   (ii) If $\Delta = 0$, then f has one real root;

  (iii) If $\Delta < 0$, then f has no real roots.

Given the polynomial $f(x) = x^2 - 68 + 1156$, it would be a pain to go through the process of solving the equation $f(x) = 0$ in order to determine how many real roots f has. However, each of the propositions (i), (ii) and (iii) take the form $p \Rightarrow q$, so Strategy 1.1.25 reduces the problem of finding how many real roots f has to that of evaluating Δ and comparing it with 0. And indeed, $(-68)^2 - 4 \times 1156 = 0$, so the implication (ii) together with ($\Rightarrow$E) tell us that f has one real root. ◁

A common task faced by mathematicians is to prove that two conditions are equivalent. For example, given a polynomial $f(x) = x^2 + ax + b$ with $a, b \in \mathbb{R}$, we

know that *if $a^2 - 4b > 0$ then f has two real roots*, but is it also true that *if f has two real roots then $a^2 - 4b > 0$?* (The answer is 'yes'.) The relationship between these two implications is that each is the *converse* of the other.

Definition 1.1.27
The **converse** of a proposition of the form $p \Rightarrow q$ is the proposition $q \Rightarrow p$.

A quick remark on terminology is pertinent. The following table summarises some common ways of referring to the propositions '$p \Rightarrow q$' and '$q \Rightarrow p$'.

$p \Rightarrow q$	$q \Rightarrow p$
if p, then q	if q, then p
p only if q	p if q
p is sufficient for q	p is necessary for q

We so often encounter the problem of proving both an implication and its converse that we introduce a new logical operator that represents the conjunction of both.

Definition 1.1.28
The **biconditional** operator is the logical operator $\Leftrightarrow$ (LATEX code: `\Leftrightarrow`), defined by declaring $p \Leftrightarrow q$ to mean $(p \Rightarrow q) \wedge (q \Rightarrow p)$. The expression $p \Leftrightarrow q$ represents 'p if and only if q'.

Many examples of biconditional statements come from solving equations; indeed, to say that the values $\alpha_1, \ldots, \alpha_n$ are the solutions to a particular equation is precisely to say that

$$x \text{ is a solution} \quad \Leftrightarrow \quad x = \alpha_1 \text{ or } x = \alpha_2 \text{ or } \cdots \text{ or } x = \alpha_n$$

Example 1.1.29
We find all real solutions x to the equation

$$\sqrt{x-3} + \sqrt{x+4} = 7$$

Let's rearrange the equation to find out what the possible solutions may be.

$$\sqrt{x-3}+\sqrt{x+4}=7$$

$\Rightarrow (x-3)+2\sqrt{(x-3)(x+4)}+(x+4)=49 \qquad$ squaring

$\Rightarrow 2\sqrt{(x-3)(x+4)}=48-2x \qquad$ rearranging

$\Rightarrow 4(x-3)(x+4)=(48-2x)^2 \qquad$ squaring

$\Rightarrow 4x^2+4x-48=2304-192x+4x^2 \qquad$ expanding

$\Rightarrow 196x=2352 \qquad$ rearranging

$\Rightarrow x=12 \qquad$ dividing by 196

You might be inclined to stop here. Unfortunately, all we have proved is that, given a real number x, *if* x solves the equation $\sqrt{x-3}+\sqrt{x+4}=7$, *then* $x=12$. This narrows down the set of possible solutions to just one candidate—but we still need to check the converse, namely that *if* $x=12$, *then* x is a solution to the equation.

As such, to finish off the proof, note that

$$\sqrt{12-3}+\sqrt{12+4}=\sqrt{9}+\sqrt{16}=3+4=7$$

and so the value $x=12$ is indeed a solution to the equation. $\qquad\qquad \triangleleft$

The last step in Example 1.1.29 may have seemed a little bit silly; but Example 1.1.30 demonstrates that proving the converse when solving equations truly is necessary.

Example 1.1.30
We find all real solutions x to the equation

$$x+\sqrt{x}=0$$

We proceed as before, rearranging the equation to find all possible solutions.

$$x+\sqrt{x}=0$$

$\Rightarrow x=-\sqrt{x} \qquad$ rearranging

$\Rightarrow x^2=x \qquad$ squaring

$\Rightarrow x(x-1)=0 \qquad$ rearranging

$\Rightarrow x=0 \text{ or } x=1$

Now certainly 0 is a solution to the equation, since

$$0+\sqrt{0}=0+0=0$$

However, 1 is *not* a solution, since

$$1 + \sqrt{1} = 1 + 1 = 2$$

Hence it is actually the case that, given a real number x, we have

$$x + \sqrt{x} = 0 \quad \Leftrightarrow \quad x = 0$$

Checking the converse here was vital to our success in solving the equation!   ◁

A slightly more involved example of a biconditional statement arising from the solution to an equation—in fact, a class of equations—is the proof of the quadratic formula.

Theorem 1.1.31 (Quadratic formula)
Let $a, b \in \mathbb{C}$. A complex number α is a root of the polynomial $x^2 + ax + b$ if and only if

$$\alpha = \frac{-a + \sqrt{a^2 - 4b}}{2} \quad \text{or} \quad \alpha = \frac{-a - \sqrt{a^2 - 4b}}{2}$$

Proof
First we prove that *if* α is a root, *then* α is one of the values given in the statement of the proposition. So suppose α be a root of the polynomial $x^2 + ax + b$. Then

$$\alpha^2 + a\alpha + b = 0$$

The algebraic technique of 'completing the square' tells us that

$$\alpha^2 + a\alpha = \left(\alpha + \frac{a}{2}\right)^2 - \frac{a^2}{4}$$

and hence

$$\left(\alpha + \frac{a}{2}\right)^2 - \frac{a^2}{4} + b = 0$$

Rearranging yields

$$\left(\alpha + \frac{a}{2}\right)^2 = \frac{a^2}{4} - b = \frac{a^2 - 4b}{4}$$

Taking square roots gives

$$\alpha + \frac{a}{2} = \frac{\sqrt{a^2 - 4b}}{2} \quad \text{or} \quad \alpha + \frac{a}{2} = \frac{-\sqrt{a^2 - 4b}}{2}$$

and, finally, subtracting $\frac{a}{2}$ from both sides gives the desired result.

The proof of the converse is Exercise 1.1.32.   □

Exercise 1.1.32

Complete the proof of the quadratic formula. That is, for fixed $a, b \in \mathbb{C}$, prove that if

$$\alpha = \frac{-a + \sqrt{a^2 - 4b}}{2} \quad \text{or} \quad \alpha = \frac{-a - \sqrt{a^2 - 4b}}{2}$$

then α is a root of the polynomial $x^2 + ax + b$. ◁

Another class of examples of biconditional propositions arise in finding necessary and sufficient criteria for an integer n to be divisible by some number—for example, that an integer is divisible by 10 if and only if its base-10 expansion ends with the digit 0.

Example 1.1.33

Let $n \in \mathbb{N}$. We will prove that n is divisible by 8 if and only if the number formed of the last three digits of the base-10 expansion of n is divisible by 8.

First, we will do some 'scratch work'. Let $d_r d_{r-1} \ldots d_1 d_0$ be the base-10 expansion of n. Then

$$n = d_r \cdot 10^r + d_{r-1} \cdot 10^{r-1} + \cdots + d_1 \cdot 10 + d_0$$

Define

$$n' = d_2 d_1 d_0 \quad \text{and} \quad n'' = n - n' = d_r d_{r-1} \ldots d_4 d_3 000$$

Now $n - n' = 1000 \cdot d_r d_{r-1} \ldots d_4 d_3$ and $1000 = 8 \cdot 125$, so it follows that 8 divides n''.

Our goal is now to prove that 8 divides n if and only if 8 divides n'.

- ($\Rightarrow$) Suppose 8 divides n. Since 8 divides n'', it follows from Exercise 0.16 that 8 divides $an + bn''$ for all $a, b \in \mathbb{Z}$. But

$$n' = n - (n - n') = n - n'' = 1 \cdot n + (-1) \cdot n''$$

  so indeed 8 divides n', as required.

- ($\Leftarrow$) Suppose 8 divides n'. Since 8 divides n'', it follows from Exercise 0.16 that 8 divides $an' + bn''$ for all $a, b \in \mathbb{Z}$. But

$$n = n' + (n - n') = n' + n'' = 1 \cdot n' + 1 \cdot n''$$

  so indeed 8 divides n, as required.

◁

Exercise 1.1.34

Prove that a natural number n is divisible by 3 if and only if the sum of its base-10 digits is divisible by 3. ◁

Negation ('not', ¬)

So far we only officially know how to prove that true propositions are *true*. The negation operator makes precise what we mean by 'not', which allows us to prove that false propositions are *false*.

Definition 1.1.35
A **contradiction** is a proposition that is known or assumed to be false. We will use the symbol ⊥ (LATEX code: `\bot`) to represent an arbitrary contradiction.

Example 1.1.36
Some examples of contradictions include the assertion that $0 = 1$, or that $\sqrt{2}$ is rational, or that the equation $x^2 = -1$ has a solution $x \in \mathbb{R}$. ◁

Definition 1.1.37
The **negation** operator is the logical operator ¬ (LATEX code: `\neg`), defined according to the following rules:

- (¬I) If a contradiction can be derived from the assumption that p is true, then $\neg p$ is true;

- (¬E) If $\neg p$ and p are both true, then a contradiction may be derived.

The expression $\neg p$ represents 'not p' (or 'p is false').

$$\frac{\overset{\displaystyle [p]}{\overset{\displaystyle \wr}{\bot}}}{\neg p}\,(\neg \text{I}) \qquad\qquad \frac{\neg p \quad p}{\bot}\,(\neg \text{E})$$

Aside
The rules (¬I) and (¬E) closely resemble (⇒I) and (⇒E)—indeed, we could simply define $\neg p$ to mean '$p \Rightarrow \bot$', where $\bot$ represents an arbitrary contradiction, but it will be easier later on to have a primitive notion of negation. ◁

The introduction rule for negation (¬I) gives rise to a proof strategy called *proof by contradiction*, which turns out to be extremely useful.

Strategy 1.1.38 (Proving negations—proof by contradiction)
In order to prove a proposition p is false (that is, that $\neg p$ is true), it suffices to assume that p is true and derive a contradiction. ◁

The following proposition has a classic proof by contradiction.

Proposition 1.1.39

Let r be a rational number and let a be an irrational number. Then $r + a$ is irrational.

Proof

By Definition 0.28, we need to prove that $r + a$ is real and not rational. It is certainly real, since r and a are real, so it remains to prove that $r + a$ is not rational.

Suppose $r + a$ is rational. Since r is rational, it follows from Proposition 1.1.23 that a is rational, since
$$a = (r + a) - r$$
This contradicts the assumption that a is irrational. It follows that $r + a$ is not rational, and is therefore irrational. □

Now you can try proving some elementary facts by contradiction.

Exercise 1.1.40

Let $x \in \mathbb{R}$. Prove by contradiction that if x is irrational then $-x$ and $\frac{1}{x}$ are irrational. ◁

Exercise 1.1.41

Prove by contradiction that there is no least positive real number. That is, prove that there is not a positive real number a such that $a \leqslant b$ for all positive real numbers b. ◁

A proof need not be a 'proof by contradiction' in its entirety—indeed, it may be that only a small portion of the proof uses contradiction. This is exhibited in the proof of the following proposition.

Proposition 1.1.42

Let a be an integer. Then a is odd if and only if $a = 2b + 1$ for some integer b.

Proof

Suppose a is odd. By the division theorem (Theorem 0.18), either $a = 2b$ or $a = 2b + 1$, for some $b \in \mathbb{Z}$. If $a = 2b$, then 2 divides a, contradicting the assumption that a is odd; so it must be the case that $a = 2b + 1$.

Conversely, suppose $a = 2b + 1$. Then a leaves a remainder of 1 when divided by 2. However, by the division theorem, the even numbers are precisely those that leave a remainder of 0 when divided by 2. It follows that a is not even, so is odd. □

The elimination rule for the negation operator ($\neg$E) simply says that a proposition can't be true and false at the same time.

Strategy 1.1.43 (Assuming negations)
If an assumption in a proof has the form $\neg p$, then any derivation of p leads to a contradiction.   ◁

The main use of Strategy 1.1.43 is for obtaining the contradiction in a proof by contradiction—in fact, we have already used it in our examples of proof by contradiction! As such, we will not dwell on it further.

Logical axioms

We wrap up this section by introducing a couple of additional logical rules (*axioms*) that we will use in our proofs.

The first is the so-called *law of excluded middle*, which appears so obvious that it is not even worth stating (let alone naming)—what it says is that every proposition is either true or false. But beware, as looks can be deceiving; the law of excluded middle is a non-constructive axiom, meaning that it should not be accepted in settings it is important to keep track of how a proposition is proved—simply knowing that a proposition is either true or false tells us nothing about how it might be proved or refuted. In most mathematical contexts, though, it is accepted without a second's thought.

Axiom 1.1.44 (Law of excluded middle)
Let p be a propositional formula. Then $p \vee (\neg p)$ is true.

The law of excluded middle can be represented diagramatically as follows; there are no premises above the line, since we are simply asserting that it is true.

$$\frac{}{p \vee (\neg p)} \text{ LEM}$$

Strategy 1.1.45 (Using the law of excluded middle)
In order to prove a proposition q is true, it suffices to split into cases based on whether some other proposition p is true or false, and prove that q is true in each case.   ◁

The proof of Proposition 1.1.46 below makes use of the law of excluded middle—note that we defined 'odd' to mean 'not even' (Definition 0.17).

Proposition 1.1.46

Let $a, b \in \mathbb{Z}$. If ab is even, then either a is even or b is even (or both).

Proof

Suppose $a, b \in \mathbb{Z}$ with ab even.

- Suppose a is even—then we're done.

- Suppose a is odd. If b is also odd, then by Proposition 1.1.42 can write

$$a = 2k + 1 \quad \text{and} \quad b = 2\ell + 1$$

for some integers k, ℓ. This implies that

$$ab = (2k+1)(2\ell+1) = 4k\ell + 2k + 2\ell + 1 = 2(\underbrace{2k\ell + k + \ell}_{\in \mathbb{Z}}) + 1$$

so that ab is odd. This contradicts the assumption that ab is even, and so b must in fact be even.

In both cases, either a or b is even. □

Exercise 1.1.47

Reflect on the proof of Proposition 1.1.46. Where in the proof did we use the law of excluded middle? Where in the proof did we use proof by contradiction? What was the contradiction in this case? Prove Proposition 1.1.46 twice more, once using contradiction and not using the law of excluded middle, and once using the law of excluded middle and not using contradiction. ◁

Exercise 1.1.48

Let a and b be irrational numbers. By considering the number $\sqrt{2}^{\sqrt{2}}$, prove that it is possible that a^b be rational. ◁

Another logical rule that we will use is the *principle of explosion*, which is also known by its Latin name, *ex falso sequitur quodlibet*, which approximately translates to '*from falsity follows whatever you like*'.

Axiom 1.1.49 (Principle of explosion)

If a contradiction is assumed, any consequence may be derived.

$$\frac{\perp}{p} \text{ Expl}$$

The principle of explosion is a bit confusing on first sight. To shed a tiny bit of intuition on it, think of it as saying that both true and false propositions are consequences of a contradictory assumption. For instance, suppose that $-1 = 1$. From this we can obtain consequences that are false, such as $0 = 2$ by adding 1 to both sides of the equation, and consequences that are true, such as $1 = 1$ by squaring both sides of the equation.

We will rarely use the principle of explosion directly in our mathematical proofs, but we will use it in Section 1.3 for proving logical formulae are equivalent.

Section 1.2
Variables and quantifiers

Free and bound variables

Everything we did in Section 1.1 concerned *propositions* and the logical rules concerning their proofs. Unfortunately if all we have to work with is propositions then our ability to do mathematical reasoning will be halted pretty quickly. For example, consider the following statement:

$$x \text{ is divisible by } 7$$

This statement seems like the kind of thing we should probably be able to work with if we're doing mathematics. It makes sense if x is a integer, such as 28 or 41; but it doesn't make sense at all if x is a parrot called Alex.[a] In any case, even when it does make sense, its truth value depends on x; indeed, '28 is divisible by 7' is a true proposition, but '41 is divisible by 7' is a false proposition.

This means that the statement 'x is divisible by 7' isn't a proposition—*quel horreur*! But it *almost* is a proposition: if we know that x refers somehow to an integer, then it becomes a proposition as soon as a particular numerical value of x is specified. The symbol x is called a *free variable*.

Definition 1.2.1
Let x be a variable that is understood to refer to an element of a set X. In a statement involving x, we say x is **free** if it makes sense to substitute particular elements of X in the statement; otherwise, we say x is **bound**.

To represent statements that have free variables in them abstractly, we generalise the notion of a propositional variable (Definition 1.1.2) to that of a *predicate*.

---
[a] Alex the parrot is the only non-human animal to have ever been observed to ask an existential question; he died in September 2007 so we may never know if he was divisible by 7, but it is unlikely. According to *Time*, his last words were 'you be good, see you tomorrow, I love you'. The reader is advised to finish crying before they continue reading about variables and quantifiers.

Definition 1.2.2

A **predicate** is a symbol p together with a specified list of free variables $x_1, x_2, \ldots, x_n$ (where $n \in \mathbb{N}$) and, for each free variable x_i, a specification of a set X_i called the **domain of discourse** (or **range**) of x_i. We will typically write $p(x_1, x_2, \ldots, x_n)$ in order to make the variables explicit.

The statements represented by predicates are those that become propositions when specific values are substituted for their free variables from their respective domains of discourse. For example, 'x is divisible by 7' is not a proposition, but it becomes a proposition when specific integers (such as 28 or 41) are substituted for x.

This is a lot to take in, so let's look at some examples.

Example 1.2.3

(i) We can represent the statement 'x is divisible by 7' discussed above by a predicate $p(x)$ whose only free variable x has $\mathbb{Z}$ as its domain of discourse. Then $p(28)$ is the true proposition '28 is divisible by 7' and $p(41)$ is the false proposition '41 is divisible by 7'.

(ii) A predicate with no free variables is precisely a propositional variable. This means that the notion of a predicate generalises that of a propositional variable.

(iii) The expression '$2^n - 1$ is prime' can be represented by a predicate $p(n)$ with one free variable n, whose domain of discourse is the set $\mathbb{N}$ of natural numbers. Then $p(3)$ is the true proposition '$2^3 - 1$ is prime' and $p(4)$ is the false proposition '$2^4 - 1$ is prime'.

(iv) The expression '$x - y$ is rational' can be represented by a predicate $q(x, y)$ with free variables x and y, whose domain of discourse is the set $\mathbb{R}$ of real numbers.

(v) The expression 'there exist integers a and b such that $x = a^2 + b^2$' has free variable x and bound variables a, b. It can be represented by a predicate $r(x)$ with one free variable x, whose domain of discourse is $\mathbb{Z}$.

(vi) The expression 'every even natural number $n \geqslant 2$ is divisible by k' has free variable k and bound variable n. It can be represented by a predicate $s(k)$ with one free variable k, whose domain of discourse is $\mathbb{N}$.

$\triangleleft$

Quantifiers

Look again at the statements in parts (v) and (vi) of Example 1.2.3. Both contained bound variables, which were so because we used words like 'there exists' and 'every'—had we not used these words, those variables would be free, as in '$x = a^2 + b^2$' and 'n is divisible by k'.

Expressions that refer to *how many* elements of a set make a statement true, such as 'there exists' and 'every', turn free variables into bound variables. We represent such expressions using symbols called *quantifiers*, which are the central objects of study of this section.

The two main quantifiers used throughout mathematics are the *universal* quantifier $\forall$ and the *existential* quantifier $\exists$. We will define these quantifiers formally later in this section, but for now, the following informal definitions suffice:

- The expression '$\forall x \in X, \ldots$' denotes 'for all $x \in X, \ldots$' and will be defined formally in Definition 1.2.9;

- The expression '$\exists x \in X, \ldots$' denotes 'there exists $x \in X$ such that $\ldots$' and will be defined formally in Definition 1.2.17.

Note that we always place the quantifier *before* the statement, so even though we might write or say things like '$n = 2k$ for some integer k' or '$x^2 \geqslant 0$ for all $x \in \mathbb{R}$', we would express these statements symbolically as '$\exists k \in \mathbb{Z}, n = 2k$' and '$\forall x \in \mathbb{R}, x^2 \geqslant 0$', respectively.

We will define a third quantifier $\exists!$ in terms of $\forall$ and $\exists$ to say that there is *exactly one* element of a set making a statement true. There are plenty of other quantifiers out there, but they tend to be specific to particular fields—examples include 'almost everywhere' in measure theory, 'almost surely' in probability theory, 'for all but finitely many' in set theory and related disciplines, and 'for fresh' in the theory of nominal sets.

Using predicates, logical formulae and quantifiers, we are able to build up more complicated expressions, called *logical formulae*. Logical formulae generalise propositional formulae (Definition 1.1.3) in by allowing (free and bound) variables and quantification to occur.

> **Definition 1.2.4**
> A **logical formula** is an expression that is built from predicates using logical operators and quantifiers; it may have both free and bound variables. The truth value of a logical formula depends on its free variables according to the rules for logical operators and quantifiers.

Translating between plain English statements and purely symbolic logical formulae is an important skill to obtain:

- The plain English statements are easier to understand and are the kinds of things you would speak aloud or write down when discussing the mathematical ideas involved.

- The symbolic logical formulae are what provide the precision needed to guide a proof of the statement being discussed—we will see strategies for proving statements involving quantifiers soon.

The following examples and exercise concern translating between plain English statements and purely symbolic logical formulae.

Example 1.2.5

Recall that an integer n is even if and only if it is divisible by 2. According to Definition 0.12, that is to say that 'n is even' means '$n = 2k$ for some integer k'. Using quantifiers, we can express 'n is even' as '$\exists k \in \mathbb{Z}, n = 2k$'.

The (false) proposition 'every integer is even' can then be written symbolically as follows. First introduce a variable n to refer to an integer; to say 'every integer is even' is to say '$\forall n \in \mathbb{Z}, n$ is even', and so using the symbolic representation of 'n is even', we can express 'every integer is even' as $\forall n \in \mathbb{Z}, \exists k \in \mathbb{Z}, n = 2k$'.   ◁

Exercise 1.2.6

Find logical formulae that represent each of the following English statements.

(a) There is an integer that is divisible by every integer.

(b) There is no greatest odd integer.

(c) Between any two distinct rational numbers is a third distinct rational number.

(d) If an integer has a rational square root, then that root is an integer.

◁

Example 1.2.7

Consider the following logical formula.

$$\forall a \in \mathbb{R}, (a \geqslant 0 \Rightarrow \exists b \in \mathbb{R}, a = b^2)$$

If we translate this expression symbol-for-symbol, what it says is:

For every real number a, if a is non-negative,
then there exists a real number b such that $a = b^2$.

Read in this way, it is not a particularly enlightening statement. However, we can distill the robotic nature of the symbol-for-symbol reading by thinking more carefully about what the statement *really* means.

Indeed, to say '$a = b^2$ for some real number b' is exactly to say that a has a real square root—after all, what is a square root of a if not a real number whose square is equal to a? This translation eliminates explicit reference to the bound variable b, so that the statement now reads:

For every real number a, if a is non-negative, then a has a real square root.

We're getting closer. Next note that instead of the clunky expression 'for every real number a, if a is non-negative, then …', we could just say 'for every non-negative real number a, …'.

For every non-negative real number a, a has a real square root.

Finally, we can eliminate the bound variable a by simply saying:

Every non-negative real number has a real square root.

This is now a meaningful expression that is much easier to understand than the logical formula we started with. ◁

Exercise 1.2.8

Find statements in plain English, involving as few variables as possible, that are represented by each of the following logical formulae. (The domains of discourse of the free variables are indicated in each case.)

(a) $\exists q \in \mathbb{Z}, a = qb$ — free variables $a, b \in \mathbb{Z}$

(b) $\exists a \in \mathbb{Z}, \exists b \in \mathbb{Z}, (b \neq 0 \wedge bx = a)$ — free variable $x \in \mathbb{R}$

(c) $\forall d \in \mathbb{N}, [(\exists q \in \mathbb{Z}, n = qd) \Rightarrow (d = 1 \vee d = n)]$ — free variable $n \in \mathbb{N}$

(d) $\forall a \in \mathbb{R}, [a > 0 \Rightarrow \exists b \in \mathbb{R}, (b > 0 \wedge a < b)]$ — no free variables

◁

Now that we have a better understanding of how to translate between plain English statements and logical formulae, we are ready to give a precise mathematical

treatment of quantifiers. Just like with logical operators in Section 1.1, quantifiers will be defined according to *introduction rules*, which tell us how to prove a quantified formula, and *elimination rules*, which tell us how to use an assumption that involves a quantifier.

Universal quantification ('for all', $\forall$)

The universal quantifier makes precise what we mean when we say 'for all', or '$p(x)$ is always true no matter what value x takes'.

Definition 1.2.9

The **universal quantifier** is the quantifier $\forall$ (LATEX code: \forall); if $p(x)$ is a logical formula with free variable x with range X, then $\forall x \in X,\, p(x)$ is the logical formula defined according to the following rules:

- ($\forall$I) If $p(x)$ can be derived from the assumption that x is an arbitrary element of X, then $\forall x \in X,\, p(x)$;

- ($\forall$E) If $a \in X$ and $\forall x \in X,\, p(x)$ is true, then $p(a)$ is true.

The expression $\forall x \in X,\, p(x)$ represents 'for all $x \in X,\, p(x)$'.

$$\begin{array}{cc} [x \in X] & \\ \vdots & \\ \dfrac{p(x)}{\forall x \in X,\, p(x)} & \dfrac{\forall x \in X,\, p(x) \qquad a \in X}{p(a)} \end{array}$$

Strategy 1.2.10 (Proving universally quantified statements)

To prove a proposition of the form $\forall x \in X,\, p(x)$, it suffices to prove $p(x)$ for an arbitrary element $x \in X$—in other words, prove $p(x)$ whilst assuming nothing about the variable x other than that it is an element of X.   ◁

Useful phrases for introducing an arbitrary variable of a set X in a proof include 'fix $x \in X$' or 'let $x \in X$' or 'take $x \in X$'—more on this is discussed in Appendix A.2.

The proofs of the following propositions illustrate how a proof of a universally quantified statement might look.

Proposition 1.2.11

The square of every odd integer is odd.

Proof

Let n be an odd integer. Then $n = 2k+1$ for some $k \in \mathbb{Z}$ by the division theorem (Theorem 0.18), and so

$$n^2 = (2k+1)^2 = 4k^2 + 4k + 1 = 2(2k^2 + 2k) + 1$$

Since $2k^2 + 2k \in \mathbb{Z}$, we have that n^2 is odd, as required. $\qquad\square$

Note that in the proof of Proposition 1.2.11, we did not assume anything about n other than that it is an odd integer.

Proposition 1.2.12

The base-10 expansion of the square of every natural number ends in one of the digits 0, 1, 4, 5, 6 or 9.

Proof

Fix $n \in \mathbb{N}$, and let

$$n = d_r d_{r-1} \ldots d_0$$

be its base-10 expansion. Write

$$n = 10m + d_0$$

where $m \in \mathbb{N}$—that is, m is the natural number obtained by removing the final digit from n. Then

$$n^2 = 100m^2 + 20md_0 + d_0^2 = 10m(10m + 2d_0) + d_0^2$$

Hence the final digit of n^2 is equal to the final digit of d_0^2. But the possible values of d_0^2 are

$$0 \quad 1 \quad 4 \quad 9 \quad 16 \quad 25 \quad 36 \quad 49 \quad 64 \quad 81$$

all of which end in one of the digits 0, 1, 4, 5, 6 or 9. $\qquad\square$

Exercise 1.2.13

Prove that every integer is rational. $\qquad\triangleleft$

Exercise 1.2.14

Prove that every linear polynomial over $\mathbb{Q}$ has a rational root. $\qquad\triangleleft$

Exercise 1.2.15

Prove that, for all real numbers x and y, if x and y are irrational, then $x+y$ and $x-y$ are not both rational. $\qquad\triangleleft$

Before advancing too much further, beware of the following common error that arises when dealing with universal quantifiers.

Common error

Consider the following (non-)proof of the proposition $\forall n \in \mathbb{Z}, n^2 \geqslant 0$.

Let n be an arbitrary integer, say $n = 17$. Then $17^2 = 289 \geqslant 0$, so the statement is true.

The error made here is that the *writer* has picked an arbitrary value of n, not the *reader*. (In fact, the above argument actually proves $\exists n \in \mathbb{Z}, n^2 \geqslant 0$.)

The proof should make no assumptions about the value of n other than that it is an integer. Here is a correct proof:

Let n be an arbitrary integer. Either $n \geqslant 0$ or $n < 0$. If $n \geqslant 0$ then $n^2 \geqslant 0$, since the product of two nonnegative numbers is nonnegative; if $n < 0$ then $n^2 \geqslant 0$, since the product of two negative numbers is positive.

<div align="right">◁</div>

The strategy suggested by the elimination rule for the universal quantifier is one that we use almost without thinking about it.

> **Strategy 1.2.16** (Assuming universally quantified statements)
> If an assumption in a proof has the form $\forall x \in X, p(x)$, then we may assume that $p(a)$ is true whenever a is an element of X.   ◁

Existential quantification ('there exists', $\exists$)

> **Definition 1.2.17**
> The **existential quantifier** is the quantifier $\exists$ (LaTeX code: `\exists`)(LaTeX code: `\exists`)$\exists$; if $p(x)$ is a logical formula with free variable x with range X, then $\exists x \in X, p(x)$ is the logical formula defined according to the following rules:
>
> - ($\exists$I) If $a \in X$ and $p(a)$ is true, then $\exists x \in X, p(x)$;
>
> - ($\exists$E) If $\exists x \in X, p(x)$ is true, and q can be derived from the assumption that $p(a)$ is true for some fixed $a \in X$, then q is true.
>
> The expression $\exists x \in X, p(x)$ represents 'there exists $x \in X$ such that $p(x)$'.

$$\frac{a \in X \qquad p(a)}{\exists x \in X, p(x)} \; (\exists\mathrm{I}) \qquad \frac{\exists x \in X, p(x) \qquad \begin{array}{c}[a \in X], [p(a)] \\ \vdots \\ q\end{array}}{q} \; (\exists\mathrm{E})$$

Strategy 1.2.18 (Proving existentially quantified statements)
To prove a proposition of the form $\exists x \in X, p(x)$, it suffices to prove $p(a)$ for some specific element $a \in X$, which should be explicitly defined. ◁

Example 1.2.19
We prove that there is a natural number that is a perfect square and is one more than a perfect cube. That is, we prove

$$\exists n \in \mathbb{N}, ([\exists k \in \mathbb{Z}, n = k^2] \wedge [\exists \ell \in \mathbb{Z}, n = \ell^3 + 1])$$

So define $n = 9$. Then $n = 3^2$ and $n = 2^3 + 1$, so that n is a perfect square and is one more than a perfect cube, as required. ◁

The following proposition involves an existentially quantified statement—indeed, to say that a polynomial $f(x)$ has a real root is to say $\exists x \in \mathbb{R}, f(x) = 0$.

Proposition 1.2.20
Fix $a \in \mathbb{R}$. The cubic polynomial $x^3 + (1 - a^2)x - a$ has a real root.

Proof
Let $f(x) = x^3 + (1 - a^2)x - a$. Define $x = a$; then

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

Hence a is a root of $f(x)$. Since a is real, $f(x)$ has a real root. □

The following exercises require you to prove existentially quantified statements.

Exercise 1.2.21
Prove that there is a real number which is irrational but whose square is rational. ◁

Exercise 1.2.22
Prove that there is an integer which is divisible by zero. ◁

Example 1.2.23
Prove that, for all $x, y \in \mathbb{Q}$, if $x < y$ then there is some $z \in \mathbb{Q}$ with $x < z < y$. ◁

The elimination rule for the existential quantifier gives rise to the following proof strategy.

Strategy 1.2.24 (Assuming existentially quantified statements)
If an assumption in the proof has the form $\exists x \in X, p(x)$, then we may introduce a new variable $a \in X$ and assume that $p(a)$ is true. ◁

It ought to be said that when using existential elimination in a proof, the variable a used to denote a particular element of X for which $p(a)$ is true should not already be in use earlier in the proof.

Strategy 1.2.24 is very useful in proofs of divisibility, since the expression 'a divides b' is an existentially quantified statement—this was Exercise 1.2.8(a).

Proposition 1.2.25
Let $n \in \mathbb{Z}$. If n^3 is divisible by 3, then $(n+1)^3 - 1$ is divisible by 3.

Proof
Suppose n^3 is divisible by 3. Take $q \in \mathbb{Z}$ such that $n^3 = 3q$. Then

$$
\begin{aligned}
&(n+1)^3 - 1 \\
&= (n^3 + 3n^2 + 3n + 1) - 1 && \text{expanding} \\
&= n^3 + 3n^2 + 3n && \text{simplifying} \\
&= 3q + 3n^2 + 3n && \text{since } n^3 = 3q \\
&= 3(q + n^2 + n) && \text{factorising}
\end{aligned}
$$

Since $q + n^2 + n \in \mathbb{Z}$, we have proved that $(n+1)^3 - 1$ is divisible by 3, as required. $\qquad\square$

Uniqueness

The concept of uniqueness arises whenever we want to use the word 'the'. For example, in Definition 0.6 we defined the base-b expansion of a natural number n to be *the* string $d_r d_{r-1} \ldots d_1 d_0$ satisfying some properties. The issue with the word 'the' here is that we don't know ahead of time whether a natural number n may have base-b expansions other than $d_r d_{r-1} \ldots d_1 d_0$—this fact actually requires proof. To prove this fact, we would need to assume that $e_s e_{s-1} \ldots e_1 e_0$ were another base-b expansion of n, and prove that the strings $d_r d_{r-1} \ldots d_1 d_0$ and $e_s e_{s-1} \ldots e_1 e_0$ are the same—this is done in Theorem 4.3.56.

Uniqueness is typically coupled with *existence*, since we usually want to know if there is *exactly one* object satisfying a property. This motivates the definition of the *unique existential* quantifier, which encodes what we mean when we say 'there is exactly one $x \in X$ such that $p(x)$ is true'. The 'existence' part ensures that at least one $x \in X$ makes $p(x)$ true; the 'uniqueness' part ensures that x is the only element of X making $p(x)$ true.

Definition 1.2.26

The **unique existential quantifier** is the quantifier ∃! ((LATEX code: `\exists!`)) defined such that $\exists!x \in X,\, p(x)$ is shorthand for

$$\underbrace{(\exists x \in X,\, p(x))}_{\text{existence}} \wedge \underbrace{(\forall a \in X,\, \forall b \in X,\, [p(a) \wedge p(b) \Rightarrow a = b])}_{\text{uniqueness}}$$

Example 1.2.27

Every positive real number has a unique positive square root. We can write this symbolically as

$$\forall a \in \mathbb{R},\, (a > 0 \Rightarrow \exists!b \in \mathbb{R},\, (b > 0 \wedge b^2 = a))$$

Reading this from left to right, this says: for every real number a, if a is positive, then there exists a unique real number b, which is positive and whose square is a. ◁

Discussion 1.2.28

Explain why Definition 1.2.26 captures the notion of there being 'exactly one' element $x \in X$ making $p(x)$ true. Can you think of any other ways that $\exists!x \in X,\, p(x)$ could be defined? ◁

Strategy 1.2.29 (Proving unique-existentially quantified statements)

A proof of a statement of the form $\exists!x \in X,\, p(x)$, consists of two parts:

- **Existence** — prove that $\exists x \in X,\, p(x)$ is true (e.g. using Strategy 1.2.18);

- **Uniqueness** — let $a, b \in X$, assume that $p(a)$ and $p(b)$ are true, and derive $a = b$.

Alternatively, prove existence to obtain a fixed $a \in X$ such that $p(a)$ is true, and then prove $\forall x \in X,\, [p(x) \Rightarrow x = a]$. ◁

Example 1.2.30

We prove Example 1.2.27, namely that for each real $a > 0$ there is a unique $b > 0$ such that $b^2 = a$. So first fix $a > 0$.

- (**Existence**) The real number $\sqrt{a}$ is positive and satisfies $(\sqrt{a})^2 = a$ by definition. Its existence will be deferred to a later time, but an informal argument for its existence could be provided using 'number line' arguments as in Chapter 0.

- (**Uniqueness**) Let $y, z > 0$ be real numbers such that $y^2 = a$ and $z^2 = a$. Then $y^2 = z^2$. Rearranging and factorising yields

$$(y - z)(y + z) = 0$$

so either $y - z = 0$ or $y + z = 0$. If $y + z = 0$ then $z = -y$, and since $y > 0$, this means that $z < 0$. But this contradicts the assumption that $z > 0$. As such, it must be the case that $y - z = 0$, and hence $y = z$, as required.

<div align="right">◁</div>

Exercise 1.2.31

For each of the propositions, write it out as a logical formula involving the $\exists!$ quantifier and then prove it, using the structure of the logical formula as a guide.

(a) For each real number a, the equation $x^2 + 2ax + a^2 = 0$ has exactly one real solution x.

(b) There is a unique real number a for which the equation $x^2 + a^2 = 0$ has a real solution x.

(c) There is a unique natural number with exactly one positive divisor.

<div align="right">◁</div>

The unique existential quantifier will play a large role when we study functions in Section 2.2.

Quantifier alternation

Compare the following two statements:

(i) For every door, there is a key that can unlock it.

(ii) There is a key that can unlock every door.

Letting the variables x and y refer to doors and keys, respectively, and letting $p(x,y)$ be the statement 'door x can be unlocked by key y', we can formulate these statements as:

(i) $\forall x, \exists y, p(x,y)$

(ii) $\exists y, \forall x, p(x,y)$

This is a typical 'real-world' example of what is known as *quantifier alternation*—the two statements differ only by the order of the front-loaded quantifiers, and yet they say very different things. Statement (i) requires every door to be unlockable, but the keys might be different for different doors; statement (ii), however, implies the existence of some kind of 'master key' that can unlock all the doors.

Here's another example with a more mathematical nature:

Exercise 1.2.32
Let $p(x,y)$ be the statement '$x+y$ is even'.

- Prove that $\forall x \in \mathbb{Z}, \exists y \in \mathbb{Z}, p(x,y)$ is true.

- Prove that $\exists y \in \mathbb{Z}, \forall x \in \mathbb{Z}, p(x,y)$ is false.

◁

In both of the foregoing examples, you might have noticed that the '$\forall\exists$' statement says something *weaker* than the '$\exists\forall$' statement—in some sense, it is easier to make a $\forall\exists$ statement true than it is to make an $\exists\forall$ statement true.

This idea is formalised in Theorem 1.2.33 below, which despite its abstract nature, has an extremely simple proof.

Theorem 1.2.33
Let $p(x,y)$ be a logical formula with free variables $x \in X$ and $y \in Y$. Then

$$\exists y \in Y, \forall x \in X, p(x,y) \Rightarrow \forall x \in X, \exists y \in Y, p(x,y)$$

Proof
Suppose $\exists y \in Y, \forall x \in X, p(x,y)$ is true. We need to prove $\forall x \in X, \exists y \in Y, p(x,y)$, so fix $a \in X$—our goal is now to prove $\exists y \in Y, p(a,y)$.

Using our assumption $\exists y \in Y, \forall x \in X, p(x,y)$, we may choose $b \in Y$ such that $\forall x, p(x,b)$ is true. But then $p(a,b)$ is true, so we have proved $\exists y \in Y, p(a,y)$, as required. □

Statements of the form $\exists y \in Y, \forall x \in X, p(x,y)$ imply some kind of *uniformity*: a value of y making $\forall x \in X, p(x,y)$ true can be thought of as a 'one size fits all' solution to the problem of proving $p(x,y)$ for a given $x \in X$. Later in your studies, it is likely that you will encounter the word 'uniform' many times—it is precisely this notion of quantifier alternation that the word 'uniform' refers to.

Logical equivalence

We motivate the content of this section with an example.

Example 1.3.1
Consider the following two logical formulae, where P denotes the set of all prime numbers.

(1) $\forall n \in P, (n > 2 \Rightarrow [\exists k \in \mathbb{Z}, n = 2k + 1])$;

(2) $\neg \exists n \in P, (n > 2 \wedge [\exists k \in \mathbb{Z}, n = 2k])$.

The logical formula (1) translates to 'every prime number greater than two is odd', and the logical formula (2) translates to 'there does not exist an even prime number greater than two'. These statements are evidently *equivalent*—they mean the same thing—but they suggest different proof strategies:

(1) Fix a prime number n, assume that $n > 2$, and then prove that $n = 2k + 1$ for some $k \in \mathbb{Z}$.

(2) Assume that there is some prime number n such that $n > 2$ and $n = 2k$ for some $k \in \mathbb{Z}$, and derive a contradiction.

While statement (1) more directly translates the plain English statement 'every prime number greater than two is odd', it is the proof strategy suggested by (2) that is easier to use. Indeed, if n is a prime number such that $n > 2$ and $n = 2k$ for some $k \in \mathbb{Z}$, then 2 is a divisor of n other than 1 and n (since $1 < 2 < n$), contradicting the assumption that n is prime.   ◁

The notion of *logical equivalence*, captures precisely the sense in which the logical formulae in (1) and (2) in Example 1.3.1 'mean the same thing'. Being able to transform a logical formula into a different (but equivalent) form allows us to identify a wider range of feasible proof strategies.

Definition 1.3.2
Let p and q be logical formulae. We say that p and q are **logically equivalent**, and write $p \equiv q$ (LaTeX code: `\equiv`), if q can be derived from p and p can be derived from q.

Logical equivalence of propositional formulae

While Definition 1.3.2 defines logical equivalence between arbitrary logical formulae, we will start by focusing our attention on logical equivalence between *propositional* formulae, like those we saw in Section 1.1.

First, let's look at a couple of examples of what proofs of logical equivalence might look like. Be warned—they're not very nice to read! But there is light at the end of the tunnel. After struggling through Examples 1.3.3 and 1.3.4 and Exercise 1.3.5, we will introduce a very quick and easy tool for proving propositional formulae are logically equivalent.

Example 1.3.3
We demonstrate that $p \land (q \lor r) \equiv (p \land q) \lor (p \land r)$, where p, q and r are propositional variables.

- First assume that $p \land (q \lor r)$ is true. Then p is true and $q \lor r$ is true by definition of conjunction. By definition of disjunction, either q is true or r is true.

  ◇ If q is true, then $p \land q$ is true by definition of conjunction.

  ◇ If r is true, then $p \land r$ is true by definition of conjunction.

  In both cases we have that $(p \land q) \lor (p \land r)$ is true by definition of disjunction.

- Now assume that $(p \land q) \lor (p \land r)$ is true. Then either $p \land q$ is true or $p \land r$ is true, by definition of disjunction.

  ◇ If $p \land q$ is true, then p is true and q is true by definition of conjunction.

  ◇ If $p \land r$ is true, then p is true and r is true by definition of conjunction.

  In both cases we have that p is true, and that $q \lor r$ is true by definition of disjunction. Hence $p \land (q \lor r)$ is true by definition of conjunction.

Since we can derive $(p \land q) \lor (p \land r)$ from $p \land (q \lor r)$ and vice versa, it follows that
$$p \land (q \lor r) \equiv (p \land q) \lor (p \land r)$$
as required. ◁

Example 1.3.4
We prove that $p \Rightarrow q \equiv (\neg p) \lor q$, where p, q and r are propositional variables.

- First assume that $p \Rightarrow q$ is true. By the law of exluded middle (Axiom 1.1.44), either p is true or $\neg p$ is true—we derive $(\neg p) \lor q$ in each case.

  ◇ If p is true, then since $p \Rightarrow q$ is true, it follows from ($\Rightarrow$E) that q is true, and so $(\neg p) \lor q$ is true by ($\lor$I$_2$);

⬦ If $\neg p$ is true, then $(\neg p) \vee q$ is true by $(\vee I_1)$.

In both cases, we see that $(\neg p) \vee q$ is true.

- Now assume that $(\neg p) \vee q$ is true. To prove that $p \Rightarrow q$ is true, it suffices by $(\Rightarrow I)$ to assume that p is true and derive q. So assume p is true. Since $(\neg p) \vee q$ is true, we have that either $\neg p$ is true or q is true.

  ⬦ If $\neg p$ is true, then we obtain a contradiction from the assumption that p is true, and so q is true by the principle of explosion (Axiom 1.1.49).

  ⬦ If q is true... well, then q is true—there is nothing more to prove!

  In both cases we have that q is true. Hence $p \Rightarrow q$ is true.

We have derived $(\neg p) \vee q$ from $p \Rightarrow q$ and vice versa, and so the two formulae are logically equivalent.   ◁

Exercise 1.3.5

Let p, q and r be propositional variables. Prove that the propositional formula $(p \vee q) \Rightarrow r$ is logically equivalent to $(p \Rightarrow r) \wedge (q \Rightarrow r)$.   ◁

Working through the derivations each time we want to prove logical equivalence can become cumbersome even for small examples like Examples 1.3.3 and 1.3.4 and Exercise 1.3.5.

The following theorem reduces the problem of proving logical equivalence between *propositional* formulae to the purely algorithmic task of checking when the formulae are true and when they are false in a (relatively) small list of cases. We will streamline this process even further using *truth tables* (Definition 1.3.7).

Theorem 1.3.6

Two propositional formulae are logically equivalent if and only if their truth values are the same under any assignment of truth values to their constituent propositional variables.

Idea of proof

A formal proof of this fact is slightly beyond our reach at this point, although we will be able to prove it formally by *structural induction*, introduced in Section 5.3.

The idea of the proof is that, since propositional formulae are built up from simpler propositional formulae using logical operators, the truth value of a more complex propositional formula is determined by the truth values of its simpler subformulae. If we keep 'chasing' these subformulae, we end up with just propositional variables.

For example, the truth value of $(p \Rightarrow r) \wedge (q \Rightarrow r)$ is determined by the truth values of $p \Rightarrow r$ and $q \Rightarrow r$ according to the rules for the conjunction operator $\wedge$. In turn, the truth value of $p \Rightarrow r$ is determined by the truth values of p and r according to the implication operator $\Rightarrow$, and the truth value of $q \Rightarrow r$ is determined by the truth values of q and r according to the implication operator again. It follows that the truth value of the whole propositional formula $(p \Rightarrow r) \wedge (q \Rightarrow r)$ is determined by the truth values of p, q, r according to the rules for $\wedge$ and $\Rightarrow$.

If some assignment of truth values to propositional variables makes one propositional formula true but another false, then it must be impossible to derive one from the other—otherwise we'd obtain a contradiction. Hence both propositional formulae must have the same truth values no matter what assignment of truth values is given to their constituent propositional variables. □

We now develop a systematic way of checking the truth values of a propositional formula under each assignment of truth values to its constituent propositional variables.

Definition 1.3.7
The **truth table** of a propositional formula is the table with one row for each possible assignment of truth values to its constituent propositional variables, and one column for each subformula (starting with the propositional variables themselves, and ending with the formula itself). The entries of the truth table are the truth values of the subformulae.

Example 1.3.8
The following are the truth tables for $\neg p$, $p \wedge q$, $p \vee q$ and $p \Rightarrow q$.

p	$\neg p$
✓	✕
✕	✓

p	q	$p \wedge q$
✓	✓	✓
✓	✕	✕
✕	✓	✕
✕	✕	✕

p	q	$p \vee q$
✓	✓	✓
✓	✕	✓
✕	✓	✓
✕	✕	✕

p	q	$p \Rightarrow q$
✓	✓	✓
✓	✕	✕
✕	✓	✓
✕	✕	✓

◁

In Example 1.3.8 we have used the symbol ✓ (LATEX code: \checkmark) to mean 'true' and ✕ (LATEX code: \times) to mean 'false'. Some authors adopt other conventions, such as T, F or $\top, \bot$ (LATEX code: \top, \bot) or $1, 0$ or $0, 1$—the possibilites are endless!

Exercise 1.3.9
Use the definitions of $\wedge$, $\vee$ and $\Rightarrow$ to justify the truth tables in Example 1.3.8. ◁

The next example shows how the truth tables for the individual logical operators (as in Example 1.3.8) may be combined to form a truth table for a more complicated propositional formula that involves three propositional variables.

Example 1.3.10

The following is the truth table for $(p \wedge q) \vee (p \wedge r)$.

p	q	r	$p \wedge q$	$p \wedge r$	$(p \wedge q) \vee (p \wedge r)$
✓	✓	✓	✓	✓	✓
✓	✓	×	✓	×	✓
✓	×	✓	×	✓	✓
✓	×	×	×	×	×
×	✓	✓	×	×	×
×	✓	×	×	×	×
×	×	✓	×	×	×
×	×	×	×	×	×

propositional variables   subformulae   main formula

Some comments about the construction of this truth table are pertinent:

- The propositional variables appear first. Since there are three of them, there are $2^3 = 8$ rows. The column for p contains four ✓s followed by four ×s; the column for q contains two ✓s, two ×s, and then repeats; and the column for r contains one ✓, one ×, and then repeats.

- The next group of columns are the next-most complicated subformulae. Each is constructed by looking at the relevant columns further to the left and comparing with the truth table for conjunction.

- The final column is the main formula itself, which again is constructed by looking at the relevant columns further to the left and comparing with the truth table for disjunction.

Our choices of where to put the vertical bars and what order to put the rows in were not the only choices that could have been made, but when constructing truth tables for more complex logical formulae, it is useful to develop a system and stick to it.   ◁

Returning to Theorem 1.3.6, we obtain the following strategy for proving that two propositional formulae are logically equivalent.

> **Strategy 1.3.11** (Logical equivalence using truth tables)
> In order to prove that propositional formulae are logically equivalent, it suffices
> to show that they have the identical columns in a truth table. ◁

Example 1.3.12

In Example 1.3.3 we proved that $p \wedge (q \vee r) \equiv (p \wedge q) \vee (p \wedge r)$. We prove this
again using truth tables. First we construct the truth table for $p \wedge (q \vee r)$:

p	q	r	$q \vee r$	$p \wedge (q \vee r)$
✓	✓	✓	✓	✓
✓	✓	✗	✓	✓
✓	✗	✓	✓	✓
✓	✗	✗	✗	✗
✗	✓	✓	✓	✗
✗	✓	✗	✓	✗
✗	✗	✓	✓	✗
✗	✗	✗	✗	✗

Note that the column for $p \wedge (q \vee r)$ is identical to that of $(p \wedge q) \vee (p \wedge r)$ in
Example 1.3.10. Hence the two formulae are logically equivalent. ◁

To avoid having to write out two truth tables, it can be helpful to combine them
into one. For example, the following truth table exhibits that $p \wedge (q \vee r)$ is logic-
ally equivalent to $(p \wedge q) \vee (p \wedge r)$:

p	q	r	$q \vee r$	$p \wedge (q \vee r)$	$p \wedge q$	$p \wedge r$	$(p \wedge q) \vee (p \wedge r)$
✓	✓	✓	✓	✓	✓	✓	✓
✓	✓	✗	✓	✓	✓	✗	✓
✓	✗	✓	✓	✓	✗	✓	✓
✓	✗	✗	✗	✗	✗	✗	✗
✗	✓	✓	✓	✗	✗	✗	✗
✗	✓	✗	✓	✗	✗	✗	✗
✗	✗	✓	✓	✗	✗	✗	✗
✗	✗	✗	✗	✗	✗	✗	✗

In the following exercises, we use truth tables to repeat the proofs of logical
equivalence from Example 1.3.4 and Exercise 1.3.5.

Exercise 1.3.13

Use a truth table to prove that $p \Rightarrow q \equiv (\neg p) \vee q$. ◁

Exercise 1.3.14

Let p, q and r be propositional variables. Use a truth table to prove that the

propositional formula $(p \vee q) \Rightarrow r$ is logically equivalent to $(p \Rightarrow r) \wedge (q \Rightarrow r)$.
◁

Some proof strategies

We are now in good shape to use logical equivalence to derive some more sophisticated proof strategies.

Theorem 1.3.15 (Law of double negation)
Let p be a propositional variable. Then $p \equiv \neg\neg p$.

Proof
The proof is almost trivialised using truth tables. Indeed, consider the following truth table.

p	$\neg p$	$\neg\neg p$
✓	✗	✓
✗	✓	✗

The columns for p and $\neg\neg p$ are identical, and so $p \equiv \neg\neg p$. □

The law of double negation is important because it suggests a second way that we can prove statements by contradiction. Strategy 1.1.38 says that to prove that a proposition p is *false*, it suffices to assume that p is *true* and derive a contradiction. Using Theorem 1.3.15, we obtain the following similar (but fundamentally different) proof strategy.

Strategy 1.3.16 (Proof by contradiction (indirect version))
In order to prove a proposition p is true, it suffices to assume that p is false and derive a contradiction. ◁

Example 1.3.17
We prove that if a, b and c are non-negative real numbers satisfying $a^2 + b^2 = c^2$, then $a + b \geqslant c$.

Indeed, let $a, b, c \in \mathbb{R}$ with $a, b, c \geqslant 0$, and assume that $a^2 + b^2 = c^2$. Towards a contradiction, assume that it is not the case that $a + b \geqslant c$. Then we must have $a + b < c$. But then

$$(a+b)^2 = (a+b)(a+b) < (a+b)c < c \cdot c = c^2$$

and so
$$c^2 > (a+b)^2 = a^2 + 2ab + b^2 = c^2 + 2ab \geqslant c^2$$
This implies that $c^2 > c^2$, which is a contradiction. So it must be the case that $a + b \geqslant c$, as required. ◁

The next proof strategy we derive concerns proving implications.

Definition 1.3.18
The **contrapositive** of a proposition of the form $p \Rightarrow q$ is the proposition $\neg q \Rightarrow \neg p$.

Theorem 1.3.19 (Law of contraposition)
Let p and q be propositional variables. Then $p \Rightarrow q \equiv (\neg q) \Rightarrow (\neg p)$.

Proof
We build the truth tables for $p \Rightarrow q$ and $(\neg q) \Rightarrow (\neg p)$.

p	q	$p \Rightarrow q$	$\neg q$	$\neg p$	$(\neg q) \Rightarrow (\neg p)$
✓	✓	✓	✗	✗	✓
✓	✗	✗	✓	✗	✗
✗	✓	✓	✗	✓	✓
✗	✗	✓	✓	✓	✓

The columns for $p \Rightarrow q$ and $(\neg q) \Rightarrow (\neg p)$ are identical, so they are logically equivalent. □

Theorem 1.3.19 suggests the following proof strategy.

Strategy 1.3.20 (Proof by contraposition)
In order to prove a proposition of the form $p \Rightarrow q$, it suffices to assume that q is false and derive that p is false. ◁

Example 1.3.21
Fix two natural numbers m and n. We will prove that if $mn > 64$, then either $m > 8$ or $n > 8$.

By contraposition, it suffices to assume that it is *not* the case that $m > 8$ or $n > 8$, and derive that it is not the case that $mn > 64$.

So assume that neither $m > 8$ nor $n > 8$. Then $m \leqslant 8$ and $n \leqslant 8$, so that $mn \leqslant 64$, as required. ◁

Exercise 1.3.22

Use the law of contraposition to prove that $p \Leftrightarrow q \equiv (p \Rightarrow q) \wedge ((\neg p) \Rightarrow (\neg q))$, and use the proof technique that this equivalence suggests to prove that an integer is even if and only if its square is even.   ◁

It feels good to invoke impressive-sounding results like *proof by contraposition*, but in practice, the logical equivalence between *any* two different propositional formulae suggests a new proof technique, and not all of these techniques have names. And indeed, the proof strategy in the following exercise, while useful, has no slick-sounding name—at least, not one that would be widely understood.

Exercise 1.3.23

Prove that $p \vee q \equiv (\neg p) \Rightarrow q$. Use this logical equivalence to suggest a new strategy for proving propositions of the form $p \vee q$, and use this strategy to prove that if two integers sum to an even number, then either both integers are even or both are odd.   ◁

Negation

In pure mathematics it is common to ask whether or not a certain property holds of a mathematical object. For example, in Section 7.2, we will look at convergence of sequences of real numbers: to say that a sequence $x_0, x_1, x_2, \ldots$ of real numbers *converges* (Definition 7.2.15) is to say

$$\exists a \in \mathbb{R}, \forall \varepsilon \in \mathbb{R}, (\varepsilon > 0 \Rightarrow \exists N \in \mathbb{N}, \forall n \in \mathbb{N}, [n \geqslant N \Rightarrow |x_n - a| < \varepsilon])$$

This is already a relatively complicated logical formula. But what if we wanted to prove that a sequence *does not* converge? Simply assuming the logical formula above and deriving a contradiction might work sometimes, but it is not particularly enlightening.

Our next goal is to develop a systematic method for negating complicated logical formulae. With this done, we will be able to negate the logical formula expressing 'the sequence $x_0, x_1, x_2, \ldots$ converges' as follows

$$\forall a \in \mathbb{R}, \exists \varepsilon \in \mathbb{R}, (\varepsilon > 0 \wedge \forall N \in \mathbb{N}, \exists n \in \mathbb{N}, [n \geqslant N \wedge |x_n - a| \geqslant \varepsilon])$$

Granted, this is still a complicated expression, but when broken down element by element, it provides useful information about how it may be proved.

The rules for negating conjunctions and disjunctions are instances of *de Morgan's laws*, which exhibit a kind of duality between $\wedge$ and $\vee$.

> **Theorem 1.3.24** (de Morgan's laws for logical operators)
> Let p and q be logical formulae. Then:
>
> (a) $\neg(p \wedge q) \equiv (\neg p) \vee (\neg q)$; and
>
> (b) $\neg(p \vee q) \equiv (\neg p) \wedge (\neg q)$.

Proof of (a)
Consider the following truth table.

p	q	$p \wedge q$	$\neg(p \wedge q)$	$\neg p$	$\neg q$	$(\neg p) \vee (\neg q)$
✓	✓	✓	✗	✗	✗	✗
✓	✗	✗	✓	✗	✓	✓
✗	✓	✗	✓	✓	✗	✓
✗	✗	✗	✓	✓	✓	✓

The columns for $\neg(p \wedge q)$ and $(\neg p) \vee (\neg q)$ are identical, so they are logically equivalent. □

Exercise 1.3.25
Prove Theorem 1.3.24(b) thrice, once using the definition of logical equivalence directly (like we did in Examples 1.3.3 and 1.3.4 and Exercise 1.3.5), once using a truth table, and once using part (a) together with the law of double negation. ◁

Example 1.3.26
We often use de Morgan's laws for logical operators without thinking about it. For example to say that 'neither 3 nor 7 is even' is equivalent to saying '3 is odd and 7 is odd'. The former statement translates to

$$\neg[(3 \text{ is even}) \vee (7 \text{ is even})]$$

while the second statement translates to

$$[\neg(3 \text{ is even})] \wedge [\neg(7 \text{ is even})]$$

◁

Exercise 1.3.27
Prove that $\neg(p \Rightarrow q) \equiv p \wedge (\neg q)$ twice, once using a truth table, and once using Exercise 1.3.13 together with de Morgan's laws and the law of double negation. ◁

De Morgan's laws for logical operators generalise to statements about quantifiers, expressing a similar duality between $\forall$ and $\exists$ as we have between $\wedge$ and $\vee$.

Theorem 1.3.28 (de Morgan's laws for quantifiers)
let $p(x)$ be a logical formula with free variable x ranging over a set X. Then:

(a) $\neg \forall x \in X, p(x) \equiv \exists x \in X, \neg p(x)$; and

(b) $\neg \exists x \in X, p(x) \equiv \forall x \in X, \neg p(x)$.

Proof
Unfortunately, since these logical formulae involve quantifiers, we do not have truth tables at our disposal, so we must assume each formula and derive the other.

We start by proving the equivalence in part (b), and then we derive (a) as a consequence.

- Assume $\neg \exists x \in X, p(x)$. To prove $\forall x \in X, \neg p(x)$, fix some $x \in X$. If $p(x)$ were true, then we'd have $\exists x \in X, p(x)$, which contradicts our main assumption; so we have $\neg p(x)$. But then $\forall x \in X, \neg p(x)$ is true.

- Assume $\forall x \in X, \neg p(x)$. For the sake of contradiction, assume $\exists x \in X, p(x)$ were true. Then we obtain some $a \in X$ for which $p(a)$ is true. But $\neg p(a)$ is true by the assumption that $\forall x \in X, \neg p(a)$, so we obtain a contradiction. Hence $\neg \exists x \in X, p(x)$ is true.

This proves that $\neg \exists x \in X, p(x) \equiv \forall x \in X, \neg p(x)$.

Now (a) follows from (b) using the law of double negation (Theorem 1.3.15):

$$\exists x \in X, \neg p(x) \equiv \neg \neg \exists x \in X, \neg p(x) \stackrel{(b)}{\equiv} \neg \forall x \in X, \neg \neg p(x) \equiv \neg \forall x \in X, p(x)$$

as required.   □

The proof strategy suggested by the logical equivalence in Theorem 1.3.28(b) is so important that it has its own name.

Strategy 1.3.29 (Proof by counterexample)
To prove that a proposition of the form $\forall x \in X, p(x)$ is false, it suffices to find a single element $a \in X$ such that $p(a)$ is false. The element a is called a **counterexample** to the proposition $\forall x \in X, p(x)$.   ◁

Example 1.3.30
We prove by counterexample that not every integer is divisible by a prime number. Indeed, let $x = 1$. The only integral factors of 1 are 1 and -1, neither of which are prime, so that 1 is not divisible by any primes.   ◁

Exercise 1.3.31

Prove by counterexample that not every rational number can be expressed as $\dfrac{a}{b}$ where $a \in \mathbb{Z}$ is even and $b \in \mathbb{Z}$ is odd. ◁

We have now seen how to negate the logical operators $\neg$, $\wedge$, $\vee$ and $\Rightarrow$, as well as the quantifiers $\forall$ and $\exists$.

Definition 1.3.32

A logical formula is **maximally negated** if the only instances of the negation operator $\neg$ appear immediately before a predicate (or other proposition not involving logical operators or quantifiers).

Example 1.3.33

The following propositional formula is maximally negated:

$$[p \wedge (q \Rightarrow (\neg r))] \Leftrightarrow (s \wedge (\neg t))$$

Indeed, all instances of $\neg$ appear immediately before propositional variables.

However the following propositional formula is *not* mmaximally negated:

$$(\neg\neg q) \Rightarrow q$$

Here the subformula $\neg\neg q$ contains a negation operator immediately before another negation operator $(\neg\neg q)$. However by the law of double negation, this is equivalent to $q \Rightarrow q$, which is maximally negated trivially since there are no negation operators to speak of. ◁

Exercise 1.3.34

Determine which of the following logical formulae are maximally negated.

(a) $\forall x \in X, (\neg p(x)) \Rightarrow \forall y \in X, \neg(r(x,y) \wedge s(x,y))$;

(b) $\forall x \in X, (\neg p(x)) \Rightarrow \forall y \in X, (\neg r(x,y)) \vee (\neg s(x,y))$;

(c) $\forall x \in \mathbb{R}, [x > 1 \Rightarrow (\exists y \in \mathbb{R}, [x < y \wedge \neg(x^2 \leqslant y)])]$;

(d) $\neg\exists x \in \mathbb{R}, [x > 1 \wedge (\forall y \in \mathbb{R}, [x < y \Rightarrow x^2 \leqslant y])]$.

◁

The following theorem allows us to replace logical formulae by maximally negated ones, which in turn suggests proof strategies that we can use for proving that complicated-looking propositions are *false*.

> **Theorem 1.3.35**
> Every logical formula (built using only the logical operators and quantifiers we
> have seen so far) is logically equivalent to a maximally negated logical formula.

Idea of proof

Much like Theorem 1.3.6, a precise proof of Theorem 1.3.35 requires some form
of induction argument, so instead we will give an idea of the proof.

Every logical formula we have seen so far is built from predicates using the lo-
gical operators $\wedge, \vee, \Rightarrow$ and $\neg$ and the quantifiers $\forall$ and $\exists$—indeed, the logical
operator $\Leftrightarrow$ was defined in terms of $\wedge$ and $\Rightarrow$, and the quantifier $\exists$ was defined in
terms of the quantifiers $\forall$ and $\exists$ and the logical operators $\wedge$ and $\Rightarrow$.

But the results in this section allow us to push negations 'inside' each of these
logical operators and quantifiers, as summarised in the following table.

Negation outside		Negation inside	Proof
$\neg(p \wedge q)$	$\equiv$	$(\neg p) \vee (\neg q)$	Theorem 1.3.24(a)
$\neg(p \vee q)$	$\equiv$	$(\neg p) \wedge (\neg q)$	Theorem 1.3.24(b)
$\neg(p \Rightarrow q)$	$\equiv$	$p \wedge (\neg q)$	Exercise 1.3.27
$\neg(\neg p)$	$\equiv$	p	Theorem 1.3.15
$\neg \forall x \in X, \, p(x)$	$\equiv$	$\exists x \in X, \, \neg p(x)$	Theorem 1.3.28(a)
$\neg \exists x \in X, \, p(x)$	$\equiv$	$\forall x \in X, \, \neg p(x)$	Theorem 1.3.28(b)

Repeatedly applying these rules to a logical formula eventually yields a logically
equivalent, maximally negated logical formula.   □

Example 1.3.36

Recall the logical formula expressing the assertion that a sequence $x_0, x_1, x_2, \ldots$
of real numbers converges:

$$\exists a \in \mathbb{R}, \, \forall \varepsilon \in \mathbb{R}, \, (\varepsilon > 0 \Rightarrow \exists N \in \mathbb{N}, \, \forall n \in \mathbb{N}, \, [n \geqslant N \Rightarrow |x_n - a| < \varepsilon])$$

We will maximally negate this to obtain a logical formula expressing the assertion
that the sequence does not converge.

Let's start at the beginning. The negation of the formula we started with is:

$$\neg \exists a \in \mathbb{R}, \, \forall \varepsilon \in \mathbb{R}, \, (\varepsilon > 0 \Rightarrow \exists N \in \mathbb{N}, \, \forall n \in \mathbb{N}, \, [n \geqslant N \Rightarrow |x_n - a| < \varepsilon])$$

The key to maximally negating a logical formula is to ignore information that
is not immediately relevant. Here, the expression that we are negating takes the
form $\neg \exists a \in \mathbb{R}$, (stuff). It doesn't matter what the 'stuff' is just yet; all that matters
is that we are negating an existentially quantified statement, and so de Morgan's

laws for quantifiers tells us that this is logically equivalent to $\forall a \in \mathbb{R},\, \neg(\text{stuff})$. We apply this rule and just re-write the 'stuff', to obtain:

$$\forall a \in \mathbb{R},\, \neg\forall \varepsilon \in \mathbb{R},\, (\varepsilon > 0 \Rightarrow \exists N \in \mathbb{N},\, \forall n \in \mathbb{N},\, [n \geqslant N \Rightarrow |x_n - a| < \varepsilon])$$

Now we are negating a universally quantified statement, $\neg\forall \varepsilon \in \mathbb{R},\, (\text{stuff})$ which, by de Morgan's laws for quantifiers, is equivalent to $\exists \varepsilon \in \mathbb{R},\, (\text{stuff})$:

$$\forall a \in \mathbb{R},\, \exists \varepsilon \in \mathbb{R},\, \neg(\varepsilon > 0 \Rightarrow \exists N \in \mathbb{N},\, \forall n \in \mathbb{N},\, [n \geqslant N \Rightarrow |x_n - a| < \varepsilon])$$

At this point, the statement being negated is of the form $(\text{stuff}) \Rightarrow (\text{junk})$, which by Exercise 1.3.27 negates to $(\text{stuff}) \wedge \neg(\text{junk})$. Here, 'stuff' is $\varepsilon > 0$ and 'junk' is $\exists N \in \mathbb{N}, \forall n \in \mathbb{N},\, [n \geqslant N \Rightarrow |x_n - a| < \varepsilon]$. So performing this negation yields:

$$\forall a \in \mathbb{R},\, \exists \varepsilon \in \mathbb{R},\, (\varepsilon > 0 \wedge \neg \exists N \in \mathbb{N},\, \forall n \in \mathbb{N},\, [n \geqslant N \Rightarrow |x_n - a| < \varepsilon])$$

Now we are negating an existentially quantified formula again, so using de Morgan's laws for quantifiers gives:

$$\forall a \in \mathbb{R},\, \exists \varepsilon \in \mathbb{R},\, (\varepsilon > 0 \wedge \forall N \in \mathbb{N},\, \neg\forall n \in \mathbb{N},\, [n \geqslant N \Rightarrow |x_n - a| < \varepsilon])$$

The formula being negated here is univerally quantified, so using de Morgan's laws for quantifiers *again* gives:

$$\forall a \in \mathbb{R},\, \exists \varepsilon \in \mathbb{R},\, (\varepsilon > 0 \wedge \forall N \in \mathbb{N},\, \exists n \in \mathbb{N},\, \neg[n \geqslant N \Rightarrow |x_n - a| < \varepsilon])$$

We're almost there! The statement being negated here is an implication, so applying the rule $\neg(p \Rightarrow q) \equiv p \wedge (\neg q)$ again yields:

$$\forall a \in \mathbb{R},\, \exists \varepsilon \in \mathbb{R},\, (\varepsilon > 0 \wedge \forall N \in \mathbb{N},\, \exists n \in \mathbb{N},\, [n \geqslant N \wedge \neg(|x_n - a| < \varepsilon)])$$

At this point, strictly speaking, the formula is maximally negated, since the statement being negated does not involve any other logical opreators or quantifiers. However, since $\neg(|x_n - a| < \varepsilon)$ is equivalent to $|x_n - a| \geqslant \varepsilon$, we can go one step further to obtain:

$$\forall a \in \mathbb{R},\, \exists \varepsilon \in \mathbb{R},\, (\varepsilon > 0 \wedge \forall N \in \mathbb{N},\, \exists n \in \mathbb{N},\, [n \geqslant N \wedge |x_n - a| \geqslant \varepsilon])$$

This is as negated as we could ever dream of, and so we stop here. ◁

Exercise 1.3.37
Find a maximally negated propositional formula that is logically equivalent to $\neg(p \Leftrightarrow q)$. ◁

Exercise 1.3.38
Maximally negate the following logical formula, then prove that it is true or prove that it is false.

$$\exists x \in \mathbb{R},\, [x > 1 \wedge (\forall y \in \mathbb{R},\, [x < y \Rightarrow x^2 \leqslant y])]$$

◁

Tautologies

The final concept that we introduce in this chapter is that of a *tautology*, which can be thought of as the opposite of a contradiction. The word 'tautology' has other implications when used colloquially, but in the context of symbolic logic it has a precise definition.

Definition 1.3.39
A **tautology** is a proposition or logical formula that is true, no matter how truth values are assigned to its component propositional variables and predicates.

The reason we are interested in tautologies is that tautologies can be used as assumptions at any point in a proof, for any reason.

Strategy 1.3.40 (Assuming tautologies)
Let p be a proposition and let t be a tautology. In order to prove p, it suffices to prove that p is true from the assumption that t is true.   ◁

Example 1.3.41
The law of excluded middle (Axiom 1.1.44) says precisely that $p \vee (\neg p)$ is a tautology.   ◁

Example 1.3.42
The formula $p \Rightarrow (q \Rightarrow p)$ is a tautology.

A direct proof of this fact is as follows. In order to prove $p \Rightarrow (q \Rightarrow p)$ is true, it suffices to assume p and derive $q \Rightarrow p$. So assume p. Now in order to prove $q \Rightarrow p$, it suffices to assume q and derive p. So assume q. But we're already assuming that p is true! So $q \Rightarrow p$ is true, and hence $p \Rightarrow (q \Rightarrow p)$ is true.

A proof using truth tables is as follows:

p	q	$q \Rightarrow p$	$p \Rightarrow (q \Rightarrow p)$
✓	✓	✓	✓
✓	✗	✓	✓
✗	✓	✗	✓
✗	✗	✓	✓

We see that $p \Rightarrow (q \Rightarrow p)$ is true regardless of the truth values of p and q.   ◁

Exercise 1.3.43
Prove that each of the following is a tautology:

(a) $[(p \Rightarrow q) \wedge (q \Rightarrow r)] \Rightarrow (p \Rightarrow r)$;

(b) $[p \Rightarrow (q \Rightarrow r)] \Rightarrow [(p \Rightarrow q) \Rightarrow (p \Rightarrow r)]$;

(c) $\exists y \in Y, \forall x \in X, p(x,y) \Rightarrow \forall x \in X, \exists y \in Y, p(x,y)$;

(d) $[\neg(p \wedge q)] \Leftrightarrow [(\neg p) \vee (\neg q)]$;

(e) $(\neg \forall x \in X, p(x)) \Leftrightarrow (\exists x \in X, \neg p(x))$.

$\triangleleft$

You may have noticed parallels between de Morgan's laws for logical operators and quantifiers, and parts (d) and (e) of Exercise 1.3.43, respectively. They almost seem to say the same thing, except that in Exercise 1.3.43 we used '$\Leftrightarrow$' and in Theorems 1.3.24 and 1.3.28 we used '$\equiv$'. There is an important difference, though: if p and q are logical formulae, then $p \Rightarrow q$ is itself a logical formula, which we may study as a mathematical object in its own right. However, $p \equiv q$ is not a logical formula: it is an assertion *about* logical formulae, namely that the logical formulae p and q are equivalent.

There is, nonetheless, a close relationship between $\Leftrightarrow$ and $\equiv$—this relationship is summarised in the following theorem.

Theorem 1.3.44
Let p and q be logical formulae.

(a) q can be derived from p if and only if $p \Rightarrow q$ is a tautology;

(b) $p \equiv q$ if and only if $p \Leftrightarrow q$ is a tautology.

Proof
For (a), note that a derivation of q from p is sufficient to establish the truth of $p \Rightarrow q$ by the introduction rule for conjunction ($\Rightarrow$I), and so if q can be derived from p, then $p \Rightarrow q$ is a tautology. Conversely, if $p \Rightarrow q$ is a tautology, then q can be derived from p using the elimination rule for conjunction ($\Rightarrow$E) together with the (tautological) assumption that $p \Rightarrow q$ is true.

Now (b) follows from (a), since logical equivalence is defined in terms of derivation in each direction, and $\Leftrightarrow$ is simply the conjunction of two implications. $\quad \square$

Aaand breathe! All this new notation can be overwhelming at first, but it will be worth it in the end. This chapter was all about teaching you a new language— new symbols, new terminology—because without it, our future pursuits will be

impossible. If you're stuck now, then don't worry: you'll soon get the hang of it, especially when we start using this new language in context. You can, of course, refer back to the results in this chapter for reference at any point in the future.

Section 1.Q
Chapter 1 exercises

Under construction!
The end-of-chapter exercise sections are new and in an incomplete state.

1. For fixed $n \in \mathbb{N}$, let p represent the proposition 'n is even', let q represent the proposition 'n is prime' and let r represent the proposition '$n = 2$'. For each of the following propositional formulae, translate it into plain English and determine whether it is true for all $n \in \mathbb{N}$, true for some values of n and false for some values of n, or false for all $n \in \mathbb{N}$.

(a) $(p \wedge q) \Rightarrow r$

(b) $q \wedge (\neg r) \Rightarrow (\neg p)$

(c) $(\neg p) \vee (\neg q) \vee (\neg r)$

(d) $p \wedge q \wedge (\neg r)$

2. Find a statement in plain English, involving no variables at all, that is equivalent to the logical formula $\forall a \in \mathbb{Q}, \forall b \in \mathbb{Q}, (a < b \Rightarrow \exists c \in \mathbb{R}, [a < c < b \wedge \neg(c \in \mathbb{Q})])$. Then prove this statement, using the structure of the logical formula as a guide.

3. Find a purely symbolic logical formula that is equivalent to the following statement, and then prove it: "*No matter which integer you may choose, there will be an integer greater than it.*"

4. Prove that
$$p \Leftrightarrow q \equiv (p \Rightarrow q) \wedge ((\neg p) \Rightarrow (\neg q))$$
How might this logical equivalence help you to prove statements of the form 'p if and only if q'?

5. Prove using truth tables that $p \Rightarrow q \not\equiv q \Rightarrow p$. Give an example of propositions p and q such that $p \Rightarrow q$ is true but $q \Rightarrow p$ is false.

6. A new logical operator $\uparrow$ is defined by the following rules:

(i) If a contradiction can be derived from the assumption that p is true, then $p \uparrow q$ is true;

(ii) If a contradiction can be derived from the assumption that q is true, then $p \uparrow q$ is true;

(iii) If r is any proposition, and if $p \uparrow q$, p and q are all true, then r is true.

This question explores this curious new logical operator.

(a) Prove that $p \uparrow p \equiv \neg p$, and deduce that $((p \uparrow p) \uparrow (p \uparrow p)) \equiv p$.

(b) Prove that $p \vee q \equiv (p \uparrow p) \uparrow (q \uparrow q)$ and $p \wedge q \equiv (p \uparrow q) \uparrow (p \uparrow q)$.

(c) Find a propositional formula using only the logical operator $\uparrow$ that is equivalent to $p \Rightarrow q$.

7. Let X be $\mathbb{Z}$ or $\mathbb{Q}$, and define a logical formula p by:

$$\forall x \in X, \exists y \in X, (x < y \wedge [\forall z \in X, \neg(x < z \wedge z < y)])$$

Write out $\neg p$ as a maximally negated logical formula. Prove that p is true when $X = \mathbb{Z}$, and p is false when $X = \mathbb{Q}$.

8. Use Definition 1.2.26 to write out a maximally negated logical formula that is equivalent to $\neg \exists! x \in X, p(x)$. Describe the strategy that this equivalence suggests for proving that there is not a unique $x \in X$ such that $p(x)$ is true, and use this strategy to prove that, for all $a \in \mathbb{R}$, if $a \neq -1$ then there is not a unique $x \in \mathbb{R}$ such that $x^4 - 2ax^2 + a^2 - 1 = 0$.

Sets and functions

Now that we have a precise way of reasoning mathematically, it's time to start doing some mathematics!

In Chapter 0 we gave a preliminary definition of a 'set' as a collection of objects (Definition 0.3), but then we focused almost exclusively on the number sets $\mathbb{N}$, $\mathbb{Z}$, $\mathbb{Q}$, $\mathbb{R}$ and $\mathbb{C}$ in Chapter 0 and Chapter 1.

Our first task in this chapter, in Section 2.1, is to make the notion of a set slightly more precise, and to get comfortable with reasoning about sets in the abstract—this is extremely important, as sets are the building blocks of pure mathematics. We will study how different sets relate to one another, and how to build new sets out of old.

Just as fundamental as sets, the concept of a *function* is central to almost every mathematical field. We will use functions heavily throughout the book—they are so important that we have devoted not one, but *two* sections to them. Our first exposure to functions in in Section 2.2, where we will define the notion of a function and explore their basic properties.

In Section 2.3 we study two properties that functions might have: *injectivity* and *surjectivity*. These conditions are used to compare sizes of sets, amongst other things, and they arise frequently in areas of mathematics where functions are used.

Sets and set operations

We begin by redefining the notion of a *set* with a notch more precision than we provided in Chapter 0. At their core, sets seem extremely simple—sets are just collections of objects—except that if not kept in check, this characterisation of a set leads to logical inconsistencies, such as the infamous *Russell's paradox*.

These logical paradoxes can be overcome by restricting ourselves to working inside a *universe* $\mathscr{U}$, which we consider to be a set which is so big that it contains all of the mathematical objects that we want to talk about. This is a subtle issue, which is well beyond the scope of this section, but is discussed further in Section B.1.

Definition 2.1.1

A **set** is a collection of **elements** from a specified **universe of discourse**. The collection of everything in the universe of discourse is called the **universal set**, denoted by $\mathscr{U}$ (LaTeX code: `\mathcal{U}`).

The expression $x \in X$ (LaTeX code: `\in`) denotes the statement that x is an element of X; we write $x \notin X$ (LaTeX code: `\not\in`) to mean $\neg(x \in X)$, that is that x is not an element of X.

Example 2.1.2
In Chapter 0, we introduced five sets: the set $\mathbb{N}$ of natural numbers, the set $\mathbb{Z}$ of integers, the set $\mathbb{Q}$ of rational numbers, the set $\mathbb{R}$ of real numbers and the set $\mathbb{C}$ of complex numbers. ◁

Exercise 2.1.3
Which of the following propositions are true, and which are false?

$$\frac{1}{2} \in \mathbb{Z} \qquad \frac{1}{2} \in \mathbb{Q} \qquad \mathbb{Z} \in \mathbb{Q} \qquad \mathbb{Z} \in \mathscr{U} \qquad \frac{1}{2} \in \mathscr{U}$$

◁

We will avoid referring explicitly to the universal set $\mathscr{U}$ whenever possible, but it will always be there in the background. This is convenient because we no longer need to worry about the domain of discourse of free variables (as we did in Definition 1.2.2), so that we can abbreviate '$\forall x \in \mathscr{U}, p(x)$' by '$\forall x, p(x)$', and '$\exists x \in \mathscr{U}, p(x)$' by '$\exists x, p(x)$'.

Note that under this convention:

- $\forall x \in X$, $p(x)$ is logically equivalent to $\forall x$, $(x \in X \Rightarrow p(x))$; and

- $\exists x \in X$, $p(x)$ is logically equivalent to $\exists x$, $(x \in X \land p(x))$.

Specifying a set

One way of defining a set is simply to describe it in words, like we have done up to now. There are other, more concise ways of specifying sets, which also remove such ambiguity from the process.

Lists. One way is simply to provide a **list** of the elements of the set. To specify that the list denotes a set, we enclose the list with curly brackets $\{,\}$ (LaTeX code: `\{,\}`). For example, the following is a specification of a set X, whose elements are the natural numbers between 0 and 5 (inclusive):

$$X = \{0,1,2,3,4,5\}$$

Implied lists. Sometimes a list might be too long to write out—maybe even infinite—or the length of the list might depend on a variable. In these cases it will be convenient to use an **implied list**, in which some elements of the list are written, and the rest are left implicit by writing an ellipsis '...' (LaTeX code: `\dots`). For example, the statement

$$X = \{1,4,9,\ldots,n^2\}$$

means that X is the set whose elements are all the square numbers from 1 to n^2, where n is some number. Implied lists can be ambiguous, since they rely on the reader's ability to infer the pattern being followed, so use with caution!

Set-builder notation. In general, implied lists can be ambiguous, so in practice they are avoided unless the implied list is very simple, such as a set of consecutive numbers like $\{3,4,\ldots,9\}$. In fact, many sets can't even be listed in this way.

To get around this, we can use *set-builder notation*, which is a means of specifying a set in terms of the properties its elements satisfy. Given a set X, the set of elements of X satisfying some property $p(x)$ is denoted

$$\{x \in X \mid p(x)\}$$

The bar '|' (LaTeX code: `\mid`) separates the variable name from the formula that they make true—some authors use a colon instead (as in $\{x \in X : p(x)\}$).

The set $\{x \in X \mid p(x)\}$ is read aloud as 'the set of $x \in X$ such that $p(x)$', but beware—neither the bar '|' nor the colon ':' mean 'such that' in other contexts.

Example 2.1.4

The set of all even integers can be written in set-builder notation as

$$\{n \in \mathbb{Z} \mid n \text{ is even}\}$$

For comparison, the set of all even natural numbers can be written as

$$\{n \in \mathbb{N} \mid n \text{ is even}\} = \{0, 2, 4, 6, \dots\}$$

Note that -6 is an element of the former set but not of the latter set, since -6 is an integer but is not a natural number.

Note moreover that the expression

$$\{n \in \mathbb{Q} \mid n \text{ is even}\}$$

is meaningless, since we have not defined a notion of 'evenness' for rational numbers.   ◁

Strategy 2.1.5

Let X be a set and let $p(x)$ be a logical formula with free variable $x \in X$. In order to prove $a \in \{x \in X \mid p(x)\}$, it suffices to prove $a \in X$ and that $p(a)$ is true.   ◁

Exercise 2.1.6

A **dyadic rational** is a rational number that can be expressed as an integer divided by a power of 2. Express the set of all dyadic rationals using set-builder notation.
◁

An alternate form of set-builder notation uses an expression involving one or more variables to the left of the vertical bar, and the range of the variable(s) to the right. The elements of the set are then the values of the expression as the variable(s) vary as indicated—that is:

$$\{\text{expr}(x) \mid x \in X\} \text{ is defined to mean } \{y \mid \exists x \in X, y = \text{expr}(x)\}$$

where $\text{expr}(x)$ is the expression in question.

Example 2.1.7

The expression $\{3k+2 \mid k \in \mathbb{Z}\}$ denotes the set of all integers of the form $3k+2$, where $k \in \mathbb{Z}$. It is shorthand for $\{n \in \mathbb{Z} \mid \exists k \in \mathbb{Z}, n = 3k+2\}$. In implied list notation, we could write this set as $\{\dots, -4, -1, 2, 5, 8, \dots\}$.   ◁

Exercise 2.1.8

Express the set of dyadic rationals (defined in Exercise 2.1.6) in this alternate form of set-builder notation.   ◁

Set-builder notation is useful for defining sets based on the properties they satisfy, as in Definitions 2.1.9 and 2.1.11 below.

Definition 2.1.9
Let $n \in \mathbb{N}$. The set $[n]$ is defined by $[n] = \{k \in \mathbb{N} \mid 1 \leqslant k \leqslant n\}$.

Example 2.1.10
In implied list notation, $[n] = \{1, 2, \ldots, n\}$. For example, $[4] = \{1, 2, 3, 4\}$. Note that $[0]$ has no elements (it is *empty*—see Definition 2.1.26), since there are no natural numbers k satisfying the inequality $1 \leqslant k \leqslant 0$. ◁

While not particularly interesting yet, sets of the form $[n]$ will be fundamental throughout Chapter 3, as they are used to define the notion of a *finite set*, as well as the *size* of a finite set.

Intervals are particular subsets of $\mathbb{R}$ that are ubiquitous in mathematics, particularly in analysis and topology.

Definition 2.1.11 (Intervals of the real line)
Let $a, b \in \mathbb{R}$. The **open interval** (a, b), the **closed interval** $[a, b]$, and the **half-open intervals** $[a, b)$ and $(a, b]$ from a to b are defined by

$$(a,b) = \{x \in \mathbb{R} \mid a < x < b\} \qquad (a,b] = \{x \in \mathbb{R} \mid a < x \leqslant b\}$$
$$[a,b) = \{x \in \mathbb{R} \mid a \leqslant x < b\} \qquad [a,b] = \{x \in \mathbb{R} \mid a \leqslant x \leqslant b\}$$

We further define the **unbounded intervals** $(-\infty, a)$, $(-\infty, a]$, $[a, \infty)$ and (a, ∞) (LaTeX code: `\infty`) by

$$(-\infty,a) = \{x \in \mathbb{R} \mid x < a\} \qquad (a,\infty) = \{x \in \mathbb{R} \mid x > a\}$$
$$(-\infty,a] = \{x \in \mathbb{R} \mid x \leqslant a\} \qquad [a,\infty) = \{x \in \mathbb{R} \mid x \geqslant a\}$$

Example 2.1.12
The following illustration depicts the open interval $(-2, 5)$.

The hollow circles ○ indicate that the endpoints are not included in the interval. ◁

Be warned that the use of the symbol ∞ is misleading, since it suggests that the symbol ∞ on its own has a specific meaning (or, worse, that it refers to a

real number). It doesn't—it is just a symbol that suggests unboundedness of the interval in question. A less misleading way of writing $[a, \infty)$, for instance, might be $[a, \rightarrow)$ or $\mathbb{R}^{\geqslant a}$; however, $[a, \infty)$ is standard, so it is what we will write.

Exercise 2.1.13

For each of the following illustrations, find the interval that it depicts. A filled circle • indicates that an end-point is included in the interval, whereas a hollow circle ◦ indicates that an end-point is not included in the interval.

Subsets

It is often the case that everything that is also an element of one set is an element of another set. For example, every integer is a rational number; that is

$$\forall n \in \mathbb{Z}, n \in \mathbb{Q}$$

We can say this more concisely by saying that $\mathbb{Z}$ is a *subset* of $\mathbb{Q}$.

Definition 2.1.14
Let X be a set. A **subset** of X is a set U such that

$$\forall a, (a \in U \Rightarrow a \in X)$$

We write $U \subseteq X$ (LaTeX code: `\subseteq`) for the assertion that U is a subset of X.

Additionally, the notation $U \nsubseteq X$ (LaTeX code: `\nsubseteq`) means that U is not a subset of X, and the notation $U \subsetneqq X$ (LaTeX code: `\subsetneqq`) means that U is a **proper subset** of X, that is a subset of X that is not equal to X.

> **Strategy 2.1.15** (Proving a subset containment)
> In order to prove that a set U is a subset of a set X, it suffices to take an arbitrary element $a \in U$ and prove that $a \in X$. ◁

Example 2.1.16
Every set is a subset of itself—that is, $X \subseteq X$ for all sets X. The proof of this is extremely simple: we must prove $\forall x \in X, x \in X$. But then this is trivial: let $x \in X$, then $x \in X$ by assumption. Done! ◁

Example 2.1.17
Let $a, b, c, d \in \mathbb{R}$ with $a < c < d < b$. Then $[c, d] \subseteq (a, b)$. Indeed, let $x \in [c, d]$. Then $c \leqslant x \leqslant d$. But then

$$a < c \leqslant x \leqslant d < b \quad \Rightarrow \quad a < x < b$$

so that $[c, d] \subseteq (a, b)$, as required. ◁

Exercise 2.1.18
Let $a, b, c, d \in \mathbb{R}$ with $a < b$ and $c < d$. Prove that $[a, b) \subseteq (c, d]$ if and only if $a \geqslant c$ and $b \leqslant d$. ◁

Example 2.1.19
The number sets from Chapter 0 are related by the following chain of subset inclusions.

$$\mathbb{N} \subseteq \mathbb{Z} \subseteq \mathbb{Q} \subseteq \mathbb{R} \subseteq \mathbb{C}$$

◁

The following proposition proves a property of subsethood known as *transitivity*—we'll revisit this property in Sections 5.1 and 5.2.

Proposition 2.1.20
Let X, Y, Z be sets. If $X \subseteq Y$ and $Y \subseteq Z$, then $X \subseteq Z$.

Proof
Suppose that $X \subseteq Y$ and $Y \subseteq Z$. We need to prove $X \subseteq Z$.

So let $a \in X$. Since $X \subseteq Y$, it follows from Definition 2.1.14 that $a \in Y$; and since $Y \subseteq Z$, it follows again from Definition 2.1.14 that $a \in Z$.

Hence $X \subseteq Z$, as required. ☐

Set equality

This section is all about defining sets, comparing sets, and building new sets from old, and so to make much more progress, we first need to establish what we mean when we say that two sets are *equal*.

Discussion 2.1.21

Let X and Y be sets. What should it mean to say that X and Y are equal? Try to provide a precise definition of equality of sets before reading on.   ◁

There are different possible notions of 'sameness' for sets: we might want to say that two sets X and Y are equal when they have quite literally the same definition; or we might want to say that X and Y are equal when they contain the same objects as elements. For instance, suppose X is 'the set of all odd natural numbers' and Y is 'the set of all integers that are differences of consecutive perfect squares'— in this case, the first of these characterisations of equality might lead us to say $X \neq Y$, whereas the second would lead us to say $X = Y$.

Clearly, we have to state our terms at some point. And that point is now.

Axiom 2.1.22 (Set extensionality)

Let X and Y be sets. Then $X = Y$ if and only if $\forall a, (a \in X \Leftrightarrow a \in Y)$, or equivalently, if $X \subseteq Y$ and $Y \subseteq X$.

This characterisation of set equality suggests the following strategy for proving that two sets are equal.

Strategy 2.1.23 (Proof by double containment)

In order to prove that a set X is equal to a set Y, it suffices to:

- Prove $X \subseteq Y$, i.e. let $a \in X$ be an arbitrary element, and derive $a \in Y$; and then

- Prove $X \supseteq Y$, i.e. let $a \in Y$ be an arbitrary element, and derive $a \in X$.

We often write '($\subseteq$)' and '($\supseteq$)' to indicate the direction of the containment being proved.   ◁

Example 2.1.24

We prove that $\{x \in \mathbb{R} \mid x^2 \leqslant 1\} = [-1, 1]$ by double containment.

- ($\subseteq$) Let $a \in \{x \in \mathbb{R} \mid x^2 \leqslant 1\}$. Then $a \in \mathbb{R}$ and $a^2 \leqslant 1$, so that $(1-a)(1+a) = 1 - a^2 \geqslant 0$. It follows that either:

  ◇ $1 - a \geqslant 0$ and $1 + a \geqslant 0$, in which case $a \leqslant 1$ and $a \geqslant -1$, so that $a \in [-1, 1]$.

  ◇ $1 - a \leqslant 0$ and $1 + a \leqslant 0$, in which case $a \geqslant 1$ and $a \leqslant -1$, which is a contradiction since $-1 < 1$.

  So we must have $a \in [-1, 1]$, as required.

- ($\supseteq$) Let $a \in [-1, 1]$. Then $-1 \leqslant a \leqslant 1$, so $|a| \leqslant 1$, and hence $a^2 = |a|^2 \leqslant 1$, so that $a \in \{x \in \mathbb{R} \mid x^2 \leqslant 1\}$, as required.

$\lhd$

Exercise 2.1.25
Prove that $\{x \in \mathbb{R} \mid x^2 < x\} = (0, 1)$. $\lhd$

Inhabitation and emptiness

Another fundamental example of a set is the *empty set*, which is the set with no elements. But we have to be slightly careful about how we use the word 'the', since it implies *uniqueness*, and we don't know (yet) that two sets with no elements are necessarily equal. So first we will define what it means for a set to be empty, and then we'll show that there is exactly one empty set.

Definition 2.1.26
A set X is **inhabited** (or **nonempty**) if it has at least one element; otherwise, it is **empty**.

The assertion that X is inhabited is equivalent to the logical formula $\exists a, a \in X$, and the assertion that X is empty is equivalent to the logical formula $\neg \exists a, a \in X$. This suggests the following strategy for proving that a set is inhabited, or that it is empty.

Strategy 2.1.27 (Proving that a set is inhabited or empty)
In order to prove a set X is inhabited, it suffices to exhibit an element. In order to prove a set X is empty, assume that X is inhabited—that is, that there is some element $a \in X$—and derive a contradiction. $\lhd$

In other texts, the term *nonempty* is more common than *inhabited*, but there are reasons to prefer latter. Indeed, the statement 'X is non-empty' translates more directly to $\neg(\neg \exists a, a \in X)$, which has an unnecessary double-negative and suggests a proof of inhabitation by contradiction. For this reason, we use the term *inhabited* in this book.

Emptiness may seem like a trivial condition—and it is—but owing to its canonicity, it arises all over the place.

Example 2.1.28
The set $\{x \in \mathbb{R} \mid x^2 = 2\}$ is inhabited since, for example $\sqrt{2} \in \mathbb{R}$ and $\sqrt{2}^2 = 2$.

However, the set $\{x \in \mathbb{Q} \mid x^2 = 2\}$ is empty since, if it were inhabited, then there would be a rational number x such that $x^2 = 2$, contrary to Proposition 0.29. ◁

Example 2.1.29

We observed in Example 2.1.10 that the set $[0]$ is empty; here's a more formal proof. Towards a contradiction, suppose $[0]$ is inhabited. Then there is some $k \in \mathbb{N}$ such that $1 \leqslant k \leqslant 0$. It follows that $1 \leqslant 0$, which contradicts the fact that $0 < 1$. Hence $[0]$ is empty, after all. ◁

Exercise 2.1.30

Let $a, b \in \mathbb{R}$. Prove that $[a, b]$ is empty if and only if $a > b$, and that (a, b) is empty if and only if $a \geqslant b$. ◁

The next exercise is a logical technicality, which is counterintuitive for the same reason that makes the principle of explosion (Axiom 1.1.49) difficult to grasp. However, it is extremely useful for proving facts about the empty set, as we will see soon in Theorem 2.1.32.

Exercise 2.1.31

Let E be an empty set and let $p(x)$ be a predicate with one free variable x with domain of discourse E. Show that the proposition $\forall x \in E, p(x)$ is true, and that the proposition $\exists x \in E, p(x)$ is false. What does the proposition $\forall x \in E, x \neq x$ mean in English? Is it true? ◁

Thanks to the axiom of extensionality (Axiom 2.1.22), any two empty sets must be equal since they both contain the same elements—namely, no elements at all! This is made formal in the following theorem.

Theorem 2.1.32

Let E and E' be sets. If E and E' are empty, then $E = E'$.

Proof. Suppose that E and E' are empty. The assertion that $E = E'$ is equivalent to

$$(\forall a \in E, a \in E') \wedge (\forall a \in E', a \in E)$$

But $\forall a \in E, a \in E'$ and $\forall a \in E', a \in E$ are both true by Exercise 2.1.31 since E and E' are empty. So $E = E'$, as claimed. □

Knowing that there is one and only one empty set means that we may now make the following definition, without worrying about whether the word 'the' is problematic.

Definition 2.1.33

The **empty set** (also known as the **null set**) is the set with no elements, and is denoted by $\varnothing$ (LaTeX code: `\varnothing`).

Some authors write $\{\}$ instead of $\varnothing$, since $\{\}$ is simply the empty set expressed in list notation.

Exercise 2.1.34

Let X be a set. Prove that $\varnothing \subseteq X$. ◁

Set operations

In Example 2.1.24 we noted that $[0, \infty)$ is the set of all non-negative real numbers. What if we wanted to talk about the set of all non-negative rational numbers instead? It would be nice if there was some expression in terms of $[0, \infty)$ and $\mathbb{Q}$ to denote this set.

This is where *set operations* come in—they allow us to use previously defined sets to introduce new sets.

Intersection ($\cap$)

The *intersection* of two sets is the set of things which are elements of both sets.

