A Starters Guide to Mathematics

J.M. Swart

January 23, 2020

Contents

Chapter 1

The real numbers

1.1 Informal definition of the real numbers

Informally, a real number is any number that can be written in decimal notation, including negative numbers. So, for example,

$$0, \quad 7, \quad 100.5, \quad 1/3 = 0.33333\ldots, \quad -1.125, \quad \text{and} \quad \pi = 3.14159265\ldots$$

are real numbers. A bit more formally, a real number consists of three ingredients:

- a sign which can be + or -,
- a finite number of digits before the decimal dot,
- an infinite number of digits after the decimal dot.

Here, a *digit* is any of the numbers $0, 1, 2, 3, 4, 5, 6, 7, 8, 9$. A few special rules apply:

- $+0 = -0$,
- we usually do not write the sign $+$,
- if from a certain point on, all digits after the decimal dot are 0, then we do not write these digits,
- we do not allow real numbers that have the property that from a certain point on, all digits after the decimal dot are 9.

Concerning this last point, if we would allow such numbers, then the only logical interpretation of a number like

$$0.1249999999999\ldots$$

would be that it is equal to 0.125. In order to have a unique way of writing down such a number, we do not allow the notation with the repeating 9's.

We can imagine the real numbers as points on an infinitely long straight line:

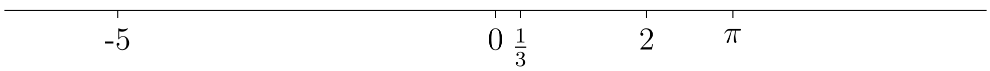

We use the symbols $<$ to indicate that one real number is strictly smaller than another. For example, $3 < \pi < 4$, or more precisely $3.14 < \pi < 3.15$. The symbol $>$ has the opposite meaning: $3.15 > \pi > 3.14$. The symbol $=$ means that two real numbers are equal and $\leq$ (respectively, $\geq$) means that $<$ or $=$ (respectively, $>$ or $=$). The symbol $\neq$ indicates that two numbers are not equal. We will use the symbol $:=$ to indicate that something is equal by definition. For example, we may write $x := 2\pi$ to define a new number x in terms of the numbers 2 and π that we already know.

The numbers $0, 1, 2, \ldots$ are called the *natural numbers*. The numbers $\ldots, -2, -1, 0, 1, 2, \ldots$ are called the *integers*. We say that a number a is *positive* if $a > 0$ and we say that a is *nonnegative* if $a \geq 0$.

1.2 The field operations

The two basic operations with real numbers that everybody learns how to do in school is to *add* them and to *multiply* them. We denote addition with a $+$ and multiplication with a $\cdot$. We use brackets to indicate what operation should be carried out first. So, for example $2 + (3 \cdot 4)$ means that we should first multiply 3 with 4 and then add the result to 2, while $(2 + 3) \cdot 4$ means that we should first add 2 with 3 and then multiply the result with 4. As a result

$$2 + (3 \cdot 4) = 14 \quad \text{and} \quad (2 + 3) \cdot 4 = 20.$$

Mathematicians often use roman or Greek letters to denote real numbers. (As we will later see, they also use letters to denote all kind of other things.) Sometimes, such as in the case of π, these letters signify one particular real number. More often, they signify arbitrary numbers, about which a priori nothing is known. This allows us to write down all kind of general rules. For example:

(i) $a + b = b + a$,

(ii) $a + (b + c) = (a + b) + c$,

(iii) $a + 0 = a$,

(iv) $a \cdot b = b \cdot a$,

(v) $a \cdot (b \cdot c) = (a \cdot b) \cdot c$,

(vi) $a \cdot 1 = a$,

(vii) $a \cdot (b + c) = (a \cdot b) + (a \cdot c)$,

These rules mean that whatever real numbers we write instead of the letters a, b, c, the (in)equality will always be true.[1] A few special rules apply:

- When we denote numbers by letters, we usually do not write the $\cdot$. So ab is the same as $a \cdot b$. Of course, when we write down explicit numbers, we need to write the $\cdot$ since $1 \cdot 3 \neq 13$.
- When we write down a sum of three or more numbers, we do not use brackets, since the order does not matter because of rule (ii).
- When we write down a product of three or more numbers, we do not use brackets, since the order does not matter because of rule (v).
- We often do not write brackets at all. In this case, products should be evaluated first. So, for example, rule (viii) can be written more concisely as $a(b + c) = ab + ac$.