Definition 2.1.35

Let X and Y be sets. The (**pairwise**) **intersection** of X and Y, denoted $X \cap Y$ (LaTeX code: `\cap`), is defined by

$$X \cap Y = \{a \mid a \in X \land a \in Y\}$$

Example 2.1.36

By definition of intersection, we have $x \in [0, \infty) \cap \mathbb{Q}$ if and only if $x \in [0, \infty)$ and $x \in \mathbb{Q}$. Since $x \in [0, \infty)$ if and only if x is a non-negative real number (see Example 2.1.24), it follows that $[0, \infty) \cap \mathbb{Q}$ is the set of all non-negative rational numbers. ◁

Exercise 2.1.37

Prove that $[0, \infty) \cap \mathbb{Z} = \mathbb{N}$. ◁

Exercise 2.1.38

Write down the elements of the set

$$\{0, 1, 4, 7\} \cap \{1, 2, 3, 4, 5\}$$

◁

Exercise 2.1.39

Express $[-2,5) \cap [4,7)$ as a single interval. ◁

Proposition 2.1.40

Let X and Y be sets. Prove that $X \subseteq Y$ if and only if $X \cap Y = X$.

Proof

Suppose that $X \subseteq Y$. We prove $X \cap Y = X$ by double containment.

- ($\subseteq$) Suppose $a \in X \cap Y$. Then $a \in X$ and $a \in Y$ by definition of intersection, so in particular we have $a \in X$.

- ($\supseteq$) Suppose $a \in X$. Then $a \in Y$ since $X \subseteq Y$, so that $a \in X \cap Y$ by definition of intersection.

Conversely, suppose that $X \cap Y = X$. To prove that $X \subseteq Y$, let $a \in X$. Then $a \in X \cap Y$ since $X = X \cap Y$, so that $a \in Y$ by definition of intersection, as required. $\square$

Exercise 2.1.41

Let X be a set. Prove that $X \cap \varnothing = \varnothing$. ◁

Union ($\cup$)

The *union* of two sets is the set of things which are elements of at least one of the sets.

> **Definition 2.1.42**
> Let X and Y be sets. The (**pairwise**) **union** of X and Y, denoted $X \cup Y$ (LaTeX code: `\cup`), is defined by
> $$X \cup Y = \{a \mid a \in X \vee a \in Y\}$$

Example 2.1.43

Let E be the set of even integers and O be the set of odd integers. Since every integer is either even or odd, $E \cup O = \mathbb{Z}$. Note that $E \cap O = \varnothing$, thus $\{E, O\}$ is an example of a *partition* of $\mathbb{Z}$—see Definition 3.3.25. ◁

Exercise 2.1.44

Write down the elements of the set

$$\{0, 1, 4, 7\} \cup \{1, 2, 3, 4, 5\}$$

◁

Exercise 2.1.45

Express $[-2,5) \cup [4,7)$ as a single interval. ◁

The union operation allows us to define the following class of sets that will be particularly useful for us when studying counting principles in Section 3.3.

Exercise 2.1.46

Let X and Y be sets. Prove that $X \subseteq Y$ if and only if $X \cup Y = Y$. ◁

Example 2.1.47

Let X, Y, Z be sets. We prove that $X \cap (Y \cup Z) = (X \cap Y) \cup (X \cap Z)$.

- ($\subseteq$) Let $x \in X \cap (Y \cup Z)$. Then $x \in X$, and either $x \in Y$ or $x \in Z$. If $x \in Y$ then $x \in X \cap Y$, and if $x \in Z$ then $x \in X \cap Z$. In either case, we have $x \in (X \cap Y) \cup (X \cap Z)$.

- ($\supseteq$) Let $x \in (X \cap Y) \cup (X \cap Z)$. Then either $x \in X \cap Y$ or $x \in X \cap Z$. In both cases we have $x \in X$ by definition of intersection. In the first case we have $x \in Y$, and in the second case we have $x \in Z$; in either case, we have $x \in Y \cup Z$, so that $x \in X \cap (Y \cup Z)$.

◁

Exercise 2.1.48

Let X, Y, Z be sets. Prove that $X \cup (Y \cap Z) = (X \cup Y) \cap (X \cup Z)$. ◁

Indexed families of sets

We will often have occasion to take the intersection or union not of just two sets, but of an arbitrary collection of sets (even of infinitely many sets). For example, we might want to know which real numbers are elements of $[0, 1 + \frac{1}{n})$ for each $n \geqslant 1$, and which real numbers are elements of at least one of such sets.

Our task now is therefore to generalise our pairwise notions of intersection and union to arbitrary collections of sets, called *indexed families* of sets.

Definition 2.1.49
An (**indexed**) **family of sets** is a specification of a set X_i for each element i of some **indexing set** I. We write $\{X_i \mid i \in I\}$ for the indexed family of sets.

Example 2.1.50

The sets $[0, 1 + \frac{1}{n})$ mentioned above assemble into an indexed family of sets, whose indexing set is $\{n \in \mathbb{N} \mid n \geqslant 1\}$. We can abbreviate this family of sets by

$$\{[0, 1 + \tfrac{1}{n}) \mid n \geqslant 1\}$$

Observe that we have left implicit the fact that the variable n is ranging over the natural numbers and just written '$n \geqslant 1$' on the right of the vertical bar, rather than separately defining $I = \{n \in \mathbb{N} \mid n \geqslant 1\}$ and writing $\{[0, 1 + \frac{1}{n}) \mid n \in I\}$.   ◁

Definition 2.1.51

The (**indexed**) **intersection** of an indexed family $\{X_i \mid i \in I\}$ is defined by

$$\bigcap_{i \in I} X_i = \{a \mid \forall i \in I, a \in X_i\} \quad \text{(LATEX code: \bigcap_{i \in I})}$$

The (**indexed**) **union** of $\{X_i \mid i \in I\}$ is defined by

$$\bigcup_{i \in I} X_i = \{a \mid \exists i \in I, a \in X_i\} \quad \text{(LATEX code: \bigcup_{i \in I})}$$

Example 2.1.52

We prove that the intersection of the half-open intervals $[0, 1 + \frac{1}{n})$ for $n \geqslant 1$ is $[0, 1]$. We will use the notation $\bigcap_{n \geqslant 1}$ as shorthand for $\bigcap_{n \in \{x \in \mathbb{N} \mid x \geqslant 1\}}$.

- ($\subseteq$) Let $x \in \bigcap_{n \geqslant 1} [0, 1 + \frac{1}{n})$.

  Then $x \in [0, 1 + \frac{1}{n})$ for all $n \geqslant 1$. In particular, $x \geqslant 0$.

  To see that $x \leqslant 1$, assume that $x > 1$—we will derive a contradiction. Since $x > 1$, we have $x - 1 > 0$. Let $N \geqslant 1$ be some natural number greater or equal to $\frac{1}{x-1}$, so that $\frac{1}{N} \leqslant x - 1$. Then $x \geqslant 1 + \frac{1}{N}$, and hence $x \notin [0, 1 + \frac{1}{N})$, contradicting the assumption that $x \in [0, 1 + \frac{1}{n})$ for all $n \geqslant 1$.

  So we must have $x \leqslant 1$ after all, and hence $x \in [0, 1]$.

- ($\supseteq$) Let $x \in [0, 1]$.

  To prove that $x \in \bigcap_{n \geqslant 1} [0, 1 + \frac{1}{n})$, we need to show that $x \in [0, 1 + \frac{1}{n})$ for all $n \geqslant 1$. So fix $n \geqslant 1$. Since $x \in [0, 1]$, we have $x \geqslant 0$ and $x \leqslant 1 < 1 + \frac{1}{n}$, so that $x \in [0, 1 + \frac{1}{n})$, as required.

Hence $\bigcap_{n \geqslant 1} [0, 1 + \frac{1}{n}) = [0, 1]$ by double containment.   ◁

Exercise 2.1.53

Express $\bigcup_{n \geqslant 1} [0, 1 + \frac{1}{n})$ as an interval.   ◁

Exercise 2.1.54

Prove that $\bigcap_{n\in\mathbb{N}} [n] = \varnothing$ and $\bigcup_{n\in\mathbb{N}} [n] = \{k \in \mathbb{N} \mid k \geqslant 1\}$. ◁

Indexed intersections and unions generalise their pairwise counterparts, as the following exercise proves.

Exercise 2.1.55

Let X_1 and X_2 be sets. Prove that

$$X_1 \cap X_2 = \bigcap_{k\in[2]} X_k \quad \text{and} \quad X_1 \cup X_2 = \bigcup_{k\in[2]} X_k$$

◁

Exercise 2.1.56

Find a family of sets $\{X_n \mid n \in \mathbb{N}\}$ such that:

(i) $\bigcup_{n\in\mathbb{N}} X_n = \mathbb{N}$;

(ii) $\bigcap_{n\in\mathbb{N}} X_n = \varnothing$; and

(iii) $X_i \cap X_j \neq \varnothing$ for all $i, j \in \mathbb{N}$.

◁

Relative complement ($\setminus$)

Definition 2.1.57

Let X and Y be sets. The **relative complement** of Y in X, denoted $X \setminus Y$ (LaTeX code: `\setminus`), is defined by

$$X \setminus Y = \{x \in X \mid x \notin Y\}$$

Example 2.1.58

Let E be the set of all even integers. Then $n \in \mathbb{Z} \setminus E$ if and only if n is an integer and n is not an even integer; that is, if and only if n is odd. Thus $\mathbb{Z} \setminus E$ is the set of all odd integers.

Moreover, $n \in \mathbb{N} \setminus E$ if and only if n is a natural number and n is not an even integer. Since the even integers which are natural numbers are precisely the even natural numbers, $\mathbb{N} \setminus E$ is precisely the set of all odd natural numbers. ◁

Exercise 2.1.59

Write down the elements of the set

$$\{0,1,4,7\} \setminus \{1,2,3,4,5\}$$

◁

Exercise 2.1.60

Express $[-2,5) \setminus [4,7)$ and $[4,7) \setminus [-2,5)$ as intervals.   ◁

Exercise 2.1.61

Let X and Y be sets. Prove that $Y \setminus (Y \setminus X) = X \cap Y$, and deduce that $X \subseteq Y$ if and only if $Y \setminus (Y \setminus X) = X$.   ◁

Comparison with logical operators and quantifiers

The astute reader will have noticed some similarities between set operations and the logical operators and quantifiers that we saw in Chapter 1.

Indeed, this can be summarised in the following table. In each row, the expressions in both columns are equivalent, where p denotes '$a \in X$', q denotes '$a \in Y$', and $r(i)$ denotes '$a \in X_i$'.

sets	logic
$a \notin X$	$\neg p$
$a \in X \cap Y$	$p \wedge q$
$a \in X \cup Y$	$p \vee q$
$a \in \bigcap_{i \in I} X_i$	$\forall i \in I, r(i)$
$a \in \bigcup_{i \in I} X_i$	$\exists i \in I, r(i)$
$a \in X \setminus Y$	$p \wedge (\neg q)$

This translation between logic and set theory does not stop there; in fact, as the following theorem shows, De Morgan's laws for the logical operators (Theorem 1.3.24) and for quantifiers (Theorem 1.3.28) also carry over to the set operations of union and intersection.

Theorem 2.1.62 (De Morgan's laws for sets)
Given sets A, X, Y and a family $\{X_i \mid i \in I\}$, we have

(a) $A \setminus (X \cup Y) = (A \setminus X) \cap (A \setminus Y)$;

(b) $A \setminus (X \cap Y) = (A \setminus X) \cup (A \setminus Y)$;

(c) $A \setminus \bigcup_{i \in I} X_i = \bigcap_{i \in I} (A \setminus X_i)$;

(d) $A \setminus \bigcap_{i \in I} X_i = \bigcup_{i \in I} (A \setminus X_i)$.

Proof of (a)
Let a be arbitrary. By definition of union and relative complement, the assertion that $a \in A \setminus (X \cup Y)$ is equivalent to the logical formula

$$a \in A \wedge \neg (a \in X \vee a \in Y)$$

By de Morgan's laws for logical operators, this is equivalent to

$$a \in A \wedge (a \notin X \wedge a \notin Y)$$

which, in turn, is equivalent to

$$(a \in A \wedge a \notin X) \wedge (a \in A \wedge a \notin Y)$$

But then by definition of intersection and relative complement, this is equivalent to

$$a \in (A \setminus X) \cap (A \setminus Y)$$

Hence $A \setminus (X \cup Y) = (A \setminus X) \cap (A \setminus Y)$, as required. $\square$

Exercise 2.1.63
Complete the proof of de Morgan's laws for sets. $\triangleleft$

Power sets

Definition 2.1.64
Let X be a set. The **power set** of X, written $\mathscr{P}(X)$ (LaTeX code: `\mathcal{P}`), is the set of all subsets of X.

Example 2.1.65
There are four subsets of $\{1, 2\}$, namely

$$\varnothing, \quad \{1\}, \quad \{2\}, \quad \{1, 2\}$$

so $\mathscr{P}(X) = \{\varnothing, \{1\}, \{2\}, \{1, 2\}\}$. $\triangleleft$

Exercise 2.1.66
Write out the elements of $\mathscr{P}(\{1,2,3\})$.   ◁

Exercise 2.1.67
Let X be a set. Show that $\varnothing \in \mathscr{P}(X)$ and $X \in \mathscr{P}(X)$.   ◁

Exercise 2.1.68
Write out the elements of $\mathscr{P}(\varnothing)$, $\mathscr{P}(\mathscr{P}(\varnothing))$ and $\mathscr{P}(\mathscr{P}(\mathscr{P}(\varnothing)))$.   ◁

Power sets are often a point of confusion because they bring the property of being a *subset* of one set to that of being an *element* of another, in the sense that for all sets U and X we have

$$U \subseteq X \quad \Leftrightarrow \quad U \in \mathscr{P}(X)$$

This distinction looks easy to grasp, but when the sets U and X look alike, it's easy to fall into various traps. Here's a simple example.

Example 2.1.69
It is true that $\varnothing \subseteq \varnothing$, but false that $\varnothing \in \varnothing$. Indeed,

- $\varnothing \subseteq \varnothing$ means $\forall x \in \varnothing, x \in \varnothing$; but propositions of the form $\forall x \in \varnothing, p(x)$ are always true, as discussed in Exercise 2.1.31.

- The empty set has no elements; if $\varnothing \in \varnothing$ were true, it would mean that $\varnothing$ had an element (that element being $\varnothing$). So it must be the case that $\varnothing \notin \varnothing$.

◁

The following exercise is intended to help you overcome similar potential kinds of confusion by means of practice. Try to think precisely about what the definitions involved are.

Exercise 2.1.70
Determine, with proof, whether or not each of the following statements is true.

(a) $\mathscr{P}(\varnothing) \in \mathscr{P}(\mathscr{P}(\varnothing))$;

(b) $\varnothing \in \{\{\varnothing\}\}$;

(c) $\{\varnothing\} \in \{\{\varnothing\}\}$;

(d) $\mathscr{P}(\mathscr{P}(\varnothing)) \in \{\varnothing, \{\varnothing, \{\varnothing\}\}\}$.

Repeat the exercise with all instances of '$\in$' replaced by '$\subseteq$'.   ◁

Product ($\times$)

> **Definition 2.1.71**
> Let X and Y be sets. The (**pairwise**) **cartesian product** of X and Y is the set $X \times Y$ (LATEX code: `\times`) defined by
>
> $$X \times Y = \{(a,b) \mid x \in X \wedge y \in Y\}$$
>
> The elements $(a,b) \in X \times Y$ are called **ordered pairs**, whose defining property is that, for all $a, x \in X$ and all $b, y \in Y$, we have $(a,b) = (x,y)$ if and only if $a = x$ and $b = y$.

Example 2.1.72
If you have ever taken calculus, you will probably be familiar with the set $\mathbb{R} \times \mathbb{R}$.

$$\mathbb{R} \times \mathbb{R} = \{(x,y) \mid x, y \in \mathbb{R}\}$$

Formally, this is the set of ordered pairs of real numbers. Geometrically, if we interpret $\mathbb{R}$ as an infinite line, the set $\mathbb{R} \times \mathbb{R}$ is the (real) plane: an element $(x,y) \in \mathbb{R} \times \mathbb{R}$ describes the point in the plane with coordinates (x,y).

We can investigate this further. For example, the following set:

$$\mathbb{R} \times \{0\} = \{(x,0) \mid x \in \mathbb{R}\}$$

is precisely the x-axis. We can describe graphs as subsets of $\mathbb{R} \times \mathbb{R}$. Indeed, the graph of $y = x^2$ is given by

$$G = \{(x,y) \in \mathbb{R} \times \mathbb{R} \mid y = x^2\} = \{(x,x^2) \mid x \in \mathbb{R}\} \subseteq \mathbb{R} \times \mathbb{R}$$

$\triangleleft$

Exercise 2.1.73
Write down the elements of the set $\{1,2\} \times \{3,4,5\}$. $\quad\triangleleft$

Exercise 2.1.74
Let X be a set. Prove that $X \times \varnothing = \varnothing$. $\quad\triangleleft$

Exercise 2.1.75
Let X, Y and Z be sets. Under what conditions is it true that $X \times Y = Y \times X$? Under what conditions is it true that $(X \times Y) \times Z = X \times (Y \times Z)$? $\quad\triangleleft$

We might have occasion to take cartesian products of more than two sets. For example, whatever the set $\mathbb{R} \times \mathbb{R} \times \mathbb{R}$ is, its elements *should* be ordered triples (a,b,c) consisting of elements $a, b, c \in \mathbb{R}$. This is where the following definition comes in handy.

Definition 2.1.76

Let $n \in \mathbb{N}$ and let $X_1, X_2, \ldots, X_n$ be sets. The (**n-fold**) **cartesian product** of $X_1, X_2, \ldots, X_n$ is the set $\prod_{k=1}^{n} X_k$ (LaTeX code: `\prod_{k=1}^{n}`) defined by

$$\prod_{k=1}^{n} X_k = \{(a_1, a_2, \ldots, a_n) \mid a_k \in X_k \text{ for all } 1 \leqslant k \leqslant n\}$$

The elements $(a_1, a_2, \ldots, a_n) \in \prod_{k=1}^{n} X_k$ are called **ordered k-tuples**, whose defining property is that, for all $1 \leqslant k \leqslant n$ and all $a_k, b_k \in X_k$, we have $(a_1, a_2, \ldots, a_n) = (b_1, b_2, \ldots, b_n)$ if and only if $a_k = b_k$ for all $1 \leqslant k \leqslant n$.

Given a set X, write X^n to denote the set $\prod_{k=1}^{n} X$. We might on occasion also write

$$X_1 \times X_2 \times \cdots \times X_n = \prod_{k=1}^{n} X_k$$

Example 2.1.77

In Exercise 2.1.75 you might have noticed that the sets $(X \times Y) \times Z$ and $X \times (Y \times Z)$ are not always equal—Definition 2.1.76 introduces a *third* potentially non-equal cartesian product of X, Y and Z. For example, consider when $X = Y = Z = \mathbb{R}$. Then

- The elements of $(\mathbb{R} \times \mathbb{R}) \times \mathbb{R}$ are ordered pairs $((a,b),c)$, where (a,b) is itself an ordered pair of real numbers and c is a real number.

- The elements of $\mathbb{R} \times (\mathbb{R} \times \mathbb{R})$ are ordered pairs $(a,(b,c))$, where a is a real number and (b,c) is an ordered pair of real numbers.

- The elements of $\mathbb{R} \times \mathbb{R} \times \mathbb{R}$ ($= \mathbb{R}^3$) are ordered triples (a,b,c), where a, b and c are real numbers.

So, although these three sets *appear* to be the same, zooming in closely on the definitions reveals that there are subtle differences between them. A sense in which they are the same is that there are *bijections* between them—the notion of a bijection will be introduced in Section 2.3. ◁

Section 2.2
Functions

One way of studying interactions between sets is by studying *functions* between them, which we will define informally in Definition 2.2.1. Functions are mathematical objects which assign to each element of one set exactly one element of another. Almost every branch of mathematics studies functions, be it directly or indirectly, and almost every application of mathematics arises from a translation of the abstract notion of a function to the real world. Just one example of this is the theory of computation—functions provide precisely the language necessary to describe the deterministic input-output behaviour of algorithms.

You might have come across the notion of a function before now. In schools, functions are often introduced as being like *machines*—they have inputs and outputs, and on a given input they always return the same output. For instance, there is a function which takes integers as inputs and gives integers as outputs, which on the input x returns the integer $x+3$.

This characterisation of functions, however, is clearly not precise enough for the purposes of mathematical proof. A next approximation to a precise definition of a function might look something like this:

Definition 2.2.1

A **function** f from a set X to a set Y is a specification of elements $f(x) \in Y$ for $x \in X$, such that
$$\forall x \in X, \exists! y \in Y, y = f(x)$$
Given $x \in X$, the (unique!) element $f(x) \in Y$ is called the **value** of f at x.

The set X is called the **domain** (or **source**) of f, and Y is called the **codomain** (or **target**) of f. We write $f : X \to Y$ (LaTeX code: `f : X \to Y`) to denote the assertion that f is a function with domain X and codomain Y.

This is better—we're now talking about sets (and not mysterious 'machines'), which we have explored in Section 2.1.

Moreover, this definition establishes a close relationship between functions and the $\exists!$ quantifier: indeed, to say that f assigns to each element of X a unique element of Y is to say precisely that

$$\forall x \in X, \exists! y \in Y, y = f(x)$$

Conversely, any true proposition of the form $\forall x \in X, \exists! y \in Y, p(x,y)$ defines a function $f : X \to Y$: the function f assigns to each $x \in X$ the unique $y \in Y$ such that $p(x,y)$ is true. In other words, $\forall x \in X, p(x, f(x))$ is true!

We can use this to generate some examples of functions.

Example 2.2.2

Example 1.2.27 said that every positive real number has a unique positive square root; we proved this in Example 1.2.30. What this means is that there is a function

$$r : \mathbb{R}^{>0} \to \mathbb{R}^{>0} \qquad \text{where } \mathbb{R}^{>0} = \{x \in \mathbb{R} \mid x > 0\}$$

defined by letting $r(x)$ be the (unique) positive square root of x, for each $x \in \mathbb{R}^{>0}$. That is, we have a function r defined by $r(x) = \sqrt{x}$. ◁

Exercise 2.2.3

Recall Exercise 1.2.31. Which of the statements (a), (b) or (c) is of the form $\forall x \in X, \exists! y \in Y, p(x,y)$? For each statement of this form, determine the domain and codomain of the corresponding function, and write an expression defining this function. ◁

Specifying a function

Just like with sets, there are many ways to specify a function $f : X \to Y$, but when we do so, we must be careful that what we write really *does* define a function!

This correctness of specification is known as *well-definedness*, and ultimately amounts to verifying that the condition $\forall x \in X, \exists! y \in Y, f(x) = y$ holds for the specification of f. Namely *totality*, *existence* and *uniqueness*:

- **Totality.** A value $f(x)$ should be specified for each $x \in X$—this corresponds to the '$\forall x \in X$' quantifier in the definition of functions.

- **Existence.** For each $x \in X$, the specified value $f(x)$ should actually exist, and should be an element of Y—this corresponds to the *existence* part of the '$\exists! y \in Y$' quantifier in the definition of functions.

- **Uniqueness.** For each $x \in X$, the specified value $f(x)$ should refer to only one element of Y—this corresponds to the *uniqueness* part of the '$\exists! y \in Y$' quantifier in the definition of functions.

When specifying a function, you should justify each of these components of well-definedness unless they are extremely obvious. You will probably find that, in most cases, the only component in need of justification is uniqueness, but keep all three in mind.

Lists. If X is finite, then we can specify a function $f : X \to Y$ by simply listing the values of f at all possible elements $x \in X$. For example, we can define a function

$$f : \{1,2,3\} \to \{\text{red}, \text{yellow}, \text{green}, \text{blue}, \text{purple}\}$$

by declaring

$$f(1) = \text{red}, \quad f(2) = \text{purple}, \quad f(3) = \text{green}$$

Note that the function is at this point completely specified: we know its values at all elements of the domain $\{1,2,3\}$. It doesn't matter that some of the elements of the codomain (yellow and blue) are unaccounted for—all that matters is that each element of the domain is associated with exactly one element of the codomain.

Unfortunately, most of the sets that we work with will be infinite, or of an unspecified finite size; in these cases, simply writing a list of values isn't sufficient. Fortunately for us, there are other ways of specifying functions.

Formulae. In many cases, particularly when the domain X and codomain Y are number sets, we can define a function by giving a formula for the value of $f(x)$ for each $x \in X$. For example, we can define a function $f : \mathbb{R} \to \mathbb{R}$ by letting

$$f(x) = x^2 + 3 \text{ for all } x \in \mathbb{R}$$

By cases. It will at times be convenient to define a function using different specifications for different elements of the domain. A very simple example is the *absolute value function* $|-| : \mathbb{R} \to \mathbb{R}$, defined for $x \in \mathbb{R}$

$$|x| = \begin{cases} x & \text{if } x \geqslant 0 \\ -x & \text{if } x \leqslant 0 \end{cases}$$

Here we have split into two cases based on the conditions $x \geqslant 0$ and $x \leqslant 0$.

When specifying a function $f : X \to Y$ by cases, it is important that the conditions be:

- **exhaustive**: given $x \in X$, at least one of the conditions on X must hold; and

- **compatible**: if any $x \in X$ satisfies more than one condition, the specified value must be the same no matter which condition is picked.

For the absolute value function defined above, these conditions are satisfied. Indeed, for $x \in \mathbb{R}$, it is certainly the case that $x \geqslant 0$ or $x \leqslant 0$, so the conditions are exhaustive. Moreover, given $x \in \mathbb{R}$, if both $x \geqslant 0$ and $x \leqslant 0$, then $x = 0$—so we need to check that the specification yields the same value when $x = 0$ regardless of which condition we pick. The $x \geqslant 0$ condition yields the value 0, and the $x \leqslant 0$

condition yields the value -0, which is equal to 0—so the conditions are compatible. We could have used $x < 0$ instead of $x \leqslant 0$; in this case the conditions are *mutually exclusive*, so certainly compatible because they do not overlap.

Algorithms. You might, on first exposure to functions, have been taught to think of a function as a *machine* which, when given an *input*, produces an *output*. This 'machine' is defined by saying what the possible inputs and outputs are, and then providing a list of instructions (an *algorithm*) for the machine to follow, which on any input produces an output—and, moreover, if fed the same input, the machine always produces the same output.

For example, we might instruct a machine to take rational numbers as inputs and give rational numbers as outputs, and to follow the following sequence of steps on a given input

$$\text{multiply by } 2 \rightarrow \text{add } 5 \rightarrow \text{square the result} \rightarrow \text{divide by } 6$$

This 'machine' defines a function $M : \mathbb{Q} \rightarrow \mathbb{Q}$ which, in equation form, is specified by

$$M(x) = \frac{(2x+5)^2}{6} \text{ for all } x \in \mathbb{Q}$$

In our more formal set-up, therefore, we can define a function $M : I \rightarrow O$ by specifying:

- a set I of all **inputs**;

- a set O of potential **outputs**; and

- a deterministic[a] algorithm which describes how an input $x \in I$ is transformed into an output $M(x) \in O$.

That is, the domain is the set I of all possible 'inputs', the codomain is a set O containing all the possible 'outputs', and the function M is a rule specifying how an input is associated with the corresponding output.

For now, we will use algorithmic specifications of functions only sparingly—this is because it is much harder to make formal what is meant by an 'algorithm', and it is important to check that a given algorithm is deterministic.

---

[a]The word 'deterministic' just means that the algorithm always produces the same output on a single input.

Function equality

In Section 2.1 we discussed how there may be many different possible ways of characterising equality of sets. This matter was resolved by declaring that two sets are equal if and only if they have the same elements (this was Axiom 2.1.22).

A similar matter arises for functions. For example, consider the function $f : \mathbb{R} \to \mathbb{R}$ defined by $f(x) = 2x$ for all $x \in \mathbb{R}$, and the function $g : \mathbb{R} \to \mathbb{R}$, defined by letting $g(x)$ be the result of taking x, multiplying it by three, dividing the result by four, dividing the result by six, and then multiplying the result by sixteen. It so happens that $g(x) = 2x$ for all $x \in \mathbb{R}$ as well, but that is not how g is defined; moreover, if f and g were implemented as algorithms, then it would take longer to compute the values of g than it would take to compute the values of f.

Should we consider f and g to be *equal*? If we are only interested in whether f and g have the same values on each argument, then the answer should be 'yes'; if we are interested in the algorithmic behaviour of f and g, then the answer should be 'no'.

We resolve this dilemma with the following axiom. By adopting this axiom, we are stating that the functions f and g discussed above are equal.

Axiom 2.2.4 (Function extensionality)
Let $f : X \to Y$ and $g : A \to B$ be functions. Then $f = g$ if and only if f and g have the same domain and codomain, and $f(x) = g(x)$ for all $x \in X$.

Strategy 2.2.5 (Proving two functions are equal)
Given functions $f, g : X \to Y$ with the same domain and codomain, in order to prove that $f = g$, it suffices to prove that $f(x) = g(x)$ for all $x \in X$. ◁

A consequence of Axiom 2.2.4 is that, for fixed sets X and Y, a function $X \to Y$ is uniquely determined by its input-output pairs. This set is called the *graph* of the function; the proof of the equivalence between functions and their graphs is the content of Theorem 2.2.9.

Definition 2.2.6
Let $f : X \to Y$ be a function. The **graph** of f is the subset $\mathrm{Gr}(f) \subseteq X \times Y$ (LaTeX code: `\mathrm{Gr}`) defined by

$$\mathrm{Gr}(f) = \{(x, f(x)) \mid x \in X\} = \{(x, y) \in X \times Y \mid y = f(x)\}$$

Example 2.2.7

Given a (sufficiently well-behaved) function $f : \mathbb{R} \to \mathbb{R}$, we can represent $\mathrm{Gr}(f) \subseteq \mathbb{R} \times \mathbb{R}$ by plotting it on a pair of axes using Cartesian coordinates in the usual way. For example, if f is defined by $f(x) = \frac{x}{2}$ for all $x \in \mathbb{R}$, then its graph

$$\mathrm{Gr}(f) = \left\{ \left(x, \frac{x}{2} \right) \,\middle|\, x \in \mathbb{R} \right\}$$

can be represented by graph plot in Figure 2.1.


Figure 2.1: Graph of the function $f : \mathbb{R} \to \mathbb{R}$ defined by $f(x) = \frac{x}{2}$ for all $x \in \mathbb{R}$

◁

Exercise 2.2.8

Find a function $f : \mathbb{Z} \to \mathbb{Z}$ whose graph is equal to the set

$$\{\ldots, (-2, -5), (-1, -2), (0, 1), (1, 4), (2, 7), (3, 10), \ldots\}$$

◁

Theorem 2.2.9 below provides a way of verifying that a function is well-defined by characterising their graphs.

Theorem 2.2.9

Let X and Y be sets. A subset $G \subseteq X \times Y$ is the graph of a function if and only if

$$\forall x \in X, \, \exists! y \in Y, \, (x, y) \in G$$

Proof

($\Rightarrow$). Suppose $G \subseteq X \times Y$ is the graph of a function, say $G = \mathrm{Gr}(f)$ for some

$f : X \to Y$. Then for each $x \in X$, it follows from well-definedness of f that $f(x)$ is the unique element $y \in Y$ for which $(x, y) \in G$. That is, $(x, f(x)) \in G$, and if $y \in Y$ with $(x, y) \in G$, then $y = f(x)$.

($\Leftarrow$). Suppose $G \subseteq X \times Y$ satisfies $\forall x \in X, \exists! y \in Y, (x, y) \in G$. Define a function $f : X \to Y$ by, for each $x \in X$, defining the value $f(x)$ to be the unique element $y \in Y$ for which $(x, y) \in G$. Well-definedness of f is then immediate from our assumption of the existence and uniqueness of such a value of y for each $x \in X$. $\square$

Example 2.2.10
The set G defined by

$$G = \{(1, \text{red}), (2, \text{red}), (3, \text{green})\}$$

is the graph of a function $f : \{1, 2, 3\} \to \{\text{red}, \text{green}, \text{blue}\}$. The function f is defined by

$$f(1) = \text{red}, \quad f(2) = \text{red}, \quad f(3) = \text{green}$$

However, G is *not* the graph of a function $\{1, 2, 3, 4\} \to \{\text{red}, \text{green}, \text{blue}\}$, since G contains no elements of the form $(4, y)$ for $y \in \{\text{red}, \text{green}, \text{blue}\}$. Moreover, the set G' defined by

$$G' = \{(1, \text{red}), (2, \text{red}), (2, \text{blue}), (3, \text{green})\}$$

does not define the graph of a function $\{1, 2, 3\} \to \{\text{red}, \text{green}, \text{blue}\}$, since there is not a *unique* element of the form $(2, y)$ in G'—rather, there are two of them! $\triangleleft$

Exercise 2.2.11
For each of the following specifications of sets X, Y, G, determine whether or not G is the graph of a function from X to Y.

(a) $X = \mathbb{R}, Y = \mathbb{R}, G = \{(a, a^2) \mid a \in \mathbb{R}\}$;

(b) $X = \mathbb{R}, Y = \mathbb{R}, G = \{(a^2, a) \mid a \in \mathbb{R}\}$;

(c) $X = \mathbb{R}^{\geqslant 0}, Y = \mathbb{R}^{\geqslant 0}, G = \{(a^2, a) \mid a \in \mathbb{R}^{\geqslant 0}\}$, where $\mathbb{R}^{\geqslant 0} = \{x \in \mathbb{R} \mid x \geqslant 0\}$;

(d) $X = \mathbb{Q}, Y = \mathbb{Q}, G = \{(x, y) \in \mathbb{Q} \times \mathbb{Q} \mid xy = 1\}$.

(e) $X = \mathbb{Q}, Y = \mathbb{Q}, G = \{(a, a) \mid a \in \mathbb{Z}\}$;

$\triangleleft$

Aside
In light of Theorem 2.2.9, some people choose to define functions $X \to Y$ as particular subsets of $X \times Y$—that is, they identify functions with their graphs. This is particularly useful when studying the logical foundations of mathematics.

We avoid this practice here, because it is not conceptually necessary, and it would preclude other possible ways of encoding functions.   ◁

We will now look at some more examples (and non-examples) of functions.

Example 2.2.12

Example 1.2.27 gives a prime example of a function: it says that for every positive real number a there is a unique positive real number b such that $b^2 = a$. This unique b is precisely the positive square root $\sqrt{a}$ of a. Writing $\mathbb{R}^{>0}$ for the set of positive real numbers, we have thus established that taking the positive square root defines a function $\mathbb{R}^{>0} \to \mathbb{R}^{>0}$.   ◁

There is a class of functions called *identity functions* that, despite being very simple, are so important that we will give them a numbered definition!

Definition 2.2.13

Let X be a set. The **identity function** on X is the function $\mathrm{id}_X : X \to X$ (LaTeX code: `\mathrm{id}_X`) defined by $\mathrm{id}_X(x) = x$ for all $x \in X$.

You should convince yourself that the specification of id_X given in Definition 2.2.13 is well-defined.

Another interesting example of a function is the *empty function*, which is useful in coming up with counterexamples and proving combinatorial identities (see Section 3.3).

Definition 2.2.14

Let X be a set. The **empty function** with codomain X is the (unique!) function $\varnothing \to X$. It has no values, since there are no elements of its domain.

Again, you should convince yourself that this specification is well-defined. Conceptually, convincing yourself of this is not easy; but writing down the proof of well-definedness is extremely easy—you will find that there is simply nothing to prove!

Example 2.2.15

Define $f : \mathbb{R} \to \mathbb{R}$ by the equation $f(x)^2 = x$ for all $x \in \mathbb{R}$. This is not well-defined for a few reasons. First, if $x < 0$ then there is no real number y such that $y^2 = x$, so for $x < 0$ there are no possible values of $f(x)$ in the codomain of f, so *existence* fails. Second, if $x > 0$ then there are in fact *two* real numbers y such that $y^2 = x$, namely the positive square root $\sqrt{x}$ and the negative square root

$-\sqrt{x}$. The specification of f does not indicate which of these values to take, so *uniqueness* fails.

Notice that the function $r : \mathbb{R}^{>0} \to \mathbb{R}^{>0}$ from Example 2.2.2 *is* (well-)defined by the equation $r(x)^2 = x$ for all $x \in \mathbb{R}^{>0}$. This illustrates why it is very important to specify the domain and codomain when defining a function. ◁

Exercise 2.2.16
Which of the following specifications of functions are well-defined?

(a) $g : \mathbb{Q} \to \mathbb{Q}$ defined by the equation $(x+1)g(x) = 1$ for all $x \in \mathbb{Q}$;

(b) $h : \mathbb{N} \to \mathbb{Q}$ defined by $(x+1)h(x) = 1$ for all $x \in \mathbb{N}$;

(c) $k : \mathbb{N} \to \mathbb{N}$ defined by $(x+1)k(x) = 1$ for all $x \in \mathbb{N}$;

(d) $\ell : \mathbb{N} \to \mathbb{N}$ defined by $\ell(x) = \ell(x)$ for all $x \in \mathbb{N}$.

◁

Exercise 2.2.17
Find a condition on sets X and Y such that the specification of a function $i : X \cup Y \to \{0,1\}$ given by

$$i(z) = \begin{cases} 0 & \text{if } z \in X \\ 1 & \text{if } z \in Y \end{cases}$$

to be well-defined. ◁

Composition of functions

In our section on sets, we talked about various operations that can be performed on sets—union, intersection, and so on. There are also operations on functions, by far the most important of which is *composition*. To understand how composition works, let's revisit the algorithmically defined function $M : \mathbb{Q} \to \mathbb{Q}$ from page 100:

$$\text{multiply by } 2 \to \text{add } 5 \to \text{square the result} \to \text{divide by } 6$$

The function M is, in some sense, a *sequence* of functions, performed one-by-one until the desired result is reached. This is precisely *composition of functions*.

> **Definition 2.2.18**
> Given functions $f : X \to Y$ and $g : Y \to Z$, their **composite** $g \circ f$ (LaTeX code:
> `g \circ f`) (read 'g composed with f' or 'g after f' or even just '$g\,f$') is the
> function $g \circ f : X \to Z$ defined by
>
> $$(g \circ f)(x) = g(f(x)) \text{ for all } x \in X$$

Intuitively, $g \circ f$ is the function resulting from first applying f, and then applying
g, to the given input.

Common error
Function composition is in some sense written 'backwards': in the expression
$g \circ f$, the function which is applied *first* is written *last*—there is a good reason
for this: the argument to the function is written after the function! However, this
mis-match often trips students up on their first exposure to function composition,
so be careful!   ◁

Example 2.2.19
The function M from page 100 can be defined as the composite

$$M = ((k \circ h) \circ g) \circ f$$

where

- $f : \mathbb{Q} \to \mathbb{Q}$ is defined by $f(x) = 2x$ for all $x \in \mathbb{Q}$;

- $g : \mathbb{Q} \to \mathbb{Q}$ is defined by $g(x) = x + 5$ for all $x \in \mathbb{Q}$;

- $h : \mathbb{Q} \to \mathbb{Q}$ is defined by $h(x) = x^2$ for all $x \in \mathbb{Q}$;

- $k : \mathbb{Q} \to \mathbb{Q}$ is defined by $k(x) = \frac{x}{6}$ for all $x \in \mathbb{Q}$.

                                                                                      ◁

Exercise 2.2.20
Let $f, g, h, k : \mathbb{Q} \to \mathbb{Q}$ be as in Example 2.2.19. Compute equations defining the
following composites:

(a) $f \circ g$;

(b) $g \circ f$;

(c) $((f \circ g) \circ h) \circ k$;

(d) $f \circ (g \circ (h \circ k))$;

(e) $(g \circ g) \circ (g \circ g)$.

◁

Example 2.2.21
Let $f : X \to Y$ be any function. Then

$$\mathrm{id}_Y \circ f = f = f \circ \mathrm{id}_X$$

To see this, let $x \in X$. Then

$$
\begin{aligned}
(\mathrm{id}_Y \circ f)(x) &= \mathrm{id}_Y(f(x)) &&\text{by definition of composition}\\
&= f(x) &&\text{by definition of } \mathrm{id}_Y\\
&= f(\mathrm{id}_X(x)) &&\text{by definition of } \mathrm{id}_X\\
&= (f \circ \mathrm{id}_X)(x) &&\text{by definition of composition}
\end{aligned}
$$

Equality of the three functions in question follows. ◁

Exercise 2.2.22
Prove that composition of functions is *associative*, that is, if $f : X \to Y$, $g : Y \to Z$ and $h : Z \to W$ are functions, then

$$h \circ (g \circ f) = (h \circ g) \circ f : X \to W$$

As a consequence of associativity, when we want to compose more than two functions, it doesn't matter what order we compose the functions in. As such, we can just write $h \circ g \circ f$. ◁

Exercise 2.2.23
Let $f : X \to Y$ and $g : Z \to W$ be functions, and suppose that $Y \subsetneq Z$. Note that there is a function $h : X \to W$ defined by $h(x) = g(f(x))$ for all $x \in X$. Write h as a composite of functions involving f and g. ◁

Images and preimages

Definition 2.2.24
Let $f : X \to Y$ be a function and let $U \subseteq X$. The **image of U under** f is the subset $f[U] \subseteq Y$ (also written $f_*(U)$ (LaTeX code: `f_*`) or even just $f(U)$) is defined by

$$f[U] = \{ f(x) \mid x \in U \} = \{ y \in Y \mid \exists x \in U, y = f(x) \}$$

That is, $f[U]$ is the set of values that the function f takes when given inputs from U.

The **image of** f is the image of the entire domain, i.e. the set $f[X]$.

Example 2.2.25
Let $f : \mathbb{R} \to \mathbb{R}$ be defined by $f(x) = x^2$. The image of f is the set $\mathbb{R}^{\geq 0}$ of all nonnegative real numbers. Let's prove this:

- $(f[\mathbb{R}] \subseteq \mathbb{R}^{\geqslant 0})$. Let $y \in f[\mathbb{R}]$. Then $y = x^2$ for some $x \in \mathbb{R}$. But $x^2 \geqslant 0$, so we must have $y \in \mathbb{R}^{\geqslant 0}$, as required.

- $(\mathbb{R}^{\geqslant 0} \subseteq f[\mathbb{R}])$. Let $y \in \mathbb{R}^{\geqslant 0}$. Then $\sqrt{y} \in \mathbb{R}$, and $y = (\sqrt{y})^2 = f(\sqrt{y})$. Hence $y \in f[\mathbb{R}]$, as required.

We have shown by double containment that $f[\mathbb{R}] = \mathbb{R}^{\geqslant 0}$. ◁

Exercise 2.2.26

For each of the following functions f and subsets U of their domain, describe the image $f[U]$.

(a) $f : \mathbb{Z} \to \mathbb{Z}$ defined by $f(n) = 3n$, with $U = \mathbb{N}$;

(b) $f : X \to X \times X$ (where X is any set) defined by $f(x) = (x, x)$ with $U = X$;

(c) $f : \{a, b, c\} \to \{1, 2, 3\}$ defined by $f(a) = 1$, $f(b) = 3$ and $f(c) = 1$, with $U = \{a, b, c\}$.

◁

Exercise 2.2.27

Prove that $f[\varnothing] = \varnothing$ for all functions f. ◁

Example 2.2.28

Let $f : X \to Y$ be a function and let $U, V \subseteq X$. Then $f[U \cap V] \subseteq f[U] \cap f[V]$. To see this, let $y \in f[U \cap V]$. Then $y = f(x)$ for some $x \in U \cap V$. By definition of intersection, $x \in U$ and $x \in V$. Since $x \in U$ and $y = f(x)$, we have $y \in f[U]$; likewise, since $x \in V$, we have $y \in f[V]$. But then by definition of intersection, we have $y \in f[U] \cap f[V]$. ◁

Exercise 2.2.29

Let $f : X \to Y$ be a function and let $U, V \subseteq X$. We saw in Example 2.2.28 that $f[U \cap V] \subseteq f[U] \cap f[V]$. Determine which of the following is true, and for each, provide a proof of its truth or falsity:

(a) $f[U] \cap f[V] \subseteq f[U \cap V]$;

(b) $f[U \cup V] \subseteq f[U] \cup f[V]$;

(c) $f[U] \cup f[V] \subseteq f[U \cup V]$.

◁

> **Definition 2.2.30**
>
> Let $f : X \to Y$ be a function and let $V \subseteq Y$. The **preimage of V under** f is the subset $f^{-1}[V]$ (LATEX code: `f^{-1}`) (also written $f^*(V)$ (LATEX code: `f^*`), or just $f^{-1}(V)$) is defined by
>
> $$f^{-1}[V] = \{x \in X \mid f(x) \in V\} = \{x \in X \mid \exists y \in V, f(x) = y\}$$
>
> That is, $f^{-1}[V]$ is the set of all the elements of its domain X that the function f sends to elements of V.

Example 2.2.31

Let $f : \mathbb{Z} \to \mathbb{Z}$ be the function defined by $f(x) = x^2$ for all $x \in X$. Then

- $f^{-1}[\{1, 4, 9\}] = \{-3, -2, -1, 1, 2, 3\}$;

- $f^{-1}[\{1, 2, 3, 4, 5, 6, 7, 8, 9\}] = \{-3, -2, -1, 1, 2, 3\}$ too, since the other elements of $[9]$ are not perfect squares, and hence not of the form $f(x)$ for $x \in \mathbb{Z}$;

- $f^{-1}[\mathbb{N}] = \mathbb{Z}$, since for any $x \in \mathbb{Z}$ we have $f(x) \geqslant 0$, so that $f(x) \in \mathbb{N}$.

$\triangleleft$

Example 2.2.32

Let $f : X \to Y$ be a function, let $U \subseteq X$ and let $V \subseteq Y$. Then $f[U] \subseteq V$ if and only if $U \subseteq f^{-1}[V]$. The proof is as follows.

($\Rightarrow$). Suppose $f[U] \subseteq V$; we'll prove $U \subseteq f^{-1}[V]$. So fix $x \in U$. Then $f(x) \in f[U]$ by definition of image. But then $f(x) \in V$ by our assumption that $f[U] \subseteq V$, and so $x \in f^{-1}[V]$ by definition of preimage. Since x was arbitrarily chosen from U, it follows that $U \subseteq f^{-1}[V]$.

($\Leftarrow$). Suppose $U \subseteq f^{-1}[V]$; we'll prove $f[U] \subseteq V$. So fix $y \in f[U]$. Then $y = f(x)$ for some $x \in U$ by definition of image. But then $x \in f^{-1}[V]$ by our assumption that $U \subseteq f^{-1}[V]$, and so $f(x) \in V$ by definition of preimage. But $y = f(x)$, so $y \in V$, and since y was arbitrarily chosen, it follows that $f[U] \subseteq V$. $\triangleleft$

The following exercise demonstrates that preimages interact very nicely with the basic set operations (intersection, union and relative complement):

Exercise 2.2.33

Let $f : X \to Y$ be a function and let $U, V \subseteq Y$. Prove that:

(a) $f^{-1}[U \cap V] = f^{-1}[U] \cap f^{-1}[V]$;

(b) $f^{-1}[U \cup V] = f^{-1}[U] \cup f^{-1}[V]$; and

(c) $f^{-1}[Y \setminus U] = X \setminus f^{-1}[U]$.

◁

Exercise 2.2.34

Let $f : X \to Y$ be a function. Prove that $f^{-1}[\varnothing] = \varnothing$ and $f^{-1}[Y] = X$.   ◁

Exercise 2.2.35

Let $f : X \to Y$ be a function. Provide a proof of the truth or falsity of each of the following statements:

(a) $U \subseteq f^{-1}[f[U]]$ for all $U \subseteq X$;

(b) $f^{-1}[f[U]] \subseteq U$ for all $U \subseteq X$;

(c) $V \subseteq f[f^{-1}[V]]$ for all $V \subseteq Y$;

(d) $f[f^{-1}[V]] \subseteq V$ for all $V \subseteq Y$.

◁

Injections and surjections

To motivate some of the definitions to come, look at the dots (•) and stars (⋆) below. Are there more dots or more stars?

Pause for a second and think about how you knew the answer to this question.

Indeed, there are more dots than stars. There are a couple of ways to arrive at this conclusion:

(i) You could count the number of dots, count the number of stars, and then compare the two numbers; or

(ii) You could notice that the dots and the stars are evenly spaced, but that the line of dots is longer than the line of stars.

It is likely that you chose method (ii). In fact, it is likely that you haven't even counted the number of dots or the number of stars yet—and you don't need to! We can conclude that there are more dots than stars by simply pairing up dots with stars—we eventually run out of stars, and there are still dots left over, so there must have been more dots than stars.

Injectivity

One way of formalising this act of pairing up stars with dots mathematically is to define a function $f : S \to D$ from the set S of stars to the set D of dots, where the value of f at each star is the dot that it is paired with. We of course must do this in such a way that each dot is paired with at most one star:

It is a property of this function—called *injectivity*—that allows us to deduce that there are more dots than stars.

Intuitively, a function $f : X \to Y$ is injective if it puts the elements of X in one-to-one correspondence with the elements of a subset of Y—just like how the stars are in one-to-one correspondence with a subset of the dots in the example above.

Definition 2.3.1
A function $f : X \to Y$ is **injective** (or **one-to-one**) if

$$\forall a, b \in X, \, f(a) = f(b) \Rightarrow a = b$$

An injective function is said to be an **injection**.

Strategy 2.3.2 (Proving a function is injective)
In order to prove that a function $f : X \to Y$ is injective, it suffices to fix $a, b \in X$, assume that $f(a) = f(b)$, and then derive $a = b$.   ◁

By contraposition, $f : X \to Y$ being injective is equivalent to saying, for all $a, b \in X$, if $a \neq b$, then $f(a) \neq f(b)$. This is usually less useful for *proving* that a function is injective, but it does provide a good intuition—it says that f sends distinct inputs to distinct outputs.

The following is a very simple example from elementary arithmetic:

Example 2.3.3
Define $f : \mathbb{Z} \to \mathbb{Z}$ by letting $f(x) = 2n + 1$ for all $n \in \mathbb{Z}$. We'll prove that f is injective. Fix $m, n \in \mathbb{Z}$, and assume that $f(m) = f(n)$. By definition of f, we have $2m + 1 = 2n + 1$. Subtracting 1 yields $2m = 2n$, and dividing by 2 yields $m = n$. Hence f is injective.   ◁

The following example is slightly more sophisticated.

Proposition 2.3.4
Let $f : X \to Y$ and $g : Y \to Z$ be functions. If f and g are injective, then $g \circ f$ is injective.

Proof
Suppose that f and g are injective and let $a, b \in X$. We need to prove that

$$(g \circ f)(a) = (g \circ f)(b) \quad \Rightarrow \quad a = b$$

So assume $(g \circ f)(a) = (g \circ f)(b)$. By definition of function composition, this implies that $g(f(a)) = g(f(b))$. By injectivity of g, we have $f(a) = f(b)$; and by injectivity of f, we have $a = b$.   □

Exercise 2.3.5

Let $f : X \to Y$ and $g : Y \to Z$ be functions. Prove that if $g \circ f$ is injective, then f is injective. ◁

Exercise 2.3.6

Write out what it means to say a function $f : X \to Y$ is *not* injective, and say how you would prove that a given function is not injective. Give an example of a function which is not injective, and use your proof technique to write a proof that it is not injective. ◁

Exercise 2.3.7

For each of the following functions, determine whether it is injective or not injective.

- $f : \mathbb{N} \to \mathbb{Z}$, defined by $f(n) = n^2$ for all $n \in \mathbb{N}$.

- $g : \mathbb{Z} \to \mathbb{N}$, defined by $g(n) = n^2$ for all $n \in \mathbb{Z}$.

- $h : \mathbb{N} \times \mathbb{N} \times \mathbb{N} \to \mathbb{N}$, defined by $h(x,y,z) = 2^x \cdot 3^y \cdot 5^z$ for all $x, y, z \in \mathbb{N}$.

◁

Surjectivity

Let's revisit the rows of dots and stars that we saw earlier. Beforehand, we made our idea that there are more dots than stars formal by proving the existence of an injection $f : S \to D$ from the set S of stars to the set D of dots.

However, we could have drawn the same conclusion instead from defining a function $D \to S$, which in some sense *covers* the stars with dots—that is, every star is paired up with at least one dot.

This property is called *surjectivity*—a function $f : X \to Y$ is surjective if every element of Y is a value of f. This is made precise in Definition 2.3.8.

Definition 2.3.8

A function $f : X \to Y$ is **surjective** (or **onto**) if

$$\forall y \in Y, \ \exists x \in X, \ f(x) = y$$

A surjective function is said to be a **surjection**.

Strategy 2.3.9

To prove that a function $f : X \to Y$ is surjective, it suffices to take an arbitrary element $y \in Y$ and, in terms of y, find an element $x \in X$ such that $f(x) = y$.

In order to find x, it is often useful to start from the equation $f(x) = y$ and derive some possible values of x. But be careful—in order to complete the proof, it is necessary to verify that the equation $f(x) = y$ is true for the chosen value of x.   ◁

Example 2.3.10

Fix $n \in \mathbb{N}$ with $n > 0$, and define a function $r : \mathbb{Z} \to \{0, 1, \ldots, n-1\}$ by letting $r(a)$ be the remainder of a when divided by n (see Theorem 0.18). This function is surjective, since for each $k \in \{0, 1, \ldots, n-1\}$ we have $r(k) = k$.   ◁

Exercise 2.3.11

For each of the following pairs of sets (X, Y), determine whether the function $f : X \to Y$ defined by $f(x) = 2x + 1$ is surjective.

  (a) $X = \mathbb{Z}$ and $Y = \{x \in \mathbb{Z} \mid x \text{ is odd}\}$;

  (b) $X = \mathbb{Z}$ and $Y = \mathbb{Z}$;

  (c) $X = \mathbb{Q}$ and $Y = \mathbb{Q}$;

  (d) $X = \mathbb{R}$ and $Y = \mathbb{R}$.

                                                                                              ◁

Exercise 2.3.12

Let $f : X \to Y$ be a function. Find a subset $V \subseteq Y$ and a surjection $g : X \to V$ agreeing with f (that is, such that $g(x) = f(x)$ for all $x \in X$).   ◁

Exercise 2.3.13

Let $f : X \to Y$ be a function. Prove that f is surjective if and only if $Y = f[X]$ ◁

Exercise 2.3.14

Let $f : X \to Y$ be a function. Prove that there is a set Z and functions

$$p : X \to Z \quad \text{and} \quad i : Z \to Y$$

such that p is surjective, i is injective, and $f = i \circ p$.   ◁

Exercise 2.3.15

Let $f : X \to \mathscr{P}(X)$ be a function. By considering the set $B = \{x \in X \mid x \notin f(x)\}$, prove that f is not surjective. ◁

Bijectivity

Bijective functions formalise the idea of putting sets into one-to-one correspondence—each element of one set is paired with exactly one element of another.

Definition 2.3.16

A function $f : X \to Y$ is **bijective** if it is injective and surjective. A bijective function is said to be a **bijection**.

Proof tip

To prove that a function f is bijective, prove that it is injective and surjective. ◁

Example 2.3.17

Let $D \subseteq \mathbb{Q}$ be the set of *dyadic rational numbers*, that is

$$D = \left\{ x \in \mathbb{Q} \ \middle| \ x = \frac{a}{2^n} \text{ for some } a \in \mathbb{Z} \text{ and } n \in \mathbb{N} \right\}$$

Let $k \in \mathbb{N}$, and define $f : D \to D$ by $f(x) = \frac{x}{2^k}$. We will prove that f is a bijection.

- (**Injectivity**) Fix $x, y \in D$ and suppose that $f(x) = f(y)$. Then $\frac{x}{2^k} = \frac{y}{2^k}$, so that $x = y$, as required.

- (**Surjectivity**) Fix $y \in D$. We need to find $x \in D$ such that $f(x) = y$. Well certainly if $2^k y \in D$ then we have

$$f(2^k y) = \frac{2^k y}{2^k} = y$$

  so it suffices to prove that $2^k y \in D$. Since $y \in D$, we must have $y = \frac{a}{2^n}$ for some $n \in \mathbb{N}$.

  ◇ If $k \leqslant n$ then $n - k \in \mathbb{N}$ and so $2^k y = \frac{a}{2^{n-k}} \in D$.

  ◇ If $k > n$ then $k - n > 0$ and $2^k y = 2^{k-n} a \in \mathbb{Z}$; but $\mathbb{Z} \subseteq D$ since if $a \in \mathbb{Z}$ then $a = \frac{a}{2^0}$. So again we have $2^k y \in D$.

  In any case we have $2^k y \in D$ and $f(2^k y) = y$, so that f is surjective.

Since f is both injective and surjective, it is bijective. ◁

Exercise 2.3.18

Let X be a set. Prove that the identity function $\mathrm{id}_X : X \to X$ is a bijection.   ◁

Exercise 2.3.19

Let $n \in \mathbb{N}$ and let $\{X_k \mid 1 \leqslant k \leqslant n\}$ be a family of sets. Prove by induction on n

that there is a bijection $\displaystyle\prod_{k=1}^{n+1} X_k \to \left(\prod_{k=1}^{n} X_k \right) \times X_n$.   ◁

Exercise 2.3.20

Let $f : X \to Y$ and $g : Y \to Z$ be bijections. Prove that $g \circ f$ is a bijection.   ◁

Inverses

Our next goal is to characterise injections, surjections and bijections in terms of other functions, called *inverses*.

Recall Definition 2.3.1, which says that a function $f : X \to Y$ is injective if, for all $a, b \in X$, if $f(a) = f(b)$ then $a = b$.

Exercise 2.3.21

Let $f : X \to Y$ be a function. Prove that f is injective if and only if

$$\forall y \in f[X], \exists! x \in X, y = f(x)$$

◁

Thinking back to Section 2.2, you might notice that this means that the logical formula '$y = f(x)$' defines a function $f[X] \to X$—specifically, if f is injective then there is a function $g : f[X] \to X$ which is (well-)defined by specifying $x = g(f(x))$ for all $x \in X$. Thinking of f as an *encoding* function, we then have that g is the corresponding *decoding* function—decoding is possible by injectivity of f. (If f were not injective then distinct elements of X might have the same encoding, in which case we're stuck if we try to decode them!)

Exercise 2.3.22

Define a function $e : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ by $e(m, n) = 2^m \cdot 3^n$. Prove that e is injective. We can think of e as encoding *pairs* of natural numbers as single natural numbers— for example, the pair $(4, 1)$ is encoded as $2^4 \cdot 3^1 = 48$. For each of the following natural numbers k, find the pairs of natural numbers encoded by e as k:

$$1 \qquad 24 \qquad 7776 \qquad 59049 \qquad 396718580736$$

◁

In Exercise 2.3.22, we were able to decode any natural number of the form $2^m \cdot 3^n$ for $m, n \in \mathbb{N}$. This process of decoding yields a function

$$d : \{k \in \mathbb{N} \mid k = 2^m \cdot 3^n \text{ for some } m, n \in \mathbb{N}\} \to \mathbb{N} \times \mathbb{N}$$

What would happen if we tried to decode a natural number not of the form $2^m \cdot 3^n$ for $m, n \in \mathbb{N}$, say 5 or 100? Well... it doesn't really matter! All we need to be true is that $d(e(m,n)) = (m,n)$ for all $(m,n) \in \mathbb{N} \times \mathbb{N}$; the value of d on other natural numbers is irrelevant.

Definition 2.3.23
Let $f : X \to Y$ be a function. A **left inverse** (or **post-inverse**) for f is a function $g : Y \to X$ such that $g \circ f = \mathrm{id}_X$.

Example 2.3.24
Let $e : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ be as in Exercise 2.3.22. Define a function $d : \mathbb{N} \to \mathbb{N} \times \mathbb{N}$ by

$$d(k) = \begin{cases} (m,n) & \text{if } k = 2^m \cdot 3^n \text{ for some } m, n \in \mathbb{N} \\ (0,0) & \text{otherwise} \end{cases}$$

Note that d is well-defined by the fundamental theorem of arithmetic (Theorem 4.2.12). Moreover, given $m, n \in \mathbb{N}$, we have

$$d(e(m,n)) = d(2^m \cdot 3^n) = (m,n)$$

and so d is a left inverse for e. ◁

Exercise 2.3.25
Let $f : X \to Y$ be a function. Prove that if f has a left inverse, then f is injective. ◁

Exercise 2.3.25 gives us a new strategy for proving that a function is injective.

Strategy 2.3.26 (Proving a function is injective by finding a left inverse)
In order to prove that a function $f : X \to Y$ is injective, it suffices to find a function $g : Y \to X$ such that $g(f(x)) = x$ for all $x \in X$. ◁

It would be convenient if the converse to Exercise 2.3.25 were true—and it is, provided that we impose the condition that the domain of the function be inhabited.

Proposition 2.3.27
Let $f : X \to Y$ be a function. If f is injective and X is inhabited, then f has a left inverse.

Proof

Suppose that f is injective and X is inhabited. Fix $x_0 \in X$—note that this element exists since X is inhabited—and define $g : Y \to X$ as follows.

$$g(y) = \begin{cases} x & \text{if } y = f(x) \text{ for some } x \in X \\ x_0 & \text{otherwise} \end{cases}$$

The only part of the specification of g that might cause it to fail to be well-defined is the case when $y = f(x)$ for some $x \in X$. The reason why g is well-defined is precisely injectivity of f: if $y = f(x)$ for some $x \in X$, then the value of $x \in X$ for which $y = f(x)$ is unique. (Indeed, if $a \in X$ satisfied $y = f(a)$, then we'd have $a = x$ by injectivity of f.)

So g is indeed well-defined. To see that g is a left inverse for f, let $x \in X$. Letting $y = f(x)$, we see that y falls into the first case in the specification of g, so that $g(f(x)) = g(y) = a$ for the value of $a \in X$ for which $y = f(a)$—but as noted above, we have $a = x$ by injectivity of f. □

Exercise 2.3.28

Let $f : X \to Y$ be a function with left inverse $g : Y \to X$. Prove that g is a surjection. ◁

We established a relationship between injections and left inverses in Exercise 2.3.25 and proposition 2.3.27, so it might come as no surprise that there is a relationship between surjections and *right* inverses.

Definition 2.3.29

Let $f : X \to Y$ be a function. A **right inverse** (or **pre-inverse**) for f is a function $g : Y \to X$ such that $f \circ g = \mathrm{id}_Y$.

Example 2.3.30

Define $f : \mathbb{R} \to \mathbb{R}^{\geq 0}$ by $f(x) = x^2$. Note that f is surjective, since for each $y \in \mathbb{R}^{\geq 0}$ we have $\sqrt{y} \in \mathbb{R}$ and $f(\sqrt{y}) = y$. However f is not injective; for instance

$$f(-1) = 1 = f(1)$$

Here are three right inverses for f:

- The positive square root function $g : \mathbb{R}^{\geq 0} \to \mathbb{R}$ defined by $g(y) = \sqrt{y}$ for all $y \in \mathbb{R}^{\geq 0}$. Indeed, for each $y \in \mathbb{R}^{\geq 0}$ we have

$$f(g(y)) = f(\sqrt{y}) = (\sqrt{y})^2 = y$$

- The negative square root function $h : \mathbb{R}^{\geq 0} \to \mathbb{R}$ defined by $h(y) = -\sqrt{y}$ for all $y \in \mathbb{R}^{\geq 0}$. Indeed, for each $y \in \mathbb{R}^{\geq 0}$ we have

$$f(h(y)) = f(-\sqrt{y}) = (-\sqrt{y})^2 = y$$

- The function $k : \mathbb{R}^{\geq 0} \to \mathbb{R}$ defined by

$$k(y) = \begin{cases} \sqrt{y} & \text{if } 2n \leq y < 2n+1 \text{ for some } n \in \mathbb{N} \\ -\sqrt{y} & \text{otherwise} \end{cases}$$

Note that k is well-defined, and again $f(k(y)) = y$ for all $y \in \mathbb{R}^{\geq 0}$ since no matter what value $k(y)$ takes, it is equal to either $\sqrt{y}$ or $-\sqrt{y}$.

There are many more right inverses for f—in fact, there are infinitely many more! ◁

Exercise 2.3.31
Let $f : X \to Y$ be a function. Prove that if f has a right inverse, then f is surjective.
◁

Strategy 2.3.32 (Proving a function is surjective by finding a right inverse)
In order to prove that a function $f : X \to Y$ is surjective, it suffices to find a function $g : Y \to X$ such that $f(g(y)) = y$ for all $y \in Y$. ◁

Interlude: the axiom of choice

It would be convenient if the converse to Exercise 2.3.31 were true—that is, if $f : X \to Y$ is surjective, then it has a right inverse. Let's examine what a proof of this fact would entail. The fact that $f : X \to Y$ is surjective can be expressed as

$$\forall y \in Y, \exists x \in X, f(x) = y$$

A right inverse would be a function $g : Y \to X$, so by Definition 2.2.1, it must satisfy the following condition

$$\forall y \in Y, \exists! x \in X, g(y) = x$$

The temptation is therefore to construct $g : Y \to X$ as follows. Let $y \in Y$. By definition of surjectivity, there exists some $x \in X$ such that $f(x) = y$—define $g(y)$ to be such an element x. Then we have $f(g(y)) = f(x) = y$, as required.

There is an extremely subtle—but important—issue with this construction.

By choosing $g(y)$ to be a fixed element of X such that $f(x) = y$, we are assuming ahead of time that there is a mechanism for choosing, for each $y \in Y$, a unique element of $f^{-1}[\{y\}]$ to be the value of $g(y)$. In other words we are assuming that some statement $R(x,y)$ satisfies the property

$$\forall y \in Y, \exists! x \in X, [x \in f^{-1}[\{y\}] \wedge R(x,y)]$$

But by Definition 2.2.1, this assumption is saying exactly that there exists a function $Y \to X$ that associates to each $y \in Y$ an element $x \in X$ such that $f(x) = y$.

To state this in plainer terms: we tried to prove that there exists a right inverse for f by assuming that there exists a right inverse for f. Evidently, this is not a valid proof strategy.

Surprisingly, it turns out that neither the assumption that every surjection has a right inverse, nor the assumption that there exists a surjection with no right inverse, leads to a contradiction. As such, the assertion that every surjection has a right inverse is *provably unprovable*, although the proof that it is unprovable is far beyond the scope of this textbook.

Nonetheless, the construction of a right inverse that we gave above didn't *feel* like we were abusing the fabric of mathematics and logic.

The essence of the proof is that if a statement of the form $\forall a \in A, \exists b \in B, p(a,b)$ is true, then we should be able to define a function $h : A \to B$ such that $p(a, h(a))$ is true for all $a \in A$: the function h 'chooses' for each $a \in A$ a particular element $b = h(a) \in B$ such that $p(a, b)$ is true.

What makes this possible is to *axiom of choice*, stated precisely below.

Axiom 2.3.33 (Axiom of choice)

Let $\{X_i \mid i \in I\}$ be a family of inhabited sets. Then there is a function $h : I \to \bigcup_{i \in I} X_i$ such that $h(i) \in X_i$ for each $i \in I$.

There are reasons to keep track of the axiom of choice:

- The axiom of choice is perhaps the *strangest* assumption that we make—most of the other axioms that we have stated have been 'evidently true', but this is not the case for the axiom of choice;

- There are fields of mathematics which require the translation of results about sets into results about other kinds of objects—knowing whether the axiom of choice is necessary to prove a result tells us whether this is possible;

- The axiom of choice is highly non-constructive: if a proof of a result that does not use the axiom of choice is available, it usually provides more information than a proof of the same result that does use the axiom of choice.

With this in mind, when we need to invoke the axiom of choice to prove a result, we will mark the result with the letters **AC**. This can be freely ignored on first

reading, but readers may find it useful when using this book as a reference at a later date.

Proposition^{AC} 2.3.34

Let X and Y be sets and let $p(x,y)$ be a logical formula with free variables $x \in X$ and $y \in Y$. If $\forall x \in X$, $\forall y \in Y$, $p(x,y)$ is true, then there exists a function $h : X \to Y$ such that $\forall x \in X$, $p(x,h(x))$ is true.

Proof

For each $a \in X$, define $Y_a = \{b \in Y \mid p(a,b)\}$. Note that Y_a is inhabited for each $a \in X$ by the assumption that $\forall x \in X$, $\exists y \in Y$, $p(x,y)$ is true. Since $Y_a \subseteq Y$ for each $a \in X$, by the axiom of choice there exists a function $h : X \to Y$ such that $h(a) \in Y_a$ for all $a \in X$. But then $p(a,h(a))$ is true for each $a \in X$ by definition of the sets Y_a. $\qquad\square$

In light of Proposition 2.3.34, the axiom of choice manifests itself in proofs as the following proof strategy.

Strategy^{AC} 2.3.35 (Making choices)
If an assumption in a proof has the form $\forall x \in X$, $\exists y \in Y$, $p(x,y)$, then we may make a choice, for each $a \in A$, of a particular element $b = b_a \in B$ for which $p(a,b)$ is true. $\qquad\triangleleft$

Back to inverses

We now return to the converse of Exercise 2.3.31.

Proposition^{AC} 2.3.36

Every surjection has a right inverse.

Proof

Let $f : X \to Y$ be a surjection, and define $g : Y \to X$ as follows. Given $y \in Y$, define $g(y)$ to be a particular choice of $x \in X$ such that $f(x) = y$—note that there exists such an element $x \in X$ since f is surjective, so g exists by Strategy 2.3.35. But then by definition of g we have $f(g(y)) = y$ for all $y \in Y$, so that g is a surjection. $\qquad\square$

It seems logical that we might be able to classify bijections as being those functions which have a left inverse and a right inverse. We can actually say something stronger—the left and right inverse can be taken to be the same function! (In fact, Proposition 2.3.42 establishes that they are necessarily the same function.)

Definition 2.3.37

Let $f : X \to Y$ be a function. A (**two-sided**) **inverse** for f is a function $g : Y \to X$ which is both a left inverse and a right inverse for f.

It is customary to simply say 'inverse' rather than 'two-sided inverse'.

Example 2.3.38

Let D be the set of dyadic rational numbers, as defined in Example 2.3.17. There, we defined a function $f : D \to D$ defined by $f(x) = \frac{x}{2^k}$ for all $x \in D$, where k is some fixed natural number. We find an inverse for f.

Define $g : D \to D$ by $g(x) = 2^k x$. Then

- g is a left inverse for f. To see this, note that for all $x \in D$ we have

$$g(f(x)) = g\left(\frac{x}{2^k}\right) = 2^k \cdot \frac{x}{2^k} = x$$

- g is a right inverse for f. To see this, note that for all $y \in D$ we have

$$f(g(y)) = f(2^k y) = \frac{2^k y}{2^k} = y$$

Since g is a left inverse for f and a right inverse for f, it is a two-sided inverse for f. ◁

Exercise 2.3.39

The following functions have two-sided inverses. For each, find its inverse and prove that it is indeed an inverse.

(a) $f : \mathbb{R} \to \mathbb{R}$ defined by $f(x) = \frac{2x+1}{3}$.

(b) $g : \mathscr{P}(\mathbb{N}) \to \mathscr{P}(\mathbb{N})$ defined by $g(X) = \mathbb{N} \setminus X$.

(c) $h : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ defined by $h(m,n) = 2^m(2n+1) - 1$ for all $m, n \in \mathbb{N}$.

◁

In light of the correspondences between injections and left inverses, and surjections and right inverses, it may be unsurprising that there is a correspondence between *bijections* and *two-sided inverses*.

Exercise 2.3.40

Let $f : X \to Y$ be a function. Then f is bijective if and only if f has an inverse. ◁

Strategy 2.3.41 (Proving a function is bijective by finding an inverse)

In order to prove that a function $f : X \to Y$ is bijective, it suffices to find a function $g : Y \to X$ such that $g(f(x)) = x$ for all $x \in X$ and $f(g(y)) = y$ for all $y \in Y$. ◁

When proving a function $f : X \to Y$ is bijective by finding an inverse $g : Y \to X$, it is important to check that g is *both* a left inverse *and* a right inverse for f. If you only prove that g is a left inverse for f, for example, then you have only proved that f is injective!

It turns out that if a function has both a left and a right inverse, then they must be equal. This is the content of the following proposition.

Proposition 2.3.42

Let $f : X \to Y$ be a function and suppose $\ell : Y \to X$ is a left inverse for f and $r : Y \to X$ is a right inverse for f. Then $\ell = r$.

Proof

The proof is deceptively simple:

$$\begin{aligned}
\ell &= \ell \circ \mathrm{id}_Y && \text{by definition of identity functions} \\
&= \ell \circ (f \circ r) && \text{since } r \text{ is a right inverse for } f \\
&= (\ell \circ f) \circ r && \text{by Exercise 2.2.22} \\
&= \mathrm{id}_X \circ r && \text{since } \ell \text{ is a left inverse for } f \\
&= r && \text{by definition of identity functions}
\end{aligned}$$

$\square$

There is some intuition behind why the left and right inverses of a function $f : X \to Y$ should be equal if they both exist.

- A left inverse $\ell : Y \to X$ exists only if f is injective. It looks at each element $y \in Y$ and, if it is in the image of f, returns the (unique) value $x \in X$ for which $f(x) = y$.

- A right inverse $r : Y \to X$ exists only if f is surjective. It looks at each element $y \in Y$ and picks out one of the (possibly many) values $x \in X$ for which $f(x) = y$.

When f is a bijection, every element of Y is in the image of f (by surjectivity), and is a value of f at a unique element of X (by injectivity), and so the left and right inverses are forced to return the same value on each input—hence they are equal.

It follows from Proposition 2.3.42 that, for any function $f : X \to Y$, any two inverses for f are equal—that is, every bijective function has a *unique* inverse!

Notation 2.3.43

Let $f : X \to Y$ be a function. Write $f^{-1} : Y \to X$ to denote the (unique) inverse for f, if it exists.

Proposition 2.3.44

Let $f : X \to Y$ be a bijection. A function $g : Y \to X$ is a left inverse for f if and only if it is a right inverse for f.

Proof

We will prove the two directions separately.

- ($\Rightarrow$) Suppose $g : Y \to X$ is a left inverse for f—that is, $g(f(x)) = x$ for all $x \in X$. We prove that $f(g(y)) = y$ for all $y \in Y$, thus establishing that g is a right inverse for f. So let $y \in Y$. Since f is a bijection, it is in particular a surjection, so there exists $x \in X$ such that $y = f(x)$. But then

$$
\begin{aligned}
f(g(y)) &= f(g(f(x))) && \text{since } y = f(x) \\
&= f(x) && \text{since } g(f(x)) = x \\
&= y && \text{since } y = f(x)
\end{aligned}
$$

  So indeed g is a right inverse for f.

- ($\Leftarrow$) Suppose $g : Y \to X$ is a right inverse for f—that is, $f(g(y)) = y$ for all $y \in Y$. We prove that $g(f(x)) = x$ for all $x \in X$, thus establishing that g is a left inverse for f. So let $x \in X$. Letting $y = f(x)$, we have $f(g(y)) = y$ since g is a right inverse for f. This says precisely that $f(g(f(x))) = f(x)$, since $y = f(x)$. By injectivity of f, we have $g(f(x)) = x$, as required.

$\square$

Exercise 2.3.45

Let $f : X \to Y$ be a bijection. Prove that $f^{-1} : Y \to X$ is a bijection. ◁

Exercise 2.3.46

Let $f : X \to Y$ and $g : Y \to Z$ be bijections. Prove that $g \circ f : X \to Z$ is a bijection, and write an expression for its inverse in terms of f^{-1} and g^{-1}. ◁

At the beginning of this section we motivated the definitions of injections, surjections and bijections by using them to compare two quantities (of dots and stars)—however, as you might have noticed, we have not yet actually proved that thais intuition aligns with reality. For example, how do we know that if there is an injection $f : X \to Y$ then Y has at least as many elements as X?

Answering this seemingly simple question is surprisingly difficult and has different answers depending on whether the sets involved are finite or infinite. In fact, the words 'finite', 'infinite' and 'size' are themselves defined in terms of injections, surjections and bijections! We therefore leave this task to future sections.

In Section 3.2, we define what it means for a set to be finite and what the size of a finite set is (Definition 3.2.1), and then prove that the sizes of finite sets

can be compared by finding an injection, surjection or bijection between them
Theorem 3.2.6.

Comparing the sizes of infinite sets, and even defining what 'size' means for infinite sets, is another can of worms entirely and leads to some fascinating mathematics. We begin this journey in Section 6.1, where we prove some counterintuitive results, such as the set $\mathbb{N}$ of natural numbers and the set $\mathbb{Q}$ of rational numbers have the same size (Theorem 6.1.8).

Section 2.Q
Chapter 2 exercises

Under construction!
The end-of-chapter exercise sections are new and in an incomplete state.

Set notation

1. Express the following sets in the indicated form of notation.

(a) $\{n \in \mathbb{Z} \mid n^2 < 20\}$ in list notation;

(b) $\{4k+3 \mid k \in \mathbb{N}\}$ in implied list notation;

(c) The set of all odd multiples of six in set-builder notation;

(d) The set $\{1, 2, 5, 10, 17, \ldots, n^2 + 1, \ldots\}$ in set-builder notation.

Set operations

2. For each of the following statements, determine whether it is true for all sets X, Y, false for all sets X, Y, or true for some choices of X and Y and false for others.

(a) $\mathscr{P}(X \cup Y) = \mathscr{P}(X) \cup \mathscr{P}(Y)$ (c) $\mathscr{P}(X \times Y) = \mathscr{P}(X) \times \mathscr{P}(Y)$

(b) $\mathscr{P}(X \cap Y) = \mathscr{P}(X) \cap \mathscr{P}(Y)$ (d) $\mathscr{P}(X \setminus Y) = \mathscr{P}(X) \setminus \mathscr{P}(Y)$

Images and preimages

3. Let $f : X \to Y$ be a function. For each of the following statements, either prove it is true or find a counterexample.

(a) $U \subseteq f^{-1}[f[U]]$ for all $U \subseteq X$;   (c) $V \subseteq f[f^{-1}[V]]$ for all $V \subseteq Y$;

(b) $f^{-1}[f[U]] \subseteq U$ for all $U \subseteq X$;   (d) $f[f^{-1}[V]] \subseteq V$ for all $V \subseteq Y$.

Injections, surjections and bijections

4. (a) Prove that, for all functions $f : X \to Y$ and $g : Y \to Z$, if $g \circ f$ is bijective, then f is injective and g is surjective.

(b) Find an example of a function $f : X \to Y$ and a function $g : Y \to Z$ such that $g \circ f$ is bijective, f is not surjective and g is not injective.

5. For each of the following pairs (U, V) of subsets of $\mathbb{R}$, determine whether the specification '$f(x) = x^2 - 4x + 7$ for all $x \in U$' defines a function $f : U \to V$ and, if it does, determine whether f is injective and whether f is surjective.

(a) $U = \mathbb{R}$ and $V = \mathbb{R}$;

(b) $U = (1, 4)$ and $V = [3, 7)$;

(c) $U = [3, 4)$ and $V = [4, 7)$;

(d) $U = (3, 4]$ and $V = [4, 7)$;

(e) $U = [2, \infty)$ and $V = [3, \infty)$;

(f) $U = [2, \infty)$ and $V = \mathbb{R}$.

6. For each of the following pairs of sets X and Y, find (with proof) a bijection $f : X \to Y$.

(a) $X = \mathbb{Z}$ and $Y = \mathbb{N}$;

(b) $X = \mathbb{R}$ and $Y = (-1, 1)$;

(c) $X = [0, 1]$ and $Y = (0, 1)$;

(d) $X = [a, b]$ and $Y = (c, d)$, where $a, b, c, d \in \mathbb{R}$ with $a < b$ and $c < d$.

Chapter 3

Finite sets

We all have a sense of what the word 'finite' means, but when it comes to giving a precise mathematical definition, it is not entirely obvious how to do it.

The definition that we will eventually give (Definition 3.2.1) amounts to saying that a set X is finite if, for some natural number n, the elements of X can be paired up with the natural numbers from 1 to n—this definition is useful because it tells us that the set X has n elements.

This close relationship between the natural numbers and finite sets means that, if we're going to get very far in proving propositions involving finite sets, we must first obtain a good understanding of what the natural numbers are and how we can prove things about them.

In Section 3.1, we begin by isolating the fundamental properties that the natural numbers satisfy—these properties are called *Peano's axioms*. We then use Peano's axioms to develop several strategies for proving propositions about natural numbers. These strategies are all examples of a more general proof schema called *induction*, which is studied in more generality in Section 5.3.

We get around to defining what a finite set is in Section 3.2, and we prove some useful strategies for proving that a set is finite. What is often more useful than knowing that a set is finite, though, is knowing how many elements it has— answering this question is the topic of Section 3.3.

The natural numbers

The purpose of this section is to forget everything we think we know about the natural numbers, and reconstruct our former knowledge (and more!) using the following fundamental property:

> *Every natural number can be obtained in a unique way by*
> *starting from zero and adding one some finite number of times.*

This is slightly imprecise—it is not clear what is meant by 'adding one some finite number of times', for example. Worse still, we are going to define what 'finite' means in terms of natural numbers in Section 3.2, so we'd better not refer to finiteness in our definition of natural numbers!

The following definition captures precisely the properties that we need in order to characterise the idea of $\mathbb{N}$ that we have in our minds. To begin with, $\mathbb{N}$ should be a set. Whatever the elements of this set $\mathbb{N}$ actually *are*, we will think about them as being natural numbers. One of the elements, in particular, should play the role of the natural number 0—this will be the *zero element* $z \in \mathbb{N}$; and there should be a notion of 'adding one'—this will be the *successor function* $s : \mathbb{N} \to \mathbb{N}$. Thus given an element $n \in \mathbb{N}$, though of as a natural number, we think about the element $s(n)$ as the natural number '$n + 1$'. Note that this is strictly for the purposes of intuition: we will define '$+$' and '1' in terms of z and s, not vice versa.

Definition 3.1.1
A **notion of natural numbers** is a set $\mathbb{N}$, together with an element $z \in \mathbb{N}$, called a **zero element**, and a function $s : \mathbb{N} \to \mathbb{N}$ called a **successor function**, satisfying the following properties:

(i) $z \notin s[\mathbb{N}]$; that is, $z \ne s(n)$ for any $n \in \mathbb{N}$.

(ii) s is injective; that is, for all $m, n \in \mathbb{N}$, if $s(m) = s(n)$, then $m = n$.

(iii) $\mathbb{N}$ is generated by z and s; that is, for all sets X, if $z \in X$ and $s(n) \in X$ for all $n \in \mathbb{N}$, then $\mathbb{N} \subseteq X$.

The properties (i), (ii) and (iii) are called **Peano's axioms**.

Note that Definition 3.1.1 does not specify what $\mathbb{N}$, z and s actually are; it just specifies the properties that they must satisfy. It turns out that it doesn't really matter

what notion of natural numbers we use, since any two notions are essentially the same. We will not worry about the specifics here—that task is left to Section B.2: a particular notion of natural numbers is defined in Construction B.2.5, and the fact that all notions of natural numbers are 'essentially the same' is made precise and proved in Theorem B.2.8.

We can define all the concepts involving natural numbers that we are familiar with, and prove all the properties that we take for granted, just from the element $z \in \mathbb{N}$ and the successor function $s : \mathbb{N} \to \mathbb{N}$.

For instance, we define '0' to mean z, define '1' to mean $s(z)$, define '2' to mean $s(s(z))$, and so on. For instance, '12' is defined to mean

$$s(s(s(s(s(s(s(s(s(s(s(s(z))))))))))))$$

From now on, then, let's write 0 instead of z for the zero element of $\mathbb{N}$. It would be nice if we could write '$n+1$' instead of $s(n)$, but we must first define what '+' means. In order to do this, we need a way of defining expressions involving natural numbers; this is what the *recursion theorem* allows us to do.

Theorem 3.1.2 (Recursion theorem)

Let X be a set. For all $a \in X$ and all $h : \mathbb{N} \times X \to X$, there is a unique function $f : \mathbb{N} \to X$ such that $f(0) = a$ and $f(s(n)) = h(n, f(n))$ for all $n \in \mathbb{N}$.

Proof

Let $a \in X$ and $h : \mathbb{N} \times X \to X$. We prove existence and uniqueness of f separately.

- Define $f : \mathbb{N} \to X$ by specifying $f(0) = a$ and $f(s(n)) = h(n, f(n))$. Since h is a function and s is injective, existence and uniqueness of $x \in X$ such that $f(n) = x$ is guaranteed, provided that $f(n)$ is defined, so we need only verify totality.

  So let $D = \{n \in \mathbb{N} \mid f(n) \text{ is defined}\}$. Then:

  $\diamond$ $0 \in D$, since $f(0)$ is defined to be equal to a.

  $\diamond$ Let $n \in \mathbb{N}$ and suppose $n \in D$. Then $f(n)$ is defined and $f(s(n)) = h(n, f(n))$, so that $f(s(n))$ is defined, and hence $s(n) \in D$.

  By condition (iii) of Definition 3.1.1, we have $\mathbb{N} \subseteq D$, so that $f(n)$ is defined for all $n \in \mathbb{N}$, as required.

- To see that f is unique, suppose $g : \mathbb{N} \to X$ were another function such that $g(0) = a$ and $g(s(n)) = h(n, g(n))$ for all $n \in \mathbb{N}$.

  To see that $f = g$, let $E = \{n \in \mathbb{N} \mid f(n) = g(n)\}$. Then

◇ $0 \in E$, since $f(0) = a = g(0)$.

◇ Let $n \in \mathbb{N}$ and suppose that $n \in E$. Then $f(n) = g(n)$, and so

$$f(s(n)) = h(n, f(n)) = h(n, g(n)) = g(s(n))$$

and so $s(n) \in E$.

Again, condition (iii) of Definition 3.1.1 is satisfied, so that $\mathbb{N} \subseteq E$. It follows that $f(n) = g(n)$ for all $n \in \mathbb{N}$, and so $f = g$.

Thus we have established the existence and uniqueness of a function $f : \mathbb{N} \to X$ such that $f(0) = a$ and $f(s(n)) = h(n, f(n))$ for all $n \in \mathbb{N}$.   □

The recursion theorem allows us to define expressions involving natural numbers *by recursion*; this is Strategy 3.1.3.

Strategy 3.1.3 (Definition by recursion)
In order to specify a function $f : \mathbb{N} \to X$, it suffices to define $f(0)$ and, for given $n \in \mathbb{N}$, assume that $f(n)$ has been defined, and define $f(s(n))$ in terms of n and $f(n)$.   ◁

Example 3.1.4

We can use recursion to define addition on the natural numbers as follows.

For fixed $m \in \mathbb{N}$, we can define a function $\mathrm{add}_m : \mathbb{N} \to \mathbb{N}$ by recursion by:

$$\mathrm{add}_m(0) = m \quad \text{and} \quad \mathrm{add}_m(s(n)) = s(\mathrm{add}_m(n)) \text{ for all } n \in \mathbb{N}$$

In more familiar notation, for $m, n \in \mathbb{N}$, define the expression '$m + n$' to mean $\mathrm{add}_m(n)$. Another way of expressing the recursive definition of $\mathrm{add}_m(n)$ is to say that, for each $m \in \mathbb{N}$, we are defining $m + n$ by recursion on n as follows:

$$m + 0 = m \quad \text{and} \quad m + s(n) = s(m + n) \text{ for all } n \in \mathbb{N}$$

◁

We can use the recursive definition of addition to prove familiar equations between numbers. The following proposition is a proof that $2 + 2 = 4$. This may seem silly, but notice that the expression '$2 + 2 = 4$' is actually shorthand for the following:

$$\mathrm{add}_{s(s(0))}(s(s(0))) = s(s(s(s(0))))$$

We must therefore be careful to apply the definitions in its proof.

Proposition 3.1.5

$2 + 2 = 4$

Proof

We use the recursive definition of addition.

$$
\begin{aligned}
2+2 &= 2+s(1) && \text{since } 2 = s(1) \\
&= s(2+1) && \text{by definition of } + \\
&= s(2+s(0)) && \text{since } 1 = s(0) \\
&= s(s(2+0)) && \text{by definition of } + \\
&= s(s(2)) && \text{by definition of } + \\
&= s(3) && \text{since } 3 = s(2) \\
&= 4 && \text{since } 4 = s(3)
\end{aligned}
$$

as required. □

The following result allows us to drop the notation '$s(n)$' and just write '$n+1$' instead.

Proposition 3.1.6

For all $n \in \mathbb{N}$, we have $s(n) = n+1$.

Proof

Let $n \in \mathbb{N}$. Then by the recursive definition of addition we have

$$
n+1 = n+s(0) = s(n+0) = s(n)
$$

as required. □

In light of Proposition 3.1.6, we will now abandon the notation $s(n)$, and write $n+1$ instead.

We can define the arithmetic operations of multiplication and exponentiation by recursion, too.

Example 3.1.7

Fix $m \in \mathbb{N}$. Define $m \cdot n$ for all $n \in \mathbb{N}$ by recursion on n as follows:

$$
m \cdot 0 = 0 \quad \text{and} \quad m \cdot (n+1) = (m \cdot n) + m \text{ for all } n \in \mathbb{N}
$$

Formally, what we have done is define a function $\mathrm{mult}_m : \mathbb{N} \to \mathbb{N}$ recursively by $\mathrm{mult}_m(z) = z$ and $\mathrm{mult}_m(s(n)) = \mathrm{add}_{\mathrm{mult}_m(n)}(m)$ for all $n \in \mathbb{N}$. But the definition we provided is easier to understand. ◁

Proposition 3.1.8

$2 \cdot 2 = 4$

Proof

We use the recursive definitions of addition and recursion.

$$
\begin{aligned}
2 \cdot 2 &= 2 \cdot (1+1) && \text{since } 2 = 1+1 \\
&= (2 \cdot 1) + 2 && \text{by definition of } \cdot \\
&= (2 \cdot (0+1)) + 2 && \text{since } 1 = 0+1 \\
&= ((2 \cdot 0) + 2) + 2 && \text{by definition of } \cdot \\
&= (0+2) + 2 && \text{by definition of } \cdot \\
&= (0 + (1+1)) + 2 && \text{since } 2 = 1+1 \\
&= ((0+1)+1) + 2 && \text{by definition of } + \\
&= (1+1) + 2 && \text{since } 1 = 0+1 \\
&= 2+2 && \text{since } 2 = 1+1 \\
&= 4 && \text{by Proposition 3.1.5}
\end{aligned}
$$

as required.   □

Exercise 3.1.9

Given $m \in \mathbb{N}$, define m^n for all $n \in \mathbb{N}$ by recursion on n, and prove that $2^2 = 4$ using the recursive definitions of exponentiation, multiplication and addition.   ◁

We could spend the rest of our lives doing long computations involving recursively defined arithmetic operations, so at this point we will stop, and return to taking for granted the facts that we know about arithmetic operations.

There are, however, a few more notions that we need to define by recursion so that we can use them in our proofs later.

Definition 3.1.10

The **sum** of real numbers $a_1, a_2, \ldots, a_n$ is the real number $\sum\limits_{k=1}^{n} a_k$ defined by recursion on $n \in \mathbb{N}$ by

$$
\sum_{k=1}^{0} a_k = 0 \quad \text{and} \quad \sum_{k=1}^{n+1} a_k = \left(\sum_{k=0}^{n} a_k \right) + a_{n+1} \text{ for all } n \in \mathbb{N}
$$

Definition 3.1.11

The **product** of real numbers $a_1, a_2, \ldots, a_n$ is the real number $\prod\limits_{k=1}^{n} a_k$ defined by recursion on $n \in \mathbb{N}$ by

$$
\prod_{k=1}^{0} a_k = 1 \quad \text{and} \quad \prod_{k=1}^{n+1} a_k = \left(\prod_{k=0}^{n} a_k \right) \cdot a_{n+1} \text{ for all } n \in \mathbb{N}
$$

Example 3.1.12

Let $x_i = i^2$ for each $i \in \mathbb{N}$. Then

$$\sum_{i=1}^{5} x_i = 1 + 4 + 9 + 16 + 25 = 55$$

and

$$\prod_{i=1}^{5} x_i = 1 \cdot 4 \cdot 9 \cdot 16 \cdot 25 = 14400$$

◁

Exercise 3.1.13

Let $x_1, x_2 \in \mathbb{R}$. Working strictly from the definitions of indexed sum and indexed product, prove that

$$\sum_{i=1}^{2} x_i = x_1 + x_2 \quad \text{and} \quad \prod_{i=1}^{2} x_i = x_1 \cdot x_2$$

◁

Proof by induction

Just as recursion exploited the structure of the natural numbers to *define expressions* involving natural numbers, induction exploits the very same structure to *prove results* about natural numbers.

Theorem 3.1.14 (Weak induction principle)
Let $p(n)$ be logical formula with free variable $n \in \mathbb{N}$, and let $n_0 \in \mathbb{N}$. If

(i) $p(n_0)$ is true; and

(ii) For all $n \geqslant n_0$, if $p(n)$ is true, then $p(n+1)$ is true;

then $p(n)$ is true for all $n \geqslant n_0$.

Proof
Define $X = \{n \in \mathbb{N} \mid p(n_0 + n) \text{ is true}\}$; that is, given a natural number n, we have $n \in X$ if and only if $p(n_0 + n)$ is true. Then

- $0 \in X$, since $n_0 + 0 = n_0$ and $p(n_0)$ is true by (i).

- Let $n \in \mathbb{N}$ and assume $n \in X$. Then $p(n_0 + n)$ is true. Since $n_0 + n \geqslant n_0$ and $p(n_0 + n)$ is true, we have $p(n_0 + n + 1)$ is true by (ii). But then $n_0 + n + 1 \in X$.

So by Definition 3.1.1(iii) we have $\mathbb{N} \subseteq X$. Hence $p(n_0 + n)$ is true for all $n \in \mathbb{N}$. But this is equivalent to saying that $p(n)$ is true for all $n \geqslant n_0$. $\qquad\square$


Strategy 3.1.15 (Proof by (weak) induction)
In order to prove a proposition of the form $\forall n \in \mathbb{N}$, $p(n)$, it suffices to prove that $p(0)$ is true and that, for all $n \in \mathbb{N}$, if $p(n)$ is true, then $p(n+1)$ is true. $\qquad\triangleleft$

Some terminology has evolved for proofs by induction, which we mention now:

- The proof of $p(n_0)$ is called the **base case**;

- The proof of $\forall n \geqslant n_0$, $(p(n) \Rightarrow p(n+1))$ is called the **induction step**;

- In the induction step, the assumption $p(n)$ is called the **induction hypothesis**;

- In the induction step, the proposition $p(n+1)$ is called the **induction goal**.

The following diagram illustrates the weak induction principle.

To interpret this diagram:

- The shaded diamond represents the base case $p(n_0)$;

- The square represents the induction hypothesis $p(n)$;

- The dashed circle represents the induction goal $p(n+1)$;

- The arrow represents the implication we must prove in the induction step.

We will use analogous diagrams to illustrate the other induction principles in this section.

Proposition 3.1.16
Let $n \in \mathbb{N}$. Then $\displaystyle\sum_{k=1}^{n} k = \frac{n(n+1)}{2}$

Proof
We proceed by induction on $n \geqslant 0$.

- (**Base case**) We need to prove $\sum_{k=1}^{0} k = \dfrac{0(0+1)}{2}$.

  This is true, since $\dfrac{0(0+1)}{2} = 0$, and $\sum_{k=1}^{0} k = 0$ by Definition 3.1.10.

- (**Induction step**) Let $n \geqslant 0$ and suppose that $\sum_{k=1}^{n} k = \dfrac{n(n+1)}{2}$; this is the induction hypothesis.

  We need to prove that $\sum_{k=1}^{n+1} k = \dfrac{(n+1)(n+2)}{2}$; this is the induction goal.

  We proceed by calculation:

  $$\begin{aligned}
  \sum_{k=1}^{n+1} k &= \left(\sum_{k=1}^{n} k\right) + (n+1) && \text{by Definition 3.1.10} \\
  &= \frac{n(n+1)}{2} + (n+1) && \text{by induction hypothesis} \\
  &= (n+1)\left(\frac{n}{2} + 1\right) && \text{factorising} \\
  &= \frac{(n+1)(n+2)}{2} && \text{rearranging}
  \end{aligned}$$

The result follows by induction. □

Before moving on, let's reflect on the proof of Proposition 3.1.16 to highlight some effective ways of writing a proof by induction.

- We began the proof by indicating that it was a proof by induction. While it is clear in this section that most proofs will be by induction, that will not always be the case, so it is good practice to indicate the proof strategy at hand.

- The base case and induction step are clearly labelled in the proof. This is not strictly *necessary* from a mathematical perspective, but it helps the reader to navigate the proof and to identify what the goal is at each step.

- We began the induction step by writing, 'Let $n \geqslant n_0$ and suppose that [...*induction hypothesis goes here*...]'. This is typically how your induction step should begin, since the proposition being proved in the induction step is of the form $\forall n \geqslant n_0, (p(n) \Rightarrow \cdots)$.

- Before proving anything in the base case or induction step, we wrote out what it was that we were trying to prove in that part of the proof. This is helpful because it helps to remind us (and the person reading the proof) what we are aiming to achieve.

Look out for these features in the proof of the next proposition, which is also by induction on $n \geqslant 0$.

Proposition 3.1.17
The natural number $n^3 - n$ is divisible by 3 for all $n \in \mathbb{N}$.

Proof

We proceed by induction on $n \geqslant 0$.

- (**Base case**) We need to prove that $0^3 - 0$ is divisible by 3. Well

$$0^3 - 0 = 0 = 3 \times 0$$

  so $0^3 - 0$ is divisible by 3.

- (**Induction step**) Let $n \in \mathbb{N}$ and suppose that $n^3 - n$ is divisible by 3. Then $n^3 - n = 3k$ for some $k \in \mathbb{Z}$.

  We need to prove that $(n+1)^3 - (n+1)$ is divisible by 3; in other words, we need to find some natural number ℓ such that

$$(n+1)^3 - (n+1) = 3\ell$$

  We proceed by computation.

$$
\begin{aligned}
&(n+1)^3 - (n+1) \\
&= (n^3 + 3n^2 + 3n + 1) - n - 1 && \text{expand brackets} \\
&= n^3 - n + 3n^2 + 3n + 1 - 1 && \text{rearrange terms} \\
&= n^3 - n + 3n^2 + 3n && \text{since } 1 - 1 = 0 \\
&= 3k + 3n^2 + 3n && \text{by induction hypothesis} \\
&= 3(k + n^2 + n) && \text{factorise}
\end{aligned}
$$

  Thus we have expressed $(n+1)^3 - (n+1)$ in the form 3ℓ for some $\ell \in \mathbb{Z}$; specifically, $\ell = k + n^2 + n$.