Recall that each real number has a sign, which can be $+$ or $-$. We let $-a$ denote the number that has the opposite sign as a, i.e., if a has the sign $+$, then $-a$ has the sign $-$, and if a has the sign $-$, then $-a$ has the sign $+$ with $-0 = 0$. Then, in addition to the rules above, we have the following general rules:

(viii) For each real number a, there exists a unique real number $-a$ such that $a + (-a) = 0$.

(ix) For each real number $a \neq 0$, there exists a unique real number $1/a$ such that $a(1/a) = 1$.

[1] Of course, it is understood that if a letter occurs more than once in a formula, it should always be replaced by the same number. Different letters may be different numbers, but they need not be. For example, $a + b = b + a$ is also true if $a = b$.

Again, some special rules apply:

- Instead of $a + (-b)$ we write $a - b$. *Subtracting a number is the same as adding minus that number.*

- Instead of $a(1/b)$ we write a/b. *Dividing by a number is the same as multiplying with the inverse of that number.*

Note that $1/0$ is not defined and more generally, division by zero is not defined.

From the properties (i)–(ix), other properties can be deduced, for example:

- $a \cdot 0 = 0$,

- $-a = (-1) \cdot a$,

- $ab = 0$ if and only if $a = 0$ or $b = 0$,

- $\frac{a}{b} \cdot \frac{c}{d} = \frac{ac}{bd}$.

One may wonder if *all* properties of the real numbers can be deduced from the properties (i)–(ix). This is not the case. We will later see that there exist other "numbers" (for example, the rational numbers or the complex numbers) that also satisfy (i)–(ix). In general, any collection of "numbers" that satisfies (i)–(ix) is called a *field*.

If a is a real number and n a positive natural number, then we call

$$a^n := \underbrace{a \cdot \dots \cdot a}_{n \text{ times}}$$

the *n-th power* of a. If we do not write brackets, then taking powers takes precedence over multiplication,. i.e., $ab^n = a \cdot (b^n)$. For each nonnegative real number a and positive natural number n, there exists a unique nonnegative real number, denoted by $\sqrt[n]{a}$, such that

$$(\sqrt[n]{a})^n = a.$$

The number $\sqrt[n]{a}$ is called the *n-th root* of a. In particular, we simply call $\sqrt{a} := \sqrt[2]{a}$ the *root* of a. It is easy to see that

$$a^n \cdot a^m = a^{n+m}. \tag{1.1}$$

Mathematicians usually define $a^0 := 1$, which has the consequence that our previous formula holds even when m or n are zero. With the additional convention that

$$a^{-n} := \frac{1}{a^n},$$

formula (1.1) is even true for when $a \neq 0$ and n, m are general integers.

1.3 Prime numbers and rational numbers

If m and n are positive natural numbers, then there exist unique natural numbers $k > 0$ and $r \geq 0$ such that $m = kn + r$. We call r the *remainder* of m after division by n. If $r = 0$, then we say that n is a *divisor* of m. A *prime number* is a positive natural number n that has precisely *two* divisors, namely the numbers 1 and n itself. Note that according to this definition, 1 is not a prime number. The first few prime numbers are:

$$2, 3, 5, 7, 11, 13, 17, 19, 23, 29, \ldots$$

It is a mathematical theorem that there are infinitely many prime numbers. Let us formally state and prove this.