The result follows by induction. □

Exercise 3.1.18
Prove by induction that $\sum_{k=0}^{n} 2^k = 2^{n+1} - 1$ for all $n \in \mathbb{N}$. ◁

The following proposition has a proof by induction in which the base case is not zero.

Proposition 3.1.19
For all $n \geqslant 4$, we have $3n < 2^n$.

Proof

We proceed by induction on $n \geqslant 4$.

- (**Base case**) $p(4)$ is the statement $3 \cdot 4 < 2^4$. This is true, since $12 < 16$.

- (**Induction step**) Suppose $n \geqslant 4$ and that $3n < 2^n$. We want to prove $3(n+1) < 2^{n+1}$. Well,

$$
\begin{aligned}
3(n+1) &= 3n+3 && \text{expanding} \\
&< 2^n +3 && \text{by induction hypothesis} \\
&< 2^n + 2^4 && \text{since } 3 < 16 = 2^4 \\
&\leqslant 2^n + 2^n && \text{since } n \geqslant 4 \\
&= 2 \cdot 2^n && \text{simplifying} \\
&= 2^{n+1} && \text{simplifying}
\end{aligned}
$$

So we have proved $3(n+1) < 2^{n+1}$, as required.

The result follows by induction. $\qquad\qquad\square$

Note that the proof in Proposition 3.1.19 says nothing about the truth or falsity of $p(n)$ for $n = 0, 1, 2, 3$. In order to assert that these cases are false, you need to show them individually; indeed:

- $3 \times 0 = 0$ and $2^0 = 1$, hence $p(0)$ is true;

- $3 \times 1 = 3$ and $2^1 = 2$, hence $p(1)$ is false;

- $3 \times 2 = 6$ and $2^2 = 4$, hence $p(2)$ is false;

- $3 \times 3 = 9$ and $2^3 = 8$, hence $p(3)$ is false.

So we deduce that $p(n)$ is true when $n = 0$ or $n \geqslant 4$, and false when $n \in \{1, 2, 3\}$.

Exercise 3.1.20

Find all natural numbers n such that $n^5 < 5^n$. ◁

Sometimes a 'proof' by induction might appear to be complete nonsense. The following is a classic example of a 'fail by induction':

Example 3.1.21

The following argument supposedly proves that every horse is the same colour.

- (**Base case**) Suppose there is just one horse. This horse is the same colour as itself, so the base case is immediate.

- (**Induction step**) Suppose that every collection of n horses is the same colour. Let X be a set of $n+1$ horses. Removing the first horse from X, we see that the last n horses are the same colour by the induction hypothesis. Removing the last horse from X, we see that the first n horses are the same colour. Hence all the horses in X are the same colour.

By induction, we're done.   ◁

Exercise 3.1.22

Write down the statement $p(n)$ that Example 3.1.21 attempted to prove for all $n \geqslant 1$. Convince yourself that the proof of the base case is correct, then write down—with quantifiers—exactly the proposition that the induction step is meant to prove. Explain why the argument in the induction step failed to prove this proposition.   ◁

There are several ways to avoid situations like that of Example 3.1.21 by simply putting more thought into writing the proof. Some tips are:

- State $p(n)$ explicitly. In the statement 'all horses are the same colour' it is not clear exactly what the induction variable is. However, we could have said:

  Let $p(n)$ be the statement 'every set of n horses has the same colour'.

- Refer explicitly to the base case n_0 in the induction step. In Example 3.1.21, our induction hypothesis simply stated 'assume every set of n horses has the same colour'. Had we instead said:

  Let $n \geqslant 1$ and assume every set of n horses has the same colour.

  We may have spotted the error in what was to come.

What follows are a couple more examples of proofs by weak induction.

Proposition 3.1.23

For all $n \in \mathbb{N}$, we have $\displaystyle\sum_{k=0}^{n} k^3 = \left(\sum_{k=0}^{n} k \right)^2$.

Proof

We proved in Proposition 3.1.19 that $\displaystyle\sum_{k=0}^{n} k = \frac{n(n+1)}{2}$ for all $n \in \mathbb{N}$, thus it suffices to prove that

$$\sum_{k=0}^{n} k^3 = \frac{n^2(n+1)^2}{4}$$

for all $n \in \mathbb{N}$.

We proceed by induction on $n \geqslant 0$.

- **(Base case)** We need to prove that $0^3 = \dfrac{0^2(0+1)^2}{4}$. This is true since both sides of the equation are equal to 0.

- **(Induction step)** Fix $n \in \mathbb{N}$ and suppose that $\displaystyle\sum_{k=0}^{n} k^3 = \dfrac{n^2(n+1)^2}{4}$. We need to prove that $\displaystyle\sum_{k=0}^{n+1} k^3 = \dfrac{(n+1)^2(n+2)^2}{4}$. This is true since:

$$
\begin{aligned}
\sum_{i=0}^{n+1} k^3 &= \sum_{i=0}^{n} k^3 + (n+1)^3 && \text{by definition of sum} \\
&= \frac{n^2(n+1)^2}{4} + (n+1)^3 && \text{by induction hypothesis} \\
&= \frac{n^2(n+1)^2 + 4(n+1)^3}{4} && (\textit{algebra}) \\
&= \frac{(n+1)^2(n^2 + 4(n+1))}{4} && (\textit{algebra}) \\
&= \frac{(n+1)^2(n+2)^2}{4} && (\textit{algebra})
\end{aligned}
$$

By induction, the result follows. $\qquad\square$

In the next proposition, we prove the correctness of a well-known formula for the sum of an *arithmetic progression* of real numbers.

Proposition 3.1.24
Let $a, d \in \mathbb{R}$. Then
$$
\sum_{k=0}^{n}(a+kd) = \frac{(n+1)(2a+nd)}{2}
$$
for all $n \in \mathbb{N}$.

Proof
We proceed by induction on $n \geqslant 0$.

- **(Base case)** We need to prove that $\displaystyle\sum_{k=0}^{0}(a+kd) = \dfrac{(0+1)(2a+0d)}{2}$. This is true, since

$$
\sum_{k=0}^{0}(a+kd) = a+0d = a = \frac{2a}{2} = \frac{1\cdot(2a)}{2} = \frac{(0+1)(2a+0d)}{2}
$$

- (**Induction step**) Fix $n \in \mathbb{N}$ and suppose that $\sum_{k=0}^{n}(a+kd) = \dfrac{(n+1)(2a+nd)}{2}$.

We need to prove:

$$\sum_{k=0}^{n+1}(a+kd) = \frac{(n+2)(2a+(n+1)d)}{2}$$

This is true, since

$$\sum_{k=0}^{n+1}(a+kd)$$

$$= \sum_{k=0}^{n}(a+kd) + (a+(n+1)d) \qquad \text{by definition of sum}$$

$$= \frac{(n+1)(2a+nd)}{2} + (a+(n+1)d) \qquad \text{by induction hypothesis}$$

$$= \frac{(n+1)(2a+nd) + 2a + 2(n+1)d}{2} \qquad (algebra)$$

$$= \frac{(n+1)\cdot 2a + (n+1)\cdot nd + 2a + 2(n+1)d}{2} \qquad (algebra)$$

$$= \frac{2a(n+1+1) + (n+1)(nd+2d)}{2} \qquad (algebra)$$

$$= \frac{2a(n+2) + (n+1)(n+2)d}{2} \qquad (algebra)$$

$$= \frac{(n+2)(2a+(n+1)d)}{2} \qquad (algebra)$$

By induction, the result follows. □

The following exercises generalises Exercise 3.1.18 to prove the correctness of a formula for the sum of a *geometric progression* of real numbers.

Exercise 3.1.25

Let $a, r \in \mathbb{R}$ with $r \neq 1$. Then

$$\sum_{k=0}^{n} ar^{n} = \frac{a(1-r^{n+1})}{1-r}$$

for all $n \in \mathbb{N}$. ◁

When attempting the following exercise, you might find that your induction step requires an auxiliary result, which itself has a proof by induction.

Exercise 3.1.26

Prove by induction that $7^{n} - 2\cdot 4^{n} + 1$ is divisible by 18 for all $n \in \mathbb{N}$. ◁

A first look at binomials and factorials

In Section 3.3, two kinds of natural number will turn out to be extremely useful, namely *factorials* and *binomial coefficients*. These numbers allow us to count the number of elements of certain kinds of sets, and correspond with the 'real-world' processes of *permutation* and *selection*, respectively. Everything we do here will be re-defined and re-proved combinatorially in Section 3.3. In this section, we will overlook the combinatorial nature, and instead characterise them recursively. We will prove that the combinatorial and recursive definitions of binomial coefficients and factorials are equivalent in Section 3.3.

Definition 3.1.27 (to be redefined in Definition 3.3.10)
Let $n \in \mathbb{N}$. The **factorial** of n, written $n!$, is defined recursively by

$$0! = 1 \quad \text{and} \quad (n+1)! = (n+1) \cdot n! \text{ for all } n \geqslant 0$$


Put another way, we have

$$n! = \prod_{i=1}^{n} i$$

for all $n \in \mathbb{N}$—recall Definition 3.1.11 to see why these definitions are really just two ways of wording the same thing.

Definition 3.1.28 (to be redefined in Definition 3.3.4)
Let $n, k \in \mathbb{N}$. The **binomial coefficient** $\binom{n}{k}$ (LaTeX code: `\binom{n}{k}`) (read 'n choose k') is defined by recursion on n *and* on k by

$$\binom{n}{0} = 1, \quad \binom{0}{k+1} = 0, \quad \binom{n+1}{k+1} = \binom{n}{k} + \binom{n}{k+1}$$

This definition gives rise to an algorithm for computing binomial coefficients: they fit into a diagram known as **Pascal's triangle**, with each binomial coefficient computed as the sum of the two lying above it (with zeroes omitted):

$$
\begin{array}{ccccccc}
 & & & \binom{0}{0} & & & \\
 & & \binom{1}{0} & & \binom{1}{1} & & \\
 & \binom{2}{0} & & \binom{2}{1} & & \binom{2}{2} & \\
\binom{3}{0} & & \binom{3}{1} & & \binom{3}{2} & & \binom{3}{3} \\
\end{array}
\qquad = \qquad
\begin{array}{ccccccc}
 & & & 1 & & & \\
 & & 1 & & 1 & & \\
 & 1 & & 2 & & 1 & \\
1 & & 3 & & 3 & & 1 \\
\end{array}
$$

with rows $\binom{4}{0}\,\binom{4}{1}\,\binom{4}{2}\,\binom{4}{3}\,\binom{4}{4}$ and $\binom{5}{0}\,\binom{5}{1}\,\binom{5}{2}\,\binom{5}{3}\,\binom{5}{4}\,\binom{5}{5}$ equalling $1\ 4\ 6\ 4\ 1$ and $1\ 5\ 10\ 10\ 5\ 1$ respectively.

Exercise 3.1.29

Write down the next two rows of Pascal's triangle.   ◁

We can prove lots of identities concerning binomial coefficients and factorials by induction.

Example 3.1.30

We prove that $\sum_{i=0}^{n} \binom{n}{i} = 2^n$ by induction on n.

- (**Base case**) We need to prove $\binom{0}{0} = 1$ and $2^0 = 1$. These are both true by the definitions of binomial coefficients and exponents.

- (**Induction step**) Fix $n \geqslant 0$ and suppose that

$$\sum_{i=0}^{n} \binom{n}{i} = 2^n$$

We need to prove

$$\sum_{i=0}^{n+1} \binom{n+1}{i} = 2^{n+1}$$

This is true, since

$$\sum_{i=0}^{n+1} \binom{n+1}{i}$$

$$= \binom{n+1}{0} + \sum_{i=1}^{n+1} \binom{n+1}{i} \qquad \text{splitting the sum}$$

$$= 1 + \sum_{j=0}^{n} \binom{n+1}{j+1} \qquad \text{letting } j = i - 1$$

$$= 1 + \sum_{j=0}^{n} \left(\binom{n}{j} + \binom{n}{j+1} \right) \qquad \text{by Definition 3.1.28}$$

$$= 1 + \sum_{j=0}^{n} \binom{n}{j} + \sum_{j=0}^{n} \binom{n}{j+1} \qquad \text{separating the sums}$$

Now $\sum_{j=0}^{n} \binom{n}{j} = 2^n$ by the induction hypothesis. Moreover, reindexing the sum using $k = j+1$ yields

$$\sum_{j=0}^{n} \binom{n}{j+1} = \sum_{k=1}^{n+1} \binom{n}{k} = \sum_{k=1}^{n} \binom{n}{k} + \binom{n}{n+1}$$

By the induction hypothesis, we have

$$\sum_{k=1}^{n} \binom{n}{k} = \sum_{k=0}^{n} \binom{n}{k} - \binom{n}{0} = 2^n - 1$$

and $\binom{n}{n+1} = 0$, so that $\sum_{j=0}^{n} \binom{n}{j+1} = 2^n - 1$.

Putting this together, we have

$$1 + \sum_{j=0}^{n} \binom{n}{j} + \sum_{j=0}^{n} \binom{n}{j+1} = 1 + 2^n + (2^n - 1)$$
$$= 2 \cdot 2^n$$
$$= 2^{n+1}$$

so the induction step is finished.

By induction, we're done. ◁

Exercise 3.1.31
Prove by induction on $n \geqslant 1$ that

$$\sum_{i=0}^{n} (-1)^i \binom{n}{i} = 0$$

◁

Theorem 3.1.32
Let $n, k \in \mathbb{N}$. Then
$$\binom{n}{k} = \begin{cases} \frac{n!}{k!(n-k)!} & \text{if } k \leqslant n \\ 0 & \text{if } k > n \end{cases}$$

Proof
We proceed by induction on n.

- (**Base case**) When $n = 0$, we need to prove that $\binom{0}{k} = \frac{0!}{k!(-k)!}$ for all $k \leqslant 0$, and that $\binom{0}{k} = 0$ for all $k > 0$.

  If $k \leqslant 0$ then $k = 0$, since $k \in \mathbb{N}$. Hence we need to prove

  $$\binom{0}{0} = \frac{0!}{0!0!}$$

  But this is true since $\binom{0}{0} = 1$ and $\frac{0!}{0!0!} = \frac{1}{1 \times 1} = 1$.

  If $k > 0$ then $\binom{0}{k} = 0$ by Definition 3.1.28.

- (**Induction step**) Fix $n \in \mathbb{N}$ and suppose that $\binom{n}{k} = \frac{n!}{k!(n-k)!}$ for all $k \leqslant n$ and $\binom{n}{k} = 0$ for all $k > n$.

  We need to prove that, for all $k \leqslant n+1$, we have

  $$\binom{n+1}{k} = \frac{(n+1)!}{k!(n+1-k)!}$$

  and that $\binom{n+1}{k} = 0$ for all $k > n+1$.

  So fix $k \in \mathbb{N}$. There are four possible cases: either (i) $k = 0$, or (ii) $0 < k \leqslant n$, or (iii) $k = n+1$, or (iv) $k > n+1$. In cases (i), (ii) and (iii), we need to prove the factorial formula for $\binom{n+1}{k}$; in case (iv), we need to prove that $\binom{n+1}{k} = 0$.

  (i) Suppose $k = 0$. Then $\binom{n+1}{0} = 1$ by Definition 3.1.28, and

  $$\frac{(n+1)!}{k!(n+1-k)!} = \frac{(n+1)!}{0!(n+1)!} = 1$$

  since $0! = 1$. So $\binom{n+1}{0} = \frac{(n+1)!}{0!(n+1)!}$.

  (ii) If $0 < k \leqslant n$ then $k = \ell + 1$ for some natural number $\ell < n$. Then $\ell + 1 \leqslant n$, so we can use the induction hypothesis to apply factorial formula to both $\binom{n}{\ell}$ and $\binom{n}{\ell+1}$. Hence

  $$
  \begin{aligned}
  & \binom{n+1}{k} \\
  & = \binom{n+1}{\ell+1} && \text{since } k = \ell + 1 \\
  & = \binom{n}{\ell} + \binom{n}{\ell+1} && \text{by Definition 3.1.28} \\
  & = \frac{n!}{\ell!(n-\ell)!} + \frac{n!}{(\ell+1)!(n-\ell-1)!} && \text{by induction hypothesis}
  \end{aligned}
  $$

  Now note that

  $$\frac{n!}{\ell!(n-\ell)!} = \frac{n!}{\ell!(n-\ell)!} \cdot \frac{\ell+1}{\ell+1} = \frac{n!}{(\ell+1)!(n-\ell)!} \cdot (\ell+1)$$

  and

  $$\frac{n!}{(\ell+1)!(n-\ell-1)!} = \frac{n!}{(\ell+1)!(n-\ell-1)!} \cdot \frac{n-\ell}{n-\ell} = \frac{n!}{(\ell+1)!(n-\ell)!} \cdot (n-\ell)$$

Piecing this together, we have

$$\frac{n!}{\ell!(n-\ell)!} + \frac{n!}{(\ell+1)!(n-\ell-1)!}$$
$$= \frac{n!}{(\ell+1)!(n-\ell)!} \cdot [(\ell+1)+(n-\ell)]$$
$$= \frac{n!(n+1)}{(\ell+1)!(n-\ell)!}$$
$$= \frac{(n+1)!}{(\ell+1)!(n-\ell)!}$$

so that $\binom{n+1}{\ell+1} = \frac{(n+1)!}{(\ell+1)!(n-\ell)!}$. Now we're done; indeed,

$$\frac{(n+1)!}{(\ell+1)!(n-\ell)!} = \frac{(n+1)!}{k!(n+1-k)!}$$

since $k = \ell + 1$.

(iii) If $k = n + 1$, then

$$\binom{n+1}{k} = \binom{n+1}{n+1} \qquad \text{since } k = n+1$$
$$= \binom{n}{n} + \binom{n}{n+1} \qquad \text{by Definition 3.1.28}$$
$$= \frac{n!}{n!0!} + 0 \qquad \text{by induction hypothesis}$$
$$= 1$$

and $\frac{(n+1)!}{(n+1)!0!} = 1$, so again the two quantities are equal.

(iv) If $k > n + 1$, then $k = \ell + 1$ for some $\ell > n$, and so by Definition 3.1.28 and the induction hypothesis we have

$$\binom{n+1}{k} = \binom{n+1}{\ell+1} \overset{\text{IH}}{=} \binom{n}{\ell} + \binom{n}{\ell+1} = 0 + 0 = 0$$

$\square$

On first reading, this proof is long and confusing, especially in the induction step where we are required to split into four cases. We will give a much simpler proof in Section 3.3 (see Theorem 3.3.40), where we prove the statement *combinatorially* by putting the elements of two sets in one-to-one correspondence.

We can use Theorem 3.1.32 to prove useful identities involving binomial coefficients.

Example 3.1.33

Let $n, k, \ell \in \mathbb{N}$ with $\ell \leqslant k \leqslant n$ then

$$\binom{n}{k}\binom{k}{\ell} = \binom{n}{\ell}\binom{n-\ell}{k-\ell}$$

Indeed:

$$
\begin{aligned}
&\binom{n}{k}\binom{k}{\ell} \\
&= \frac{n!}{k!(n-k)!} \cdot \frac{k!}{\ell!(k-\ell!)} && \text{by Theorem 3.1.32} \\
&= \frac{n!k!}{k!\ell!(n-k)!(k-\ell)!} && \text{combine fractions} \\
&= \frac{n!}{\ell!(n-k)!(k-\ell)!} && \text{cancel } k! \\
&= \frac{n!(n-\ell)!}{\ell!(n-k)!(k-\ell)!(n-\ell)!} && \text{multiply by } \frac{(n-\ell)!}{(n-\ell)!} \\
&= \frac{n!}{\ell!(n-\ell!)} \cdot \frac{(n-\ell)!}{(k-\ell)!(n-k)!} && \text{separate fractions} \\
&= \frac{n!}{\ell!(n-\ell!)} \cdot \frac{(n-\ell)!}{(k-\ell)!((n-\ell)-(k-\ell))!} && \text{rearranging} \\
&= \binom{n}{\ell}\binom{n-\ell}{k-\ell} && \text{by Theorem 3.1.32}
\end{aligned}
$$

◁

Exercise 3.1.34

Prove that $\binom{n}{k} = \binom{n}{n-k}$ for all $n, k \in \mathbb{N}$ with $k \leqslant n$.   ◁

A very useful application of binomial coefficients in elementary algebra is to the binomial theorem.

Theorem 3.1.35 (Binomial theorem)

Let $n \in \mathbb{N}$ and $x, y \in \mathbb{R}$. Then

$$(x+y)^n = \sum_{k=0}^{n} \binom{n}{k} x^k y^{n-k}$$

Proof

In the case when $y = 0$ we have $y^{n-k} = 0$ for all $k < n$, and so the equation reduces to

$$x^n = x^n y^{n-n}$$

which is true, since $y^0 = 1$. So for the rest of the proof, we will assume that $y \neq 0$.

We will now reduce to the case when $y = 1$; and extend to arbitrary $y \neq 0$ afterwards.

We prove $(1+x)^n = \sum_{k=0}^{n} \binom{n}{k} x^k$ by induction on n.

- (**Base case**) $(1+x)^0 = 1$ and $\binom{0}{0} x^0 = 1 \cdot 1 = 1$, so the statement is true when $n = 0$.

- (**Induction step**) Fix $n \in \mathbb{N}$ and suppose that

$$(1+x)^n = \sum_{k=0}^{n} \binom{n}{k} x^k$$

We need to show that $(1+x)^{n+1} = \sum_{k=0}^{n+1} \binom{n+1}{k} x^k$. Well,

$$
\begin{aligned}
&(1+x)^{n+1} \\
&= (1+x)(1+x)^n &&\text{by laws of indices} \\
&= (1+x) \cdot \sum_{k=0}^{n} \binom{n}{k} x^k &&\text{by induction hypothesis} \\
&= \sum_{k=0}^{n} \binom{n}{k} x^k + x \cdot \sum_{k=0}^{n} \binom{n}{k} x^k &&\text{by expanding } (x+1) \\
&= \sum_{k=0}^{n} \binom{n}{k} x^k + \sum_{k=0}^{n} \binom{n}{k} x^{k+1} &&\text{distributing } x \\
&= \sum_{k=0}^{n} \binom{n}{k} x^k + \sum_{k=1}^{n+1} \binom{n}{k-1} x^k &&k \to k-1 \text{ in second sum} \\
&= \binom{n}{0} x^0 + \sum_{k=1}^{n} \left(\binom{n}{k} + \binom{n}{k-1} \right) x^k + \binom{n}{n} x^{n+1} &&\text{splitting the sums} \\
&= \binom{n}{0} x^0 + \sum_{k=1}^{n} \binom{n+1}{k} x^k + \binom{n}{n} x^{n+1} &&\text{by Definition 3.1.28} \\
&= \binom{n+1}{0} x^0 + \sum_{k=1}^{n} \binom{n+1}{k} x^k + \binom{n+1}{n+1} x^{n+1} &&\text{see } (\ast) \text{ below} \\
&= \sum_{k=0}^{n+1} \binom{n+1}{k} x^k
\end{aligned}
$$

The step labelled $(\ast)$ holds because

$$\binom{n}{0} = 1 = \binom{n+1}{0} \qquad \text{and} \qquad \binom{n}{n} = 1 = \binom{n+1}{n+1}$$

By induction, we've shown that $(1+x)^n = \sum_{i=0}^{n} \binom{n}{k} x^k$ for all $n \in \mathbb{N}$.

When $y \neq 0$ is not necessarily equal to 1, we have that

$$(x+y)^n = y^n \cdot \left(1+\frac{x}{y}\right)^n = y^n \cdot \sum_{k=0}^{n} \binom{n}{k}\left(\frac{x}{y}\right)^k = \sum_{k=0}^{n} \binom{n}{k} x^k y^{n-k}$$

The middle equation follows by what we just proved; the leftmost and rightmost equations are simple algebraic rearrangements. $\square$

Example 3.1.36
In Example 3.1.30 we saw that

$$\sum_{k=0}^{n} \binom{n}{k} = 2^n$$

This follows quickly from the binomial theorem, since

$$2^n = (1+1)^n = \sum_{k=0}^{n} \binom{n}{k} \cdot 1^k \cdot 1^{n-k} = \sum_{k=0}^{n} \binom{n}{k}$$

Likewise, in Exercise 3.1.31 you proved that the alternating sum of binomial coefficients is zero; that is, for $n \in \mathbb{N}$, we have

$$\sum_{k=0}^{n} (-1)^k \binom{n}{k} = 0$$

The proof is greatly simplified by applying the binomial theorem. Indeed, by the binomial theorem, we have

$$0 = 0^n = (-1+1)^n = \sum_{k=0}^{n} \binom{n}{k} (-1)^k 1^{n-k} = \sum_{k=0}^{n} (-1)^k \binom{n}{k}$$

Both of these identities can be proved much more elegantly, quickly and easily using *enumerative combinatorics*. This will be the topic covered in Section 3.3.
$\triangleleft$

Strong induction

Consider the following example, which we will attempt to prove by induction.

Example 3.1.37
Define a sequence recursively by

$$b_0 = 1 \quad \text{and} \quad b_{n+1} = 1 + \sum_{k=0}^{n} b_k \quad \text{for all } n \in \mathbb{N}$$

We will attempt to prove by induction that $b_n = 2^n$ for all $n \in \mathbb{N}$.

- (**Base case**) By definition of the sequence we have $b_0 = 1 = 2^0$. So far so good.

- (**Induction step**) Fix $n \in \mathbb{N}$, and suppose that $b_n = 2^n$. We need to show that $b_{n+1} = 2^{n+1}$.

  Well, $b_{n+1} = 1 + \sum_{k=0}^{n} b_k = \ldots$ uh oh.

Here's what went wrong. If we could replace each b_k by 2^k in the sum, then we'd be able to complete the proof. However we cannot justify this substitution: our induction hypothesis only gives us information about b_n, not about a general term b_k for $k < n$. ◁

The *strong* induction principle looks much like the weak induction principle, except that its induction hypothesis is more powerful. Despite its name, strong induction is no stronger than weak induction; the two principles are equivalent. In fact, we'll prove the strong induction principle *by weak induction*!

Theorem 3.1.38 (Strong induction principle)
Let $p(x)$ be a statement about natural numbers and let $n_0 \in \mathbb{N}$. If

  (i) $p(n_0)$ is true; and

  (ii) For all $n \in \mathbb{N}$, if $p(k)$ is true for all $n_0 \leqslant k \leqslant n$, then $p(n+1)$ is true;

then $p(n)$ is true for all $n \geqslant n_0$.

Proof
For each $n \geqslant n_0$, let $q(n)$ be the assertion that $p(k)$ is true for all $n_0 \leqslant k \leqslant n$.

Notice that $q(n)$ implies $p(n)$ for all $n \geqslant n_0$, since given $n \geqslant n_0$, if $p(k)$ is true for all $n_0 \leqslant k \leqslant n$, then in particular $p(k)$ is true when $k = n$.

So it suffices to prove $q(n)$ is true for all $n \geqslant n_0$. We do so by weak induction.

- (**Base case**) $q(n_0)$ is equivalent to $p(n_0)$, since the only natural number k with $n_0 \leqslant k \leqslant n_0$ is n_0 itself; hence $q(n_0)$ is true by condition (i).

- (**Induction step**) Let $n \geqslant n_0$ and suppose $q(n)$ is true. Then $p(k)$ is true for all $n_0 \leqslant k \leqslant n$.

  We need to prove that $q(n+1)$ is true—that is, that $p(k)$ is true for all $n_0 \leqslant k \leqslant n+1$. But we know $p(k)$ is true for all $n_0 \leqslant k \leqslant n$—this is the induction hypothesis—and then $p(n+1)$ is true by condition (ii). So we have that $p(k)$ is true for all $n_0 \leqslant k \leqslant n+1$ after all.

By induction, $q(n)$ is true for all $n \geqslant n_0$. Hence $p(n)$ is true for all $n \geqslant n_0$. $\qquad\square$

Strategy 3.1.39 (Proof by strong induction)
In order to prove a proposition of the form $\forall n \geqslant n_0,\ p(n)$, it suffices to prove that $p(n_0)$ is true and that, for all $n \geqslant n_0$, if $p(k)$ is true for all $n_0 \leqslant k \leqslant n$, then $p(n+1)$ is true. $\qquad\triangleleft$

Like with weak induction, we can illustrate how strong induction works diagrammatically. The induction hypothesis, represented by the large square, now encompasses $p(k)$ for all $n_0 \leqslant k \leqslant n$, where $p(n_0)$ is the base case.

Observe that the only difference from weak induction is the induction hypothesis.

- **Weak induction step:** Fix $n \geqslant n_0$, $\boxed{\text{assume } p(n) \text{ is true}}$, derive $p(n+1)$;

- **Strong induction step:** Fix $n \geqslant n_0$, $\boxed{\text{assume } p(k) \text{ is true for all } n_0 \leqslant k \leqslant n}$, derive $p(n+1)$.

We now use strong induction to complete the proof of Example 3.1.37.

Example 3.1.40 (Example 3.1.37 revisited)
Define a sequence recursively by

$$b_0 = 1 \quad \text{and} \quad b_{n+1} = 1 + \sum_{k=0}^{n} b_k \text{ for all } n \in \mathbb{N}$$

We will prove by strong induction that $b_n = 2^n$ for all $n \in \mathbb{N}$.

- (**Base case**) By definition of the sequence we have $b_0 = 1 = 2^0$.

- (**Induction step**) Fix $n \in \mathbb{N}$, and suppose that $b_k = 2^k$ for all $k \leqslant n$. We need to

show that $b_{n+1} = 2^{n+1}$. This is true, since

$$b_{n+1} = 1 + \sum_{k=0}^{n} b_k \qquad \text{by the recursive formula for } b_{n+1}$$

$$= 1 + \sum_{k=0}^{n} 2^k \qquad \text{by the induction hypothesis}$$

$$= 1 + (2^{n+1} - 1) \qquad \text{by Exercise 3.1.18}$$

$$= 2^{n+1}$$

By induction, it follows that $b_n = 2^n$ for all $n \in \mathbb{N}$. ◁

The following theorem adapts the strong induction principle to proofs where we need to refer to a *fixed* number of previous steps in our induction step.

Theorem 3.1.41 (Strong induction principle (multiple base cases))
Let $p(n)$ be a logical formula with free variable $n \in \mathbb{N}$ and let $n_0 < n_1 \in \mathbb{N}$. If

(i) $p(n_0), p(n_0 + 1), \ldots, p(n_1)$ are all true; and

(ii) For all $n \geqslant n_1$, if $p(k)$ is true for all $n_0 \leqslant k \leqslant n$, then $p(n+1)$ is true;

then $p(n)$ is true for all $n \geqslant n_0$.

Proof
The fact that $p(n)$ is true for all $n \geqslant n_1$ follows from strong induction. Indeed:

- $p(n_1)$ is true by (i);

- Fix $n \geqslant n_1$ and assume $p(k)$ is true for all $n_1 \leqslant k \leqslant n$. Then in fact $p(k)$ is true for all $n_0 \leqslant k \leqslant n$, since $p(n_0), p(n_0 + 1), \ldots p(n_1 - 1)$ are true by (i). So $p(n+1)$ is true by (ii)

So $p(n)$ is true for all $n \geqslant n_1$. But then $p(n)$ is true for all $n \geqslant n_0$, again by (i). □


Strategy 3.1.42 (Proof by strong induction with multiple base cases)
In order to prove a statement of the form $\forall n \geqslant n_0,\ p(n)$, it suffices to prove $p(k)$ for all $k \in \{n_0, n_0 + 1, \ldots, n_1\}$, where $n_1 > n_0$, and then given $n \geqslant n_1$, assuming $p(k)$ is true for all $n_0 \leqslant k \leqslant n$, prove that $p(n+1)$ is true. ◁

This kind of strong induction differs from the usual kind in two main ways:

- There are multiple base cases $p(n_0), p(n_0 + 1), \ldots, p(n_1)$, not just one.

- The induction step refers to both the least base case n_0 and the greatest base case n_1: the variable n in the induction step is taken to be greater than or equal to n_1, while the induction hypothesis assumes $p(k)$ for all $n_0 \leqslant k \leqslant n$.

The following diagram illustrates how strong induction with multiple base cases works.

Note the difference in quantification of variables in the induction step between regular strong induction and strong induction with multiple base cases:

- **One base case.** Fix $n \geqslant \boxed{n_0}$ and assume $p(k)$ is true for all $\boxed{n_0} \leqslant k \leqslant n$.

- **Multiple base cases.** Fix $n \geqslant \boxed{n_1}$ and assume $p(k)$ is true for all $\boxed{n_0} \leqslant k \leqslant n$.

Getting the quantification of the variables n and k in the strong induction step is crucial, since the quantification affects what may be assumed about n and k.

The need for multiple base cases often arises when proving results about recursively defined sequences, where the definition of a general term depends on the values of a fixed number of previous terms.

Example 3.1.43
Define the sequence

$$a_0 = 0, \quad a_1 = 1, \quad a_n = 3a_{n-1} - 2a_{n-2} \text{ for all } n \geqslant 2$$

We find and prove a general formula for a_n in terms of n. Writing out the first few terms in the sequence establishes a pattern that we might attempt to prove:

n	0	1	2	3	4	5	6	7	8
a_n	0	1	3	7	15	31	63	127	255

It appears that $a_n = 2^n - 1$ for all $n \geqslant 0$. We prove this by strong induction, taking the cases $n = 0$ and $n = 1$ as our base cases.

- (**Base cases**) By definition of the sequence, we have:

⋄ $a_0 = 0 = 2^0 - 1$; and

⋄ $a_1 = 1 = 2^1 - 1$;

so the claim is true when $n = 0$ and $n = 1$.

- (**Induction step**) Fix $n \geqslant 1$ and assume that $a_k = 2^k - 1$ for all $0 \leqslant k \leqslant n$. We need to prove that $a_{n+1} = 2^{n+1} - 1$.

  Well since $n \geqslant 1$, we have $n + 1 \geqslant 2$, so we can apply the recursive formula to a_{n+1}. Thus

$$
\begin{aligned}
a_{n+1} &= 3a_n - 2a_{n-1} && \text{by definition of } a_{n+1} \\
&= 3(2^n - 1) - 2(2^{n-1} - 1) && \text{by induction hypothesis} \\
&= 3 \cdot 2^n - 3 - 2 \cdot 2^{n-1} + 2 && \text{expanding} \\
&= 3 \cdot 2^n - 3 - 2^n + 2 && \text{using laws of indices} \\
&= 2 \cdot 2^n - 1 && \text{simplifying} \\
&= 2^{n+1} - 1 && \text{using laws of indices}
\end{aligned}
$$

So the result follows by induction. ◁

The following exercises have proofs by strong induction with multiple base cases.

Exercise 3.1.44

Define a sequence recursively by $a_0 = 4$, $a_1 = 9$ and $a_n = 5a_{n-1} - 6a_{n-2}$ for all $n \geqslant 2$. Prove that $a_n = 3 \cdot 2^n + 3^n$ for all $n \in \mathbb{N}$. ◁

Exercise 3.1.45

The *Tribonacci sequence* is the sequence $t_0, t_1, t_2, \ldots$ defined by

$$
t_0 = 0, \quad t_1 = 0, \quad t_2 = 1, \quad t_n = t_{n-1} + t_{n-2} + t_{n-3} \text{ for all } n \geqslant 3
$$

Prove that $t_n \leqslant 2^{n-3}$ for all $n \geqslant 3$. ◁

Exercise 3.1.46

The *Frobenius coin problem* asks when a given amount of money can be obtained from coins of given denominations. For example, a value of 7 dubloons cannot be obtained using only 3 dubloon and 5 dubloon coins, but for all $n \geqslant 8$, a value of n dubloons *can* be obtained using only 3 dubloon and 5 dubloon coins. Prove this. ◁

Well-ordering principle

In a way that we will make precise in Section 5.2, the underlying reason why we can perform induction and recursion on the natural numbers is because of the

way they are ordered. Specifically, the natural numbers satisfy the *well-ordering principle*: if a set of natural numbers has at least one element, then it has a least element. This sets the natural numbers apart from the other number sets; for example, $\mathbb{Z}$ has no least element, since if $a \in \mathbb{Z}$ then $a - 1 \in \mathbb{Z}$ and $a - 1 < a$.

Theorem 3.1.47 (Well-ordering principle)
Let X be a set of natural numbers. If X is inhabited, then X has a least element.

Idea of proof
Under the assumption that X is a set of natural numbers, the proposition we want to prove has the form $p \Rightarrow q$. Namely

$$X \text{ is inhabited} \quad \Rightarrow \quad X \text{ has a least element}$$

Assuming X is inhabited doesn't really give us much to work with, so let's try the contrapositive:

$$X \text{ has no least element} \quad \Rightarrow \quad X \text{ is empty}$$

The assumption that X has no least element *does* give us something to work with. Under this assumption we need to deduce that X is empty.

We will do this by 'forcing X up' by strong induction. Certainly $0 \notin X$, otherwise it would be the least element. If none of the numbers $0, 1, \ldots, n$ are elements of X then neither can $n + 1$ be, since if it were then *it* would be the least element of X. Let's make this argument formal.   □

Proof
Let X be a set of natural numbers containing no least element. We prove by strong induction that $n \notin X$ for all $n \in \mathbb{N}$.

- **(BC)** $0 \notin X$ since if $0 \in X$ then 0 must be the least element of X, as it is the least natural number.

- **(IS)** Suppose $k \notin X$ for all $0 \leqslant k \leqslant n$. If $n + 1 \in X$ then $n + 1$ is the least element of X; indeed, if $\ell < n + 1$ then $0 \leqslant \ell \leqslant n$, so $\ell \notin X$ by the induction hypothesis. This contradicts the assumption that X has no least element, so $n + 1 \notin X$.

By strong induction, $n \notin X$ for each $n \in \mathbb{N}$. Since X is a set of natural numbers, and it contains no natural numbers, it follows that X is empty.   □

Aside
In Section 5.2 we will encounter more general sets with a notion of 'less than', for which any inhabited subset has a 'least' element. Any such set has an induction principle, the proof of which is more or less identical to the proof of

Theorem 3.1.38. This has powerful applications in computer science, where it can be used to formally verify that a computer program containing various loops will terminate: termination of a program corresponds to a particular set having a 'least' element. ◁

The following proof that $\sqrt{2}$ is irrational is a classic application of the well-ordering principle.

Proposition 3.1.48
The number $\sqrt{2}$ is irrational.

To prove Proposition 3.1.48 we will use the following lemma, which uses the well-ordering principle to prove that fractions can be 'cancelled to lowest terms'.

Lemma 3.1.49
Let q be a positive rational number. There is a pair of nonzero natural numbers a, b such that $q = \frac{a}{b}$ and such that the only natural number which divides both a and b is 1.

Proof
First note that we can express q as the ratio of two nonzero natural numbers, since q is a positive rational number. By the well-ordering principle, there is a *least* natural number a such that $q = \frac{a}{b}$ for some positive $b \in \mathbb{N}$.

Suppose that some natural number d other than 1 divides both a and b. Note that $d \neq 0$, since if $d = 0$ then that would imply $a = 0$. Since $d \neq 1$, it follows that $d \geqslant 2$.

Since d divides a and b, there exist natural numbers a', b' such that $a = a'd$ and $b = b'd$. Moreover, $a', b' > 0$ since $a, b, d > 0$. It follows that

$$q = \frac{a}{b} = \frac{a'd}{b'd} = \frac{a'}{b'}$$

But $d \geqslant 2$, and hence

$$a' = \frac{a}{d} \leqslant \frac{a}{2} < a$$

contradicting minimality of a. Hence our assumption that some natural number d other than 1 divides both a and b was false—it follows that the only natural number dividing both a and b is 1. □

We are now ready to prove that $\sqrt{2}$ is irrational.

Proof of Proposition 3.1.48
Suppose $\sqrt{2}$ is rational. Since $\sqrt{2} > 0$, this means that we can write

$$\sqrt{2} = \frac{a}{b}$$

where a and b are both positive natural numbers. By Lemma 3.1.49, we may assume that the only natural number dividing a and b is 1.

Multiplying the equation $\sqrt{2} = \frac{a}{b}$ and squaring yields

$$a^2 = 2b^2$$

Hence a^2 is even. By Proposition 1.1.46, a is even, so we can write $a = 2c$ for some $c > 0$. Then $a^2 = (2c)^2 = 4c^2$, and hence

$$4c^2 = 2b^2$$

Dividing by 2 yields

$$2c^2 = b^2$$

and hence b^2 is even. By Proposition 1.1.46 again, b is even.

But if a and b are both even, the natural number 2 divides both a and b. This contradicts the fact that the only natural number dividing both a and b is 1. Hence our assumption that $\sqrt{2}$ is rational is incorrect, and $\sqrt{2}$ is irrational. □

Writing tip

In the proof of Proposition 3.1.48 we could have separately proved that a^2 being even implies a is even, and that b^2 being even implies b is even. This would have required us to repeat the same proof twice, which is somewhat tedious! Proving auxiliary results separately (as in Lemma 3.1.49) and then quoting them in later theorems can improve the readability of the main proof, particularly when the auxiliary results are particularly technical. Doing so also helps emphasise the important steps.   ◁

Exercise 3.1.50

What goes wrong in the proof of Proposition 3.1.48 if we try instead to prove that $\sqrt{4}$ is irrational?   ◁

Exercise 3.1.51

Prove that $\sqrt{3}$ is irrational.   ◁

Section 3.2
Finite sets

As its title suggests, this section is all about exploring the properties of finite sets, and to do this we must first define what we mean by 'finite'. We certainly know a finite set when we see one—for example:

- The set $\{\text{red}, \text{orange}, \text{yellow}, \text{green}, \text{blue}, \text{purple}\}$ is finite.

- The set $[0, 1]$ is infinite, but it has finite length.

- The set $[0, \infty)$ is infinite and has infinite length.

- The set $\mathscr{P}(\mathbb{N})$ is infinite, but has no notion of 'length' to speak of.

- The empty set $\varnothing$ is finite.

If we are to make a definition of 'finite set', we must first figure out what the finite sets above have in common but the infinite sets do not.

It is difficult to define 'finite' without being imprecise. A first attempt at a definition might be something like the following:

> *A set X is finite if the elements of X don't go on forever.*

This is good intuition, but isn't good enough as a mathematical definition, because 'go on' and 'forever' are not precise terms (unless they themselves are defined). So let's try to make this more precise:

> *A set X is finite if the elements of X can be listed one by one*
> *in such a way that the list has both a start and an end.*

This is better but is still not entirely precise—it is not entirely clear what is meant by 'listed one by one'. But we can make this precise: to list the elements of X one-by-one means that we are specifying a 'first element', a 'second element', a 'third element', and so on. To say that this list has an end means that we eventually reach the 'n^{th} element', for some $n \in \mathbb{N}$, and there is no '$(n+1)^{\text{st}}$ element'. In other words, for some natural number n, we are pairing up the elements of X with the natural numbers from 1 to n.

Recall that, for each $n \in \mathbb{N}$, the set of natural numbers from 1 up to n has its own notation:

> **Definition 2.1.9**
> Let $n \in \mathbb{N}$. The set $[n]$ is defined by $[n] = \{k \in \mathbb{N} \mid 1 \leqslant k \leqslant n\}$.

Since 'pairing up' really means 'finding a bijection', we are now ready to define what it means for a set to be finite.

> **Definition 3.2.1**
> A set X is **finite** if there exists a bijection $f : [n] \to X$ for some $n \in \mathbb{N}$. The function f is called an **enumeration** of X. If X is not finite we say it is **infinite**.

This definition suggests the following strategy for proving that a set is finite.

> **Strategy 3.2.2** (Proving that a set is finite)
> In order to prove that a set X is finite, it suffices to find a bijection $[n] \to X$ for some $n \in \mathbb{N}$. ◁

Example 3.2.3
Let $X = \{\text{red}, \text{orange}, \text{yellow}, \text{green}, \text{blue}, \text{purple}\}$. We said above that X is finite; now we can prove it. Define $f : [6] \to X$ by

$$f(1) = \text{red} \quad f(2) = \text{orange} \quad f(3) = \text{yellow}$$
$$f(4) = \text{green} \quad f(5) = \text{blue} \quad f(6) = \text{purple}$$

The function f is evidently a bijection, since each element of X can be expressed uniquely as $f(k)$ for some $k \in [6]$. So X is finite. ◁

Exercise 3.2.4
Prove that $[n]$ is finite for each $n \in \mathbb{N}$. ◁

Note that Exercise 3.2.4 implies, in particular, that $\varnothing$ is finite, since $\varnothing = [0]$.

The size of a finite set

Whilst it might sometimes be useful just to know *that* set is finite, it will be even more useful to know how many elements it has. This quantity is called the *size* of the set. Intuitively, the size of the set should be the length of the list of its elements, but for this to be well-defined, we first need to know that the number of elements in the list is independent of the order in which we list them.

The 'list of elements' of a finite set X is the bijection $[n] \to X$ given by Definition 3.2.1, and n is the length of the list, this means that we need to prove that if $[m] \to X$ and $[n] \to X$ are bijections, then $m = n$. This will be Theorem 3.2.8.

To be able to prove this, we must first prove some technical results that we will use in the proof.

Lemma 3.2.5

Let X be an inhabited set. There is a bijection $X \setminus \{a\} \to X \setminus \{b\}$ for all $a, b \in X$.

Proof

Let $a, b \in X$. First note that if $a = b$ then $X \setminus \{a\} = X \setminus \{b\}$, and so the identity function $\mathrm{id}_{X \setminus \{a\}}$ is the desired bijection.

So assume $a \neq b$, and define $f : X \setminus \{a\} \to X \setminus \{b\}$ by

$$f(x) = \begin{cases} a & \text{if } x = b \\ x & \text{otherwise} \end{cases}$$

Note that f is well-defined since it ensures that $f(x) \neq b$ for any $x \in X \setminus \{a\}$.

We prove that f is a bijection by finding an inverse.

So define $g : X \setminus \{b\} \to X \setminus \{a\}$ by

$$g(x) = \begin{cases} b & \text{if } x = a \\ x & \text{otherwise} \end{cases}$$

Again, g is well-defined since we have ensured that $g(x) \neq a$ for any $x \in X \setminus \{b\}$.

Given $x \in X$, if $x \neq a$ and $x \neq b$, then $f(x) \neq a$ and $g(x) \neq b$, so that

$$g(f(x)) = g(x) = x \quad \text{and} \quad f(g(x)) = f(x) = x$$

Moreover $g(f(b)) = g(a) = b$ and $f(g(a)) = f(b) = a$.

This proves that $g \circ f = \mathrm{id}_{X \setminus \{a\}}$ and $f \circ g = \mathrm{id}_{X \setminus \{b\}}$, so that g is an inverse for f, as required. □

Theorem 3.2.6

Let $m, n \in \mathbb{N}$.

(a) If there exists an injection $f : [m] \to [n]$, then $m \leqslant n$.

(b) If there exists a surjection $g : [m] \to [n]$, then $m \geqslant n$.

(c) If there exists a bijection $h : [m] \to [n]$, then $m = n$.

Proof of (a)

For fixed $m \in \mathbb{N}$, let $p(m)$ be the assertion that, for all $n \in \mathbb{N}$, if there exists an injection $[m] \to [n]$, then $m \leqslant n$. We prove that $p(m)$ is true for all $m \in \mathbb{N}$ by induction.

- (**Base case**) We need to prove that, for all $n \in \mathbb{N}$ if there exists an injection $[0] \to [n]$, then $0 \leqslant n$. This is automatically true, since $0 \leqslant n$ for all $n \in \mathbb{N}$.

- (**Induction step**) Fix $m \in \mathbb{N}$ and suppose that, for all $n \in \mathbb{N}$, if there exists an injection $[m] \to [n]$, then $m \leqslant n$.

  Now let $n \in \mathbb{N}$ and suppose that there is an injection $f : [m+1] \to [n]$. We need to prove that $m+1 \leqslant n$.

  First note that $n \geqslant 1$. Indeed, since $m+1 \geqslant 1$, we have $1 \in [m+1]$, and so $f(1) \in [n]$. This means that $[n]$ is inhabited, and so $n \geqslant 1$. In particular, $n-1 \in \mathbb{N}$ and so the set $[n-1]$ is well-defined. It suffices to prove that $m \leqslant n-1$.

  Let $a = f(m+1) \in [n]$ and define $f^- : [m] \to [n] \setminus \{a\}$ by $f^-(k) = f(k)$ for all $k \in [m]$. Note that f^- is well-defined; indeed, $f(k) \neq a$ for all $k \in [m]$ since $a = f(m+1)$ and f is injective.

  The function f^- is injective. To see this, let $k, \ell \in [m]$ and suppose $f^-(k) = f^-(\ell)$. Then $f(k) = f(\ell)$ by definition of f^-, and so $k = \ell$ by injectivity of f.

  Since $[n-1] = [n] \setminus \{n\}$, there is a bijection $s : [n] \setminus \{a\} \to [n-1]$ by Lemma 3.2.5. In particular, s is injective, and so $s \circ f^-$ is an injection $[m] \to [n-1]$ by Proposition 2.3.4.

  By the induction hypothesis, we have $m \leqslant n-1$, and so $m+1 \leqslant n$ as required.

The result now follows by induction.   □

Exercise 3.2.7
Prove parts (b) and (c) of Theorem 3.2.6.   ◁

Phew! That was fun. With these technical results proved, we can now prove the theorem we needed for the size of a finite set to be well-defined.

Theorem 3.2.8 (Uniqueness of size)
Let X be a finite set and let $f : [m] \to X$ and $g : [n] \to X$ be enumerations of X, where $m, n \in \mathbb{N}$. Then $m = n$.

Proof
Since $f : [m] \to X$ and $g : [n] \to X$ are bijections, the function $g^{-1} \circ f : [m] \to [n]$ is a bijection by Exercises 2.3.20 and 2.3.45. Hence $m = n$ by Theorem 3.2.6(c).   □

As we mentioned above, Theorem 3.2.8 tells us that any two ways of listing (enumerating) the elements of a finite set yield the same number of elements. We may now make the following definition.

Definition 3.2.9
Let X be a finite set. The **size** of X, written $|X|$, is the unique natural number n for which there exists a bijection $[n] \to X$.

Example 3.2.10
Example 3.2.3 showed that $|\{\text{red}, \text{orange}, \text{yellow}, \text{green}, \text{blue}, \text{purple}\}| = 6$, and provided the proof was correct, Exercise 3.2.4 showed that $|[n]| = n$ for all $n \in \mathbb{N}$; in particular, $|\varnothing| = 0$. ◁

Example 3.2.11
Fix $n \in \mathbb{N}$ and let $X = \{a \in \mathbb{Z} \mid -n \leqslant a \leqslant n\}$. There is a bijection $f : [2n+1] \to X$ defined by $f(k) = k - n - 1$. Indeed:

- f **is well-defined.** Given $k \in [2n+1]$, we have $1 \leqslant k \leqslant 2n+1$, and so

$$-n = 1 - (n+1) \leqslant \underbrace{k - (n+1)}_{=f(k)} \leqslant (2n+1) - (n+1) = n$$

  so that $f(k) \in X$ as claimed.

- f **is injective.** Let $k, \ell \in [2n+1]$ and assume $f(k) = f(\ell)$. Then $k - n - 1 = \ell - n - 1$, and so $k = \ell$.

- f **is surjective.** Let $a \in X$ and define $k = a + n + 1$. Then

$$1 = (-n) + n + 1 \leqslant \underbrace{a + n + 1}_{=k} \leqslant n + n + 1 = 2n + 1$$

  and so $k \in [2n+1]$, and moreover $f(k) = (a + n + 1) - n - 1 = a$.

Since f is a bijection, we have $|X| = 2n + 1$ by Definition 3.2.9. ◁

Exercise 3.2.12
Let X be a finite set with $|X| = n > 1$. Let $a \in X$ and let $b \notin X$. Prove that

(a) $X \setminus \{a\}$ is finite and $|X \setminus \{a\}| = n - 1$; and

(b) $X \cup \{b\}$ is finite and $|X \cup \{b\}| = n + 1$.

Identify where in your proofs you make use of the hypotheses that $a \in X$ and $b \notin X$. ◁

Comparing the sizes of finite sets

When we used dots and stars to motivate the definitions of injective and surjective functions at the beginning of Section 2.3, we suggested the following intuition:

- If there is an injection $f : X \to Y$, then X has 'at most as many elements as Y'; and

- If there is a surjection $g : X \to Y$, then X has 'at least as many elements as Y'.

We are now in a position to prove this, at least when X and Y are finite. The following theorem is a generalisation of Theorem 3.2.6.

Theorem 3.2.13

Let X and Y be sets.

(a) If Y is finite and there is an injection $f : X \to Y$, then X is finite and $|X| \leqslant |Y|$.

(b) If X is finite and there is a surjection $f : X \to Y$, then Y is finite and $|X| \geqslant |Y|$.

(c) If one of X or Y is finite and there is a bijection $f : X \to Y$, then X and Y are both finite and $|X| = |Y|$.

Proof of (a)

We prove by induction that, for all $n \in \mathbb{N}$, if Y is a finite set of size n and there is an injection $f : X \to Y$, then X is finite and $|X| \leqslant n$.

- (**Base case**) Let Y be a finite set of size 0—that is, Y is empty. Suppose there is an injection $f : X \to Y$. If X is inhabited, then there exists an element $a \in X$, so that $f(a) \in Y$. This contradicts emptiness of Y, so that X must be empty. Hence $|X| = 0 \leqslant 0$, as required.

- (**Induction step**) Fix $n \in \mathbb{N}$ and assume that, if Y is a finite set of size n and there is an injection $f : X \to Y$, then X is finite and $|X| \leqslant n$.

  Fix a finite set Y of size $n+1$ and an injection $f : X \to Y$. We need to prove that X is finite and $|X| \leqslant n+1$.

  If X is empty, then $|X| = 0 \leqslant n+1$ as required. So assume that X is inhabited, and fix an element $a \in X$.

  Define $f^{\vee} : X \setminus \{a\} \to Y \setminus \{f(a)\}$ by $f^{\vee}(x) = f(x)$ for all $x \in X \setminus \{a\}$. Note that $f^{\vee}$ is well-defined since $f(x) \neq f(a)$ for any $x \in X \setminus \{a\}$ by injectivity of f. Moreover $f^{\vee}$ is injective; indeed, let $x, y \in X \setminus \{a\}$ and assume $f^{\vee}(x) = f^{\vee}(y)$. Then
  $$f(x) = f^{\vee}(x) = f^{\vee}(y) = f(y) \quad \Rightarrow \quad x = y$$
  by injectivity of f. So $f^{\vee}$ is an injection.

  By Exercise 3.2.12, $Y \setminus \{f(a)\}$ is finite and $|Y \setminus \{f(a)\}| = (n+1) - 1 = n$.

By the induction hypothesis, $X \setminus \{a\}$ is finite and $|X \setminus \{a\}| \leqslant (n+1) - 1$. But $|X \setminus \{a\}| = |X| - 1$ by Exercise 3.2.12, and so $|X| \leqslant n+1$, as required.

The result now follows by induction. □

Exercise 3.2.14
Prove parts (b) and (c) of Theorem 3.2.13. ◁

Theorem 3.2.13 suggests the following strategies for comparing the sizes of finite sets:

Strategy 3.2.15 (Comparing the sizes of finite sets)
Let X and Y be finite sets.

(a) In order to prove that $|X| \leqslant |Y|$, it suffices to find an injection $X \rightarrow Y$.

(b) In order to prove that $|X| \geqslant |Y|$, it suffices to find a surjection $X \rightarrow Y$.

(c) In order to prove that $|X| = |Y|$, it suffices to find a bijection $X \rightarrow Y$.

Strategy (c) is commonly known as **bijective proof**. ◁

Closure properties of finite sets

We now use Strategy 3.2.15 to prove some *closure properties* of finite sets—that is, operations we can perform on finite sets to ensure that the result remains finite.

Exercise 3.2.16
Let X be a finite set. Prove that every subset $U \subseteq X$ is finite and $|U| \leqslant |X|$. ◁

Exercise 3.2.17
Let X and Y be finite sets. Prove that $X \cap Y$ is finite. ◁

Proposition 3.2.18
Let X and Y be finite sets. Then $X \cup Y$ is finite, and moreover

$$|X \cup Y| = |X| + |Y| - |X \cap Y|$$

Proof
We will prove this in the case when X and Y are disjoint. The general case, when X and Y are not assumed to be disjoint, will be Exercise 3.2.19.

Let $m = |X|$ and $n = |Y|$, and let $f : [m] \rightarrow X$ and $g : [n] \rightarrow Y$ be bijections.

Since X and Y are disjoint, we have $X \cap Y = \varnothing$. Define $h : [m+n] \rightarrow X \cup Y$ as

follows; given $k \in [m+n]$, let

$$h(k) = \begin{cases} f(k) & \text{if } k \leqslant m \\ g(k-m) & \text{if } k > m \end{cases}$$

Note that h is well-defined: the cases $k \leqslant m$ and $k > m$ are mutually exclusive, they cover all possible cases, and $k - m \in [n]$ for all $m+1 \leqslant k \leqslant n$ so that $g(k-m)$ is defined. It is then straightforward to check that h has an inverse $h^{-1} : X \cup Y \to [m+n]$ defined by

$$h^{-1}(z) = \begin{cases} f^{-1}(z) & \text{if } z \in X \\ g^{-1}(z)+m & \text{if } z \in Y \end{cases}$$

Well-definedness of h^{-1} relies fundamentally on the assumption that $X \cap Y = \varnothing$, as this is what ensures that the cases $x \in X$ and $x \in Y$ are mutually exclusive.

Hence $|X \cup Y| = m+n = |X| + |Y|$, which is as required since $|X \cap Y| = 0$.   □

Exercise 3.2.19
The following steps complete the proof of Proposition 3.2.18:

(a) Given sets A and B, prove that the sets $A \times \{0\}$ and $B \times \{1\}$ are disjoint, and find bijections $A \to A \times \{0\}$ and $B \to B \times \{1\}$. Write $A \sqcup B$ (LaTeX code: \sqcup) to denote the set $(A \times \{0\}) \cup (B \times \{1\})$. The set $A \sqcup B$ is called the **disjoint union** of A and B.

(b) Prove that, if A and B are finite then $A \sqcup B$ is finite and

$$|A \sqcup B| = |A| + |B|$$

(c) Let X and Y be sets. Find a bijection

$$(X \cup Y) \sqcup (X \cap Y) \to X \sqcup Y$$

(d) Complete the proof of Proposition 3.2.18—that is, prove that if X and Y are finite sets, not necessarily disjoint, then $X \cup Y$ is finite and

$$|X \cup Y| = |X| + |Y| - |X \cap Y|$$

◁

Exercise 3.2.20
Let X be a finite set and let $U \subseteq X$. Prove that $X \setminus U$ is finite, and moreover $|X \setminus U| = |X| - |U|$.   ◁

Exercise 3.2.21
Let $m, n \in \mathbb{N}$. Prove that $|[m] \times [n]| = mn$.   ◁

Proposition 3.2.22

Let X and Y be finite sets. Then $X \times Y$ is finite, and moreover

$$|X \times Y| = |X| \cdot |Y|$$

Proof

Let X and Y be finite sets, let $m = |X|$ and $n = |Y|$, and let $f : [m] \to X$ and $g : [n] \to Y$ be bijections. Define a function $h : [m] \times [n] \to X \times Y$ by

$$h(k, \ell) = (f(k), g(\ell))$$

for each $k \in [m]$ and $\ell \in [n]$. It is easy to see that this is a bijection, with inverse defined by

$$h^{-1}(x, y) = (f^{-1}(x), g^{-1}(y))$$

for all $x \in X$ and $y \in Y$. By Exercise 3.2.21 there is a bijection $p : [mn] \to [m] \times [n]$, and by Exercise 2.3.20 the composite $h \circ p : [mn] \to X \times Y$ is a bijection. Hence $|X \times Y| = mn$. $\qquad\square$

In summary, we have proved that the property of finiteness is preserved by taking subsets, pairwise unions, pairwise intersections, pairwise cartesian products, and relative complements.

Infinite sets

We conclude this section by proving that not all sets are finite—specifically, we'll prove that $\mathbb{N}$ is infinite. *Intuitively* this seems extremely easy: of *course* $\mathbb{N}$ is infinite! But in mathematical practice, this isn't good enough: we need to use our definition of 'infinite' to prove that $\mathbb{N}$ is infinite. Namely, we need to prove that there is no bijection $[n] \to \mathbb{N}$ for any $n \in \mathbb{N}$. We will use Lemma 3.2.23 below in our proof.

Lemma 3.2.23

Every inhabited finite set of natural numbers has a greatest element.

Proof

We'll prove by induction on $n \geqslant 1$ that every subset $U \subseteq \mathbb{N}$ of size n has a greatest element.

- (**Base case**) Take $U \subseteq \mathbb{N}$ with $|U| = 1$. then $U = \{m\}$ for some $m \in \mathbb{N}$. Since m is the only element of U, it is certainly the greatest element of U!

- (**Induction step**) Fix $n \geqslant 1$ and suppose that every set of natural numbers of size n has a greatest element (**IH**).

Let $U \subseteq \mathbb{N}$ with $|U| = n+1$. We wish to show that U has a greatest element.

Since $|U| = n+1$, we may write $U = \{m_1, m_2, \ldots, m_n, m_{n+1}\}$ for distinct natural numbers m_k. But then $|U \setminus \{m_{n+1}\}| = n$ by Exercise 3.2.12, and so by the induction hypothesis, $U \setminus \{m_{n+1}\}$ has a greatest element, say m_k. Now:

◇ If $m_k < m_{n+1}$, then m_{n+1} is the greatest element of U.

◇ If $m_k > m_{n+1}$, then m_k is the greatest element of U.

In any case, U has a greatest element. This completes the induction step. $\qquad\square$

Theorem 3.2.24
The set $\mathbb{N}$ is infinite.

Proof
We proceed by contradiction. Suppose $\mathbb{N}$ is finite. Then $|\mathbb{N}| = n$ for some $n \in \mathbb{N}$, and hence $\mathbb{N}$ is either empty (nonsense, since $0 \in \mathbb{N}$) or, by Lemma 3.2.23, it has a greatest element g. But $g+1 \in \mathbb{N}$ since every natural number has a successor, and $g+1 > g$, so this contradicts maximality of g. Hence $\mathbb{N}$ is infinite. $\qquad\square$

Counting principles

In Section 3.2 we were interested in establishing conditions under which a set is finite, and proving that we may perform certain operations on finite sets—such as unions and cartesian products—without losing the property of finiteness.

In this section, our attention turns to the task of finding the size of a set that is known to be finite. This process is called *counting* and is at the core of the mathematical field of combinatorics.

Binomials and factorials revisited

We defined binomial coefficients $\binom{n}{k}$ and factorials $n!$ *recursively* in Section 3.1, and proved elementary facts about them by induction. We will now re-define them *combinatorially*—that is, we give them meaning in terms of sizes of particular finite sets. We will prove that the combinatorial and recursive definitions are equivalent, and prove facts about them using combinatorial arguments.

The reasons for doing so are manifold. The combinatorial definitions allow us to reason about binomials and factorials with direct reference to descriptions of finite sets, which will be particularly useful when we prove identities about them using *double counting*. Moreover, the combinatorial definitions remove the seemingly arbitrary nature of the recursive definitions—for example, they provide a reason why it makes sense to define $0! = 1$ and $\binom{0}{0} = 1$.

Definition 3.3.1
Let X be a set and let $k \in \mathbb{N}$. A k-**element subset** of X is a subset $U \subseteq X$ such that $|U| = k$. The set of all k-element subsets of X is denoted $\binom{X}{k}$ (read: 'X choose k') (LaTeX code: `\binom{X}{k}`).

Intuitively, $\binom{X}{k}$ is the set of ways of picking k elements from X, without repetitions, in such a way that order doesn't matter. (If order mattered, the elements would be *sequences* instead of *subsets*.)

Example 3.3.2
We find $\binom{[4]}{k}$ for all $k \in \mathbb{N}$.

- $\binom{[4]}{0} = \{\varnothing\}$ since the only set with 0 elements is $\varnothing$;

- $\binom{[4]}{1} = \{\{1\},\{2\},\{3\},\{4\}\}$;

- $\binom{[4]}{2} = \{\{1,2\},\{1,3\},\{1,4\},\{2,3\},\{2,4\},\{3,4\}\}$;

- $\binom{[4]}{3} = \{\{1,2,3\},\{1,2,4\},\{1,3,4\},\{2,3,4\}\}$;

- $\binom{[4]}{4} = \{\{1,2,3,4\}\}$;

- If $k \geqslant 5$ then $\binom{[4]}{k} = \varnothing$, since by Exercise 3.2.16, no subset of $[4]$ can have more than 4 elements.

$\triangleleft$

Proposition 3.3.3
If X is a finite set, then $\mathscr{P}(X) = \bigcup_{k \leqslant |X|} \binom{X}{k}$.

Proof
Let $U \subseteq X$. By Exercise 3.2.16, U is finite and $|U| \leqslant |X|$. Thus $U \in \binom{X}{|U|}$, and hence $U \in \bigcup_{k \leqslant |X|} \binom{X}{k}$. This proves that $\mathscr{P}(X) \subseteq \bigcup_{k \leqslant |X|} \binom{X}{k}$.

The fact that $\bigcup_{k \leqslant |X|} \binom{X}{k} \subseteq \mathscr{P}(X)$ is immediate, since elements of $\binom{X}{k}$ are defined to be subsets of X, and hence elements of $\mathscr{P}(X)$. $\square$

Definition 3.3.4
Let $n, k \in \mathbb{N}$. Denote by $\binom{n}{k}$ (read: 'n choose k') (LaTeX code: `\binom{n}{k}`) the number of k-element subsets of $[n]$. That is, we define $\binom{n}{k} = \left|\binom{[n]}{k}\right|$. The numbers $\binom{n}{k}$ are called **binomial coefficients**.[a]

---
[a]Some authors use the notation $_nC_k$ or ^nC_k instead of $\binom{n}{k}$. We avoid this, as it is unnecessarily clunky.

Intuitively, $\binom{n}{k}$ is the number of ways of selecting k things from n, without repetitions, in such a way that order doesn't matter.

The value behind this notation is that it allows us to express huge numbers in a concise and meaningful way. For example,

$$\binom{4000}{11} = 103\,640\,000\,280\,154\,258\,645\,590\,429\,564\,000$$

Although these two numbers are equal, their *expressions* are very different; the expression on the left is meaningful, but the expression on the right is completely meaningless out of context.

Writing tip

When expressing the sizes of finite sets described combinatorially, it is best to *avoid* evaluating the expression; that is, leave in the powers, products, sums, binomial coefficients and factorials! The reason for this is that performing the calculations takes the meaning away from the expressions. ◁

Example 3.3.5

In Example 3.3.2 we proved that:

$$\binom{4}{0} = 1, \ \binom{4}{1} = 4, \ \binom{4}{2} = 6, \ \binom{4}{3} = 4, \ \binom{4}{4} = 1$$

and that $\binom{4}{k} = 0$ for all $k \geqslant 5$. ◁

Exercise 3.3.6

Fix $n \in \mathbb{N}$. Prove that $\binom{n}{0} = 1$, $\binom{n}{1} = n$ and $\binom{n}{n} = 1$. ◁

Definition 3.3.7

Let X be a set. A **permutation** of X is a bijection $X \to X$. Denote the set of all permutations of X by S_X (LaTeX code: S_X),[a] and write $S_{[n]} = S_n$ for $n \in \mathbb{N}$.

---

[a]The 'S' comes from 'symmetry'. The set S_X comes with the natural structure of a *group*.

Example 3.3.8

There are six permutations of the set $[3]$. Representing each $f \in S_{[3]}$ by the ordered triple $(f(1), f(2), f(3))$, these permutations are thus given by

$$(1,2,3), \ (1,3,2), \ (2,1,3), \ (2,3,1), \ (3,1,2), \ (3,2,1)$$

For example, $(2,3,1)$ represents the permutation $f : [3] \to [3]$ defined by $f(1) = 2$, $f(2) = 3$ and $f(3) = 1$. ◁

Exercise 3.3.9

List all the permutations of the set $[4]$. ◁

Definition 3.3.10

Let $n \in \mathbb{N}$. Denote by $n!$ (read: 'n factorial') the number of permutations of a set of size n. That is, $n! = |S_n|$. The numbers $n!$ are called **factorials**.

Example 3.3.11

Example 3.3.8 shows that $3! = 6$. ◁

Products and partitions

We saw in Proposition 3.2.22 and Proposition 3.2.18 that, given two finite sets X and Y, the product $X \times Y$ and the union $X \cup Y$ are finite. We also found formulae for their size. The *multiplication principle* (Strategy 3.3.21) and *addition principle* (Strategy 3.3.28) generalise these formulae, extending to products and (disjoint) unions of any finite number of finite sets.

Lemma 3.3.12

Let $\{X_1, \ldots, X_n\}$ be a family of finite sets, with $n \geqslant 1$. Then $\prod_{i=1}^{n} X_i$ is finite, and

$$\left| \prod_{i=1}^{n} X_i \right| = |X_1| \cdot |X_2| \cdot \cdots \cdot |X_n|$$

Proof

We proceed by induction on n.

- **(BC)** When $n = 1$, an element of $\prod_{i=1}^{1} X_i$ is 'officially' a sequence (x_1) with $x_1 \in X_1$. This is the same as an element of X_1, in the sense that the assignments $(x_1) \mapsto x_1$ and $x_1 \mapsto (x_1)$ define mutually inverse (hence bijective) functions between $\prod_{i=1}^{1} X_i$ and X_1, and so

$$\left| \prod_{i=1}^{1} X_i \right| = |X_1|$$

- **(IS)** Fix $n \in \mathbb{N}$, and suppose that

$$\left| \prod_{i=1}^{n} X_i \right| = |X_1| \cdot |X_2| \cdot \cdots \cdot |X_n|$$

for all sets X_i for $i \in [n]$. This is our induction hypothesis.

Now let $X_1, \ldots, X_n, X_{n+1}$ be sets. We define a function

$$F : \prod_{i=1}^{n+1} X_i \to \left(\prod_{i=1}^{n} X_i \right) \times X_{n+1}$$

by letting $F((x_1, \ldots, x_n, x_{n+1})) = ((x_1, \ldots, x_n), x_{n+1})$. It is again easy to check that F is a bijection, and hence

$$\left| \prod_{i=1}^{n+1} X_i \right| = \left| \prod_{i=1}^{n} X_i \right| \cdot |X_{n+1}|$$

by Proposition 3.2.22. Applying the induction hypothesis, we obtain the desired result, namely

$$\left| \prod_{i=1}^{n+1} X_i \right| = |X_1| \cdot |X_2| \cdot \cdots \cdot |X_n| \cdot |X_{n+1}|$$

By induction, we're done. □

Lemma 3.3.12 gives rise to a useful strategy for computing the size of a finite set X—see Strategy 3.3.13. Intuitively, by devising a step-by-step procedure for specifying an element of X, we are constructing a cartesian product $\prod_{k=1}^{n} X_k$, where X_k is the set of choices to be made in the k^{th} step. This establishes a bijection $\prod_{k=1}^{n} X_k \to X$, which by bijective proof (Strategy 3.2.15(c)) lets us compute $|X|$ as the product of the numbers of choices that can be made in each step.

Strategy 3.3.13 (Multiplication principle (independent version))
Let X be a finite set. In order to compute $|X|$, it suffices to find a step-by-step procedure for specifying elements of X, such that:

- Each element is specified by a unique sequence of choices;

- Each step in the procedure is independent of the previous step;

- There are finitely many choices to be made at each step.

If there are $n \in \mathbb{N}$ steps and $m_k \in \mathbb{N}$ possible choices in the k^{th} step, then $|X| = \prod_{k=1}^{n} m_k$. ◁

Example 3.3.14
You go to an ice cream stand selling the following flavours:

vanilla, strawberry, chocolate, rum and raisin, mint choc chip, toffee crunch

You can have your ice cream in a tub, a regular cone or a *choco-cone*. You can have one, two or three scoops. We will compute how many options you have.

To select an ice cream, you follow the following procedure:

- **Step 1.** Choose a flavour. There are 6 ways to do this.

- **Step 2.** Choose whether you'd like it in a tub, regular cone or choco-cone. There are 3 ways to do this.

- **Step 3.** Choose how many scoops you'd like. There are 3 ways to do this.

Hence there are $6 \times 3 \times 3 = 54$ options in total.   ◁

This may feel informal, but really what we are doing is establishing a bijection. Letting X be the set of options, the above procedure defines a bijection

$$F \times C \times S \to X$$

where F is the set of flavours, $C = \{\text{tub}, \text{regular cone}, \text{choco-cone}\}$ and $S = [3]$ is the set of possible numbers of scoops.

Example 3.3.15

We will prove that $|\mathscr{P}(X)| = 2^{|X|}$ for all finite sets X.[a]

Let X be a finite set and let $n = |X|$. Write

$$X = \{x_k \mid k \in [n]\} = \{x_1, x_2, \dots, x_n\}$$

Intuitively, specifying an element of $\mathscr{P}(X)$—that is, a subset $U \subseteq X$—is equivalent to deciding, for each $k \in [n]$, whether $x_k \in U$ or $x_k \notin U$ ('in or out'), which in turn is equivalent to specifying an element of $\{\text{in}, \text{out}\}^n$.

For example, taking $X = [7]$, the subset $U = \{1, 2, 6\}$ corresponds with the choices

$$1 \text{ in, } 2 \text{ in, } 3 \text{ out, } 4 \text{ out, } 5 \text{ out, } 6 \text{ in, } 7 \text{ out}$$

and hence the element $(\text{in}, \text{in}, \text{out}, \text{out}, \text{out}, \text{in}, \text{out}) \in \{\text{in}, \text{out}\}^7$.

This defines a function $i : \mathscr{P}(X) \to \{\text{in}, \text{out}\}^n$. This function is injective, since different subsets determine different sequences; and it is surjective, since each sequence determines a subset.

The above argument is sufficient for most purposes, but since this is the first bijective proof we have come across, we now give a more formal presentation of the details.

Define a function

$$i : \mathscr{P}(X) \to \{\text{in}, \text{out}\}^n$$

by letting the k^{th} component of $i(U)$ be 'in' if $x_k \in U$ or 'out' if $x_k \notin U$, for each $k \in [n]$.

We prove that i is a bijection.

---

[a]Some authors write 2^X to refer to the power set of a set X. This is justified by Example 3.3.15.

- *i* **is injective.** To see this, take $U, V \subseteq X$ and suppose $i(U) = i(V)$. We prove that $U = V$. So fix $x \in X$ and let $k \in [n]$ be such that $x = x_k$. Then

$$x \in U \Leftrightarrow \text{the } k^{\text{th}} \text{ component of } i(U) \text{ is 'in'} \qquad \text{by definition of } i$$
$$\Leftrightarrow \text{the } k^{\text{th}} \text{ component of } i(V) \text{ is 'in'} \qquad \text{since } i(U) = i(V)$$
$$\Leftrightarrow x \in V \qquad \text{by definition of } i$$

so indeed we have $U = V$, as required.

- *i* **is surjective.** To see this, let $v \in \{\text{in}, \text{out}\}^n$, and let

$$U = \{x_k \mid \text{the } k^{\text{th}} \text{ component of } v \text{ is 'in'}\}$$

Then $i(U) = v$, since for each $k \in [n]$ we have $x_k \in U$ if and only if the k^{th} component of v is 'in', which is precisely the definition of $i(U)$.

Hence

$$|\mathscr{P}(X)| = |\{\text{in}, \text{out}\}|^n = 2^n$$

as required. ◁

Exercise 3.3.16
Let X and Y be finite sets, and recall that Y^X denotes the set of functions from X to Y. Prove that $|Y^X| = |Y|^{|X|}$. ◁

Example 3.3.17
We count the number of ways we can shuffle a standard deck of cards in such a way that the colour of the cards alternate between red and black.

A procedure for choosing the order of the cards is as follows:

(i) Choose the colour of the first card. There are 2 such choices. This then determines the colours of the remaining cards, since they have to alternate.

(ii) Choose the order of the red cards. There are 26! such choices.

(iii) Choose the order of the black cards. There are 26! such choices.

By the multiplication principle, there are $2 \cdot (26!)^2$ such rearrangements. This number is huge, and in general there is no reason to write it out. Just for fun, though:

325 288 005 235 264 929 014 077 766 819 257 214 042 112 000 000 000 000

◁

Exercise 3.3.18

Since September 2001, car number plates on the island of Great Britain have taken the form XX NN XXX, where each X can be any letter of the alphabet except for 'I', 'Q' or 'Z', and NN is the last two digits of the year.[b] How many possible number plates are there? Number plates of vehicles registered in the region of Yorkshire begin with the letter 'Y'. How many Yorkshire number plates can be issued in a given year? ◁

The multiplication principle in the form of Strategy 3.3.13 does not allow for steps later in a procedure to depend on those earlier in the procedure. To see why this is a problem, suppose we want to count the size of the set $X = \{(a,b) \in [n] \times [n] \mid a \neq b\}$. A step-by-step procedure for specifying such an element is as follows:

- **Step 1.** Select an element $a \in [n]$. There are n choices.

- **Step 2.** Select an element $b \in [n]$ with $b \neq a$. There are $n-1$ choices.

We would like to use Strategy 3.3.13 to deduce that $|X| = n(n-1)$. Unfortunately, this is not valid because the possible choices available to us in Step 2 depend on the choice made in Step 1. Elements of cartesian products do not depend on one another, and so the set of sequences of choices made cannot necessarily be expressed as a cartesian product of two sets. Thus we cannot apply Lemma 3.3.12. Oh no!

However, provided that the *number* of choices in each step remains constant, in spite of the choices themselves changing, it turns out that we can still compute the size of the set in question by multiplying together the numbers of choices.

This is what we prove next. We begin with a pairwise version (analogous to Exercise 3.2.21) and then prove the general version by induction (like in Lemma 3.3.12).

Lemma 3.3.19

Fix $m, n \in \mathbb{N}$. Let X be a finite set with $|X| = m$, and for each $a \in X$, let Y_a be a finite set with $|Y_a| = n$. Then the set

$$P = \{(a,b) \mid a \in X \text{ and } b \in Y_a\}$$

is finite and $|P| = mn$.

Proof

Fix bijections $f : [m] \to X$ and $g_a : [n] \to Y_a$ for each $a \in X$. Define $h : [m] \times [n] \to P$ by letting $h(i, j) = (f(i), g_{f(i)}(j))$ for each $(i, j) \in [m] \times [n]$. Then:

---

[b]This is a slight simplification of what is really the case, but let's not concern ourselves with *too* many details!

- h is well-defined, since for all $i \in [m]$ and $j \in [n]$ we have $f(i) \in X$ and $g_{f(i)}(j) \in Y_{f(i)}$.

- h is injective. To see this, fix $(i,j), (k,\ell) \in [m] \times [n]$ and assume that $h(i,j) = h(k,\ell)$. Then $(f(i), g_{f(i)}(j)) = (f(k), g_{f(k)}(\ell))$, so that $f(i) = f(k)$ and $g_{f(i)}(j) = g_{f(k)}(\ell)$. Since f is injective, we have $i = k$—therefore $g_{f(i)}(j) = g_{f(i)}(\ell)$, and then since $g_{f(i)}$ is injective, we have $j = \ell$. Thus $(i,j) = (k,\ell)$, as required.

- h is surjective. To see this, let $(a,b) \in P$. Since f is surjective and $a \in X$, we have $a = f(i)$ for some $i \in [m]$. Since g_a is surjective and $b \in Y_a$, we have $b = g_a(j)$ for some $j \in [n]$. But then

$$(a,b) = (f(i), g_a(j)) = (f(i), g_{f(i)}(j)) = h(i,j)$$

so that h is surjective.

Since h is a bijection, we have $|P| = |[m] \times [n]|$ by [Theorem 3.2.13](iii), which is equal to mn by [Proposition 3.2.22](). □

We are now ready to state and prove the theorem giving rise to the multiplication principle in its full generality.

Theorem 3.3.20

Let $n \geqslant 1$ and $m_1, m_2, \ldots, m_n \in \mathbb{N}$. Suppose for each $i \in [n]$ that we are given finite sets $X^{(i)}_{a_1, \ldots, a_{i-1}}$ with $|X^{(i)}_{a_1, \ldots, a_{i-1}}| = m_i$, where $a_j \in X^{(i)}_{a_1, \ldots, a_{j-1}}$ for each $j < i$. Define

$$P = \{(a_1, a_2, \ldots, a_n) \mid a_1 \in X^{(1)}, \ a_2 \in X^{(2)}_{a_1}, \ \ldots, \ a_n \in X^{(n)}_{a_1, \ldots, a_{n-1}}\}$$

Then P is finite and $|P| = m_1 \times m_2 \times \cdots \times m_n$.

Proof

We proceed by induction on $n \geqslant 1$.

- **(Base case)** When $n = 1$, the statement says that given $m_1 \in \mathbb{N}$ and a finite set $X^{(1)}$ with $|X^{(1)}| = m_1$, then $P = \{(a_1) \mid a_1 \in X^{(1)}\}$ is finite and $|P| = m_1$. This is true, since the function $X^{(1)} \to P$ defined by $a \mapsto (a)$ is evidently a bijection.

- **(Induction step)** Fix $n \geqslant 1$ and assume that the statement is true for this value of n.

Let $m_1, m_2, \ldots, m_n, m_{n+1} \in \mathbb{N}$ and suppose that we are given finite sets $X^{(i)}_{a_1, \ldots, a_{i-1}}$ for each $i \in [n+1]$ just as in the statement of the theorem, and let

$$P = \{(a_1, a_2, \ldots, a_{n+1}) \mid a_1 \in X^{(1)}, \ a_2 \in X^{(2)}_{a_1}, \ \ldots, \ a_{n+1} \in X^{(n+1)}_{a_1, \ldots, a_{n-1}, a_n}\}$$

We need to prove that $|P| = m_1 \times m_2 \times \cdots \times m_n \times m_{n+1}$.

So define

$$Q = \{(a_1, a_2, \ldots, a_n) \mid a_1 \in X^{(1)},\ a_2 \in X^{(2)}_{a_1},\ \ldots,\ a_n \in X^{(n)}_{a_1,\ldots,a_{n-1}}\}$$

and, given $q = (a_1, \ldots, a_n) \in Q$, define $Y_q = X^{(n+1)}_{a_1,\ldots,a_n}$. Observe that there is an evident bijection

$$\{(q, a_{n+1}) \mid q \in Q,\ a_{n+1} \in Y_q\} \to P$$

defined by $((a_1, a_2, \ldots, a_n), a_{n+1}) \mapsto (a_1, a_2, \ldots, a_n, a_{n+1})$.

Now $|Q| = m_1 \times m_2 \times \cdots \times m_n$, and $|Y_q| = m_{n+1}$ for each $q \in Q$, so it follows from Lemma 3.3.19 that

$$|P| = (m_1 \times m_2 \times \cdots \times m_n) \times m_{n+1} = m_1 \times m_2 \times \cdots \times m_n \times m_{n+1}$$

as required.

$\square$

Strategy 3.3.21 summarises how Theorem 3.3.20 is useful in our proofs.

Strategy 3.3.21 (Multiplication principle)

Let X be a finite set. In order to compute $|X|$, it suffices to find a step-by-step procedure for specifying elements of X, such that:

- Each element is specified by a unique sequence of choices;

- The choices available in each step depend only on choices made in previous steps;

- There are finitely many choices available in each step;

- The *number* of choices available in each step does not depend on choices made in previous steps;

If there are $n \in \mathbb{N}$ steps and $m_k \in \mathbb{N}$ possible choices in the k^{th} step, then $|X| = \prod_{k=1}^{n} m_k$. $\triangleleft$

Example 3.3.22

We prove that there are six bijections $[3] \to [3]$. We can specify a bijection $f : [3] \to [3]$ according to the following procedure.

- **Step 1.** Choose the value of $f(1)$. There are 3 choices.

- **Step 2.** Choose the value of $f(2)$. The values $f(2)$ can take depend on the chosen value of $f(1)$.

  ◇ If $f(1) = 1$, then $f(2)$ can be equal to 2 or 3.

  ◇ If $f(1) = 2$, then $f(2)$ can be equal to 1 or 3.

  ◇ If $f(1) = 3$, then $f(2)$ can be equal to 1 or 2.

  In each case, there are 2 choices for the value of $f(2)$.

- **Step 3.** Choose the value of $f(3)$. The values $f(3)$ can take depend on the values of $f(1)$ and $f(2)$. We could split into the (six!) cases based on the values of $f(1)$ and $f(2)$ chosen in Steps 1 and 2; but we won't. Instead, note that $\{f(1), f(2)\}$ has two elements, and by injectivity $f(3)$ must have a distinct value, so that $[3] \setminus \{f(1), f(2)\}$ has one element. This element must be the value of $f(3)$. Hence there is only possible choice of $f(3)$.

By the multiplication principle, there are $3 \times 2 \times 1 = 6$ bijections $[3] \to [3]$. ◁

Exercise 3.3.23
Count the number of injections $[3] \to [4]$. ◁

The *addition principle* says that if we can *partition* a set into smaller chunks, then the size of the set is the sum of the sizes of the chunks. We will first make this notion of 'partition' precise.

Definition 3.3.24
Sets X and Y are **disjoint** if $X \cap Y = \varnothing$. More generally, given $n \in \mathbb{N}$, a collection of sets $X_1, X_2, \ldots, X_n$ is **pairwise disjoint** if $X_i \cap X_j = \varnothing$ for all $i, j \in [n]$ with $i \neq j$.

Definition 3.3.25
A (**finite**) **partition** of a set X is, for some $n \in \mathbb{N}$, a collection $\{U_i \mid i \in [n]\}$ of subsets of X such that:

 (a) Each U_i is inhabited;

 (b) The sets $U_1, U_2, \ldots, U_n$ are pairwise disjoint; and

 (c) $\bigcup_{i=1}^{n} U_i = X$.

For the purposes of proving Theorem 3.3.26 and stating the addition principle (Strategy 3.3.28), we may dispense with the requirement that the sets U_i in the

partition be inhabited, since if any of them are empty, they contribute a value of 0 to the sum. Thus when we say 'partition' in this section, we will secretly allow the sets in the partition to be empty. Be warned, though—when we discuss partitions in contexts other than the addition principle (for example in Section 5.1), we will require the sets in the partition to be inhabited.

Theorem 3.3.26
Let X be a set and let $\{U_1,\ldots,U_n\}$ be a partition of X for some $n \in \mathbb{N}$, such that each set U_i is finite. Then X is finite, and

$$|X| = |U_1| + |U_2| + \cdots + |U_n|$$

Exercise 3.3.27
Prove Theorem 3.3.26. The proof follows the same pattern as that of Lemma 3.3.12. Be careful to make sure you identify where you use the hypothesis that the sets U_i are pairwise disjoint! ◁

Strategy 3.3.28 (Addition principle)
Let X be a finite set. In order to compute $|X|$, it suffices to find a partition $U_1, U_2, \ldots, U_n$ of X; it then follows that $|X| = \sum_{k=1}^{n} |X_i|$. ◁

Example 3.3.29
We will count the number of inhabited subsets of $[7]$ which either contain only even numbers, or contain only odd numbers.

Let O denote the set of inhabited subsets of $[7]$ containing only odd numbers, and let E denote the set of inhabited subsets of $[7]$ containing only even numbers. Note that $\{O,E\}$ forms a partition of the set we are counting, and so our set has $|O| + |E|$ elements.

- An element of O must be a subset of $\{1,3,5,7\}$. By Example 3.3.15 there are $2^4 = 16$ such subsets. Thus the number of *inhabited* subsets of $[7]$ containing only odd numbers is 15, since we must exclude the empty set. That is, $|O| = 15$.

- A subset containing only even numbers must be a subset of $\{2,4,6\}$. Again by Example 3.3.15 there are $2^3 = 8$ such subsets. Hence there are 7 inhabited subsets of $[7]$ containing only even numbers. That is, $|E| = 7$.

Hence there are $15 + 7 = 22$ inhabited subsets of $[7]$ containing only even or only odd numbers. And here they are:

$\{1\}$	$\{3\}$	$\{5\}$	$\{7\}$	$\{1,3\}$	$\{2\}$	$\{4\}$	$\{6\}$
$\{1,5\}$	$\{1,7\}$	$\{3,5\}$	$\{3,7\}$	$\{5,7\}$	$\{2,4\}$	$\{2,6\}$	$\{4,6\}$
$\{1,3,5\}$	$\{1,3,7\}$	$\{1,5,7\}$	$\{3,5,7\}$	$\{1,3,5,7\}$	$\{2,4,6\}$		

◁

Exercise 3.3.30

Pick your favourite integer $n > 1000$. For this value of n, how many inhabited subsets of $[n]$ contain either only even or only odd numbers? (You need not evaluate exponents.)

◁

We now consider some examples of finite sets which use both the multiplication principle and the addition principle.

Example 3.3.31

A city has $6n$ inhabitants. The favourite colour of n of the inhabitants is orange, the favourite colour of $2n$ of the inhabitants is pink, and the favourite colour of $3n$ of the inhabitants is turquoise. The city government wishes to form a committee with equal representation from the three colour preference groups to decide how the new city hall should be painted. We count the number of ways this can be done.

Let X be the set of possible committees. First note that

$$X = \bigcup_{k=0}^{n} X_k$$

where X_k is the set of committees with exactly k people from each colour preference group. Indeed, we must have $k \leqslant n$, since it is impossible to have a committee with more than n people from the orange preference group.

Moreover, if $k \neq \ell$ then $X_k \cap X_\ell = \varnothing$, since if $k \neq \ell$ then a committee cannot simultaneously have exactly k people and exactly ℓ people from each preference group.

By the addition principle, we have

$$|X| = \sum_{k=0}^{n} |X_k|$$

We count X_k for fixed k using the following procedure:

- **Step 1.** Choose k people from the orange preference group to be on the committee. There are $\binom{n}{k}$ choices.

- **Step 2.** Choose k people from the pink preference group to be on the committee. There are $\binom{2n}{k}$ choices.

- **Step 3.** Choose k people from the turquoise preference group to be on the committee. There are $\binom{3n}{k}$ choices.

By the multiplication principle, it follows that $|X_k| = \binom{n}{k}\binom{2n}{k}\binom{3n}{k}$. Hence

$$|X| = \sum_{k=0}^{n} \binom{n}{k}\binom{2n}{k}\binom{3n}{k}$$

◁

Exercise 3.3.32
In Example 3.3.31, how many ways could a committee be formed with a *representative* number of people from each colour preference group? That is, the proportion of people on the committee which prefer any of the three colours should be equal to the corresponding proportion of the population of the city.   ◁

Double counting

Double counting (also known as *counting in two ways*) is a proof technique that allows us to prove that two natural numbers are equal by establishing they are two expressions for the size of the same set. (More generally, by Theorem 3.2.13(iii), we can relate them to the sizes of two sets which are in bijection.)

The proof of Proposition 3.3.33 illustrates this proof very nicely. We proved it already by induction in Example 3.1.30; the combinatorial proof we now provide is much shorter and cleaner.

Proposition 3.3.33
Let $n \in \mathbb{N}$. Then $2^n = \sum_{k=0}^{n} \binom{n}{k}$.

Proof
We know that $|\mathcal{P}([n])| = 2^n$ by Example 3.3.15 and that $\mathcal{P}([n]) = \bigcup_{k=0}^{n} \binom{[n]}{k}$ by Proposition 3.3.3. Moreover, the sets $\binom{[n]}{k}$ are pairwise disjoint, so by the addition principle it follows that

$$2^n = |\mathcal{P}([n])| = \left| \bigcup_{k=0}^{n} \binom{[n]}{k} \right| = \sum_{k=0}^{n} \left| \binom{[n]}{k} \right| = \sum_{k=0}^{n} \binom{n}{k}$$

□

Strategy 3.3.34 (Double counting)
In order to prove that two expressions involving natural numbers are equal, it suffices to define a set X and devise two counting arguments to show that $|X|$ is equal to both expressions.   ◁

The next example counts elements of *different* sets and puts them in bijection to establish an identity.

Proposition 3.3.35

Let $n, k \in \mathbb{N}$ with $n \geqslant k$. Then

$$\binom{n}{k} = \binom{n}{n-k}$$

Proof

First note that $\binom{n}{k} = \left| \binom{[n]}{k} \right|$ and $\binom{n}{n-k} = \left| \binom{[n]}{n-k} \right|$, so it suffices to find a bijection $f : \binom{[n]}{k} \to \binom{[n]}{n-k}$. Intuitively, this bijection arises because choosing k elements from $[n]$ to *put into* a subset is equivalent to choosing $n - k$ elements from $[n]$ to *leave out of* the subset. Specifically, we define

$$f(U) = [n] \setminus U \text{ for all } U \in \binom{[n]}{k}$$

Note first that f is well-defined, since if $U \subseteq [n]$ with $|U| = k$, then $[n] \setminus U \subseteq [n]$ and $|[n] \setminus U| = |[n]| - |U| = n - k$ by Exercise 3.2.20. We now prove f is a bijection:

- *f is injective.* Let $U, V \subseteq [n]$ and suppose $[n] \setminus U = [n] \setminus V$. Then for all $k \in [n]$, we have

$$\begin{aligned} k \in U &\Leftrightarrow k \notin [n] \setminus U &&\text{by definition of set difference} \\ &\Leftrightarrow k \notin [n] \setminus V &&\text{since } [n] \setminus U = [n] \setminus V \\ &\Leftrightarrow k \in V &&\text{by definition of set difference} \end{aligned}$$

  so $U = V$, as required.

- *f is surjective.* Let $V \in \binom{[n]}{n-k}$. Then $|[n] \setminus V| = n - (n - k) = k$ by Exercise 3.2.20, so that $[n] \setminus V \in \binom{[n]}{k}$. But then

$$f([n] \setminus V) = [n] \setminus ([n] \setminus V) = V$$

  by Exercise 2.1.61.

Since f is a bijection, we have

$$\binom{n}{k} = \left| \binom{[n]}{k} \right| = \left| \binom{[n]}{n-k} \right| = \binom{n}{n-k}$$

as required. □

We put a lot of detail into this proof. A slightly less formal proof might simply say that $\binom{n}{k} = \binom{n}{n-k}$ since choosing k elements from $[n]$ to put into a subset is equivalent to choosing $n - k$ elements from $[n]$ to leave out of the subset. This

would be fine as long as the members of the intended audience of your proof could reasonably be expected to construct the bijection by themselves.

The proof of Proposition 3.3.36 follows this more informal format.

Proposition 3.3.36

Let $n, k, \ell \in \mathbb{N}$ with $n \geqslant k \geqslant \ell$. Then

$$\binom{n}{k}\binom{k}{\ell} = \binom{n}{\ell}\binom{n-\ell}{k-\ell}$$

Proof

Let's home in on the left-hand side of the equation. By the multiplication principle, $\binom{n}{k}\binom{k}{\ell}$ is the number of ways of selecting a k-element subset of $[n]$ and an ℓ-element subset of $[k]$. Equivalently, it's the number of ways of selecting a k-element subset of $[n]$ and then an ℓ-element subset *of the k-element subset that we just selected*. To make this slightly more concrete, let's put it this way:

> $\binom{n}{k}\binom{k}{\ell}$ is the number of ways of painting k balls red from a bag of n balls, and painting ℓ of the red balls blue. This leaves us with ℓ blue balls and $k - \ell$ red balls.

Now we need to find an equivalent interpretation of $\binom{n}{\ell}\binom{n-\ell}{k-\ell}$. Well, suppose we pick the ℓ elements to be coloured blue first. To make up the rest of the k-element subset, we now have to select $k - \ell$ elements, and there are now $n - \ell$ to choose from. Thus

> $\binom{n}{\ell}\binom{n-\ell}{k-\ell}$ is the number of ways of painting ℓ balls from a bag of n balls blue, and painting $k - \ell$ of the remaining balls red.

Thus, both numbers represent the number of ways of painting ℓ balls blue and $k - \ell$ balls red from a bag of n balls. Hence they are equal. □

Exercise 3.3.37

Make the proof of Proposition 3.3.36 more formal by defining a bijection between sets of the appropriate sizes. ◁

Exercise 3.3.38

Provide a combinatorial proof that if $n, k \in \mathbb{N}$ with $n \geqslant k$, then

$$\binom{n+1}{k+1} = \binom{n}{k} + \binom{n}{k+1}$$

Deduce that the combinatorial definition of binomial coefficients (Definition 3.3.4) is equivalent to the recursive definition (Definition 3.1.28). ◁

The following proposition demonstrates that the combinatorial definition of factorials (Definition 3.3.10) is equivalent to the recursive definition (Definition 3.1.27).

> **Theorem 3.3.39**
> $0! = 1$ and if $n \in \mathbb{N}$ then $(n+1)! = (n+1) \cdot n!$.

Proof
The only permutation of $\varnothing$ is the empty function $e : \varnothing \to \varnothing$. Hence $S_0 = \{e\}$ and $0! = |S_0| = 1$.

Let $n \in \mathbb{N}$. A permutation of $[n+1]$ is a bijection $f : [n+1] \to [n+1]$. Specifying such a bijection is equivalent to carrying out the following procedure:

- Choose the (unique!) element $k \in [n+1]$ such that $f(k) = n+1$. There are $n+1$ choices for k.

- Choose the values of f at each $\ell \in [n+1]$ with $\ell \neq k$. This is equivalent to finding a bijection $[n+1] \setminus \{k\} \to [n]$. Since $|[n+1] \setminus \{k\}| = |[n]| = n$, there are $n!$ such choices.

By the multiplication principle, we have

$$(n+1)! = |S_{n+1}| = (n+1) \cdot n!$$

so we're done. $\qquad\square$

We now revisit Theorem 3.1.32; this time, our proof will be combinatorial, rather than inductive.

> **Theorem 3.3.40**
> Let $n, k \in \mathbb{N}$. Then
> $$\binom{n}{k} = \begin{cases} \dfrac{n!}{k!(n-k)!} & \text{if } k \leqslant n \\ 0 & \text{if } k > n \end{cases}$$

Proof
Suppose $k > n$. By Exercise 3.2.16, if $U \subseteq [n]$ then $|U| \leqslant n$. Hence if $k > n$, then $\binom{[n]}{k} = \varnothing$, and so $\binom{n}{k} = 0$, as required.

Now suppose $k \leqslant n$. We will prove that $n! = \binom{n}{k} \cdot k! \cdot (n-k)!$; the result then follows by dividing through by $k!(n-k)!$. We prove this equation by counting the number of elements of S_n.

A procedure for defining an element of S_n is as follows:

(i) Choose which elements will appear in the first k positions of the list. There are $\binom{n}{k}$ such choices.

(ii) Choose the order of these k elements. There are $k!$ such choices.

(iii) Choose the order of the remaining $n-k$ elements. There are $(n-k)!$ such choices.

By the multiplication principle, $n! = \binom{n}{k} \cdot k! \cdot (n-k)!$. $\square$

Note that the proof of Theorem 3.3.40 relied only on the combinatorial definitions of binomial coefficients and factorials; we didn't need to know how to compute them at all! The proof was *much* shorter, cleaner and, in some sense, more meaningful, than the inductive proof we gave in Section 3.1—see Theorem 3.1.32.

We conclude this section with some more examples and exercises in which double counting can be used.

Exercise 3.3.41

Let $n, k \in \mathbb{N}$ with $k \leqslant n+1$. Prove that

$$k\binom{n}{k} = (n-k+1)\binom{n}{k-1}$$

$\triangleleft$

Example 3.3.42

Let $m, n, k \in \mathbb{N}$. We prove that

$$\sum_{\ell=0}^{k} \binom{m}{\ell}\binom{n}{k-\ell} = \binom{m+n}{k}$$

by finding a procedure for counting the number of k-element subsets of an appropriate $(m+n)$-element set. Specifically, let X be a set containing m cats and n dogs. Then $\left|\binom{m+n}{k}\right|$ is the number of k-element subsets $U \subseteq X$. We can specify such a subset according to the following procedure.

• **Step 1.** Split into cases based on the number ℓ of cats in U. Note that we must have $0 \leqslant \ell \leqslant k$, since the number of cats must be a natural number and cannot exceed k as $|U| = k$. Moreover, these cases are mutually exclusive. Hence by the addition principle we have

$$\binom{m+n}{k} = \sum_{\ell=0}^{k} a_\ell$$

where a_ℓ is the number of subsets of X containing ℓ cats and $k-\ell$ dogs.

- **Step 2.** Choose ℓ cats from the m cats in X to be elements of U. There are $\binom{[m]}{\ell}$ such choices.

- **Step 3.** Choose $k - \ell$ dogs from the n dogs in X to be elements of U. There are $\binom{[n]}{k-\ell}$ such choices.

The multiplication principle shows that $a_\ell = \binom{m}{\ell}\binom{n}{k-\ell}$. Hence

$$\binom{m+n}{k} = \sum_{\ell=0}^{k} \binom{m}{\ell}\binom{n}{k-\ell}$$

as required. ◁

Exercise 3.3.43

Given natural numbers n, a, b, c with $a + b + c = n$, define the **trinomial coefficient** $\binom{n}{a,b,c}$ to be the number of ways of partitioning $[n]$ into three sets of sizes a, b and c, respectively. That is, $\binom{n}{a,b,c}$ is the size of the set

$$\left\{ (A,B,C) \;\middle|\; \begin{array}{l} A \subseteq [n], \quad B \subseteq [n], \quad C \subseteq [n], \\ |A| = a, \quad |B| = b, \quad |C| = c, \\ \text{and } A \cup B \cup C = [n] \end{array} \right\}$$

By considering trinomial coefficients, prove that if $a, b, c \in \mathbb{N}$, then $(a + b + c)!$ is divisible by $a! \cdot b! \cdot c!$. ◁

Inclusion–exclusion principle

The addition principle is useful only for counting unions of *pairwise disjoint* sets, i.e. sets that do not overlap. We saw in Proposition 3.2.18 how to compute the size of a union of two sets which *do* overlap:

$$|X \cup Y| = |X| + |Y| - |X \cap Y|$$

So far so good. But what if we have three or four sets instead of just two?

Exercise 3.3.44

Let X, Y, Z be sets. Show that

$$|X \cup Y \cup Z| = |X| + |Y| + |Z| - |X \cap Y| - |X \cap Z| - |Y \cap Z| + |X \cap Y \cap Z|$$

Let W be another set. Derive a similar formula for $|W \cup X \cup Y \cup Z|$. ◁

The inclusion–exclusion principle is a generalisation of Exercise 3.3.44 to arbitrary finite collections of finite sets, but it is stated in a slightly different way in order to make the proof more convenient.

Theorem 3.3.45 (Inclusion–exclusion principle)
Let $n \in \mathbb{N}$, let X_i be a finite set for each $i \in [n]$, and let $X = X_1 \cup X_2 \cup \cdots \cup X_n$. Then
$$\sum_{I \subseteq [n]} (-1)^{|I|} |X_I| = 0$$
where $X_I = \{a \in X \mid a \in X_i \text{ for all } i \in I\}$.

The statement of Theorem 3.3.45 looks fairly abstract, so before we prove it, let's examine its content. The sum is over all subsets $I \subseteq [n]$, and then the power $(-1)^{|I|}$ is equal to 1 if I has an even number of elements, and -1 if I has an odd number of elements. Moreover, if I is inhabited then X_I is the intersection of the sets X_i for $i \in I$—for example $X_{\{2,3,5\}} = X_2 \cap X_3 \cap X_5$; on the other hand, a careful examination of the definition of X_I reveals that $X_\varnothing = X$.

Thus when $n = 3$, the sum $\sum_{I \subseteq [3]} (-1)^{|I|} |X_I|$ can be evaluated as

$$|X| - |X_1| - |X_2| - |X_3| + |X_1 \cap X_2| + |X_1 \cap X_3| + |X_2 \cap X_3| - |X_1 \cap X_2 \cap X_3|$$

The theorem says that this sum is equal to zero, and solving for $|X| = |X_1 \cup X_2 \cup X_3|$ yields an equivalent equation to that in Exercise 3.3.44.

Proof of *Theorem 3.3.45*
Define sets $\mathscr{E}$ and $\mathscr{O}$ as follows:

- $\mathscr{E} = \{(I, x) \mid I \subseteq [n], x \in X_I, |I| \text{ is even}\}$;

- $\mathscr{O} = \{(I, x) \mid I \subseteq [n], x \in X_I, |I| \text{ is odd}\}$.

We first prove that $|\mathscr{E}| = |\mathscr{O}|$.

Given $x \in X$, define $i_x = \min\{i \in [n] \mid x \in X_i\}$. Note that i_x is well-defined since $X = X_1 \cup X_2 \cup \cdots \cup X_n$.

Now define $f : \mathscr{E} \to \mathscr{O}$ by

$$f(I, x) = \begin{cases} (I \cup \{i_x\}, x) & \text{if } i_x \notin I \\ (I \setminus \{i_x\}, x) & \text{if } i_x \in I \end{cases}$$

To see that f is well-defined, note that $f(I, x) \in \mathscr{O}$ for each $(I, x) \in \mathscr{E}$; indeed:

- If $|I|$ is even, then $|I \cup \{i_x\}| = |I| + 1$ if $i_x \notin I$, and $|I \setminus \{i_x\}| = |I| - 1$ if $i_x \in I$, and both $|I| + 1$ and $|I| - 1$ are odd.

- Suppose $(I, x) \in \mathscr{E}$. Then $x \in X_I$, which is to say that $\forall i \in I$, $x \in X_i$. Thus $x \in X_J$ for all $J \subseteq I$—in particular, for $J = I \setminus \{i_x\}$. Moreover, $X_{I \cup \{i_x\}} = X_I \cap X_{i_x}$. Since $x \in X_i$ and $x \in X_I$, we do indeed have $x \in X_{I \cup \{i_x\}}$.

To see that f is a bijection, note that it has an inverse $g : \mathscr{O} \to \mathscr{E}$ defined just like f:

$$g(I, x) = \begin{cases} (I \cup \{i_x\}, x) & \text{if } i_x \notin I \\ (I \setminus \{i_x\}, x) & \text{if } i_x \in I \end{cases}$$

Well-definedness of g follows from the same argument as that of f, and the fact that $g \circ f = \mathrm{id}_{\mathscr{E}}$ and $f \circ g = \mathrm{id}_{\mathscr{O}}$ can be checked by using the respective definitions of f and g.

Hence $|\mathscr{E}| = |\mathscr{O}|$.

For each $I \subseteq [n]$, define $\mathscr{S}_I = \{(I, x) \mid x \in X_I\}$. Observe that:

- For each $I \subseteq [n]$, we have $|\mathscr{S}_I| = |X_I|$, since the function $X_I \to \mathscr{S}_I$ defined by $x \mapsto (I, x)$ is a bijection.

- The sets $\mathscr{S}_I$ with $|I|$ even form a partition of $\mathscr{E}$, and the sets $\mathscr{S}_I$ with $|I|$ odd form a partition of $\mathscr{O}$.

Therefore by the addition principle we have

$$|\mathscr{E}| = \sum_{I \subseteq [n],\, |I| \text{ even}} |X_I| \quad \text{and} \quad |\mathscr{O}| = \sum_{I \subseteq [n],\, |I| \text{ odd}} |X_I|$$

Also observe that $(-1)^{|I|} = 1$ if $|I|$ is even, and $(-1)^{|I|} = -1$ if $|I|$ is odd.

Putting this all together, we have

$$\sum_{I \subseteq [n]} (-1)^{|I|} |X_I|$$

$$= \sum_{I \subseteq [n],\, |I| \text{ even}} (-1)^{|I|} |X_I| + \sum_{I \subseteq [n],\, |I| \text{ odd}} (-1)^{|I|} |X_I| \quad \text{splitting up the sum}$$

$$= \sum_{I \subseteq [n],\, |I| \text{ even}} |X_I| - \sum_{I \subseteq [n],\, |I| \text{ odd}} |X_I| \quad \text{evaluating } (-1)^{|I|} \text{ in each sum}$$

$$= |\mathscr{E}| - |\mathscr{O}| \quad \text{by the addition principle}$$

$$= 0 \quad \text{since } |\mathscr{E}| = |\mathscr{O}|$$

as required. □

It is more common to see the inclusion–exclusion principle stated in one two equivalent forms, stated here as Corollaries 3.3.46 and 3.3.47.

Corollary 3.3.46

Let $X_1, X_2, \ldots, X_n$ be sets. Then

$$\left| \bigcup_{i=1}^{n} X_i \right| = \sum_{k=1}^{n} \left(\sum_{1 \leqslant i_1 < i_2 < \cdots < i_k \leqslant n} (-1)^{k-1} |X_{i_1} \cap X_{i_2} \cap \cdots \cap X_{i_k}| \right)$$

Proof

Moving all terms to the left-hand side of the equation and observing that $-(-1)^{k-1} = (-1)^k$, the statement is equivalent to

$$\left| \bigcup_{i=1}^{n} X_i \right| - \sum_{k=1}^{n} \left(\sum_{1 \leqslant i_1 < i_2 < \cdots < i_k \leqslant n} (-1)^k |X_{i_1} \cap X_{i_2} \cap \cdots \cap X_{i_k}| \right) = 0$$

But using the notation of Theorem 3.3.45, we have

$$\left| \bigcup_{i=1}^{n} X_i \right| = |X| = (-1)^{|\varnothing|} |X_{\varnothing}|$$

and for all $1 \leqslant i_1 < i_2 < \cdots < i_k \leqslant n$, we have

$$(-1)^k |X_{i_1} \cap X_{i_2} \cap \cdots \cap X_{i_k}| = (-1)^{|\{i_1, i_2, \ldots, i_k\}|} |X_{\{i_1, i_2, \ldots, i_k\}}|$$

and so we see that this is just a restatement of Theorem 3.3.45. □

Corollary 3.3.47

Let X be a set and let $U_1, U_2, \ldots, U_n \subseteq X$. Then

$$\left| X \setminus \bigcup_{i=1}^{n} U_i \right| = |X| + \sum_{k=1}^{n} \left(\sum_{1 \leqslant i_1 < i_2 < \cdots < i_k \leqslant n} (-1)^k |U_{i_1} \cap U_{i_2} \cap \cdots \cap U_{i_k}| \right)$$

Proof

Since $\bigcup_{i=1}^{n} U_i \subseteq X$, we have

$$\left| X \setminus \bigcup_{i=1}^{n} U_i \right| = |X| - \left| \bigcup_{i=1}^{n} U_i \right|$$

The result then follows immediately from Corollary 3.3.46. □

Proof tip

To find the size of a union of $\bigcup_{i=1}^{n} X_i$:

- Add the sizes of the individual sets X_i;

- Subtract the sizes of the double-intersections $X_i \cap X_j$;

- Add the sizes of the triple-intersections $X_i \cap X_j \cap X_k$;

- Subtract the sizes of the quadruple-intersections $X_i \cap X_j \cap X_k \cap X_\ell$;

- ... and so on ...

Keep alternating until the intersection of all the sets is covered. ◁

Example 3.3.48
We count how many subsets of $[12]$ contain a multiple of 3. Precisely, we count the number of elements of the set

$$X_3 \cup X_6 \cup X_9 \cup X_{12}$$

where $X_k = \{S \subseteq [12] \mid k \in S\}$. We will apply the inclusion–exclusion principle:

(i) An element $S \in X_3$ is precisely a set of the form $\{3\} \cup S'$, where $S' \subseteq [12] \setminus \{3\}$. Since $[12] \setminus \{3\}$ has 11 elements, there are 2^{11} such subsets. So $|X_3| = 2^{11}$, and likewise $|X_6| = |X_9| = |X_{12}| = 2^{11}$.

(ii) An element $S \in X_3 \cap X_6$ is a set of the form $\{3,6\} \cup S'$, where $S' \subseteq [12] \setminus \{3,6\}$. Thus there are 2^{10} such subsets, so $|X_3 \cap X_6| = 2^{10}$. And likewise

$$|X_3 \cap X_9| = |X_3 \cap X_{12}| = |X_6 \cap X_9| = |X_6 \cap X_{12}| = |X_9 \cap X_{12}| = 2^{10}$$

(iii) Reasoning as in the last two cases, we see that

$$|X_3 \cap X_6 \cap X_9| = |X_3 \cap X_6 \cap X_{12}| = |X_3 \cap X_9 \cap X_{12}| = |X_6 \cap X_9 \cap X_{12}| = 2^9$$

(iv) ... and $|X_3 \cap X_6 \cap X_9 \cap X_{12}| = 2^8$.

Thus the number of subsets of $[12]$ which contain a multiple of 3 is

$$\underbrace{4 \times 2^{11}}_{\text{by (i)}} - \underbrace{6 \times 2^{10}}_{\text{by (ii)}} + \underbrace{4 \times 2^9}_{\text{by (iii)}} - \underbrace{2^8}_{\text{by (iv)}}$$

which is equal to 3840. ◁

Exercise 3.3.49
How many natural numbers less than 1000 are multiples of 2, 3, 5 or 7? ◁

Section 3.Q
Chapter 3 exercises

Under construction!
The end-of-chapter exercise sections are new and in an incomplete state.

Finite sets

1. Let $n \in \mathbb{N}$ and let $f : [n] \to [n]$ be a function. Prove that f is injective if and only if f is surjective.

Double counting

2. Let $a, b, m, n \in \mathbb{N}$. Prove each of the following by double counting.

(a) $a(m+n) = am + an$ (c) $(a^m)^n = a^{mn}$

(b) $a^{m+n} = a^m \cdot a^n$ (d) $(ab)^n = a^n \cdot b^n$

3. Prove that $\displaystyle\sum_{k=0}^{n} \binom{n}{k}^2 = \binom{2n}{n}$ for all $n \in \mathbb{N}$

4. Prove that $\displaystyle\sum_{k=m}^{n} \binom{n}{k}\binom{k}{m} = 2^{n-m}\binom{n}{m}$ for all $m, n \in \mathbb{N}$ with $m \leqslant n$.

5. Prove that $\displaystyle\sum_{j=0}^{k} \binom{n-j}{k-j} = \binom{n+1}{k}$ for all $n, k \in \mathbb{N}$.

6. Prove that $\displaystyle\sum_{k=1}^{n}\sum_{\ell=0}^{k} k\binom{n}{k}\binom{n-k}{\ell} = n \cdot 3^{n-1}$ for all $n \in \mathbb{N}$.

7. Prove that $\displaystyle\binom{n}{r+s+1} = \sum_{k=r+1}^{n-s} \binom{k-1}{r}\binom{n-k}{s}$ for all $n, r, s \in \mathbb{N}$.

8. Let $a_1, a_2, \ldots, a_r \in \mathbb{N}$ and let $n = a_1 + a_2 + \cdots + a_r$. Prove that

$$\binom{n}{a_1, a_2, \ldots, a_r} = \prod_{k=0}^{r-1} \binom{n - \sum_{i=1}^{k} a_i}{a_{k+1}}$$

where $\begin{pmatrix} n \\ a_1,a_2,\ldots,a_r \end{pmatrix}$ is the number of ordered r-tuples $(U_1,U_2,\ldots,U_r)$ such that $U_1,U_2,\ldots,U_r$ is a partition of $[n]$ and $|U_k| = a_k$ for all $k \in [r]$.bi

Inclusion–exclusion principle

9. Find the number of subsets of $[100]$ that do not contain a multiple of 8.

Chapter 4

Number theory

Section 4.1
Division

This section introduces the notion of *divisibility*. As we have already mentioned, it is not always the case that one integer can divide another. As you read through this section, note that we never use fractions; everything we do is *internal* to $\mathbb{Z}$, and does not require that we 'spill over' to $\mathbb{Q}$ at any point. This will help you when you study ring theory in the future, and is a good practice to mimic in your own work.

The following theorem, called the division theorem, is the crux of everything that is to follow.

> **Theorem 4.1.1** (Division theorem)
> Let $a, b \in \mathbb{Z}$ with $b \neq 0$. There exist unique $q, r \in \mathbb{Z}$ such that
>
> $$a = qb + r \quad \text{and} \quad 0 \leqslant r < |b|$$

Strategy
Let's look at the simple case when $a \geqslant 0$ and $b > 0$. We can always find q, r such that $a = qb + r$, for example $q = 0$ and $r = a$. Moreover, by increasing q we can reduce r, since

$$qb + r = (q+1)b + (r - b)$$

We will keep doing this until the 'remainder' is as small as it can be without being negative. As an example, consider the case when $a = 14$ and $b = 5$. This procedure gives

$$
\begin{aligned}
14 &= 0 \times 5 + 14 \\
&= 1 \times 5 + 9 \\
&= 2 \times 5 + 4 \qquad\qquad \leftarrow \text{least nonnegative remainder} \\
&= 3 \times 5 + (-1) \\
&= \cdots
\end{aligned}
$$

This procedure shows that in this case we should take $q = 2$ and $r = 4$, since $14 = 2 \times 5 + 4$ and $0 \leqslant 4 < |5|$.

We can show that such a descending sequence of remainders terminates using the well-ordering principle, and then we must argue that the quotient and remainder that we obtain are unique. ◁

Proof ⋆ *Proof*
We may assume that $b > 0$: if not, replace b by $-b$ and q by $-q$. We may also assume that $a \geqslant 0$. Otherwise, replace a by $-a$, q by $-(q+1)$ and r by $b - r$.

Thus, what follows assumes that $a \geqslant 0$ and $b > 0$.

- **Existence.** We prove that such integers q, r exist by the well-ordering principle. Namely, we define a sequence $(r_n)_{n \in \mathbb{N}}$ such that $a = nb + r_n$ and $r_0 > r_1 > r_2 > \cdots$, and use this sequence to find the values of q, r.

  ◇ Let $r_0 = a$. Then $a = 0b + r_0$, as required.
  ◇ Suppose r_n has been defined, and let $r_{n+1} = r_n - b$. Then

  $$(n+1)b + r_{n+1} = (n+1)b + r_n - b$$
  $$= nb + b + r_n - b$$
  $$= nb + r = a$$

  Since $b > 0$, we must have $r_{n+1} < r_n$ for all n.

  Let $R = \mathbb{N} \cap \{r_n \mid n \in \mathbb{N}\}$. That is, R is the set of terms of the sequence which are non-negative. Since $r_0 = a \geqslant 0$, we have that $r_0 \in R$ and hence R is inhabited. By the well-ordering principle, R has a least element r_k for some $k \in \mathbb{N}$.

  Define $q = k$ and $r = r_k$. By construction we have $a = qb + r$ and $r \geqslant 0$, so it remains to show that $r < b$. Well, if $r \geqslant b$ then $r - b \geqslant 0$, but $r - b = r_{k+1}$, so this would imply $r_{k+1} \in R$, contradicting minimality of r. Hence $r < b$, so q, r are as required.

- **Uniqueness.** Suppose q', r' also satisfy $a = q'b + r'$ and $0 \leqslant r' < b$. If we can show that $r' = r$ then this proves that $q = q'$: indeed, if $qb + r = q'b + r$ then we can subtract r and then divide by b, since $b > 0$.

  First note that $q' \geqslant 0$. If $q' < 0$ then $q' \leqslant -1$, so

  $$a = q'b + r' \leqslant -b + r'$$

  and hence $r' \geqslant a + b \geqslant b$ since $a \geqslant 0$. This contradicts the assumption that $r < b$. So $q' \geqslant 0$.

  Since $q' \geqslant 0$, we also know that $a = q'b + r_{q'}$, and hence $r' = r_{q'} \in R$. By minimality of r we have $r \leqslant r'$. It remains to show that $r = r'$. If not then $r < r'$. Thus

  $$qb + r = q'b + r' > q'b + r \quad \Rightarrow \quad qb > q'b \quad \Rightarrow \quad q > q'$$

  and hence $q = q' + t$ for some $t \geqslant 1$. But then

  $$q'b + r' = a = qb + r = (q' + t)b + r = q'b + (tb + r)$$

  so $r' = tb + r \geqslant b$, contradicting $r' < b$. So $r = r'$ as desired, and hence $q = q'$.

At long last, we are done.   □

Definition 4.1.2

Let $a, b \in \mathbb{Z}$ with $b \neq 0$, and let q, r be the unique integers such that

$$a = qb + r \quad \text{and} \quad 0 \leqslant r < |b|$$

We say q is the **quotient** and r is the **remainder** of a divided by b.

Example 4.1.3

Some examples of division include:

$$14 = 2 \times 5 + 4, \qquad -14 = -3 \times 5 + 1, \qquad 15 = 3 \times 5 + 0$$

◁

Definition 4.1.4

Let $a, b \in \mathbb{Z}$. We say b **divides** a, or that b is a **divisor** (or **factor**) of a, if there exists $q \in \mathbb{Z}$ such that $a = qb$. To denote the fact that b divides a we write $b \mid a$ (LaTeX code: `\mid`). For the negation $\neg(b \mid a)$ write $b \nmid a$ (LaTeX code: `\nmid`).

Thus, when $b \neq 0$, saying $b \mid a$ is equivalent to saying that the remainder of a divided by b is 0.

Example 4.1.5

5 divides 15 since $15 = 3 \times 5$. However, 5 does not divide 14: we know that the remainder of 14 divided by 5 is 4, not 0—and it can't be both since we proved in the division theorem that remainders are unique!   ◁

Exercise 4.1.6

Show that if $a \in \mathbb{Z}$ then $1 \mid a$, $-1 \mid a$ and $a \mid 0$. For which integers a does $a \mid 1$? For which integers a does $0 \mid a$?   ◁

We now introduce the very basic notion of a *unit*. This notion is introduced to rule out trivialities. Units become interesting when talking about general rings, but in $\mathbb{Z}$, the units are very familiar.

Definition 4.1.7

Let $u \in \mathbb{Z}$. We say u is a **unit** if $u \mid 1$; that is, u is a unit if there exists $v \in \mathbb{Z}$ such that $uv = 1$.

Proposition 4.1.8

The only units in $\mathbb{Z}$ are 1 and -1.

Proof

First note that 1 and -1 are units, since $1 \cdot 1 = 1$ and $(-1) \cdot (-1) = 1$. Now suppose that $u \in \mathbb{Z}$ is a unit, and let $v \in \mathbb{Z}$ be such that $uv = 1$. Certainly $u \neq 0$, since $0v = 0 \neq 1$. If $u > 1$ or $u < -1$ then $v = \frac{1}{u} \notin \mathbb{Z}$. So we must have $u \in \{-1, 1\}$. □

Exercise 4.1.6 shows that -1, 0 and 1 are, from the point of view of divisibility, fairly trivial. For this reason, most of the results we discuss regarding divisibility will concern **nonzero nonunits**, i.e. all integers except -1, 0 or 1.

Greatest common divisors

> **Definition 4.1.9**
>
> Let $a, b \in \mathbb{Z}$. An integer d is a **greatest common divisor** of a and b if:
>
> (a) $d \mid a$ and $d \mid b$;
>
> (b) If q is another integer such that $q \mid a$ and $q \mid b$, then $q \mid d$.

Example 4.1.10

2 is a greatest common divisor of 4 and 6; indeed:

(a) $4 = 2 \times 2$, and $6 = 3 \times 2$, so $2 \mid 4$ and $2 \mid 6$;

(b) Suppose $q \mid 4$ and $q \mid 6$. The divisors of 4 are $\pm 1, \pm 2, \pm 4$ and the divisors of 6 are $\pm 1, \pm 2, \pm 3, \pm 6$. Since q divides both, it must be the case that $q \in \{-2, -1, 1, 2\}$; in any case, $q \mid 2$.

Likewise, -2 is a greatest common divisor of 4 and 6. ◁

Exercise 4.1.11

There are two greatest common divisors of 6 and 15; find both. ◁

We will now prove that greatest common divisors *exist*—that is, any two integers have a greatest common divisor—and that they are *unique up to sign*.

> **Theorem 4.1.12**
>
> Every pair of integers a, b has a greatest common divisor.

Proof

First note that if $a = b = 0$, then 0 is a greatest common divisor for a and b. Moreover, we may take a, b to be non-negative, since divisibility is insensitive to sign. So suppose that $a, b \geqslant 0$ and that a, b are not both zero.

Define a set $X \subseteq \mathbb{Z}$ by

$$X = \{au + bv \mid u, v \in \mathbb{Z}, \ au + bv > 0\}$$

That is, X is the set of positive integers of the form $au + bv$.

X is inhabited. To see this, note that $a^2 > 0$ or $b^2 > 0$ since $a \neq 0$ or $b \neq 0$, so letting $u = a$ and $v = b$ in the expression $au + bv$, we see that

$$au + bv = a^2 + b^2 > 0 \quad \Rightarrow \quad a^2 + b^2 \in X$$

By the well-ordering principle, X has a least element d, and by definition of X there exist $u, v \in \mathbb{Z}$ such that $d = au + bv$.

We will prove that d is a greatest common divisor for a and b.

- $d \mid a$. If $a = 0$, then this is immediate, so suppose that $a > 0$. Let $q, r \in \mathbb{Z}$ be such that

$$a = qd + r \quad \text{and} \quad 0 \leqslant r < d$$

  Now $a = a \cdot 1 + b \cdot 0$, so $a \in X$, and hence $d \leqslant a$. Moreover

$$r = a - qd = a - q(au + bv) = a(1 - qu) + b(-qv)$$

  If $r > 0$ then this implies that $r \in X$; but this would contradict minimality of d, since $r < d$. So we must have $r = 0$ after all.

- $d \mid b$. The proof of this is identical to the proof that $d \mid a$.

- Suppose q is an integer dividing both a and b. Then $q \mid au + bv$ by Exercise 0.16. Since $au + bv = d$, we have $q \mid d$.

So d is a greatest common divisor of a and b after all.   □

Exercise 4.1.13

Let $a, b \in \mathbb{Z}$. If d and d' are two greatest common divisors of a and b, then either $d = d'$ or $d = -d'$.   ◁

Aside

A consequence of Theorem 4.1.12 and Exercise 4.1.13 is that every pair of integers has a unique non-negative greatest common divisor! Written symbolically, we can say

$$\forall (a, b) \in \mathbb{Z} \times \mathbb{Z}, \ \exists ! d \in \mathbb{Z}, \ \begin{pmatrix} d \geqslant 0 \text{ and } d \text{ is a greatest} \\ \text{common divisor for } a \text{ and } b \end{pmatrix}$$

As discussed in Section 2.2, since this is a formula of the form 'for all ... there exists a unique ...', this defines a function $\gcd : \mathbb{Z} \times \mathbb{Z} \to \mathbb{Z}$. We won't explicitly refer to the fact that gcd is a function; rather, we'll just concern ourselves with its values, as in Notation 4.1.14.   ◁

Exercise 4.1.13 justifies our use of the following notation to refer to greatest common divisors.

Notation 4.1.14
Let $a, b \in \mathbb{Z}$. Denote by $\gcd(a,b)$ (LATEX code: `\mathrm{gcd}`) the (unique!) non-negative greatest common divisor of a and b.

Example 4.1.15
In Example 4.1.10, we saw that both 2 and -2 are greatest common divisors of 4 and 6. Using Notation 4.1.14, we can now write $\gcd(4,6) = 2$. ◁

Exercise 4.1.16
For each $n \in \mathbb{Z}$, let $D_n \subseteq \mathbb{Z}$ be the set of divisors of n. Prove that $D_a \cap D_b = D_{\gcd(a,b)}$ for all $a, b \in \mathbb{Z}$. ◁

Our goal for the rest of this subsection is to investigate the behaviour of greatest common divisors, find out how to compute them, and look into the implications they have for solutions to certain kinds of equations.

Theorem 4.1.17
Let $a, b, q, r \in \mathbb{Z}$, and suppose that $a = qb + r$. Then

$$\gcd(a,b) = \gcd(b,r)$$

Proof
Let $d = \gcd(a,b)$. We check that d satisfies the conditions required to be a greatest common divisor of b and r.

Note that $d \mid a$ and $d \mid b$, so let $s, t \in \mathbb{Z}$ be such that $a = sd$ and $b = td$.

- $d \mid b$ by definition, and $d \mid r$ since

$$r = a - qb = sd - qtd = (s - qt)d$$

- Suppose $d' \mid b$ and $d' \mid r$; say $b = ud'$ and $r = vd'$ with $u, v \in \mathbb{Z}$. Then $d' \mid a$, since
$$a = qb + r = qud' + vd' = (qu + v)d'$$
so $d' \mid d$ since $d = \gcd(a,b)$.

So d is a greatest common divisor of b and r. Since $d > 0$, the result is shown. □

Combined with the division theorem (Theorem 4.1.1), Theorem 4.1.17 gives a relatively fast algorithm for computing the greatest common divisor of two integers, known as the **Euclidean algorithm**.

Proof tip
Euclidean algorithm. Let $a, b \in \mathbb{Z}$. To find $\gcd(a, b)$, proceed as follows.

- Set $r_0 = |a|$ and $r_1 = |b|$.

- Given r_{n-2} and r_{n-1}, define r_n to be the remainder of r_{n-2} divided by r_{n-1}.

- Stop when $r_n = 0$; then $r_{n-1} = \gcd(a, b)$.

$\triangleleft$

Example 4.1.18
We will find the greatest common divisor of 148 and 28.

$$148 = 5 \times 28 + 8$$
$$28 = 3 \times 8 + 4$$
$$8 = 2 \times \boxed{4} + 0 \qquad\qquad \leftarrow \text{Stop!}$$

Hence $\gcd(148, 28) = 4$. Here the sequence of remainders is given by:

$$r_0 = 148, \quad r_1 = 28, \quad r_2 = 8, \quad r_3 = 4, \quad r_4 = 0$$

$\triangleleft$

Example 4.1.19
The Euclidean algorithm works surprisingly quickly, even for relatively large numbers. Consider the problem of computing $\gcd(1311, 5757)$ for example:

$$5757 = 4 \times 1311 + 513$$
$$1311 = 2 \times 513 + 285$$
$$513 = 1 \times 285 + 228$$
$$285 = 1 \times 228 + 57$$
$$228 = 4 \times \boxed{57} + 0 \qquad\qquad \leftarrow \text{Stop!}$$

Hence $\gcd(1311, 5757) = 57$. Here the sequence of remainders is given by:

$$r_0 = 5757, \quad r_1 = 1311, \quad r_2 = 513, \quad r_3 = 285, \quad r_4 = 228, \quad r_5 = 57, \quad r_6 = 0$$

$\triangleleft$

Example 4.1.20
Here's an example where one of the numbers is negative: we compute the value of $\gcd(-420, 76)$:

$$-420 = (-6) \times 76 + 36$$
$$76 = 2 \times 36 + 4$$
$$36 = 9 \times \boxed{4} + 0 \qquad\qquad \leftarrow \text{Stop!}$$

Hence $\gcd(-420, 76) = 4$. $\triangleleft$

Example 4.1.21

Use the Euclidean algorithm to compute the greatest common divisors of the following pairs of integers

$$(12,9), \quad (100,35), \quad (7125,1300), \quad (1010,101010), \quad (-4,14)$$

◁

The following theorem will be useful when we study modular arithmetic in Section 4.3; it is called a 'lemma' for historical reasons, and is really an important result in its own right.

Theorem 4.1.22 (Bézout's lemma)

Let $a,b,c \in \mathbb{Z}$, and let $d = \gcd(a,b)$. The equation

$$ax + by = c$$

has a solution $(x,y) \in \mathbb{Z} \times \mathbb{Z}$ if and only if $d \mid c$.

Proof

($\Rightarrow$) Write $a = a'd$ and $b = b'd$, for $a',b' \in \mathbb{Z}$. If there exist $x,y \in \mathbb{Z}$ such that $ax + by = c$, then

$$c = ax + by = a'dx + b'dy = (a'x + b'y)d$$

and so $d \mid c$.

($\Leftarrow$) Suppose $d \mid c$, and let $c = kd$ for some $k \in \mathbb{Z}$.

If $c = 0$, then a solution is $x = y = 0$. If $c < 0$, then $ax + by = c$ if and only if $a(-x) + b(-y) = -c$; so we may assume that $c > 0$.

We proved in Theorem 4.1.12 that a greatest common divisor of a and b is a least element of the set

$$X = \{au + bv \mid u,v \in \mathbb{Z}, \ au + bv > 0\}$$

So let $u,v \in \mathbb{Z}$ be such that $au + bv = d$. Then

$$a(ku) + b(kv) = k(au + bv) = kd = c$$

and so letting $x = ku$ and $y = kv$, we see that the equation $ax + by = c$ has a solution $(x,y) \in \mathbb{Z} \times \mathbb{Z}$. □

Bézout's lemma completely characterises when the equation $ax + by = c$ has a solution. An easy generalisation of Bézout's lemma provides a complete characterisation of when solutions to **linear Diophantine equations** exist, that is equations of the form

$$ax + by = c$$

where $a, b, c \in \mathbb{Z}$. We will soon develop an algorithm for computing *all* solutions to these equations.

Example 4.1.23

Here are some examples of applications of Bézout's lemma.

- Consider the equation $1311x + 5757y = 12963$. We computed in Example 4.1.19 that $\gcd(1311, 5757) = 57$. But $57 \nmid 12963$ since $12963 = 227 \times 57 + 24$. By Bézout's lemma, the equation $1311x + 5757y = 12963$ has no integer solutions.

- For fixed z, the equation $4u + 6v = z$ has solutions exactly when z is even, since $\gcd(4, 6) = 2$.

- For fixed a, b, the equation $au + bv = 0$ always has solution. Indeed, setting $u = b$ and $v = -a$ gives a solution; but we knew one had to exist since by Exercise 4.1.6 we know that $d \mid 0$ for all $d \in \mathbb{Z}$.

◁

Exercise 4.1.24

Which of the following equations have solutions?

(a) $12u + 9v = -18$

(b) $12u + 9v = 1$

(c) $100u + 35v = 125$

(d) $7125u + 1300v = 0$

(e) $1010u + 101010v = 1010101010101010$

(f) $14u - 4v = 12$

◁

Coprimality

<div style="border: 1px solid; padding: 10px;">

Definition 4.1.25
Let $a, b \in \mathbb{Z}$. We say a and b are **coprime** (or **relatively prime**), and write $a \perp b$ (LaTeX code: `\perp`) (read 'a is coprime to b'), if $\gcd(a, b) = 1$.

</div>

Example 4.1.26
$4 \perp 9$. To see this, note that if $d \mid 4$ then $d \in \{-4, -2, -1, 1, 2, 4\}$, and if $d \mid 9$ then $d \in \{-9, -3, -1, 1, 3, 9\}$. Hence if $d \mid 4$ and $d \mid 9$, then $d = 1$ or $d = -1$. It follows that $\gcd(4, 9) = 1$. ◁

Exercise 4.1.27
Which integers in the set $[15]$ are coprime to 15? ◁

Proposition 4.1.28
Let $a, b \in \mathbb{Z}$. The following are equivalent:

(1) a and b are coprime;

(2) If $d \in \mathbb{Z}$ with $d \mid a$ and $d \mid b$, then d is a unit.

Proof
We prove that condition (1) implies condition (2), and vice versa.

- (1)⇒(2). Suppose a and b are coprime, and fix $d \in \mathbb{Z}$ with $d \mid a$ and $d \mid b$. Then $d \mid \gcd(a, b) = 1$, so d is a unit.

- (2)⇒(1). Suppose condition (2) above holds. We prove that 1 satisfies the conditions required to be a greatest common divisor of a and b. The fact that $1 \mid a$ and $1 \mid b$ is automatic; and the fact that if $d \mid a$ and $d \mid b$ implies $d \mid 1$ is precisely the condition (2) that we are assuming.

Hence the two conditions are equivalent. □

Exercise 4.1.29
Let a and b be integers, not both zero, and let $d = \gcd(a, b)$. The integers $\frac{a}{d}$ and $\frac{b}{d}$ are coprime. ◁

The following corollary is a specialisation of Bézout's lemma to the case when a and b are coprime.

Corollary 4.1.30
Let $a, b \in \mathbb{Z}$. The equation $au + bv = 1$ has a solution if and only if a and b are coprime. Moreover, if a and b are coprime, then the equation $au + bv = z$ has a solution for all $z \in \mathbb{Z}$.

Proof

By Bézout's lemma (Theorem 4.1.22), the equation $au + bv = 1$ has a solution if and only if $\gcd(a,b) \mid 1$. But the only positive divisor of 1 is 1, so a solution exists if and only if $\gcd(a,b) = 1$, which is precisely the assertion that a and b are coprime.

If a and b are coprime, then $1 = \gcd(a,b) \mid z$ for all $z \in \mathbb{Z}$. So by Bézout's lemma again, the equation $au + bv = z$ has a solution for all $z \in \mathbb{Z}$. $\square$

A useful consequence of Bézout's lemma is the following result:

Proposition 4.1.31
Let $a, b, c \in \mathbb{Z}$. If a and b are coprime and $a \mid bc$, then $a \mid c$.

Proof

By Bézout's lemma (Theorem 4.1.22) there exist integers u and v such that $au + bv = 1$. Multiplying by c gives $acu + bcv = c$. Since $a \mid bc$, we can write $bc = ka$ for some $k \in \mathbb{Z}$, and so $acu + kav = c$. But then

$$(cu + kv)a = c$$

which proves that $a \mid c$. $\square$

Linear Diophantine equations

We have now seen two important results:

- The **Euclidean algorithm**, which was a procedure for computing the greatest common divisor of two integers.

- **Bézout's lemma**, which provides a necessary and sufficient condition for equations of the form $ax + by = c$ to have an integer solution.

We will now develop the **reverse Euclidean algorithm**, which provides a method for computing a solutions to (bivariate) linear Diophantine equations, when such a solution exists. Then we will prove a theorem that characterises *all* integer solutions in terms of a given solution.

Example 4.1.32
Suppose we want to find integers x and y such that $327x + 114y = 18$. Running the Euclidean algorithm yields that $\gcd(327, 114) = 3$ — see below. For reasons soon to become apparent, we rearrange each equation to express the remainder

on its own.

$$327 = 2 \times 114 + 99 \qquad \Rightarrow \qquad 99 = 327 - 2 \times 114 \qquad (1)$$
$$114 = 1 \times 99 + 15 \qquad \Rightarrow \qquad 15 = 114 - 1 \times 99 \qquad (2)$$
$$99 = 6 \times 15 + 9 \qquad \Rightarrow \qquad 9 = 99 - 6 \times 15 \qquad (3)$$
$$15 = 1 \times 9 + 6 \qquad \Rightarrow \qquad 6 = 15 - 1 \times 9 \qquad (4)$$
$$9 = 1 \times 6 + 3 \qquad \Rightarrow \qquad 3 = 9 - 1 \times 6 \qquad (5)$$
$$6 = 2 \times 3 + 0$$

We can then express 3 in the form $327u + 114v$ by successively substituting the equations into each other:

- Equation (5) expresses 3 as a linear combination of 6 and 9. Substituting equation (4) yields:

$$3 = 9 - 1 \times (15 - 1 \times 9) \quad \Rightarrow \quad 3 = 2 \times 9 - 1 \times 15$$

- This now expresses 3 as a linear combination of 9 and 15. Substituting equation (3) yields:

$$3 = 2 \times (99 - 6 \times 15) - 1 \times 15 \quad \Rightarrow \quad 3 = (-13) \times 15 + 2 \times 99$$

- This now expresses 3 as a linear combination of 15 and 99. Substituting equation (2) yields:

$$3 = (-13) \times (114 - 1 \times 99) + 2 \times 99 \quad \Rightarrow \quad 3 = 15 \times 99 - 13 \times 114$$

- This now expresses 3 as a linear combination of 99 and 114. Substituting equation (1) yields:

$$3 = 15 \times (327 - 2 \times 114) - 13 \times 114 \quad \Rightarrow \quad 3 = (-43) \times 114 + 15 \times 327$$

Now that we've expressed 3 as a linear combination of 114 and 327, we're nearly done: we know that $18 = 6 \times 3$, so multiplying through by 6 gives

$$18 = (-258) \times 114 + 90 \times 327$$

Hence $(x, y) = (90, -258)$ is a solution to the equation $327x + 114y = 18$. ◁

Proof tip

Let $a, b \in \mathbb{Z}$ and let $d = \gcd(a, b)$. To find integers x, y such that $ax + by = d$:

(i) Run the Euclidean algorithm on the pair (a, b), keeping track of all quotients and remainders.

(ii) Rearrange each equation of the form $r_{n-2} = q_n r_{n-1} + r_n$ to isolate r_n.

(iii) Substitute for the remainders r_k in reverse order until $\gcd(a,b)$ is expressed in the form $ax + by$ for some $x, y \in \mathbb{Z}$.

This process is called the **reverse Euclidean algorithm**.   ◁

Exercise 4.1.33
Find a solution $(x,y) \in \mathbb{Z} \times \mathbb{Z}$ to the equation $630x + 385y = 4340$.   ◁

Now that we have a procedure for computing *one* solution to the equation $ax + by = c$, we need to come up with a procedure for computing *all* solutions. This can be done by proving the following theorem.

Theorem 4.1.34
Let $a, b, c \in \mathbb{Z}$, where a and b are not both zero. Suppose that x_0 and y_0 are integers such that $ax_0 + by_0 = c$. Then, $(x,y) \in \mathbb{Z} \times \mathbb{Z}$ is another solution to the equation $ax + by = c$ if and only if

$$x = x_0 + k \cdot \frac{b}{\gcd(a,b)} \qquad \text{and} \qquad y = y_0 - k \cdot \frac{a}{\gcd(a,b)}$$

for some $k \in \mathbb{Z}$.

Thus, as soon as we've found one solution $(x,y) = (x_0, y_0)$ to the equation $ax + by = c$, this theorem tells us what all other solutions must look like.

Proof of Theorem 4.1.34
We prove the two directions separately.

($\Rightarrow$). First suppose that (x_0, y_0) is an integer solution to the equation $ax + by = c$. Let $k \in \mathbb{Z}$ and let

$$x = x_0 + k \cdot \frac{b}{\gcd(a,b)} \qquad \text{and} \qquad y = y_0 - k \cdot \frac{a}{\gcd(a,b)}$$

Then

$$ax + by$$

$$= a \left(x_0 + k \cdot \frac{b}{\gcd(a,b)} \right) + b \left(y_0 - k \cdot \frac{a}{\gcd(a,b)} \right) \quad \text{by definition of } x \text{ and } y$$

$$= (ax_0 + by_0) + ak \cdot \frac{b}{\gcd(a,b)} - kb \cdot \frac{a}{\gcd(a,b)} \quad \text{rearranging}$$

$$= (ax_0 + by_0) + \frac{kab - kab}{\gcd(a,b)} \quad \text{combining the fractions}$$

$$= ax_0 + by_0 \quad \text{since } kab - kab = 0$$

$$= c \quad \text{since } (x_0, y_0) \text{ is a solution}$$

so (x,y) is indeed a solution to the equation.

($\Leftarrow$). First suppose that $a \perp b$. Fix a solution (x_0, y_0) to the equation $ax + by = c$, and let (x,y) be another solution. Then

$$a(x - x_0) + b(y - y_0) = (ax_0 + by_0) - (ax + by) = c - c = 0$$

so that

$$a(x - x_0) = b(y_0 - y)$$

Now a and b are coprime, so by [Proposition 4.1.31](), we have $a \mid y_0 - y$ and $b \mid x - x_0$. Let $k, \ell \in \mathbb{Z}$ be such that $x - x_0 = kb$ and $y_0 - y = \ell a$. Then substituting into the above equation yields

$$a \cdot kb = b \cdot \ell a$$

and hence $(k - \ell)ab = 0$. Since $ab \neq 0$, we have $k = \ell$, so that

$$x = x_0 + kb \quad \text{and} \quad y = y_0 - ka$$

Now we drop the assumption that $a \perp b$. Let $\gcd(a,b) = d \geq 1$. We know that $d \mid c$, by Bézout's lemma ([Theorem 4.1.22]()), and so

$$\frac{a}{d}x + \frac{b}{d}y = \frac{c}{d}$$

is another linear Diophantine equations, and moreover $\frac{a}{d} \perp \frac{b}{d}$ by [Exercise 4.1.29](). By what we proved above, we have

$$x = x_0 + k \cdot \frac{b}{d} \quad \text{and} \quad y = y_0 - k \cdot \frac{a}{d}$$

for some $k \in \mathbb{Z}$. But this is exactly what we sought to prove! $\qquad\square$

Example 4.1.35

We know that $(x,y) = (90, -258)$ is a solution to the equation $327x + 114y = 18$, and

$$\frac{327}{\gcd(327,114)} = \frac{327}{3} = 109 \quad \text{and} \quad \frac{114}{\gcd(327,114)} = \frac{114}{3} = 38$$

so this theorem tells us that $(x,y) \in \mathbb{Z} \times \mathbb{Z}$ is a solution to the equation $327x + 114y = 18$ if and only if

$$x = 90 + 38k \quad \text{and} \quad y = -258 - 109k$$

for some $k \in \mathbb{Z}$.   ◁

Exercise 4.1.36

Find all integers x, y such that

$$630x + 385y = 4340$$

◁

Least common multiples

You would be forgiven for wondering why so much of the foregoing section was devoted to greatest common divisors, with no mention of least common multiples. We will now give the latter some attention.

Definition 4.1.37

Let $a, b \in \mathbb{Z}$. An integer m is a **least common multiple** of a and b if:

(a) $a \mid m$ and $b \mid m$;

(b) If n is another integer such that $a \mid n$ and $b \mid n$, then $m \mid n$.

In a sense that can be made precise, the definition of least common multiple is *dual* to that of greatest common divisor (Definition 4.1.9).[a] This means that many properties of greatest common divisors have corresponding 'dual' properties, which hold of least common multiples. As such, we will not say much here about least common multiples, and that which we *do* say is in the form of exercises.

---

[a]Specifically, we refer here to the dual of a *preorder*, i.e. a reflexive, transitive relation—see Chapter 5 for more on this!

Exercise 4.1.38

Let $a, b \in \mathbb{Z}$. Prove that a and b have a least common multiple. Furthermore, prove that least common multiples are unique up to sign, in the sense that if m, m' are two least common multiples of a and b, then $m = m'$ or $m = -m'$. ◁

As with greatest common divisors, Exercise 4.1.38 justifies the following definition.

Definition 4.1.39

Given $a, b \in \mathbb{Z}$, denote by $\mathrm{lcm}(a, b)$ (LaTeX code: `\mathrm{lcm}`) the non-negative least common multiple of a and b.

Exercise 4.1.40

Let $a, b \in \mathbb{Z}$. Prove that $\gcd(a, b) \cdot \mathrm{lcm}(a, b) = |ab|$. ◁

Section 4.2
Prime numbers

Thinking of divisibility as a way of *breaking down* an integer, for example $12 = 2 \times 2 \times 3$, our goal now is to show that:

- There are numbers which are *atomic*, in the sense that they can't be broken down any further by division;

- ... and every nonzero nonunit can be written as a product of these atomic numbers;

- ... *and* this product is essentially unique.

There are a couple of fairly vague terms used here: 'atomic' and 'essentially unique'. We will soon make these precise; the atomic numbers will be the *irreducible* and *prime* numbers (two notions which coincide for the integers), and 'essentially unique' will mean unique up to reordering and multiplication by units.

Primes and irreducibles

Definition 4.2.1
Let p be a nonzero nonunit. We say p is **prime** if for all $a, b \in \mathbb{Z}$, if $p \mid ab$ then $p \mid a$ or $p \mid b$.

Example 4.2.2
Here are some examples of prime and non-prime numbers:

- 2 is prime. Suppose not; then there exist $a, b \in \mathbb{Z}$ such that $2 \mid ab$ but 2 divides neither a nor b. Thus a and b are both odd, meaning that ab is odd... but this contradicts the assumption that $2 \mid ab$.

- 6 is not prime. Indeed, $6 \mid 2 \times 3$ but 6 divides neither 2 nor 3.

$\triangleleft$

Exercise 4.2.3
Using Definition 4.2.1, prove that 3 and 5 are prime and that 4 is not prime. $\triangleleft$

For the following example and exercise, you will need to recall the definitions of *binomial coefficients* $\binom{n}{k}$ and *factorials* $n!$, which we studied from the point of view of induction in Section 3.1, and then redefined and studied from the point of

view of combinatorics in Section 3.3. In case you skipped over Section 3.3, we provide references to the relevant results in both sections.

Example 4.2.4

Let $k \in \mathbb{Z}$ with $0 < k < 5$. We'll show that $5 \mid \binom{5}{k}$.

Well, by Theorems 3.1.32 and 3.3.40 we know that

$$5! = \binom{5}{k} k!(5-k)!$$

By Definition 3.1.27 and Theorem 3.3.39, we have

$$\underbrace{5 \times 4!}_{=5!} = \binom{5}{k} \times \underbrace{1 \times \cdots \times k}_{=k!} \times \underbrace{1 \times \cdots \times (5-k)}_{=(5-k)!}$$

Since 5 is prime, it must divide one of the factors on the right-hand side of this equation. Thus, either 5 divides $\binom{5}{k}$, or it divides ℓ for some $1 \leqslant \ell \leqslant k$ or $1 \leqslant \ell \leqslant 5 - k$. But $k < 5$ and $5 - k < 5$, so it cannot divide any of these values of ℓ—if it did, it would imply $5 \leqslant \ell \leqslant k$ or $5 \leqslant \ell \leqslant 5 - k$, which is nonsense. Hence 5 must divide $\binom{5}{k}$. ◁

Exercise 4.2.5

Let $p \in \mathbb{Z}$ be a positive prime and let $0 < k < p$. Show that $p \mid \binom{p}{k}$. ◁

Aside

Most people are introduced to primes with a definition along the lines of 'p is prime if p has exactly two positive divisors'. We have avoided this to elucidate the fact that the integers together with their arithmetic structure are the canonical example of a mathematical object called a *ring*. The notion of a *prime element* can be defined in any ring as in Definition 4.2.1. Secondly, these two definitions are equivalent in $\mathbb{Z}$, but not in all rings. ◁

Definition 4.2.6

Let a be a nonzero nonunit. We say a is **reducible** if $a = mn$ for some nonunits m, n; otherwise it is **irreducible**.

Proposition 4.2.7

A nonzero nonunit p is irreducible if and only if the only divisors of p are p, $-p$, 1 and -1.

Proof

Suppose p is irreducible and that $a \mid p$. Then $p = ab$ for some $b \in \mathbb{Z}$. Since p is irreducible, either a or b is a unit. If a is a unit then $b = \pm p$, and if b is a unit then $a = \pm p$. So the only divisors of p are ± 1 and $\pm p$.

Conversely, suppose that the only divisors of p are ± 1 and $\pm p$, and let $a, b \in \mathbb{Z}$ with $p = ab$. We want to prove that a or b is a unit. Since $a \mid p$, we have $a \in \{1, -1, p, -p\}$. If $a = \pm 1$, then a is a unit; if $a = \pm p$, then $b = \pm 1$, so that b is a unit. In any case, either a or b is a unit, and hence p is irreducible. □

Example 4.2.8

A couple of examples of reducible and irreducible numbers are:

- 2 is irreducible: if $2 = mn$ then either m or n is even, otherwise we'd be expressing an even number as the product of two odd numbers. We may assume m is even, say $m = 2k$; then $2 = 2kn$, so $kn = 1$ and hence n is a unit.

- 6 is reducible since $6 = 2 \times 3$ and both 2 and 3 are nonzero nonunits.

◁

Exercise 4.2.9

Prove that if $p \in \mathbb{Z}$ is prime then p is irreducible.   ◁

Lemma 4.2.10

Let $a \in \mathbb{Z}$ be a nonzero nonunit. Then there are irreducibles $p_1, \ldots, p_n$ such that $a = p_1 \times \cdots \times p_n$.

Proof

We may assume $a > 0$, since if $a < 0$ we can just multiply by -1.

We proceed by strong induction on $a \geqslant 2$. The base case has $a = 2$ since we consider only nonunits.

- **(BC)** We have shown that 2 is irreducible, so setting $p_1 = 2$ yields a product of primes.

- **(IS)** Let $a \geqslant 2$ and suppose that each integer k with $2 \leqslant k \leqslant a$ has an expression as a product of irreducibles. If $a + 1$ is irreducible then we're done; otherwise we can write $a + 1 = st$, where $s, t \in \mathbb{Z}$ are nonzero nonunits. We may assume further that s and t are positive. Moreover, $s < a + 1$ and $t < a + 1$ since $s, t \geqslant 2$.

  By the induction hypothesis, s and t have expressions as products of irreducibles. Write
  $$s = p_1 \times \cdots \times p_m, \qquad t = q_1 \times \cdots \times q_n$$

  This gives rise to an expression of a as a product of irreducibles:
  $$a = st = \underbrace{p_1 \times \cdots \times p_m}_{=s} \times \underbrace{q_1 \times \cdots \times q_n}_{=t}$$

By induction, we're done.   □

> **Theorem 4.2.11**
> Let $p \in \mathbb{Z}$. Then p is prime if and only if p is irreducible.

Proof
We prove the two directions separately.

- **Prime $\Rightarrow$ irreducible.** This was Exercise 4.2.9.

- **Irreducible $\Rightarrow$ prime.** Suppose p is irreducible. Let $a, b \in \mathbb{Z}$ and suppose $p \mid ab$. We need to show that $p \mid a$ or $p \mid b$. It suffices to show that if $p \nmid a$ then $p \mid b$.

  So suppose $p \nmid a$. Let $d = \gcd(p, a)$. Since $d \mid p$ and p is irreducible, we must have $d = 1$ or $d = p$ by Proposition 4.2.7. Since $p \nmid a$ and $d \mid a$, we must therefore have $d = 1$.

  By Bézout's lemma (Theorem 4.1.22), there exist $u, v \in \mathbb{Z}$ such that $au + pv = 1$. Multiplying by b gives $abu + pbv = b$. Since $p \mid ab$, there exists $k \in \mathbb{Z}$ such that $pk = ab$. Then

$$b = abu + pbv = pku + pbv = p(ku + bv)$$

  so $p \mid b$, as required.

So we're done. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

Since primes and irreducibles are the same thing in $\mathbb{Z}$, we will refer to them as 'primes', unless we need to emphasise a particular aspect of them.

Prime factorisation

Having described prime numbers in two ways, each of which emphasises their nature of being 'unbreakable' by multiplication, we will extend Lemma 4.2.10 to prove that every integer can be expressed as a product of primes in an essentially unique way.

> **Theorem 4.2.12** (Fundamental theorem of arithmetic)
> Let $a \in \mathbb{Z}$ be a nonzero nonunit. There exist primes $p_1, \ldots, p_k \in \mathbb{Z}$ such that
>
> $$a = p_1 \times \cdots \times p_k$$
>
> Moreover, this expression is essentially unique: if $a = q_1 \times \cdots \times q_\ell$ is another expression of a as a product of primes, then $k = \ell$ and, re-ordering the q_i if necessary, for each i there is a unit u_i such that $q_i = u_i p_i$.

Proof

We showed that such a factorisation exists in Lemma 4.2.10, with the word 'prime' replaced by the word 'irreducible'. It remains to prove (essential) uniqueness.

Let k be least such that there is an expression of a as a product of k primes, namely $a = p_1 \times \cdots \times p_k$. Let $a = q_1 \times \cdots \times q_\ell$ be any other such expression. We prove by induction on k that $\ell = k$ and, after re-ordering if necessary, for each i there is a unit u_i such that $q_i = u_i p_i$.

- **(BC)** If $k = 1$ then $a = p_1$ is itself prime. Then we have $p_1 = q_1 \times \cdots \times q_\ell$. Since p_1 is prime, $p_1 \mid q_j$ for some j; by relabelling q_1 and q_j we may assume that $j = 1$, so that $p_1 \mid q_1$. By irreducibility of q_1 we have $q_1 = u_1 p_1$ for some unit u_1.

- **(IS)** Let $k \geqslant 1$ and suppose that any integer which can be expressed as a product of k primes is (essentially) uniquely expressible in such a way. Suppose a has an expression as a product of $k + 1$ primes, and that $k + 1$ is the least such number. Suppose also that

$$a = p_1 \times \cdots \times p_k \times p_{k+1} = q_1 \times \cdots \times q_\ell$$

  Note that $\ell \geqslant k + 1$. Since p_{k+1} is prime we must have $p_{k+1} \mid q_j$ for some j; by relabelling q_j and q_ℓ if necessary, we may assume that $j = \ell$, so that $p_{k+1} \mid q_\ell$. As before, $q_\ell = u_{k+1} p_{k+1}$ for some unit u_{k+1}. Dividing through by p_{k+1} gives

$$p_1 \times \cdots \times p_k = q_1 \times \cdots \times q_{\ell-1} \times u_{k+1}$$

  Replacing $q_{\ell-1}$ by $q_{\ell-1} u_{k+1}$, which is still prime, we can apply the induction hypothesis to obtain $k = \ell - 1$, so $k + 1 = \ell$, and, after reordering if necessary $q_i = u_i p_i$ for all $i \leqslant k$. Since this also holds for $i = k + 1$, we're done.

By induction, we're done. □

Example 4.2.13

Here are some examples of numbers written as products of primes:

- $12 = 2 \times 2 \times 3$. We could also write this as $2 \times 3 \times 2$ or $(-2) \times (-3) \times 2$, and so on.

- $53 = 53$ is an expression of 53 as a product of primes.

- $-1000 = 2 \times 5 \times (-2) \times 5 \times 2 \times 5$.

- We can view any unit as a product of *no* primes. (Don't dwell on this point for too long as it will not arise very often!)

◁

Exercise 4.2.14

Express the following numbers as products of primes:

$$16 \quad -240 \quad 5050 \quad 111111 \quad -123456789$$

◁

To make things slightly more concise, we introduce a standard way of expressing a number as a product of primes:

Definition 4.2.15

The **canonical prime factorisation** of a nonzero integer a is the expression in the form

$$a = up_1^{j_1} \cdots p_r^{j_r}$$

where $r \geqslant 0$ and:

- $u = 1$ if $a > 0$, and $u = -1$ if $a < 0$;

- The numbers p_i are all positive primes;

- $p_1 < p_2 < \cdots < p_r$;

- $j_i \geqslant 1$ for all i.

We call j_i the **multiplicity** of p_i in a, and we call u the **sign** of a.

Typically we omit u if $u = 1$ (unless $a = 1$), and just write a minus sign $(-)$ if $u = -1$.

Example 4.2.16

The canonical prime factorisations of the integers given in Example 4.2.13 are:

- $12 = 2^2 \cdot 3$.

- $53 = 53$.

- $-1000 = -2^3 \cdot 5^3$.

- $1 = 1$.

◁

Exercise 4.2.17

Write out the canonical prime factorisations of the numbers from Exercise 4.2.14, which were:

$$16 \qquad -240 \qquad 5050 \qquad 111111 \qquad -123456789$$

◁

The following exercise provides another tool for computing greatest common divisors of pairs of integers by looking at their prime factorisations.

Exercise 4.2.18

Let $p_1, p_2, \ldots, p_r$ be distinct primes, and let $k_i, \ell_i \in \mathbb{N}$ for all $1 \leqslant i \leqslant r$. Define

$$m = p_1^{k_1} \times p_2^{k_2} \times \cdots \times p_r^{k_r} \quad \text{and} \quad n = p_1^{\ell_1} \times p_2^{\ell_2} \times \cdots \times p_r^{\ell_r}$$

Prove that

$$\gcd(m,n) = p_1^{u_1} \times p_2^{u_2} \times \cdots \times p_r^{u_r}$$

where $u_i = \min\{k_i, \ell_i\}$ for all $1 \leqslant i \leqslant r$. ◁

Example 4.2.19

We use Exercise 4.2.18 to compute the greatest common divisor of 17640 and 6468.

First we compute the prime factorisations of 17640 and 6468:

$$17640 = 2^3 \cdot 3^2 \cdot 5 \cdot 7^2 \quad \text{and} \quad 6468 = 2^2 \cdot 3 \cdot 7^2 \cdot 11$$

It now follows from Exercise 4.2.18 that

$$\gcd(17640, 6468) = 2^2 \cdot 3^1 \cdot 5^0 \cdot 7^2 \cdot 11^0$$
$$= 4 \cdot 3 \cdot 1 \cdot 49 \cdot 1$$
$$= 588$$

◁

Exercise 4.2.18 allows us to provide a concise proof of the following result.

Corollary 4.2.20

Let $p \in \mathbb{Z}$ be prime, let $a \in \mathbb{Z}$ be nonzero, and let $k \geqslant 1$. Then $a \perp p^k$ if and only if $p \nmid a$.

Proof

By the fundamental theorem of arithmetic, we can write

$$a = p^j \times p_1^{j_1} \times \cdots \times p_r^{j_r}$$

where $p_1, \ldots, p_r$ are the primes other than p appearing in the prime factorisation of a, and $j, j_i \in \mathbb{N}$ for all $1 \leqslant i \leqslant r$. Note that $p \mid a$ if and only if $j \geqslant 1$.

Furthermore we have

$$p^k = p^k \times p_1^0 \times \cdots \times p_r^0$$

By Exercise 4.2.18 it follows that

$$\gcd(a, p^k) = p^{\min\{j,k\}} \times p_1^0 \times \cdots \times p_r^0 = p^{\min\{j,k\}}$$

Now:

- If $\min\{j,k\} = 0$ then $j = 0$, in which case $p \nmid a$, and $\gcd(a, p^k) = p^0 = 1$;

- If $\min\{j,k\} > 0$ then $j \geqslant 1$, in which case $p \mid a$, and $p \mid \gcd(a, p^k)$, so $\gcd(a, p^k) \neq 1$.

In particular, $p \nmid a$ if and only if $a \perp p^k$. □

Distribution of primes

So far we have seen several examples of prime numbers; to name a few, we've seen 2, 3, 5 and 53. It might seem like the prime numbers go on forever, but proving this is less than obvious.

Exercise 4.2.21
Let P be an inhabited finite set of positive prime numbers and let m be the product of all the elements of P. That is, for some $n \geqslant 1$ let

$$P = \{p_1, \ldots, p_n\} \quad \text{and} \quad m = p_1 \times \cdots \times p_n$$

where each $p_k \in P$ is a positive prime. Using the fundamental theorem of arithmetic, show that $m + 1$ has a positive prime divisor which is not an element of P. ◁

Theorem 4.2.22
There are infinitely many primes.

Proof
We prove that there are infinitely many *positive* prime numbers—the result then follows immediately. Let P be the set of all positive prime numbers. Then P is inhabited, since $2 \in P$, for example. If P were finite, then by Exercise 4.2.21, there would be a positive prime which is not an element of P—but P contains all positive primes, so that is impossible. Hence there are infinitely many positive primes. □

This is one proof of many and is attributed to Euclid, who lived around 2300 years ago. We might hope that a proof of the infinitude of primes gives some insight into how the primes are *distributed*. That is, we might ask questions like: how frequently do primes occur? How fast does the sequence of primes grow? How likely is there to be a prime number in a given set of integers?

As a starting point, Euclid's proof gives an algorithm for writing an infinite list of primes:

- Let $p_1 = 2$; we know that 2 is prime;

- Given $p_1, \ldots, p_n$, let p_{n+1} be the smallest positive prime factor of $p_1 \times \cdots \times p_n + 1$.

The first few terms produced would be:

- $p_1 = 2$ by definition;

- $2 + 1 = 3$, which is prime, so $p_2 = 3$;

- $2 \times 3 + 1 = 7$, which is prime, so $p_3 = 7$;

- $2 \times 3 \times 7 + 1 = 43$, which is prime, so $p_4 = 43$;

- $2 \times 3 \times 7 \times 43 + 1 = 1807 = 13 \times 139$, so $p_5 = 13$;

- $2 \times 3 \times 7 \times 43 \times 13 + 1 = 23479 = 53 \times 443$, so $p_6 = 53$;

- ... and so on.

The sequence obtained, called the *Euclid–Mullin sequence*, is a bit bizarre:

$$2, 3, 7, 43, 13, 53, 5, 6221671, 38709183810571, 139, 2801, 11, 17, 5471, \ldots$$

Big primes like 38709183810571 often appear before small primes like 11. It remains unknown whether or not every positive prime number appears in this list!

The chaotic nature of this sequence makes it difficult to extract information about how the primes are distributed: the numbers $p_1 \times \cdots \times p_n + 1$ grow very quickly— indeed, it must be the case that $p_1 \times \cdots \times p_n + 1 > 2^n$ for all n—so the upper bounds for the sequence grow at least exponentially.

Another proof of the infinitude of primes that gives a (slightly) tighter bound can be obtained using the following exercise.

Exercise 4.2.23
Let $n \in \mathbb{Z}$ with $n > 2$. Prove that the set $\{k \in \mathbb{Z} \mid n < k < n!\}$ contains a prime number. ◁

Section 4.3
Modular arithmetic

It turns out that much arithmetic can be done by considering only the *remainders* of integers when divided by a fixed integer. Here is a simple example:

Example 4.3.1

Suppose a_1 has remainder r_1 and a_2 has remainder r_2 when divided by 7. That is, there exist $q_1, q_2 \in \mathbb{Z}$ such that

$$a_1 = 7q_1 + r_1 \qquad \text{and} \qquad a_2 = 7q_2 + r_2$$

Then $a_1 + a_2$ has the same remainder as $r_1 + r_2$ when divided by 7. Indeed, suppose $a_1 + a_2 = 7q + r$, where $0 \leqslant r < 7$. Then

$$\begin{aligned}
r_1 + r_2 &= (a_1 - 7q_1) + (a_2 - 7q_2) \\
&= (a_1 + a_2) + 7(-q_1 - q_2) \\
&= (7q + r) + 7(-q_1 - q_1) \\
&= 7(q - q_1 - q_2) + r
\end{aligned}$$

An example of this in action: $41 = 5 \times 7 + 6$ and $240 = 34 \times 7 + 2$, so the remainders of 41 and 240 when divided by 7 are 6 and 2, respectively. Now

$$41 + 240 = 281 = 40 \times 7 + 1 \qquad \text{and} \qquad 6 + 2 = 8 = 1 \times 7 + 1$$

which demonstrates that $41 + 240$ and $6 + 2$ have the same remainder when divided by 7. ◁

In this section we will study the extent to which we can do arithmetic with integers knowing only their remainders upon division by a given integer.

Definition 4.3.2

Fix $n \in \mathbb{Z}$. Given integers $a, b \in \mathbb{Z}$, we say a is **congruent** to b **modulo** n, and write

$$a \equiv b \bmod n \ (\text{\LaTeX code: } \texttt{a \textbackslash equiv b \textbackslash bmod\{n\}})$$

if $n \mid a - b$. If a is not congruent to b modulo n, write

$$a \not\equiv b \bmod n \ (\text{\LaTeX code: } \texttt{\textbackslash not\textbackslash equiv})$$

The number n is called the **modulus** of the congruence.

Convention 4.3.3

When talking about modular arithmetic, we will restrict our attention to *positive* integers. This is because for any integers a, b, n we have

$$a \equiv b \bmod n \quad \Leftrightarrow \quad a \equiv b \bmod (-n)$$

and $a \equiv b \bmod 0$ if and only if $a = b$. Thus, whenever we write ' mod n' or specify that a variable n is a 'modulus', it is implicit that n is an integer and $n > 0$. This will shorten some of our proofs. ◁

Example 4.3.4

Some examples of congruence modulo n are as follows:

- $16 \equiv 30 \bmod 2$ since $30 - 16 = 14$, which is a multiple of 2.

- $44 \equiv 20 \bmod 6$ since $20 - 44 = -24$, which is a multiple of 6.

◁

Exercise 4.3.5

Show that if $a, b \in \mathbb{Z}$ with $a, b \geqslant 0$ then $a \equiv b \bmod 10$ if and only if the decimal expressions of a and b end in the same digit. What happens when a and b are allowed to be negative? ◁

It is important from the outset to point out that, although congruence is written with a symbol that looks like that of equality ('$\equiv$' vs. '$=$'), we can only treat congruence like equality inasmuch as we have proved we can. Specifically, the ways in which congruence *can* be treated like equality will be proved in two theorems:

- Theorem 4.3.6 tells us that congruence satisfies three extremely basic properties of equality.[b] One useful consequence of this is that it is valid to use strings of congruences, for example

$$-5 \equiv 18 \equiv 41 \equiv 64 \bmod 23 \quad \Rightarrow \quad -5 \equiv 64 \bmod 23$$

- Theorem 4.3.9 tells us that we can treat congruence like equality for the purposes of addition, multiplication and subtraction. Thus it will be valid to write things like
$$x \equiv 7 \bmod 12 \quad \Rightarrow \quad 2x + 5 \equiv 19 \bmod 12$$

and we'll be able to replace values by congruent values in congruences, provided they're only being added, subtracted or multiplied. For example,

---

[b]Using the language of Definition 5.1.30, Theorem 4.3.6 says precisely that congruence is an *equivalence relation*.

from the knowledge that $2^{60} \equiv 1 \bmod 61$ and $60! \equiv -1 \bmod 61$, we will be able to deduce

$$2^{60} \cdot 3 \equiv 60! \cdot x \bmod 61 \quad \Rightarrow \quad 3 \equiv -x \bmod 61$$

Don't let these properties shared by congruence and equality lull you into a false sense of security! We will soon see that for other purposes, such as division and various other algebraic operations, congruence does *not* behave like equality.

Theorem 4.3.6

Let $a, b, c \in \mathbb{Z}$ and let n be a modulus. Then

(a) $a \equiv a \bmod n$;

(b) If $a \equiv b \bmod n$, then $b \equiv a \bmod n$;

(c) If $a \equiv b \bmod n$ and $b \equiv c \bmod n$, then $a \equiv c \bmod n$.

Proof

(a) Note that $a - a = 0$, which is divisible by n since $0 = 0 \times n$, and hence $a \equiv a \bmod n$.

(b) Suppose $a \equiv b \bmod n$. Then $n \mid a - b$, so that $a - b = kn$ for some $k \in \mathbb{Z}$. Hence $b - a = -kn$, and so $n \mid b - a$, so that $b \equiv a \bmod n$ as required.

(c) Suppose $a \equiv b \bmod n$ and $b \equiv c \bmod n$. Then $n \mid a - b$ and $n \mid b - c$, so there exist $k, \ell \in \mathbb{Z}$ such that

$$a - b = kn \quad \text{and} \quad b - c = \ell n$$

Hence $a - c = (a - b) + (b - c) = (k + \ell)n$, so that $n \mid a - c$. Hence $a \equiv c \bmod n$, as required.

$\square$

There is a slightly simpler characterisation of congruence modulo n, as seen in Proposition 4.3.7 below.

Proposition 4.3.7

Fix a modulus n and let $a, b \in \mathbb{Z}$. The following are equivalent:

(i) a and b leave the same remainder when divided by n;

(ii) $a = b + kn$ for some $k \in \mathbb{Z}$;

(iii) $a \equiv b \bmod n$.

Proof

We prove (i) $\Leftrightarrow$ (iii) and (ii) $\Leftrightarrow$ (iii).

- (i) $\Rightarrow$ (iii). Suppose a and b leave the same remainder when divided by n, and let $q_1, q_2, r \in \mathbb{Z}$ be such that

$$a = q_1 n + r, \quad b = q_2 n + r \quad \text{and} \quad 0 \leqslant r < n$$

  Then $a - b = (q_1 - q_2)n$, which proves that $n \mid a - b$, and so $a \equiv b \bmod n$.

- (iii) $\Rightarrow$ (i). Suppose that $a \equiv b \bmod n$, so that $b - a = qn$ for some $q \in \mathbb{Z}$. Write

$$a = q_1 n + r_1, \quad b = q_2 n + r_2 \quad \text{and} \quad 0 \leqslant r_1, r_2 < n$$

  We may further assume that $r_1 \leqslant r_2$. (If not, swap the roles of a and b—this is fine, since $n \mid b - a$ if and only if $n \mid a - b$.) Now we have

$$\begin{aligned} b - a = qn &\Rightarrow (q_2 n + r_2) - (q_1 n + r_1) = qn \\ &\Rightarrow (q_2 - q_1 - q)n + (r_2 - r_1) = 0 \qquad \text{rearranging} \end{aligned}$$

  since $0 \leqslant r_1 \leqslant r_2 < n$ we have $0 \leqslant r_2 - r_1 < n$, so that $r_2 - r_1$ is the remainder of 0 when divided by n. That is, $r_2 - r_1 = 0$, so $r_1 = r_2$. Hence a and b have the same remainder when divided by n.

- (ii) $\Leftrightarrow$ (iii). We unpack the definitions of (ii) and (iii) to see that they are equivalent. Indeed

$$\begin{aligned} \text{(ii)} &\Leftrightarrow a = b + kn \text{ for some } k \in \mathbb{Z} \\ &\Leftrightarrow a - b = kn \text{ for some } k \in \mathbb{Z} \qquad \text{rearranging} \\ &\Leftrightarrow n \mid a - b \qquad\qquad\qquad\quad \text{by definition of divisibility} \\ &\Leftrightarrow a \equiv b \bmod n \qquad\qquad\quad\; \text{by definition of congruence} \\ &\Leftrightarrow \text{(iii)} \end{aligned}$$

$\square$

Discussion 4.3.8

Where in the proof of Proposition 4.3.7 did we rely on the convention that the modulus n is positive? Is the result still true if n is negative? $\triangleleft$

The following theorem tells us that, in a very limited sense, the $\equiv$ symbol can be treated as a $=$ symbol for the purposes of doing addition, subtraction and multiplication. Emphatically, it does *not* say that we can treat '$\equiv$' like '$=$' for the purposes of doing *division*.

Theorem 4.3.9 (Modular arithmetic)

Fix a modulus n, and let $a_1, a_2, b_1, b_2 \in \mathbb{Z}$ be such that

$$a_1 \equiv b_1 \bmod n \qquad \text{and} \qquad a_2 \equiv b_2 \bmod n$$

Then the following congruences hold:

(a) $a_1 + a_2 \equiv b_1 + b_2 \bmod n$;

(b) $a_1 a_2 \equiv b_1 b_2 \bmod n$;

(c) $a_1 - a_2 \equiv b_1 - b_2 \bmod n$.

Proof

By Definition 4.3.2 that $n \mid a_1 - b_1$ and $n \mid a_2 - b_2$, so there exist $q_1, q_2 \in \mathbb{Z}$ such that

$$a_1 - b_1 = q_1 n \qquad \text{and} \qquad a_2 - b_2 = q_2 n$$

This implies that

$$(a_1 + a_2) - (b_1 + b_2) = (a_1 - b_1) + (a_2 - b_2) = q_1 n + q_2 n = (q_1 + q_2) n$$

so $n \mid (a_1 + a_2) - (b_1 + b_2)$. This proves (a).

The algebra for (b) is slightly more involved:

$$\begin{aligned}
a_1 a_2 - b_1 b_2 &= (q_1 n + b_1)(q_2 n + b_2) - b_1 b_2 \\
&= q_1 q_2 n^2 + b_1 q_2 n + b_2 q_1 n + b_1 b_2 - b_1 b_2 \\
&= q_1 q_2 n^2 + b_1 q_2 n + b_2 q_1 n \\
&= (q_1 q_2 n + b_1 q_2 + b_2 q_1) n
\end{aligned}$$

This shows that $n \mid a_1 a_2 - b_1 b_2$, thus proving (b).

Now (a) and (b) together imply (c). Indeed, we know that $-1 \equiv -1 \bmod n$ and $b_1 \equiv b_2 \bmod n$, so by (b) we have $-b_1 \equiv -b_2 \bmod n$. We also know that $a_1 \equiv a_2 \bmod n$, and hence $a_1 - b_1 \equiv a_2 - b_2 \bmod n$ by (a). $\qquad\square$

Theorem 4.3.9 allows us to perform algebraic manipulations with congruences as if they were equations, provided all we're doing is adding, multiplying and subtracting.

Example 4.3.10

We will solve the congruence $3x - 5 \equiv 2x + 3 \bmod 7$ for x:

$$3x - 5 \equiv 2x + 3 \bmod 7$$

$\Leftrightarrow$	$x - 5 \equiv 3 \bmod 7$	$(\Rightarrow)$ subtract $2x$	$(\Leftarrow)$ add $2x$
$\Leftrightarrow$	$x \equiv 8 \bmod 7$	$(\Rightarrow)$ add 5	$(\Leftarrow)$ subtract 5
$\Leftrightarrow$	$x \equiv 1 \bmod 7$	since $8 \equiv 1 \bmod 7$	

So the integers x for which $3x - 5$ and $2x + 3$ leave the same remainder when divided by 7, are precisely the integers x which leave a remainder of 1 when divided by 7:

$$3x - 5 \equiv 2x + 3 \bmod 7 \qquad \Leftrightarrow \qquad x = 7q + 1 \text{ for some } q \in \mathbb{Z}$$

$\triangleleft$

Exercise 4.3.11
For which integers x does the congruence $5x + 1 \equiv x + 8 \bmod 3$ hold? Characterise such integers x in terms of their remainder when divided by 3. $\triangleleft$

So far this all feels like we haven't done very much: we've just introduced a new symbol $\equiv$ which behaves just like equality... but does it really? The following exercises should expose some more ways in which congruence *does* behave like equality, and some in which it *doesn't*.

Exercise 4.3.12
Fix a modulus n. Is it true that

$$a \equiv b \bmod n \quad \Rightarrow \quad a^k \equiv b^k \bmod n$$

for all $a, b \in \mathbb{Z}$ and $k \in \mathbb{N}$? If so, prove it; if not, provide a counterexample. $\triangleleft$

Exercise 4.3.13
Fix a modulus n. Is it true that

$$k \equiv \ell \bmod n \quad \Rightarrow \quad a^k \equiv a^\ell \bmod n$$

for all $k, \ell \in \mathbb{N}$ and $a \in \mathbb{Z}$? If so, prove it; if not, provide a counterexample. $\triangleleft$

Exercise 4.3.14
Fix a modulus n. Is it true that

$$qa \equiv qb \bmod n \quad \Rightarrow \quad a \equiv b \bmod n$$

for all $a, b, q \in \mathbb{Z}$ with $q \not\equiv 0 \bmod n$? If so, prove it; if not, provide a counterexample. $\triangleleft$

Common error

The false sense of security that Theorem 4.3.9 induces often leads students new to all this to the belief that $\equiv$ and $=$ are interchangeable concepts. This is emphatically *not* the case. In particular:

- Fractions don't make sense in modular arithmetic; for instance, it is invalid to say $2x \equiv 1 \bmod 5$ implies $x \equiv \frac{1}{2} \bmod 5$.

- Square roots don't make sense in modular arithmetic; for instance, it is invalid to say $x^2 \equiv 3 \bmod 4$ implies $x \equiv \pm\sqrt{3} \bmod 4$.

- Numbers in exponents cannot be replaced by congruent numbers; for instance, it is invalid to say $x^3 \equiv 2^3 \bmod 4$ implies $x \equiv 2 \bmod 4$.

$\triangleleft$

Multiplicative inverses

We made a big deal about the fact that fractions don't make sense in modular arithmetic. That is, it is invalid to say

$$2x \equiv 1 \bmod 5 \quad \Rightarrow \quad x \equiv \frac{1}{2} \bmod 5$$

Despite this, we can still make sense of 'division', provided we change what we mean when we say 'division'. Indeed, the congruence $2x \equiv 1 \bmod 5$ has a solution:

$$2x \equiv 1 \bmod 5$$
$$\Leftrightarrow \quad 6x \equiv 3 \bmod 5 \qquad (\Rightarrow) \text{ multiply by 3} \qquad (\Leftarrow) \text{ subtract 3}$$
$$\Leftrightarrow \quad x \equiv 3 \bmod 5 \qquad \text{since } 6 \equiv 1 \bmod 5$$

Here we didn't divide by 2, but we still managed to cancel the 2 by instead multiplying through by 3. For the purposes of solving the equation this had the same effect as division by 2 would have had if we were allowed to divide. The key here was that $2 \times 3 \equiv 1 \bmod 5$.

Definition 4.3.15
Fix a modulus n. Given $a \in \mathbb{Z}$, a **multiplicative inverse** for a modulo n is an integer u such that $au \equiv 1 \bmod n$.

Example 4.3.16
Some examples of multiplicative inverses are as follows:

- 2 is a multiplicative inverse of itself modulo 3, since $2 \times 2 \equiv 4 \equiv 1 \bmod 3$.

- 2 is a multiplicative inverse of 3 modulo 5, since $2 \times 3 \equiv 6 \equiv 1 \bmod 5$.

- 7 is also a multiplicative inverse of 3 modulo 5, since $3 \times 7 \equiv 21 \equiv 1 \bmod 5$.

- 3 has no multiplicative inverse modulo 6. Indeed, suppose $u \in \mathbb{Z}$ with $3u \equiv 1 \bmod 6$. Then $6 \mid 3u - 1$, so $3u - 1 = 6q$ for some $q \in \mathbb{Z}$. But then

$$1 = 3u - 6q = 3(u - 2q)$$

  which implies that $3 \mid 1$, which is nonsense.

◁

Knowing when multiplicative inverses exist is very important for solving congruences: if u is a multiplicative inverse for a modulo n, then we can solve equations of the form $ax \equiv b \bmod n$ extremely easily:

$$ax \equiv b \bmod n \quad \Rightarrow \quad x \equiv ub \bmod n$$

Exercise 4.3.17
For $n = 7, 8, 9, 10, 11, 12$, either find a multiplicative inverse for 6 modulo n, or show that no multiplicative inverse exists. Can you spot a pattern? ◁

Some authors write a^{-1} to denote multiplicative inverses. We refrain from this, since it suggests that multiplicative inverses are unique—but they're not, as you'll see in the following exercise.

Exercise 4.3.18
Let n be a modulus and let $a \in \mathbb{Z}$. Suppose that u is a multiplicative inverse for a modulo n. Prove that, for all $k \in \mathbb{Z}$, $u + kn$ is a multiplicative inverse for a modulo n. ◁

Proposition 4.3.19
Let $a \in \mathbb{Z}$ and let n be a modulus. Then a has a multiplicative inverse modulo n if and only if $a \perp n$.

Proof
Note that a has a multiplicative inverse u modulo n if and only if there is a solution (u, v) to the equation $au + nv = 1$. Indeed, $au \equiv 1 \bmod n$ if and only if $n \mid au - 1$, which occurs if and only if there is some $q \in \mathbb{Z}$ such that $au - 1 = nq$. Setting $q = -v$ and rearranging yields the desired equivalence.

By Bézout's lemma (Theorem 4.1.22), such a solution (u, v) exists if and only if $\gcd(a, n) \mid 1$. This occurs if and only if $\gcd(a, n) = 1$, i.e. if and only if $a \perp n$. □

Proof tip

To solve a congruence of the form $ax \equiv b \bmod n$ when $a \perp n$, first find a multiplicative inverse u for a modulo n, and then simply multiply through by u to obtain $x \equiv ub \bmod n$. ◁

Corollary 4.3.20

Let $a, p \in \mathbb{Z}$, where p is a positive prime. If $p \nmid a$ then a has a multiplicative inverse modulo p.

Proof

Suppose $p \nmid a$, and let $d = \gcd(a, p)$. Since $d \mid p$ and p is prime we have $d = 1$ or $d = p$. Since $d \mid a$ and $p \nmid a$ we can't have $d = p$; therefore $d = 1$. By Proposition 4.3.19, a has a multiplicative inverse modulo p. □

Example 4.3.21

11 is prime, so each of the integers a with $1 \leqslant a \leqslant 10$ should have a multiplicative inverse modulo 11. And indeed, the following are all congruent to 1 modulo 11:

$$1 \times 1 = 1 \quad 2 \times 6 = 12 \quad 3 \times 4 = 12 \quad 4 \times 3 = 12 \quad 5 \times 9 = 45$$
$$6 \times 2 = 12 \quad 7 \times 8 = 56 \quad 8 \times 7 = 56 \quad 9 \times 5 = 45 \quad 10 \times 10 = 100$$

◁

Exercise 4.3.22

Find all integers x such that $25x - 4 \equiv 4x + 3 \bmod 13$. ◁

Orders and totients

For any modulus n, there are only finitely many possible remainders modulo n. A nice consequence of this finiteness is that, when $a \perp n$, we can choose some power of a to be its multiplicative inverse, as proved in the following exercise.

Exercise 4.3.23

Let n be a modulus and let $a \in \mathbb{Z}$ with $a \perp n$. Prove that there exists $k \geqslant 1$ such that $a^k \equiv 1 \bmod n$. ◁

Exercise 4.3.23, together with the well-ordering principle, justify the following definition.

Definition 4.3.24

Let n be a modulus and let $a \in \mathbb{Z}$ with $a \perp n$. The **order** of a modulo n is the least $k \geqslant 1$ such that $a^k \equiv 1 \bmod n$.

Note that this definition makes sense by Exercise 4.3.23 and the well-ordering principle.

Example 4.3.25

The powers of 7 modulo 100 are:

- $7^1 = 7$, so $7^1 \equiv 7 \bmod 100$;

- $7^2 = 49$, so $7^2 \equiv 49 \bmod 100$;

- $7^3 = 343$, so $7^3 \equiv 43 \bmod 100$;

- $7^4 = 2401$, so $7^4 \equiv 1 \bmod 100$.

Hence the order of 7 modulo 100 is 4, and 7^3 and 43 are multiplicative inverses of 7 modulo 100.   ◁

Our focus turns to computing specific values of k such that $a^k \equiv 1 \bmod n$, whenever $a \in \mathbb{Z}$ and $a \perp n$. We first focus on the case when n is prime; then we develop the machinery of *totients* to study the case when n is not prime.

Lemma 4.3.26

Let $a, b \in \mathbb{Z}$ and let $p \in \mathbb{Z}$ be a positive prime. Then $(a+b)^p \equiv a^p + b^p \bmod p$.

Proof

By the binomial theorem (Theorem 3.1.35), we have

$$(a+b)^p = \sum_{k=0}^{p} \binom{p}{k} a^k b^{p-k}$$

By Exercise 4.2.5, $p \mid \binom{p}{k}$ for all $0 < k < p$, and hence $\binom{p}{k} a^k b^{p-k} \equiv 0 \bmod p$ for all $0 < k < p$. Thus

$$(a+b)^p \equiv \binom{p}{0} a^0 b^{p-0} + \binom{p}{p} a^p b^{p-p} \equiv a^p + b^p \bmod p$$

as desired.   □

Theorem 4.3.27 (Fermat's little theorem)

Let $a, p \in \mathbb{Z}$ with p a positive prime. Then $a^p \equiv a \bmod p$.

Proof

We may assume that $a \geqslant 0$, otherwise replace a by its remainder modulo p.

We will prove that $a^p \equiv a \bmod p$ by induction on a.

- (**BC**) Since $p > 0$ we have $0^p = 0$, hence $0^p \equiv 0 \bmod p$.

- (**IS**) Fix $a \geqslant 0$ and suppose $a^p \equiv a \bmod p$. Then $(a+1)^p \equiv a^p + 1^p \bmod p$ by Lemma 4.3.26. Now $a^p \equiv a \bmod p$ by the induction hypothesis, and $1^p = 1$, so we have $(a+1)^p \equiv a+1 \bmod p$.

By induction, we're done. □

The following consequence of Theorem 4.3.27 is often also referred to as 'Fermat's little theorem', but is slightly less general since it requires an additional hypothesis. In keeping with the wider mathematical community, we will refer to both Theorem 4.3.27 and Corollary 4.3.28 as 'Fermat's little theorem'.

Corollary 4.3.28 (Fermat's little theorem — alternative version)
Let $a, p \in \mathbb{Z}$ with p a positive prime and $p \nmid a$. Then $a^{p-1} \equiv 1 \bmod p$.

Proof
Since $p \nmid a$, it follows that $a \perp p$. Theorem 4.3.27 tells us that $a^p \equiv a \bmod p$. By Proposition 4.3.19, a has a multiplicative inverse b modulo p. Hence

$$a^p b \equiv ab \bmod p$$

But $a^p b \equiv a^{p-1} ab \bmod p$, and $ab \equiv 1 \bmod p$, so we get

$$a^{p-1} \equiv 1 \bmod p$$

as required. □

Corollary 4.3.28 can be useful for computing remainders of humongous numbers when divided by smaller primes.

Example 4.3.29
We compute the remainder of 2^{1000} when divided by 7. Since $7 \nmid 2$, it follows from Fermat's little theorem (Corollary 4.3.28) that $2^6 \equiv 1 \bmod 7$. Now $1000 = 166 \times 6 + 4$, so

$$2^{1000} \equiv 2^{166 \times 6 + 4} \equiv (2^6)^{166} \cdot 2^4 \equiv 2^4 \equiv 16 \equiv 2 \bmod 7$$

so the remainder of 2^{1000} when divided by 7 is 2. ◁

Exercise 4.3.30
Find the remainder of 3^{244886} when divided by 13. ◁

Unfortunately, the hypothesis that p is prime in Fermat's little theorem cannot be disposed of. For example, 6 is not prime, and $5^{6-1} = 5^5 = 3125 = 520 \times 6 + 5$, so $5^5 \equiv 5 \bmod 6$. Our next order of business is to generalise Corollary 4.3.28 by removing the requirement that the modulus p be prime, and replacing $p - 1$ by the *totient* of the modulus.

Definition 4.3.31

Let $n \in \mathbb{Z}$. The **totient** of n is the natural number $\varphi(n)$ (LATEX code: `\varphi(n)`) defined by

$$\varphi(n) = |\{k \in [|n|] \mid k \perp n\}|$$

That is, $\varphi(n)$ is the number of integers from 1 up to $|n|$ which are coprime to n. The function $\varphi : \mathbb{Z} \to \mathbb{N}$ is called **Euler's totient function**.

Example 4.3.32

Here are some examples of totients:

- The elements of $[6]$ which are coprime to 6 are 1 and 5, so $\varphi(6) = 2$.

- If p is a positive prime, then by Corollary 4.2.20, every element of $[p]$ is coprime to p except for p itself. Hence if p is a positive prime then $\varphi(p) = p - 1$. More generally, if p is prime then $\varphi(p) = |p| - 1$.

$\triangleleft$

Exercise 4.3.33

Let $n \in \mathbb{Z}$ and let $p > 0$ be prime. Prove that if $p \mid n$, then $\varphi(pn) = p \cdot \varphi(n)$. Deduce that $\varphi(p^k) = p^k - p^{k-1}$ for all prime $p > 0$ and all $k \geqslant 1$. $\triangleleft$

Exercise 4.3.34

Let p and q be distinct positive primes. Prove that $\varphi(pq) = (p-1)(q-1)$. $\triangleleft$

Together, Exercises 4.3.33 and 4.3.34 allow us to compute the totient of any integer with up to two primes in its prime factorisation.

Example 4.3.35

We compute $\varphi(100)$. The prime factorisation of 100 is $2^2 \times 5^2$. Applying Exercise 4.3.33 twice

$$\varphi(2^2 \times 5^2) = 2 \times 5 \times \varphi(2 \times 5) = 10\varphi(10)$$

Finally, Exercise 4.3.34 tells us that $\varphi(10) = \varphi(2 \times 5) = 1 \times 4 = 4$. Hence $\varphi(100) = 40$. $\triangleleft$

Exercise 4.3.36

Prove that $\varphi(100) = 40$, this time using the inclusion–exclusion principle. $\triangleleft$

Euler's theorem uses totients to generalise Fermat's little theorem (Theorem 4.3.27) to arbitrary moduli, not just prime ones.

Theorem 4.3.37 (Euler's theorem)

Let n be a modulus and let $a \in \mathbb{Z}$ with $a \perp n$. Then

$$a^{\varphi(n)} \equiv 1 \bmod n$$

Proof

By definition of totient, the set X defined by

$$X = \{k \in [n] \mid k \perp n\}$$

has $\varphi(n)$ elements. List the elements as

$$X = \{x_1, x_2, \ldots, x_{\varphi(n)}\}$$

Note that $ax_i \perp n$ for all i, so let y_i be the (unique) element of X such that $ax_i \equiv y_i \bmod n$.

Note that if $i \neq j$ then $y_i \neq y_j$. We prove this by contraposition; indeed, since $a \perp n$, by Proposition 4.3.19, a has a multiplicative inverse, say b. Then

$$y_i \equiv y_j \bmod n \implies ax_i \equiv ax_j \bmod n \implies bax_i \equiv bax_j \bmod n \implies x_i \equiv x_j \bmod n$$

and $x_i \equiv x_j \bmod n$ if and only if $i = j$. Thus

$$X = \{x_1, x_2, \ldots, x_{\varphi(n)}\} = \{y_1, y_2, \ldots, y_{\varphi(n)}\}$$

This means that the product of the 'x_i's is equal to the product of the 'y_i's, and hence

$$
\begin{aligned}
x_1 \cdot \ldots \cdot x_{\varphi(n)} & \\
\equiv y_1 \cdot \ldots \cdot y_{\varphi(n)} \bmod n & \qquad \text{since } \{x_1, \ldots\} = \{y_1, \ldots\} \\
\equiv (ax_1) \cdot \ldots \cdot (ax_{\varphi(n)}) \bmod n & \qquad \text{since } y_i \equiv ax_i \bmod n \\
\equiv a^{\varphi(n)} \cdot x_1 \cdot \ldots \cdot x_{\varphi(n)} \bmod n & \qquad \text{rearranging}
\end{aligned}
$$

Since each x_i is coprime to n, we can cancel the x_i terms (by multiplying by their multiplicative inverses) to obtain

$$a^{\varphi(n)} \equiv 1 \bmod n$$

as required. □

Example 4.3.38

Some examples of Euler's theorem in action are as follows:

- We have seen that $\varphi(6) = 2$, and we know that $5 \perp 6$. And, indeed,

$$5^{\varphi(6)} = 5^2 = 25 = 4 \times 6 + 1$$

so $5^{\varphi(6)} \equiv 1 \bmod 6$.

- By Exercise 4.3.33, we have

$$\varphi(121) = \varphi(11^2) = 11^2 - 11^1 = 121 - 11 = 110$$

Moreover, given $a \in \mathbb{Z}$, $a \perp 121$ if and only if $11 \nmid a$ by Corollary 4.2.20. Hence $a^{110} \equiv 1 \bmod 121$ whenever $11 \nmid a$.

◁

Exercise 4.3.39
Use Euler's theorem to prove that the last two digits of 3^{79} are '67'.   ◁

Example 4.3.40
Let n be a modulus and let $a \in \mathbb{Z}$ with $a \perp n$. Prove that the order of a modulo n divides $\varphi(n)$.   ◁

A formula for the totient of an arbitrary nonzero integer is proved in Theorem 4.3.64—its proof is an application of the Chinese remainder theorem Theorem 4.3.51, and uses the techniques for counting finite sets discussed in Section 3.3.

Wilson's theorem

We conclude this chapter on number theory with *Wilson's theorem*, which is a nice result that completely characterises prime numbers in the sense that we can tell when a number is prime by computing the remainder of $(n-1)!$ when divided by n.

Let's test a few numbers first:

n	$(n-1)!$	remainder
2	1	1
3	2	2
4	6	2
5	24	4
6	120	0
7	720	6
8	5040	0

n	$(n-1)!$	remainder
9	40320	0
10	362880	0
11	3628800	10
12	39916800	0
13	479001600	12
14	6227020800	0
15	87178291200	0

It's tempting to say that an integer $n > 1$ is prime if and only if $n \nmid (n-1)!$, but this isn't true since it fails when $n = 4$. But it's extremely close to being true.

Theorem 4.3.41 (Wilson's theorem)
Let $n > 1$ be a modulus. Then n is prime if and only if $(n-1)! \equiv -1 \bmod n$.

The following sequence of exercises will piece together into a proof of Wilson's theorem.

Exercise 4.3.42
Let $n \in \mathbb{Z}$ be composite. Prove that if $n > 4$, then $n \mid (n-1)!$. ◁

Exercise 4.3.43
Let p be a positive prime and let $a \in \mathbb{Z}$. Prove that, if $a^2 \equiv 1 \bmod p$, then $a \equiv 1 \bmod p$ or $a \equiv -1 \bmod p$. ◁

Exercise 4.3.43 implies that the only elements of $[p-1]$ that are their own multiplicative inverses are 1 and $p-1$; this morsel of information allows us to deduce result in the following exercise.

Exercise 4.3.44
Let p be a positive prime. Prove that $(p-1)! \equiv -1 \bmod p$. ◁

Proof of Wilson's theorem (*Theorem 4.3.41*)
Let $n > 1$ be a modulus.

- If n is prime, then $(n-1)! \equiv -1 \bmod n$ by Exercise 4.3.44.

- If n is composite, then either $n = 4$ or $n > 4$. If $n = 4$ then

$$(n-1)! = 3! = 6 \equiv 2 \bmod 4$$

  and so $(n-1)! \not\equiv -1 \bmod n$. If $n > 4$, then

$$(n-1)! \equiv 0 \bmod n$$

  by Exercise 4.3.42.

Hence $(n-1)! \equiv -1 \bmod n$ if and only if n is prime, as desired.   □

Since Wilson's theorem completely characterises the positive prime numbers, we could have defined 'n is prime', for $n > 1$, to mean that $(n-1)! \equiv -1 \bmod n$. We don't do this because, although this is an interesting result, it is not particularly useful in applications. We might even hope that Wilson's theorem gives us an easy way to test whether a number is prime, but unfortunately even this is a bust: computing the remainder $(n-1)!$ on division by n is not particularly efficient.

However, there are some nice applications of Wilson's theorem, which we will explore now.

Example 4.3.45
We'll compute the remainder of $3^{45} \cdot 44!$ when divided by 47. Note that $3^{45} \cdot 44!$ is equal to a monstrous number with 76 digits; I don't recommend doing the long division! Anyway. . .

- 47 is prime, so we can apply both Fermat's little theorem (Theorem 4.3.27) and Wilson's theorem (Theorem 4.3.41).

- By Fermat's little theorem, we know that $3^{46} \equiv 1 \bmod 47$. Since $3 \cdot 16 = 48 \equiv 1 \bmod 47$, we have

$$3^{45} \equiv 3^{45} \cdot (3 \cdot 16) \equiv 3^{46} \cdot 16 \equiv 16 \bmod 47$$

- By Wilson's theorem, we have $46! \equiv -1 \bmod 47$. Now
  ◇ $46 \equiv -1 \bmod 47$, so 46 is its own multiplicative inverse modulo 47.
  ◇ The extended Euclidean algorithm yields $45 \cdot 23 \equiv 1 \bmod 47$.
  So we have

$$44! = 44! \cdot (45 \cdot 23) \cdot (46 \cdot 46) \equiv 46! \cdot 23 \cdot 46 \equiv (-1) \cdot 23 \cdot (-1) \equiv 23 \bmod 47$$

Putting this information together yields

$$3^{45} \cdot 44! \equiv 16 \cdot 23 = 368 \equiv 39 \bmod 47$$

So the remainder left when $3^{45} \cdot 44!$ is divided by 47 is 39.   ◁

Exercise 4.3.46
Let p be an odd positive prime. Prove that

$$\left[\left(\frac{p-1}{2} \right)! \right]^2 \equiv (-1)^{\frac{p+1}{2}} \bmod p$$

◁

Chinese remainder theorem

We introduce the Chinese remainder theorem with an example.

Example 4.3.47

We find all integer solutions x to the system of congruences

$$x \equiv 2 \bmod 5 \quad \text{and} \quad x \equiv 4 \bmod 8$$

Note that $x \equiv 4 \bmod 8$ if and only if $x = 4 + 8k$ for some $k \in \mathbb{Z}$. Now, for all $k \in \mathbb{Z}$ we have

$\quad x \equiv 2 \bmod 5$

$\quad \Leftrightarrow 4 + 8k \equiv 2 \bmod 5 \quad$ since $x = 4 + 8k$

$\quad \Leftrightarrow 8k \equiv -2 \bmod 5 \qquad$ subtracting 4

$\quad \Leftrightarrow 3k \equiv 3 \bmod 5 \qquad$ since $8 \equiv -2 \equiv 3 \bmod 5$

$\quad \Leftrightarrow k \equiv 1 \bmod 5 \qquad$ multiplying by a multiplicative inverse for 3 modulo 5

So $4 + 8k \equiv 2 \bmod 5$ if and only if $k = 1 + 5\ell$ for some $\ell \in \mathbb{Z}$.

Combining this, we see that x satisfies both congruences if and only if

$$x = 4 + 8(1 + 5\ell) = 12 + 40\ell$$

for some $\ell \in \mathbb{Z}$.

Hence the integers x for which both congruences are satisfied are precisely those integers x such that $x \equiv 12 \bmod 40$. ◁

Exercise 4.3.48

Find all integer solutions x to the system of congruences:

$$\begin{cases} x \equiv -1 \bmod 4 \\ x \equiv 1 \bmod 9 \\ x \equiv 5 \bmod 11 \end{cases}$$

Express your solution in the form $x \equiv a \bmod n$ for suitable $n > 0$ and $0 \leqslant a < n$.
◁

Exercise 4.3.49

Let m, n be coprime moduli and let $a, b \in \mathbb{Z}$. Let $u, v \in \mathbb{Z}$ be such that

$$mu \equiv 1 \bmod n \quad \text{and} \quad nv \equiv 1 \bmod m$$

In terms of a, b, m, n, u, v, find an integer x such that

$$x \equiv a \bmod m \quad \text{and} \quad x \equiv b \bmod n$$

◁

Exercise 4.3.50

Let m, n be coprime moduli and let $x, y \in \mathbb{Z}$. Prove that if $x \equiv y \bmod m$ and $x \equiv y \bmod n$, then $x \equiv y \bmod mn$.   ◁

Theorem 4.3.51 (Chinese remainder theorem)

Let m, n be moduli and let $a, b \in \mathbb{Z}$. If m and n are coprime, then there exists an integer solution x to the simultaneous congruences

$$x \equiv a \bmod m \quad \text{and} \quad x \equiv b \bmod n$$

Moreover, if $x, y \in \mathbb{Z}$ are two such solutions, then $x \equiv y \bmod mn$.

Proof

Existence of a solution x is precisely the content of Exercise 4.3.49.

Now let $x, y \in \mathbb{Z}$ be two solutions to the two congruences. Then

$$\begin{cases} x \equiv a \bmod m \\ y \equiv a \bmod m \end{cases} \Rightarrow \quad x \equiv y \bmod m$$

$$\begin{cases} x \equiv b \bmod n \\ y \equiv b \bmod n \end{cases} \Rightarrow \quad x \equiv y \bmod n$$

so by Exercise 4.3.50, we have $x \equiv y \bmod mn$, as required.   □

We now generalise the Chinese remainder theorem to the case when the moduli m, n are not assumed to be coprime. There are two ways we could make this generalisation: either we could reduce the more general version of the theorem to the version we proved in Theorem 4.3.51, or we could prove the more general version from scratch. We opt for the latter approach, but you might want to consider what a 'reductive' proof would look like.

Theorem 4.3.52

Let m, n be moduli and let $a, b \in \mathbb{Z}$. There exists an integer solution x to the system of congruences

$$x \equiv a \bmod m \quad \text{and} \quad x \equiv b \bmod n$$

if and only if $a \equiv b \bmod \gcd(m, n)$.

Moreover, if $x, y \in \mathbb{Z}$ are two such solutions, then $x \equiv y \bmod \operatorname{lcm}(m, n)$

Proof

Let $d = \gcd(m, n)$, and write $m = m'd$ and $n = n'd$ for some $m', n' \in \mathbb{Z}$.

We prove that an integer solution x to the system of congruences exists if and only if $a \equiv b \bmod d$.

- ($\Rightarrow$) Suppose an integer solution x to the system of congruences exists. Then there exist integers k, ℓ such that

$$x = a + mk = b + n\ell$$

But $m = m'd$ and $n = n'd$, so we have $a + m'dk = b + n'd\ell$, and so

$$a - b = (n'\ell - m'k)d$$

so that $a \equiv b \bmod d$, as required.

- ($\Leftarrow$) Suppose $a \equiv b \bmod d$, and let $t \in \mathbb{Z}$ be such that $a - b = td$. Let $u, v \in \mathbb{Z}$ be solutions to the congruence $mu + nv = d$, which exists by Bézout's lemma (Theorem 4.1.22). Note also that, since $m = m'd$ and $n = n'd$, dividing through by d yields $m'u + n'v = 1$.

Define

$$x = an'v + bm'u$$

Now we have

$$
\begin{aligned}
x &= an'v + bm'u && \text{by definition of } x \\
&= an'v + (a - td)m'u && \text{since } a - b = td \\
&= a(m'u + n'v) - tdm'u && \text{rearranging} \\
&= a - tdm'u && \text{since } m'u + n'v = 1 \\
&= a - tum && \text{since } m = m'd
\end{aligned}
$$

so $x \equiv a \bmod m$. Likewise

$$
\begin{aligned}
x &= an'v + bm'u && \text{by definition of } x \\
&= (b + td)n'v + bm'u && \text{since } a - b = td \\
&= b(m'u + n'v) + tdn'v && \text{rearranging} \\
&= b + tdn'v && \text{since } m'u + n'v = 1 \\
&= b + tvn && \text{since } n = n'd
\end{aligned}
$$

so $x \equiv b \bmod n$.

Hence $x = an'v + bm'u$ is a solution to the system of congruences.

We now prove that if x, y are two integer solutions to the system of congruences, then they are congruent modulo $\mathrm{lcm}(a, b)$. First note that we must have

$$x \equiv y \bmod m \quad \text{and} \quad x \equiv y \bmod n$$

so that $x = y + km$ and $x = y + \ell n$ for some $k, \ell \in \mathbb{Z}$. But then

$$x - y = km = \ell n$$

Writing $m = m'd$ and $n = n'd$, we see that $km'd = \ell n'd$, so that $km' = \ell n'$. But m', n' are coprime by Exercise 4.1.29, and hence $m' \mid \ell$ by Proposition 4.1.31. Write $\ell = \ell'm'$ for some $\ell' \in \mathbb{Z}$. Then we have

$$x - y = \ell n = \ell'm'n$$

and hence $x \equiv y \bmod m'n$. But $m'n = \mathrm{lcm}(m, n)$ by Exercise 4.1.40. $\quad\square$

This theorem is in fact *constructive*, in that it provides an algorithm for finding all integer solutions x to a system of congruences

$$x \equiv a \bmod m \quad \text{and} \quad x \equiv b \bmod n$$

as follows:

- Use the Euclidean algorithm to compute $d = \gcd(m, n)$.

- If $d \nmid a - b$ then there are no solutions, so stop. If $d \mid a - b$, then proceed to the next step.

- Use the extended Euclidean algorithm to compute $u, v \in \mathbb{Z}$ such that $mu + nv = d$.

- The integer solutions x to the system of congruences are precisely those of the form
$$x = \frac{anv + bmu + kmn}{d} \quad \text{for some } k \in \mathbb{Z}$$

Exercise 4.3.53
Verify that the algorithm outlined above is correct. Use it to compute the solutions to the system of congruences

$$x \equiv 3 \bmod 12 \quad \text{and} \quad x \equiv 15 \bmod 20$$

$\triangleleft$

⋆ Exercise 4.3.54
Generalise the Chinese remainder theorem to systems of arbitrarily (finitely) many congruences. That is, given $r \in \mathbb{N}$, find precisely the conditions on moduli

$n_1, n_2, \ldots, n_r$ and integers $a_1, a_2, \ldots, a_r$ such that an integer solution exists to the congruences

$$x \equiv a_1 \bmod n_1, \quad x \equiv a_2 \bmod n_2, \quad \cdots \quad x_r \equiv a_r \bmod n_r$$

Find an explicit formula for such a value of x, and find a suitable modulus n in terms of $n_1, n_2, \ldots, n_r$ such that any two solutions to the system of congruences are congruent modulo n. ◁

Exercise 4.3.55
Prove that gaps between consecutive primes can be made arbitrarily large. That is, prove that for all $n \in \mathbb{N}$, there exists an integer a such that the numbers

$$a, \ a+1, \ a+2, \ \ldots, \ a+n$$

are all composite. ◁

Application: tests for divisibility

The language of modular arithmetic provides a practical setting for proving tests for divisibility using number bases. Number bases were introduced in Chapter 0, and we gave a preliminary definition in Definition 0.6 of what a number base is. Our first job will be to justify why this definition makes sense at all—that is, we need to prove that every natural number *has* a base-b expansion, and moreover, that it only has one of them. Theorem 4.3.56 says exactly this.

Theorem 4.3.56
Let $n \in \mathbb{N}$ and let $b \in \mathbb{N}$ with $b \geqslant 2$. Then there exist unique $r \in \mathbb{N}$ and $d_0, d_1, \ldots, d_r \in \{0, 1, \ldots, b-1\}$ such that

$$n = \sum_{i=0}^{r} d_i b^i$$

and such that $d_r \neq 0$, except $n = 0$, in which case $r = 0$ and $d_0 = 0$.

Proof
We proceed by strong induction on n.

- **(BC)** We imposed the requirement that if $n = 0$ then $r = 0$ and $d_0 = 0$; and this evidently satisfies the requirement that $n = \sum_{i=0}^{r} d_i b^i$.

- **(IS)** Fix $n \geqslant 0$ and suppose that the requirements of the theorem are satisfied for all the natural numbers up to and including n.

By the division theorem ([Theorem 4.1.1](#)), there exist unique $u, v \in \mathbb{N}$ such that

$$n+1 = ub+v \quad \text{and} \quad v \in \{0, 1, \ldots, b-1\}$$

Since $b \geqslant 2$, we have $u < n+1$, and so $u \leqslant n$. It follows from the induction hypothesis that there exist unique $r \in \mathbb{N}$ and $d_1, \ldots, d_r \in \{0, 1, \ldots, b-1\}$ such that

$$u = \sum_{i=0}^{r} d_{i+1} b^i$$

and $d_r \neq 0$. Writing $d_0 = v$ yields

$$n = ub+v = \sum_{i=0}^{r} d_{i+1} b^{i+1} + d_0 = \sum_{i=0}^{r} d_i b^i$$

Since $d_r \neq 0$, this proves existence.

For uniqueness, suppose that there exists $s \in \mathbb{N}$ and $e_0, \ldots, e_s \in \{0, 1, \ldots, b-1\}$ such that

$$n+1 = \sum_{j=0}^{s} e_j b^j$$

and $e_s \neq 0$. Then

$$n+1 = \left(\sum_{j=1}^{s} e_j b^{j-1} \right) b + e_0$$

so by the division theorem we have $e_0 = d_0 = v$. Hence

$$u = \frac{n+1-v}{b} = \sum_{j=1}^{s} e_j b^{j-1} = \sum_{i=1}^{r} d_i b^{j-1}$$

so by the induction hypothesis, it follows that $r = s$ and $d_i = e_i$ for all $1 \leqslant i \leqslant r$. This proves uniqueness.

By induction, we're done. □

We now re-state the definition of base-b expansion, confident in the knowledge that this definition makes sense.

Definition 4.3.57

Let $n \in \mathbb{N}$. The **base-b expansion** of n is the unique string $d_r d_{r-1} \ldots d_0$ such that the conditions in [Theorem 4.3.56](#) are satisfied. The base-2 expansion is also known as the **binary expansion**, and the base-10 expansion is called the **decimal expansion**.

Example 4.3.58

Let $n \in \mathbb{N}$. Then n is divisible by 3 if and only if the sum of the digits in the decimal expansion of n is divisible by 3. Likewise, n is divisible by 9 if and only if the sum of the digits in the decimal expansion n is divisible by 9.

We prove this for divisibility by 3. Let

$$n = d_r d_{r-1} \cdots d_1 d_0$$

be the decimal expansion of n, and let $s = \sum_{i=0}^{r} d_i$ be the sum of the digits of n.

Then we have

$$n \equiv \sum_{i=0}^{r} d_i 10^i \bmod 3 \qquad \text{since } n = \sum_{i} d_i 10^i$$

$$\equiv \sum_{i=0}^{r} d_i 1^i \bmod 3 \qquad \text{since } 10 \equiv 1 \bmod 3$$

$$\equiv \sum_{i=0}^{r} d_i \qquad \text{since } 1^i = 1 \text{ for all } i$$

$$\equiv s \qquad \text{by definition of } s$$

Since $n \equiv s \bmod 3$, it follows that n is divisible by 3 if and only if s is divisible by 3. ◁

Exercise 4.3.59

Let $n \in \mathbb{N}$. Prove that n is divisible by 5 if and only if the final digit in the decimal expansion of n is 5 or 0.

More generally, fix $k \geqslant 1$ and let m be the number whose decimal expansion is given by the last k digits of that of n. Prove that n is divisible by 5^k if and only if m is divisible by 5^k. For example, we have

$$125 \mid 9\,550\,828\,230\,495\,875 \qquad \Leftrightarrow \qquad 125 \mid 875$$

◁

Exercise 4.3.60

Let $n \in \mathbb{N}$. Prove that n is divisible by 11 if and only if the *alternating sum* of the digits of n is divisible by 11. That is, prove that if the decimal expansion of n is $d_r d_{r-2} \cdots d_0$, then

$$11 \mid n \qquad \Leftrightarrow \qquad 11 \mid d_0 - d_1 + d_2 - \cdots + (-1)^r d_r$$

◁

Exercise 4.3.61

Let $n \in \mathbb{N}$. Find a method for testing if n is divisible by 7 based on the decimal expansion of n. ◁

Application: public-key cryptography

Public-key cryptography is a method of encryption and decryption that works according to the following principles:

- Encryption is done using a *public key* that is available to anyone.

- Decryption is done using a *private key* that is only known to the recipient.

- Knowledge of the private key should be extremely difficult to derive from knowledge of the public key.

Specifically, suppose that Alice wants to securely send Bob a message. As the recipient of the message, Bob has a public key and a private key. So:

- Bob sends the *public key* to Alice.

- Alice uses the public key to encrypt the message.

- Alice sends the encrypted message, which is visible (but encrypted) to anyone who intercepts it.

- Bob keeps the private key secret, and uses it upon receipt of the message to decrypt the message.

Notice that, since the public key can only be used to *encrypt* messages, a hacker has no useful information upon intercepting the message or the public key.

RSA encryption is an algorithm which provides one means of doing public-key cryptography using the theory of modular arithmetic. It works as follows.

Step 1. Let p and q be distinct positive prime numbers, and let $n = pq$. Then $\varphi(n) = (p-1)(q-1)$.

Step 2. Choose $e \in \mathbb{Z}$ such that $1 < e < \varphi(n)$ and $e \perp \varphi(n)$. The pair (n,e) is called the **public key**.

Step 3. Choose $d \in \mathbb{Z}$ such that $de \equiv 1 \bmod \varphi(n)$. The pair (n,d) is called the **private key**.

Step 4. To encrypt a message M (which is encoded as an integer), compute $K \in [n]$ such that $K \equiv M^e \bmod n$. Then K is the encrypted message.

Step 5. The original message M can be recovered since $M \equiv K^d \bmod n$.

Computing the private key (n,d) from the knowledge of (n,e) would allow a hacker to decrypt an encrypted message. However, doing so is typically very

difficult when the prime factors of n are large. So if we choose p and q to be very large primes—which we can do without much hassle at all—then it becomes computationally infeasible for a hacker to compute the private key.

Example. Suppose I want to encrypt the message M, which I have encoded as the integer 32. Let $p = 13$ and $q = 17$. Then $n = 221$ and $\varphi(n) = 192$. Let $e = 7$, and note that $7 \perp 192$. Now $7 \times 55 \equiv 1$ mod 192, so we can define $d = 55$.

- The public key is $(221, 7)$, which Bob sends to Alice. Now Alice can encrypt the message:
$$32^7 \equiv 59 \text{ mod } 221$$

  Alice then sends Bob the encrypted message 59.

- The private key is $(221, 55)$, so Bob can decrypt the message:

$$59^{55} \equiv 32 \text{ mod } 221$$

  so Bob has received Alice's message 32.

Exercise 4.3.62
Prove that the RSA algorithm is correct. Specifically, prove:

(a) If $n = pq$, for distinct positive primes p and q, then $\varphi(n) = (p-1)(q-1)$;

(b) Given $1 < e < \varphi(n)$ with $e \perp \varphi(n)$, there exists $d \in \mathbb{Z}$ with $de \equiv 1$ mod $\varphi(n)$.

(c) Given $M, K \in \mathbb{Z}$ with $K \equiv M^e$ mod n, it is indeed the case that $K^d \equiv M$ mod n.

◁

Application: Euler's totient function

We now derive a formula for computing the totient of an arbitrary integer using the tools from Section 3.3—in particular, if you chose to read this section *before* learning about the multiplication principle, you should skip over this material.

Theorem 4.3.63 (Multiplicativity of Euler's totient function)
Let $m, n \in \mathbb{Z}$ and let $\varphi : \mathbb{Z} \to \mathbb{N}$ be Euler's totient function (see Definition 4.3.31). If m and n are coprime, then $\varphi(mn) = \varphi(m)\varphi(n)$.

Proof
Since $\varphi(-k) = \varphi(k)$ for all $k \in \mathbb{Z}$, we may assume that $m \geqslant 0$ and $n \geqslant 0$.

Moreover, if $m = 0$ or $n = 0$, then $\varphi(m)\varphi(n) = 0$ and $\varphi(mn) = 0$, so the result is immediate. Hence we may assume that $m > 0$ and $n > 0$.

Given $k \in \mathbb{Z}$, define
$$C_k = \{a \in [k] \mid a \perp k\}$$
By definition of Euler's totient function, we thus have $|C_k| = \varphi(k)$ for all $k \in \mathbb{Z}$. We will define a bijection
$$f : C_m \times C_n \to C_{mn}$$
using the Chinese remainder theorem (Theorem 4.3.51).

Given $a \in C_m$ and $b \in C_n$, let $f(a,b)$ be the element $x \in [mn]$ such that
$$\begin{cases} x \equiv a \bmod m \\ x \equiv b \bmod n \end{cases}$$

- *f* **is well-defined.** We check the properties of totality, existence and uniqueness.
  - ◇ **Totality.** We have accounted for all the elements of $C_m \times C_n$ in our specification of f.
  - ◇ **Existence.** By the Chinese remainder theorem, there exists $x \in \mathbb{Z}$ such that $x \equiv a \bmod m$ and $x \equiv b \bmod n$. By adding an appropriate integer multiple of mn to x, we may additionally require $x \in [mn]$. It remains to check that $x \perp mn$.
    So let $d = \gcd(x, mn)$. If $d > 1$, then there is a positive prime p such that $p \mid x$ and $p \mid mn$. But then $p \mid m$ or $p \mid n$, meaning that either $p \mid \gcd(x,m)$ or $p \mid \gcd(x,n)$. But $x \equiv a \bmod m$, so $\gcd(x,m) = \gcd(a,m)$; and likewise $\gcd(x,n) = \gcd(b,n)$. So this contradicts the assumption that $a \perp m$ and $b \perp n$. Hence $x \perp mn$ after all.
  - ◇ **Uniqueness.** Suppose $x, y \in C_{mn}$ both satisfy the two congruences in question. By the Chinese remainder theorem, we have $x \equiv y \bmod mn$, and hence $x = y + kmn$ for some $k \in \mathbb{Z}$. Since $x, y \in [mn]$, we have
    $$|k|mn = |kmn| = |x - y| \leqslant mn - 1 < mn$$
    This implies $|k| < 1$, so that $k = 0$ and $x = y$.
  so f is well-defined.

- *f* **is injective.** Let $a, a' \in C_m$ and $b, b' \in C_n$, and suppose that $f(a,b) = f(a',b')$. Then there is an element $x \in C_{mn}$ such that
$$\begin{cases} x \equiv a \bmod m \\ x \equiv a' \bmod m \\ x \equiv b \bmod n \\ x \equiv b' \bmod n \end{cases}$$

Hence $a \equiv a'$ mod m and $b \equiv b'$ mod n. Since $a, a' \in [m]$ and $b, b' \in [n]$, we must have $a = a'$ and $b = b'$.

- *f* **is surjective.** Let $x \in C_{mn}$. Let $a \in [m]$ and $b \in [n]$ be the (unique) elements such that $x \equiv a$ mod m and $x \equiv b$ mod n, respectively. If $a \in C_m$ and $b \in C_n$, then we'll have $f(a, b) = x$ by construction, so it remains to check that $a \perp m$ and $b \perp n$.

  Suppose $d \in \mathbb{Z}$ with $d \mid a$ and $d \mid m$. We prove that $d = 1$. Since $x \equiv a$ mod m, we have $d \mid x$ by Theorem 4.1.17. Since $m \mid mn$, we have $d \mid mn$. By definition of greatest common divisors, it follows that $d \mid \gcd(x, mn)$. But $\gcd(x, mn) = 1$, so that d is a unit, and so $a \perp m$ as required.

  The proof that $b \perp n$ is similar.

It was a lot of work to check that it worked, but we have defined a bijection $f : C_m \times C_n \to C_{mn}$. By the multiplication principle, we have

$$\varphi(m)\varphi(n) = |C_m| \cdot |C_n| = |C_m \times C_n| = |C_{mn}| = \varphi(mn)$$

as required. □

It turns out that Theorem 4.3.63 and Exercise 4.3.33 are precisely the ingredients we need to find a general formula for the totient of a nonzero integer.

Theorem 4.3.64 (Formula for Euler's totient function)
Let n be a nonzero integer. Then

$$\varphi(n) = |n| \cdot \prod_{p \mid n} \left(1 - \frac{1}{p}\right)$$

where the product is indexed over positive primes p dividing n

---

Proof
Since $\varphi(n) = \varphi(-n)$ for all $n \in \mathbb{Z}$, we may assume that $n > 0$. Moreover

$$\varphi(1) = 1 = 1 \cdot \prod_{p \mid 1} \left(1 - \frac{1}{p}\right)$$

Note that the product here is empty, and hence equal to 1, since there are no positive primes p which divide 1. So now suppose $n > 1$.

Using the fundamental theorem of arithmetic (Theorem 4.2.12), we can write

$$n = p_1^{k_1} p_2^{k_2} \cdots p_r^{k_r}$$

for primes $0 < p_1 < p_2 < \cdots < p_r$ and natural numbers $k_1, k_2, \ldots, k_r \geqslant 1$.

By repeated application of Theorem 4.3.63, we have

$$\varphi(n) = \prod_{i=1}^{r} \varphi(p_i^{k_i})$$

By Exercise 4.3.33, we have

$$\varphi(p_i^{k_i}) = p_i^{k_i} - p_i^{k_i-1} = p_i^{k_i}\left(1 - \frac{1}{p_i}\right)$$

Combining these two results, it follows that

$$\varphi(n) = \prod_{i=1}^{r} p_i^{k_i}\left(1 - \frac{1}{p_i}\right) = \left(\prod_{i=1}^{r} p_i^{k_i}\right)\left(\prod_{i=1}^{r}\left(1 - \frac{1}{p_i}\right)\right) = n \cdot \prod_{i=1}^{r}\left(1 - \frac{1}{p_i}\right)$$

which is as required. $\square$

Section 4.Q
Chapter 4 exercises

Under construction!
The end-of-chapter exercise sections are new and in an incomplete state.

Relations

Section 5.1
Relations

When studying a kind of mathematical object, such as numbers, sets or proposi-tions, the most interesting results arise from observing how the objects in question relate to one another, or to other objects. For example:

- The most powerful proof techniques that we derived in Chapter 1 arose from studying *logical equivalence*, which exploited what happens when two logical formulae can be derived from one another.

- Although Section 3.2 was dedicated to studying finite sets, we were not even able to define what it means for a set X to be finite until we were comfortable with the notion of a *bijection* between sets.

- All of Chapter 4 concerned *divisibility* between integers, rather than studying integers in isolation—for example, we cannot define what it means for 5 to be prime without saying how the number 5 interacts with other integers.

The notion of a *relation* is the mathematical abstraction of these ideas.

Definition 5.1.1

Let X and Y be sets. A (**binary**) **relation** from X to Y is a logical formula $R(x,y)$ with two free variables x, y, where x has range X and y has range Y. We call X the **domain** of R and Y the **codomain** of R.

Given $x \in X$ and $y \in Y$, if $R(x,y)$ is true then we say 'x is **related** to y by R', and write $x \mathbin{R} y$ (LaTeX code: `x \mathbin{R} y`).

In more human terms, a relation from X to Y is a statement about a generic element $x \in X$ and a generic element $y \in Y$, which is either true or false depending on the values of x and y.

Example 5.1.2
We have seen many examples of relations so far. For example:

- Every function $f : X \to Y$ defines a relation R_f from X to Y by letting

$$x \mathbin{R_f} y \quad \Leftrightarrow \quad f(x) = y$$

- Given a set X, equality between elements of X ('$x = y$') is a relation from X to X.

- Divisibility ('$x \mid y$') is a relation from $\mathbb{Z}$ to $\mathbb{Z}$.

- For fixed $n \in \mathbb{Z}$, congruence modulo n ('$x \equiv y$ mod n') is a relation from $\mathbb{Z}$ to $\mathbb{Z}$.

- Order ('$x \leqslant y$') is a relation from $\mathbb{N}$ to $\mathbb{N}$, or from $\mathbb{Z}$ to $\mathbb{Z}$, or from $\mathbb{Q}$ to $\mathbb{Q}$, and so on.

- Given sets X and Y, there is an **empty relation** $\varnothing_{X,Y}$ from X to Y, which is defined simply by declaring $\varnothing_{X,Y}(x,y)$ to be false for all $x \in X$ and $y \in Y$.

$\triangleleft$

Exercise 5.1.3

Define a relation R from $\mathbb{Z}$ to $\mathbb{Z}$ which is not on the list given in Example 5.1.2.

$\triangleleft$

It is possible, and extremely useful, to represent relations as sets. We do this by defining the *graph* of a relation, which is the set of all pairs of elements which are related by the relation. You might recognise this as being similar to the graph of a *function* (Definition 2.2.6).

Definition 5.1.4

Let X and Y be sets, and let R be a relation from X to Y. The **graph** of R is the set $\mathrm{Gr}(R)$ (LaTeX code: `\mathrm{Gr}{R}`) of pairs $(x,y) \in X \times Y$ for which $x \, R \, y$. That is

$$\mathrm{Gr}(R) = \{(x,y) \in X \times Y \mid x \, R \, y\} \subseteq X \times Y$$

Example 5.1.5

Consider the relation of divisibility from $\mathbb{Z}$ to $\mathbb{Z}$, that is $R(x,y)$ is the statement $x \mid y$. The graph $\mathrm{Gr}(R)$ of R is the set whose elements are all pairs (m,n) where $m,n \in \mathbb{Z}$ and $m \mid n$. For example, $(2,6) \in \mathrm{Gr}(R)$ since $2 \mid 6$, but $(2,7) \notin \mathrm{Gr}(R)$ since $2 \nmid 7$.

Since $m \mid n$ if and only if $n = qm$ for some $q \in \mathbb{Z}$, we thus have

$$\mathrm{Gr}(R) = \{(m,qm) \mid m,q \in \mathbb{Z}\} \subseteq \mathbb{Z} \times \mathbb{Z}$$

$\triangleleft$

Exercise 5.1.6

Let X and Y be sets. What is the graph of the empty relation from X to Y? $\triangleleft$

Exercise 5.1.7

Let $f : X \to Y$ be a function, and define the relation R_f from X to Y as in Example 5.1.2. Prove that $\mathrm{Gr}(R_f) = \mathrm{Gr}(f)$—that is, the graph of the *relation* R_f is equal to the graph of the *function* f. $\triangleleft$

As with functions, the graph of a relation R from a set X to a set Y can often be represented graphically: draw a pair of axes, with the horizontal axis representing the elements of X and the vertical axis representing the elements of Y, and plot the point (x, y) if and only if $R(x, y)$ is true.

Example 5.1.8

Consider the relation S from $\mathbb{R}$ to $\mathbb{R}$ defined by $x \, S \, y \Leftrightarrow x^2 + y^2 = 1$. Then

$$\text{Gr}(S) = \{(x, y) \in \mathbb{R} \times \mathbb{R} \mid x^2 + y^2 = 1\}$$

Plotting $\text{Gr}(S)$ on a standard pair of axes yields a circle with radius 1 centred at the point $(0, 0)$. Note that $\text{Gr}(S)$ is *not* the graph of a function $s : [0, 1] \to \mathbb{R}$. Indeed, since for example both $0 \, S \, 1$ and $0 \, S \, -1$, the value $s(0)$ would not be uniquely defined.   ◁

Example 5.1.9

Let X be a set. The graph of the equality relation from X to X is very simple:

$$\text{Gr}(=) = \{(x, y) \in X \times X \mid x = y\} = \{(x, x) \mid x \in X\} \subseteq X \times X$$

This set is often denoted Δ_X (LaTeX code: \Delta_{X}), and called the **diagonal subset** of $X \times X$. The reason for the word 'diagonal' is because—provided the horizontal and vertical axes have the same ordering of the elements of X—the points plotted are precisely those on the diagonal line.   ◁

Since we defined relations as particular logical formulae, and we have not defined a notion of equality between logical formulae, if we want to say that two relations are equal then first we need to define what we mean by *equal*. As with sets, this raises some subtleties: should two relations be equal when they're described by the same formula? Or should two relations be equal when they relate the same elements, even if their underlying descriptions are somewhat different? As with equality between sets (Axiom 2.1.22), our notion of equality between relations will be *extensional*: for the purposes of deciding whether two relations are equal, we forget their descriptions and look only at whether or not they relate the same pairs elements.

Axiom 5.1.10 (Relation extensionality)
Let X and Y be sets, and let R and S be relations from X to Y. Then $R = S$ if and only if

$$\forall x \in X, \forall y \in Y, (x \, R \, y \Leftrightarrow x \, S \, y)$$

That is, $R = S$ if they relate exactly the same pairs of elements.

Note that two relations R and S from a set X to a set Y are equal as relations if and only if their graphs $\text{Gr}(R)$ and $\text{Gr}(S)$ are equal as sets. This fact, together

with the correspondence between relations from X to Y and subsets of $X \times Y$ (Theorem 5.1.11 below) is incredibly convenient, because it makes the notion of a relation more concrete.

Theorem 5.1.11

Let X and Y be sets. Any subset $G \subseteq X \times Y$ is the graph of exactly one relation R from X to Y.

Proof

Fix $G \subseteq X \times Y$. Define a relation R by

$$\forall x \in X, \forall y \in Y, x \, R \, y \Leftrightarrow (x, y) \in G$$

Then certainly $G = \mathrm{Gr}(R)$.

Moreover, if S is a relation from X to Y such that $G = \mathrm{Gr}(S)$, then, for all $x \in X$ and $y \in Y$

$$x \, S \, y \Leftrightarrow (x, y) \in \mathrm{Gr}(S) \Leftrightarrow (x, y) \in \mathrm{Gr}(R) \Leftrightarrow x \, R \, y$$

so $S = R$. Hence there is exactly one relation from X to Y whose graph is G. $\square$

Theorem 5.1.11 allows us to use the counting principles from Section 3.3 to find the number of relations from one finite set to another.

Exercise 5.1.12

Let X and Y be finite sets with $|X| = m$ and $|Y| = n$. Prove that there are 2^{mn} relations from X to Y. ◁

Aside

It is very common to identify a relation with its graph, saying that a relation from a set X to a set Y 'is' a subset of $X \times Y$. This practice is justified by Theorem 5.1.11, which says precisely that there is a correspondence between relations from X to Y and subsets of $X \times Y$. ◁

Relations on a set

In most of the examples of relations we've seen so far, the domain of the relation is equal to its codomain. The remainder of this section—in fact, the remainder of this *chapter*—is dedicated to such relations. So let's simplify the terminology slightly.

Definition 5.1.13

Let X be a set. A **relation on X** is a relation from X to X.

We have seen many such relations so far, such as: equality on any set, congru-ence modulo n on $\mathbb{Z}$, divisibility on $\mathbb{Z}$, inclusion of subsets ($\subseteq$) on $\mathscr{P}(X)$, and comparison of size ($\leqslant$) on $\mathbb{N}$, $\mathbb{Z}$, $\mathbb{Q}$ or $\mathbb{R}$. Remarkably, each of these relations can be characterised in one of two ways: either as an *equivalence relation* or as a *partial order*.

Equivalence relations are those that behave in some sense like equality, and partial orders are those that behave in some way like $\leqslant$.

- **Equality.** If X is any set, then equality on X satisfies:
  - ◇ Given $x \in X$, we have $x = x$;
  - ◇ Given $x, y \in X$, if $x = y$, then $y = x$;
  - ◇ Given $x, y, z \in X$, if $x = y$ and $y = z$, then $x = z$.

  Note that these are all true if we replace X by $\mathbb{Z}$ and $\cdot = \cdot$ by $\cdot \equiv \cdot \bmod n$ for some fixed $n > 0$.

- **Order.** If $X = \mathbb{N}$ (or $\mathbb{Z}$ or $\mathbb{Q}$ or $\mathbb{R}$), then the order relation $\leqslant$ on X satisfies:
  - ◇ Given $x \in X$, we have $x \leqslant x$;
  - ◇ Given $x, y \in X$, if $x \leqslant y$ and $y \leqslant x$, then $x = y$;
  - ◇ Given $x, y, z \in X$, if $x \leqslant y$ and $y \leqslant z$, then $x \leqslant z$.

  Note that these are all true if we replace $(X, \leqslant)$ by $(\mathscr{P}(X), \subseteq)$ or $(\mathbb{N}, \mid)$.

For both equality and order, the first condition states that every element is related to itself, and the third condition states that in some sense we can cut out interme-diate steps. These conditions are known as *reflexivity* and *transitivity*. The second condition for equality states that the direction of the relation doesn't matter; this condition is called *symmetry*. The second condition for the order relation states that the only way two objects can be related to each other in both directions is if they are equal; this condition is called *antisymmetry*.

The remainder of this section will develop the language needed to talk about equi-valence relations and partial orders. We will finish the section with a discussion of equivalence relations, and then study partial orders in depth in Section 5.2.

Reflexive relations are those that relate everything to itself.

Definition 5.1.14
Let X be a set. A relation R on X is **reflexive** if $x \mathrel{R} x$ for all $x \in X$.

Example 5.1.15
Given a set X, the equality relation on X is reflexive since $x = x$ for all $x \in X$.   ◁

Example 5.1.16

The divisibility relation on $\mathbb{N}$, or on $\mathbb{Z}$, is reflexive. Given $n \in \mathbb{Z}$ we have $n = 1 \times n$, and so $n \mid n$. ◁

The following exercise demonstrates the importance of specifying the (co)domain of a relation: it shows that a logical formula may define a reflexive relation on one set, but not on another.

Exercise 5.1.17

Prove that coprimality ('$x \perp y$') is not a reflexive relation on $\mathbb{Z}$, but that it is a reflexive relation on the set $\{-1, 1\}$.

As such, it doesn't make sense to say 'coprimality is a reflexive relation' or 'coprimality is not a reflexive relation': we must specify on which set we are considering the coprimality relation. ◁

The result of the next exercise characterises reflexive relations in terms of their graph.

Exercise 5.1.18

Let X be a set and let R be a relation on X. Prove that R is reflexive if and only if $\Delta_X \subseteq \mathrm{Gr}(R)$, where Δ_X is the diagonal subset of $X \times X$ (see Example 5.1.9). Deduce that if X is finite and $|X| = n$, then there are $2^{n(n-1)}$ reflexive relations on X. ◁

Symmetric relations are those for which the *direction* of the relation doesn't matter.

Definition 5.1.19

Let X be a set. A relation R on X is **symmetric** if, for all $x, y \in X$, $x R y$ implies $y R x$.

Example 5.1.20

Some examples of symmetric relations include:

- Equality is a symmetric relation on any set X. Indeed, if $x, y \in X$ and $x = y$, then $y = x$.

- Coprimality is a symmetric relation on $\mathbb{Z}$, since if $a, b \in \mathbb{Z}$ then $a \perp b$ if and only if $b \perp a$.

- Divisibility is not a symmetric relation on $\mathbb{Z}$, since for instance $1 \mid 2$ but $2 \nmid 1$. However, divisibility *is* a symmetric relation on $\{-1, 1\}$, since $1 \mid -1$ and $-1 \mid 1$.

◁

Exercise 5.1.21

Let X be a finite set with $|X| = n$. Prove that there are $2^{\binom{n}{2}} \cdot 2^n$ symmetric relations on X.

◁

A related condition a relation may possess is *antisymmetry*.

Definition 5.1.22

Let X be a set. A relation R on X is **antisymmetric** if, for all $x, y \in X$, if $x \, R \, y$ and $y \, R \, x$, then $x = y$.

A word of warning here is that 'antisymmetric' does not mean the same thing as 'not symmetric'—indeed, we we will see, equality is both symmetric and anti-symmetric, and many relations are neither symmetric nor antisymmetric.[a]

Example 5.1.23

Some examples of antisymmetric relations include are as follows.

- Let X be a set. The equality relation on X is antisymmetric: it is immediate that if $x, y \in X$ and $x = y$ and $y = x$, then $x = y$.

- The relation $\leqslant$ on the set $\mathbb{N}$ (or $\mathbb{Z}$ or $\mathbb{Q}$ or $\mathbb{R}$) is antisymmetric: if $m, n \in \mathbb{N}$ and $m \leqslant n$ and $n \leqslant m$, then $m = n$.

- The divisibility relation on $\mathbb{N}$ is antisymmetric. Indeed, let $m, n \in \mathbb{N}$ and suppose $m \mid n$ and $n \mid m$. Then $n = km$ for some $k \in \mathbb{Z}$ and $m = \ell n$ for some $\ell \in \mathbb{Z}$. It follows that $n = k\ell n$. If $n = 0$ then $m = n$ trivially; otherwise, we have $k\ell = 1$. Hence k is a unit; moreover, since $m, n \geqslant 0$ and $n = km$, we must have $k = 1$. Hence $m = n$.

◁

Exercise 5.1.24

Show that the divisibility relation on $\mathbb{Z}$ is not antisymmetric.   ◁

Exercise 5.1.25

Let X be a set and let R be a relation on X. Prove that R is both symmetric and antisymmetric if and only if $\mathrm{Gr}(R) \subseteq \Delta_X$, where Δ_X is the diagonal subset of $X \times X$ (see Example 5.1.9). Deduce that the only reflexive, symmetric and antisymmetric relation on a set X is the equality relation on X.   ◁

---

[a]Even more confusingly, there is a notion of *asymmetric relation*, which also does not mean 'not symmetric'.

Exercise 5.1.26

Let X be a finite set with $|X| = n$. Prove that there are $3^{\binom{n}{2}} \cdot 2^n$ antisymmetric relations on X. ◁

Transitivity is the property of $\leqslant$ that allows us to deduce, for example, that $0 \leqslant 4$, from the information that $0 \leqslant 1 \leqslant 2 \leqslant 3 \leqslant 4$.

Definition 5.1.27
Let X be a set. A relation R on X is **transitive** if, for all $x, y, z \in X$, if $x \mathrel{R} y$ and $y \mathrel{R} z$, then $x \mathrel{R} z$.

Example 5.1.28

Some examples of transitive relations include:

- Equality is a transitive relation on any set X, since it is immediate that if $x, y, z \in X$ with $x = y$ and $y = z$, then $x = z$.

- Divisibility is a transitive relation on $\mathbb{N}$, or on $\mathbb{Z}$. Indeed, if $a, b, c \in \mathbb{N}$ with $a \mid b$ and $b \mid c$, then there exist $k, \ell \in \mathbb{Z}$ such that $b = ka$ and $c = \ell b$. Then $c = (k\ell)a$, so $a \mid c$.

- Inclusion is a transitive relation on $\mathscr{P}(X)$, for any set X. Indeed, Proposition 2.1.20 implies that if $U, V, W \subseteq X$ with $U \subseteq V$ and $V \subseteq W$, then $U \subseteq W$.

◁

A fundamental property of transitive relations is that we can prove two elements a and b are related by finding a chain of related elements starting at a and finishing at b. This is the content of the following proposition.

Proposition 5.1.29

Let R be a relation on a set X. Then R is transitive if and only if, for any finite sequence $x_0, x_1, \ldots, x_n$ of elements of X such that $x_{i-1} \mathrel{R} x_i$ for all $i \in [n]$, we have $x_0 \mathrel{R} x_n$.

Proof
For the sake of abbreviation, let $p(n)$ be the assertion that, for any $n \geqslant 1$ and any sequence $x_0, x_1, \ldots, x_n$ of elements of X such that $x_{i-1} \mathrel{R} x_i$ for all $i \in [n]$, we have $x_0 \mathrel{R} x_n$.

We prove the two directions of the proposition separately.

- ($\Rightarrow$) Suppose R is transitive. For $n \geqslant 1$. We prove $p(n)$ is true for all $n \geqslant 1$ by induction.

    ◇ (**BC**) When $n = 1$ this is immediate, since we assume that $x_0 \mathrel{R} x_1$.

⋄ **(IS)** Fix $n \geqslant 1$ and suppose $p(n)$ is true. Let $x_0, \ldots, x_n, x_{n+1}$ is a sequence such that $x_{i-1} \mathrel{R} x_i$ for all $i \in [n+1]$. We need to prove that $x_0 \mathrel{R} x_{n+1}$.

By the induction hypothesis we know that $x_0 \mathrel{R} x_n$. By definition of the sequence we have $x_n \mathrel{R} x_{n+1}$. By transitivity, we have $x_0 \mathrel{R} x_{n+1}$.

So by induction, we have proved the $\Rightarrow$ direction.

- ($\Leftarrow$) Suppose $p(n)$ is true for all $n \geqslant 1$. Then in particular $p(2)$ is true, which is precisely the assertion that R is transitive.

So we're done.   □

That is, Proposition 5.1.29 states that for a transitive relation R on a set X, if $x_0, x_1, \ldots, x_n \in X$, then

$$x_0 \mathrel{R} x_1 \mathrel{R} \cdots \mathrel{R} x_n \quad \Rightarrow \quad x_0 \mathrel{R} x_n$$

where '$x_0 \mathrel{R} x_1 \mathrel{R} \cdots \mathrel{R} x_n$' abbreviates the assertion that $x_i \mathrel{R} x_{i+1}$ for each $i < n$.

Equivalence relations

We will now study what it is for a relation to be *equality-like*.

> **Definition 5.1.30**
> A relation R on a set X is an **equivalence relation** if R is reflexive, symmetric and transitive.

When we talk about arbitrary equivalence relations, we usually use a symbol like '$\sim$' (LATEX code: `\sim`) or '$\equiv$' (LATEX code: `\equiv`) or '$\approx$' (LATEX code: `\approx`) instead of 'R'.

Example 5.1.31
Recall Theorem 4.3.6. With our new language of relations, we could succinctly re-state it as follows:

Let n be a modulus. Congruence modulo n is an equivalence relation on $\mathbb{Z}$.

Indeed, part (a) of Theorem 4.3.6 proved reflexivity, part (b) proved symmetry, and part (c) proved transitivity.   ◁

Exercise 5.1.32
Prove that equality of sets (Axiom 2.1.22) is an equivalence relation on the universe of discourse $\mathscr{U}$.   ◁

Exercise 5.1.33

Given a function $f : X \to Y$, define a relation $\sim_f$ on X by

$$a \sim_f b \quad \Leftrightarrow \quad f(a) = f(b)$$

for all $a, b \in X$. Prove that $\sim_f$ is an equivalence relation on X. ◁

In the following exercise, we construct a particular equivalence relation $\sim_R$ out of an arbitrary relation R and prove that $\sim_R$ is, in a suitable sense, the 'smallest' equivalence relation extending R.

Exercise 5.1.34

Let R be any relation on a set X. Define a new relation $\sim_R$ on X as follows. Given $x, y \in X$, say $x \sim_R y$ if and only if for some $k \in \mathbb{N}$ there is a sequence $(a_0, a_1, \ldots, a_k)$ of elements of X such that $a_0 = x$, $a_k = y$ and, for all $0 \leqslant i < k$, either $a_i R a_{i+1}$ or $a_{i+1} R a_i$.

First we'll work out a couple of examples.

(a) Fix a modulus n and let R be the relation on $\mathbb{Z}$ defined by xRy if and only if $y = x + n$. Prove that $\sim_R$ is the relation of congruence modulo n.

(b) Let X be a set and let R be the subset relation on $\mathscr{P}(X)$. Prove that $U \sim_R V$ for all $U, V \subseteq X$.

(c) Let X be a set, fix two distinct elements $a, b \in X$, and define a relation R on X by declaring aRb only—that is, for all $x, y \in X$, we have xRy if and only if $x = a$ and $y = b$. Prove that the relation $\sim_R$ is defined by $x \sim_R y$ if and only if either $x = y$ or $\{x, y\} = \{a, b\}$. (Intuitively, $\sim_R$ 'glues' the elements a and b together.)

Next we prove the fundamental facts about $\sim_R$ that we mentioned before the statement of this exercise.

(d) Prove that $\sim_R$ is an equivalence relation on X

(e) Prove that $xRy \Rightarrow x \sim_R y$ for all $x, y \in X$.

(f) Prove that, furthermore, if $\approx$ is any equivalence relation on X and $xRy \Rightarrow x \approx y$ for all $x, y \in X$, then $x \sim_R y \Rightarrow x \approx y$ for all $x, y \in X$.

(g) Use parts (e) and (f) to prove that if R is already an equivalence relation, then the relation $\sim_R$ is equal to R.

We say that the relation $\sim_R$ is the equivalence relation on X **generated by** R. ◁

Equivalence relations are useful because they allow us to ignore irrelevant information about elements of a set. As an example, suppose we want to prove that, for $a \in \mathbb{Z}$, if $3 \nmid a$ then a^2 leaves a remainder of 1 when divided by 3. Before we learnt about modular arithmetic in Section 4.3, in order to prove this, we would have written $a = 3k \pm 1$ for some $k \in \mathbb{Z}$ and done some tedious algebra to deduce that $a^2 = 3(3k^2 \pm 2k) + 1$. This required us to use more information than we need: the value of k doesn't make any difference to the truth of the result, the expression $3(3k^2 \pm 2k) + 1$ is ugly and, more importantly, keeping track of k made the proof longer and more difficult than it has to be. When we learnt modular arithmetic, everything was simplified: if $3 \nmid a$ then $a \equiv \pm 1 \bmod 3$, so that $a^2 \equiv (\pm 1)^2 \equiv 1 \bmod 3$. This proof was shorter and simpler because we didn't need to keep track of exactly which integer a was—all we cared about was its value modulo 3. We could just as well have replaced a with any other integer which leaves the same remainder modulo 3.

This motivates the following definition, which provides a means of identifying two elements of a set that are related by an equivalence relation.

Definition 5.1.35

Let X be a set and let $\sim$ be an equivalence relation on X. The $\sim$-**equivalence class** of $x \in X$ is the set $[x]_\sim$ (LaTeX code: `[x]_{\sim}`) defined by

$$[x]_\sim = \{y \in X \mid x \sim y\}$$

The **quotient** of X by $\sim$ is the set $X/\sim$ (LaTeX code: `X/{\sim}`) of all $\sim$-equivalence classes of elements of X; that is

$$X/\sim = \{[x]_\sim \mid x \in X\}$$

LaTeX tip

Putting braces (`{` and `}`) around a symbol like $\sim$ tells LaTeX to consider the symbol *as a symbol*, rather than as a connective. Compare:

	LaTeX code:	output:
Without braces:	`X/\sim = Y`	$X/\sim = Y$
With braces:	`X/{\sim} = Y`	$X/\sim = Y$

This is because, without braces, LaTeX thinks you're saying 'X-forward-slash is related to is equal to Y', which clearly makes no sense; putting braces around `\sim` signifies to LaTeX that the $\sim$ symbol is being considered as an object in its own right, rather than as a connective.   ◁

Example 5.1.36

Let $f : X \to Y$ be a function, and let $\sim_f$ be the equivalence relation on X that we defined in Exercise 5.1.33. Given $a \in X$, we have

$$[a]_{\sim_f} = \{x \in X \mid a \sim_f x\} = \{x \in X \mid f(a) = f(x)\}$$

Thus we have $[a]_{\sim_f} = f^{-1}[\{f(a)\}]$. ◁

Exercise 5.1.37

Let $f : X \to Y$ be a function. Prove that f is injective if and only if $|[a]_{\sim_f}| = 1$ for all $a \in X$. ◁

Example 5.1.38

Let $\sim$ be the relation of congruence modulo 5 on the set of integers. Then

$$[0]_{\sim} = \{a \in \mathbb{Z} \mid a \sim 0\}$$

Now, $a \sim 0$ if and only if $5 \mid a$, so we can also write

$$[0]_{\sim} = \{\ldots, -10, -5, 0, 5, 10, \ldots\} = \{5k \mid k \in \mathbb{Z}\}$$

So in fact $[0]_{\sim} = [5k]_{\sim}$ for any $k \in \mathbb{Z}$. And likewise

$$[r]_{\sim} = [r + 5k]_{\sim}$$

for all $r, k \in \mathbb{Z}$. It follows that $\mathbb{Z}/\sim = \{[0]_{\sim}, [1]_{\sim}, [2]_{\sim}, [3]_{\sim}, [4]_{\sim}\}$. ◁

Definition 5.1.39

Consider the relation of congruence modulo n on the set $\mathbb{Z}$ of integers. We call the equivalence class of $a \in \mathbb{Z}$ the **congruence class** of a modulo n, denoted $[a]_n$, and we write $\mathbb{Z}/n\mathbb{Z}$ to denote the quotient of $\mathbb{Z}$ by the relation of congruence modulo n.

Example 5.1.40

The set $\mathbb{Z}/5\mathbb{Z}$ has five elements:

$$\mathbb{Z}/5\mathbb{Z} = \{[0]_5, [1]_5, [2]_5, [3]_5, [4]_5\}$$

Example 5.1.38 demonstrates that for all $n \in \mathbb{Z}$ and all $0 \leqslant r < 5$, we have $[n]_5 = [r]_5$ if and only if n leaves a remainder of r when divided by 5. For example, $[7]_5 = [2]_5$. ◁

Exercise 5.1.41

Let n be a modulus. Prove that $\mathbb{Z}/n\mathbb{Z}$ is finite and $|\mathbb{Z}/n\mathbb{Z}| = n$. ◁

Exercise 5.1.41 doesn't tell us much more than we already know: namely, that there are only finitely many possible remainders modulo n. But it makes our lives significantly easier for doing modular arithmetic, because now there are only finitely many objects to work with.

One last word on equivalence relations is that they are essentially the same thing as partitions (see Definition 3.3.25).

Exercise 5.1.42
If $\sim$ be an equivalence relation on X, then $X/\!\!\sim$ is a partition X. Deduce that, for $x, y \in X$, we have $x \sim y$ if and only if $[x]_\sim = [y]_\sim$.   ◁

In fact, the converse of Exercise 5.1.42 is also true, as we prove next.

Proposition 5.1.43
Let X be a set and let $\mathscr{U}$ be a partition of X. Then $\mathscr{U} = X/\!\!\sim$ for exactly one equivalence relation $\sim$ on X.

Proof
Define a relation $\sim$ by

$$x \sim y \quad \Leftrightarrow \quad \exists U \in \mathscr{U}, x \in U \text{ and } y \in U$$

for all $x, y \in X$. That is, $x \sim y$ if and only if x and y are elements of the same set of the partition. We check that $\sim$ is an equivalence relation.

- **Reflexivity.** Let $x \in X$. Then $x \in U$ for some $U \in \mathscr{U}$ since $\bigcup_{U \in \mathscr{U}} U = X$. Hence $x \sim x$.

- **Symmetry.** Let $x, y \in X$ and suppose $x \sim y$. Then there is some $U \in \mathscr{U}$ with $x \in U$ and $y \in U$. But then it is immediate that $y \sim x$.

- **Transitivity.** Let $x, y, z \in X$ and suppose that $x \sim y$ and $y \sim z$. Then there exist $U, V \in \mathscr{U}$ with $x, y \in U$ and $y, z \in V$. Thus $y \in U \cap V$. Since $\mathscr{U}$ is a partition of X, its elements are pairwise disjoint; thus if $U \neq V$ then $U \cap V = \varnothing$. Hence $U = V$. Thus $x \in U$ and $z \in U$, so $x \sim z$.

The definition of $\sim$ makes it immediate that $X/\!\!\sim = \mathscr{U}$.

To prove that $\sim$ is the only such relation, suppose $\approx$ is another equivalence relation on X for which $X/\!\!\approx = \mathscr{U}$. Then, given $x, y \in X$, we have:

$$
\begin{aligned}
x \sim y &\Leftrightarrow \exists U \in \mathscr{U}, x \in U \wedge y \in U &&\text{by definition of } \sim \\
&\Leftrightarrow \exists z \in X, x \in [z]_\approx \wedge y \in [z]_\approx &&\text{since } \mathscr{U} = X/\!\!\approx \\
&\Leftrightarrow \exists z \in X, x \approx z \wedge y \approx z &&\text{by definition of } [z]_\approx \\
&\Leftrightarrow x \approx y &&\text{by symmetry and transitivity}
\end{aligned}
$$

So $\sim \, = \, \approx$. $\qquad \square$

Orders and lattices

We saw in Section 5.1 how equivalence relations behave like '=', in the sense that they are reflexive, symmetric and transitive.

This section explores a new kind of relation which behaves like '⩽'. This kind of relation proves to be extremely useful for making sense of mathematical structures, and has powerful applications throughout mathematics, computer science and even linguistics.

Definition 5.2.1

A relation R on a set X is a **partial order** if R is reflexive, antisymmetric and transitive. That is, if:

- (Reflexivity) $x \, R \, x$ for all $x \in X$;

- (Antisymmetry) For all $x, y \in X$, if $x \, R \, y$ and $y \, R \, x$, then $x = y$;

- (Transitivity) For all $x, y, z \in X$, if $x \, R \, y$ and $y \, R \, z$, then $x \, R \, z$.

A set X together with a partial order R on X is called a **partially ordered set**, or **poset** for short, and is denoted (X, R).

When we talk about partial orders, we usually use a suggestive symbol like '⪯' (LaTeX code: `\preceq`) or '⊑' (LaTeX code: `\sqsubseteq`).

Example 5.2.2
We have seen many examples of posets so far:

- Any of the sets $\mathbb{N}, \mathbb{Z}, \mathbb{Q}$ or $\mathbb{R}$, with the usual order relation $\leqslant$.

- Given a set X, its power set $\mathscr{P}(X)$ is partially ordered by $\subseteq$. Indeed:
    - ◇ **Reflexivity.** If $U \in \mathscr{P}(X)$ then $U \subseteq U$.
    - ◇ **Antisymmetry.** If $U, V \in \mathscr{P}(X)$ with $U \subseteq V$ and $V \subseteq U$, then $U = V$ by definition of set equality.
    - ◇ **Transitivity.** If $U, V, W \in \mathscr{P}(X)$ with $U \subseteq V$ and $V \subseteq W$, then $U \subseteq W$ by Proposition 2.1.20.

- The set $\mathbb{N}$ of natural numbers is partially ordered by the divisibility relation—see Examples 5.1.16, 5.1.23 and 5.1.28. However, by Exercise 5.1.24, the set $\mathbb{Z}$ of integers is not partially ordered by divisibility, since divisibility is not antisymmetric on $\mathbb{Z}$.

- Any set X is partially ordered by its equality relation. This is called the **discrete order** on X.

◁

Much like the difference between the relations $\leqslant$ and $<$ on $\mathbb{N}$, or between $\subseteq$ and $\subsetneq$ on $\mathscr{P}(X)$, every partial order can be *strictified*, in a precise sense outlined in the following definition and proposition.

Definition 5.2.3

A relation R on a set X is a **strict partial order** if it is irreflexive, asymmetric and transitive. That is, if:

- (Irreflexivity) $\neg(x \, R \, x)$ for all $x \in X$;

- (Asymmetry) For all $x, y \in X$, if $x \, R \, y$, then $\neg(y \, R \, x)$;

- (Transitivity) For all $x, y, z \in X$, if $x \, R \, y$ and $y \, R \, z$, then $x \, R \, z$.

Proposition 5.2.4

Let X be a set. Partial orders $\preccurlyeq$ on X are in natural correspondence with strict partial orders $\prec$ on X, according to the rule:

$$x \preccurlyeq y \Leftrightarrow (x \prec y \vee x = y) \quad \text{and} \quad x \prec y \Leftrightarrow (x \preccurlyeq y \wedge x \neq y)$$

Proof

Let P be the set of all partial orders on X and let S be the set of all strict partial orders on X. Define functions

$$f : P \to S \quad \text{and} \quad g : S \to P$$

as in the statement of the proposition, namely:

- Given a partial order $\preccurlyeq$, let $f(\preccurlyeq)$ be the relation $\prec$ defined for $x, y \in X$ by letting $x \prec y$ be true if and only if $x \preccurlyeq y$ and $x \neq y$;

- Given a strict partial order $\prec$, let $g(\prec)$ be the relation $\preccurlyeq$ defined for $x, y \in X$ by letting $x \preccurlyeq y$ be true if and only if $x \prec y$ or $x = y$.

We'll prove that f and g are mutually inverse functions. Indeed:

- f is well-defined. To see this, fix $\preccurlyeq$ and $\prec = f(\preccurlyeq)$ and note that:
  - ◇ $\prec$ is irreflexive, since for $x \in X$ if $x \prec x$ then $x \neq x$, which is a contradiction.
  - ◇ $\prec$ is asymmetric. To see this, let $x, y \in X$ and suppose $x \prec y$. Then $x \preccurlyeq y$ and $x \neq y$. If also $y \prec x$, then we'd have $y \preccurlyeq x$, so that $x = y$ by antisymmetry of $\preccurlyeq$. But $x \neq y$, so this is a contradiction.

◇ ≺ is transitive. To see this, let $x, y, z \in X$ and suppose $x \prec y$ and $y \prec z$. Then $x \preccurlyeq y$ and $y \preccurlyeq z$, so that $x \preccurlyeq z$. Moreover, if $x = z$ then we'd also have $z \preccurlyeq x$ by reflexivity of $\preccurlyeq$, so $z \preccurlyeq y$ by transitivity of $\preccurlyeq$, and hence $y = z$ by antisymmetry of $\preccurlyeq$. But this contradicts $y \prec z$.

So ≺ is a strict partial order on X.

- g is well-defined. To see this, fix $\prec$ and $\preccurlyeq = g(\prec)$ and note that:

  ◇ ≼ is reflexive. This is built into the definition of $\preccurlyeq$.

  ◇ ≼ is symmetric. To see this, fix $x, y \in X$ and suppose $x \preccurlyeq y$ and $y \preccurlyeq x$. Now if $x \ne y$ then $x \prec y$ and $y \prec x$, but this contradicts asymmetry of $\prec$. Hence $x = y$.

  ◇ ≼ is transitive. To see this, fix $x, y, z \in X$ and suppose $x \preccurlyeq y$ and $y \preccurlyeq z$. Then one of the following four cases must be true:

  * $x = y = z$. In this case, $x = z$, so $x \preccurlyeq z$.
  * $x = y \prec z$. In this case, $x \prec z$, so $x \preccurlyeq z$.
  * $x \prec y = z$. In this case, $x \prec z$, so $x \preccurlyeq z$.
  * $x \prec y \prec z$. In this case, $x \prec z$ by transitivity of $\prec$, so $x \preccurlyeq z$.

  In any case, we have that $x \preccurlyeq z$.

  So ≼ is a partial order on X.

- $g \circ f = \mathrm{id}_P$. To see this, let $\prec = f(\preccurlyeq)$ and $\sqsubseteq = g(\prec)$. For $x, y \in X$, we have $x \sqsubseteq y$ if and only if $x \prec y$ or $x = y$, which in turn occurs if and only if $x = y$ or both $x \preccurlyeq y$ and $x \ne y$. This is equivalent to $x \preccurlyeq y$, since if $x = y$ then $x \preccurlyeq y$ by reflexivity. Hence $\sqsubseteq$ and $\preccurlyeq$ are equal relations, so $g \circ f = \mathrm{id}_P$.

- $f \circ g = \mathrm{id}_S$. To see this, let $\preccurlyeq = g(\prec)$ and $\sqsubset = f(\preccurlyeq)$. For $x, y \in X$, we have $x \sqsubset y$ if and only if $x \preccurlyeq y$ and $x \ne y$, which in turn occurs if and only if $x \ne y$ and either $x \prec y$ or $x = y$. Since $x \ne y$ precludes $x = y$, this is equivalent to $x \prec y$. Hence $\prec$ and $\sqsubset$ are equal relations, so $f \circ g = \mathrm{id}_S$.

So f and g are mutually inverse functions, and we have established the required bijection.   □

In light of Proposition 5.2.4, we will freely translate between partial orders and strict partial orders wherever necessary. When we do so, we will use $\prec$ (LaTeX code: `\prec`) to denote the 'strict' version, and $\preccurlyeq$ to denote the 'weak' version. (Likewise for $\sqsubset$ (LaTeX code: `\sqsubset`).)

<div>

Definition 5.2.5

Let $(X, \preccurlyeq)$ be a poset. A $\preccurlyeq$-**least element** of X (or a **least element of X with respect to** $\preccurlyeq$) is an element $\bot \in X$ (LATEX code: `\bot`) such that $\bot \preccurlyeq x$ for all $x \in X$. A $\preccurlyeq$-**greatest element** of X (or a **greatest element of X with respect to** $\preccurlyeq$) is an element $\top \in X$ (LATEX code: `\top`) such that $x \preccurlyeq \top$ for all $x \in X$.

</div>

Example 5.2.6

Some examples of least and greatest elements that we have already seen are:

- In $(\mathbb{N}, \leqslant)$, 0 is a least element; there is no greatest element.

- Let $n \in \mathbb{N}$ with $n > 0$. Then 1 is a least element of $([n], \leqslant)$, and n is a greatest element.

- $(\mathbb{Z}, \leqslant)$ has no greatest or least elements.

$\triangleleft$

Proposition 5.2.7 says that least and greatest elements of posets are unique, if they exist. This allows us to talk about 'the' least or 'the' greatest element of a poset.

Proposition 5.2.7

Let $(X, \preccurlyeq)$ be a poset. If X has a least element, then it is unique; and if X has a greatest element, then it is unique.

Proof

Suppose X has a least element ℓ. We prove that if ℓ' is another least element, then $\ell' = \ell$.

So take another least element ℓ'. Since ℓ is a least element, we have $\ell \preccurlyeq \ell'$. Since ℓ' is a least element, we have $\ell' \preccurlyeq \ell$. By antisymmetry of $\preccurlyeq$, it follows that $\ell = \ell'$.

Hence least elements are unique. The proof for greatest elements is similar, and is left as an exercise. $\square$

Exercise 5.2.8

Let X be a set. The poset $(\mathscr{P}(X), \subseteq)$ has a least element and a greatest element; find both. $\triangleleft$

Exercise 5.2.9

Prove that the least element of $\mathbb{N}$ with respect to divisibility is 1, and the greatest element is 0. $\triangleleft$

Definition 5.2.10 (Supremum)

Let $(X, \preccurlyeq)$ be a poset and let $A \subseteq X$. A $\preccurlyeq$-**supremum** of A is an element $s \in X$ such that

- $a \preccurlyeq s$ for each $a \in A$; and

- If $s' \in X$ with $a \preccurlyeq s'$ for all $a \in A$, then $s \preccurlyeq s'$.

A $\preccurlyeq$-**infimum** of A is an element $i \in X$ such that

- $i \preccurlyeq a$ for each $a \in A$; and

- If $i' \in X$ with $i' \preccurlyeq a$ for all $a \in A$, then $i' \preccurlyeq i$.

Example 5.2.11

The well-ordering principle states that if $U \subseteq \mathbb{N}$ is inhabited then U has a $\leqslant$-infimum, and moreover the infimum of U is an element of U. ◁

Exercise 5.2.12

Let X be a set, and let $U, V \in \mathscr{P}(X)$. Prove that the $\subseteq$-supremum of $\{U, V\}$ is $U \cup V$, and the $\subseteq$-infimum of $\{U, V\}$ is $U \cap V$. ◁

Exercise 5.2.13

Let $a, b \in \mathbb{N}$. Show that $\gcd(a, b)$ is an infimum of $\{a, b\}$ and that $\mathrm{lcm}(a, b)$ is a supremum of $\{a, b\}$ with respect to divisbility. ◁

Example 5.2.14

Define $U = [0, 1) = \{x \in \mathbb{R} \mid 0 \leqslant x < 1\}$. We prove that U has both an infimum and a supremum in the poset $(\mathbb{R}, \leqslant)$.

- **Infimum.** 0 is an infimum for U. Indeed:

    (i) Let $x \in U$. Then $0 \leqslant x$ by definition of U.

    (ii) Let $y \in \mathbb{R}$ and suppose that $y \leqslant x$ for all $x \in U$. Then $y \leqslant 0$, since $0 \in U$.

  so 0 is as required.

- **Supremum.** 1 is a supremum for U. Indeed:

    (i) Let $x \in U$. Then $x < 1$ by definition of U, so certainly $x \leqslant 1$.

    (ii) Let $y \in \mathbb{R}$ and suppose that $x \leqslant y$ for all $x \in U$. We prove that $1 \leqslant y$ by contradiction. So suppose it is not the case that $1 \leqslant y$. Then $y < 1$. Since $x \leqslant y$ for all $x \in U$, we have $0 \leqslant y$. But then

$$0 \leqslant y = \frac{y + y}{2} < \frac{y + 1}{2} < \frac{1 + 1}{2} = 1$$

  But then $\frac{y+1}{2} \in U$ and $y < \frac{y+1}{2}$. This contradicts the assumption that $x \leqslant y$ for all $x \in U$. So it must in fact have been the case that $1 \leqslant y$.

so 1 is as required.

◁

The following proposition proves that suprema and infima are unique, provided they exist.

Proposition 5.2.15

Let $(X, \preccurlyeq)$ is a poset, and let $A \subseteq X$.

(i) If $s, s' \in X$ are suprema of A, then $s = s'$;

(ii) If $i, i' \in X$ are infima of A, then $i = i'$.

Proof

Suppose s, s' are suprema of A. Then:

- $a \preccurlyeq s'$ for all $a \in A$, so $s' \preccurlyeq s$ since s is a supremum of A;

- $a \preccurlyeq s$ for all $a \in A$, so $s \preccurlyeq s'$ since s' is a supremum of A.

Since $\preccurlyeq$ is antisymmetric, it follows that $s = s'$. This proves (i).

The proof of (ii) is almost identical and is left as an exercise to the reader. □

Notation 5.2.16

Let $(X, \preccurlyeq)$ be a poset and let $U \subseteq X$. Denote the $\preccurlyeq$-infimum of U, if it exists, by $\bigwedge U$ (LaTeX code: `\bigwedge`); and denote the $\preccurlyeq$-supremum of U, if it exists, by $\bigvee U$ (LaTeX code: `\bigvee`). Moreover, for $x, y \in X$, write

$$\bigwedge \{x, y\} = x \wedge y \text{ (LaTeX code: } \texttt{\textbackslash wedge}), \qquad \bigvee \{x, y\} = x \vee y \text{ (LaTeX code: } \texttt{\textbackslash vee})$$

Example 5.2.17

Some examples of Notation 5.2.16 are as follows.

- Let X be a set. In $(\mathscr{P}(X), \subseteq)$ we have $U \wedge V = U \cap V$ and $U \vee V = U \cup V$ for all $U, V \in \mathscr{P}(X)$.

- We have seen that, in $(\mathbb{N}, |)$, we have $a \wedge b = \gcd(a, b)$ and $a \vee b = \mathrm{lcm}(a, b)$ for all $a, b \in \mathbb{N}$.

- In $(\mathbb{R}, \leqslant)$, we have $a \wedge b = \min\{a, b\}$ and $a \vee b = \max\{a, b\}$.

◁

Definition 5.2.18

A **lattice** is a poset $(X, \preccurlyeq)$ such that every pair of elements of X has a $\preccurlyeq$-supremum and a $\preccurlyeq$-infimum.

Example 5.2.19

We have seen that $(\mathscr{P}(X), \subseteq)$, $(\mathbb{R}, \leqslant)$ and $(\mathbb{N}, |)$ are lattices.   ◁

Proposition 5.2.20 (Associativity laws for lattices)

Let $(X, \preccurlyeq)$ be a lattice, and let $x, y, z \in X$. Then

$$x \wedge (y \wedge z) = (x \wedge y) \wedge z \quad \text{and} \quad x \vee (y \vee z) = (x \vee y) \vee z$$

Proof

We prove $x \wedge (y \wedge z) = (x \wedge y) \wedge z$; the other equation is dual and is left as an exercise. We prove that the sets $\{x, y \wedge z\}$ and $\{x \wedge y, z\}$ have the same sets of lower bounds, and hence the same infima. So let

$$L_1 = \{i \in X \mid i \preccurlyeq x \text{ and } i \preccurlyeq y \wedge z\} \quad \text{and} \quad L_2 = \{i \in X \mid i \preccurlyeq x \wedge y \text{ and } i \preccurlyeq z\}$$

We prove $L_1 = L = L_2$, where

$$L = \{i \in X \mid i \preccurlyeq x, \ i \preccurlyeq y \text{ and } i \preccurlyeq z\}$$

First we prove $L_1 = L$. Indeed:

- $L_1 \subseteq L$. To see this, suppose $i \in L_1$. Then $i \preccurlyeq x$ by definition of L_1. Since $i \preccurlyeq y \wedge z$, and $y \wedge z \preccurlyeq y$ and $y \wedge z \preccurlyeq z$, we have $i \preccurlyeq y$ and $i \preccurlyeq z$ by transitivity of $\preccurlyeq$.

- $L \subseteq L_1$. To see this, suppose $i \in L$. Then $i \preccurlyeq x$ by definition of L. Moreover, $i \preccurlyeq y$ and $i \preccurlyeq z$ by definition of L, so that $i \preccurlyeq y \wedge z$ by definition of $\wedge$. Hence $i \in L$.

The proof that $L_2 = L$ is similar. Hence $L_1 = L_2$. But $x \wedge (y \wedge z)$ is, by definition of $\wedge$, the $\preccurlyeq$-greatest element of L_1, which exists since $(X, \preccurlyeq)$ is a lattice. Likewise, $(x \wedge y) \wedge z$ is the $\preccurlyeq$-greatest element of L_2.

Since $L_1 = L_2$, it follows that $x \wedge (y \wedge z) = (x \wedge y) \wedge z$, as required.   □

Exercise 5.2.21 (Commutativity laws for lattices)

Let $(X, \preccurlyeq)$ be a lattice. Prove that, for all $x, y \in X$, we have

$$x \wedge y = y \wedge x \quad \text{and} \quad x \vee y = y \vee x$$

◁

Exercise 5.2.22 (Absorption laws for lattices)

Let $(X, \preccurlyeq)$ be a lattice. Prove that, for all $x, y \in X$, we have

$$x \vee (x \wedge y) = x \quad \text{and} \quad x \wedge (x \vee y) = x$$

◁

Example 5.2.23

It follows from what we've proved that if $a, b, c \in \mathbb{Z}$ then

$$\gcd(a, \gcd(b,c)) = \gcd(\gcd(a,b), c)$$

For example, take $a = 882$, $b = 588$ and $c = 252$. Then

- $\gcd(b,c) = 84$, so $\gcd(a, \gcd(b,c)) = \gcd(882, 84) = 42$;
- $\gcd(a,b) = 294$, so $\gcd(\gcd(a,b), c) = \gcd(294, 252) = 42$.

These are indeed equal. ◁

Distributive lattices and Boolean algebras

One particularly important class of lattice is that of a *distributive lattice*, in which suprema and infima interact in a particularly convenient way. This makes algebraic manipulations of expressions involving suprema and infima particularly simple.

Definition 5.2.24
A lattice $(X, \preccurlyeq)$ is **distributive** if

$$x \wedge (y \vee z) = (x \wedge y) \vee (x \wedge z) \quad \text{and} \quad x \vee (y \wedge z) = (x \vee y) \wedge (x \vee z)$$

for all $x, y, z \in X$.

Example 5.2.25

For any set X, the power set lattice $(\mathscr{P}(X), \subseteq)$ is distributive. That is to say that for all $U, V, W \subseteq X$ we have

$$U \cap (V \cup W) = (U \cap V) \cup (U \cap W) \quad \text{and} \quad U \cup (V \cap W) = (U \cup V) \cap (U \cup W)$$

This was the content of Example 2.1.47 and Exercise 2.1.48. ◁

Exercise 5.2.26
Prove that $(\mathbb{N}, |)$ is a distributive lattice. ◁

Definition 5.2.27
Let $(X, \preccurlyeq)$ be a lattice with a greatest element $\top$ and a least element $\bot$, and let $x \in X$. A **complement** for x is an element y such that

$$x \wedge y = \bot \quad \text{and} \quad x \vee y = \top$$

Example 5.2.28

Let X be a set. We show that every element $U \in \mathscr{P}(X)$ has a complement.   ◁

Exercise 5.2.29

Let $(X, \preccurlyeq)$ be a distributive lattice with a greatest element and a least element, and let $x \in X$. Prove that, if a complement for x exists, then it is unique; that is, prove that if $y, y' \in X$ are complements for X, then $y = y'$.   ◁

Exercise 5.2.29 justifies the following notation.

Notation 5.2.30

Let $(X, \preccurlyeq)$ be a distributive lattice with greatest and least elements. If $x \in X$ has a complement, denote it by $\neg x$.

Definition 5.2.31

A lattice $(X, \preccurlyeq)$ is **complemented** if every element $x \in X$ has a complement. A **Boolean algebra** is a complemented distributive lattice with a greatest element and a least element.

The many preceding examples and exercises concerning $(\mathscr{P}(X), \subseteq)$ piece together to provide a proof of the following theorem.

Theorem 5.2.32

Let X be a set. Then $(\mathscr{P}(X), \subseteq)$ is a Boolean algebra.

Another extremely important example of a Boolean algebra is known as the *Lindenbaum–Tarski algebra*, which we define in Definition 5.2.35. In order to define it, we need to prove that the definition will make sense. First of all, we fix some notation.

Definition 5.2.33

Let P be a set, thought of as a set of propositional variables. Write $L(P)$ to denote the set of propositional formulae with propositional variables in P—that is, the elements of $L(P)$ are strings built from the elements of P, using the operations of conjunction ($\wedge$), disjunction ($\vee$) and negation ($\neg$).

Lemma 5.2.34

Logical equivalence $\equiv$ is an equivalence relation on $L(P)$.

Proof

This is immediate from definition of equivalence relation, since for $s, t \in L(P)$, $s \equiv t$ is defined to mean that s and t have the same truth values for all assignments of truth values to their propositional variables.   □

In what follows, the set P of propositional variables is fixed; we may moreover take it to be countably infinite, since all strings in $L(P)$ are finite.

Definition 5.2.35

The **Lindenbaum–Tarski algebra** (**for propositional logic**) over P is the pair $(A, \vdash)$, where $A = L(P)/\equiv$ and $\vdash$ is the relation on A defined by $[s]_{\equiv} \vdash [t]_{\equiv}$ if and only if $s \Rightarrow t$ is a tautology.

In what follows, we will simply write $[-]$ for $[-]_{\equiv}$.

Theorem 5.2.36

The Lindenbaum–Tarski algebra is a Boolean algebra.

Proof *Sketch proof*

There is lots to prove here! Indeed, we must prove:

- $\vdash$ is a well-defined relation on A; that is, if $s \equiv s'$ and $t \equiv t'$ then we must have $[s] \vdash [t]$ if and only if $[s'] \vdash [t']$.

- $\vdash$ is a partial order on A; that is, it is reflexive, antisymmetric and transitive.

- The poset $(A, \vdash)$ is a lattice; that is, it has suprema and infima.

- The lattice $(A, \vdash)$ is distributive, has a greatest element and a least element, and is complemented.

We will omit most of the details, which are left as an exercise; instead, we outline what the components involved are.

The fact that $\vdash$ is a partial order can be proved as follows.

- Reflexivity of $\vdash$ follows from the fact that $s \Rightarrow s$ is a tautology for all propositional formulae s.

- Symmetry of $\vdash$ follows from the fact that, for all propositional formulae s, t, if $s \Leftrightarrow t$ is a tautology then s and t are logically equivalent.

- Transitivity of $\vdash$ follows immediately from transitivity of $\Rightarrow$.

The fact that $(A, \vdash)$ is a lattice can be proved by verifying that:

- Given $[s], [t] \in A$, the infimum $[s] \wedge [t]$ is given by conjunction, namely $[s] \wedge [t] = [s \wedge t]$.

- Given $[s], [t] \in A$, the supremum $[s] \vee [t]$ is given by disjunction, namely $[s] \vee [t] = [s \vee t]$.

Finally, distributivity of suprema and infima in $(A, \vdash)$ follows from the corresponding properties of conjunction and disjunction; $(A, \vdash)$ has greatest element $[p \Rightarrow p]$ and least element $[\neg(p \Rightarrow p)]$, where p is some fixed propositional variable; and the complement of $[s] \in A$ is given by $[\neg s]$. □

We finish this section on orders and lattices with a general version of de Morgan's laws for Boolean algebras, which by Theorems Theorems 5.2.32 and 5.2.36 implies the versions we proved for logical formulae (Theorem 1.3.24) and for sets (Theorem 2.1.62(a)–(b)).

Theorem 5.2.37 (De Morgan's laws)
Let $(X, \preccurlyeq)$ be a Boolean algebra, and let $x, y \in X$. Then

$$\neg(x \wedge y) = (\neg x) \vee (\neg y) \quad \text{and} \quad \neg(x \vee y) = (\neg x) \wedge (\neg y)$$

Proof
To do: □

Well-foundedness and structural induction

Warning!
This section is not yet finished—do not rely on its correctness or completeness.

Section 3.1 introduced induction as a technique for proving statements which are true of all natural numbers. We saw induction in three flavours: weak induction, strong induction and the well-ordering principle.

- The **weak induction principle** (Theorem 3.1.14) exploited the *inductively defined* structure of $\mathbb{N}$. Every natural number can be obtained from 0 by repeatedly applying the successor ('plus one') operation, so if a statement $p(n)$ is true of 0, and its truth is preserved by the successor operation (i.e. if $p(n) \Rightarrow p(n+1)$ is true for all $n \in \mathbb{N}$), then it must be true of all natural numbers

- The **well-ordering principle** (Theorem 3.1.47) exploited the *well-founded* nature of the order relation $<$ on $\mathbb{N}$. It says that every inhabited subset of $\mathbb{N}$, so that any proposition $p(n)$ which is *not* true of all natural numbers n must have a least counterexample—this led to the technique of proof by infinite descent.

In this section, we will generalise these techniques to other sets with an *inductively defined* or a *well-founded* structure.

- An *inductively defined set* will, intuitively, be a set X built from some set of *basic elements* (like zero) using a set of *constructors* (like the successor operation). We will be able to perform induction on these sets to prove that a statement $p(x)$ is true for all $x \in X$ by proving that it is true for the basic elements, and then proving that its truth is preserved by the constructors. This proof technique generalises weak induction and is called *structural induction*.

- A set X with a *well-founded relation R* will allow us to generalise proof by infinite descent: if there is a counterexample to a logical formula $p(x)$, then there must be one which is 'minimal' with respect to R. This leads to a proof technique called *well-founded induction*, which has similarities with strong induction.

Structural induction is conceptually easier to comprehend than well-founded induction, so we will introduce it first. However, we will not be able to prove that it is a valid proof technique until after we have introduced well-founded induction.

Inductively defined sets

In Section 3.1, we formalised the idea that the set of natural numbers should be what is obtained by starting with zero and repeating the successor ('plus one') operation. In a sense, zero was a *basic element*—we posited its existence from the outset—and the successor operation *constructed* the remaining elements.

Although hidden beneath the surface, this method of defining a set was implicitly used in Chapter 1 when defining propositional formulae. Here, our *basic elements* were propositional variables $p, q, r, s, \ldots$, and the remaining propositional formulae could be *constructed* by repeatedly applying the logical connectives $\wedge$, $\vee$, $\neg$ and $\Rightarrow$.

To do:

Definition 5.3.1

An **inductively defined set** is a set A together with a set C_A of functions

$$\sigma : A^{n(\sigma)} \to A$$

where $n(\sigma) \in \mathbb{N}$ for each $\sigma \in C_A$, such that:

(i) For each $a \in A$, there exists a unique $\sigma \in C_A$ and unique elements $a_1, a_2, \ldots, a_{n(\sigma)} \in A$ such that $a = \sigma(a_1, a_2, \ldots, a_{n(\sigma)})$; and

(ii) For all sets X, if $\sigma(a_1, a_2, \ldots, a_n) \in A$ for all $\sigma \in C_A$ and all $a_1, a_2, \ldots, a_{n(\sigma)} \in A$, then $A \subseteq X$.

The elements $\sigma \in C_A$ are called the **constructors** of A, and the natural number $n(\sigma)$ is called the **arity** of $\sigma \in C_A$.

A quick note on terminology: a constructor of arity $n \in \mathbb{N}$ is called an *n-ary* constructor. For $n = 0, 1, 2$, we may say *nullary*, *unary* and *binary*, respectively.

Note that $A^0 = \{()\}$, where $()$ is the empty list of elements of A. Since A^0 only has one element, specifying a function $\sigma : A^0 \to A$ is equivalent to specifying an element $a = \sigma(()) \in A$. With this in mind, instead of thinking of a constructor $\sigma : A^0 \to A$ as being a function, we think of σ as *being* an element of A.

Definition 5.3.2

Given an inductive set (A, C_A), a nullary constructor $\sigma \in C_A$—considered as an element of A as mentioned above—is called a **basic element** of A.

Example 5.3.3

The set $\mathbb{N}$ of natural numbers is an inductively defined set. The set $C_{\mathbb{N}}$ of constructors is given by $C_{\mathbb{N}} = \{0, s\}$, where $0 \in \mathbb{N}$ is a basic element $s : \mathbb{N} \to \mathbb{N}$ is a constructor of arity 1 defined by $s(n) = n + 1$ for all $n \in \mathbb{N}$.

To see this, we will verify the conditions of Definition 5.3.1 by observing that they are essentially just restatements of the Peano axioms (Definition 3.1.1). Indeed:

(i) Let $n \in \mathbb{N}$.

- If $n = 0$, then $n \neq s(m)$ for any $m \in \mathbb{N}$ by Definition 3.1.1(i), so '0' is the unique expression for 0 as a constructor applied to elements of $\mathbb{N}$.

- If $n > 0$, then $n - 1 \in \mathbb{N}$ and $n = s(n-1)$. To see that this expression is unique, note that if $m \in \mathbb{N}$ and $s(m) = n$, then $m = n - 1$ by Definition 3.1.1(ii).

(ii) We need to prove that if X is a set such that $0 \in X$ and $s(n) \in X$ for all $n \in \mathbb{N}$, then $\mathbb{N} \subseteq X$. But this is exactly Definition 3.1.1(iii).

Hence $\mathbb{N}$ is indeed an inductively defined set. ◁

Exercise 5.3.4

Prove that the set $A = \{1, 2, 4, 8, 16 \dots\}$ of (natural number) powers of 2 is inductively defined by taking $C_A = \{1, d\}$, where 1 is basic and $d : A \to A$ is defined by $d(n) = 2n$ for all $n \in \mathbb{N}$. ◁

The next exercise gives a different way of inductively defining $\mathbb{N}$—it demonstrates that we can consider a set to be inductively defined in more than one way.

Exercise 5.3.5

Prove that $\mathbb{N}$ is inductively defined by taking $C_{\mathbb{N}} = \{0, \sigma\}$, where $\sigma : \mathbb{N} \to \mathbb{N}$ is defined by

$$\sigma(n) = \begin{cases} n + 2 & \text{if } n \equiv 0 \text{ or } 1 \bmod 3 \\ n - 1 & \text{if } n \equiv 2 \bmod 3 \end{cases}$$

for all $n \in \mathbb{N}$. ◁

To do: Example: propositional formulae

Theorem 5.3.6 (Principle of structural induction)

Let X be an inductively defined set, and let $p(x)$ be a logical formula concerning elements of X. Suppose that

- $p(b)$ is true for all basic elements $b \in X$; and

- For all constructors f of arity n and all $x_1, x_2, \ldots, x_n \in X$, if $p(x_1), p(x_2), \ldots, p(x_n)$ are all true, then $p(f(x_1, x_2, \ldots, x_n))$ is true.

Then $p(x)$ is true for all $x \in X$.

We will prove Theorem 5.3.6 on page 287.

Example 5.3.7

To do: Structural induction on $\mathbb{N}$ is weak induction.   ◁

To do: Disjunctive normal form

To do: Generalise to quotients of inductive structures ⤳ induction on $\mathbb{Z}$ using 0 and $+, -$ and on $\mathbb{Z}^{>0}$ using 1 and $p \times (-)$.

We saw in Proposition 5.3.14 that the relation R on the set $\mathbb{Z}^{>0}$ of positive integers defined for $m, n \in \mathbb{Z}^{>0}$ by

$$m \, R \, n \quad \Leftrightarrow \quad n = pm \text{ for some prime } p > 0$$

is well-founded. We can use well-founded induction to prove a general formula for the totient of an integer n.

Theorem 5.3.8 (Formula for Euler's totient function)

Let $n \in \mathbb{Z}$ be nonzero, and let $\varphi : \mathbb{Z} \to \mathbb{N}$ be Euler's totient function (see Definition 4.3.31). Then

$$\varphi(n) = |n| \cdot \prod_{p \mid n \text{ prime}} \left(1 - \frac{1}{p}\right)$$

where the product is indexed over the distinct positive prime factors p of n.

Proof

If $n < 0$ then $\varphi(n) = \varphi(-n)$, $|n| = -n$ and $p \mid n$ if and only if $p \mid -n$, so the theorem holds for negative integers if and only if it holds for positive integers.

We prove the formula for $n > 0$ by well-founded induction on $\mathbb{Z}^{>0}$ with respect to the relation R defined in Proposition 5.3.14.

- **(BC)** $\varphi(1) = 1$ and, since no prime p divides 1, we have $\prod_{p|1 \text{ prime}} \left(1 - \frac{1}{p}\right) = 1$.

  Hence

  $$1 \cdot \prod_{p|1 \text{ prime}} \left(1 - \frac{1}{p}\right) = 1 \cdot 1 = 1$$

  as erquired.

- **(IS)** Fix $n \geqslant 1$ and suppose that

  $$\varphi(n) = n \cdot \prod_{p|n \text{ prime}} \left(1 - \frac{1}{p}\right)$$

  Let $q > 0$ be prime. We prove that

  $$\varphi(qn) = qn \cdot \prod_{p|qn \text{ prime}} \left(1 - \frac{1}{p}\right)$$

  ◇ Suppose $q \mid n$. Then by we have

  $$\varphi(qn) = q\varphi(n) \qquad \text{by Exercise 4.3.33}$$

  $$= qn \cdot \prod_{p|n \, prime} \left(1 - \frac{1}{p}\right) \qquad \text{by induction hypothesis}$$

  $$= qn \cdot \prod_{p|qn \, prime} \left(1 - \frac{1}{p}\right)$$

  The last equation holds because the fact that $q \mid n$ implies that, for all positive primes p, we have $p \mid n$ if and only if $p \mid qn$.

  ◇ Suppose $q \nmid n$. Then $q \perp n$, so we have

  $$\varphi(qn) = \varphi(q)\varphi(n) \qquad \text{by Theorem 4.3.63}$$

  $$= \varphi(q) \cdot n \cdot \prod_{p|n \text{ prime}} \left(1 - \frac{1}{p}\right) \qquad \text{by induction hypothesis}$$

  $$= (q-1) \cdot n \cdot \prod_{p|n \text{ prime}} \left(1 - \frac{1}{q}\right) \qquad \text{by Example 4.3.32}$$

  $$= q\left(1 - \frac{1}{p}\right) n \cdot \prod_{p|n \text{ prime}} \left(1 - \frac{1}{p}\right) \qquad \text{rearranging}$$

  $$= qn \cdot \left(\prod_{p|n \text{ prime}} \left(1 - \frac{1}{p}\right)\right) \cdot \left(1 - \frac{1}{q}\right) \qquad \text{rearranging}$$

  $$= qn \cdot \prod_{p|qn} \left(1 - \frac{1}{p}\right) \qquad \text{reindexing the product}$$

  In both cases, we have shown that the formula holds.

By induction, we're done. □

Well-founded relations

First, we introduce the notion of a *well-founded relation*.

Definition 5.3.9
Let X be a set. A relation R on X is **well-founded** if every inhabited subset of X has an **R-minimal** element, in the following sense: for each inhabited $U \subseteq X$, there exists $m \in U$ such that $\neg(x \, R \, m)$ for all $x \in U$. A relation that is not well-founded is called **ill-founded**.

Example 5.3.10
The relation $<$ on $\mathbb{N}$ is well-founded—this is just a fancy way of stating the well-ordering principle (Theorem 3.1.47). Indeed, let $U \subseteq \mathbb{N}$ be an inhabited subset. By the well-ordering principle, there exists an element $m \in U$ such that $m \leqslant x$ for all $x \in U$. But this says precisely that $\neg(x < m)$ for all $x \in U$.   ◁

Example 5.3.11
However, the relation $<$ on $\mathbb{Z}$ is not well-founded—indeed, $\mathbb{Z}$ is an inhabited subset of $\mathbb{Z}$ with no $<$-least element.   ◁

Exercise 5.3.12
Let $<^1$ be the relation on $\mathbb{N}$ defined for $m, n \in \mathbb{N}$ by

$$m <^1 n \quad \Leftrightarrow \quad n = m+1$$

Prove that $<^1$ is a well-founded relation on $\mathbb{N}$.   ◁

Proposition 5.3.13
Let X be a set and let R be a relation on X. R is well-founded if and only if there is no infinite R-descending chains; that is, there does not exist a sequence $(x_n)_{n \in \mathbb{N}}$ of elements of X such that $x_{n+1} \, R \, x_n$ for all $n \in \mathbb{N}$.

Proof
We prove the contrapositives of the two directions; that is, R is ill-founded if and only if R has an infinite descending R-chain.

- ($\Rightarrow$) Suppose that R is ill-founded, and let $U \subseteq X$ be an inhabited subset with no R-minimal element. Define a sequence $(x_n)_{n \in \mathbb{N}}$ of elements of X—in fact, of U—recursively as follows:

  ◇ Let $x_0 \in U$ be arbitrarily chosen.
  ◇ Fix $n \in \mathbb{N}$ and suppose $x_0, x_1, \ldots, x_n \in U$ have been defined. Since U has no R-minimal element, it contains an element which is related to x_n by R; define x_{n+1} to be such an element.

  Then $(x_n)_{n \in \mathbb{N}}$ is an infinite R-descending chain

- ($\Leftarrow$) Suppose there is an infinite R-descending chain $(x_n)_{n \in \mathbb{N}}$. Define $U = \{x_n \mid n \in \mathbb{N}\}$ to be the set of elements in this sequence. Then U has no R-minimal element. Indeed, given $m \in U$, we must have $m = x_n$ for some $n \in \mathbb{N}$; but then $x_{n+1} \in U$ and $x_{n+1} \, R \, m$. Hence R is ill-founded.

$\square$

Proposition 5.3.14

Let $\mathbb{Z}^{>0}$ be the set of positive integers and define a relation R on $\mathbb{Z}^{>0}$ by

$$m \, R \, n \quad \Leftrightarrow \quad n = pm \text{ for some prime } p > 0$$

for all $m, n > 0$. Then R is a well-founded relation on $\mathbb{Z}^{>0}$.

Proof

Suppose that $(x_n)_{n \in \mathbb{N}}$ is an infinite R-descending chain in $\mathbb{Z}^{>0}$. Since $x_{n+1} \, R \, x_n$ for all $n \in \mathbb{N}$, we have $x_n = px_{n+1}$ for some positive prime p for all $n \in \mathbb{N}$. Since all positive primes are greater than or equal to 2, this implies that $x_n \geqslant 2x_{n+1}$ for all $n \in \mathbb{N}$.

We prove by strong induction on $n \in \mathbb{N}$ that $x_0 > 2^n x_{n+1}$ for all $n \in \mathbb{N}$.

- **(BC)** We proved above that $x_0 \geqslant 2x_1$. Hence $x_0 > x_1 = 2^0 x_1$, as required.

- **(IS)** Fix $n \in \mathbb{N}$ and suppose $x_0 > 2^n x_{n+1}$. We want to show $x_0 > 2^{n+1} x_{n+2}$. Well $x_{n+1} > 2x_{n+2}$, as proved above, and hence

$$x_0 \overset{\text{IH}}{>} 2^n x_{n+1} > 2^n \cdot 2x_{n+2} = 2^{n+1} x_{n+2}$$

  as required.

By induction, we've shown that $x_0 > 2^n x_{n+1}$ for all $n \in \mathbb{N}$. But $x_{n+1} > 0$ for all $n \in \mathbb{N}$, so $x_0 > 2^n$ for all $n \in \mathbb{N}$. This implies that x_0 is greater than every integer, which is a contradiction.

So such a sequence $(x_n)_{n \in \mathbb{N}}$ cannot exist, and by Proposition 5.3.13, the relation R is well-founded.

$\square$

Exercise 5.3.15

Let X be a set and let R be a well-founded relation on X. Given $x, y \in X$, prove that not both $x \, R \, y$ and $y \, R \, x$ are true.

$\triangleleft$

Theorem 5.3.16 (Principle of well-founded induction)

Let X be a set, let R be a well-founded relation on X, and let $p(x)$ be a logical formula concerning elements of X. Suppose that for each $x \in X$, the following is true:

If $p(y)$ is true for all R-predecessors y of x, then $p(x)$ is true.

That is, suppose for each $x \in X$ that

$$[\forall y \in X, (y \, R \, x \Rightarrow p(y))] \Rightarrow p(x)$$

Then $p(x)$ is true for all $x \in X$.

Proof

Suppose that, for each $x \in X$, if $p(y)$ is true for all R-predecessors y of x, then $p(x)$ is true. Let

$$U = \{x \in X \mid \neg p(x)\}$$

Towards a contradiction, suppose that $p(x)$ is false for some $x \in X$. Then U is inhabited. Since R is well-founded, U has an R-minimal element $m \in U$. Now

(i) $p(m)$ is false, since $m \in U$.

(ii) $p(x)$ is true for all $x \in X$ with $x \, R \, m$. To see this, note that if $p(x)$ is false and $x \, R \, m$, then $x \in U$, so that $m \, R \, x$ by R-minimality of m in U. Since also $x \, R \, m$, this contradicts Exercise 5.3.15.

Since $p(x)$ is true for all $x \in X$ with $x \, R \, m$, by assumption we also have that $p(m)$ is true. But this contradicts our assumption that $m \in U$.

So it must in fact be the case that $U = \emptyset$, so that $p(x)$ is true for all $x \in X$.   □

Exercise 5.3.17

Prove that the principle of $<$-induction on $\mathbb{N}$ is precisely strong induction. Specifically, prove that the following two statements are equivalent:

(i) $p(0)$ is true and, for all $n \in \mathbb{N}$, if $p(k)$ is true for all $k \leqslant n$, then $p(n+1)$ is true;

(ii) For all $n \in \mathbb{N}$, if $p(k)$ is true for all $k < n$, then $p(n)$ is true.

Strong induction says that we can deduce that $p(n)$ is true for all $n \in \mathbb{N}$ from the knowledge that (i) is true for all $n \in \mathbb{N}$; and $<$-induction tells us that $p(n)$ is true for all $n \in \mathbb{N}$ from the knowledge that (ii) is true for all $n \in \mathbb{N}$. You should prove that (i) and (ii) are equivalent.   ◁

Example 5.3.18

Let $<^1$ be the relation on $\mathbb{N}$ defined in Exercise 5.3.12. We prove that the principle of $<^1$-induction on $\mathbb{N}$ is precisely strong induction. Specifically, prove that the following two statements are equivalent:

  (i) $p(0)$ is true and, for all $n \in \mathbb{N}$, if $p(n)$ is true then $p(n+1)$ is true;

  (ii) For all $n \in \mathbb{N}$, if $p(k)$ is true for all $k \in \mathbb{N}$ with $k+1 = n$, then $p(n)$ is true.

Weak induction says that we can deduce that $p(n)$ is true for all $n \in \mathbb{N}$ from the knowledge that (i) is true for all $n \in \mathbb{N}$; and $<^1$-induction tells us that $p(n)$ is true for all $n \in \mathbb{N}$ from the knowledge that (ii) is true for all $n \in \mathbb{N}$. We prove that (i) and (ii) are equivalent.

- (i) $\Rightarrow$ (ii). Suppose that $p(0)$ and, for all $n \in \mathbb{N}$, if $p(n)$ is true then $p(n+1)$ is true. We will prove that

$$[\forall m \in \mathbb{N}, (n = m+1 \Rightarrow p(m))] \Rightarrow p(n)$$

  is true for all $n \in \mathbb{N}$.

  So fix $n \in \mathbb{N}$, and assume $\forall m \in \mathbb{N}, (n = m+1 \Rightarrow p(m))$. We prove $p(n)$ is true.

  ◇ If $n = 0$ then we're done, since $p(0)$ is true by assumption.

  ◇ If $n > 0$ then $n = m+1$ for some $m \in \mathbb{N}$. By our assumption, we have $\forall m \in \mathbb{N}, (n = m+1 \Rightarrow p(m))$, and so in particular, $p(m)$ is true. By the weak induction step, we have $p(m) \Rightarrow p(m+1)$ is true. But then $p(m+1)$ is true. Since $n = m+1$, we have that $p(n)$ is true.

  In any case, we've proved that $p(n)$ is true, as required.

- (ii) $\Rightarrow$ (i). For $n \in \mathbb{N}$, denote the following statement by $H(n)$

$$[\forall m \in \mathbb{N}, (n = m+1 \Rightarrow p(m))] \Rightarrow p(n)$$

  Assume $H(n)$ is true for all $n \in \mathbb{N}$. We prove that $p(0)$ is true and, for all $n \in \mathbb{N}$, if $p(n)$ is true then $p(n+1)$ is true.

  ◇ $p(0)$ is true. Indeed, for any $m \in \mathbb{N}$ we have that $0 = m+1$ is false, so the statement $0 = m+1 \Rightarrow p(m)$ is true. Hence $\forall m \in \mathbb{N}, (0 = m+1 \Rightarrow p(m))$ is true. Since $H(0)$ is true, it follows that $p(0)$ is true.

  ◇ Fix $n \in \mathbb{N}$ and suppose $p(n)$ is true. By $H(n+1)$, we have that if $p(n+1)$ is true for all $m \in \mathbb{N}$ with $m+1 = n+1$, then $p(n+1)$ is true. But the only $m \in \mathbb{N}$ such that $m+1 = n+1$ is n itself, and $p(n)$ is true by assumption; so by $H(n+1)$, we have $p(n+1)$, as required.

Hence the two induction principles are equivalent. ◁

Example 5.3.19

◁

Structural induction from well-founded induction

We will now derive the principle of structural induction in terms of the principle of well-founded induction. To do this, we need to associate to each inductively defined set X a corresponding well-founded relation R_X, such that well-founded induction on R_X corresponds with structural induction on X.

Definition 5.3.20

Let X be an inductively defined set. Define a relation R_X on X as follows: for all $x, y \in X$, $x \, R_X \, y$ if and only if

$$y = f(x_1, x_2, \ldots, x_n)$$

for some constructor f of arity n and elements $x_1, x_2, \ldots, x_n$, such that $x_i = x$ for some $i \in [n]$.

Example 5.3.21
Let $\mathbb{N}$ be the set of natural numbers, taken to be inductively defined in the usual way. Since the only constructor is the successor operation, we must have for $m, n \in \mathbb{N}$ that

$$m \, R_\mathbb{N} \, n \quad \Leftrightarrow \quad n = m + 1$$

This is precisely the relation $<^1$ from Exercise 5.3.12. We already established that structural induction on $\mathbb{N}$ is precisely weak induction (Example 5.3.7), and that well-founded induction on $<^1$ is also precisely weak induction (Example 5.3.18).

$\triangleleft$

Example 5.3.22
Let P be a set of propositional variables and let $L(P)$ be the set of propositional formulae built from variables in P and the logical operators $\wedge$, $\vee$, $\Rightarrow$ and $\neg$.

Then $R = R_{L(P)}$ is the relation defined for $s, t \in L(P)$ by letting $s \, R \, t$ if and only if

$$t \in \{s \wedge u, \; u \wedge s, \; s \vee u, \; u \vee s, \; s \Rightarrow u, \; u \Rightarrow s, \; \neg s\}$$

for some $u \in L(P)$.

$\triangleleft$

The plan for the rest of this section is to demonstrate that structural induction follows from well-founded induction. To do this, we prove that the relation R_X associated with an inductively defined set X is well-founded, and then we prove that structural induction on X is equivalent to well-founded induction on R_X.

To simplify our proofs, we introduce the notion of *rank*. The rank of an element x of an inductively defined set X is a natural number which says how many constructors need to be applied in order to obtain x.

Definition 5.3.23

Let X be an inductively defined set. The function rank $: X \to \mathbb{N}$ is defined recursively as follows:

- If b is a basic element of X, then $\mathrm{rank}(b) = 0$.

- Let f be a constructor of arity n and let $x_1, x_2, \ldots, x_n \in X$. Then

$$\mathrm{rank}(f(x_1, x_2, \ldots, x_n)) = \max\{\mathrm{rank}(x_1), \mathrm{rank}(x_2), \ldots, \mathrm{rank}(x_n)\} + 1$$

Note that rank $: X \to \mathbb{N}$ is a well-defined function, since by the conditions listed in Definition 5.3.1, every element of X is either basic or has a unique representation in the form $f(x_1, x_2, \ldots, x_n)$ for some constructor f and elements $x_1, x_2, \ldots, x_n \in X$.

Example 5.3.24

The rank function on the inductively defined set of natural numbers is fairly boring. Indeed, it tells us that

- $\mathrm{rank}(0) = 0$; and

- $\mathrm{rank}(n+1) = \mathrm{rank}(n) + 1$ for all $n \in \mathbb{N}$.

It can easily be seen that $\mathrm{rank}(n) = n$ for all $n \in \mathbb{N}$. This makes sense, since n can be obtained from 0 by iterating the successor operation n times. ◁

Lemma 5.3.25

Let X be an inductively defined set. The relation R_X defined in Definition 5.3.20 is well-founded.

Proof

☐

Proof of *Theorem 5.3.6*
To do: Write proof ☐

To do: Examples and exercises

Section 5.Q
Chapter 5 exercises

Under construction!
The end-of-chapter exercise sections are new and in an incomplete state.

1. For each of the eight subsets

$$P \subseteq \{\text{reflexive}, \text{symmetric}, \text{transitive}\}$$

find a relation satisfying (only) the properties in P.

2. Prove that if R is a symmetric, antisymmetric relation on a set X, then it is a subrelation of the equality relation—that is, $\text{Gr}(R) \subseteq \text{Gr}(=)$.

3. A relation R on a set X is **left-total** if for all $x \in X$, there exists some $y \in X$ such that $x R y$. Prove that every left-total, symmetric, transitive relation is reflexive.

Chapter 6

Infinite sets

In Section 3.2 we characterised *finiteness*, and defined a notion of *size* for finite sets, in terms of bijections of the form $[n] \to X$. This turned out to be extremely fruitful, as we were then able to compare sizes of finite sets by finding injections, surjections and bijections between them. For example, we showed that for any two finite sets X and Y, then $|X| \leqslant |Y|$ if and only if there is an injection $X \to Y$.

This chapter is dedicated to removing the requirement that the sets in question be finite, and then seeing what happens.

Our first step will be to characterise what can be thought of as the *smallest* size of infinity—countable infinity—in Section 6.1. Countable sets behave particularly nicely and satisfy some useful closure properties; we will also develop some techniques for finding when a set has too many elements to be countable.

Section 6.2 introduces the general concept of *cardinality* for comparing the sizes of infinite sets. This allows us to make finer distinctions between infinite sets than just 'countable' and 'uncountable'.

Finally, in Section 6.3, we study the interactions between the concept of infinity and the axiom of choice (see Axiom 2.3.33).

Section 6.1
Countable and uncountable sets

To do:

Definition 6.1.1
A set X is **countably infinite** if there exists a bijection $f : \mathbb{N} \to X$. The bijection f is called an **enumeration** of X. We say X is **countable** if it is finite or countably infinite.

Thus a set X is countably infinite if its elements can be *listed*, with one entry in the list for each natural number.

Example 6.1.2
We have already seen many examples of countably infinite sets.

- The set $\mathbb{N}$ is countably infinite, since by Exercise 2.3.18, $\mathrm{id}_{\mathbb{N}} : \mathbb{N} \to \mathbb{N}$ is a bijection.

- The function $f : \mathbb{Z} \to \mathbb{N}$ defined for $x \in \mathbb{Z}$ by

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

  is a bijection. Indeed, it has an inverse is given by

$$f^{-1}(x) = \begin{cases} \frac{x}{2} & \text{if } x \text{ is even} \\ -\frac{x+1}{2} & \text{if } x \text{ is odd} \end{cases}$$

  Hence the set of integers $\mathbb{Z}$ is countably infinite. The corresponding list of integers is given by

$$0, -1, 1, -2, 2, -3, 3, -4, 4, \ldots$$

  The fact that f^{-1} is a bijection means that each integer appears on this list exactly once.

$\triangleleft$

Exercise 6.1.3
Let $f : X \to Y$ be a bijection. Prove that X is countably infinite if and only if Y is countably infinite.

$\triangleleft$

Exercise 6.1.4

Prove that the function $p : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ defined by $p(x,y) = 2^x(2y+1) - 1$ is a bijection. Deduce that if X and Y are countably infinite sets, then $X \times Y$ is countably infinite. ◁

Exercise 6.1.4 allows us to prove that the product of finitely many countably infinite sets are countably infinite.

Proposition 6.1.5

Let $n \geqslant 1$ and let $X_1, \ldots, X_n$ be countably infinite sets. Then the product $\prod_{i=1}^{n} X_i$ is countably infinite.

Proof
We proceed by induction on n.

- **(BC)** When $n = 1$ the assertion is trivial: if X_1 is countably infinite then X_1 is countably infinite.

- **(IS)** Fix $n \geqslant 1$ and suppose that for any sets $X_1, \ldots, X_n$, the product $\prod_{i=1}^{n} X_i$ is countably infinite. Fix sets $X_1, \ldots, X_{n+1}$. Then $\prod_{i=1}^{n} X_i$ is countably infinite by the induction hypothesis, and X_{n+1} is countably infinite by assumption, so by Exercise 6.1.4, the set

$$\left(\prod_{i=1}^{n} X_i \right) \times X_{n+1}$$

is countably infinite. But by Exercise 2.3.19 there is a bijection

$$\prod_{i=1}^{n+1} X_i \to \left(\prod_{i=1}^{n} X_i \right) \times X_{n+1}$$

and so by Exercise 6.1.3 we have that $\prod_{i=1}^{n+1} X_i$ is countably infinite, as required.

By induction, we're done. □

Finding a *bijection* $\mathbb{N} \to X$, or equivalently $X \to \mathbb{N}$, can be a bit of a hassle. However, in order to prove that a set X is countable, it suffices to find either a surjection $\mathbb{N} \to X$ or an injection $X \to \mathbb{N}$.

> **Theorem 6.1.6**
>
> Let X be an inhabited set. The following are equivalent:
>
>   (i) X is countable;
>
> (ii) There exists a surjection $f : \mathbb{N} \to X$;
>
> (iii) There exists an injection $f : X \to \mathbb{N}$.

Proof

We'll prove (i)⇔(ii) and (i)⇔(iii).

- (i)⇒(ii). Suppose X is countable. If X is countably infinite, then there exists a bijection $f : \mathbb{N} \to X$, which is a surjection. If X is finite then there exists a bijection $g : [m] \to X$, where $m = |X| \geqslant 1$. Define $f : \mathbb{N} \to X$ by

$$f(n) = \begin{cases} g(n) & \text{if } 1 \leqslant n \leqslant m \\ g(1) & \text{if } n = 0 \text{ or } n > m \end{cases}$$

  Then f is surjective: if $x \in X$ then there exists $n \in [m]$ such that $g(n) = x$, and then $f(n) = g(n) = x$.

- (ii)⇒(i). Suppose there exists a surjection $f : \mathbb{N} \to X$. To prove that X is countable, it suffices to prove that if X is infinite then it is countably infinite. So suppose X is infinite, and define a sequence recursively by

  ◇ $a_0 = 0$;

  ◇ Fix $n \in \mathbb{N}$ and suppose $a_0, \ldots, a_n$ have been defined. Define a_{n+1} to be the least natural number for which $f(a_{n+1}) \notin \{f(a_0), f(a_1), \ldots, f(a_n)\}$.

  Define $g : \mathbb{N} \to X$ by $g(n) = f(a_n)$ for all $n \in \mathbb{N}$. Then

  ◇ g is injective, since if $m \leqslant n$ then $f(a_m) \neq f(a_n)$ by construction of the sequence $(a_n)_{n \in \mathbb{N}}$.

  ◇ g is surjective. Indeed, given $x \in X$, by surjectivity there exists $m \in \mathbb{N}$ which is least such that $f(m) = x$, and we must have $a_n = m$ for some $n \leqslant m$ by construction of the sequence $(a_n)_{n \in \mathbb{N}}$. So $x = f(a_n) = g(n)$, and hence g is surjective.

  So g is a bijection, and X is countable.

- (i)⇒(iii). Suppose X is countable. If X is countably infinite, then there exists a bijection $f : \mathbb{N} \to X$, so $f^{-1} : X \to \mathbb{N}$ is bijective and hence injective. If X is finite then there exists a bijection $g : [m] \to X$, where $m = |X| \geqslant 1$. Then $g^{-1} : X \to [m]$ is injective. Let $i : [m] \to \mathbb{N}$ be defined by $i(k) = k$ for all $k \in [m]$. Then $i \circ g^{-1}$ is injective; indeed, for $x, x' \in X$ we have

$$i(g^{-1}(x)) = i(g^{-1}(x')) \Rightarrow g^{-1}(x) = g^{-1}(x') \Rightarrow x = x'$$

The first implication is by definition of i, and the second is by injectivity of g^{-1}. So there exists an injection $X \to \mathbb{N}$.

- (iii)$\Rightarrow$(i). Suppose there exists an injection $f : X \to \mathbb{N}$. To prove that X is countable, it suffices to prove that if X is infinite then it is countably infinite. Define a sequence $(a_n)_{n \in \mathbb{N}}$ recursively as follows:

  ◇ Let a_0 be the least element of $f[X]$;

  ◇ Fix $n \in \mathbb{N}$ and suppose $a_0, \dots, a_n$ have been defined. Let a_{n+1} be the least element of $f[X] \setminus \{a_0, \dots, a_n\}$. This exists since f is injective, so $f[X]$ is infinite.

  Define $g : \mathbb{N} \to X$ by, for each $n \in \mathbb{N}$, letting $g(n)$ be the unique value of x for which $f(x) = a_n$. Then

  ◇ g is injective. By construction $a_m \ne a_n$ whenever $m \ne n$. Let $x, y \in X$ be such that $f(x) = a_m$ and $f(y) = a_n$. Since f is injective, we must have $x \ne y$, and so $g(m) = x \ne y = g(n)$.

  ◇ g is surjective. Fix $x \in X$. Then $f(x) \in f[X]$, so there exists $m \in \mathbb{N}$ such that $f(x) = a_m$. Hence $g(m) = x$.

  So g is a bijection, and X is countably infinite.

Hence the equivalences have been proved. $\square$

In fact, we needn't even use $\mathbb{N}$ as the domain of the surjection or the codomain of the injection; we can in fact use any countable set C.

Exercise 6.1.7
Let X be an inhabited set. The following are equivalent:

(a) X is countable;

(b) There exists a surjection $f : C \to X$ for some countable set C;

(c) There exists an injection $f : X \to C$ for some countable set C.

◁

Exercise 6.1.7 is useful for proving the countability of many other sets: as we build up our repertoire of countable sets, all we need to do in order to prove a set X is countable is find a surjection from a set we already know is countable to X, or an injection from X into a set we already know is countable.

This proof technique yields an incredibly short proof of the following counter-intuitive result, which can be interpreted to mean that there are exactly as many rational numbers as there are natural numbers.

Theorem 6.1.8
The set $\mathbb{Q}$ of rational numbers is countable.

Proof
Define a function $q : \mathbb{Z} \times (\mathbb{Z} \setminus \{0\}) \to \mathbb{Q}$ by letting $q(a,b) = \frac{a}{b}$ for all $a, b \in \mathbb{Z}$ with $b \neq 0$. By Example 6.1.2 and Exercise 6.1.4, the set $\mathbb{Z} \times (\mathbb{Z} \setminus \{0\})$ is countable. The function q is surjective by definition of $\mathbb{Q}$. By Exercise 6.1.7, it follows that $\mathbb{Q}$ is countable. $\square$

Exercise 6.1.9
Let X be a countable set. Prove that $\binom{X}{k}$ is countable for each $k \in \mathbb{N}$. $\triangleleft$

Theorem 6.1.10
A countable union of countable sets is countable. More precisely, let $\{X_n \mid n \in \mathbb{N}\}$ be a family of countable sets. Then the set X defined by

$$X = \bigcup_{n \in \mathbb{N}} X_n$$

is countable.

Proof
We may assume that the sets X_n are all inhabited, since the empty set does not contribute to the union.

For each $n \in \mathbb{N}$ there is a surjection $f_n : \mathbb{N} \to X_n$. Define $f : \mathbb{N} \times \mathbb{N} \to X$ by $f(m,n) = f_m(n)$ for all $m, n \in \mathbb{N}$. Then f is surjective: if $x \in X$ then $x \in X_m$ for some $m \in \mathbb{N}$. Since f_m is surjective, it follows that $x = f_m(n)$ for some $n \in \mathbb{N}$. But then $x = f(m,n)$. Since $\mathbb{N} \times \mathbb{N}$ is countable, it follows from Exercise 6.1.7 that X is countable. $\square$

Example 6.1.11
Let X be a countable set. The set of all finite subsets of X is countable. Indeed, the set of all finite subsets of X is equal to $\bigcup_{k \in \mathbb{N}} \binom{X}{k}$, which is a union of countably many countable sets by Exercise 6.1.9, so is countable by Theorem 6.1.10. $\triangleleft$

We can also use some clever trickery to prove that certain sets are *uncountable*. The proof of the following theorem is known as **Cantor's diagonal argument**.

Theorem 6.1.12
The set $\{0,1\}^{\mathbb{N}}$ is uncountable.

Proof

Let $f : \mathbb{N} \to \{0,1\}^{\mathbb{N}}$ be a function. We will prove that f is not surjective by constructing a sequence which is not contained in the image of $\mathbb{N}$ under f.

Define an element $b \in \{0,1\}^{\mathbb{N}}$, i.e. a function $b : \mathbb{N} \to \{0,1\}$, by

$$b(n) = 1 - f(n)(n)$$

Then $b(n) \ne f(n)(n)$ for all $n \in \mathbb{N}$. If $b = f(m)$ for some m, then by definition of function equality we must have $b(m) = f(m)(m)$; but we just saw that this is necessarily false. Hence $b \notin f[\mathbb{N}]$, so f is not surjective.

Hence there does not exist a surjective function $\mathbb{N} \to \{0,1\}^{\mathbb{N}}$. By Theorem 6.1.6, the set $\{0,1\}^{\mathbb{N}}$ is uncountable. $\square$

This result can be used to show that the set $\mathbb{R}$ of all real numbers is uncountable, though this relies on features of the real numbers that we have not developed so far in this course.

Exercise 6.1.13

Let X be a set. Prove that $\mathscr{P}(X)$ is either finite or uncountable. $\triangleleft$

Section 6.2
Cardinal arithmetic

Relative cardinality

To do: Relate to size for finite sets

Definition 6.2.1
Let X and Y be sets. We say X and Y are **equinumerous**, and write $|X| = |Y|$ or $X \cong Y$ (LaTeX code: `\cong`), if there exists a bijection $X \to Y$. Write $|X| \leqslant |Y|$ if there is an injection $X \to Y$. The notation $|X|$ denotes the **cardinality** of X.

To do: Examples, etc

To do:

Exercise 6.2.2
Prove that $|\mathbb{R}| = |(0,1)| = |[0,1)| = |(0,1]| = |[0,1]|$.   ◁

To do:

Theorem 6.2.3 (Cantor's theorem)
Let X be a set. Then $|X| < |\mathscr{P}(X)|$.

Proof
The function $x \mapsto \{x\}$ evidently defines an injection $X \to \mathscr{P}(X)$, so $|X| \leqslant |\mathscr{P}(X)|$. The fact that $|X| \neq |\mathscr{P}(X)|$ is then immediate from Exercise 2.3.15.   □

To do:

Lemma 6.2.4
Let X be a set and $h : X \to X$ be an injection. The relation $\preccurlyeq$ on X defined for $a, b \in X$ by
$$a \preccurlyeq b \quad \Leftrightarrow \quad h^n(a) = b \text{ for some } n \in \mathbb{N}$$
is a partial order relation (Definition 5.2.1), where $h^0 = \mathrm{id}_X$ (Definition 2.2.13) and where $h^n = \underbrace{h \circ h \circ \cdots \circ h}_{n \text{ copies}}$ for all $n > 0$.

Proof
We need to prove that $\preccurlyeq$ is reflexive, antisymmetric and transitive.

- $\preccurlyeq$ is reflexive. To see this, let $a \in X$. Then $a = h^0(a)$, and so $a \preccurlyeq a$.

- $\preccurlyeq$ is antisymmetric. To see this, let $a, b \in X$ and suppose $a \preccurlyeq b$ and $b \preccurlyeq a$. Then there exist $m, n \in \mathbb{N}$ such that $b = h^m(a)$ and $a = h^n(b)$, and so

$$h^m(a) = h^n(b)$$

Using the well-ordering principle (Theorem 3.1.47), take $m \in \mathbb{N}$ to be least such that $h^m(a) = h^k(b)$ for some $k \in \mathbb{N}$. We show that $m = 0$.

So suppose $m > 0$. Note that $n \geqslant m > 0$ by minimality of m. Since h is injective, we have

$$h(h^{m-1}(a)) = h^m(a) = h^n(b) = h(h^{n-1}(b))$$

and so $h^{m-1}(a) = h^{n-1}(b)$. But then this contradicts minimality of m.

So we must have $m = 0$, so that $b = h^0(a) = a$.

Hence $\preccurlyeq$ is antisymmetric.

- $\preccurlyeq$ is transitive. To see this, let $a, b, c \in X$ and suppose $a \preccurlyeq b$ and $b \preccurlyeq c$. Then there exist $m, n \in \mathbb{N}$ such that $b = h^n(a)$ and $c = h^m(b)$. But then

$$c = h^m(b) = h^m(h^n(a)) = (h^m \circ h^n)(a) = h^{m+n}(a)$$

and so $a \preccurlyeq c$, as required.

So $\preccurlyeq$ is a partial order. $\qquad\square$

Theorem 6.2.5 (Cantor–Schröder–Bernstein theorem)
Let X and Y be sets. If $|X| \leqslant |Y|$ and $|Y| \leqslant |X|$, then $|X| = |Y|$.

Idea of proof
To do: $\qquad\square$

Proof
Fix injections $f : X \to Y$ and $g : Y \to X$; we will construct a bijection $h : X \to Y$.

Define a relation $\preccurlyeq$ on X by

$$a \preccurlyeq b \quad \Leftrightarrow \quad b = (g \circ f)^n(a) \text{ for some } n \in \mathbb{N}$$

Note that $\preccurlyeq$ is a partial order relation by Lemma 6.2.4, since $g \circ f$ is injective.

Now given $a \in X$, we define $h(a)$ according to which of the following scenarios holds.

- **Scenario 1.** There is no $\preccurlyeq$-minimal $m \in X$ such that $m \preccurlyeq a$. In this case, let $h(a) = f(a)$.

- **Scenario 2.** There is some $\preccurlyeq$-minimal $m \in X$ such that $m \preccurlyeq a$ and $m \neq g(y)$ for any $y \in Y$. In this case, let $h(a) = f(a)$.

- **Scenario 3.** There is some $\preccurlyeq$-minimal $m \in X$ such that $m \preccurlyeq a$ and $m = g(y)$ for some $y \in Y$. In this case, we must have $a = g(c)$ for some $c \in Y$—otherwise we'd be in Scenario 2 with $m = a$—and the element $c \in Y$ for which $a = g(c)$ is unique since g is injective. So let $h(a) = c$ for this uniquely determined $c \in Y$.

Note also that if $a, b \in X$ with $a \preccurlyeq b$, then a and b are in the same scenario.

It remains to prove that h is a bijection.

- h is injective. To see this, let $a, b \in X$ and assume that $h(a) = h(b)$.

  If a and b both fall in Scenario 1 or 2, then $h(a) = f(a)$ and $h(b) = f(b)$, so that $a = b$ since f is injective. Likewise, if a and b both fall in Scenario 3, then $a = g(c)$ and $b = g(d)$ for some $c, d \in Y$; but then $c = h(a) = h(b) = d$, and so $a = g(c) = g(d) = b$.

  The only other possibility is that a falls in Scenario 1 or 2, and b falls in Scenario 3; or vice versa. We prove that this is impossible. Without loss of generality, assume that a falls in Scenario 1 or 2 and b falls in Scenario 3—otherwise swap the roles of a and b in what follows. Then $h(a) = f(a)$, and $h(b)$ is the unique element of Y such that $g(h(b)) = b$. Therefore

  $$g(f(a)) = g(h(a)) = g(h(b)) = b$$

  so that $a \preccurlyeq b$. But then a and b are in the same scenario—this contradicts the assumption that a and b are in different scenarios.

  Thus in the only possible cases, we have proved that $a = b$, so that h is injective.

- h is surjective. To see this, let $c \in Y$ and define $a = g(c)$.

  ◇ If a falls in Scenario 1, then $c = f(b)$ for some $b \in X$, and b also falls in Scenario 1, so that $h(b) = f(b) = c$.

  ◇ If a falls in Scenario 2, then a is not $\preccurlyeq$-minimal—if it were, then we'd be in Scenario 3 since $a = g(c)$. So let $b \in X$ be such that $f(b) = c$. Then $h(b) = f(b) = c$.

  ◇ If a falls in Scenario 3, then since c is the unique element of Y with $a = g(c)$, we have $h(a) = c$.

  In each case, we have found some $x \in X$ such that $h(x) = c$. So h is surjective.

Since h is a bijection, we have $|X| = |Y|$, as required.   □

To do:

Absolute cardinality

To do: Introduce, return to notion of universe, refer to appendix

Exercise 6.2.6
Prove that the relation $\cong$ (Definition 6.2.1) is an equivalence relation on the universe $\mathscr{U}$. $\lhd$

Definition 6.2.7
A **cardinal number** is an element of the set $\mathsf{Card} = \mathscr{U}/{\cong}$ (LaTeX code: `\mathsf{Card}`). Given a set X, the **cardinality** of X is the cardinal number $|X| = [X]_{\cong} \in \mathsf{Card}$.

We will identify the cardinal number $|[n]| \in \mathsf{Card}$ with the natural number $n \in \mathbb{N}$, so that we can view $\mathbb{N}$ as a subset of Card. Thus *cardinality* generalises the notion of *size* defined in Definition 3.2.9. This is a dangerous thing to do, though: we will soon be defining arithmetic operations for cardinal numbers, so we must be careful that the operations we define generalise those for natural numbers.

Cardinal numbers will usually be denoted by lower-case Greek letters $\kappa, \lambda, \mu, \ldots$ (see Appendix A.1).

To do:

Definition 6.2.8
The cardinal number $|\mathbb{N}|$ is called **aleph naught** and is written $\aleph_0$ (LaTeX code: `\aleph_0`).

The symbol $\aleph$ is the Hebrew letter *aleph*. It is the first in a hierarchy of cardinal numbers $\aleph_0, \aleph_1, \aleph_2, \ldots$.

To do:

Definition 6.2.9
The **cardinality of the continuum** is the cardinal number $\mathfrak{c}$ defined by $\mathfrak{c} = |\mathbb{R}|$ (LaTeX code: `\mathfrak{c}`).

To do:

Arithmetic with cardinal numbers

Definition 6.2.10

Let κ and λ be cardinal numbers. The (**cardinal**) **sum** $\kappa + \lambda$ is defined by

$$\kappa + \lambda = |X \cup Y|$$

where X and Y are disjoint sets with $|X| = \kappa$ and $|Y| = \lambda$.

Theorem[AC] **6.2.1**

Let κ and λ be cardinal numbers. If $\kappa, \lambda \geqslant \aleph_0$, then $\kappa + \lambda = \max\{\kappa, \lambda\}$.

Proof
$\square$

Definition 6.2.11

Let κ and λ be cardinal numbers. The (**cardinal**) **product** $\kappa \cdot \lambda$ (LaTeX code: \cdot) of κ and λ is defined by

$$\kappa \cdot \lambda = |X \cdot Y|$$

where X and Y are sets with $|X| = \kappa$ and $|Y| = \lambda$.

Theorem[AC] **6.2.2**

Let κ and λ be cardinal numbers. If $\kappa, \lambda \geqslant \aleph_0$, then $\kappa \cdot \lambda = \max\{\kappa, \lambda\}$.

Proof
$\square$

> **Definition 6.2.12**
> Let κ and λ be cardinal numbers. The (**cardinal**) **exponential** λ^{κ} is defined by
> $$\lambda^{\kappa} = |Y^X|$$
> where X and Y are sets with $|X| = \kappa$ and $|Y| = \lambda$.

To do:

Example 6.2.13
We have $|\mathscr{P}(\mathbb{N})| = 2^{\aleph_0}$. Indeed, there is a bijection $\mathscr{P}(\mathbb{N}) \to \{0,1\}^{\mathbb{N}}$ defined by $U \mapsto i_U$, where $i_U : \mathbb{N} \to \{0,1\}$ is defined by
$$i_U(n) = \begin{cases} 0 & \text{if } n \notin U \\ 1 & \text{if } n \in U \end{cases}$$

More generally, $|\mathscr{P}(X)| = 2^{|X|}$ for all sets X. ◁

In light of Example 6.2.13, we can interpret Cantor's theorem (Theorem 6.2.3) as saying that $\kappa < 2^{\kappa}$ for all cardinal numbers κ.

Exercise 6.2.14
Prove that
$$\mu^{\lambda+\kappa} = \mu^{\lambda} \cdot \mu^{\kappa} \quad \text{and} \quad (\mu \cdot \lambda)^{\kappa} = \mu^{\kappa} \cdot \lambda^{\kappa}$$
for all cardinal numbers κ, λ, μ. ◁

Exercise 6.2.15
Prove that for all cardinal numbers μ, ν, κ, if $\mu \leqslant \nu$, then $\mu^{\kappa} \leqslant \nu^{\kappa}$. ◁

To do:

Example 6.2.16
We prove that $\kappa^{\kappa} = 2^{\kappa}$ for all $\kappa \geqslant \aleph_0$. Indeed:

- $\kappa < 2^{\kappa}$ by Cantor's theorem (Theorem 6.2.3), so that by Exercise 6.2.15 and Theorem 6.2.2 we have
$$\kappa^{\kappa} \leqslant (2^{\kappa})^{\kappa} = 2^{\kappa \cdot \kappa} = 2^{\max\{\kappa,\kappa\}} = 2^{\kappa}$$

- Since $2 < \aleph_0 \leqslant \kappa$, we have $2^{\kappa} \leqslant \kappa^{\kappa}$ by Exercise 6.2.15.

By the Cantor–Schröder–Bernstein theorem (Theorem 6.2.5), it follows that $\kappa^{\kappa} = 2^{\kappa}$. ◁

To do:

Theorem 6.2.17

$\mathfrak{c} = 2^{\aleph_0}$

Proof

We have $|[0,1)| = \mathfrak{c}$ by Exercise 6.2.2 and $|\mathscr{P}(\mathbb{N})| = 2^{\aleph_0}$ by Example 6.2.13, so by the Cantor–Schröder–Bernstein theorem (Theorem 6.2.5), it suffices to find injections $f : [0,1) \to \mathscr{P}(\mathbb{N})$ and $g : \mathscr{P}(\mathbb{N}) \to \mathbb{R}$.

Define $f : [0,1) \to \mathscr{P}(\mathbb{N})$ as follows. Given a real number $x \in [0,1)$ let $(d_n)_{n \geqslant 1}$ be the (unique) sequence of 0s and 1s defined by:

(i) $x = 0.d_1 d_2 d_3 \cdots = \sum\limits_{n=1}^{\infty} d_n \cdot 2^{-n}$; and

(ii) For all $n \geqslant 1$, there exists $r \geqslant n$ such that $d_r = 0$.

This sequence is uniquely defined by (**To do:** Prove this in Chapter 7), so that f is well-defined. Define

$$f(x) = \{n \in \mathbb{N} \mid d_n = 1\}$$

To do: Injectivity of f

Define $g : \mathscr{P}(\mathbb{N}) \to \mathbb{R}$ by

$$g(U) = \sum_{n \in U} 3^{-u}$$

To do: Injectivity of g

Since f and g are injective, we have **To do:** Finish   □

Section 6.3
Ordinal numbers and the axiom of choice

To do:

Axiom 6.3.1 (Axiom of choice)
For any family of inhabited sets $\{X_i \mid i \in I\}$, there is a function $f : I \to \bigcup_{i \in I} X_i$ such that $f(i) \in X_i$ for each $i \in I$. The function f is called a **choice function** for $\{X_i \mid i \in I\}$.

To do:

Well-ordered sets

Definition 6.3.2
Let X be a set. A **well-order** on X is a well-founded total order relation $\preccurlyeq$.

Theorem 6.3.3 (Well-ordering principle)

To do:

Ordinal numbers

To do:

Ordinal arithmetic

To do:

Cardinal numbers as ordinal numbers

To do:

Definition 6.3.4

Chapter 6 exercises

Under construction!

The end-of-chapter exercise sections are new and in an incomplete state.

The real numbers

Inequalities and means

We first encountered the real numbers in Chapter 0, when the real numbers were introduced using a vague (but intuitive) notion of an *infinite number line* (Definition 0.26):

This section will scrutinise the set of real numbers in its capacity as a *complete ordered field*. Decomposing what this means:

- A *field* is a set with a notion of 'zero' and 'one', in which it makes sense to talk about addition, subtraction, multiplication, and division by everything except zero. Examples are $\mathbb{Q}$, $\mathbb{R}$, and $\mathbb{Z}/p\mathbb{Z}$ when p is a prime number (but not when p is composite). However, $\mathbb{Z}$ is not a field, since we can't freely divide by nonzero elements—for example, $1 \in \mathbb{Z}$ and $2 \in \mathbb{Z}$, but no integer n satisfies $2n = 1$.

- An *ordered field* is a field which is equipped with a well-behaved notion of order. Both $\mathbb{Q}$ and $\mathbb{R}$ are ordered fields, but $\mathbb{Z}/p\mathbb{Z}$ is not. We'll see why soon.

- A *complete ordered field* is an ordered field in which every set with an upper bound has a *least* upper bound. As we will see, $\mathbb{Q}$ is not a complete ordered field, but $\mathbb{R}$ is.

This is made (extremely) precise in Section B.2.

Magnitude and scalar product

In this part of the section, we home in on sets of the form $\mathbb{R}^n$, for $n \in \mathbb{N}$. Elements of $\mathbb{R}^n$ are sequences of the form $(x_1, x_2, \ldots, x_n)$, where each $x_i \in \mathbb{R}$. With our interpretation of the reals $\mathbb{R}$ as a *line*, we can interpret a sequence $(x_1, x_2, \ldots, x_n)$ as a point in *n-dimensional space*:

- 0-dimensional space is a single point. The set $\mathbb{R}^0$ has one element, namely the empty sequence $()$, so this makes sense.

- 1-dimensional space is a line. This matches our intuition that $\mathbb{R} = \mathbb{R}^1$ forms a line.

- 2-dimensional space is a *plane*. The elements of $\mathbb{R}^2$ are pairs (x, y), where x and y are both real numbers. We can interpret the pair (x, y) as *coordinates* for a point which is situated x units to the right of $(0,0)$ and y units above $(0,0)$ (where negative values of x or y reverse this direction)—see Figure 7.1.

Figure 7.1: Some points in $\mathbb{R}^2$

With this intuition in mind, we set up the following notation.

Notation 7.1.1

Let $n \in \mathbb{N}$. Elements of $\mathbb{R}^n$ will be denoted $\vec{x}, \vec{y}, \vec{z}, \ldots$ (LaTeX code: `\vec`) and called (*n*-**dimensional**) **vectors**. Given a vector $\vec{x} \in \mathbb{R}^n$, we write x_i for the i^{th} **component** of $\vec{x}$, so that

$$\vec{x} = (x_1, x_2, \ldots, x_n)$$

The element $(0, 0, \ldots, 0) \in \mathbb{R}^n$ is called the **origin** or **zero vector** of $\mathbb{R}^n$, and is denoted by $\vec{0}$.

Moreover, if $\vec{x}, \vec{y} \in \mathbb{R}^n$ and $a \in \mathbb{R}$ we write

$$\vec{x} + \vec{y} = (x_1 + y_1, x_2 + y_2, \ldots, x_n + y_n) \quad \text{and} \quad a\vec{x} = (ax_1, ax_2, \ldots, ax_n)$$

Example 7.1.2

For all $\vec{x} \in \mathbb{R}^n$, we have

$$\vec{x} + \vec{0} = \vec{x} \quad \text{and} \quad 1\vec{x} = \vec{x}$$

◁

Definition 7.1.3

Let $\vec{x} \in \mathbb{R}^n$. The **magnitude** of $\vec{x}$ is the real number $\|\vec{x}\|$ (LaTeX code: `\lVert \vec x \rVert`) defined by

$$\|\vec{x}\| = \sqrt{\sum_{i=1}^{n} x_i^2} = \sqrt{x_1^2 + x_2^2 + \cdots + x_n^2}$$

Given vectors $\vec{x}, \vec{y} \in \mathbb{R}^n$, the **distance** from $\vec{x}$ to $\vec{y}$ is defined to be $\|\vec{y} - \vec{x}\|$. Thus the magnitude of a vector can be thought of as the distance from that vector to the origin.

Example 7.1.4

In $\mathbb{R}^2$, Definition 7.1.3 says that

$$\|(x, y)\| = \sqrt{x^2 + y^2}$$

This matches the intuition obtained from the Pythagorean theorem on the sides of right-hand triangles. For example, consider the triangle with vertices $(0,0)$, $(4,0)$ and $(4,3)$:

The hypotenuse of the triangle has magnitude

$$\|(4, 3)\| = \sqrt{4^2 + 3^2} = \sqrt{25} = 5$$

◁

Exercise 7.1.5

Let $\vec{x}, \vec{y} \in \mathbb{R}^n$. Prove that $\|\vec{x} - \vec{y}\| = \|\vec{y} - \vec{x}\|$. That is, the distance from $\vec{x}$ to $\vec{y}$ is equal to the distance from $\vec{y}$ to $\vec{x}$. ◁

Exercise 7.1.6

Prove that if $x \in \mathbb{R}$ then the magnitude $\|(x)\|$ is equal to the absolute value $|x|$. ◁

Exercise 7.1.7

Let $\vec{x} \in \mathbb{R}^n$. Prove that $\|\vec{x}\| = 0$ if and only if $\vec{x} = \vec{0}$. ◁

The triangle inequality and the Cauchy–Schwarz inequality

The first, and simplest, inequality that we investigate is the (one-dimensional version of the) *triangle inequality* (Theorem 7.1.9). It is so named because of a generalisation to higher dimensions (Theorem 7.1.19), which can be interpreted geometrically as saying that the sum of two side lengths of a triangle is greater than or equal to the third side length.

The triangle inequality is used very frequently in mathematical proofs—you will encounter it repeatedly in this chapter—yet its proof is surprisingly simple.

Before we can prove the triangle inequality, we need the following fact about square roots of real numbers.

Lemma 7.1.8
Let $x, y \in \mathbb{R}$. If $0 \leqslant x \leqslant y$, then $\sqrt{x} \leqslant \sqrt{y}$.

Proof
Suppose $0 \leqslant x \leqslant y$. Note that, by definition of the square root symbol, we have $\sqrt{x} \geqslant 0$ and $\sqrt{y} \geqslant 0$.

Suppose $\sqrt{x} > \sqrt{y}$. By two applications of Theorem B.2.30(d), we have

$$y = \sqrt{y} \cdot \sqrt{y} < \sqrt{x} \cdot \sqrt{y} < \sqrt{x} \cdot \sqrt{x} = x$$

so that $y < x$. But this contradicts the assumption that $x \leqslant y$. Hence $\sqrt{x} \leqslant \sqrt{y}$, as required. □

Theorem 7.1.9 (Triangle inequality in one dimension)
Let $x, y \in \mathbb{R}$. Then $|x+y| \leqslant |x| + |y|$. Moreover, $|x+y| = |x| + |y|$ if and only if x and y have the same sign.

Proof
Note first that $xy \leqslant |xy|$; indeed, xy and $|xy|$ are equal if xy is non-negative, and otherwise we have $xy < 0 < |xy|$. Also $x^2 = |x|^2$ and $y^2 = |y|^2$. Hence

$$(x+y)^2 = x^2 + 2xy + y^2 \leqslant |x|^2 + 2|xy| + |y|^2 = (|x| + |y|)^2$$

Taking (nonnegative) square roots yields

$$|x+y| \leqslant ||x| + |y||$$

by Lemma 7.1.8. But $|x| + |y| \geqslant 0$, so $||x| + |y|| = |x| + |y|$. This completes the first part of the proof.

Equality holds in the above if and only if $xy = |xy|$, which occurs if and only if $xy \geqslant 0$. But this is true if and only if x and y are both non-negative or both non-positive—that is, they have the same sign. □

Example 7.1.10

Let $x, y \in \mathbb{R}$. We prove that

$$\frac{|x+y|}{1+|x+y|} \leqslant \frac{|x|}{1+|x|} + \frac{|y|}{1+|y|}$$

First note that, if $0 \leqslant s \leqslant t$, then

$$\frac{s}{1+s} \leqslant \frac{t}{1+t}$$

To see this, note that

$$
\begin{aligned}
s \leqslant t \Rightarrow 1+s \leqslant 1+t & \qquad \text{rearranging} \\
\Rightarrow \frac{1}{1+t} \leqslant \frac{1}{1+s} & \qquad \text{since } 1+s, 1+t > 0 \\
\Rightarrow 1 - \frac{1}{1+s} \leqslant 1 - \frac{1}{1+t} & \qquad \text{rearranging} \\
\Rightarrow \frac{s}{1+s} \leqslant \frac{t}{1+t} & \qquad \text{rearranging}
\end{aligned}
$$

Now letting $s = |x+y|$ and $t = |x| + |y|$, we have $s \leqslant t$ by the triangle inequality, and hence

$$\frac{|x+y|}{1+|x+y|} \leqslant \frac{|x|}{1+|x|+|y|} + \frac{|y|}{1+|x|+|y|} \leqslant \frac{|x|}{1+|x|} + \frac{|y|}{1+|y|}$$

as required. ◁

Exercise 7.1.11

Let $n \in \mathbb{N}$ and let $x_i \in \mathbb{R}$ for each $i \in [n]$. Prove that

$$\left| \sum_{i=1}^{n} x_i \right| \leqslant \sum_{i=1}^{n} |x_i|$$

with equality if and only if the numbers x_i are either all non-positive or all non-negative. ◁

Exercise 7.1.12

Let $x, y \in \mathbb{R}$. Prove that
$$||x| - |y|| \leqslant |x - y|$$

◁

We will generalise the triangle inequality to arbitrary dimensions in Theorem 7.1.19. Our proof will go via the *Cauchy–Schwarz inequality* (Theorem 7.1.16). To motivate the Cauchy–Schwarz inequality, we introduce another geometric notion called the *scalar product* of two vectors.

<div style="border:1px solid #c080a0; padding:10px;">

Definition 7.1.13

Let $\vec{x}, \vec{y} \in \mathbb{R}^n$. The **scalar product** (or **dot product**) of $\vec{x}$ with $\vec{y}$ is the real number $\vec{x} \cdot \vec{y}$ (LATEX code: `\cdot`) defined by

$$\vec{x} \cdot \vec{y} = \sum_{i=1}^{n} x_i y_i = x_1 y_1 + x_2 y_2 + \cdots + x_n y_n$$

</div>

Example 7.1.14

Let $\vec{x} \in \mathbb{R}^n$. Then $\vec{x} \cdot \vec{x} = \|\vec{x}\|^2$. Indeed

$$\vec{x} \cdot \vec{x} = \sum_{i=1}^{n} x_i^2 = \|\vec{x}\|^2$$

◁

Exercise 7.1.15

Let $\vec{x}, \vec{y}, \vec{z}, \vec{w} \in \mathbb{R}^n$ and let $a, b, c, d \in \mathbb{R}$. Prove that

$$(a\vec{x} + b\vec{y}) \cdot (c\vec{z} + d\vec{w}) = ac(\vec{x} \cdot \vec{z}) + ad(\vec{x} \cdot \vec{w}) + bc(\vec{y} \cdot \vec{z}) + bd(\vec{y} \cdot \vec{w})$$

◁

Intuitively, the scalar product of two vectors $\vec{x}$ and $\vec{y}$ measures the extent to which $\vec{x}$ and $\vec{y}$ fail to be *orthogonal*. In fact, if θ is the acute angle formed between the lines ℓ_1 and ℓ_2, where ℓ_1 passes through $\vec{0}$ and $\vec{x}$ and ℓ_2 passes through $\vec{0}$ and $\vec{y}$, then a formula for the scalar product of $\vec{x}$ and $\vec{y}$ is given by

$$\vec{x} \cdot \vec{y} = \|\vec{x}\| \|\vec{y}\| \cos \theta$$

Evidently, $\vec{x}$ and $\vec{y}$ are orthogonal if and only if $\cos \theta = 0$, in which case $\vec{x} \cdot \vec{y} = 0$ as well. We cannot prove this yet, though, as we have not yet defined trigonometric functions or explored their properties, but hopefully this provides some useful intuition.

The Cauchy–Schwarz inequality provides a useful comparison of the size of a scalar product of two vectors with the magnitudes of the vectors.

Theorem 7.1.16 (Cauchy–Schwarz inequality)
Let $n \in \mathbb{N}$ and let $x_i, y_i \in \mathbb{R}$ for each $i \in [n]$. Then

$$|\vec{x} \cdot \vec{y}| \leqslant \|\vec{x}\| \|\vec{y}\|$$

with equality if and only if $a\vec{x} = b\vec{y}$ for some $a, b \in \mathbb{R}$ which are not both zero.

Proof
If $\vec{y} = \vec{0}$, then this is trivial: both sides of the equation are equal to zero! So assume that $\vec{y} \neq \vec{0}$. In particular, by Exercise 7.1.7, we have $\|\vec{y}\| > 0$.

Define $k = \dfrac{\vec{x} \cdot \vec{y}}{\|\vec{y}\|^2}$. Then

$$
\begin{aligned}
0 &\leqslant \|\vec{x} - k\vec{y}\|^2 && \text{since squares are nonnegative} \\
&= (\vec{x} - k\vec{y}) \cdot (\vec{x} - k\vec{y}) && \text{by Example 7.1.14} \\
&= (\vec{x} \cdot \vec{x}) - 2k(\vec{x} \cdot \vec{y}) + k^2(\vec{y} \cdot \vec{y}) && \text{by Exercise 7.1.15} \\
&= \|\vec{x}\|^2 - \frac{(\vec{x} \cdot \vec{y})^2}{\|y\|^2} && \text{by definition of } k
\end{aligned}
$$

Multiplying through by $\|\vec{y}\|^2$, which is non-negative and therefore doesn't change the sign of the inequality, yields

$$0 \leqslant \|\vec{x}\|^2 \|\vec{y}\|^2 - (\vec{x} \cdot \vec{y})^2$$

which is equivalent to what was to be proved.

Evidently, equality holds if and only if $\|\vec{x} - k\vec{y}\| = 0$, which by Exercise 7.1.7 occurs if and only if $\vec{x} - k\vec{y} = 0$. Now:

- If $\vec{x} - k\vec{y} = 0$, then we have

$$
\begin{aligned}
&\vec{x} - k\vec{y} = 0 \\
\Leftrightarrow\ &\vec{x} - \frac{\vec{x} \cdot \vec{y}}{\|\vec{y}\|^2}\vec{y} = 0 && \text{by definition of } k \\
\Leftrightarrow\ &\|\vec{y}\|^2 \vec{x} = (\vec{x} \cdot \vec{y})\vec{y} && \text{rearranging}
\end{aligned}
$$

If $\vec{y} \neq \vec{0}$ then let $a = \|\vec{y}\|^2$ and $b = \vec{x} \cdot \vec{y}$; otherwise, let $a = 0$ and $b = 1$. In both cases, we have $a\vec{x} = b\vec{y}$ and a, b are not both zero.

If $a\vec{x} = b\vec{y}$ for some $a, b \in \mathbb{R}$ not both zero, then either:

◇ $a = 0$ and $b \neq 0$, in which case $\vec{y} = 0$ and we have equality in the Cauchy–Schwarz inequality; or

◇ $a \neq 0$, in which case $\vec{y} = \frac{b}{a}\vec{x}$. Write $c = \frac{b}{a}$. Then

$$
\begin{aligned}
|\vec{x} \cdot \vec{y}| &= |\vec{x} \cdot (c\vec{x})| \\
&= |c(\vec{x} \cdot \vec{x})| && \text{by Exercise 7.1.15} \\
&= |c| \|\vec{x}\|^2 && \text{by Example 7.1.14} \\
&= \|\vec{x}\| \|c\vec{x}\| && \text{rearranging} \\
&= \|\vec{x}\| \|\vec{y}\|
\end{aligned}
$$

In either case, we have equality in the Cauchy–Schwarz inequality.

So equality holds if and only if $a\vec{x} = b\vec{y}$ for some $a, b \in \mathbb{R}$ not both zero. □

Example 7.1.17

Let $a, b, c \in \mathbb{R}$. We'll prove that

$$ab + bc + ca \leqslant a^2 + b^2 + c^2$$

and examine when equality holds.

Letting $\vec{x} = (a, b, c)$ and $\vec{y} = (b, c, a)$ yields

$$\vec{x} \cdot \vec{y} = ab + bc + ca$$

and

$$\|\vec{x}\| = \sqrt{a^2 + b^2 + c^2} = \sqrt{b^2 + c^2 + a^2} = \|\vec{y}\|$$

Hence $\|\vec{x}\| \|\vec{y}\| = a^2 + b^2 + c^2$. By the Cauchy–Schwarz inequality, it follows that

$$\vec{x} \cdot \vec{y} = ab + bc + ca \leqslant a^2 + b^2 + c^2 = \|\vec{x}\| \|\vec{y}\|$$

as required. Equality holds if and only if $k(a, b, c) = \ell(b, c, a)$ for some $k, \ell \in \mathbb{R}$ not both zero. We may assume $k \neq 0$—otherwise, swap the vectors $\vec{x}$ and $\vec{y}$ in what follows. Then, letting $t = \frac{\ell}{k}$, we have

$$
\begin{aligned}
&k(a, b, c) = \ell(b, c, a) \\
\Leftrightarrow\ &(a, b, c) = (tb, tc, ta) && \text{by definition of } t \\
\Leftrightarrow\ &(a, b, c) = (t^2 c, t^2 a, t^2 b) && \text{substituting } a = tb \text{ etc.} \\
\Leftrightarrow\ &(a, b, c) = (t^3 a, t^3 b, t^3 c) && \text{substituting } a = tb \text{ etc. again} \\
\Leftrightarrow\ &\vec{x} = t^3 \vec{x}
\end{aligned}
$$

This occurs if and only if either $(a, b, c) = (0, 0, 0)$, or $t = 1$, in which case

$$(a, b, c) = (tb, tc, ta) = (b, c, a)$$

So equality holds if and only if $a = b = c$. ◁

Exercise 7.1.18

Let $r \in \mathbb{N}$ and let $a_1, a_2, \ldots, a_r \in \mathbb{R}$ be such that $a_1^2 + a_2^2 + \cdots + a_n^2 = 6$. Prove that

$$(a_1 + 2a_2 + \cdots + na_n)^2 \leqslant n(n+1)(2n+1)$$

and determine when equality holds. ◁

We now use the Cauchy–Schwarz inequality to generalise the one-dimensional version of the triangle inequality (Theorem 7.1.9) to arbitrary (finite) dimensions.

Theorem 7.1.19 (Triangle inequality)

Let $\vec{x}, \vec{y} \in \mathbb{R}^n$. Then

$$\|\vec{x} + \vec{y}\| \leqslant \|\vec{x}\| + \|\vec{y}\|$$

with equality if and only if $a\vec{x} = b\vec{y}$ for some real numbers $a, b \geqslant 0$.

Proof

We proceed by calculation:

$$
\begin{aligned}
\|\vec{x} + \vec{y}\|^2 &= (\vec{x} + \vec{y}) \cdot (\vec{x} + \vec{y}) && \text{by Example 7.1.14} \\
&= (\vec{x} \cdot \vec{x}) + 2(\vec{x} \cdot \vec{y}) + (\vec{y} \cdot \vec{y}) && \text{by Exercise 7.1.15} \\
&\leqslant (\vec{x} \cdot \vec{x}) + 2|\vec{x} \cdot \vec{y}| + (\vec{y} \cdot \vec{y}) && \text{since } a \leqslant |a| \text{ for all } a \in \mathbb{R} \\
&\leqslant \|\vec{x}\|^2 + 2\|x\|\|y\| + \|\vec{y}\|^2 && \text{by Example 7.1.14 and Cauchy–Schwarz} \\
&= (\|\vec{x}\| + \|\vec{y}\|)^2 && \text{rearranging}
\end{aligned}
$$

Taking (nonnegative) square roots of both sides yields

$$\|\vec{x} + \vec{y}\| \leqslant \|\vec{x}\| + \|\vec{y}\|$$

by Lemma 7.1.8, as required.

Equality holds if and only if the two '$\leqslant$' symbols in the above derivation are in fact '$=$' symbols.

- The first inequality is equality if and only if $\vec{x} \cdot \vec{y} = |\vec{x} \cdot \vec{y}|$, which holds if and only if $\vec{x} \cdot \vec{y} \geqslant 0$.

- The second inequality is equality if and only if equality holds in the Cauchy–Schwarz inequality. In turn, this occurs if and only if $a\vec{x} = b\vec{y}$ for some $a, b \in \mathbb{R}$. We may, moreover, assume that $a \geqslant 0$—if not, replace a and b by their negatives.

If $a = 0$ then we can take $b = 0$. If $a > 0$, then by Example 7.1.14 and Exercise 7.1.15, we have

$$\vec{x} \cdot \left(\frac{b}{a}\vec{x}\right) = \frac{b}{a}\|\vec{x}\|^2$$

which is non-negative if and only if $b \geqslant 0$, since we are assuming that $a \geqslant 0$.

Thus, equality holds in the triangle inequality if and only if $a\vec{x} = b\vec{y}$ for some $a, b \geqslant 0$. $\qquad\square$

This general version of the triangle inequality has a geometric interpretation in terms of—you guessed it—triangles. Any three points $\vec{a}, \vec{b}, \vec{c} \in \mathbb{R}^n$ form a (potentially flat) triangle:

The side lengths u, v, w are given by the following equations:

$$u = \|\vec{b} - \vec{a}\|, \quad v = \|\vec{c} - \vec{b}\|, \quad w = \|\vec{a} - \vec{c}\|$$

The triangle inequality says tells us that $u \leqslant v + w$. Indeed:

$$
\begin{aligned}
u &= \|\vec{b} - \vec{a}\| && \text{by definition of } u \\
&= \|(\vec{b} - \vec{c}) + (\vec{c} - \vec{a})\| && \text{rearranging} \\
&\leqslant \|\vec{b} - \vec{c}\| + \|\vec{c} - \vec{a}\| && \text{by the triangle inequality} \\
&= \|\vec{c} - \vec{b}\| + \|\vec{a} - \vec{c}\| && \text{by Exercise 7.1.5} \\
&= v + w && \text{by definition of } v \text{ and } w
\end{aligned}
$$

That is, the triangle inequality says that the sum of two side lengths of a triangle is greater than or equal to the third side length. Moreover, it tells us $u = v + w$ precisely when $k(\vec{a} - \vec{c}) = \ell(\vec{c} - \vec{b})$ for some $k, \ell \geqslant 0$. If $k = 0$ then

$$\vec{c} \;=\; \vec{b} \;=\; 0\vec{a} + (1-0)\vec{b}$$

if $k > 0$, then $k + \ell > 0$, so we have

$$\vec{c} = \frac{k}{k+\ell}\vec{a} + \frac{\ell}{k+\ell}\vec{b} = \frac{k}{k+\ell}\vec{a} + \left(1 - \frac{k}{k+\ell}\right)\vec{b}$$

Examining this a bit more closely yields that $u = v + w$ if and only if

$$\vec{c} = t\vec{a} + (1-t)\vec{b}$$

for some $0 \leqslant t \leqslant 1$, which is to say precisely that $\vec{c}$ lies on the line segment between $\vec{a}$ and $\vec{b}$. In other words, equality holds in the triangle inequality only if the three vertices of the triangle are *collinear*, which is to say that the triangle whose vertices are the points $\vec{a}$, $\vec{b}$ and $\vec{c}$, is flat.

Inequalities of means

Our goal now is to explore different kinds of average—specifically, *means*—of finite sets of non-negative real numbers. We will compare the relative sizes of these means with respect to one-another, with emphasis on three particular kinds of mean: the *arithmetic mean* (Definition 7.1.20), the *geometric mean* (Definition 7.1.21) and the *harmonic mean* (Definition 7.1.29). These means in fact assemble into a continuum of means, called *generalised means* (Definition 7.1.37), all of which can be compared with one another.

Definition 7.1.20
Let $n \geqslant 1$. The (**arithmetic**) **mean** of real numbers $x_1, \ldots, x_n$ is

$$\frac{1}{n}\sum_{i=1}^{n} x_i = \frac{x_1 + x_2 + \cdots + x_n}{n}$$

Definition 7.1.21
Let $n \geqslant 1$. The **geometric mean** of non-negative real numbers $x_1, \ldots, x_n$ is

$$\sqrt[n]{\prod_{i=1}^{n} x_i} = \sqrt[n]{x_1 \cdot x_2 \cdot \cdots \cdot x_n}$$

The following theorem is commonly known as the **AM–GM inequality**.

Theorem 7.1.22 (Inequality of arithmetic and geometric means)
Let $n \in \mathbb{N}$ and $x_1, x_2, \ldots, x_n \geqslant 0$. Then

$$\underbrace{\sqrt[n]{x_1 \cdots x_n}}_{\text{geometric mean}} \leqslant \underbrace{\frac{x_1 + \cdots + x_n}{n}}_{\text{arithmetic mean}}$$

with equality if and only if $x_1 = \cdots = x_n$.

Proof when $n = 2$
We need to show that, if $x, y \in \mathbb{R}$ with $x, y \geqslant 0$, then

$$\sqrt{xy} \leqslant \frac{x+y}{2}$$

with equality if and only if $x = y$.

First note that the square roots of x and y exist since they are non-negative. Now

$$
\begin{aligned}
0 &\leqslant (\sqrt{x} - \sqrt{y})^2 && \text{since squares are nonnegative} \\
&= (\sqrt{x})^2 - 2\sqrt{x}\sqrt{y} + (\sqrt{y})^2 && \text{expanding} \\
&= x - 2\sqrt{xy} + y && \text{rearranging}
\end{aligned}
$$

Rearranging the inequality $0 \leqslant x - 2\sqrt{xy} + y$ yields the desired result.

If $\sqrt{xy} = \frac{x+y}{2}$, then we can rearrange the equation as follows:

$$
\begin{aligned}
\sqrt{xy} = \frac{x+y}{2} &\Rightarrow 2\sqrt{xy} = x + y && \text{multiplying by 2} \\
&\Rightarrow 4xy = x^2 + 2xy + y^2 && \text{squaring both sides} \\
&\Rightarrow x^2 - 2xy + y^2 = 0 && \text{rearranging} \\
&\Rightarrow (x - y)^2 = 0 && \text{factorising} \\
&\Rightarrow x - y = 0 && \text{since } a^2 = 0 \Rightarrow a = 0 \text{ for } a \in \mathbb{R} \\
&\Rightarrow x = y && \text{rearranging}
\end{aligned}
$$

So we have proved both parts of the theorem. $\qquad\square$

Example 7.1.23

We use the AM–GM inequality to prove that the area of a rectangle with fixed perimeter is maximised when the rectangle is a square.

Indeed, fix a perimeter $p > 0$, and let $x, y > 0$ be side lengths of a rectangle with perimeter p—that is, x and y satisfy the equation $2x + 2y = p$. The area a of the rectangle satisfies $a = xy$. By the AM–GM inequality, we have

$$a = xy \leqslant \left(\frac{x+y}{2}\right)^2 = \frac{p^2}{16}$$

Equality holds if and only if $x = y$, in other words, if and only if the rectangle is a square. ◁

Exercise 7.1.24

Let $a, b > 0$ be real numbers. Prove that $\dfrac{a^2 + b^2}{2} \geqslant ab$. ◁

Example 7.1.25

Let $x > 0$. We find the minimum possible value of $x + \frac{9}{x}$. By the AM–GM inequality, we have

$$x + \frac{9}{x} \geqslant 2\sqrt{x \cdot \frac{9}{x}} = 2\sqrt{9} = 6$$

with equality if and only if $x = \frac{9}{x}$, which occurs if and only if $x = 3$. Hence the minimum value of $x + \frac{9}{x}$ when $x > 0$ is 6. ◁

Exercise 7.1.26

Let $x > 0$ and let $n \in \mathbb{N}$. Find the minimum possible value of $\displaystyle\sum_{k=-n}^{n} x^k$. ◁

Exercises 7.1.27 and 7.1.28 complete the proof of the AM–GM inequality (Theorem 7.1.22). Before proceeding with the exercises, let's fix some notation: for each $n \in \mathbb{N}$, let $p_{\text{AM–GM}}(n)$ be the assertion that the AM–GM inequality holds for collections of n numbers; that is, $p_{\text{AM–GM}}(n)$ is the assertion:

For all $x_1, x_2, \ldots, x_n \geqslant 0$, we have

$$\sqrt[n]{\prod_{i=1}^{n} x_i} \leqslant \frac{1}{n} \sum_{i=1}^{n} x_i$$

with equality if and only if $x_1 = x_2 = \cdots = x_n$.

Note that we already proved $p_{\text{AM–GM}}(2)$.

Exercise 7.1.27

Let $r \in \mathbb{N}$ and let $x_1, x_2, \ldots, x_{2r} \in \mathbb{R}$. Write

$$a = \frac{1}{r} \sum_{i=1}^{r} x_i \quad \text{and} \quad g = \sqrt[r]{\prod_{i=1}^{r} x_i}$$

for the arithmetic and geometric means, respectively, of the numbers $x_1, \ldots, x_r$; write

$$a' = \frac{1}{r} \sum_{i=r+1}^{2r} x_i \quad \text{and} \quad g' = \sqrt[r]{\prod_{i=r+1}^{2r} x_i}$$

for the arithmetic and geometric means, respectively, of the numbers $x_{r+1}, \ldots, x_{2r}$; and write

$$A = \frac{1}{2r} \sum_{i=1}^{2r} x_i \quad \text{and} \quad G = \sqrt[2r]{\prod_{i=1}^{2r} x_i}$$

for the arithmetic and geometric means, respectively, of all the numbers $x_1, \ldots, x_{2r}$.

Prove that

$$A = \frac{a + a'}{2} \quad \text{and} \quad G = \sqrt{gg'}$$

Deduce that, for each $r \in \mathbb{N}$, if $p_{\text{AM--GM}}(r)$ is true then $p_{\text{AM--GM}}(2r)$ is true. Deduce further than $p_{\text{AM--GM}}(2^m)$ is true for all $m \in \mathbb{N}$. ◁

Exercise 7.1.28

Let $r \geqslant 2$ and let $x_1, \ldots, x_{r-1} \in \mathbb{N}$. Define

$$x_r = \frac{1}{r-1} \sum_{i=1}^{r-1} x_i$$

Prove that

$$\frac{1}{r} \sum_{i=1}^{r} x_i = x_r$$

Assuming $p_{\text{AM--GM}}(r)$, deduce that

$$x_r^r \geqslant \prod_{i=1}^{r} x_i = \left(\prod_{i=1}^{r-1} x_i \right) \cdot x_r$$

with equality if and only if $x_1 = x_2 = \cdots = x_r$. Deduce that $p_{\text{AM--GM}}(r)$ implies $p_{\text{AM--GM}}(r-1)$. Use Exercise 7.1.27 to deduce further that $p_{\text{AM--GM}}(n)$ is true for all $n \geqslant 1$. ◁

We now introduce another kind of mean, called the *harmonic mean*.

Definition 7.1.29

Let $n \in \mathbb{N}$. The **harmonic mean** of nonzero real numbers $x_1, x_2, \ldots, x_n$ is

$$\left(\frac{1}{n} \sum_{i=1}^{n} x_i^{-1} \right)^{-1} = \frac{n}{\frac{1}{x_1} + \frac{1}{x_2} + \cdots + \frac{1}{x_n}}$$

The harmonic mean of two nonzero real numbers x and y has a simpler expression:

$$\left(\frac{x^{-1} + y^{-1}}{2} \right)^{-1} = \frac{2xy}{x+y}$$

The harmonic mean arises naturally when considering

Example 7.1.30

The cities of York and Leeds are located $d > 0$ miles apart. Two cars drive from York to Leeds, then immediately turn around and drive back. The two cars depart from York at the same time and arrive back in York at the same time.

- The first car drives from York to Leeds at a constant speed of u miles per hour, and drives back to York at a constant speed of v miles per hour.

- The second car drives from York to Leeds and back again at the same constant speed of w miles per hour.

According to the following formula from physics:

$$\text{speed} \times \text{time} = \text{distance}$$

the time spent driving by the first car is $\frac{d}{u} + \frac{d}{v}$, and the time spent driving by the second car is $\frac{2d}{w}$.

Since the cars spend the same amount of time driving, it follows that

$$\frac{2d}{w} = \frac{d}{u} + \frac{d}{v} \qquad \Rightarrow \qquad w = \frac{2uv}{u+v}$$

That is, the second car's speed is the harmonic mean of the two speeds driven by the first car. ◁

As might be expected, we now prove a theorem relating the harmonic means with the other means we have established so far—this theorem is known as the **GM–HM inequality**.

Theorem 7.1.31 (Inequality of geometric and harmonic means)
Let $n \in \mathbb{N}$ and $x_1, x_2, \ldots, x_n > 0$. Then

$$\underbrace{\frac{n}{\frac{1}{x_1} + \frac{1}{x_2} + \cdots + \frac{1}{x_n}}}_{\text{harmonic mean}} \leqslant \underbrace{\sqrt[n]{x_1 x_2 \cdots x_n}}_{\text{geometric mean}}$$

with equality if and only if $x_1 = \cdots = x_n$.

Proof *when $n = 2$*
We need to prove that if $x, y > 0$, then

$$\frac{2}{\frac{1}{x} + \frac{1}{y}} \leqslant \sqrt{xy}$$

This is almost immediate from the AM–GM inequality (Theorem 7.1.22). Indeed, since all numbers in sight are positive, we can take reciprocals to see that this inequality is equivalent to the assertion that

$$\frac{1}{\sqrt{xy}} \leqslant \frac{x^{-1} + y^{-1}}{2}$$

But $\frac{1}{\sqrt{xy}} = \sqrt{x^{-1}y^{-1}}$, so this is immediate from the AM–GM inequality. □

Exercise 7.1.32
Prove the GM–HM inequality for general values of $n \in \mathbb{N}$. ◁

Another example of a mean that has applications in probability theory and statistics is that of the *quadratic mean*.

Definition 7.1.33
Let $n \in \mathbb{N}$. The **quadratic mean** (or **root-mean-square**) of real numbers $x_1, x_2, \ldots, x_n$ is

$$\left(\frac{1}{n} \sum_{i=1}^{n} x_i^2 \right)^{\frac{1}{2}} = \sqrt{\frac{x_1^2 + x_2^2 + \cdots + x_n^2}{n}}$$

The following theorem is, predictably, known as the **QM–AM inequality** (or **RMS–AM inequality**); it is a nice application of the Cauchy–Schwarz inequality.

Theorem 7.1.34 (Inequality of quadratic and arithmetic means)
Let $n > 0$ and $x_1, x_2, \ldots, x_n \geqslant 0$. Then

$$\underbrace{\frac{x_1 + \cdots + x_n}{n}}_{\text{arithmetic mean}} \leqslant \underbrace{\sqrt{\frac{x_1^2 + x_2^2 + \cdots + x_n^2}{n}}}_{\text{quadratic mean}}$$

with equality if and only if $x_1 = \cdots = x_n$.

Proof
Define

$$\vec{x} = (x_1, x_2, \ldots, x_n) \quad \text{and} \quad \vec{y} = (1, 1, \ldots, 1)$$

Then

$$
\begin{aligned}
x_1 + x_2 + \cdots + x_n &= \vec{x} \cdot \vec{y} && \text{by definition of scalar product} \\
&\leqslant \|\vec{x}\| \|\vec{y}\| && \text{by Cauchy–Schwarz} \\
&= \sqrt{x_1^2 + x_2^2 + \cdots + x_n^2} \cdot \sqrt{n} && \text{evaluating the magnitudes}
\end{aligned}
$$

Dividing through by n yields

$$\frac{x_1 + x_2 + \cdots + x_n}{n} \leqslant \sqrt{\frac{x_1^2 + x_2^2 + \cdots + x_n^2}{n}}$$

as required. Equality holds if and only if equality holds in the Cauchy–Schwarz inequality, which occurs if and only if

$$(ax_1, ax_2, \ldots, ax_n) = (b, b, \ldots, b)$$

for some $a, b \in \mathbb{R}$ not both zero. If $a = 0$ then $b = 0$, so we must have $a \neq 0$. Hence equality holds if and only if $x_i = \frac{b}{a}$ for all $i \in [n]$—in particular, if and only if $x_1 = x_2 = \cdots = x_n$. $\qquad\square$

To summarise, what we have proved so far is

harmonic mean	$\overset{(7.1.31)}{\leqslant}$	geometric mean	$\overset{(7.1.22)}{\leqslant}$	arithmetic mean	$\overset{(7.1.34)}{\leqslant}$	quadratic mean

with equality in each case if and only if the real numbers whose means we are taking are all equal.

The following exercise allows us to bookend our chain of inequalities with the minimum and maximum of the collections of numbers.

Exercise 7.1.35

Let $n > 0$ and let $x_1, x_2, \ldots, x_n$ be positive real numbers. Prove that

$$\min\{x_1, x_2, \ldots, x_n\} \leqslant \left(\frac{1}{n} \sum_{i=1}^{n} x_i^{-1}\right)^{-1} \quad \text{and} \quad \max\{x_1, x_2, \ldots, x_n\} \geqslant \left(\frac{1}{n} \sum_{i=1}^{n} x_i^2\right)^{\frac{1}{2}}$$

with equality in each case if and only if $x_1 = x_2 = \cdots = x_n$. $\qquad\triangleleft$

⋆ Generalised means

We conclude this section by mentioning a generalisation of the results we have proved about means. We are not yet ready to prove the results that we mention; they are only here for the sake of interest.

Definition 7.1.36

The **extended real number line** is the (ordered) set $[-\infty, \infty]$, defined by

$$[-\infty, \infty] = \mathbb{R} \cup \{-\infty, \infty\}$$

where $\mathbb{R}$ is the set of real numbers with its usual ordering, and $-\infty, \infty$ are new elements ordered in such a way that $-\infty < x < \infty$ for all $x \in \mathbb{R}$.

Note that the extended real line does *not* form a field—the arithmetic operations are not defined on $-\infty$ or ∞, and we will at no point treat $-\infty$ and ∞ as real numbers; they are merely elements which extend the reals to add a least element and a greatest element.

Definition 7.1.37

Let $p \in [-\infty,\infty]$, let $n \in \mathbb{N}$, and let $x_1,x_2,\ldots,x_n$ be positive real numbers. The **generalised mean with exponent** p (or simply p-**mean**) $x_1,x_2,\ldots,x_n$ is the real number $M_p(x_1,x_2,\ldots,x_n)$ defined by

$$M_p(x_1,x_2,\ldots,x_n) = \left(\frac{1}{n} \sum_{i=1}^{n} x_i^p \right)^{\frac{1}{p}}$$

if $p \notin \{-\infty,0,\infty\}$, and by

$$M_p(x_1,x_2,\ldots,x_n) = \lim_{q \to p} M_q(x_1,x_2,\ldots,x_n)$$

if $p \in \{-\infty,0,\infty\}$, where the notation $\lim_{q \to p}$ refers to the *limit* as q tends to p. (We have not yet defined this notion.)

We can see immediately that the harmonic, arithmetic and quadratic means of a finite set of positive real numbers are the p-means for a suitable value of p: the harmonic mean is the (-1)-mean, the arithmetic mean is the 1-mean, and the quadratic mean is the 2-mean. Furthermore, Proposition 7.1.38 exhibits the *minimum* as the $(-\infty)$-mean, the *geometric mean* as the 0-mean, and the *maximum* as the ∞-mean.

Proposition 7.1.38

Let $n > 0$ and let $x_1,x_2,\ldots,x_n \geqslant 0$. Then

- $M_{-\infty}(x_1,x_2,\ldots,x_n) = \min\{x_1,x_2,\ldots,x_n\}$;

- $M_0(x_1,x_2,\ldots,x_n) = \sqrt[n]{x_1 x_2 \cdots x_n}$; and

- $M_\infty(x_1,x_2,\ldots,x_n) = \min\{x_1,x_2,\ldots,x_n\}$. $\qquad\qquad\square$

All of the inequalities of means we have seen so far will be subsumed by Theorem 7.1.39, which compares the p-mean and q-mean of a set of numbers for all values of $p,q \in [-\infty,\infty]$.

Theorem 7.1.39

Let $n > 0$, let $x_1, x_2, \ldots, x_n \geqslant 0$ and let $p, q \in [-\infty, \infty]$ with $p < q$. Then

$$M_p(x_1, x_2, \ldots, x_n) \leqslant M_q(x_1, x_2, \ldots, x_n)$$

with equality if and only if $x_1 = x_2 = \cdots = x_n$.   □

Theorem 7.1.39 implies each of the following:

- **HM–min inequality** (Exercise 7.1.35): take $p = -\infty$ and $q = -1$;

- **GM–HM inequality** (Theorem 7.1.31): take $p = -1$ and $q = 0$;

- **AM–GM inequality** (Theorem 7.1.22): take $p = 0$ and $q = 1$;

- **QM–AM inequality** (Theorem 7.1.34): take $p = 1$ and $q = 2$;

- **max–QM inequality** (Exercise 7.1.35): take $p = 2$ and $q = \infty$.

Section 7.2
Completeness and convergence

For most of the results that we proved in Section 7.1, it did not matter that we were talking about real numbers. We could just as well have been working with any other ordered field, such as the rational numbers—that is, most of the results in Section 7.1 remain true by replacing $\mathbb{R}$ by $\mathbb{Q}$ (or any other ordered field) throughout.

From here onwards, we isolate the property of $\mathbb{R}$ that separates it from $\mathbb{Q}$— namely, *completeness*. It is completeness that will allow us to define and explore the fundamental concepts of mathematical analysis: sequences, functions, convergence, limits, continuity, differentiability, and so on.

We first need to recall the definition of a *supremum* from Section 5.2.

Definition 7.2.1 (instance of Definition 5.2.10)
Let $A \subseteq \mathbb{R}$. A real number m is an **upper bound** for A if $a \leqslant m$ for all $a \in A$. A **supremum** of A is a *least* upper bound of A; that is, a real number m such that:

(i) m is an upper bound of A—that is, $a \leqslant m$ for all $a \in A$; and

(ii) m is least amongst all upper bounds for A—that is, for all $x \in \mathbb{R}$, if $a \leqslant x$ for all $a \in A$, then $x \leqslant m$.

Example 7.2.2
We prove that 1 is a supremum of the open interval $(0, 1)$.

(i) Let $a \in (0, 1)$. Then $a < 1$, so that 1 is an upper bound of $(0, 1)$.

(ii) Let $x \in \mathbb{R}$ be another upper bound of $(0, 1)$. If $x < 1$, then we have

$$x = \frac{x+x}{2} < \frac{x+1}{2} < \frac{1+1}{2} = 1$$

and so $x < \dfrac{x+1}{2} \in (0, 1)$. This contradicts the assumption that x is an upper bound of $(0, 1)$. It follows that $x \geqslant 1$, as required.

Hence 1 is indeed a supremum of $(0, 1)$. ◁

Exercise 7.2.3
Write down the definitions of *lower bound* and *infimum*, and find the infimum of the open interval $(0, 1)$. ◁

The following proposition provides a convenient way of testing whether a real number is a supremum of a subset.

Proposition 7.2.4

Let $A \subseteq \mathbb{R}$ and suppose $m \in \mathbb{R}$ is an upper bound of A. Then m is a supremum of A if and only if, for all $\varepsilon > 0$, there exists $a \in A$ such that $a > m - \varepsilon$.

Proof

- ($\Rightarrow$). Suppose m is a supremum of A, and let $\varepsilon > 0$. If there is no $a \in A$ such that $a > m - \varepsilon$, then $a \leqslant m - \varepsilon$ for all $a \in A$. But this contradicts the assumption that m is a supremum of a, since $m - \varepsilon$ is an upper bound of A that is less than m. So there exists $a \in A$ with $a > m - \varepsilon$, as required.

- ($\Leftarrow$). Suppose that, for all $\varepsilon > 0$, there exists $a \in A$ with $a > m - \varepsilon$, and let $x \in \mathbb{R}$ be an upper bound of A. In order to prove that m is a supremum of A, we must prove that $m \leqslant x$.

  Suppose $x < m$, and define $\varepsilon = m - x$. Then $\varepsilon > 0$, so there exists $a \in A$ such that
  $$a > m - \varepsilon = m - (m - x) = x$$
  But this contradicts the assumption that x is an upper bound of A. So we must have $m \leqslant x$, as required.

  $\square$

Theorem 7.2.5 (Uniqueness of suprema)

Let A be a subset of $\mathbb{R}$. If m_1 and m_2 are suprema of A, then $m_1 = m_2$.

Proof

Since m_1 is an upper bound of A and m_2 is a supremum of A, we have $m_2 \geqslant m_1$ by Definition 7.2.1(ii). Likewise, since m_2 is an upper bound of A and m_1 is a supremum of A, we have $m_1 \geqslant m_2$ by Definition 7.2.1(ii) again. But this implies that $m_1 = m_2$. $\square$

Definition 7.2.6

Let $A \subseteq \mathbb{R}$. The supremum of A, if it exists is denoted by $\sup(A)$ (LaTeX code: `\mathrm{sup}`); the infimum of A, if it exists, is denoted by $\inf(A)$ (LaTeX code: `\mathrm{inf}`).

Now that we are more familiar with suprema, here is the completeness axiom in its full glory.

Axiom 7.2.7 (Completeness axiom)

Let $A \subseteq \mathbb{R}$ be inhabited. If A has an upper bound, then A has a supremum.

The true power of the completeness axiom will become apparent later in the section when we discuss the existence of limits of sequences of real numbers.

Before we embark on that adventure, we first prove that the rational numbers are *not* complete, by exhibiting a subset of $\mathbb{Q}$ that has no rational supremum.

Proposition 7.2.8

Let $A = \{x \in \mathbb{Q} \mid x^2 < 2\}$. Then A does not have a rational supremum.

A quick proof of Proposition 7.2.8 would be to verify that $\sqrt{2}$, which is irrational, is a supremum of A, and use uniqueness of suprema to deduce that there can be no rational supremum. However, this is cheating. Failure of completeness is an *intrinsic* property—we should be able to prove Proposition 7.2.8 without venturing outside of the realm of rational numbers at all. That is, we cannot use irrational numbers in our proof. This makes the proof significantly longer, but significantly more satisfying.

Proof of Proposition 7.2.8

Towards a contradiction, suppose that A has a supremum q.

First note that $q > 0$. Indeed, $1^2 < 2$, so that $1 \in A$, and so $q \geqslant 1 > 0$.

Next, we prove that $q^2 = 2$. Indeed:

- Assume $q^2 < 2$, so that $2 - q^2 > 0$. For each $n \geqslant 1$, we have

$$\left(q + \frac{1}{n}\right)^2 = q^2 + \frac{2q}{n} + \frac{1}{n^2}$$

Choose n sufficiently large that $\frac{2q}{n} < \frac{2-q^2}{2}$ and $\frac{1}{n^2} < \frac{2-q^2}{2}$. Then by the above, we observe that

$$\left(q + \frac{1}{n}\right)^2 < q^2 + \frac{2-q^2}{2} + \frac{2-q^2}{2} = q^2 + (2 - q^2) = 2$$

and so $q + \frac{1}{n} \in A$. But $q + \frac{1}{n} > q$, so this contradicts the assumption that q is an upper bound of A.

- Assume $q^2 > 2$, so that $q^2 - 2 > 0$. For each $n \geqslant 1$, we have

$$\left(q - \frac{1}{n}\right)^2 = q^2 - \frac{1}{n}\left(2q - \frac{1}{n}\right)$$

Choose n sufficiently large that $\frac{1}{n} < q \ (< 2q)$ and $\frac{2q}{n} < q^2 - 2$. Then by the above work, we observe that

$$\left(q - \frac{1}{n}\right)^2 > q^2 - \frac{2q}{n} > q^2 - (q^2 - 2) = 2$$

Moreover $q - \frac{1}{n} > 0$ since $\frac{1}{n} < q$.

Suppose that $q - \frac{1}{n}$ is *not* an upper bound for A. Then there is some $x \in A$ with $x > q - \frac{1}{n} > 0$. But then $(q - \frac{1}{n})^2 < x^2 < 2$, contradicting the fact that $\left(q - \frac{1}{n}\right)^2 > 2$.

So $q - \frac{1}{n}$ is an upper bound for A, contradicing the fact that q is a supremum of A.

So we must have $q^2 = 2$. But this is impossible—the proof is identical to that of Proposition 3.1.48, but with all instances of '$\sqrt{2}$' replaced by 'q' in the proof.

So $\{x \in \mathbb{Q} \mid x^2 < 2\}$ has no rational supremum.   □

Sequences of real numbers

The rest of this chapter is dedicated to studying *convergence* of sequences of real numbers. We will use the completeness axiom to find sufficient conditions for a sequence to converge.

Definition 7.2.9

A **sequence of real numbers** is a function $x : \mathbb{N} \to \mathbb{R}$. Given a sequence x, we write x_n instead of $x(n)$ and write $(x_n)_{n \geqslant 0}$, or even just (x_n), instead of $x : \mathbb{N} \to \mathbb{R}$. The values x_n are called the **terms** of the sequence, and the variable n is called the **index** of the term x_n.

Example 7.2.10

Some very basic but very boring examples of sequences are *constant sequences*. For example, the constant sequence with value 0 is

$$(0,0,0,0,0,0,\dots)$$

More generally, for fixed $a \in \mathbb{R}$, the constant sequence with value a is defined by $x_n = a$ for all $n \in \mathbb{N}$.   ◁

Example 7.2.11

Sequences can be defined just like functions. For example, there is a sequence defined by $x_n = 2^n$ for all $n \in \mathbb{N}$. Writing out the first few terms, this sequence is

$$(1,2,4,8,16,\dots)$$

◁

Sometimes it will be convenient to start the indexing of our sequence from numbers other than 0, particularly when an expression involving a variable n isn't

defined when $n = 0$. We'll denote such sequences by $(x_n)_{n \geqslant 1}$ or $(x_n)_{n \geqslant 2}$, and so on.

Example 7.2.12

Let $(z_n)_{n \geqslant 2}$ be the sequence defined by $z_n = \frac{(n+1)(n+2)}{(n-1)n}$ for all $n \geqslant 2$:

$$\left(6, \frac{10}{3}, \frac{5}{2}, \frac{21}{10}, \ldots \right)$$

The indexing of this sequence begins at 2, rather than 0, since when $n = 0$ or $n = 1$, the expression $\frac{(n+1)(n+2)}{(n-1)n}$ is undefined. We could *reindex* the sequence: by letting $z'_n = z_{n+2}$ for all $n \geqslant 0$, we obtain a new sequence $(z'_n)_{n \geqslant 0}$ defined by $z'_n = \frac{(n+3)(n+4)}{(n+1)(n+2)}$ whose indexing starts from 0. Fortunately for us, such matters won't cause any problems—it's just important to make sure that whenever we define a sequence, we make sure the terms make sense for all of the indices. ◁

Convergence of sequences

Of particular interest to us will be sequences whose terms get closer and closer to a fixed real number. This phenomenon is called *convergence*.

Example 7.2.13

Consider the sequence $(y_n)_{n \geqslant 1}$ defined by $y_n = \frac{1}{n}$ for all $n \geqslant 1$:

$$\left(1, \frac{1}{2}, \frac{1}{3}, \frac{1}{4}, \frac{1}{5}, \ldots \right)$$

It is fairly clear that the terms y_n become closer and closer to 0 as n grows; the following diagram is a plot of y_n against n for a few values of n. ◁

Example 7.2.14

Define a sequence $(r_n)_{n \geqslant 0}$ by $r_n = \frac{2n}{n+1}$ for all $n \in \mathbb{N}$. Some of the values of this sequence are illustrated in the following table:

n	r_n	decimal expansion
0	0	0
1	1	1
2	$\frac{4}{3}$	1.333…
3	$\frac{3}{2}$	1.5
⋮	⋮	⋮
10	$\frac{20}{11}$	1.818…
⋮	⋮	⋮
100	$\frac{200}{101}$	1.980…
⋮	⋮	⋮
1000	$\frac{2000}{1001}$	1.998…
⋮	⋮	⋮

As n increases, the values of r_n become closer and closer to 2.   ◁

The precise sense in which the terms of the sequences in Examples 7.2.13 and 7.2.14 'get closer' to 0 and 2, respectively, is called *convergence*, which we will define momentarily in Definition 7.2.15.

First, let's try to work out what the definition *should be* for a sequence (x_n) to converge to a real number a.

A naïve answer might be to say that the sequence is 'eventually equal to a'—that is, after some point in the sequence, all terms are equal to a. Unfortunately, this isn't quite good enough: if it were, then the values $r_n = \frac{2n}{n+1}$ from Example 7.2.14 would be equal to 2 for sufficiently large n. However, if for some $n \in \mathbb{N}$ we have $\frac{2n}{n+1} = 2$, then it follows that $2n = 2(n+1)$; rearranging this gives $1 = 0$, which is a contradiction.

However, this answer isn't too far from giving us what we need. Instead of saying that the terms x_n are eventually *equal* to a, we might want to say that they become *infinitely close* to a, whatever that means.

We can't really make sense of an 'infinitely small positive distance' (e.g. Exercise 1.1.41), so we might instead make sense of 'infinitely close' by saying that the terms x_n eventually become as close to a as we could possibly want them to be. Spelling this out, this means that for any positive distance ε (LaTeX code:

`\varepsilon`) (read: 'epsilon')[a] no matter how small, the terms x_n are eventually within distance ε of a. In summary:

Definition 7.2.15

Let (x_n) be a sequence and let $a \in \mathbb{R}$. We say that (x_n) **converges** to a, and write $(x_n) \to a$ (LaTeX code: `\to`), if the following condition holds:

$$\forall \varepsilon > 0, \exists N \in \mathbb{N}, \forall n \geqslant N, |x_n - a| < \varepsilon$$

The value a is called a **limit** of (x_n). Moreover, we say that a sequence (x_n) **converges** if it has a limit, and diverges otherwise.

Before we move onto some examples, let's quickly digest the definition of the expression $(x_n) \to a$. The following table presents a suggestion of how you might read the expression '$\forall \varepsilon > 0, \exists N \in \mathbb{N}, \forall n \geqslant N, |x_n - a| < \varepsilon$' in English.

Symbols	English
$\forall \varepsilon > 0 \ldots$	For any positive distance ε (no matter how small)...
$\ldots \exists N \in \mathbb{N} \ldots$	... there is a stage in the sequence...
$\ldots \forall n \geqslant N \ldots$	... after which all terms in the sequence...
$\ldots |x_n - a| < \varepsilon.$	... are within distance ε of a.

Thus, a sequence (x_n) converges to a if '*for any positive distance ε (no matter how small), there is a stage in the sequence after which all terms in the sequence are within ε of a*'. After reading this a few times, you should hopefully be content that this definition captures what is meant by saying that the terms in the sequence are eventually as close to a as we could possibly want them to be.

We are now ready to see some examples of convergent (and divergent) sequences. When reading the following proofs, keep in mind the logical structure—that is, the alternating quantifiers $\forall \varepsilon \ldots \exists N \ldots \forall n \ldots$—in the definition of $(x_n) \to a$.

Example 7.2.16
The sequence (y_n) defined by $y_n = \frac{1}{n}$ for all $n \geqslant 1$ converges to 0. To see this, by Definition 7.2.15, we need to prove

$$\forall \varepsilon > 0, \exists N \in \mathbb{N}, \forall n \geqslant N, \left|\frac{1}{n} - 0\right| < \varepsilon$$

So fix $\varepsilon > 0$. Our goal is to find $N \in \mathbb{N}$ such that $\left|\frac{1}{n}\right| < \varepsilon$ for all $n \geqslant N$.

---

[a]The lower case Greek letter *epsilon* (ε) is traditionally used in analysis to denote a positive quantity whose value can be made arbitrarily small. We will encounter this letter frequently in this section and the next when discussing convergence.

Let N be any natural number which is greater than $\frac{1}{\varepsilon}$. Then for all $n \geqslant N$, we have

$$
\begin{aligned}
\left|\frac{1}{n}\right| = \frac{1}{n} && \text{since } \frac{1}{n} > 0 \text{ for all } n \geqslant 1 \\
\leqslant \frac{1}{N} && \text{since } n \geqslant N \\
< \frac{1}{1/\varepsilon} && \text{since } N > \frac{1}{\varepsilon} \\
= \varepsilon
\end{aligned}
$$

Hence $|y_n| < \varepsilon$ for all $n \geqslant N$. Thus we have proved that $(y_n) \to 0$. ◁

Remark 7.2.17
The value of N you need to find in the proof of convergence will usually depend on the parameter ε. (For instance, in Example 7.2.16, we defined N to be some natural number greater than $\frac{1}{\varepsilon}$.) This is to be expected—remember that ε is the distance away from the limit that the terms are allowed to vary after the N^{th} term. If you make this distance smaller, you'll probably have to go further into the sequence before your terms are all close enough to a. In particular, the value of N will usually grow as the value of ε gets smaller. This was the case in Example 7.2.16: note that $\frac{1}{\varepsilon}$ increases as ε decreases. ◁

Example 7.2.18
Let (r_n) be the sequence from Example 7.2.14 defined by $r_n = \frac{2n}{n+1}$ for all $n \in \mathbb{N}$. We'll prove that $(r_n) \to 2$. So fix $\varepsilon > 0$. We need to find $N \in \mathbb{N}$ such that

$$
\left|\frac{2n}{n+1} - 2\right| < \varepsilon \text{ for all } n \geqslant N
$$

To find such a value of n, we'll first do some algebra. Note first that for all $n \in \mathbb{N}$ we have

$$
\left|\frac{2n}{n+1} - 2\right| = \left|\frac{2n - 2(n+1)}{n+1}\right| = \left|\frac{-2}{n+1}\right| = \frac{2}{n+1}
$$

Rearranging the inequality $\frac{2}{n+1} < \varepsilon$ gives $\frac{n+1}{2} > \frac{1}{\varepsilon}$, and hence $n > \frac{2}{\varepsilon} - 1$.

To be clear, what we've shown so far is that a *necessary* condition for $|r_n - 2| < \varepsilon$ to hold is that $n > \frac{2}{\varepsilon} - 1$. This informs us what the desired value of N might look like—we will then verify that the desired inequality holds.

So define $N = \frac{2}{\varepsilon} - 1$. For all $n \geqslant N$, we have

$$
\begin{aligned}
\left| \frac{2n}{n+1} - 2 \right| &= \frac{2}{n+1} && \text{by the above work} \\
&\leqslant \frac{2}{N+1} && \text{since } n \geqslant N \\
&< \frac{2}{\left(\frac{2}{\varepsilon} - 1 \right) + 1} && \text{since } N > \frac{2}{\varepsilon} - 1 \\
&= \frac{2}{2/\varepsilon} && \text{rearranging} \\
&= \varepsilon && \text{rearranging}
\end{aligned}
$$

Thus, as claimed, we have $|r_n - 2| < \varepsilon$ for all $n \geqslant N$. It follows that $(r_n) \to 2$, as required. ◁

Exercise 7.2.19
Let (x_n) be the constant sequence with value $a \in \mathbb{R}$. Prove that $(x_n) \to a$. ◁

Exercise 7.2.20
Prove that the sequence (z_n) defined by $z_n = \frac{n+1}{n+2}$ converges to 1. ◁

The following proposition is a technical tool, which proves that convergence of sequences is unaffected by changing finitely many terms at the beginning of a sequence.

Proposition 7.2.21
Let (x_n) be a sequence and suppose that $(x_n) \to a$. Let (y_n) be another sequence and suppose that there is some $k \in \mathbb{N}$ such that $x_n = y_n$ for all $n \geqslant k$. Prove that $(y_n) \to a$.

Proof
Fix $\varepsilon > 0$. We need to find $N \in \mathbb{N}$ such that $|y_n - a| < \varepsilon$ for all $n \geqslant N$.

Since $(x_n) \to a$, there is some $M \in \mathbb{N}$ such that $|x_n - a| < \varepsilon$ for all $n \geqslant M$. Let N be the greater of M and k. Then for all $n \geqslant N$, we have $y_n = x_n$, since $n \geqslant k$, and hence $|y_n - a| = |x_n - a| < \varepsilon$, since $n \geqslant M$.

Hence $(y_n) \to a$, as required. □

Before we go too much further, let's see some examples of sequences which *diverge*. Recall (Definition 7.2.15) that a sequence (x_n) converges if $(x_n) \to a$ for some $a \in \mathbb{R}$. Spelling this out symbolically, to say '(x_n) converges' is to say the following:

$$ \exists a \in \mathbb{R}, \forall \varepsilon > 0, \exists N \in \mathbb{N}, \forall n \geqslant N, |x_n - a| < \varepsilon $$

We can negate this using the tools of Section 1.3: to say that a sequence (x_n) diverges is to say the following:

$$\forall a \in \mathbb{R}, \exists \varepsilon > 0, \forall N \in \mathbb{N}, \exists n \geqslant N, |x_n - a| \geqslant \varepsilon$$

In more intuitive terms: for all possible candidates for a limit $a \in \mathbb{R}$, there is a positive distance ε such that, no matter how far down the sequence you go (say x_N), you can find a term x_n beyond that point which is at distance $\geqslant \varepsilon$ away from a.

Example 7.2.22

Let (x_n) be the sequence defined by $x_n = (-1)^n$ for all $n \in \mathbb{N}$:

$$(1, -1, 1, -1, 1, -1, \ldots)$$

We'll prove that (x_n) diverges. Fix $a \in \mathbb{R}$. Intuitively, if a is non-negative, then it must be at distance $\geqslant 1$ away from -1, and if a is negative, then it must be at distance $\geqslant 1$ away from 1. We'll now make this precise.

So let $\varepsilon = 1$, and fix $N \in \mathbb{N}$. We need to find $n \geqslant N$ such that $|(-1)^n - a| \geqslant 1$. We'll split into cases based on whether a is non-negative or negative.

• Suppose $a \geqslant 0$. Then $-1 - a \leqslant -1 < 0$, so that we have

$$|-1 - a| = a - (-1) = a + 1 \geqslant 1$$

So let $n = 2N + 1$. Then $n \geqslant N$ and n is odd, so that

$$|x_n - a| = |(-1)^n - a| = |-1 - a| \geqslant 1$$

• Suppose $a < 0$. Then $1 - a > 1 > 0$, so that we have

$$|1 - a| = 1 - a > 1$$

So let $n = 2N$. Then $n \geqslant N$ and n is even, so that

$$|x_n - a| = |(-1)^n - a| = |1 - a| \geqslant 1$$

In both cases, we've found $n \geqslant N$ such that $|x_n - a| \geqslant 1$. It follows that (x_n) diverges. ◁

Example 7.2.22 is an example of a *periodic* sequence—that is, it's a sequence that repeats itself. It is difficult for such sequences to converge since, intuitively speaking, they jump up and down a lot. (In fact, the only way that a period sequence *can* converge is if it is a constant sequence!)

Exercise 7.2.23

Let (y_n) be the sequence defined by $y_n = n$ for all $n \in \mathbb{N}$:

$$(0, 1, 2, 3, \dots)$$

Prove that (y_n) diverges. ◁

Finding limits of sequences can be tricky. Theorem 7.2.25 makes it slightly easier by saying that if a sequence is built up using arithmetic operations—addition, subtraction, multiplication and division—from sequences whose limits you know, then you can simply apply those arithmetic operations to the limits.

In order to prove part of Theorem 7.2.25, however, the following lemma will be useful.

Lemma 7.2.24

Let (x_n) be a sequence of real numbers. If (x_n) converges, then (x_n) is bounded—that is, there is some real number k such that $|x_n| \leqslant k$ for all $n \in \mathbb{N}$.

Proof

Let $a \in \mathbb{R}$ be such that $(x_n) \to a$. Letting $\varepsilon = 1$ in the definition of convergence, it follows that there exists some $N \in \mathbb{N}$ such that $|x_n - a| < 1$ for all $n \geqslant N$. It follows that $-1 < x_n - a < 1$ for all $n \geqslant N$, and hence $-(1-a) < x_n < 1+a$ for all $n \geqslant N$.

Now define

$$k = \max\{|x_0|, |x_1|, \dots, |x_{N-1}|, |1-a|, |1+a|\} + 1$$

For all $n < N$, we have

$$-k < -|x_n| \leqslant x_n \leqslant |x_n| < k$$

so that $|x_n| < k$. For all $n \geqslant N$, we have

$$-k < -|1-a| \leqslant -(1-a) < x_n < 1+a \leqslant |1+a| < k$$

so that $|x_n| < k$.

Hence $|x_n| < k$ for all $n \in \mathbb{N}$, as required. □

Theorem 7.2.25

Let (x_n) and (y_n) be sequences of real numbers, let $a, b \in \mathbb{R}$, and suppose that $(x_n) \to a$ and $(y_n) \to b$. Then

(a) $(x_n + y_n) \to a + b$;

(b) $(x_n - y_n) \to a - b$;

(c) $(x_n y_n) \to ab$; and

(d) $\left(\frac{x_n}{y_n}\right) \to \frac{a}{b}$, so long as $y_n \neq 0$ for all $n \in \mathbb{N}$ and $b \neq 0$.

Proof of (a) and (c)

(a). Fix $\varepsilon > 0$. We need to prove that there is some $N \in \mathbb{N}$ such that $|(x_n + y_n) - (a + b)| < \varepsilon$ for all $n \geqslant N$.

- Since $(x_n) \to a$, there is some $N_1 \in \mathbb{N}$ such that $|x_n - a| < \frac{\varepsilon}{2}$ for all $n \geqslant N_1$;
- Since $(y_n) \to b$, there is some $N_2 \in \mathbb{N}$ such that $|x_n - b| < \frac{\varepsilon}{2}$ for all $n \geqslant N_2$.

Let N be the greatest of N_1 and N_2. Then for all $n \geqslant N$, we have $n \geqslant N_1$ and $n \geqslant N_2$; it follows from the triangle inequality (Theorem 7.1.9), that

$$|(x_n + y_n) - (a + b)| = |(x_n - a) + (y_n - b)| \leqslant |x_n - a| + |y_n - b| < \frac{\varepsilon}{2} + \frac{\varepsilon}{2}$$

as required.

(c). This one is a little harder. Fix $\varepsilon > 0$. Since (x_n) converges, it follows from Lemma 7.2.24 that there is some real number k with $|x_n| < k$ for all $n \in \mathbb{N}$.

- Since $(x_n) \to a$, there is some $N_1 \in \mathbb{N}$ such that $|x_n - a| < \frac{\varepsilon}{2|b|}$ for all $n \geqslant N_1$;
- Since $(y_n) \to b$, there is some $N_2 \in \mathbb{N}$ such that $|x_n - b| < \frac{\varepsilon}{2k}$ for all $n \geqslant N_2$.

Let N be the greatest of N_1 and N_2. Then for all $n \geqslant N$, we have

$$
\begin{aligned}
|x_n y_n - ab| &= |x_n(y_n - b) + b(x_n - a)| && \text{rearranging} \\
&\leqslant |x_n(y_n - b)| + |b(x_n - a)| && \text{by the triangle inequality} \\
&= |x_n||y_n - b| + |b||x_n - a| && \text{rearranging} \\
&< k|y_n - b| + |b||x_n - a| && \text{since } |x_n| < k \text{ for all } n \\
&< k\frac{\varepsilon}{2k} + |b|\frac{\varepsilon}{2|b|} && \text{since } n \geqslant N_1 \text{ and } n \geqslant N_2 \\
&= \varepsilon && \text{rearranging}
\end{aligned}
$$

Hence $(x_n y_n) \to ab$, as required. $\qquad\square$

Exercise 7.2.26

Prove parts (b) and (d) of Theorem 7.2.25. ◁

Theorem 7.2.25 *appears* obvious, but as you can see in the proof, it is more complicated than perhaps expected. It was worth the hard work, though, because we can now compute more complicated limits formed in terms of arithmetic operations by taking the limits of the individual components.

The following example uses Theorem 7.2.25 to prove that $\left(\frac{2n}{n+1}\right) \to 2$ in a much simpler way than we saw in Example 7.2.18.

Example 7.2.27

We provide another proof that the sequence (r_n) of Example 7.2.14, defined by $r_n = \frac{2n}{n+1}$ for all $n \in \mathbb{N}$, converges to 2.

For all $n \geqslant 1$, dividing by the top and bottom gives

$$r_n = \frac{2}{1 + \frac{1}{n}}$$

The constant sequences (2) and (1) converge to 2 and 1, respectively; and by Example 7.2.16, we know that $\left(\frac{1}{n}\right) \to 0$. It follows that

$$(r_n) \to \frac{2}{1+0} = 2$$

as required. ◁

Exercise 7.2.28

Let (x_n) be a sequence of real numbers converging to a real number a, and let $p(x) = a_0 + a_1 x + \cdots + a_d x^d$ be a polynomial function. Prove that $(p(x_n)) \to p(a)$, and that $\left(\frac{1}{p(x_n)}\right) \to \frac{1}{p(a)}$ if $p(a) \neq 0$. ◁

The so-called *squeeze theorem* provides another means of computing limits. It says that if we can eventually 'squeeze' the terms of a sequence (y_n) between terms of two other sequences that converge to the same limit, then we can deduce that (y_n) converges to the same limit.

Theorem 7.2.29 (Squeeze theorem)
Let (x_n), (y_n) and (z_n) be sequences of real numbers such that:

  (i) $(x_n) \to a$ and $(z_n) \to a$; and

  (ii) There is some $k \in \mathbb{N}$ such that $x_n \leqslant y_n \leqslant z_n$ for all $n \geqslant k$.

Then $(y_n) \to a$.

Proof

Fix $\varepsilon > 0$. We need to find $N \in \mathbb{N}$ such that $|y_n - a| < \varepsilon$ for all $n \geqslant N$.

Since $(x_n) \to a$ and $(z_n) \to a$, there exist $N_1, N_2 \in \mathbb{N}$ such that

- $|x_n - a| < \varepsilon$ for all $n \geqslant N_1$;
- $|z_n - a| < \varepsilon$ for all $n \geqslant N_2$.

Let $N = \max\{N_1, N_2, k\}$. Then for all $n \geqslant N$, we have:

- $|x_n - a| < \varepsilon$ since $n \geqslant N \geqslant N_1$;
- $|z_n - a| < \varepsilon$ since $n \geqslant N \geqslant N_2$; and
- $x_n < y_n < z_n$ since $n \geqslant N \geqslant k$.

We will prove that $|y_n - a| < \varepsilon$ for all $n \geqslant N$. To see this let $n \geqslant N$. Either $y_n \geqslant a$ or $y_n \leqslant a$.

- If $y_n \geqslant a$, then we have $a \leqslant y_n \leqslant z_n$. It follows that

$$|y_n - a| = y_n - a \leqslant z_n - a = |z_n - a| < \varepsilon$$

- If $y_n \leqslant a$, then we have $x_n \leqslant y_n \leqslant a$. It follows that

$$|y_n - a| = a - y_n \leqslant a - x_n = |x_n - a| < \varepsilon$$

Since in both cases we have proved $|y_n - a| < \varepsilon$, we may conclude that $(y_n) \to a$. $\qquad \square$

Example 7.2.30

Fix $k \geqslant 1$. We prove that the sequence $(\frac{1}{n^k})_{n \geqslant 1}$ converges to zero.

Note that $n^k > n$, so that we have $0 < \frac{1}{n^k} \leqslant \frac{1}{n}$ for all $n \in \mathbb{N}$. We know that $(\frac{1}{n}) \to 0$ by Example 7.2.13, and $(0) \to 0$ since it is a constant sequence, so the squeeze theorem implies that $(\frac{1}{n^k}) \to 0$. $\qquad \triangleleft$

Exercise 7.2.31

Use the squeeze theorem, together with Example 7.2.30, to prove that $(2^{-n}) \to 0$. $\qquad \triangleleft$

Exercise 7.2.32

Fix $d \in \mathbb{N}$, and let $p(x) = a_0 + a_1 x + \cdots + a_d x^d$ and $q(x) = b_0 + b_1 x + \cdots + b_d x^d$ be polynomials with real coefficients. Prove that if $b_d \neq 0$, then $\left(\dfrac{p(n)}{q(n)} \right) \to \dfrac{a_d}{b_d}$. $\qquad \triangleleft$

Uniqueness of limits

We now prove that a sequence can have at most one limit. This will allow us to talk about 'the' limit of a sequence, and introduce notation for the limit of a sequence.

Theorem 7.2.33 (Uniqueness of limits)
Let (x_n) be a sequence and let $a, b \in \mathbb{R}$. If $(x_n) \to a$ and $(x_n) \to b$, then $a = b$.

Proof
We'll prove that $|a - b| = 0$, which will imply that $a = b$. To do this, we'll prove that $|a - b|$ is not positive: we already know it's non-negative, so this will imply that it is equal to zero. To prove that $|a - b|$ is not positive, we'll prove that it is less than every positive number.

So fix $\varepsilon > 0$. Then also $\frac{\varepsilon}{2} > 0$. The definition of convergence (Definition 7.2.15) tells us that:

- There exists $N_1 \in \mathbb{N}$ such that $|x_n - a| < \frac{\varepsilon}{2}$ for all $n \geqslant N_1$; and
- There exists $N_2 \in \mathbb{N}$ such that $|x_n - b| < \frac{\varepsilon}{2}$ for all $n \geqslant N_2$.

Let n be the greatest of N_1 and N_2. Then $n \geqslant N_1$ and $n \geqslant N_2$, and hence

$$|x_n - a| < \frac{\varepsilon}{2} \quad \text{and} \quad |x_n - b| < \frac{\varepsilon}{2}$$

By the triangle inequality (Theorem 7.1.9), it follows that

$$
\begin{aligned}
|a - b| = |(a - x_n) + (x_n - b)| & \qquad \text{by cancelling the } x_n \text{ terms} \\
\leqslant |a - x_n| + |x_n - b| & \qquad \text{by the triangle inequality} \\
= |x_n - a| + |x_n - b| & \qquad \text{by Exercise 7.1.5} \\
< \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon & \qquad \text{since } n \geqslant N_1 \text{ and } n \geqslant N_2
\end{aligned}
$$

Since $|a - b| < \varepsilon$ for all $\varepsilon > 0$, it follows that $|a - b|$ is a non-negative real number that is less than every positive real number, so that it is equal to zero.

Since $|a - b| = 0$, we have $a - b = 0$, and so $a = b$. $\qquad \square$

Theorem 7.2.33 justifies the following notation.

Definition 7.2.34
Let (x_n) be a convergent sequence. The limit of (x_n) is denoted by $\lim\limits_{n \to \infty} (x_n)$ (LaTeX code: `\lim_{n \to \infty}`).

Take heed of the fact that the symbol '∞' in Definition 7.2.34 does not have meaning on its own—it is simply a means of suggesting that as the index n gets greater, the values x_n of the terms in the sequence get closer to the limit.

Example 7.2.35
Examples 7.2.16 and 7.2.18 tell us that

$$\lim_{n\to\infty} \frac{1}{n} = 0 \quad \text{and} \quad \lim_{n\to\infty} \frac{2n}{n+1} \to 2$$

◁

Existence of limits

It is often useful to know *that* a sequence converges, but not necessary to go to the arduous lengths of computing its limit. However, as it currently stands, we don't really have any tools for proving that a sequence converges other than finding a limit for it! The remainder of this section is dedicated to deriving tools for finding out when a sequence does or does not converge, without needing to know exactly what the limit is.

Perhaps the most fundamental result is the *monotone convergence theorem* (Theorem 7.2.40), since it underlies the proofs of all the other results that we will prove. What it says is that if the terms in a sequence always increase, or always decrease, and the set of terms in the sequence is bounded, then the sequence converges to a limit.

The sequence (r_n) from Example 7.2.14, defined by $r_n = \frac{2n}{n+1}$ for all $n \in \mathbb{N}$, is an example of such a sequence. We proved that it converged by computing its limit in Example 7.2.18 and again in Example 7.2.27. We will soon (Example 7.2.43) use the monotone convergence theorem to give *yet another proof* that it converges, but this time without going to the trouble of first finding its limit.

Before we can state the monotone convergence theorem, we must first define what we mean by a *monotonic sequence*.

Definition 7.2.36
A sequence of real numbers (x_n) is...

- ...**increasing** if $m \leqslant n$ implies $x_m \leqslant x_n$ for all $m, n \in \mathbb{N}$;

- ...**decreasing** if $m \leqslant n$ implies $x_m \geqslant x_n$ for all $m, n \in \mathbb{N}$.

If a sequence is either increasing or decreasing, we say it is **monotonic**.

Example 7.2.37

The sequence (x_n) defined by $x_n = n^2$ for all $n \in \mathbb{N}$ is increasing, since for all $m, n \in \mathbb{N}$, if $m \leqslant n$, then $m^2 \leqslant n^2$. To see this, note that if $m \leqslant n$, then $n - m \geqslant 0$ and $n + m \geqslant 0$, so that

$$n^2 - m^2 = (n - m)(n + m) \geqslant 0 \cdot 0 = 0$$

and hence $n^2 \geqslant m^2$, as required. ◁

Example 7.2.38

The sequence (r_n) from Examples 7.2.14 and 7.2.27, defined by $r_n = \frac{2n}{n+1}$ for all $n \in \mathbb{N}$, is increasing. To see this, suppose $m \leqslant n$. Then $n = m + k$ for some $k \geqslant 0$. Now

$$
\begin{array}{lll}
0 \leqslant k & & \text{by assumption} \\
\Leftrightarrow m^2 + km + m \leqslant m^2 + km + m + k & & \text{adding } m^2 + km + m \text{ to both sides} \\
\Leftrightarrow m(m + k + 1) \leqslant (m + 1)(m + k) & & \text{factorising} \\
\Leftrightarrow m(n + 1) \leqslant (m + 1)n & & \text{since } n = m + k \\
\Leftrightarrow \dfrac{m}{m + 1} \leqslant \dfrac{n}{n + 1} & & \text{dividing both sides by } (m + 1)(n + 1) \\
\Leftrightarrow r_m \leqslant r_n & & \text{by definition of } (r_n)
\end{array}
$$

Note that the step where we divided through by $(m + 1)(n + 1)$ is justified since this quantity is positive.

It is perhaps useful to add that to *come up with* this proof, it is more likely that you would start with the assumption $r_m \leqslant r_n$ and derive that $k \geqslant 0$—noting that all steps are reversible then allows us to write it in the 'correct' order. ◁

Exercise 7.2.39

Prove that the sequence $(5^n - n^5)_{n \geqslant 0}$ is *eventually* increasing—that is, there is some $k \in \mathbb{N}$ such that $(5^n - n^5)_{n \geqslant k}$ is an increasing sequence. ◁

The monotone convergence theorem underlies all of the other tools for proving convergence of sequences that are to follow. It makes essential use of the completeness axiom.

Theorem 7.2.40 (Monotone convergence theorem)
Let (x_n) be a sequence of real numbers.

(a) If (x_n) is increasing and has an upper bound, then it converges;

(b) If (x_n) is decreasing and has a lower bound, then it converges.

Proof of (a)

We prove (a) here—part (b) is Exercise 7.2.41.

So suppose (x_n) is increasing and has an upper bound. Then:

(i) $x_m \leqslant x_n$ for all $m \leqslant n$; and

(ii) There is some real number u such that $u \geqslant x_n$ for all $n \in \mathbb{N}$.

Condition (ii) tells us that the set $\{x_n \mid n \in \mathbb{N}\} \subseteq \mathbb{R}$ has an upper bound. By the completeness axiom, it has a supremum a. We prove that $(x_n) \to a$.

So let $\varepsilon > 0$. We need to find $N \in \mathbb{N}$ such that $|x_n - a| < \varepsilon$ for all $n \geqslant N$.

Since a is a supremum of $\{x_n \mid n \in \mathbb{N}\}$, there is some $N \in \mathbb{N}$ such that $x_N > a - \varepsilon$.

Since (x_n) is increasing, by (i) we have $x_N \leqslant x_n$ for all $n \geqslant N$. Moreover, since a is an upper bound of the sequence, we actually have $x_N \leqslant x_n \leqslant a$ for all $n \geqslant N$.

Putting this together, for all $n \geqslant N$, we have

$$
\begin{aligned}
|x_n - a| &= a - x_n && \text{since } x_n - a \leqslant 0 \\
&\leqslant a - x_N && \text{since } x_N \leqslant x_n \text{ for all } n \geqslant N \\
&< \varepsilon && \text{since } x_N > a - \varepsilon
\end{aligned}
$$

It follows that $(x_n) \to a$, as required.   □

Exercise 7.2.41

Prove part (b) of the monotone convergence theorem (Theorem 7.2.40). That is, prove that if a sequence (x_n) is decreasing and has a lower bound, then it converges.   ◁

Example 7.2.42

The monotone convergence theorem can be used to show that many of the sequences that we have already seen converge, although it doesn't tell us what their limit is. For example, $\left(\frac{1}{n}\right)$ converges since it is a decreasing sequence that is bounded below by 0.   ◁

Example 7.2.43

Let (r_n) be our recurring example sequence from Examples 7.2.14, 7.2.27 and 7.2.38, defined by $r_n = \frac{2n}{n+1}$ for all $n \in \mathbb{N}$. We proved in Example 7.2.38 that (r_n) is increasing. Moreover, for all $n \in \mathbb{N}$ we have

$$
r_n = \frac{2n}{n+1} < \frac{2(n+1)}{n+1} = 2
$$

and so (r_n) is bounded above by 2. By the monotone convergence theorem, the sequence (r_n) converges. Unfortunately, the monotone convergence theorem does not tell us what the limit of (r_n) is, but we have already computed it twice!   ◁

Exercise 7.2.44

Use the monotone convergence theorem to prove that the sequence $\left(\frac{n!}{n^n}\right)$ converges. ◁

Exercise 7.2.45

A sequence (x_n) is defined recursively by $x_0 = 0$ and $x_{n+1} = \sqrt{2 + x_n}$ for all $n \geqslant 0$. That is,

$$x_n = \underbrace{\sqrt{2 + \sqrt{2 + \sqrt{\cdots + \sqrt{2}}}}}_{n \text{ '2's}}$$

Prove that (x_n) converges. ◁

We now define the notion of a *subsequence* of a sequence. A subsequence of a sequence is just like a subset of a set, except we can only pick out terms in a sequence in the order they appear.

Definition 7.2.46

Let (x_n) be a sequence of real numbers. A **subsequence** of (x_n) is a sequence of the form $(x_{n_i})_{i \geqslant 0}$, where $n_i < n_j$ for all $0 \leqslant i < j$.

In Definition 7.2.46 we were careful to write $(x_{n_i})_{i \geqslant 0}$ rather than just (x_{n_i}), because we wanted to emphasise that the indexing variable is i, rather than n. This is good practice in any situation where confusion might arise over which variable is the indexing variable.

Example 7.2.47

Define a sequence (x_n) by $x_n = (-1)^n$ for all $n \geqslant 0$.

$$(x_n)_{n \geqslant 0} = (1, -1, 1, -1, 1, -1, \ldots)$$

The subsequence (x_{2i}) is the constant sequence with value 1, since for each $i \geqslant 0$ we have $x_{2i} = (-1)^{2i} = 1$, and the subsequence (x_{2i+1}) is the constant sequence with value -1, since for each $i \geqslant 0$ we have $x_{2i+1} = (-1)^{2i+1} = -1$. ◁

Theorem 7.2.48

Let (x_n) be a sequence, let $a \in \mathbb{R}$, and suppose $(x_n) \to a$. Then every subsequence of (x_n) converges to a.

Proof

Let $(x_{n_i})_{i \geqslant 0}$ be a subsequence of (x_n). We need to prove that $(x_{n_i}) \to a$ as $i \to \infty$. To this end, fix $\varepsilon > 0$. We need to find $I \geqslant 0$ such that $|x_{n_i} - a| < \varepsilon$ for all $i \geqslant I$.

Since $(x_n) \to a$ as $n \to \infty$, there exists some $N \geqslant 0$ such that $|x_n - a| < \varepsilon$ for all $n \geqslant N$. Let $I \geqslant 0$ be least such that $n_I \geqslant N$. We know that I exists since we have $0 \leqslant n_0 < n_1 < n_2 < \ldots$.

But then for all $i \geqslant I$, we have $n_i \geqslant n_I \geqslant N$, and hence $|x_{n_i} - a| < \varepsilon$ by definition of N.

Hence the subsequence (x_{n_i}) converges to a, as required.   □

Exercise 7.2.49
Prove that a subsequence of an increasing sequence is increasing, that a subsequence of a decreasing sequence is decreasing, and that a subsequence of a constant sequence is constant.   ◁

We can use the monotone convergence theorem and the squeeze theorem to prove the following very powerful result, which is related to a notion in the field of topology known as *sequential compactness*.

Theorem 7.2.50 (Bolzano–Weierstrass theorem)
Every bounded sequence of real numbers has a convergent subsequence.

Proof
Let (x_n) be a sequence of real numbers and let $a, b \in \mathbb{R}$ be such that $a < x_n < b$ for each $n \geqslant 0$—the numbers a and b exist since the sequence (x_n) is bounded.

Our strategy is as follows. The sequence (x_n) is entirely contained inside the interval $[a, b]$, which has length $\ell = b - a$. Letting $c = \frac{a+b}{2}$ be the (arithmetic) mean of a and b, we see that one of the intervals $[a, c]$ or $[c, b]$, or possibly both, must contain infinitely many terms of the sequence (x_n)—but then this defines a subsequence of (x_n) which is entirely contained inside a sub-interval of $[a, b]$ whose length is $\frac{\ell}{2}$. We iterate this process inductively, obtaining smaller and smaller intervals that contain infinitely many terms in the sequence (x_n). The end-points of these intervals are then bounded monotone sequences—the sequence of lower end-points is increasing, and the sequence of upper end-points is decreasing. The monotone convergence theorem implies that both sequences converge. We will prove that they converge to the same limit, thereby 'trapping' a subsequence of (x_n), which will converge by the squeeze theorem.

Now let's put our strategy into action. We will define the terms n_i, a_i and b_i by induction on i, and then verify that the resulting subsequence $(x_{n_i})_{i \geqslant 0}$ converges.

First, define $n_0 = 0$, $a_0 = a$ and $b_0 = b$.

Now fix $i \geqslant 0$ and suppose that the numbers n_i, a_i and b_i have been defined in such a way that:

(i) $x_{n_i} \in [a_i, b_i]$;

(ii) $x_n \in [a_i, b_i]$ for infinitely many $n > n_i$;

(iii) $a_j \leqslant a_i < b_i \leqslant b_j$ for all $j \leqslant i$; and

(iv) $b_i - a_i = \frac{\ell}{2^i}$.

Write $c_i = \frac{a_i + b_i}{2}$. By condition (ii), it must be case that infinitely many of the terms x_n, for $n > n_i$, are contained in either $[a_i, c_i]$ or in $[c_i, b_i]$. In the former case, define $a_{i+1} = a_i$ and $b_{i+1} = c_i$; and in the latter case define $a_{i+1} = c_i$ and $b_{i+1} = b_i$; and then define $n_{i+1} > n_i$ such that $x_{n_{i+1}} \in [a_{i+1}, b_{i+1}]$.

Note that conditions (i)–(iv) are satisfied, with i now replaced by $i+1$. Indeed, (i) and (ii) are satisfied by definition of a_{i+1}, b_{i+1} and n_{i+1}. Condition (iii) is satisfied since either $a_{i+1} = a_i$ or $a_{i+1} = \frac{a_i + c_i}{2} \geqslant a_i$, and likewise for b_{i+1}. Condition (iv) is satisfied since

$$c_i - a_i = \frac{a_i + b_i}{2} - a_i = \frac{b_i - a_i}{2} = \frac{\ell/2^i}{2} = \frac{\ell}{2^{i+1}}$$

and likewise $b_i - c_i = \frac{\ell}{2^{i+1}}$.

Since by construction we have $n_i < n_{i+1}$ for each $i \geqslant 0$, we have defined a subsequence $(x_{n_i})_{i \geqslant 0}$ of (x_n).

Now the sequence (a_i) is increasing and is bounded above by b, and the sequence (b_i) is decreasing and is bounded below by a. By the monotone convergence theorem $(a_i) \to a^\star$ and $(b_i) \to b^\star$ for some $a^\star, b^\star \in \mathbb{R}$. But moreover we have

$$\frac{\ell}{2^i} = b_i - a_i \to b^\star - a^\star$$

Since $\frac{\ell}{2^i} \to 0$, we have $b^\star - a^\star = 0$ by uniqueness of limits, and so $a^\star = b^\star$. Write $x^\star$ for the common value of $a^\star$ and $b^\star$.

Finally, we have $a_i \leqslant x_{n_i} \leqslant b_i$ for all $i \geqslant 0$, so that $x_{n_i} \to x^\star$ by the squeeze theorem. □

The Bolzano–Weierstrass theorem can be used to prove that a sequence converges by verifying that its terms get arbitrarily close together. Such sequences are called *Cauchy* sequences, and the fact that all Cauchy sequences converge is proved in Theorem 7.2.54.

Definition 7.2.51

A **Cauchy sequence** is a sequence (x_n) of real numbers such that, for all $\varepsilon > 0$, there exists $N \in \mathbb{N}$ such that $|x_m - x_n| < \varepsilon$ for all $m, n \geqslant N$.

Example 7.2.52

Let (r_n) be our favourite recurring example sequence from Examples 7.2.14, 7.2.27, 7.2.38 and 7.2.43, defined by $r_n = \dfrac{2n}{n+1}$ for all $n \in \mathbb{N}$. We prove that (r_n) is Cauchy.

First note that, given $m, n \geqslant 1$, we have

$$|r_m - r_n| = \left| \frac{2m}{m+1} - \frac{2n}{n+1} \right| = \frac{2|m-n|}{(m+1)(n+1)} = \frac{2\left|\frac{1}{n} - \frac{1}{m}\right|}{(1+\frac{1}{m})(1+\frac{1}{n})}$$

Now fix $\varepsilon > 0$, and let $N \in \mathbb{N}$ be such that $\frac{1}{m} < \frac{\varepsilon}{2}$ and $\frac{1}{n} < \frac{\varepsilon}{2}$ for all $m, n > N$. Note that such a value of N exists by Example 7.2.13.

Now let $m, n \geqslant N$. Then $\left|\frac{1}{n} - \frac{1}{m}\right| < \frac{\varepsilon}{2}$ since both $\frac{1}{m}$ and $\frac{1}{n}$ are elements of $(0, \frac{\varepsilon}{2})$. Moreover $1 + \frac{1}{m} > 1$ and $1 + \frac{1}{n} > 1$. It follows that, for all $m, n \geqslant N$, we have

$$|r_m - r_n| < \frac{2 \cdot \frac{\varepsilon}{2}}{1 \cdot 1} = \varepsilon$$

Hence (r_n) is Cauchy, as claimed. ◁

The following exercise generalises the previous example.

Exercise 7.2.53

Prove that every convergent sequence is a Cauchy sequence. ◁

Theorem 7.2.54

Every Cauchy sequence of real numbers converges.

Proof

Let (x_n) be a Cauchy sequence of real numbers.

First note that (x_n) is bounded. To see this, note that by definition of Cauchy sequences, there is some $N \in \mathbb{N}$ such that $|x_m - x_n| < 1$ for all $m, n \geqslant N$. In particular, $|x_m - x_N| < 1$ for all $m \geqslant N$. This means that the sequence (x_n) is bounded below by

$$a = \min\{x_0, x_1, \ldots, x_{N-1}, x_N - 1\}$$

and is bounded above by

$$b = \max\{x_0, x_1, \ldots, x_{N-1}, x_N + 1\}$$

By the Bolzano–Weierstrass theorem (Theorem 7.2.50), the sequence (x_n) has a convergent subsequence (x_{n_i}). Let $x^\star = \lim_{i \to \infty}(x_{n_i})$. We prove that $(x_n) \to x^\star$.

So let $\varepsilon > 0$. Fix M sufficiently large that:

- $|x_{n_i} - x^\star| < \frac{\varepsilon}{3}$ for all $n_i \geqslant M$; and

- $|x_n - x_m| < \frac{\varepsilon}{3}$ for all $m, n \geqslant M$.

Such a value of M exists by convergence of (x_{n_i}) and the Cauchy property of (x_n).

Fix $n \geqslant M$, and let $i \in \mathbb{N}$ be arbitrary such that $n_i \geqslant M$. Then we have

$$
\begin{aligned}
&|x_n - x^\star| \\
&= |(x_n - x_M) + (x_M - x_{n_i}) + (x_{n_i} - x^\star)| \qquad &&\text{rearranging} \\
&\leqslant |x_n - x_M| + |x_M - x_{n_i}| + |x_{n_i} - x^\star| \qquad &&\text{by the triangle inequality} \\
&< \frac{\varepsilon}{3} + \frac{\varepsilon}{3} + \frac{\varepsilon}{3} \qquad &&\text{by the above properties} \\
&= \varepsilon
\end{aligned}
$$

Hence $(x_n) \to x^\star$, as required. $\qquad\qquad\square$

Section 7.3
Series and sums

Warning!
This section is not yet finished—do not rely on its correctness or completeness.

To do:

Proposition 7.3.1

The series $\sum\limits_{n=1}^{\infty} \dfrac{1}{n}$ diverges.

Idea of proof

By rounding up denominators to the next power of 2, we get

$$\frac{1}{1}+\frac{1}{2}+\frac{1}{3}+\frac{1}{4}+\frac{1}{5}+\frac{1}{6}+\frac{1}{7}+\frac{1}{8}+\cdots \geqslant \frac{1}{1}+\frac{1}{2}+\underbrace{\frac{1}{4}+\frac{1}{4}}_{=1/2}+\underbrace{\frac{1}{8}+\frac{1}{8}+\frac{1}{8}+\frac{1}{8}}_{=1/2}+\cdots$$

This provides a lower bound on the sum, which is infinite since we are adding infinitely many multiples of $\frac{1}{2}$. The following proof makes this idea precise. $\square$

Proof

For each $k \in \mathbb{N}$ and each $2^k < r \leqslant 2^{k+1}$, we have $\frac{1}{r} \geqslant 2^{k+1}$. Therefore

$$\sum_{r=2^k+1}^{2^{k+1}} \frac{1}{r} \geqslant \sum_{r=2^k+1}^{2^{k+1}} \frac{1}{2^{k+1}} = (2^{k+1}-2^k) \cdot \frac{1}{2^{k+1}} = \frac{2^k}{2^{k+1}} = \frac{1}{2}$$

It follows that

$$\sum_{r=0}^{\infty} \frac{1}{r} = 1 + \sum_{k=0}^{\infty} \sum_{r=2^k+1}^{2^{k+1}} \frac{1}{r} \geqslant 1 + \sum_{k=0}^{\infty} \frac{1}{2} = \infty$$

and so the series diverges. $\square$

To do:

Proposition 7.3.2

Let $x \in \mathbb{R}$ with $-1 < x < 1$. Then $\sum\limits_{n\in\mathbb{N}} x^n = \dfrac{1}{1-x}$.

Proof

Given $N \in \mathbb{N}$, the N^{th} partial sum S_N of the series is given by by

$$S_N = \sum_{n=0}^{N} x^n = 1 + x + x^2 + \cdots + x^N$$

Note that
$$xS_N = \sum_{n=0}^{n} x^{n+1} = x + x^2 + \cdots + x^{N+1} = S_{N+1} - 1$$

and hence

$$(1-x)S_N = S_N - xS_N = S_N - (S_{N+1} - 1) = 1 - (S_{N+1} - S_N) = 1 - x^{N+1}$$

and hence dividing by $1-x$, which is permissible since $x \neq 1$, yields

$$S_N = \frac{1 - x^{N+1}}{1 - x}$$

To do: Finish proof $\square$

Proposition 7.3.3

Let $x \in \mathbb{R}$ with $-1 < x < 1$. Then $\displaystyle\sum_{n \in \mathbb{N}} nx^{n-1} = \frac{1}{(1-x)^2}$

Continuous functions

<div style="text-align:center">

Warning!

</div>

This section is not yet finished—do not rely on its correctness or completeness.

To do:

Open sets

To do:

Definition 7.4.1

A subset $U \subseteq \mathbb{R}$ is **open** if, for all $a \in U$, there exists some $\delta > 0$ such that $(a - \delta, a + \delta) \subseteq U$.

Example 7.4.2

The subset $\mathbb{R}$ of $\mathbb{R}$ is open, since for all $a \in \mathbb{R}$, we have $(a - 1, a + 1) \subseteq \mathbb{R}$.   ◁

Example 7.4.3

The subset $\mathbb{Z} \subseteq \mathbb{R}$ is not open. To see this, let $a \in \mathbb{Z}$. Fix $\delta > 0$ and define $\delta' = \min\{\delta, 1\}$. Then

$$a - \delta < a < a + \frac{\delta'}{2} < a + \delta' \leqslant a + \delta$$

so $a + \frac{\delta'}{2} \in (a - \delta, a + \delta)$. However, $a < a + \frac{\delta'}{2} \leqslant a + \frac{1}{2} < a + 1$, and so $a + \frac{\delta'}{2} \notin \mathbb{Z}$.

We have shown that there is no $\delta > 0$ such that $(a - \delta, a + \delta) \subseteq \mathbb{Z}$, so that $\mathbb{Z}$ is not open.   ◁

Exercise 7.4.4

Prove that $\varnothing$ and $(0, \infty)$ are open subsets of $\mathbb{R}$, and that $[0, \infty)$ and $\mathbb{Q}$ are not open subsets of $\mathbb{R}$.   ◁

Open sets are very closely related to open intervals, as we shall see in Theorem 7.4.7.

Proposition 7.4.5

Let $a, b \in \mathbb{R}$ with $a < b$. Then the open interval (a, b) is open.

Proof

Let $x \in (a,b)$, and define $\delta = \min\{x-a, b-x\}$. Note that $\delta > 0$ since $a < x < b$.

To see that $(x-\delta, x+\delta) \subseteq (a,b)$, let $y \in (x-\delta, x+\delta)$. Then since $\delta \leqslant x-a$, we have

$$y > x - \delta \geqslant x - (x-a) = a$$

and since $\delta \leqslant b - x$ we have

$$y < x + \delta \leqslant x + (b-x) = b$$

Hence $a < y < b$, so $y \in (a,b)$, as required. $\qquad\square$

Exercise 7.4.6 (Arbitrary unions of open sets are open)
Let $\{U_i \mid i \in I\}$ be a family of open subsets of $\mathbb{R}$. Prove that $\bigcup_{i \in I} U_i$ is open. $\qquad\triangleleft$

Theorem 7.4.7
A subset $U \subseteq \mathbb{R}$ is open if and only if it is a union of open intervals.

Proof

Since open intervals are open (Proposition 7.4.5) and unions of open sets are open (Exercise 7.4.6), it follows that if a subset $U \subseteq \mathbb{R}$ is a union of open intervals then it is open.

Conversely, suppose $U \subseteq \mathbb{R}$ is open. For each $a \in U$, let $\delta_a > 0$ be such that $(a-\delta_a, a+\delta_a) \subseteq U$.

We prove that $U = \bigcup_{a \in U}(a-\delta_a, a+\delta_a)$.

- ($\subseteq$) Let $x \in U$. Then $x \in (x-\delta_x, x+\delta_x)$, so $x \in \bigcup_{a \in U}(a-\delta_a, a+\delta_a)$.

- ($\supseteq$) Let $x \in \bigcup_{a \in U}(a-\delta_a, a+\delta_a)$. Then $x \in (a-\delta_a, a+\delta_a)$ for some $a \in U$, and so $x \in U$ by our assumption that $(a-\delta_a, a+\delta_a) \subseteq U$.

Hence U is a union of open intervals, as required. $\qquad\square$

To do:

Proposition 7.4.8 (Finite intersections of open sets are open)
Let $n \in \mathbb{N}$ and let $U_1, U_2, \ldots, U_n$ be open subsets of $\mathbb{R}$. Then the intersection $\bigcap_{k=1}^{n} U_k$ is open.

Proof

Define $U = \bigcap_{k=1}^{n} U_k$ and let $a \in U$. Then $a \in U_k$ for each $k \in [n]$.

Since each set U_k is open, there exist positive real numbers $\delta_k > 0$ such that $(a-\delta_k, a+\delta_k) \subseteq U_k$. for each $k \in [n]$.

Now define $\delta = \min\{\delta_k \mid k \in [n]\}$. Then $\delta > 0$. To see that $(a - \delta, a + \delta) \subseteq U$, let $x \in (a - \delta, a + \delta)$. Then for each $k \in [n]$ we have

$$a - \delta_k \leqslant a - \delta < x < a + \delta \leqslant a + \delta_k$$

so that $x \in (a - \delta_k, a + \delta_k)$. But then $x \in U_k$ since $(a - \delta_k, a + \delta_k) \subseteq U_k$.

Since $x \in U_k$ for each $k \in [n]$, we have $x \in U$. So $(a - \delta, a + \delta) \subseteq U$, as required.

$\square$

Exercise 7.4.9
Find a family $\{U_n \mid n \in \mathbb{N}\}$ of open subsets of $\mathbb{R}$ whose intersection is not open.

$\triangleleft$

To do:

Continuous functions

To do:

Convention 7.4.10
When discussing functions $f : D \to \mathbb{R}$ in this section, we assume that the domain D is an inhabited interval in $\mathbb{R}$. Thus either $D = \mathbb{R}$, or D is one of the subsets of the kind defined in Definition 2.1.11.

$\triangleleft$

Definition 7.4.11
Let $f : D \to \mathbb{R}$ be a function and let $a \in D$. Then f is **continuous at** a if, for all $\varepsilon > 0$, there exists $\delta > 0$ such that, for all $x \in D$, if $|x - a| < \delta$, then $|f(a) - f(x)| < \varepsilon$. That is:

$$\forall \varepsilon > 0, \exists \delta > 0, \forall x \in D, (|x - a| < \delta \Rightarrow |f(x) - f(a)| < \varepsilon)$$

We say f is **continuous** if f is continuous at a for all $a \in D$.

To do: Examples, etc.

Theorem 7.4.12
Let $f : D \to \mathbb{R}$ be a function. Then f is continuous if and only if, for all open subsets $U \subseteq \mathbb{R}$, we have $f^{-1}[U] = V \cap D$ for some open $V \subseteq \mathbb{R}$.

Proof

- ($\Rightarrow$) Suppose f is continuous, and let $U \subseteq \mathbb{R}$ be open. If $f^{-1}[U] = \varnothing$, then we can take $V = \varnothing$. So assume that $f^{-1}[(a,b)]$ is inhabited and fix $p \in f^{-1}[U]$.

  Since $f(p) \in U$ and U is open, there exists $\varepsilon_p > 0$ such that

  $$(f(p) - \varepsilon_p, f(p) + \varepsilon_p) \subseteq U$$

  By continuity of f, there exists $\delta_p > 0$ such that, for all $x \in D$, if $|x - p| < \delta_p$, then $|f(x) - f(p)| < \varepsilon_p$.

  But this says precisely that if $x \in (p - \delta_p, p + \delta_p) \cap D$, then $f(x) \in (f(p) - \varepsilon_p, f(p) + \varepsilon_p)$.

  Since $(f(p) - \varepsilon_p, f(p) + \varepsilon_p) \subseteq U$, it follows that $(p - \delta_p, p + \delta_p) \cap D \subseteq f^{-1}[U]$.

  Define $V \subseteq \mathbb{R}$ by

  $$V = \bigcup_{p \in f^{-1}[U]} (p - \delta_p, p + \delta_p)$$

  Then V is open by Theorem 7.4.7, and $V \cap D \subseteq f^{-1}[U]$ since each $(p - \delta_p, p + \delta_p) \subseteq f^{-1}[U]$.

  To see that $f^{-1}[U] \subseteq V \cap D$, let $p \in f^{-1}[U]$. Then $p \in D$ and $p \in (p - \delta_p, p + \delta_p)$, so that $p \in V \cap D$ as required.

  So we have $f^{-1}[U] = V \cap D$, with $V \subseteq \mathbb{R}$ open, as required.

- ($\Leftarrow$) Suppose that for all open subsets $U \subseteq \mathbb{R}$, we have $f^{-1}[U] = V \cap D$ for some open $V \subseteq \mathbb{R}$.

  Let $a \in D$ and let $\varepsilon > 0$. Let $V \subseteq \mathbb{R}$ be an open set such that

  $$f^{-1}[(f(a) - \varepsilon, f(a) + \varepsilon)] = V \cap D$$

  In particular, we have $a \in V \cap D$, so $a \in V$, and so there is some $\delta > 0$ such that $(a - \delta, a + \delta) \subseteq V$.

  Now let $x \in D$ with $|x - a| < \delta$. Then $x \in V \cap D$, and so $f(x) \in (f(a) - \varepsilon, f(a) + \varepsilon)$. But then $|f(x) - f(a)| < \varepsilon$, as required.

  So f is continuous.

  $\square$

Exercise 7.4.13

Prove that a function $f : D \to \mathbb{R}$ is continuous if and only if, for all sequences (x_n) in D such that $(x_n) \to a \in D$, we have $(f(x_n)) \to f(a)$. ◁

To do:

Limit of a function

To do:

> **Definition 7.4.14**
> Let $D \subseteq \mathbb{R}$. An **interior point** of D is an element $a \in D$ such that $(a - \delta, a + \delta) \subseteq D$ for some $\delta > 0$. Write D° (LaTeX code: `D^{\circ}`) for the set of all interior points of D, called the **interior** of D.

Example 7.4.15
Consider the half-open interval $[0, 1)$. The element $\frac{1}{2} \in [0, 1)$ is an interior point of $[0, 1)$, since $(\frac{1}{2} - \frac{1}{2}, \frac{1}{2} + \frac{1}{2}) = (0, 1) \subseteq [0, 1)$. However, the element 0 is not a limit point of $[0, 1)$, since for all $\delta > 0$ we have $-\frac{\delta}{2} \in (-\delta, \delta)$ but $-\frac{\delta}{2} \notin [0, 1)$.   ◁

Exercise 7.4.16
Prove that a subset $U \subseteq \mathbb{R}$ is open if and only if $U^\circ = U$.   ◁

To do:

> **Definition 7.4.17**
> Let $f : D \to \mathbb{R}$ be a function, let $a \in D^\circ$. A **limit of** $f(x)$ **as** x **tends to** a is a real number ℓ such that
>
> $$\forall \varepsilon > 0, \exists \delta > 0, \forall x \in D, (0 < |x - a| < \delta \Rightarrow |f(x) - \ell| < \varepsilon)$$
>
> We write '$f(x) \to \ell$ as $x \to a$' to denote the assertion that $f(x)$ tends to ℓ as x tends to a.

To do:

Exercise 7.4.18
Let $f : D \to \mathbb{R}$, $a \in D^\circ$ and $\ell_1, \ell_2 \in \mathbb{R}$. Prove that if $f(x) \to \ell_1$ as $x \to a$, and $f(x) \to \ell_2$ as $x \to a$, then $\ell_1 = \ell_2$.   ◁

To do:

Boundedness

To do:

Section 7.Q
Chapter 7 exercises

Under construction!
The end-of-chapter exercise sections are new and in an incomplete state.

Probability and measure

Section 8.1
Discrete probability spaces

Probability theory is a field of mathematics which attempts to model randomness and uncertainty in the 'real world'. The mathematical machinery it develops allows us to understand how this randomness behaves and to extract information which is useful for making predictions.

Discrete probability theory, in particular, concerns situations in which the possible outcomes form a *countable* set. This simplifies matters considerably: if there are only countably many outcomes, then the probability that any event occurs is determined entirely by the probabilities that the individual outcomes comprised by the event occur.

For example, the number N of words spoken by a child over the course of a year takes values in $\mathbb{N}$, so is discrete. To each $n \in \mathbb{N}$, we may assign a probability that $N = n$, which can take positive values in a meaningful way, and from these probabilities we can compute the probabilities of more general events occurring (e.g. the probability that the child says under a million words). However, the height H grown by the child over the same period takes values in $[0, \infty)$, which is uncountable; for each $h \in [0, \infty)$, the probability that $H = h$ is zero, so these probabilities give us no information. We must study the behaviour of H through some other means.

In this chapter, we will concern ourselves only with the discrete setting.

It is important to understand from the outset that, although we use language like *outcome*, *event*, *probability* and *random*, and although we use real-world examples, everything we do concerns mathematical objects: sets, elements of sets, and functions. If we say, for example, "the probability that a roll of a fair six-sided die shows 3 or 4 is $\frac{1}{3}$," we are actually interpreting the situation mathematically— the *outcomes* of the die rolls are interpreted as the elements of the set $[6]$; the *event* that the die shows 3 or 4 is interpreted as the subset $\{3, 4\} \subseteq [6]$; and the *probability* that this event occurs is the value of a particular function $\mathbb{P} : \mathscr{P}([6]) \to [0, 1]$ on input $\{3, 4\}$. The mathematical interpretation is called a **model** of the real-world situation.

Definition 8.1.1

A **discrete probability space** is a pair $(\Omega, \mathbb{P})$ (LATEX code: `(\Omega, \mathbb{P})`), consisting of a countable set Ω and a function $\mathbb{P} : \mathscr{P}(\Omega) \to [0,1]$, such that

   (i) $\mathbb{P}(\Omega) = 1$; and

  (ii) (**Countable additivity**) If $\{A_i \mid i \in I\}$ is any family of pairwise disjoint subsets of Ω, indexed by a countable set I, then

$$\mathbb{P}\left(\bigcup_{i \in I} A_i\right) = \sum_{i \in I} \mathbb{P}(A_i)$$

The set Ω is called the **sample space**; the elements $\omega \in \Omega$ are called **outcomes**;[a] the subsets $A \subseteq \Omega$ are called **events**; and the function $\mathbb{P}$ is called the **probability measure**. Given an event A, the value $\mathbb{P}(A)$ is called the **probability of** A.

---

[a]The symbols Ω, ω (LATEX code: `\Omega, \omega`) are the upper- and lower-case forms, respectively, of the Greek letter *omega*.

There is a general notion of a probability space, which does not require the sample space Ω to be countable. This definition is significantly more technical (Definition 8.3.10), so we restrict our attention in this section to *discrete* probability spaces. Thus, whenever we say 'probability space' in this section, the probability space can be assumed to be discrete. However, when our proofs do not specifically use countability of Ω, they typically are true of arbitrary probability spaces. As such, we will specify discreteness in the statement of results only when countability of the sample space is required.

Example 8.1.2

We model the roll of a fair six-sided die.

The possible **outcomes** of the roll are 1, 2, 3, 4, 5 and 6, so we can take $\Omega = [6]$ to be the sample space.

The **events** correspond with subsets of $[6]$. For example:

- $\{4\}$ is the event that the die roll shows 4. This event occurs with probability $\frac{1}{6}$.

- $\{1,3,5\}$ is the event that the die roll is odd. This event occurs with probability $\frac{1}{2}$

- $\{1,4,6\}$ is the event that the die roll is not prime. This event occurs with probability $\frac{1}{2}$.

- $\{3,4,5,6\}$ is the event that the die roll shows a number greater than 2. This event occurs with probability $\frac{2}{3}$.

- $\{1,2,3,4,5,6\}$ is the event that anything happens. This event occurs with probability 1.

- $\varnothing$ is the event that nothing happens. This event occurs with probability 0.

More generally, since each outcome occurs with equal probability $\frac{1}{6}$, we can define

$$\mathbb{P}(A) = \frac{|A|}{6} \text{ for all events } A$$

We will verify that $\mathbb{P}$ defines a probability measure on $[6]$ in Example 8.1.6.   ◁

Example 8.1.3
Let $(\Omega, \mathbb{P})$ be a probability space. We prove that $\mathbb{P}(\varnothing) = 0$.

Note that Ω and $\varnothing$ are disjoint, so by countable additivity, we have

$$1 = \mathbb{P}(\Omega) = \mathbb{P}(\Omega \cup \varnothing) = \mathbb{P}(\Omega) + \mathbb{P}(\varnothing) = 1 + \mathbb{P}(\varnothing)$$

Subtracting 1 throughout yields $\mathbb{P}(\varnothing) = 0$, as required.   ◁

Exercise 8.1.4
Let $(\Omega, \mathbb{P})$ be a probability space. Prove that

$$\mathbb{P}(\Omega \setminus A) = 1 - \mathbb{P}(A)$$

for all events A.   ◁

Countable additivity of probability measures—that is, condition (ii) in Definition 8.1.1—implies that probabilities of events are determined by probabilities of individual outcomes. This is made precise in Proposition 8.1.5.

Proposition 8.1.5
Let Ω be a countable set and let $\mathbb{P} : \mathscr{P}(\Omega) \to [0,1]$ be a function such that $\mathbb{P}(\Omega) = 1$. The following are equivalent:

(i) $\mathbb{P}$ is a probability measure on Ω;

(ii) $\sum_{\omega \in A} \mathbb{P}(\{\omega\}) = \mathbb{P}(A)$ for all $A \subseteq \Omega$.

Proof
Since $\mathbb{P}(\Omega) = 1$, it suffices to prove that condition (ii) of Proposition 8.1.5 is equivalent to countable additivity of $\mathbb{P}$.

- (i)$\Rightarrow$(ii). Suppose $\mathbb{P}$ is a probability measure on Ω. Let $A \subseteq \Omega$.

Note that since $A \subseteq \Omega$ and Ω is countably infinite, it follows that $\{\{\omega\} \mid \omega \in A\}$ is a countable family of pairwise disjoint sets. By countable additivity, we have

$$\mathbb{P}(A) = \mathbb{P}\left(\bigcup_{\omega \in A} \{\omega\}\right) = \sum_{\omega \in A} \mathbb{P}(\{\omega\})$$

as required. Hence condition (ii) of the proposition is satisfied.

- (ii)$\Rightarrow$(i). Suppose that $\sum_{\omega \in A} \mathbb{P}(\{\omega\}) = \mathbb{P}(A)$ for all $A \subseteq \Omega$. We prove that $\mathbb{P}$ is a probability measure on Ω.

So let $\{A_i \mid i \in I\}$ be a family of pairwise disjoint events, indexed by a countable set I. Define $A = \bigcup_{i \in I} A_i$. Since the sets A_i partition A, summing over elements of A is the same as summing over each of the sets A_i individually, and then adding those results together; specifically, for each A-tuple $(p_\omega)_{\omega \in A}$, we have

$$\sum_{\omega \in A} p_\omega = \sum_{i \in I} \sum_{\omega \in A_i} p_\omega$$

Hence

$$\begin{aligned}
\mathbb{P}(A) &= \sum_{\omega \in A} \mathbb{P}(\{\omega\}) && \text{by condition (ii) of the proposition} \\
&= \sum_{i \in I} \sum_{\omega \in A_i} \mathbb{P}(\{\omega\}) && \text{by the above observation} \\
&= \sum_{i \in I} \mathbb{P}(A_i) && \text{by condition (ii) of the proposition}
\end{aligned}$$

So $\mathbb{P}$ satisfies the countable additivity condition. Thus $\mathbb{P}$ is a probability measure on Ω.

Hence the two conditions are equivalent. $\qquad\square$

Example 8.1.6
We prove that the function $\mathbb{P}$ from Example 8.1.2 truly does define a probability measure. Indeed, let $\Omega = [6]$ and let $\mathbb{P} : \mathscr{P}(\Omega) \to [0, 1]$ be defined by

$$\mathbb{P}(A) = \frac{|A|}{6} \text{ for all events } A$$

Then $\mathbb{P}(\Omega) = \frac{6}{6} = 1$, so condition (i) in Definition 8.1.1 is satisfied. Moreover, for each $A \subseteq [6]$ we have

$$\sum_{\omega \in A} \mathbb{P}(\{\omega\}) = \sum_{\omega \in A} \frac{1}{6} = \frac{|A|}{6} = \mathbb{P}(A)$$

so, by Proposition 8.1.5, $\mathbb{P}$ defines a probability measure on $[6]$. $\qquad\triangleleft$

Proposition 8.1.5 makes defining probability measures much easier, since it implies that probability measures are determined entirely by their values on individual outcomes. This means that, in order to define a probability measure, we only need to specify its values on individual outcomes and check that the sum of these probabilities is equal to 1. This is significantly easier than defining $\mathbb{P}(A)$ on *all* events $A \subseteq \Omega$ and checking the two conditions of Definition 8.1.1.

This is made precise in Proposition 8.1.7 below.

Proposition 8.1.7

Let Ω be a countable set and, for each $\omega \in \Omega$, let $p_\omega \in [0,1]$. If $\sum_{\omega \in \Omega} p_\omega = 1$, then there is a unique probability measure $\mathbb{P}$ on Ω such that $\mathbb{P}(\{\omega\}) = p_\omega$ for each $\omega \in \Omega$.

Proof
We prove existence and uniqueness of $\mathbb{P}$ separately.

- **Existence.** Define $\mathbb{P} : \mathscr{P}(\Omega) \to [0,1]$ be defined by

$$\mathbb{P}(A) = \sum_{\omega \in A} p_\omega$$

  for all events $A \subseteq \Omega$. Then condition (ii) of Proposition 8.1.5 is automatically satisfied, and indeed $\mathbb{P}(\{\omega\}) = p_\omega$ for each $\omega \in \Omega$. Moreover

$$\mathbb{P}(\Omega) = \sum_{\omega \in \Omega} \mathbb{P}(\{\omega\}) = \sum_{\omega \in \Omega} p_\omega = 1$$

  and so condition (i) of Definition 8.1.1 is satisfied. Hence $\mathbb{P}$ defines a probability measure on Ω.

- **Uniqueness.** Suppose that $\mathbb{P}' : \mathscr{P}(\Omega) \to [0,1]$ is another probability measure such that $\mathbb{P}'(\{\omega\}) = p_\omega$ for all $\omega \in \Omega$. For each event $A \subseteq \Omega$, condition (ii) of Proposition 8.1.5 implies that

$$\mathbb{P}'(A) = \sum_{\omega \in A} \mathbb{P}'(\{\omega\}) = \sum_{\omega \in A} p_\omega = \mathbb{P}(A)$$

  hence $\mathbb{P}' = \mathbb{P}$.

So $\mathbb{P}$ is uniquely determined by the values p_ω. $\qquad\square$

The assignments of $p_\omega \in [0,1]$ to each $\omega \in \Omega$ in fact defines something that we will later defined to be a *probability mass function* (Definition 8.2.5).

With Proposition 8.1.7 proved, we will henceforth specify probability measures $\mathbb{P}$ on sample spaces Ω by specifying only the values of $\mathbb{P}(\{\omega\})$ for $\omega \in \Omega$.

Example 8.1.8

Let $p \in [0,1]$. A coin, which shows heads with probability p, is repeatedly flipped until heads shows.

The outcomes of such a sequence of coin flips all take the form

$$\underbrace{(\text{tails}, \text{tails}, \cdots, \text{tails}}_{n \text{ 'tails'}}, \text{heads})$$

for some $n \in \mathbb{N}$. Identifying such a sequence with the number n of flips before heads shows, we can take $\Omega = \mathbb{N}$ to be the sample space.

For each $n \in \mathbb{N}$, we can define

$$\mathbb{P}(\{n\}) = (1-p)^n p$$

This will define a probability measure on $\mathbb{N}$, provided these probabilities all sum to 1; and indeed by Proposition 7.3.2, we have

$$\sum_{n \in \mathbb{N}} \mathbb{P}(\{n\}) = \sum_{n \in \mathbb{N}} (1-p)^n p = p \cdot \frac{1}{1-(1-p)} = p \cdot \frac{1}{p} = 1$$

By Proposition 8.1.7, it follows that $(\Omega, \mathbb{P})$ is a probability space. ◁

Exercise 8.1.9

A fair six-sided die is rolled twice. Define a probability space $(\Omega, \mathbb{P})$ that models this situation. ◁

Exercise 8.1.10

Let $(\Omega, \mathbb{P})$ be a probability space and let A, B be events with $A \subseteq B$. Prove that $\mathbb{P}(A) \leqslant \mathbb{P}(B)$. ◁

Set operations on events

In the real world, we might want to talk about the probability that two events both happen, or the probability that an event doesn't happen, or the probability that at least one of some collection of events happens. This is interpreted mathematically in terms of set operations.

Example 8.1.11

Let $(\Omega, \mathbb{P})$ be the probability space modelling two rolls of a fair six-sided die—that is, the sample space $\Omega = [6] \times [6]$ with probability measure $\mathbb{P}$ defined by $\mathbb{P}(\{(a,b)\}) = \frac{1}{36}$ for each $(a,b) \in \Omega$.

Let *A* be the event that the sum of the die rolls is even, that is

$$A = \begin{cases} (1,1), & (1,3), & (1,5), & (2,2), & (2,4), & (2,6), \\ (3,1), & (3,3), & (3,5), & (4,2), & (4,4), & (4,6), \\ (5,1), & (5,3), & (5,5), & (6,2), & (6,4), & (6,6) \end{cases}$$

and let *B* be the event that the sum of the die rolls is greater than or equal to 9, that is

$$B = \{(3,6),(4,5),(4,6),(5,4),(5,5),(5,6),(6,3),(6,4),(6,5),(6,6)\}$$

Then

- Consider the event that the sum of the die rolls is even **or** greater than or equal to 9. An outcome ω gives rise to this event precisely when either $\omega \in A$ or $\omega \in B$; so the event in question is $A \cup B$;

- Consider the event that the sum of the die rolls is even **and** greater than or equal to 9. An outcome ω gives rise to this event precisely when both $\omega \in A$ and $\omega \in B$; so the event in question is $A \cap B$;

- Consider the event that the sum of the die rolls is **not** even. An outcome ω gives rise to this event precisely when $\omega \notin A$; so the event in question is is $([6] \times [6]) \setminus A$.

Thus we can interpret 'or' as union, 'and' as intersection, and 'not' as relative complement in the sample space.   ◁

The intuition provided by Example 8.1.11 is formalised in Exercise 8.1.13. Before we do this, we adopt a convention that simplifies notation when discussing events in probability spaces.

Definition 8.1.12
Let $(\Omega, \mathbb{P})$ be a probability space. The **complement** of an event $A \subseteq \Omega$ is the event $\Omega \setminus A \subseteq \Omega$. We write A^c (LaTeX code: `A^c`) for $\Omega \setminus A$.

That is, when we talk about the complement *of an event*, we really mean their relative complement inside the sample space.

Exercise 8.1.13
Let $(\Omega, \mathbb{P})$ be a probability space, and let $p(\omega), q(\omega)$ be logical formulae with free variable ω ranging over Ω. Let

$$A = \{\omega \in \Omega \mid p(\omega)\} \quad \text{and} \quad B = \{\omega \in \Omega \mid q(\omega)\}$$

Prove that

- $\{\omega \in \Omega \mid p(\omega) \wedge q(\omega)\} = A \cap B$;

- $\{\omega \in \Omega \mid p(\omega) \vee q(\omega)\} = A \cup B$;

- $\{\omega \in \Omega \mid \neg p(\omega)\} = A^c$.

For reference, in Example 8.1.11, we had $\Omega = [6] \times [6]$ and we defined $p(a,b)$ to be '$a+b$ is even' and $q(a,b)$ to be '$a+b \geqslant 7$'. $\triangleleft$

With this in mind, it will be useful to know how set operations on events interact with probabilities. A useful tool in this investigation is that of an *indicator function*.

Definition 8.1.14
Let Ω be a set and let $A \subseteq \Omega$. The **indicator function** of A in Ω is the function $i_A : \Omega \to \{0,1\}$ defined by

$$i_A(\omega) = \begin{cases} 1 & \text{if } \omega \in A \\ 0 & \text{if } \omega \notin A \end{cases}$$

Proposition 8.1.15
Let Ω be a set and let $A, B \subseteq \Omega$. Then for all $\omega \in \Omega$ we have

(i) $i_{A \cap B}(\omega) = i_A(\omega) i_B(\omega)$;

(ii) $i_{A \cup B}(\omega) = i_A(\omega) + i_B(\omega) - i_{A \cap B}(\omega)$; and

(iii) $i_{A^c}(\omega) = 1 - i_A(\omega)$.

Proof of (i)
Let $\omega \in \Omega$. If $\omega \in A \cap B$ then $\omega \in A$ and $\omega \in B$, so that $i_{A \cap B}(\omega) = i_A(\omega) = i_B(\omega) = 1$. Hence

$$i_A(\omega) i_B(\omega) = 1 = i_{A \cap B}(\omega)$$

If $\omega \notin A \cap B$ then either $\omega \notin A$ or $\omega \notin B$. Hence $i_{A \cap B}(\omega) = 0$, and either $i_A(\omega) = 0$ or $i_B(\omega) = 0$. Thus

$$i_A(\omega) i_B(\omega) = 0 = i_{A \cap B}(\omega)$$

In both cases, we have $i_{A \cap B}(\omega) = i_A(\omega) i_B(\omega)$, as required. $\square$

Exercise 8.1.16
Prove parts (ii) and (iii) of Proposition 8.1.15. $\triangleleft$

Exercise 8.1.17
Let $(\Omega, \mathbb{P})$ be a discrete probability space, and for each $\omega \in \Omega$ let $p_\omega = \mathbb{P}(\{\omega\})$.

Prove that, for each event A, we have

$$\mathbb{P}(A) = \sum_{\omega \in \Omega} p_\omega i_A(\omega)$$

◁

Theorem 8.1.18
Let $(\Omega, \mathbb{P})$ be a probability space and let $A, B \subseteq \Omega$. Then

$$\mathbb{P}(A \cup B) = \mathbb{P}(A) + \mathbb{P}(B) - \mathbb{P}(A \cap B)$$

Proof
For each $\omega \in \Omega$, let $p_\omega = \mathbb{P}(\{\omega\})$. Then

$$
\begin{aligned}
\mathbb{P}(A \cup B) &= \sum_{\omega \in \Omega} p_\omega i_{A \cup B}(\omega) && \text{by Exercise 8.1.17}\\
&= \sum_{\omega \in \Omega} p_\omega (i_A(\omega) + i_B(\omega) - i_{A \cap B}(\omega)) && \text{by Proposition 8.1.15}\\
&= \sum_{\omega \in \Omega} p_\omega i_A(\omega) + \sum_{\omega \in \Omega} p_\omega i_B(\omega) + \sum_{\omega \in \Omega} p_\omega i_{A \cap B}(\omega) && \text{rearranging}\\
&= \mathbb{P}(A) + \mathbb{P}(B) - \mathbb{P}(A \cap B) && \text{by Exercise 8.1.17}
\end{aligned}
$$

as required.   □

Although there are nice expressions for unions and complements of events, it is not always the case that intersection of events corresponds with multiplication of probabilities.

Example 8.1.19
Let $\Omega = [3]$ and define a probability measure $\mathbb{P}$ on Ω by letting

$$\mathbb{P}(\{1\}) = \frac{1}{4}, \qquad \mathbb{P}(\{2\}) = \frac{1}{2} \qquad \text{and} \qquad \mathbb{P}(\{3\}) = \frac{1}{4}$$

Then we have

$$\mathbb{P}(\{1,2\} \cap \{2,3\}) = \mathbb{P}(\{2\}) = \frac{1}{2} \neq \frac{9}{16} = \frac{3}{4} \cdot \frac{3}{4} = \mathbb{P}(\{1,2\}) \cdot \mathbb{P}(\{2,3\})$$

◁

This demonstrates that it is not always the case that $\mathbb{P}(A \cap B) = \mathbb{P}(A)\mathbb{P}(B)$ for events A, B in a given probability space. Pairs of events A, B for which this equation *is* true are said to be *independent*.

Definition 8.1.20

Let $(\Omega, \mathbb{P})$ be a probability space and let A, B be events. We say A and B are **independent** if $\mathbb{P}(A \cap B) = \mathbb{P}(A)\mathbb{P}(B)$; otherwise, we say they are **dependent**. More generally, events $A_1, A_2, \ldots, A_n$ are **mutually independent** if

$$\mathbb{P}(A_1 \cap A_2 \cap \cdots \cap A_n) = \mathbb{P}(A_1)\mathbb{P}(A_2) \cdots \mathbb{P}(A_n)$$

Example 8.1.21

A fair six-sided die is rolled twice. Let A be the event that the first roll shows 4, and let B be the event that the second roll is even. Then

$$A = \{(4,1), (4,2), (4,3), (4,4), (4,5), (4,6)\}$$

so $\mathbb{P}(A) = \frac{6}{36} = \frac{1}{6}$; and

$$B = \{(a,2), (a,4), (a,6) \mid a \in [6]\}$$

so $\mathbb{P}(B) = \frac{18}{36} = \frac{1}{2}$. Moreover $A \cap B = \{(4,2), (4,4), (4,6)\}$, so it follows that

$$\mathbb{P}(A \cap B) = \frac{3}{36} = \frac{1}{12} = \frac{1}{6} \cdot \frac{1}{2} = \mathbb{P}(A)\mathbb{P}(B)$$

so the events A and B are independent.

Let C be the event that the sum of the two dice rolls is equal to 5. Then

$$C = \{(1,4), (2,3), (3,2), (4,1)\}$$

so $\mathbb{P}(C) = \frac{4}{36} = \frac{1}{9}$. Moreover $A \cap C = \{(4,1)\}$, so it follows that

$$\mathbb{P}(A \cap C) = \frac{1}{36} \neq \frac{1}{54} = \frac{1}{6} \cdot \frac{1}{9} = \mathbb{P}(A)\mathbb{P}(C)$$

so the events A and C are dependent. ◁

Exercise 8.1.22

Let $(\Omega, \mathbb{P})$ be a probability space. Under what conditions is an event A independent from itself? ◁

Conditional probability

Suppose we model a real-world situation, such as the roll of a die or the flip of a coin, using a probability $(\Omega, \mathbb{P})$. When we receive new information, the situation might change, and we might want to model this new situation by updating our probabilities to reflect the fact that we know that B has occurred. This is done by defining a new probability measure $\widetilde{\mathbb{P}}$ on Ω. What follows is an example of this.

Example 8.1.23

Two cards are drawn at random, in order, without replacement, from a 52-card deck. We can model this situation by letting the sample space Ω be the set of ordered pairs of distinct cards, and letting $\mathbb{P}$ assign an equal probability (of $\frac{1}{|\Omega|}$) to each outcome. Note that $|\Omega| = 52 \cdot 51$, and so

$$\mathbb{P}(\{\omega\}) = \frac{1}{52 \cdot 51}$$

for each outcome ω.

We will compute two probabilities:

- The probability that the second card drawn is a heart.

- The probability that the second card drawn is a heart *given that* the first card drawn is a diamond.

Let $A \subseteq \Omega$ be the event that the second card drawn is a heart, and let $B \subseteq \Omega$ be the event that the first card drawn is a diamond.

To compute $\mathbb{P}(A)$, note first that $A = A' \cup A''$, where

- A' is the event that both cards are hearts, so that $|A'| = 13 \cdot 12$; and

- A'' is the event that only the second card is a heart, so that $|A''| = 39 \cdot 13$.

Since $A' \cap A'' = \varnothing$, it follows from countable additivity that

$$\mathbb{P}(A) = \mathbb{P}(A') + \mathbb{P}(A'') = \frac{13 \cdot 12 + 39 \cdot 13}{52 \cdot 51} = \frac{13 \cdot (12 + 39)}{52 \cdot 51} = \frac{1}{4}$$

Now suppose we know that first card drawn is a diamond—that is, event B has occurred—and we wish to update our probability that A occurs. We do this by defining a new probability measure

$$\widetilde{\mathbb{P}} : \mathscr{P}(\Omega) \to [0, 1]$$

such that:

(a) The outcomes that do not give rise to the event B are assigned probability zero; that is, $\widetilde{\mathbb{P}}(\{\omega\}) = 0$ for all $\omega \notin B$; and

(b) The outcomes that give rise to the event B are assigned probabilities proportional to their old probability; that is, there is some $k \in \mathbb{R}$ such that $\widetilde{\mathbb{P}}(\omega) = k\mathbb{P}(\omega)$ for all $\omega \in B$.

In order for $\widetilde{\mathbb{P}}$ to be a probability measure on Ω, we need condition (i) of Definition 8.1.1 to occur.

$$\begin{aligned}
\widetilde{\mathbb{P}}(\Omega) &= \sum_{\omega \in \Omega} \widetilde{\mathbb{P}}(\{\omega\}) && \text{by condition (ii) of Proposition 8.1.5} \\
&= \sum_{\omega \in B} \widetilde{\mathbb{P}}(\{\omega\}) && \text{since } \widetilde{\mathbb{P}}(\{\omega\}) = 0 \text{ for } \omega \notin B \\
&= \sum_{\omega \in B} k\mathbb{P}(\{\omega\}) && \text{since } \widetilde{\mathbb{P}}(\{\omega\}) = k\mathbb{P}(\{\omega\} \text{ for } \omega \in B \\
&= k\mathbb{P}(B) && \text{by condition (ii) of Proposition 8.1.5}
\end{aligned}$$

Since we need $\widetilde{\mathbb{P}}(\Omega) = 1$, we must therefore take $k = \frac{1}{\mathbb{P}(B)}$. (In particular, we need $\mathbb{P}(B) > 0$ for this notion to be well-defined.)

Recall that, before we knew that the first card was a diamond, the probability that the second card is a heart was $\frac{1}{4}$. We now calculate how this probability changes with the updated information that the first card was a diamond.

The event that the second card is a heart in the new probability space is precisely $A \cap B$, since it is the subset of B consisting of all the outcomes ω giving rise to the event A. As such, the new probability that the second card is a heart is given by

$$\widetilde{\mathbb{P}}(A) = \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)}$$

Now:

- $A \cap B$ is the event that the first card is a diamond and the second is a heart. To specify such an event, we need only specify the ranks of the two cards, so $|A \cap B| = 13 \cdot 13$ and hence $\mathbb{P}(A \cap B) = \frac{13 \cdot 13}{52 \cdot 51}$.

- B is the event that the first card is a diamond. A similar procedure as with A yields $\mathbb{P}(B) = \frac{1}{4}$.

Hence

$$\widetilde{\mathbb{P}}(A) = \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)} = \frac{13 \cdot 13 \cdot 4}{52 \cdot 51} = \frac{13}{51}$$

Thus the knowledge that the first card drawn is a diamond very slightly increases the probability that the second card is a heart from $\frac{1}{4} = \frac{13}{52}$ to $\frac{13}{51}$. ◁

Example 8.1.23 suggests the following schema: upon discovering that an event B occurs, the probability that event A occurs should change from $\mathbb{P}(A)$ to $\frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)}$. This motivates the following definition of *conditional probability*.

Definition 8.1.24

Let $(\Omega, \mathbb{P})$ be a probability space and let A, B be events such that $\mathbb{P}(B) > 0$. The **conditional probability of A given B** is the number $\mathbb{P}(A \mid B)$ (LaTeX code: `\mathbb{P}(A \mid B)`) defined by

$$\mathbb{P}(A \mid B) = \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)}$$

Example 8.1.25

A fair six-sided die is rolled twice. We compute the probability that the first roll showed a 2 given that the sum of the die rolls is less than 5.

We can model this situation by taking the sample space to be $[6] \times [6]$, with each outcome having an equal probability of $\frac{1}{36}$.

Let A be the event that the first die roll shows a 2, that is

$$A = \{(2,1),(2,2),(2,3),(2,4),(2,5),(2,6)\}$$

and let B be the event that the sum of the die rolls is less than 5, that is

$$B = \{(1,1),(1,2),(1,3),(2,1),(2,2),(3,1)\}$$

We need to compute $\mathbb{P}(A \mid B)$. Well,

$$A \cap B = \{(2,1),(2,2)\}$$

so $\mathbb{P}(A \cap B) = \frac{2}{36}$; and $\mathbb{P}(B) = \frac{6}{36}$. Hence

$$\mathbb{P}(A \mid B) = \frac{\frac{2}{36}}{\frac{6}{36}} = \frac{2}{6} = \frac{1}{3}$$

◁

Exercise 8.1.26

A fair six-sided die is rolled three times. What is the probability that the sum of the die rolls is less than or equal to 12, given that each die roll shows a power of 2? ◁

Exercise 8.1.27

Let $(\Omega, \mathbb{P})$ be a probability space and let A, B be events with $\mathbb{P}(B) > 0$. Prove that

$$\mathbb{P}(A \mid B) = \mathbb{P}(A \cap B \mid B)$$

◁

Exercise 8.1.28

Let $(\Omega, \mathbb{P})$ be a probability space and let A, B be events such that $\mathbb{P}(B) > 0$. Prove that $\mathbb{P}(A \mid B) = \mathbb{P}(A)$ if and only if A and B are independent. ◁

We will soon see some useful real-world applications of probability theory using *Bayes's theorem* (Theorem 8.1.33). Before we do so, some technical results will be useful in our proofs.

Proposition 8.1.29

Let $(\Omega, \mathbb{P})$ be a probability space and let A, B be events with $0 < \mathbb{P}(B) < 1$. Then

$$\mathbb{P}(A) = \mathbb{P}(A \mid B)\mathbb{P}(B) + \mathbb{P}(A \mid B^c)\mathbb{P}(B^c)$$

Proof

Note first that we can write

$$A = A \cap \Omega = A \cap (B \cup B^c) = (A \cap B) \cup (A \cap B^c)$$

and moreover the events $A \cap B$ and $A \cap B^c$ are mutually exclusive. Hence

$$\mathbb{P}(A) = \mathbb{P}(A \cap B) + \mathbb{P}(A \cap B^c)$$

by countable additivity. The definition of conditional probability (Definition 8.1.24) then gives

$$\mathbb{P}(A) = \mathbb{P}(A \mid B)\mathbb{P}(B) + \mathbb{P}(A \mid B^c)\mathbb{P}(B^c)$$

as required. □

Example 8.1.30

An animal rescue centre houses a hundred animals, sixty of which are dogs and forty of which are cats. Ten of the dogs and ten of the cats hate humans. We compute the probability that a randomly selected animal hates humans.

Let A be the event that a randomly selected animal hates humans, and let B be the event that the animal is a dog. Note that B^c is precisely the event that the animal is a cat. The information we are given says that:

- $\mathbb{P}(B) = \frac{60}{100}$, since 60 of the 100 animals are dogs;

- $\mathbb{P}(B^c) = \frac{40}{100}$, since 40 of the 100 animals are cats;

- $\mathbb{P}(A \mid B) = \frac{10}{60}$, since 10 of the 60 dogs hate humans;

- $\mathbb{P}(A \mid B^c) = \frac{10}{40}$, since 10 of the 40 cats hate humans.

By Proposition 8.1.29, it follows that the probability that a randomly selected animal hates humans is

$$\mathbb{P}(A) = \mathbb{P}(A \mid B)\mathbb{P}(B) + \mathbb{P}(A \mid B^c)\mathbb{P}(B^c) = \frac{60}{100} \cdot \frac{10}{60} + \frac{40}{100} \cdot \frac{10}{40} = \frac{20}{100} = \frac{1}{5}$$

◁

The following example generalises Proposition 8.1.29 to arbitrary partitions of a sample space into events with positive probabilities.

Example 8.1.31

The animal rescue centre from Example 8.1.30 acquires twenty additional rabbits, of whom sixteen hate humans. We compute the probability that a randomly selected animal hates humans, given the new arrivals.

A randomly selected animal must be either a dog, a cat or a rabbit, and each of these occurs with positive probability. Thus, letting D be the event that the selected animal is a dog, C be the event that the animal is a cat, and R be the event that the animal is a rabbit, we see that the sets D, C, R form a partition of the sample space.

Letting A be the event that the selected animal hates humans. Then

$$\mathbb{P}(A) = \mathbb{P}(A \mid D)\mathbb{P}(D) + \mathbb{P}(A \mid C)\mathbb{P}(C) + \mathbb{P}(A \mid R)\mathbb{P}(R)$$
$$= \frac{10}{60} \cdot \frac{60}{120} + \frac{10}{40} \cdot \frac{40}{120} + \frac{16}{20} \cdot \frac{20}{120}$$
$$= \frac{3}{10}$$

◁

Proposition 8.1.32 below is a technical result which proves that conditional probability truly does yield a new probability measure on a given sample space.

Proposition 8.1.32

Let $(\Omega, \mathbb{P})$ be a probability space and let B be an event such that $\mathbb{P}(B) > 0$. The function $\widetilde{\mathbb{P}} : \mathscr{P}(\Omega) \to [0, 1]$ defined by

$$\widetilde{\mathbb{P}}(A) = \mathbb{P}(A \mid B) \text{ for all } A \subseteq \Omega$$

defines a probability measure on Ω.

Proof
First note that

$$\widetilde{\mathbb{P}}(\Omega) = \mathbb{P}(\Omega \mid B) = \frac{\mathbb{P}(\Omega \cap B)}{\mathbb{P}(B)} = \frac{\mathbb{P}(B)}{\mathbb{P}(B)} = 1$$

so condition (i) of Definition 8.1.1 is satisfied.

Moreover, for each $A \subseteq \Omega$ we have

$$
\begin{aligned}
\widetilde{\mathbb{P}}(A) &= \mathbb{P}(A \mid B) && \text{by definition of } \widetilde{\mathbb{P}} \\
&= \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)} && \text{by Definition 8.1.24} \\
&= \frac{1}{\mathbb{P}(B)} \sum_{\omega \in A \cap B} \mathbb{P}(\{\omega\}) && \text{by Proposition 8.1.5} \\
&= \sum_{\omega \in A \cap B} \mathbb{P}(\{\omega\} \mid B) && \text{by Definition 8.1.24} \\
&= \sum_{\omega \in A} \mathbb{P}(\{\omega\} \mid B) && \text{since } \mathbb{P}(\{\omega\} \mid B) = 0 \text{ for } \omega \in A \setminus B \\
&= \sum_{\omega \in A} \widetilde{\mathbb{P}}(\{\omega\}) && \text{by definition of } \widetilde{\mathbb{P}}
\end{aligned}
$$

so condition (ii) of Proposition 8.1.5 is satisfied. Hence $\widetilde{\mathbb{P}}$ defines a probability measure on Ω. □

Proposition 8.1.32 implies that we can use all the results we've proved about probability measures to conditional probability given a fixed event B. For example, Theorem 8.1.18 implies that

$$\mathbb{P}(A \cup A' \mid B) = \mathbb{P}(A \mid B) + \mathbb{P}(A' \mid B) - \mathbb{P}(A \cap A' \mid B)$$

for all events A, A', B in a probability space $(\Omega, \mathbb{P})$ such that $\mathbb{P}(B) > 0$.

The next theorem we prove has a very short proof, but is extremely important in applied probability theory.

Theorem 8.1.33 (Bayes's theorem)
Let $(\Omega, \mathbb{P})$ be a probability space and let A, B be events with positive probabilities. Then
$$\mathbb{P}(B \mid A) = \frac{\mathbb{P}(A \mid B)\mathbb{P}(B)}{\mathbb{P}(A)}$$

Proof
Definition 8.1.24 gives

$$\mathbb{P}(A \mid B)\mathbb{P}(B) = \mathbb{P}(A \cap B) = \mathbb{P}(B \cap A) = \mathbb{P}(B \mid A)\mathbb{P}(A)$$

Dividing through by $\mathbb{P}(A)$ yields the desired equation. □

As stated, Bayes's theorem is not necessarily particularly enlightening, but its usefulness increases sharply when combined with Proposition 8.1.29 to express the denominator of the fraction in another way.

Corollary 8.1.34

Let $(\Omega, \mathbb{P})$ be a probability space and let A, B be events such that $\mathbb{P}(A) > 0$ and $0 < \mathbb{P}(B) < 1$. Then

$$\mathbb{P}(B \mid A) = \frac{\mathbb{P}(A \mid B)\mathbb{P}(B)}{\mathbb{P}(A \mid B)\mathbb{P}(B) + \mathbb{P}(A \mid B^c)\mathbb{P}(B^c)}$$

Proof

Bayes's theorem tells us that

$$\mathbb{P}(B \mid A) = \frac{\mathbb{P}(A \mid B)\mathbb{P}(B)}{\mathbb{P}(A)}$$

By Proposition 8.1.29 we have

$$\mathbb{P}(A) = \mathbb{P}(A \mid B)\mathbb{P}(B) + \mathbb{P}(A \mid B^c)\mathbb{P}(B^c)$$

Substituting for $\mathbb{P}(A)$ therefore yields

$$\mathbb{P}(B \mid A) = \frac{\mathbb{P}(A \mid B)\mathbb{P}(B)}{\mathbb{P}(A \mid B)\mathbb{P}(B) + \mathbb{P}(A \mid B^c)\mathbb{P}(B^c)}$$

as required.   □

The following example is particularly counterintuitive.

Example 8.1.35

A town has 10000 people, 30 of whom are infected with Disease X. Medical scientists develop a test for Disease X, which is accurate 99% of the time. A person takes the test, which comes back positive. We compute the probability that the person truly is infected with Disease X.

Let A be the event that the person tests positive for Disease X, and let B be the event that the person is infected with Disease X. We need to compute $\mathbb{P}(B \mid A)$.

By Corollary 8.1.34, we have

$$\mathbb{P}(B \mid A) = \frac{\mathbb{P}(A \mid B)\mathbb{P}(B)}{\mathbb{P}(A \mid B)\mathbb{P}(B) + \mathbb{P}(A \mid B^c)\mathbb{P}(B^c)}$$

It remains to compute the individual probabilities on the right-hand side of this equation. Well,

- $\mathbb{P}(A \mid B)$ is the probability that the person tests positive for Disease X, given that they are infected. This is equal to $\frac{99}{100}$, since the test is accurate with probability 99%.

- $\mathbb{P}(A \mid B^c)$ is the probability that the person tests positive for Disease X, given that they are *not* infected. This is equal to $\frac{1}{100}$, since the test is *inaccurate* with probability 1%.

- $\mathbb{P}(B) = \frac{30}{10000}$, since 30 of the 10000 inhabitants are infected with Disease X.

- $\mathbb{P}(B^c) = \frac{9970}{10000}$, since 9970 of the 10000 inhabitants are *not* infected with Disease X.

Piecing this together gives

$$\mathbb{P}(B \mid A) = \frac{\frac{99}{100} \cdot \frac{30}{10000}}{\frac{99}{100} \cdot \frac{30}{10000} + \frac{1}{100} \cdot \frac{9970}{10000}} = \frac{297}{1294} \approx 0.23$$

Remarkably, the probability that the person is infected with Disease X given that the test is positive is only 23%, even though the test is accurate 99% of the time!

$\triangleleft$

The following result generalises Corollary 8.1.34 to arbitrary partitions of the sample space into sets with positive probabilities.

Corollary 8.1.36
Let $(\Omega, \mathbb{P})$ be a probability space, let A be an event with $\mathbb{P}(A) > 0$, and let $\{B_i \mid i \in I\}$ be a family of mutually exclusive events indexed by a countable set I such that

$$\mathbb{P}(B_i) > 0 \text{ for all } i \in I \qquad \text{and} \qquad \bigcup_{i \in I} B_i = \Omega$$

Then

$$\mathbb{P}(B_i \mid A) = \frac{\mathbb{P}(A \mid B_i)\mathbb{P}(B_i)}{\sum\limits_{i \in I} \mathbb{P}(A \mid B_i)\mathbb{P}(B_i)}$$

for each $i \in I$.

Proof

Bayes's theorem tells us that

$$\mathbb{P}(B_i \mid A) = \frac{\mathbb{P}(A \mid B_i)\mathbb{P}(B_i)}{\mathbb{P}(A)}$$

By countable additivity, we have

$$\mathbb{P}(A) = \mathbb{P}\left(\bigcup_{i \in I} A \cap B_i\right) = \sum_{i \in I} \mathbb{P}(A \cap B_i) = \sum_{i \in I} \mathbb{P}(A \mid B_i)\mathbb{P}(B_i)$$

Substituting for $\mathbb{P}(A)$ therefore yields

$$\mathbb{P}(B_i \mid A) = \frac{\mathbb{P}(A \mid B_i)\mathbb{P}(B_i)}{\sum\limits_{i \in I} \mathbb{P}(A \mid B_i)\mathbb{P}(B_i)}$$

as required.   □

Example 8.1.37

A small car manufacturer, *Cars N'At*, makes three models of car: the *Allegheny*, the *Monongahela* and the *Ohio*. It made 3000 Alleghenys, 6500 Monongahelas, and 500 Ohios. In a given day, an Allegheny breaks down with probability $\frac{1}{100}$, a Monongahela breaks down with probability $\frac{1}{200}$, and the notoriously unreliable Ohio breaks down with probability $\frac{1}{20}$. An angry driver calls Cars N'At to complain that their car has broken down. We compute the probability that the driver was driving an Ohio.

Let A be the event that the car is an Allegheny, let B be the event that the car is a Monongahela, and let C be the event that the car is an Ohio. Then

$$\mathbb{P}(A) = \frac{3000}{10000}, \quad \mathbb{P}(B) = \frac{6500}{10000}, \quad \mathbb{P}(C) = \frac{500}{10000}$$

Let D be the event that the car broke down. Then

$$\mathbb{P}(D \mid A) = \frac{1}{100}, \quad \mathbb{P}(D \mid B) = \frac{1}{200}, \quad \mathbb{P}(D \mid C) = \frac{1}{20}$$

We need to compute $\mathbb{P}(C \mid D)$. Since the events A, B, C partition the sample space and have positive probabilities, we can use Corollary 8.1.36, which tells us that

$$\mathbb{P}(C \mid D) = \frac{\mathbb{P}(D \mid C)\mathbb{P}(C)}{\mathbb{P}(D \mid A)\mathbb{P}(A) + \mathbb{P}(D \mid B)\mathbb{P}(B) + \mathbb{P}(D \mid C)\mathbb{P}(C)}$$

Substituting the probabilities that we computed above, it follows that

$$\mathbb{P}(C \mid D) = \frac{\frac{1}{20} \cdot \frac{500}{10000}}{\frac{1}{100} \cdot \frac{3000}{10000} + \frac{1}{200} \cdot \frac{6500}{10000} + \frac{1}{20} \cdot \frac{500}{10000}} = \frac{2}{7} \approx 0.29$$

◁

Exercise 8.1.38

In Example 8.1.37, find the probabilities that the car was an Allegheny and that the car was a Monongahela.   ◁

Section 8.2
Discrete random variables

Events in a probability space are sometimes unenlightening when looked at in isolation. For example, suppose we roll a fair six-sided die twice. The outcomes are elements of the set $[6] \times [6]$, each occurring with equal probability $\frac{1}{36}$. The event that the die rolls sum to 7 is precisely the subset

$$\{(1,6),(2,5),(3,4),(4,3),(5,2),(6,1)\} \subseteq [6] \times [6]$$

and so we can say that the probability that the two rolls sum to 7 is

$$\mathbb{P}(\{(1,6),(2,5),(3,4),(4,3),(5,2),(6,1)\}) = \frac{1}{6}$$

However, it is not at all clear from the expression $\{(1,6),(2,5),(3,4),(4,3),(5,2),($ that, when we wrote it down, what we had in mind was the event that the sum of the die rolls is 7. Moreover, the expression of the event in this way does not make it clear how to generalise to other possible sums of die rolls.

Note that the sum of the die rolls defines a function $S : [6] \times [6] \to [12]$, defined by

$$S(a,b) = a+b \text{ for all } (a,b) \in [6] \times [6]$$

The function S allows us to express our event in a more enlightening way: indeed,

$$\{(1,6),(2,5),(3,4),(4,3),(5,2),(6,1)\} = \{(a,b) \in [6] \times [6] \mid a+b=7\} = S^{-1}[\{7\}]$$

(Recall the definition of *preimage* in Definition 2.2.30.) Thus the probability that the sum of the two die rolls is 7 is equal to $\mathbb{P}(S^{-1}[\{7\}])$.

If we think of S not as a function $[6] \times [6] \to [12]$, but as a $[12]$-valued *random variable*, which varies according to a random outcome in $[6] \times [6]$, then we can informally say

$$\mathbb{P}\{S = 7\} = \frac{1}{6} \quad \text{which formally means} \quad \mathbb{P}(S^{-1}[\{7\}]) = \frac{1}{6}$$

This affords us much more generality. Indeed, we could ask what the probability is that the die rolls sum to a value greater than or equal to 7. In this case, note that the die rolls (a,b) sum to a number greater than or equal to 7 if and only if $a+b \in \{7,8,9,10,11,12\}$, which occurs if and only if $(a,b) \in S^{-1}[\{7,8,9,10,11,12\}]$. Thus, we might informally say

$$\mathbb{P}\{S \geqslant 7\} = \frac{7}{12} \quad \text{which formally means} \quad \mathbb{P}(S^{-1}[\{7,8,9,10,11,12\}]) = \frac{7}{12}$$

We might also ask what the probability is that the sum of the die rolls is prime. In this case, we might informally say

$$\mathbb{P}\{S \text{ is prime}\} = \frac{5}{12} \quad \text{which formally means} \quad \mathbb{P}(S^{-1}[\{2,3,5,7,11\}]) = \frac{5}{12}$$

and so on. In each of these cases, we're defining events—which are subsets of the sample space—in terms of conditions on the values of a random variable (which is, formally, a function).

We make the above intuition formal in Definition 8.2.1.

Definition 8.2.1
Let $(\Omega, \mathbb{P})$ be a probability space and let E be a set. An E-**valued random variable** on $(\Omega, \mathbb{P})$ is a function $X : \Omega \to E$ such that the image

$$X[\Omega] = \{X(\omega) \mid \omega \in \Omega\}$$

is countable. The set E is called the **state space** of X. A random variable with countable state space is called a **discrete random variable**.

Before we proceed with examples, some notation for events regarding values of random variables will be particularly useful.

Notation 8.2.2

Let $(\Omega, \mathbb{P})$ be a probability space, let E be a set and let X be an E-valued random variable on $(\Omega, \mathbb{P})$. For each $e \in E$, write

$$\{X = e\} = \{\omega \in \Omega \mid X(\omega) = e\} = X^{-1}[\{e\}]$$

to denote the event that X takes the value e. More generally, for each logical formula $p(x)$ with free variable x ranging over E, we write

$$\{p(X)\} = \{\omega \in \Omega \mid p(X(\omega))\} = X^{-1}[\{e \in E \mid p(e)\}]$$

for the event that the value of X satisfies $p(x)$.

We will usually write $\mathbb{P}\{X = e\}$ instead of $\mathbb{P}(\{X = e\})$ for the probability that a random variable X takes a value e, and so on.

Example 8.2.3

We can model a sequence of three coin flips using the probability space $(\Omega, \mathbb{P})$, where $\Omega = \{H, T\}^3$ and $\mathbb{P}(\{\omega\}) = \frac{1}{8}$ for all $\omega \in \Omega$.

Let N be the real-valued random variable representing number of heads that show. This is formalised as a function

$$N : \Omega \to \mathbb{R} \quad \text{where} \quad N(i_1, i_2, i_3) = \text{the number of heads amongst } i_1, i_2, i_3$$

for example, $N(\mathrm{H,T,H}) = 2$. Now

- The probability that exactly two heads show is

$$
\begin{aligned}
\mathbb{P}\{N=2\} &= \mathbb{P}(N^{-1}[\{2\}]) && \text{by Notation 8.2.2}\\
&= \mathbb{P}(\{(\mathrm{H,H,T}),(\mathrm{H,T,H}),(\mathrm{T,H,H})\}) && \text{evaluating event } N^{-1}[\{2\}]\\
&= \frac{3}{2^3} = \frac{3}{8}
\end{aligned}
$$

- The probability that at least two heads show is

$$
\begin{aligned}
\mathbb{P}\{N \geqslant 2\} &= \mathbb{P}(\{\omega \in \Omega \mid N(\omega) \geqslant 2\}) && \text{by Notation 8.2.2}\\
&= \mathbb{P}\left(\left\{\begin{array}{ll}(\mathrm{H,H,T}), & (\mathrm{H,T,H}),\\ (\mathrm{T,H,H}), & (\mathrm{H,H,H})\end{array}\right\}\right) && \text{evaluating event}\\
&= \frac{4}{2^3} = \frac{1}{2}
\end{aligned}
$$

◁

Exercise 8.2.4
With probability space $(\Omega, \mathbb{P})$ and random variable N defined as in Example 8.2.3, compute $\mathbb{P}\{N \text{ is odd}\}$ and $\mathbb{P}\{N=4\}$. ◁

Each random variable comes with an associated *probability mass function*, which allows us to 'forget' the underlying probability space for the purposes of studying only the random variable.

Definition 8.2.5
Let $(\Omega, \mathbb{P})$ be a probability space, let $X : \Omega \to E$ be an E-valued random variable. The **probability mass function** of X is the function $f_X : E \to [0,1]$ defined by

$$
f_X(e) = \mathbb{P}\{X = e\} \text{ for all } e \in E
$$

Example 8.2.6
The probability mass function of the random variable N from Example 8.2.3 is the function $f_N : \mathbb{R} \to [0,1]$ defined by

$$
f_N(e) = \mathbb{P}\{N=e\} = \frac{1}{8}\binom{3}{e}
$$

for all $e \in \{0,1,2,3\}$, and $f_N(e) = 0$ otherwise. Indeed, there are $2^3 = 8$ possible outcomes, each equally likely, and $\binom{3}{e}$ of those outcomes show exactly e heads for $e \in \{0,1,2,3\}$. ◁

Exercise 8.2.7

Let $(\Omega, \mathbb{P})$ be a probability space, let E be a set, let X be an E-valued random variable and let $U \subseteq E$. Prove that the event $\{X \in U\}$ is equal to the preimage $X^{-1}[U]$. Deduce that

$$\mathbb{P}\{X \in U\} = \sum_{e \in U} f_X(e)$$

◁

In Example 8.2.6, we could have just specified the value of f_N on $\{0, 1, 2, 3\}$, with the understanding that N does not take values outside of this set and hence that $\mathbb{P}\{N = e\} = 0$ for all $e \notin \{0, 1, 2, 3\}$. This issue arises frequently when dealing with real-valued discrete random variables, and it will be useful to ignore most (or all) of those real numbers which are not values of the random variable.

As such, for $E \subseteq \mathbb{R}$, we will from now on blur the distinction between the following concepts:

  (i) E-valued random variables;

  (ii) Real-valued random variables X such that $\mathbb{P}\{X = x\} = 0$ for all $x \notin E$.

Example 8.2.8

The probability mass function of the random variable N from Example 8.2.3 can be taken to be the function $f_X : \{0, 1, 2, 3\} \to [0, 1]$ defined by

$$f_X(k) = \frac{1}{8}\binom{3}{k} \quad \text{for all } k \in \{0, 1, 2, 3\}$$

◁

Lemma 8.2.9

Let $(\Omega, \mathbb{P})$ be a probability space, let E be a set and let X be an E-valued random variable. The events $\{X = e\}$ for $e \in E$ are mutually exclusive, and their union is Ω.

Proof
If $e, e' \in E$, then for all $\omega \in \Omega$ we have

$$\omega \in \{X = e\} \cap \{X = e'\} \Leftrightarrow \omega \in X^{-1}[\{e\}] \cap X^{-1}[\{e'\}] \quad \text{by Notation 8.2.2}$$
$$\Leftrightarrow X(\omega) = e \text{ and } X(\omega) = e' \quad \text{by definition of preimage}$$
$$\Rightarrow e = e'$$

so if $e \neq e'$ then $\{X = e\} \cap \{X = e'\} = \varnothing$. So the events are mutually exclusive.

Moreover, if $\omega \in \Omega$, then $\omega \in \{X = X(\omega)\}$. Hence

$$\Omega = \bigcup_{e \in E} \{X = e\}$$

as required. □

Theorem 8.2.10

Let $(\Omega, \mathbb{P})$ be a probability space, let E be a countable set, and let X be an E-valued random variable. Then

$$\sum_{e \in E} f_X(e) = 1$$

Proof
Since $f_X(e) = \mathbb{P}\{X = e\}$ for all $e \in E$, we need to check that

$$\sum_{e \in E} \mathbb{P}\{X = e\} = 1$$

By Lemma 8.2.9, we have

$$\sum_{e \in E} \mathbb{P}\{X = e\} = \mathbb{P}\left(\bigcup_{e \in E}\{X = e\}\right) = \mathbb{P}(\Omega) = 1$$

as required. □

The following corollary follows immediately.

Corollary 8.2.11

Let $(\Omega, \mathbb{P})$ be a probability space, let E be a countable set, and let X be an E-valued random variable. The function $X_*\mathbb{P} : \mathscr{P}(E) \to [0, 1]$ defined by

$$(X_*\mathbb{P})(A) = \sum_{e \in A} f_X(e) = \mathbb{P}\{X \in A\}$$

for all $A \subseteq E$ defines a probability measure on E. The space $(E, X_*\mathbb{P})$ is called the **pushforward probability measure** of X. □

Corollary 8.2.11 implies that any statement about probability measures can be applied to the pushforward measure. For example,

$$\mathbb{P}\{X \in A \cup B\} = \mathbb{P}\{X \in A\} + \mathbb{P}\{X \in B\} - \mathbb{P}\{X \in A \cap B\}$$

for all subsets $A, B \subseteq E$.

As with events, there is a notion of independence for random variables.

Definition 8.2.12

Let $(\Omega, \mathbb{P})$ be a discrete probability space and let $X, Y : \Omega \to E$ be discrete random variables on $(\Omega, \mathbb{P})$. We say X and Y are **independent** if, for all $e, e' \in E$, the events $\{X = e\}$ and $\{Y = e'\}$ are independent. More generally, random variables $X_1, X_2, \ldots, X_n$ are **mutually independent** if, for all $e_1, e_2, \ldots, e_n \in E$, the events $\{X_i = e_i\}$ are mutually independent.

Example 8.2.13

A fair six-sided die is rolled twice. Let X be the value shown on the first roll and Y be the value shown on the second roll.

We can model this situation by letting $\Omega = [6] \times [6]$ with $\mathbb{P}(\{(a,b)\}) = \frac{1}{36}$ for all $(a,b) \in \Omega$. The random variables X,Y can thus be taken to be functions $\Omega \to [6]$ defined by

$$X(a,b) = a \text{ and } Y(a,b) = b \text{ for all } (a,b) \in \Omega$$

So let $e, e' \in [6]$. Note first that

$$\begin{aligned}
&\{X = e\} \cap \{Y = e'\} \\
&= \{(a,b) \in \Omega \mid a = e\} \cap \{(a,b) \in \Omega \mid b = e'\} \quad \text{by Notation 8.2.2} \\
&= \{(a,b) \in \Omega \mid a = e \text{ and } b = e'\} \\
&= \{(e,e')\}
\end{aligned}$$

Hence

$$\mathbb{P}(\{X = e\} \cap \{Y = e'\}) = \mathbb{P}(\{(e,e')\}) = \frac{1}{36} = \frac{1}{6} \cdot \frac{1}{6} = \mathbb{P}\{X = e\}\mathbb{P}\{Y = e'\}$$

The events $\{X = e\}$ and $\{Y = e'\}$ are independent, and so X and Y are independent. $\triangleleft$

Exercise 8.2.14

A coin which shows heads with probability $p \in [0,1]$, and tails otherwise, is flipped five times. For each $i \in [5]$, let

$$X_i = \begin{cases} 0 & \text{if the } i^{\text{th}} \text{ flip shows tails} \\ 1 & \text{if the } i^{\text{th}} \text{ flip shows heads} \end{cases}$$

Prove that the random variables X_1, X_2, X_3, X_4, X_5 are mutually independent. $\triangleleft$

One final technicality that we mention before continuing concerns performing arithmetic with random variables which assume real values.

Notation 8.2.15

Let $(\Omega, \mathbb{P})$ be a probability space, and let X, Y be real-valued random variables on $(\Omega, \mathbb{P})$. Then we can define a new real-valued random variable $X + Y$ by

$$(X + Y)(\omega) = X(\omega) + Y(\omega) \text{ for all } \omega \in \Omega$$

Likewise for multipication, scalar multiplication and constants: for each $\omega \in \Omega$, define

$$(XY)(\omega) = X(\omega)Y(\omega), \quad (aX)(\omega) = a \cdot X(\omega), \quad a(\omega) = a$$

where $a \in \mathbb{R}$. Note that the random variables $X + Y, XY, aX, a$ are all supported on a countable set.

Example 8.2.16

A coin which shows heads with probability $p \in [0,1]$, and tails otherwise, is flipped n times. For each $i \in [n]$, let

$$X_i = \begin{cases} 0 & \text{if the } i^{\text{th}} \text{ flip shows tails} \\ 1 & \text{if the } i^{\text{th}} \text{ flip shows heads} \end{cases}$$

Then each X_i is a $\{0,1\}$-valued random variable.

Define $X = X_1 + X_2 + \cdots + X_n$. Then X is a $\{0,1,\ldots,n\}$-valued random variable representing the number of heads that show in total after the coin is flipped n times. ◁

Probability distributions

Most of the random variables we are interested in are characterised by one of a few *probability distributions*. We won't define the term 'probability distribution' precisely—indeed, its use in the mathematical literature is often ambiguous and informal—instead, we will take it to mean any description of the random behaviour of a probability space or random variable.

The *uniform distribution* models the real-world situation in which any of a fixed number of outcomes occurs with equal probability.

Definition 8.2.17 (Uniform distribution)

Let $(\Omega, \mathbb{P})$ be a probability space, let E be a finite set, and let $X : \Omega \to E$ be a random variable. We say X follows the **uniform distribution on E**, or X is **uniformly distributed on E**, if f_X is constant—that is, if

$$f_X(e) = \frac{1}{|E|} \text{ for all } e \in E$$

If X is uniformly distributed on E, we write $X \sim \text{Unif}(E)$ (LaTeX code: `\sim`).

Example 8.2.18

Let $(\Omega, \mathbb{P})$ be the probability space modelling the roll of a fair six-sided die, and let X be the $[6]$-valued random variable representing the number shown. Then for each $k \in [6]$ we have

$$f_X(k) = \mathbb{P}\{X = k\} = \mathbb{P}(\{k\}) = \frac{1}{6}$$

so X is uniformly distributed on $[6]$. ◁

Exercise 8.2.19

Let $(\Omega, \mathbb{P})$ be the probability space modelling the roll of a fair six-sided die, and let X be the $\{0, 1\}$-valued random variable which is equal to 0 if the die shows an even number and 1 if the die shows an odd number. Prove that $X \sim \mathrm{Unif}(\{0, 1\})$.

◁

Before we continue, we prove that the notion of 'uniform distribution' does not make sense for countably infinite sets.

Theorem 8.2.20

Let $(\Omega, \mathbb{P})$ be a probability space and let E be a countably infinite set. There is no notion of a uniformly E-valued random variable X—that is, there is no $p \in [0, 1]$ such that $f_X(e) = p$ for all $e \in E$.

Proof

We may assume $E = \mathbb{N}$; otherwise, re-index the sums accordingly.

Let $p \in [0, 1]$. Note that

$$\sum_{n \in \mathbb{N}} f_X(n) = \sum_{n \in \mathbb{N}} p = \lim_{N \to \infty} \sum_{n=0}^{N} p = \lim_{N \to \infty} (N+1)p$$

If $p = 0$ then

$$\lim_{N \to \infty} (N+1)p = \lim_{N \to \infty} 0 = 0$$

If $p > 0$ then, for all $K > 0$, letting $N = \frac{K}{p}$ yields $(N+1)p = K + p > K$, and hence

$$\lim_{N \to \infty} (N+1)p = \infty$$

Thus $\sum_{n \in \mathbb{N}} p \neq 1$ for all $p \in [0, 1]$.

In both cases, we have contradicted Theorem 8.2.10. As such, there can be no random variable $X : \Omega \to \mathbb{N}$ such that f_X is constant. □

The *Bernoulli distribution* models real-world situations in which one of two outcomes occurs, but not necessarily with the same probability.

Definition 8.2.21 (Bernoulli distribution)

Let $(\Omega, \mathbb{P})$ be a probability space. A $\{0,1\}$-valued random variable X follows the **Bernoulli distribution with parameter** p if its probability mass function $f_X : \{0,1\} \to [0,1]$ satisfies

$$f_X(0) = 1 - p \quad \text{and} \quad f_X(1) = p$$

If X follows the Bernoulli distribution with parameter p, we write $X \sim \mathrm{B}(1, p)$.

The reason behind the notation $\mathrm{B}(1, p)$ will become clear soon—the Bernoulli distribution is a specific instance of a more general distribution, which we will see in Definition 8.2.24.

Example 8.2.22

A coin shows 'heads' with probability p and 'tails' with probability $1 - p$. Let X be the random variable which takes the value 0 if the coin shows tails and 1 if the coin shows heads. Then $X \sim \mathrm{B}(1, p)$. ◁

Exercise 8.2.23

Let X be a $\{0,1\}$-valued random variable. Prove that $X \sim \mathrm{U}(\{0,1\})$ if and only if $X \sim \mathrm{B}(1, \frac{1}{2})$. ◁

Suppose that, instead of flipping a coin just once, as in Example 8.2.22, you flip it n times. The total number of heads that show must be an element of $\{0, 1, \ldots, n\}$, and each such element occurs with some positive probability. The resulting probability distribution is called the *binomial distribution*.

Definition 8.2.24 (Binomial distribution)

Let $(\Omega, \mathbb{P})$ be a probability space. A $\{0, 1, \ldots, n\}$-valued random variable X follows the **binomial distribution with parameters** n, p if its probability mass function $f_X : \{0, 1, \ldots, n\} \to [0,1]$ satisfies

$$f_X(k) = \binom{n}{k} p^k (1-p)^{n-k}$$

for all $k \in \{0, 1, \ldots, n\}$. If X follows the binomial distribution with parameters n, p, we write $X \sim \mathrm{B}(n, p)$.

Example 8.2.25

A coin which shows heads with probability $p \in [0,1]$, and tails otherwise, is flipped n times. We will prove that the number of heads that show is binomially distributed.

We can model this situation with probability space $(\Omega, \mathbb{P})$ defined by taking $\Omega =$

$\{H, T\}^n$, and letting $\mathbb{P}(\{\omega\}) = p^h(1-p)^t$ for all $\omega \in \Omega$, where h is the number of heads that show and t is the number of tails that show in outcome ω. For example, if $n = 5$ then

$$\mathbb{P}(\{HTHHT\}) = p^3(1-p)^2 \quad \text{and} \quad \mathbb{P}(\{TTTTT\}) = (1-p)^5$$

Note in particular that $h + t = n$.

Let X be the random variable which counts the number of heads that show. Formally, we can define $X : \{H, T\}^n \to \{0, 1, \ldots, n\}$ by letting $X(\omega)$ be the number of heads that show in outcome ω. For example if $n = 5$ then

$$X(HTHHT) = 3 \quad \text{and} \quad X(TTTTT) = 0$$

The event $\{X = k\}$ is the set of n-tuples of 'H's and 'T's which contain exactly k 'H'. Hence $|\{X = k\}| = \binom{n}{k}$, since such an n-tuple can be specified by choosing the k positions of the 'H's, and putting 'T's in the remaining positions. Since each outcome in this event occurs with equal probability $p^k(1-p)^{n-k}$, it follows that

$$f_X(k) = \binom{n}{k} p^k(1-p)^{n-k}$$

for all $k \in \{0, 1, \ldots, n\}$. Hence $X \sim B(n, p)$. $\triangleleft$

The following theorem proves that the sum of Bernoulli random variables follows the binomial distribution.

Theorem 8.2.26

Let $(\Omega, \mathbb{P})$ be a probability space, let $p \in [0, 1]$ and let $X_1, X_2, \ldots, X_n : \Omega \to \{0, 1\}$ be independent random variables such that $X_i \sim B(1, p)$. Then

$$X_1 + X_2 + \cdots + X_n \sim B(n, p)$$

Proof

Let $X = X_1 + X_2 + \cdots + X_n$. For each outcome ω and each $k \in \{0, 1, \ldots, n\}$, we have $X(\omega) = k$ if and only if exactly k of the values $X_1(\omega), X_2(\omega), \ldots, X_n(\omega)$ are equal to 1.

For each specification S of *which* of the random variables X_i is equal to 1, let $A_S \subseteq \Omega$ be the event that this occurs. Formally, this is to say that, for each $S \subseteq [n]$, we define

$$A_S = \{\omega \in \Omega \mid X_i(\omega) = 0 \text{ for all } i \notin S \text{ and } X_i(\omega) = 1 \text{ for all } i \in S\}$$

Then $\mathbb{P}(A_S) = p^k(1-p)^{n-k}$, since the random variables $X_1, X_2, \ldots, X_n$ are mutually independent.

As argued above sets $\{A_S \mid U \subseteq [n], |S| = k\}$ form a partition of $\{X = k\}$, and hence

$$f_X(k) = \sum_{S \in \binom{[n]}{k}} \mathbb{P}(A_S) = \sum_{S \in \binom{[n]}{k}} p^k(1-p)^{n-k} = \binom{n}{k}p^k(1-p)^{n-k}$$

which is to say that $X \sim B(n,p)$. $\qquad\qquad\square$

We will make heavy use of Theorem 8.2.26 when we will study the *expectation* of binomially distributed random variables (Definition 8.2.34). First, let's will look at a couple of scenarios in which a binomially distributed random variable is expressed as a sum of independent Bernoulli random variables.

Example 8.2.27
In Example 8.2.25, we could have defined $\{0,1\}$-valued random variables $X_1, X_2, \ldots, X_n$ by letting

$$X_i(\omega) = \begin{cases} 0 & \text{if the } i^{\text{th}} \text{ coin flip shows tails} \\ 1 & \text{if the } i^{\text{th}} \text{ coin flip shows heads} \end{cases}$$

Then the number of heads shown in total is the random variable $X = X_1 + X_2 + \cdots + X_n$. Note that each random variable X_i follows the Bernoulli distribution with parameter p, and they are independent, so that $X \sim B(n,p)$ by Theorem 8.2.26.
$\qquad\qquad\triangleleft$

In Example 8.2.25, we flipped a coin a fixed number of times and counted how many heads showed. Now suppose that we flip a coin repeatedly until heads show, and then stop. The number of times the coin was flipped before heads shows could, theoretically, be any natural number. This situation is modelled by the *geometric distribution*.

Definition 8.2.28 (Geometric distribution on $\mathbb{N}$)
Let $(\Omega, \mathbb{P})$ be a probability space. An $\mathbb{N}$-valued random variable X follows the **geometric distribution with parameter** p if its probability mass function $f_X : \mathbb{N} \to [0,1]$ satisfies
$$f_X(k) = (1-p)^k p \text{ for all } k \in \mathbb{N}$$
If X follows the geometric distribution with parameter p, we write $X \sim \text{Geom}(p)$.

Example 8.2.29
A coin which shows heads with probability $p \in [0,1]$, and tails otherwise, is flipped repeatedly until heads shows. $\qquad\qquad\triangleleft$

Exercise 8.2.30

Let $p \in [0,1]$ and let $X \sim \text{Geom}(p)$. Prove that

$$\mathbb{P}\{X \text{ is even}\} = \frac{1}{2-p}$$

What is the probability that X is odd?   ◁

Occasionally, it will be useful to consider geometrically distributed random variables which are valued in the set

$$\mathbb{N}^+ = \{1,2,3,\dots\}$$

of all *positive* natural numbers. The probability mass function of such a random variable is slightly different.

Definition 8.2.31 (Geometric distribution on $\mathbb{N}^+$)

Let $(\Omega, \mathbb{P})$ be a probability space. An $\mathbb{N}^+$-valued random variable X follows the **geometric distribution with parameter** p if its probability mass function $f_X : \mathbb{N}^+ \to [0,1]$ satisfies

$$f_X(k) = (1-p)^{k-1}p \text{ for all } k \in \mathbb{N}^+$$

If X follows the geometric distribution with parameter p, we write $X \sim \text{Geom}(p)$.

It is to be understood from context whether a given geometric random variable is $\mathbb{N}$-valued or $\mathbb{N}^+$-valued.

Example 8.2.32

An urn contains $n \geqslant 1$ distinct coupons. Each time you draw a coupon that you have not drawn before, you get a stamp. When you get all n stamps, you win. Let X be the number of coupons drawn up to, and including, a winning draw.

For each $k \in [n]$, let X_k be the random variable representing the number of draws required to draw the k^{th} new coupon, after $k-1$ coupons have been collected. Then the total number of times a coupon must be drawn is $X = X_1 + X_2 + \cdots + X_n$.

After $k-1$ coupons have been collected, there are $n-k+1$ uncollected coupons remaining in the urn, and hence on any given draw, an uncollected coupon is drawn with probability $\frac{n-k+1}{n}$, and a coupon that has already been collected is drawn with probability $\frac{k-1}{n}$. Hence for each $r \in \mathbb{N}^+$ we have

$$\mathbb{P}[X_k = r] = \left(\frac{k-1}{n}\right)^{r-1} \left(\frac{n-k+1}{n}\right)$$

That is to say, X_k is geometrically distributed on $\mathbb{N}^+$ with parameter $\frac{n-k+1}{n}$.

We will use this in Example 8.2.47 to compute the number of times a person should expect to have to draw coupons from the urn until they win. ◁

Expectation

We motivate the definition of *expectation* (Definition 8.2.34) with the following example.

Example 8.2.33

For each $n \geqslant 1$, let X_n be the average value shown when a fair six-sided die is rolled n times.

When n is small, the value of X_n is somewhat unpredictable. For example, X_1 is uniformly distributed, since it takes each of the values $1, 2, 3, 4, 5, 6$ with equal probability. This is summarised in the following table:

e	1	2	3	4	5	6
$\mathbb{P}\{X_1 = e\}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$

The distribution of X_2 is shown in the following table:

e	1	1.5	2	2.5	3	3.5	4	4.5	5	5.5	6
$\mathbb{P}\{X_2 = e\}$	$\frac{1}{36}$	$\frac{2}{36}$	$\frac{3}{36}$	$\frac{4}{36}$	$\frac{5}{36}$	$\frac{6}{36}$	$\frac{5}{36}$	$\frac{4}{36}$	$\frac{3}{36}$	$\frac{2}{36}$	$\frac{1}{36}$

As can be seen, the probabilities increase towards the middle of the table; the extreme values occur with low probability. This effect is exaggerated as n increases. Indeed,

$$\mathbb{P}\{X_n = 1\} = \mathbb{P}\{X_n = 6\} = \frac{1}{6^n}$$

which is extremely small when n is large; however, it can be shown that for all $\varepsilon > 0$, we have

$$\mathbb{P}\{3.5 - \varepsilon < X_n < 3.5 + \varepsilon\} \to 1$$

Thus when we roll a die repeatedly, we can expect its value to approach 3.5 with arbitrary precision. This is an instance of a theorem called the *law of large numbers*, which we will not prove here. ◁

The value 3.5 in Example 8.2.33 is special because it is the average of the numbers $1, 2, 3, 4, 5, 6$. More generally, assignments of different probabilities to different values of a random variable X yields a *weighted average* of the possible

values. This weighted average, known as the *expectation* of the random variable, behaves in the same way as the number 3.5 did in Example 8.2.33.

Definition 8.2.34

Let $(\Omega, \mathbb{P})$ be a probability space, let $E \subseteq \mathbb{R}$ be countable, and let X be an E-valued random variable on $(\Omega, \mathbb{P})$. The **expectation** (or **expected value**) of X, if it exists, is the real number $\mathbb{E}[X]$ (LaTeX code: `\mathbb{E}`) defined by

$$\mathbb{E}[X] = \sum_{e \in E} e f_X(e)$$

Example 8.2.35

Let X be a random variable representing the value shown when a fair six-sided die is rolled. Then $X \sim \mathrm{U}([6])$, so that $f_X(k) = \frac{1}{6}$ for all $k \in [6]$, and hence

$$\mathbb{E}[X] = 1 \cdot \frac{1}{6} + 2 \cdot \frac{1}{6} + 3 \cdot \frac{1}{6} + 4 \cdot \frac{1}{6} + 5 \cdot \frac{1}{6} + 6 \cdot \frac{1}{6} = \frac{21}{6} = 3.5$$

so the expected value of the die roll is 3.5.   ◁

Example 8.2.36

Let $p \in [0, 1]$ and let $X \sim \mathrm{B}(1, p)$. Then

$$\mathbb{E}[X] = 0 \cdot (1 - p) + 1 \cdot p = p$$

So the expected value of a Bernoulli random variable is equal to the parameter.

◁

Exercise 8.2.37

Let $(\Omega, \mathbb{P})$ be a probability space and let $c \in \mathbb{R}$. Thinking of c as a *constant* real-valued random variable,[a] prove that $\mathbb{E}[c] = c$.   ◁

The following lemma provides an alternative method for computing the expectation of a random variable. It will be useful for proving that expectation is *linear* in Theorem 8.2.43.

Lemma 8.2.38

Let $(\Omega, \mathbb{P})$ be a probability space, let E be a countable set and let X be an E-valued random variable on $(\Omega, \mathbb{P})$. Then

$$\mathbb{E}[X] = \sum_{\omega \in \Omega} X(\omega) \mathbb{P}(\{\omega\})$$

Proof

Recall from Lemma 8.2.9 that

$$\Omega = \bigcup_{e \in E} \{X = e\}$$

---

[a]Formally, we should define $X : \Omega \to \mathbb{R}$ by letting $X(\omega) = c$ for all $\omega \in \Omega$; then compute $\mathbb{E}[X]$.

and the events $\{X = e\}$ are mutually exclusive. Hence

$$
\begin{aligned}
\sum_{\omega \in \Omega} X(\omega)\mathbb{P}(\{\omega\}) &= \sum_{e \in E} \sum_{\omega \in \{X = e\}} X(\omega)\mathbb{P}(\{\omega\}) \quad \text{by Lemma 8.2.9} \\
&= \sum_{e \in E} e\mathbb{P}\{X = e\} \quad \text{by (ii) in Proposition 8.1.5} \\
&= \sum_{e \in E} e f_X(e) \quad \text{by Definition 8.2.5}
\end{aligned}
$$

as required. $\qquad\square$

Proposition 8.2.39

Let $n \in \mathbb{N}$ and $p \in [0, 1]$, and suppose that X is a random variable such that $X \sim B(n, p)$. Then $\mathbb{E}[X] = np$.

Proof

Since $X \sim B(n, p)$, we have $f_X(k) = \binom{n}{k} p^k (1 - p)^{n-k}$ for all $0 \leqslant k \leqslant n$. Hence

$$
\begin{aligned}
\mathbb{E}[X] &= \sum_{k=0}^{n} k \cdot \binom{n}{k} p^k (1 - p)^{n-k} \quad \text{by definition of expectation} \\
&= \sum_{k=1}^{n} k \cdot \binom{n}{k} p^k (1 - p)^{n-k} \quad \text{since the } k = 0 \text{ term is zero} \\
&= \sum_{k=1}^{n} n \binom{n-1}{k-1} p^k (1 - p)^{n-k} \quad \text{by Proposition 3.3.36} \\
&= \sum_{\ell=0}^{n-1} n \binom{n-1}{\ell} p^{\ell+1} (1 - p)^{(n-1)-\ell} \quad \text{writing } \ell = k+1 \\
&= np \cdot \sum_{\ell=0}^{n-1} \binom{n-1}{\ell} p^\ell (1 - p)^{(n-1)-\ell} \quad \text{pulling out constant factors} \\
&= np(p + (1 - p))^{n-1} \quad \text{by the binomial theorem} \\
&= np \quad \text{since } p + (1 - p) = 1
\end{aligned}
$$

as required. $\qquad\square$

Example 8.2.40

A coin which shows heads with probability $\frac{1}{3}$, and tails otherwise, is tossed 12 times. Letting X be the random variable represent the number of heads that show, we see that $X \sim B(12, \frac{1}{3})$, and hence the expected number of heads that show is equal to

$$
\mathbb{E}[X] = 12 \cdot \frac{1}{3} = 4
$$

$\triangleleft$

Exercise 8.2.41

Use Proposition 7.3.3 to prove that the expectation of a $\mathbb{N}$-valued random variable

which is geometrically distributed with parameter $p \in [0,1]$ is equal to $\frac{1-p}{p}$. Use this to compute the expected number of times a coin must be flipped before the first time heads shows, given that heads shows with probability $\frac{2}{7}$. ◁

Exercise 8.2.42
Prove that the expectation of a $\mathbb{N}^+$-valued random variable which is geometrically distributed with parameter $p \in [0,1]$ is equal to $\frac{1}{p}$. ◁

Theorem 8.2.43 (Linearity of expectation)
Let $(\Omega, \mathbb{P})$ be a probability space, let $E \subseteq \mathbb{R}$ be countable, let X and Y be E-valued random variables on $(\Omega, \mathbb{P})$, and let $a, b \in \mathbb{R}$. Then

$$\mathbb{E}[aX + bY] = a\mathbb{E}[X] + b\mathbb{E}[Y]$$

Proof
This follows directly from the fact that summation is linear. Indeed,

$$
\begin{aligned}
\mathbb{E}[aX + bY] &= \sum_{\omega \in \Omega} (aX + bY)(\omega)\mathbb{P}(\{\omega\}) && \text{by Lemma 8.2.38}\\
&= \sum_{\omega \in \Omega} \Big(aX(\omega)\mathbb{P}(\{\omega\}) + bY(\omega)\mathbb{P}(\{\omega\})\Big) && \text{expanding}\\
&= a\sum_{\omega \in \Omega} X(\omega)\mathbb{P}(\{\omega\}) + b\sum_{\omega \in \Omega} Y(\omega)\mathbb{P}(\{\omega\}) && \text{by linearity of summatio}\\
&= a\mathbb{E}[X] + b\mathbb{E}[Y] && \text{by Lemma 8.2.38}
\end{aligned}
$$

as required. □

Example 8.2.44
Let X be a random variable representing the sum of the numbers shown when a fair six-sided die is rolled twice. We can write $X = Y + Z$, where Y is the value of the first die roll and Z is the value of the second die roll. By Example 8.2.35, we have $\mathbb{E}[Y] = \mathbb{E}[Z] = 3.5$. Linearity of expectation then yields

$$\mathbb{E}[X] = \mathbb{E}[Y] + \mathbb{E}[Z] = 3.5 + 3.5 = 7$$

so the expected value of the sum of the two die rolls is 7. ◁

Example 8.2.45
A coin, when flipped, shows heads with probability $p \in [0,1]$. The coin is flipped. If it shows heads, I gain \$10; if it shows tails, I lose \$20. We compute the least value of p that ensures that I do not expect to lose money.

Let X be the random variable which is equal to 0 if tails shows, and 1 if heads shows. then $X \sim \mathrm{B}(1, p)$, so that $\mathbb{E}[X] = p$ by Example 8.2.36. Let Y be the amount of money I gain. Then

$$Y = 10X - 20(1 - X) = 30X - 20$$

Hence my expected winnings are

$$\mathbb{E}[Y] = 30\mathbb{E}[X] - 20 = 30p - 20$$

In order for this number to be non-negative, we require $p \geqslant \frac{2}{3}$. ◁

Theorem 8.2.43 generalises by induction to linear combinations of countably many random variables; this is proved in the following exercise

Exercise 8.2.46
Let $(\Omega, \mathbb{P})$ be a probability space, let $E \subseteq \mathbb{R}$ be countable, let $\{X_i \mid i \in I\}$ be a family of E-valued random variables on $(\Omega, \mathbb{P})$, indexed by some countable set I, and let $\{a_n \mid n \in \mathbb{N}\}$ be an I-indexed family of real numbers. Prove that

$$\mathbb{E}\left[\sum_{i \in I} a_i X_i\right] = \sum_{i \in I} a_i \mathbb{E}[X_i]$$

◁

Example 8.2.47
Recall Example 8.2.32: an urn contains $n \geqslant 1$ distinct coupons. Each time you draw a coupon that you have not drawn before, you get a stamp. When you get all n stamps, you win. We find the expected number of times you need to draw a coupon from the urn in order to win.

For each $k \in [n]$, let X_k be the random variable representing the number of draws required to draw the k^{th} new coupon, after $k - 1$ coupons have been collected. Then the total number of times a coupon must be drawn is $X = X_1 + X_2 + \cdots + X_n$.

We already saw that $X_k \sim \text{Geom}\left(\frac{n-k+1}{n}\right)$ for each $k \in [n]$. By Exercise 8.2.42, we have $\mathbb{E}[X_k] = \frac{n}{n-k+1}$ for all $k \in [n]$. By linearity of expectation, it follows that

$$\mathbb{E}[X] = \sum_{k=1}^{n} \mathbb{E}[X_k] = \sum_{k=1}^{n} \frac{n}{n-k+1} = n \sum_{i=1}^{n} \frac{1}{i}$$

◁

Section 8.3
Measure spaces

<div align="center">

Warning!

</div>

This section is not yet finished—do not rely on its correctness or completeness.

To do:

Definition 8.3.1

Let X be a set. A **σ-algebra** on X is a collection $\mathscr{F}$ (LATEX code: \mathcal{F})
of subsets of X such that

   (i) $X \in \mathscr{F}$;

  (ii) For all $A \subseteq X$, if $A \in \mathscr{F}$, then $X \setminus A \in \mathscr{F}$; and

 (iii) If $\{A_i \subseteq X \mid i \in I\}$ is a countable family of sets in $\mathscr{F}$, then $\bigcup_{i \in I} A_i \in \mathscr{F}$.

A **measurable space** is a pair $(X, \mathscr{F})$ consisting of a set together with a σ-algebra
on X. The sets in $\mathscr{F}$ are called ($\mathscr{F}$-)**measurable sets**.

Example 8.3.2

Given any set X, the power set $\mathscr{P}(X)$ of X is a σ-algebra, since $X \subseteq X$, the
relative complement of any subset $A \subseteq X$ is a subset of X, and the union of any
family of subsets of X (countable or otherwise) is a subset of X.   ◁

Example 8.3.3

Let $\mathscr{F}$ be the set of subsets $U \subseteq \mathbb{R}$ such that either U is countable or $\mathbb{R} \setminus U$ is
countable. Then:

   (i) $\mathbb{R} \setminus \mathbb{R} = \varnothing$, which is finite (and hence countable), so $\mathbb{R} \in \mathscr{F}$;

  (ii) Let $A \in \mathscr{F}$. Either A or $\mathbb{R} \setminus A$ is countable.

      • Suppose A is countable. By Exercise 2.1.61 we have $\mathbb{R} \setminus (\mathbb{R} \setminus A) = A$,
        which is countable, so that $\mathbb{R} \setminus A \in \mathscr{F}$;

      • Suppose $\mathbb{R} \setminus A$ is countable. Then we immediately have $\mathbb{R} \setminus A \in \mathscr{F}$.

     In both cases, we see that $\mathbb{R} \setminus A \in \mathscr{F}$.

 (iii) Let $\{A_i \mid i \in I\}$ be a countable family of sets in $\mathscr{F}$.

      • If each A_i is countable, then $\bigcup_{i \in I} A_i$ is countable by Theorem 6.1.10;

- If at least one A_i is uncountable, then there is some $j \in I$ such that $\mathbb{R} \setminus A_j$ is countable. Then by de Morgan's laws for sets (Theorem 2.1.62) we have

$$\mathbb{R} \setminus \bigcup_{i \in I} A_i = \bigcap_{i \in I} (\mathbb{R} \setminus A_i) \subseteq \mathbb{R} \setminus A_j$$

Since $\mathbb{R} \setminus A_j$ is countable, we have $\mathbb{R} \setminus \bigcup_{i \in I} A_i$ is countable.

In both cases, it follows that $\bigcup_{i \in I} A_i \in \mathscr{F}$, as required.

So $\mathscr{F}$ is a σ-algebra on $\mathbb{R}$. ◁

Exercise 8.3.4

Let X be a set and let $\sim$ be an equivalence class on X. Define $\mathscr{F}$ to be the set of unions of $\sim$-equivalence classes; that is

$$\mathscr{F} = \left\{ \bigcup_{x \in U} [x]_\sim \;\middle|\; U \subseteq X \right\}$$

Prove that $\mathscr{F}$ is a σ-algebra on X. ◁

Exercise 8.3.5

Prove that σ-algebras contain the empty set and are closed under countable intersections. That is, given a σ-algebra $\mathscr{F}$ on a set X, prove that $\varnothing \in \mathscr{F}$, and that if $\{A_i \mid i \in I\}$ is a countable family of sets in $\mathscr{F}$, then $\bigcap_{i \in I} A_i \in \mathscr{F}$. ◁

It would be nice if the collection of all *open* subsets of $\mathbb{R}$ were a σ-algebra on $\mathbb{R}$. However, this is not the case.

Exercise 8.3.6

Prove that the set of all open subsets of $\mathbb{R}$ is not a σ-algebra on $\mathbb{R}$. ◁

To remedy this situation, we introduce the notion of a *Lebesgue measurable* subset of $\mathbb{R}$.

To do:

Definition 8.3.7
A **measure space** $(X, \mathscr{F}, \mu)$ consists of a set X, a σ-algebra $\mathscr{F}$ on X, and a function $\mu : X \to [0, \infty]$, satisfying the following conditions:

(i) $\mu(\varnothing) = 0$;

(ii) (**Countable additivity**) If $\{A_i \mid i \in I\}$ is any countable family of pairwise disjoint sets in $\mathscr{F}$, indexed by a countable then

$$\mu\left(\bigcup_{i \in I} A_i\right) = \sum_{i \in I} \mu(A_i)$$

The value $\mu(A)$ of an $\mathscr{F}$-measurable set is called the **measure** of A with respect to μ.

Comparing with Definition 8.1.1 reveals the following readily accessible example of a measure space.

Example 8.3.8
Every discrete probability space $(\Omega, \mathbb{P})$ defines a measure space $(\Omega, \mathscr{P}(\Omega), \mathbb{P})$. To see this, note that $\mathscr{P}(\Omega)$ is a σ-algebra on Ω by Example 8.3.2, we proved in Example 8.1.3 that $\mathbb{P}(\varnothing) = \varnothing$, and the countable additivity condition for discrete probability spaces is a direct translation of that for measure spaces.   ◁

To do:

Construction 8.3.9
To do:

To do:

Definition 8.3.10
A **probability space** is a measure space $(\Omega, \mathscr{F}, \mathbb{P})$ such that $\mathbb{P}(\Omega) = 1$. The set Ω is called the **sample space**; the elements $\omega \in \Omega$ are called **outcomes**; the measurable sets $A \subseteq \mathscr{F}$ are called **events**; and the function $\mathbb{P}$ is called a **probability measure**. Given an event A, the value $\mathbb{P}(A)$ is called the **probability of** A.

To do:

Section 8.Q
Chapter 8 exercises

Under construction!
The end-of-chapter exercise sections are new and in an incomplete state.

Communicating mathematics

Section A.1
The elements of a proof

To do: Introduction

Full sentences

To do: Discussion about how mathematics is read as sentences with words, so should be written with that in mind

The first person plural and the passive voice

To do: Customary to write mathematics in fpp and/or in the passive voice

Balancing symbols and words

To do: Balance

Propositions, theorems, lemmas and corollaries

To do: Uses, subdivision of proofs

Variables and their uses

To do: Letters of the alphabet

To do: Introduce Greek alphabet

Name	Upper	Lower	Name	Upper	Lower
Alpha	A	α	Nu	N	ν
Beta	B	β	Xi	Ξ	ξ
Gamma	Γ	γ	Omicron	O	o
Delta	Δ	δ	Pi	Π	π
Epsilon	E	ε, ϵ	Rho	P	ρ
Zeta	Z	ζ	Sigma	Σ	σ
Eta	H	η	Tau	T	τ
Theta	Θ	θ	Upsilon	Υ	υ
Iota	I	ι	Phi	Φ	φ, ϕ
Lambda	Λ	λ	Chi	X	χ
Kappa	K	κ	Psi	Ψ	ψ
Mu	M	μ	Omega	Ω	ω

LATEX tip

In order to use Greek letters as mathematical variables using LATEX:

- For the upper-case letters that are identical to a letter in the Latin alphabet, use the `\mathrm` command together with the Latin letter. For example, upper-case *rho* can be input as `\mathrm{P}`, even though *rho* corresponds phonemically with the Latin letter *R*.

- For the upper-case letters that are not identical to a letter in the Latin alphabet, the LATEX command is given by their Greek name with an upper-case first letter. For example, the command `\Gamma` produces the ouput Γ.

- For the lower-case letters (except *epsilon* and *phi*—see below), the LATEX command is given by their Greek name in lower case. For example, the command `\eta` produces the output η.

Note that ε (`\varepsilon`) and φ (`\varphi`) are preferred over ϵ (`\epsilon`) and ϕ (`\phi`) to better distinguish them from the symbols $\in$ (element symbol) and $\varnothing$ (empty set), respectively. $\triangleleft$

Seven deadly sins

Now that we have seen examples of features that constitute an *effective* proof, we conclude this section by focusing on some features that make a proof less effective or even incorrect. We will refer to these features as *deadly sins*.

Deadly Sin A.1.1 (Abuse of variables)
Using variables without proper quantification or introduction. $\triangleleft$

Deadly Sin A.1.2 (Contradiction sandwich)
Proving a proposition *p* by assuming ¬*p*, proving *p*, and saying the assumption
¬*p* is false.   ◁

Deadly Sin A.1.3 (Word salad)
Writing a jumble of words that, when strung together, do not form an intelligible
statement.   ◁

Deadly Sin A.1.4 (Proof by definition)
Justifying a claim 'by definition' when the claim does not follow immediately
from a definition.   ◁

Deadly Sin A.1.5 (Proof by intimidation)
Deducing that a result is true by saying it is 'clear'.   ◁

Deadly Sin A.1.6 (Proof by intuition)
Giving an intuitive reason for why a result is true without formal justification.   ◁

Deadly Sin A.1.7 (Proof by backwards logic—*snenop sudom*)
Deducing that a proposition *p* is true by deriving a true conclusion from *p*.   ◁

Section A.2
Writing and structuring a proof

The focus of Chapter 1 was on *how to prove* a proposition: we used symbolic logic as a tool for breaking down a proposition into simpler components, and using abstract facts from logic to uncover the steps that are needed to establish the truth of what we are trying to prove.

But this was only half the battle. A mathematician having convinced themself that they have proved a proposition is then tasked with sharing their proof with others. They must communicate their proof. This is not an easy task.

The focus of this section, therefore, is on *what to write* (and what not to write) in order to convey our mathematical ideas and logical reasoning in written form.

To do:

Logical operators

To do:

Quantifiers

To do:

Vocabulary A.2.1 (Introducing a variable)
Some sentences for introducing a variable x representing an arbitrary element of a set X include: 'let $x \in X$', 'fix $x \in X$' and 'take $x \in X$'.

To do:

To do: The rest of this section

Section A.3

Typesetting mathematics in LaTeX

Being able to type up your mathematical writing is a beneficial skill to have, and is one that is expected of anyone working in a mathematical field. Unfortunately, most Office-style WYSIWYG ('what you see is what you get') text editors are not designed for this task—it becomes quickly cumbersome to deal with complicated mathematical notation, and fast alternations between notation and prose.

LaTeX is a markup language that allows you to input both text and mathematical notation, inputting all mathematical notation, text formatting and document structure as code. What follows is a brisk introduction to LaTeX, that should suffice for the purposes of this book.

The word LaTeX is pronounced like 'LAY-tek' or 'LAH-tek', with a hard 'k' sound—the 'X' is meant to resemble the Greek letter *chi* (χ), so is pronounced by some people as such. It doesn't really matter how you say it, but do be warned that if you pronounce it like 'LAY-teks' then people will think you're talking about something somewhat different.

Finding the software

You can use LaTeX by installing it on your computer, or by using a web-based editor. There are advantages and disadvantages to both.

- On a web-based editor, everything is set up and ready to go, but you have less control over how everything compiles and you need an internet connection at all times when editing your files.

- On a computer, you have more control over how your files compile, you have access to all the logs and auxiliary files (I won't go into this) and you don't need an internet connection—but it can be harder to use different packages, you're at the mercy of your own machine's limitations, and it's more technically involved.

There are many online and computer-based options available, and the reader is encouraged to research their options, but as a starting point, I present here the LaTeX implementations used for writing this book.

Much of the book was originally written using the online editor *ShareLaTeX*,

which has merged with *Overleaf* as of September 2018 and can be accessed at the following URL:

https://www.overleaf.com/

Installing LATEX on a computer is slightly more complicated. In order to make LATEX documents on your computer, you need both a *compiler*, for turning the code into a readable document, and an *editor*, for writing the code and facilitating the compilation process.

The compiler used for this book is *TeX Live*:

https://www.tug.org/texlive/

and the editor is *TeXstudio*:

https://www.texstudio.org/

Both TeX Live and TeXstudio are free, cross-platform and open-source.

Getting started

When you have settled on an online LATEX editor or installed LATEX on your computer, you can start editing. The files containing the LATEX code are plain-text files with the file extension .tex and consists of two components: a *header* and a *body*. Worrying about the details of what goes in the header and what goes in the body is not recommended if you are new to LATEX so, with this in mind, a template can be downloaded from the book's website:

https://infinitedescent.xyz/latex/

The code is replicated at the end of this section, on page .

Text mode and math mode

Before we get into the nitty-gritty, I should mention the difference between 'text mode' and 'math mode'.

- **Text mode** is the default mode: the stuff you type will appear as text, and this is the mode you should use when writing anything that isn't mathematical notation.

- You should use **math mode** when you're typing anything which is mathematical notation, including variables, numbers, fractions, square roots, powers, sums, products, binomial coefficients, and so on.

To enter math mode, enclose whatever mathematical notation you are writing with dollar signs ($). For example, if I type `$E=mc^{2}$` then LATEX shows $E = mc^2$. Sometimes it is convenient to put longer expressions on their own line, in which case you can enclose it with double-dollar signs ($$); for example, if I type

$$a^{2}+b^{2}+c^{2}=ab+bc+ca$$

then LATEX displays

$$a^2 + b^2 + c^2 = ab + bc + ca$$

on a line all of its own.

If you need to type text inside math mode (enclosed by $ signs), you can do that using `\text{...}`, for example:

TEX code

```
$$\sum_{i=1}^n i = \frac{n(n+1)}{2} \text{ for all } n \in
\mathbb{N}$$
```

Output

$$\sum_{i=1}^{n} i = \frac{n(n+1)}{2} \text{ for all } n \in \mathbb{N}$$

Note the spaces before and after 'for all'; had I left those out of the code, they would not appear because LATEX ignores spacing in math mode. You can force a space by putting a backslash before a space, for example `$a b$` gives ab but `$a\ b$` gives $a\ b$.

All mathematical notation should be in math mode, including single variables.

Notice the difference between the following two lines:

<div align="center">

If a and b are both even then so is a+b.

</div>

<div align="center">

If a and b are both even then so is $a + b$.

</div>

While the first is written entirely in text mode, the second is written using math mode for the variables and $+$ sign. Although the differences may not seem big, spread over a whole document it is much clearer when math mode is used (as in the second example). Sometimes ambiguities appear, and in any case LATEX does a much better job at displaying mathematical notation when it is typed in math mode.

Table of mathematical symbols

The following table is a quick reference for the most commonly-used symbols in this book. A complete index of notation can be found at the end of the book.

Logic		
conjunction, disjunction	$\wedge, \vee$	`\wedge, \vee`
negation	$\neg$	`\neg`
implication, biconditional	$\Rightarrow, \Leftrightarrow$	`\Rightarrow, \Leftrightarrow`
exclusive disjunction	$\oplus$	`\oplus`
true, false (in truth table)	$\checkmark, \times$	`\checkmark, \times`
quantifiers (universal, existential)	$\forall, \exists$	`\forall, \exists`
Set theory		
element, subset	$\in, \subseteq$	`\in, \subseteq`
not equal, proper subset	$\ne, \subsetneqq$	`\ne, \subsetneqq`
intersection, (indexed)	$\cap, \bigcap_{i=1}^{n}$	`\cap, \bigcap_{i=1}^{n}`
union, (indexed)	$\cup, \bigcup_{i=1}^{n}$	`\cup, \bigcup_{i=1}^{n}`
relative complement, complement	$X \setminus Y, X^c$	`\setminus, X^c`
product, (indexed)	$\times, \prod_{i=1}^{n}$	`\times, \prod_{i=1}^{n}`
implied lists	$\{1, \dots, n\}$	`\{ 1, \dots, n \}`
indexed sets	$\{x_i \mid i \in I\}$	`\{ x_i \mid i \in I \}`
set-builder notation	$\{x \mid p(x)\}$	`\{ x \mid p(x) \}`
empty, universal set	$\varnothing, \mathcal{U}$	`\varnothing, \mathcal{U}`
number sets	$\mathbb{N}, \mathbb{Z}, \mathbb{Q}, \mathbb{R}$	`\mathbb{N}, \mathbb{Z}`, etc.
Numbers and combinatorics		
multiplication	$m \times n, m \cdot n$	`\times, \cdot`
fractions, exponents	$\frac{m}{n}, m^n$	`\frac{m}{n}, m^{n}`
order relations	$\leqslant, \geqslant$	`\le, \ge`
divisibility, (non-)	$m \mid n, m \nmid n$	`\mid, \nmid`
binomial coefficient	$\binom{n}{k}$	`\binom{n}{k}`
indexed sum, product	$\sum_{i=1}^{n} a_i, \prod_{i=1}^{n} a_i$	`\sum_{i=1}^{n} a_i, \prod`
modular arithmetic	$a \equiv b \bmod n$	`a \equiv b \bmod{n}`
Functions and relations		
functions	$f : X \to Y$	`f : X \to Y`
composition	$g \circ f$	`\circ`
isomorphism	$\cong$	`\cong`
equivalence relations	$\sim, \approx$	`\sim, \approx`
Structured sets		
order relation	$\preccurlyeq, \prec$	`\preceq, \prec`
group operations	$\cdot, \star, \circ$	`\cdot, \star, \circ`

Organisation and formatting

When typing up solutions to problem, organisation can be the difference between a masterpiece and an unreadable heap of notation. Here are some tips to help you organise your work:

Sections and paragraphs

You can split your work up into sections, subsections, subsubsections, and even subsubsubsections. To do this, use `\section{Section title}` or `\section*{Section title}`; the former includes a section number, and the latter omits it. To start a new paragraph, simply make two new lines in the code.

Bulleted and enumerated lists

Sometimes it is useful to use bullet points or give an enumerated list. For example, in these notes, I separate the base case from the induction step in proofs by induction by using bullet points.

For a bulleted list you can use the `itemize` environment:

TᴇX code	Output

```
\begin{itemize}
\item Something here\dots
\item You can also make a list
inside another list:
  \begin{itemize}
  \item Like this.
  \item Isn't it fun?
  \end{itemize}
\item Well, not that fun.
\end{itemize}
```

- Something here...

- You can also make a list inside another list:
  - ⋄ Like this.
  - ⋄ Isn't it fun?

- Well, not that fun.

For an enumerated list, you can use the `enumerate` environment. You can play around with different methods of enumeration, which you specify in square brackets `[...]`; this book most frequently uses `(i)`, `(a)` and `(1)`:

TEX code	*Output*

```
\begin{enumerate}[(a)]
\item Here's the first thing;
\item Here's the second thing;
\item Here's the third thing.
\end{enumerate}
```

(a) Here's the first thing;

(b) Here's the second thing;

(c) Here's the third thing.

Definitions, results and proofs

If you use the provided templates, you can make definitions, and state and prove results, using the following environments:

`definition`, `example`, `proposition`, `theorem`, `lemma`, `corollary`, `proof`

They are given a number, such as **Definition 3** or **Theorem 2.11**, depending on how your document is set up.

Here's an example of a theorem appearing in the third section of a document, in which five definitions, results or examples come before it:

TEX code	*Output*

```
\begin{theorem}
Let $a, b \in \mathbb{Z}$.
Then $a^2+b^2 \ge 0$.
\end{theorem}

\begin{proof}
Exercise to the reader.
\end{proof}
```

Theorem 3.6. Let $a, b \in \mathbb{Z}$. Then $a^2 + b^2 \geqslant 0$.
Proof. Exercise to the reader. $\square$

Note that the box ($\square$) designating the end of the proof is inserted automatically when you close the `proof` environment.

Labels

As you change the contents of a document, the numbering of the definitions, examples and results might change. To refer to a specific result, instead of typing the number and having to change it each time the number changes, you can use the `\label`) and `\ref`) commands.

An example of this in action is as follows:

<div style="text-align:center">*TEX code*　　　　　　　　*Output*</div>

```
\begin{definition}
\label{defDivides}
Say $a$ \textbf{divides} $b$
if there exists $k \in
\mathbb{Z}$ such that $ka = b$.
\end{definition}

We will use Definition
\ref{defDivides} for
absolutely nothing.
```

> **Definition 2.11.** Say *a* **divides** *b* if there exists $k \in \mathbb{Z}$ such that $ka = b$.
>
> We will use Definition 2.11 for absolutely nothing.

Formatting

In text mode. To put the icing on the cake, you might want to make some words **bold** or *italicised*. This is simple: for bold text type `\textbf{text here}` and for italic text type `\textit{text here}`. In TeXstudio and Overleaf you can press `Ctrl+B` and `Ctrl+I` to avoid having to type all this out. Other useful fonts include monospace (`\texttt{text here}`), sans-serif (`\textsf{text here}`) and underlined (`\underline{text here}`).

In math mode. There are also various fonts or font styles that you can use inside math mode, including:

- Roman (i.e. not italic): AaBbCc, `\mathrm{AaBbCc}`;

- Bold: **AaBbCc**, `\mathbf{AaBbCc}`;

- Sans-serif: AaBbCc, `\mathsf{AaBbCc}`;

- Blackboard bold: $\mathbb{ABCDE}$, `\mathbb{ABCDE}` — only capital letters;

- Fraktur: $\mathfrak{AaBbCc}$, `\mathfrak{AaBbCc}`;

- Calligraphic: $\mathcal{ABCDE}$, `\mathcal{ABCDE}` — only capital letters;

Tables

Tables can be created using the `tabular` environment. You can specify how columns are aligned and separated as an argument to the command `\begin{tabular}`: write `l` or `c` or `r` to specify that a column should be aligned left, centre or right, respectively. If you want columns to be separated by a single or double line, enter a single or double bar (`|` or `||`), respectively.

Columns are then separated by ampersands (`\&`) and you can move to a new row by entering a double-backslash (`\\`). To insert a horizontal line between two rows, simply enter `\hline`.

Here's an example:

TEX code	*Output*

```
\begin{tabular}{c|ccc}
$\times$ & 1 & 2 & 3 \\ \hline
1 & 1 & 2 & 3 \\
2 & 2 & 4 & 6 \\
3 & 3 & 6 & 9
\end{tabular}
```

$\times$	1	2	3
1	1	2	3
2	2	4	6
3	3	6	9

Aligned equations

Occasionally a proof may require you to demonstrate that two terms are equal by proving a sequence of intermediate equations. This can be done using the `align*` environment, which behaves much like the `tabular` environment.

New lines are introduced by inserting a double-backslash (`\\`), and alignment points are introduced with an ampersand (`&`). For example:

TEX code	*Output*

```
\begin{align*}
(n+1)! - n!
  &= (n+1)n! - n! \\
  &= n \cdot n! + n! - n! \\
  &= n \cdot n!
\end{align*}
```

$$
\begin{aligned}
(n+1)! - n! &= (n+1)n! - n! \\
&= n\cdot n! + n! - n! \\
&= n\cdot n!
\end{aligned}
$$

Note that the `align*` environment automatically enters into math mode, so no dollar signs (`$`) are needed.

Entering more ampersands will create more columns, whose alignment alternates (right, left, right, left, and so on). For example, to add annotations to each line, you can enter a double ampersand (`&&`). For example:

TEX code

```
\begin{align*}
(n+1)! - n!
  & = (n+1)n! - n! && \text{by recursive def of factorials} \\
  & = n \cdot n! + n! - n! && \text{by distributivity} \\
  & = n \cdot n! && \text{by cancellation}
\end{align*}
```

Output

$$
\begin{aligned}
(n+1)! - n! = (n+1)n! - n! \qquad & \text{by recursive def of factorials} \\
= n \cdot n! + n! - n! \qquad & \text{by distributivity} \\
= n \cdot n! \qquad & \text{by cancellation}
\end{aligned}
$$

Note again that, because the `align*` environment automatically enters math mode, any annotations must be made within the `\text{...}` command.

Graphics

Images can then be inserted using the `\includegraphics` command. The format is `\includegraphics[parameters]{filename}` where `parameters` denotes information telling LATEX how large you want the image to be, and `filename` is the name of the image file, which includes the path relative to the main `.tex` file. For example if *donkey.png* is stored in a directory called *images*, you would enter '`images/donkey.png`' instead of '`donkey.png`'.

The simplest way to control the size of the image is to enter `[width=k\textwidth]`, where `k` is a scaling factor between 0 and 1.

For example:

TEX code

```
\begin{center}
\includegraphics[width=0.3\textwidth]{includes/donkey.png}
\end{center}
```

Output

More advanced techniques

I should take a moment to emphasise that what really matters is your ability to communicate mathematical arguments clearly and correctly. The LaTeX tools discussed so far in this section are more than sufficient for our purposes.

However, if you are interested in pushing your LaTeX skills further or there is a feature you're unsure about how to implement, then I recommend browsing or searching one of the following websites:

- `http://tex.stackexchange.com` — Q&A website about LaTeX

- `https://en.wikibooks.org/wiki/LaTeX` — online LaTeX manual

Practice page

Try to recreate the following page, remembering to use `\label` and `\ref` to refer to enumerated items (such as 'Proposition 1.2').

Squarefree integers
Carl Friedrich Gauss, Wednesday 14th September 1831

Introduction

When you've written this page, you will be unstoppable, at least as far as typesetting mathematics is concerned. You will need to implement:

- Text mode stuff: sections, paragraphs, text formatting, labels and references, lists;
- Math mode stuff: definitions and results, aligned equations, etc.

So let's get on with it!

1 Squarefree integers

1.1 Definition and an elementary result

Definition 1.1. An integer a is **squarefree** if it is divisible by no perfect square other than 1. That is, if n^2 divides a then $n^2 = 1$.

Proposition 1.2. A non-zero non-unit a is squarefree if and only if

$$a = p_1 \times p_2 \times \cdots \times p_n$$

for distinct primes $p_1, p_2, \ldots, p_n$.

Proof. We leave the proof as an exercise to the reader. □

1.2 Some examples

Example 1.3. Some concrete examples include:

(i) 5610 is squarefree by Proposition 1.2, since

$$5610 = 10 \times 561$$
$$= (2 \times 5) \times (11 \times 17)$$

(ii) 12 is not squarefree since $4 \mid 12$ and $4 = 2^2$.

1

Template file

What follows is a template `.tex` file to get you started; it can be downloaded from https://infinitedescent.xyz/latex/.

```latex
 1  \documentclass[11pt]{article}
 2
 3  % Edit the following to change the title, author name and date
 4  \title{A \LaTeX{} document}
 5  \author{Firstnametina McLastnamerson}
 6  \date{Someday 0th Jantember 3000}
 7
 8  % Packages
 9  \usepackage{amsmath}
10  \usepackage{amsfonts}
11  \usepackage{amssymb}
12  \usepackage{amsthm}
13  \usepackage{enumerate}
14  \usepackage{geometry}
15  \usepackage{graphicx}
16  \usepackage{hyperref}
17  \usepackage{xcolor}
18
19  % Page setup
20  \setlength{\parskip}{10pt}
21  \setlength{\parindent}{0pt}
22  \geometry{
23      paper={letterpaper}, % Change to 'a4paper' for A4 size
24      marginratio={1:1},
25      margin={1.25in}
26  }
27
28  % Theorem environments
29  \theoremstyle{definition}
30  \newtheorem{theorem}{Theorem}
31  \newtheorem{lemma}[theorem]{Lemma}
32  \newtheorem{corollary}[theorem]{Corollary}
33  \newtheorem{proposition}[theorem]{Proposition}
34  \newtheorem{definition}[theorem]{Definition}
35  \newtheorem{example}[theorem]{Example}
36
37  \begin{document}
38  \maketitle
39
40  %%%%%%%%%%%%%%%%%%%%%%%%%%%
41  %% Start of document body %%
42  %%%%%%%%%%%%%%%%%%%%%%%%%%%
43
44  Hello world! Did you know that $3^2 + 4^2 = 5^2$?
45
46  %%%%%%%%%%%%%%%%%%%%%%%%%%%
47  %% End of document body   %%
48  %%%%%%%%%%%%%%%%%%%%%%%%%%%
49  \end{document}
```

Appendix B

Miscellany

There have been a number of times in the book where we have avoided delving too deeply into the more technical or obscure aspects of a definition or proof. Usually this was because exploring these aspects was not central to the topic at hand, or because the details involved were sufficiently messy that providing all the details would obfuscate the main ideas being discussed.

This appendix provides a home for the comments we didn't make, the theorems we didn't prove, the details we didn't provide and the obscurities we didn't explore.

We begin with a quick glance at the *foundations of mathematics* in Section B.1. We will provide the axioms for Zermelo–Fraenkel set theory (ZF), which encodes all mathematical objects as sets and allows us to derive all mathematical objects from a collection of axioms. We will also demonstrate how to encode natural numbers, integers, rational numbers and complex numbers within this framework.

Set theoretic foundations

To do:

Zermelo–Fraenkel set theory

To do:

The first two axioms of ZF set theory that we introduce are the *axiom of extensionality* and the *axiom of foundation*. They concern the behaviour of the set elementhood relation $\in$, with the axiom of extensionality describing how it relates to equality of sets (as in Axiom 2.1.22), and axiom of foundation

---

Axiom B.1.1 (Axiom of extensionality)
If two sets have the same elements, then they are equal.

$$\forall X, \forall Y, [(\forall a, a \in X \Leftrightarrow a \in Y) \Rightarrow X = Y]$$

A consequence of the axiom of extensionality is that two sets can be proved to be equal by proving that they contain the same elements.

Axiom B.1.2 (Axiom of foundation)
Every inhabited set has an element which is $\in$-minimal, in the sense of Definition 5.3.9.
$$\forall X, [(\exists y, y \in X) \Rightarrow \exists x, (x \in X \wedge \forall u \in X, u \notin x)]$$
The axiom of foundation states that $\in$ is a well-founded relation.

The axiom of foundation is mysterious at first sight, but it captures the idea that every set should built up from $\varnothing$ using set theoretic operations.

Lemma B.1.3
Under the axiom of foundation, we have $X \notin X$ for all sets X.

Proof
Let X be a set. The set $\{X\}$ is inhabited since $X \in \{X\}$, so by the axiom of foundation there is some $x \in \{X\}$ such that $u \notin x$ for all $u \in \{X\}$. But the only element of $\{X\}$ is X itself, so this says exactly that $X \notin X$. $\qquad\square$

Exercise B.1.4

Use the axiom of foundation to prove that there is no sequence of sets $X_0, X_1, X_2, \ldots$ such that $X_{n+1} \in X_n$ for all $n \in \mathbb{N}$. ◁

The next few axioms of ZF set theory posit the existence of certain sets or constructions of sets.

Axiom B.1.5 (Empty set axiom)

There is a set with no elements.

$$\exists X, \forall x, x \notin X$$

The empty set axiom asserts the existence of $\varnothing$.

Axiom B.1.6 (Pairing axiom)

For any two sets x and y, there is a set containing only x and y.

$$\forall x, \forall y, \exists X, \forall u, [u \in X \Leftrightarrow (u = x \vee u = y)]$$

The axiom of pairing asserts the existence of sets of the form $\{x, y\}$.

Axiom B.1.7 (Union axiom)

The union of any family of sets exists and is a set.

$$\forall F, \exists U, \forall x, [x \in U \Leftrightarrow \exists X, (x \in X \wedge X \in F)]$$

The axiom of union asserts that if $F = \{X_i \mid i \in I\}$ is a family of sets then the set $U = \bigcup_{i \in I} X_i$ exists.

Axiom B.1.8 (Power set axiom)

The set of all subsets of a set is a set.

$$\forall X, \exists P, \forall U, [U \in P \Leftrightarrow \forall u, (u \in U \Rightarrow u \in X)]$$

The axiom of power set asserts the existence of $\mathscr{P}(X)$ for all sets X.

Assuming only the previously stated axioms, it is entirely plausible that every set be finite. This isn't good news for us, since we want to be able to reason about infinite sets, such as the set $\mathbb{N}$ of natural numbers. The *axiom of infinity* asserts the existence of an infinite set using a clever set theoretic construction called the *successor set* operation.

Definition B.1.9
Given a set X, the **successor set** of X is the set X^+ defined by

$$X^+ = X \cup \{X\}$$

Lemma B.1.10
Let X and Y be sets. If $X^+ = Y^+$, then $X = Y$.

Proof
Assume $X^+ = Y^+$. Then

- We have $X \in X^+$, so $X \in Y^+ = Y \cup \{Y\}$, and so either $X = Y$ or $X \in Y$;

- We have $Y \in Y^+$, so $Y \in X^+ = X \cup \{X\}$, and so either $Y = X$ or $Y \in X$.

If $X = Y$ then we're done. Otherwise, we must have $X \in Y$ and $Y \in X$. But then we can define a sequence of sets by letting

$$X_n = \begin{cases} X & \text{if } n \text{ is even} \\ Y & \text{if } n \text{ is odd} \end{cases}$$

for all $n \in \mathbb{N}$. This sequence satisfies $X_{n+1} \in X_n$ for all $n \in \mathbb{N}$, since if n is even then

$$X_{n+1} = Y \in X = X_n$$

and if n is odd then

$$X_{n+1} = X \in Y = X_n$$

This contradicts Exercise B.1.4, so we must have $X = Y$, as claimed.   □

Axiom B.1.11 (Axiom of infinity)
There is an inhabited set containing successor sets of all of its elements.

$$\exists X, \, [(\exists u, \, u \in X) \wedge \forall x, \, (x \in X \Rightarrow x^+ \in X)]$$

Intuitively, the axiom of infinity tells us that there is a set X which contains (at least) a family of elements of the form u, u^+, u^{++}, u^{+++}, and so on—each of these elements must be distinct by Lemma B.1.10, so that X must be infinite.

Axiom B.1.12 (Axiom of replacement)
The image of any set under any function is a set. That is, for each logical formula $p(x, y)$ with two free variables x, y, we have

$$\forall X, \, [(\forall x \in X, \, \exists! y, \, p(x, y)) \Rightarrow \exists Y, \, \forall y, \, y \in Y \Leftrightarrow \exists x \in X, \, p(x, y)]$$

> **Axiom B.1.13** (Axiom of separation)
> For any logical formula $p(x)$ with one free variable, and any set X, there is a set consisting of the elements of X satisfying $p(x)$.
>
> $$\forall X, \exists U, \forall x, [x \in U \Leftrightarrow (x \in X \wedge p(x))]$$
>
> The axiom of separation asserts the existence of sets of the form $\{x \in X \mid p(x)\}$.

Exercise B.1.14

Prove that the axioms of infinity and separation imply the existence of an empty set. ◁

In light of Exercise B.1.14, the empty set axiom (Axiom B.1.5) is in fact redundant, in the presence of the other axioms. We keep it around for the sake of convenience.

Grothendieck universes

In Section 2.1 one of the first things we defined was a *universal set*, which we promptly forgot about and mentioned as little as possible. In this short subsection we briefly introduce the notion of a *Grothendieck universe*, named after the interesting (and influential) mathematician Alexander Grothendieck.

> **Definition B.1.15**
> A **Grothendieck universe** is a set $\mathscr{U}$ satisfying the following properties:
>
>   (i) The elements of $\mathscr{U}$ are sets;
>
> (ii) For all $X \in \mathscr{U}$, if $x \in X$, then $x \in \mathscr{U}$;
>
> (iii) $\mathbb{N}_{\mathsf{vN}} \in \mathscr{U}$ (see Construction B.2.5);
>
> (iv) For all $X \in \mathscr{U}$, we have $\mathscr{P}(X) \in \mathscr{U}$;
>
>   (v) For all $I \in \mathscr{U}$ and all $\{X_i \mid i \in I\} \subseteq \mathscr{U}$, we have $\bigcup_{i \in I} X_i \in \mathscr{U}$.

The existence of a Grothendieck universe is not implied by the axioms of Zermelo–Frankel set theory (with or without the axiom of choice)—if it were, it would violate *Gödel's incompleteness theorem*, a result that is even further beyond the scope of this book than Grothendieck universes are!

Theorem B.1.16

Let $\mathscr{U}$ be a Grothendieck universe. The axioms of Zermelo–Fraenkel set theory are satisfied relative to $\mathscr{U}$. If the axiom of choice is assumed, then that is also satisfied relative to $\mathscr{U}$. □

The upshot of Theorem B.1.16 is that although there is no universal set, if we assume the existence of a Grothendieck universe $\mathscr{U}$, then for the purposes of this book, we may relativise everything we do to $\mathscr{U}$ and *pretend* that $\mathscr{U}$ is indeed a universal set. And this is exactly what we did in Section 2.1.

Number sets and algebraic structures

To do:

The natural numbers

We can use the framework provided by Zermelo–Fraenkel set theory (Section B.1) to provide set theoretic constructions of the number sets $\mathbb{N}$, $\mathbb{Z}$, $\mathbb{Q}$, $\mathbb{R}$ and $\mathbb{C}$. Indeed, if we want to reason about mathematics within the confines of ZF, we must encode everything (including numbers) as sets!

We will begin with a set theoretic construction of the natural numbers—that is, we will construct a notion of natural numbers in the sense of Definition 3.1.1. We will encode the natural numbers as sets, called *von Neumann natural numbers*. We will identify the natural number 0 with the empty set $\varnothing$, and we will identify the successor operation s with an operation involving sets.

Definition B.2.1

A **von Neumann natural number** is any set obtainable from $\varnothing$ by repeatedly taking successor sets (see Definition B.1.9). Write $0_{\mathsf{vN}} = \varnothing$ and $(n+1)_{\mathsf{vN}} = (n_{\mathsf{vN}})^+$; that is

$$0_{\mathsf{vN}} = \varnothing, \quad 1_{\mathsf{vN}} = \varnothing^+, \quad 2_{\mathsf{vN}} = \varnothing^{++}, \quad 3_{\mathsf{vN}} = \varnothing^{+++}, \quad 4_{\mathsf{vN}} = \varnothing^{++++}, \quad \dots$$

Example B.2.2

The first three von Neumann natural numbers are:

- $0_{\mathsf{vN}} = \varnothing$;

- $1_{\mathsf{vN}} = \varnothing^+ = \varnothing \cup \{\varnothing\} = \{\varnothing\}$;

- $2_{\mathsf{vN}} = \varnothing^{++} = \{\varnothing\}^+ = \{\varnothing\} \cup \{\{\varnothing\}\} = \{\varnothing, \{\varnothing\}\}$.

◁

Exercise B.2.3

Write out the elements of 3_{vN} ($= \varnothing^{+++}$) and of 4_{vN}. ◁

Exercise B.2.4

Recall the definition of von Neumann natural numbers from Definition B.2.1. Prove that $|n_{\mathsf{vN}}| = n$ for all $n \in \mathbb{N}$. ◁

Construction B.2.5

We construct the set $\mathbb{N}_{vN}$ of all von Neumann natural numbers as follows. Let X be an arbitrary set satisfying the axiom of infinity (Axiom B.1.11), and then define $\mathbb{N}_{vN}$ to be the intersection of all subsets of X that also satisfy the axiom of infinity—that is:

$$\mathbb{N}_{vN} = \{x \in X \mid \forall U \in \mathscr{P}(X), [U \text{ satisfies the axiom of infinity } \Rightarrow x \in U]\}$$

The existence of $\mathbb{N}_{vN}$ follows from the axioms of power set (Axiom B.1.8) and separation (Axiom B.1.13).

Theorem B.2.6

The set $\mathbb{N}_{vN}$, zero element 0_{vN} and successor function $s : \mathbb{N}_{vN} \to \mathbb{N}_{vN}$ defined by $s(n_{vN}) = n_{vN}^+$ for all $n_{vN} \in \mathbb{N}_{vN}$, define a notion of natural numbers.

Proof

We must verify Peano's axioms, which are conditions (i)–(iii) of Definition 3.1.1.

To prove (i), observe that for all sets X we have $X^+ = X \cup \{X\}$, so that $X \in X^+$. In particular, we have $n_{vN} \in n_{vN}^+$ for all $n_{vN} \in \mathbb{N}_{vN}$, and hence $n_{vN}^+ \neq \varnothing = 0_{vN}$.

For (ii), let $m_{vN}, n_{vN} \in \mathbb{N}_{vN}$ and assume that $m_{vN}^+ = n_{vN}^+$. Then $m_{vN} = n_{vN}$ by Lemma B.1.10.

For (iii), let X be a set and suppose that $0_{vN} \in X$ and, for all $n_{vN} \in \mathbb{N}_{vN}$, if $n_{vN} \in X$, then $n_{vN}^+ \in X$. Then X satisfies the axiom of infinity (Axiom B.1.11), and so by Construction B.2.5 we have $\mathbb{N}_{vN} \subseteq X$. $\square$

In light of Theorem B.2.6, we may declare 'the natural numbers' to be the von Neumann natural numbers, and have done with it. As such, you can—if you want—think of all natural numbers in these notes as *being* their corresponding von Neumann natural number. With this in mind, we now omit the subscript 'vN', leaving implicit the fact that we are referring to von Neumann natural numbers.

However, there are many other possible notions of natural numbers. In Theorem B.2.8, we prove that any two notions of natural numbers are essentially the same, and so the specifics of how we actually define $\mathbb{N}$, the zero element and successor operation, are irrelevant for most purposes.

First we will prove the following handy lemma, which provides a convenient means of proving when a function is the identity function (Definition 2.2.13).

Lemma B.2.7

Let $(\mathbb{N}, z, s)$ be a notion of natural numbers, and let $j : \mathbb{N} \to \mathbb{N}$ be a function such that $j(z) = 0$ and $j(s(n)) = s(j(n))$ for all $n \in \mathbb{N}$. Then $j = \mathrm{id}_{\mathbb{N}}$.

Proof. By Theorem 3.1.2, there is a unique function $i : \mathbb{N} \to \mathbb{N}$ such that $i(z) = 0$ and $i(s(n)) = s(i(n))$ for all $n \in \mathbb{N}$. But then:

- $j = i$ by uniqueness of i, since j satisfies the same conditions as i; and

- $\mathrm{id}_{\mathbb{N}} = i$ by uniqueness of i, since $\mathrm{id}_{\mathbb{N}}(z) = z$ and $\mathrm{id}_{\mathbb{N}}(s(n)) = s(n) = s(\mathrm{id}_{\mathbb{N}}(n))$ for all $n \in \mathbb{N}$.

Hence $j = \mathrm{id}_{\mathbb{N}}$, as required. $\square$

Theorem B.2.8

Any two notions of natural numbers are essentially the same, in a very strong sense. More precisely, if $(\mathbb{N}_1, z_1, s_1)$ and $(\mathbb{N}_2, z_2, s_2)$ are notions of natural numbers, then there is a unique bijection $f : \mathbb{N}_1 \to \mathbb{N}_2$ such that $f(z_1) = z_2$ and $f(s_1(n)) = s_2(f(n))$ for all $n \in \mathbb{N}_1$.

Proof. The function f with the desired properties is obtained from Definition 3.1.1 applied to $(\mathbb{N}_1, z_1, s_1)$, with $X = \mathbb{N}_2$, $a = z_2$ and $h = s_2$. This also gives us uniqueness of f, so it remains only to prove that f is a bijection.

By applying Definition 3.1.1 to $(\mathbb{N}_2, z_2, s_2)$, with $X = \mathbb{N}_1$, $a = z_1$ and $f = s_1$, we obtain a (unique!) function $g : \mathbb{N}_2 \to \mathbb{N}_1$ such that $g(z_2) = z_1$ and $g(s_2(n)) = s_1(g(n))$ for all $n \in \mathbb{N}_2$.

But then $g(f(z_1)) = g(z_2) = z_1$ and, for all $n \in \mathbb{N}_1$, we have

$$g(f(s_1(n))) = g(s_2(f(n))) = s_2(g(f(n)))$$

and so $g \circ f = \mathrm{id}_{\mathbb{N}_1}$ by Lemma B.2.7. Likewise $f \circ g = \mathrm{id}_{\mathbb{N}_2}$. Hence g is an inverse for f, so that f is a bijection, as required. $\square$

To do:

To do: Arithmetic operations, order

To do: Define relation for the integers, prove it's well-defined, provide intuition.

Definition B.2.9
The **set of integers** is the set $\mathbb{Z}$ defined by

$$\mathbb{Z} = (\mathbb{N} \times \mathbb{N})/\sim$$

where $\sim$ is the equivalence relation on $\mathbb{N} \times \mathbb{N}$ defined by

$$(a,b) \sim (c,d) \text{ if and only if } a+d = b+c$$

for all $(a,b),(c,d) \in \mathbb{N} \times \mathbb{N}$.

To do: Arithmetic operations, order

To do: Define relation for the rationals, prove it's well-defined, provide intuition.

Definition B.2.10
The **set of rational numbers** is the set $\mathbb{Q}$ defined by

$$\mathbb{Q} = (\mathbb{Z} \times (\mathbb{Z} \setminus \{0\}))/\sim$$

where $\sim$ is the equivalence relation on $\mathbb{Z} \times (\mathbb{Z} \setminus \{0\})$ defined by

$$(a,b) \sim (c,d) \text{ if and only if } ad = bc$$

for all $(a,b),(c,d) \in \mathbb{Z} \times (\mathbb{Z} \setminus \{0\})$.

To do: Arithmetic operations, order

To do: Motivate Dedekind cuts

Definition B.2.11 (Dedekind's construction of the real numbers)
The **set of (Dedekind) real numbers** is the set $\mathbb{R}$ defined by

$$\mathbb{R} = \{D \subseteq \mathbb{Q} \mid D \text{ is bounded above and downwards-closed}\}$$

To do: Arithmetic operations, order

To do: Motivate Cauchy reals

Definition B.2.12 (Cauchy's construction of the real numbers)
The **set of** (**Cauchy**) **real numbers** is the set $\mathbb{R}$ defined by

$$\mathbb{R} = \{(x_n) \in \mathbb{Q}^{\mathbb{N}} \mid (x_n) \text{ is Cauchy}\}/\sim$$

where $\sim$ is the equivalence relation defined by

$$(x_n) \sim (y_n) \text{ if and only if } (x_n - y_n) \to 0$$

for all Cauchy sequences $(x_n), (y_n)$ of rational numbers.

To do: Arithmetic operations, order

To do: Motivate definition of complex numbers

Definition B.2.13
The **set of complex numbers** is the set $\mathbb{C} = \mathbb{R} \times \mathbb{R}$.

To do: Arithmetic operations

Algebraic structures

To do: Monoids, groups, rings

Axiomatising the real numbers

To do:

Axioms B.2.14 (Field axioms)
Let X be a set equipped with elements 0 ('zero') and 1 ('unit'), and binary operations $+$ ('addition') and $\cdot$ ('multiplication'). The structure $(X, 0, 1, +, \cdot)$ is a **field** if it satisfies the following axioms:

- **Zero and unit**

(F1) $0 \neq 1$.

- **Axioms for addition**

(F2) (Associativity) $x + (y + z) = (x + y) + z$ for all $x, y, z \in X$.

(F3) (Identity) $x + 0 = x$ for all $x \in X$.

(F4) (Inverse) For all $x \in X$, there exists $y \in X$ such that $x + y = 0$.

(F5) (Commutativity) $x + y = y + x$ for all $x, y \in X$.

- **Axioms for multiplication**

(F6) (Associativity) $x \cdot (y \cdot z) = (x \cdot y) \cdot z$ for all $x, y, z \in X$.

(F7) (Identity) $x \cdot 1 = x$ for all $x \in X$.

(F8) (Inverse) For all $x \in X$ with $x \neq 0$, there exists $y \in X$ such that $x \cdot y = 1$.

(F9) (Commutativity) $x \cdot y = y \cdot x$ for all $x, y \in X$.

- **Distributivity**

(F10) $x \cdot (y + z) = (x \cdot y) + (x \cdot z)$ for all $x, y, z \in X$.

Example B.2.15

The rationals $\mathbb{Q}$ and the reals $\mathbb{R}$ both form fields with their usual notions of zero, unit, addition and multiplication. However, the integers $\mathbb{Z}$ do not, since for example 2 has no multiplicative inverse. ◁

Example B.2.16

Let $p > 0$ be prime. The set $\mathbb{Z}/p\mathbb{Z}$ (see Definition 5.1.39) is a field, with zero element $[0]_p$ and unit element $[1]_p$, and with addition and multiplication defined by

$$[a]_p + [b]_p = [a+b]_p \quad \text{and} \quad [a]_p \cdot [b]_p = [ab]_p$$

for all $a, b \in \mathbb{Z}$. Well-definedness of these operations is immediate from Theorem 4.3.6 and the modular arithmetic theorem (Theorem 4.3.9).

The only axiom which is not easy to verify is the multiplicative inverse axiom (F8). Indeed, if $[a]_p \in \mathbb{Z}/p\mathbb{Z}$ then $[a]_p \neq [0]_p$ if and only if $p \nmid a$. But if $p \nmid a$ then $a \perp p$, so a has a multiplicative inverse u modulo p. This implies that $[a]_p \cdot [u]_p = [au]_p = [1]_p$. So (F8) holds. ◁

Exercise B.2.17

Let $n > 0$ be composite. Prove that $\mathbb{Z}/n\mathbb{Z}$ is not a field, where zero, unit, addition and multiplication are defined as in Example B.2.16. ◁

Axioms B.2.14 tell us that every element of a field has an additive inverse, and every *nonzero* element of a field has a multiplicative inverse. It would be convenient if inverses were *unique* whenever they exist. Proposition B.2.18 proves that this is the case.

Proposition B.2.18 (Uniqueness of inverses)

Let $(X, 0, 1, +, \cdot)$ be a field and let $x \in X$. Then

(a) Suppose $y, z \in X$ are such that $x + y = 0$ and $x + z = 0$. Then $y = z$.

(b) Suppose $x \neq 0$ and $y, z \in X$ are such that $x \cdot y = 1$ and $x \cdot z = 1$. Then $y = z$.

Proof of (a)

By calculation, we have

$$
\begin{aligned}
y &= y + 0 && \text{by (F3)} \\
&= y + (x + z) && \text{by definition of } z \\
&= (y + x) + z && \text{by associativity (F2)} \\
&= (x + y) + z && \text{by commutativity (F5)} \\
&= 0 + z && \text{by definition of } y \\
&= z + 0 && \text{by commutativity (F5)} \\
&= z && \text{by (F3)}
\end{aligned}
$$

so indeed $y = z$.

The proof of (b) is essentially the same and is left as an exercise. ☐

Since inverses are unique, it makes sense to have notation to refer to them.

Notation B.2.19

Let $(X, 0, 1, +, \cdot)$ be a field and let $x \in X$. Write $-x$ for the (unique) additive inverse of x and, if $x \ne 0$ write x^{-1} for the (unique) multiplicative inverse of x.

Example B.2.20

In the fields $\mathbb{Q}$ and $\mathbb{R}$, the additive inverse $-x$ of an element x is simply its negative, and the multiplicative inverse x^{-1} of some $x \ne 0$ is simply its reciprocal $\frac{1}{x}$. ◁

Example B.2.21

Let $p > 0$ be prime and let $[a]_p \in \mathbb{Z}/p\mathbb{Z}$. Then $-[a]_p = [-a]_p$ and, if $p \nmid a$, then $[a]_p^{-1} = [u]_p$, where u is any integer satisfying $au \equiv 1 \bmod p$. ◁

Exercise B.2.22

Let $(X, 0, 1, +, \cdot)$ be a field. Prove that $-(-x) = x$ for all $x \in X$, and that $(x^{-1})^{-1} = x$ for all nonzero $x \in X$. ◁

Example B.2.23

Let $(X, 0, 1, +, \cdot)$ be a field. We prove that if $x \in X$ then $x \cdot 0 = 0$. Well, $0 = 0 + 0$ by (F3). Hence $x \cdot 0 = x \cdot (0 + 0)$. By distributivity (F10), we have $x \cdot (0 + 0) = (x \cdot 0) + (x \cdot 0)$. Hence

$$
x \cdot 0 = (x \cdot 0) + (x \cdot 0)
$$

Let $y = -(x \cdot 0)$. Then

$$
\begin{aligned}
0 &= x \cdot 0 + y && \text{by (F4)} \\
&= ((x \cdot 0) + (x \cdot 0)) + y && \text{as above} \\
&= (x \cdot 0) + ((x \cdot 0) + y) && \text{by associativity (F2)} \\
&= (x \cdot 0) + 0 && \text{by (F4)} \\
&= x \cdot 0 && \text{by (F3)}
\end{aligned}
$$

so indeed we have $x \cdot 0 = 0$.   ◁

Exercise B.2.24
Let $(X, 0, 1, +, \cdot)$ be a field. Prove that $(-1) \cdot x = -x$ for all $x \in X$, and that $(-x)^{-1} = -(x^{-1})$ for all nonzero $x \in X$.   ◁

What makes the real numbers useful is not simply our ability to add, subtract, multiply and divide them; we can also compare their size—indeed, this is what gives rise to the informal notion of a *number line*. Axioms B.2.25 make precise exactly what it means for the elements of a field to be assembled into a 'number line'.

Axioms B.2.25 (Ordered field axioms)
Let X be a set, $0, 1 \in X$ be elements, $+, \cdot$ be binary operations, and $\leqslant$ be a relation on X. The structure $(X, 0, 1, +, \cdot, \leqslant)$ is an **ordered field** if it satisfies the field axioms (F1)–(F10) (see Axioms B.2.14) and, additionally, it satisfies the following axioms:

- **Linear order axioms**

(PO1) (Reflexivity) $x \leqslant x$ for all $x \in X$.

(PO2) (Antisymmetry) For all $x, y \in X$, if $x \leqslant y$ and $y \leqslant x$, then $x = y$.

(PO3) (Transitivity) For all $x, y, z \in X$, if $x \leqslant y$ and $y \leqslant z$, then $x \leqslant z$.

(PO4) (Linearity) For all $x, y \in X$, either $x \leqslant y$ or $y \leqslant x$.

- **Interaction of order with arithmetic**

(OF1) For all $x, y, z \in X$, if $x \leqslant y$, then $x + z \leqslant y + z$.

(OF2) For all $x, y \in X$, if $0 \leqslant x$ and $0 \leqslant y$, then $0 \leqslant xy$.

Example B.2.26
The field $\mathbb{Q}$ of rational numbers and and the field $\mathbb{R}$ of real numbers, with their usual notions of ordering, can easily be seen to form ordered fields.   ◁

Example B.2.27
We prove that, in any ordered field, we have $0 \leqslant 1$. Note first that either $0 \leqslant 1$

or $1 \leqslant 0$ by linearity (PO4). If $0 \leqslant 1$ then we're done, so suppose $1 \leqslant 0$. Then $0 \leqslant -1$; indeed:

$$
\begin{aligned}
0 &= 1 + (-1) & &\text{by (F4)} \\
&\leqslant 0 + (-1) & &\text{by (OF1), since } 1 \leqslant 0 \\
&= (-1) + 0 & &\text{by commutativity (F5)} \\
&= -1 & &\text{by (F3)}
\end{aligned}
$$

By (OF2), it follows that $0 \leqslant (-1)(-1)$. But $(-1)(-1) = 1$ by Exercise B.2.24, and hence $0 \leqslant 1$. Since $1 \leqslant 0$ and $0 \leqslant 1$, we have $0 = 1$ by antisymmetry (PO2). But this contradicts axiom (F1). Hence $0 \leqslant 1$. In fact, $0 < 1$ since $0 \neq 1$. ◁

We have seen that $\mathbb{Q}$ and $\mathbb{R}$ are ordered fields (Examples B.2.20 and B.2.26), and that $\mathbb{Z}/p\mathbb{Z}$ is a field for $p > 0$ prime (Example B.2.16). The following proposition is an interesting result proving that there is no notion of 'ordering' under which the field $\mathbb{Z}/p\mathbb{Z}$ can be made into an ordered field!

Proposition B.2.28

Let $p > 0$ be prime. There is no relation $\leqslant$ on $\mathbb{Z}/p\mathbb{Z}$ which satisfies the ordered field axioms.

Proof

We just showed that $[0] \leqslant [1]$. It follows that, for all $a \in \mathbb{Z}$, we have $[a] \leqslant [a] + [1]$; indeed:

$$
\begin{aligned}
[a] &= [a] + [0] & &\text{by (F3)} \\
&\leqslant [a] + [1] & &\text{by (OF1), since } [0] \leqslant [1] \\
&= [a + 1] & &\text{by definition of } + \text{ on } \mathbb{Z}/p\mathbb{Z}
\end{aligned}
$$

It is a straightforward induction to prove that $[a] \leqslant [a + n]$ for all $n \in \mathbb{N}$. But then we have

$$[1] \leqslant [1 + (p - 1)] = [p] = [0]$$

so $[0] \leqslant [1]$ and $[1] \leqslant [0]$. This implies $[0] = [1]$ by antisymmetry (PO2), contradicting axiom (F1). □

Exercise B.2.29

Let $(X, 0, 1, +, \cdot)$ be a field. Prove that if X is finite, then there is no relation $\leqslant$ on X such that $(X, 0, 1, +, \cdot, \leqslant)$ is an ordered field. ◁

Theorem B.2.30 below summarises some properties of ordered fields which are used in our proofs. Note, however, that this is certainly *not* an exhaustive list of elementary properties of ordered fields that we use—to explicitly state and prove all of these would not make for a scintillating read.

Theorem B.2.30

Let $(X, 0, 1, +, \cdot, \leqslant)$ be an ordered field. Then

(a) For all $x, y \in X$, $x \leqslant y$ if and only if $0 \leqslant y - x$;

(b) For all $x \in X$, $-x \leqslant 0 \leqslant x$ or $x \leqslant 0 \leqslant -x$;

(c) For all $x, x', y, y' \in X$, if $x \leqslant x'$ and $y \leqslant y'$, then $x + y \leqslant x' + y'$;

(d) For all $x, y, z \in X$, if $0 \leqslant x$ and $y \leqslant z$, then $xy \leqslant xz$;

(e) For all nonzero $x \in X$, if $0 \leqslant x$, then $0 \leqslant x^{-1}$.

(f) For all nonzero $x, y \in X$, if $x \leqslant y$, then $y^{-1} \leqslant x^{-1}$.

Proof of (a), (b) and (e)

(a) ($\Rightarrow$) Suppose $x \leqslant y$. Then by additivity (OF1), $x + (-x) \leqslant y + (-x)$, that is $0 \leqslant y - x$. ($\Leftarrow$) Suppose $0 \leqslant y - x$. By additivity (OF1), $0 + x \leqslant (y - x) + x$; that is, $x \leqslant y$.

(b) We know by linearity (PO4) that either $0 \leqslant x$ or $x \leqslant 0$. If $0 \leqslant x$, then by (OF1) we have $0 + (-x) \leqslant x + (-x)$, that is $-x \leqslant 0$. Likewise, if $x \leqslant 0$ then $0 \leqslant -x$.

(e) Suppose $0 \leqslant x$. By linearity (PO4), either $0 \leqslant x^{-1}$ or $x^{-1} \leqslant 0$. If $x^{-1} \leqslant 0$, then by (d) we have $x^{-1} \cdot x \leqslant 0 \cdot x$, that is $1 \leqslant 0$. This contradicts Example B.2.27, so we must have $0 \leqslant x^{-1}$.

The proofs of the remaining properties are left as an exercise. □

We wanted to characterise the reals completely, but so far we have failed to do so—indeed, Example B.2.26 showed that both $\mathbb{Q}$ and $\mathbb{R}$ are ordered fields, so the ordered field axioms do not suffice to distinguish $\mathbb{Q}$ from $\mathbb{R}$. The final piece in the puzzle is *completeness*. This single additional axiom distinguishes $\mathbb{Q}$ from $\mathbb{R}$, and in fact completely characterises $\mathbb{R}$ (see Theorem B.2.32).

Axioms B.2.31 (Complete ordered field axioms)

Let X be a set, $0, 1 \in X$ be elements, $+, \cdot$ be binary operations, and $\leqslant$ be a relation on X. The structure $(X, 0, 1, +, \cdot, \leqslant)$ is a **complete ordered field** if it is an ordered field—that is, it satisfies axioms (F1)–(F10), (PO1)–(PO4) and (OF1)–(OF2) (see Axioms B.2.14 and B.2.25)—and, in addition, it satisfies the following **completeness axiom**:

C1) Let $A \subseteq X$. If A has an upper bound, then it has a least upper bound. Specific-

ally, if there exists $u \in X$ such that $a \leqslant u$ for all $a \in A$, then there exists $s \in X$ such that

$\diamond$ $a \leqslant s$ for all $a \in A$; and

$\diamond$ If $s' \in X$ is such that $a \leqslant s'$ for all $a \in A$, then $s \leqslant s'$.

We call such a value $s \in X$ a **supremum** for A.

Theorem B.2.32
The real numbers $(\mathbb{R}, 0, 1, +, \cdot, \leqslant)$ form a complete ordered field. Moreover, any two complete ordered fields are essentially the same. $\qquad\Box$

The notion of 'sameness' alluded to in Theorem B.2.32 is more properly called *isomorphism*. A proof of this theorem is intricate and far beyond the scope of this book, so is omitted. What it tells us is that it doesn't matter exactly how we define the reals, since any complete ordered field will do. We can therefore proceed with confidence that, no matter what notion of 'real numbers' we settle on, everything we prove will be true of that notion. This is for the best, since we haven't actually defined the set $\mathbb{R}$ of real numbers at all!

The two most common approaches to constructing a set of real numbers are:

- **Dedekind reals.** In this approach, real numbers are identified with particular subsets of $\mathbb{Q}$—informally speaking, $r \in \mathbb{R}$ is identified with the set of rational numbers less than r.

- **Cauchy reals.** In this approach, real numbers are identified with equivalence classes of sequences of rational numbers—informally speaking, $r \in \mathbb{R}$ is identified with the set of sequences of rational numbers which converge to r (in the sense of Definition 7.2.15).

Discussion of Dedekind and Cauchy reals, as well as constructions of the other number sets, is relegated to Section B.1.

Appendix C

Hints for selected exercises

Hint for Exercise 1.1.34
Suppose $n = d_r \cdot 10^r + \cdots + d_1 \cdot 10 + d_0$ and let $s = d_r + \cdots + d_1 + d_0$. Start by proving that $3 \mid n - s$.

Hint for Exercise 1.1.48
Use the law of excluded middle according to whether the proposition '$\sqrt{2}^{\sqrt{2}}$' is rational' is true or false.

Hint for Exercise 1.2.22
Look carefully at the definition of divisibility (Definition 0.12).

Hint for Exercise 1.3.9
Note that you may need to use the law of excluded middle (Axiom 1.1.44) and the principle of explosion (Axiom 1.1.49).

Hint for Exercise 1.3.22
Express this statement as $\forall n \in \mathbb{Z}, (n \text{ is even}) \Leftrightarrow (n^2 \text{ is even})$, and note that the negation of 'x is even' is 'x is odd'.

Hint for Exercise 1.3.37
Start by expressing $\Leftrightarrow$ in terms of $\Rightarrow$ and $\wedge$, as in Definition 1.1.28.

Hint for Exercise 2.1.31
Recall from the beginning of Section 2.1 that $\forall x \in X, p(x)$ is equivalent to $\forall x, (x \in X \Rightarrow p(x))$ and $\exists x \in X, p(x)$ is equivalent to $\exists x, (x \in X \wedge p(x))$. What can be said about the truth value of $x \in E$ when E is empty?

Hint for Exercise 2.1.56
You need to find a family of subsets of $\mathbb{N}$ such that (i) any two of the subsets have

infinitely many elements in common, but (ii) given any natural number, you can find one of the subsets that it is *not* an element of.

Hint for Exercise 2.2.17
What is the value of $i(z)$ if $z \in X \cap Y$?

Hint for Exercise 2.2.23
Look closely at Definition 2.2.18.

Hint for Exercise 2.3.12
Recall Definition 2.2.24.

Hint for Exercise 2.3.14
If Z were a subset of Y, then we could easily define an injection $i : Z \to Y$ by $i(z) = z$ for all $z \in Z$. Are there any subsets of Y that are associated with a function whose codomain is Y?

Hint for Exercise 2.3.19
To define the bijection, think about what the elements of the two sets look like: The elements of $\prod_{k=1}^{n+1} X_k$ look like $(a_1, a_2, \ldots, a_n, a_{n+1})$, where $a_k \in X_k$ for each $1 \leqslant k \leqslant n+1$. On the other hand, the elements of $\left(\prod_{k=1}^{n} X_k \right) \times X_{n+1}$ look like $((a_1, a_2, \ldots, a_n), a_{n+1})$.

Hint for Exercise 2.3.28
This can be proved in a single sentence; if you find yourself writing a long proof, then there is an easier way.

Hint for Exercise 2.3.31
The proof is almost identical to Exercise 2.3.28.

Hint for Exercise 2.3.39
For part (c), don't try to write a formula for the inverse of h; instead, use the fundamental theorem of arithmetic.

Hint for Exercise 2.3.45
Use Exercise 2.3.40.

Hint for Exercise 3.1.26
Observe that $7^{n+1} - 2 \cdot 4^{n+1} + 1 = 4(7^n - 2 \cdot 4^n + 1) + 3(7^n - 1)$ for all $n \in \mathbb{N}$.

Hint for Exercise 3.1.46
Observe that if n dubloons can be obtained using only 3 and 5 dubloon coins,

then so can $n+3$. See how you might use this fact to exploit strong induction with multiple base cases.

Hint for Exercise 3.1.51
Prove first that if $a \in \mathbb{Z}$ and a^2 is divisible by 3, then a is divisible by 3.

Hint for Exercise 3.2.7
Part (b) has a proof by induction that looks much like the one in part (a). In the induction step, given a surjection $g : [m+1] \to [n]$, observe that we must have $n \geqslant 1$, and construct a surjection $g^- : [m+1] \setminus \{a\} \to [n-1]$ for some suitable $a \in [m+1]$. Then invoke Lemma 3.2.5, now using the fact that $[m] = [m+1] \setminus \{m+1\}$.

Hint for Exercise 3.2.16
Given $U \subseteq X$, find an injection $U \to X$ and apply Theorem 3.2.13(a).

Hint for Exercise 3.2.17
Recall that $X \cap Y \subseteq X$ for all sets X and Y.

Hint for Exercise 3.2.20
Apply Proposition 3.2.18 with $Y = U$.

Hint for Exercise 3.2.21
Prove by induction on $n \in \mathbb{N}$ that, for all $m, n \in \mathbb{N}$, there is a bijection $[mn] \to [m] \times [n]$.

Hint for Exercise 3.3.16
Any function $f : X \to Y$ with finite domain can be specified by listing its values. For each $x \in X$, how many choices do you have for the value $f(x)$?

Hint for Exercise 3.3.23
The image (Definition 2.2.24) of an injection $[3] \to [4]$ must be a subset of $[4]$ of size three.

Hint for Exercise 3.3.38
How many ways can you select $k+1$ animals from a set containing n cats and one dog?

Hint for Exercise 3.3.41
Find two procedures for counting the number of pairs (U, u), such that $U \subseteq [n]$ is a k-element subset and $u \in U$. Equivalently, count the number of ways of forming a committee of size k from a population of size n, and then appointing one member of the committee to be the chair.

Hint for Exercise 3.3.43

Find an expression for $(a+b+c)!$ in terms of $a!$, $b!$, $c!$ and $\binom{a+b+c}{a,b,c}$, following the pattern of Theorem 3.3.40.

Hint for Exercise 4.1.11
Remember that negative integers can be greatest common divisors too.

Hint for Exercise 4.1.13
Start by proving that d and d' must divide each other.

Hint for Exercise 4.1.24
Example 4.1.21 would be a good starting point.

Hint for Exercise 4.1.38
This is essentially the same as Exercise 4.1.13.

Hint for Exercise 4.1.40
Define $m = \dfrac{ab}{\gcd(a,b)}$ and prove that m satisfies the definition of being a least common multiple of a and b (Definition 4.1.37). Then apply Exercise 4.1.38.

Hint for Exercise 4.2.5
Use the factorial formula for binomial coefficients (Theorem 3.1.32).

Hint for Exercise 4.2.9
Assume $p = mn$ for some $m, n \in \mathbb{Z}$. Prove that m or n is a unit.

Hint for Exercise 4.2.23
What are the prime factors of $n! - 1$?

Hint for Exercise 4.3.23
Consider the list $a^0, a^1, a^2, \ldots$. Since there are only finitely many remainders modulo n, we must have $a^i \equiv a^j \bmod n$ for some $0 \leqslant i < j$.

Hint for Exercise 4.3.30
First find the remainder of 244886 when divided by 12.

Hint for Exercise 4.3.33
Find a bijection $[p] \times C_n \to C_{pn}$, where $C_n = \{k \in [|n|] \mid k \perp n\}$. You will need to use the techniques of Section 3.2 in your proof.

Hint for Exercise 4.3.34
Start by proving that $k \in [pq]$ is *not* coprime to pq if and only if $p \mid k$ or $q \mid k$. You will need to use the techniques of Section 3.2 in your proof.

Hint for Exercise 4.3.39
Recall that $\varphi(100) = 40$—this was Example 4.3.35.

Hint for Exercise 4.3.43
You need to use the fact that p is prime at some point in your proof.

Hint for Exercise 4.3.44
Pair as many elements of $[p-1]$ as you can into multiplicative inverse pairs modulo p.

Hint for Exercise 4.3.54
This generalisation will be tricky! You may need to generalise the definitions and results about greatest common divisors and least common multiples that we have seen so far, including Bézout's lemma. You might want to try proving this first in the case that $n_i \perp n_j$ for all $i \ne j$.

Hint for Exercise 4.3.55
Observe that if $a, k \in \mathbb{Z}$ and $k \mid a$, then $k \mid a + k$.

Hint for Exercise 5.2.26
Use the characterisation of gcd and lcm in terms of prime factorisation.

Hint for Exercise 5.2.29
Use distributivity, together with the fact that $\perp \vee y' = y'$ and $\top \wedge y' = y'$.

Hint for Exercise 6.1.4
Use prime factorisation.

Hint for Exercise 6.1.9
Suppose $X = \mathbb{N}$. By Proposition 6.1.5, the set $\mathbb{N}^k$ is countable. By Theorem 6.1.6(c), it suffices to find an injection $\binom{\mathbb{N}}{k} \to \mathbb{N}^k$.

Hint for Exercise 6.1.13
We have already proved this when X is finite. When X is countably infinite, find a bijection $\{0,1\}^X \to \mathscr{P}(X)$ and apply Theorem 6.1.12. When X is uncountably infinite, find an injection $X \to \mathscr{P}(X)$ and find a way to apply Exercise 6.1.7.

Hint for Exercise 7.2.32
Divide the numerator and denominator by n^d and apply Theorem 7.2.25 and Example 7.2.30.

Hint for Exercise 7.2.39
You might want to begin by solving Exercise 3.1.20.

Hint for Exercise 7.2.53
In the definition of a Cauchy sequence, observe that $x_m - x_n = (x_m - a) - (x_n - a)$, and apply the triangle inequality (Theorem 7.1.9).

Hint for Exercise B.1.14

Let X be the set whose existence is asserted by the axiom of infinity, and take $p(x)$ to be the formula $x \neq x$ in the axiom of separation.

Hint for Exercise B.2.22

Prove that x is an additive inverse for $-x$ (in the sense of Axioms B.2.14(F4)) and use uniqueness of additive inverses. Likewise for x^{-1}.

Index

Index of notation

Index of LaTeX commands

Math mode commands