Theorem 1.1 (Infinitely many prime numbers) *There are infinitely many prime numbers.*

Proof Imagine that there are only finitely many prime numbers, let us call them $p_1 \ldots, p_n$. Then none of the numbers $p_1 \ldots, p_n$ is a divisor of $m := p_1 \cdots p_n - 1$. By our assumption that $p_1 \ldots, p_n$ are the only prime numbers, m is not a prime number, so there exists a natural number $1 < n < m$ that divides m. Since none of the numbers $p_1 \ldots, p_n$ is a divisor of m, they cannot be divisors of n either. But then by the same argument as before, n cannot be a prime number either, so we can find a natural number $1 < k < m$ that divides n. In this way, we find an infinite sequence of positive natural numbers $m, n, k, \ldots$ such that $m > n > k > \cdots$. This is clearly impossible, so our original assumption, that there are only finitely many prime numbers, must be wrong. ∎

If a positive natural number is not a prime number, then we can write it as the product of two smaller positive natural numbers. Continuing this process until we cannot go on, we see that each positive natural number can be written as the product of prime numbers. The following theorem says that this way of writing a positive natural number is unique, i.e., whenever we decompose a positive natural number into prime factors, we find the same prime factors.

Theorem 1.2 (Decomposition into prime factors) *Each natural number $m \geq 2$ can in a unique way be written as $m = p_1^{n_1} \cdots p_k^{n_k}$ where $p_1 < \cdots < p_k$ are prime numbers and $n_1, \ldots, n_k$ are positive natural numbers.*

Recall that according to our earlier definition, 1 is not a prime number. Indeed, if we would consider 1 to be a prime number, the uniqueness claim from Theorem 1.2 would not be true, since, for example, $2 \cdot 3$, $1 \cdot 2 \cdot 3$, and $1^2 \cdot 2 \cdot 3$ are three different ways of writing 6 as a product of prime numbers and 1.

A somewhat different way of writing a positive natural number m as a product of prime factors is to write:

$$m = 2^{m_2} \cdot 3^{m_3} \cdot 5^{m_5} \cdot 7^{m_7} \cdot 11^{m_{11}} \cdots , \tag{1.2}$$

where for any prime number p, we let m_p is the maximal number such that p^{m_p} is a divisor of m, with $m_p := 0$ if p is not a divisor of m. In (1.2) we use the convention that $a^0 := 1$, so that some of the factors in (1.2) are 1. In fact, there is some maximal prime number p that divides m, and all powers m_q for prime numbers $q > p$ are zero. In view of this, even though (1.2) looks like an infinite product, only finitely many factors are different from 1, so in effect it is only a finite product. Note that (1.2) is well-defined even for $m = 1$, in which case $0 = m_2 = m_3 = m_5 = \cdots$.

A number a that can be written as $a = m/n$ where m and $n \neq 0$ are integers, is called a *rational number*. The way of writing a rational number a as m/n is not unique, since $\frac{nk}{mk} = \frac{n}{m}$ for each integer $k \neq 0$. For a positive rational number $a = m/n$, we can choose m and n both positive. We can find the simplest way of writing a by decomposing n and m into prime factors and then crossing out the prime factors that they have in common. For example:

$$\frac{600}{140} = \frac{2^3 \cdot 3 \cdot 5^2}{2^2 \cdot 5 \cdot 7} = \frac{(2^2 \cdot 5) \cdot (2 \cdot 3 \cdot 5)}{(2^2 \cdot 5) \cdot 7} = \frac{30}{7}.$$

It is well-known that a real number is rational if and only if its digits behind the decimal point start repeating from a certain point onwards. We can write down such a number by underlining the digits that repeat. For example:

$$\begin{aligned}
\frac{1}{3} &= 0.\underline{3} = 0.333333333\ldots, \\
\frac{1}{7} &= 0.\underline{142857} = 0.14285714285714\ldots, \\
\frac{1373}{3330} &= 0.4\underline{123} = 0.4123123123123\ldots
\end{aligned}$$

A real number that is not rational is called *irrational*. It is well-known that $\sqrt{2}$ and π are irrational. The proof that π is irrational is a bit complicated but the proof that $\sqrt{2}$ is, is actually quite easy.

Theorem 1.3 (Irrational roots) *The root $\sqrt{n}$ of a positive natural number is either also a positive natural number, or it is irrational.*

Proof Let n be a positive natural number. Assume that $\sqrt{n} = k/m$ for some positive natural numbers k, m. Then we can decompose k and m into prime factors as in formula (1.2):

$$k = 2^{k_2} \cdot 3^{k_3} \cdot 5^{k_5} \cdots \quad \text{and} \quad m = 2^{m_2} \cdot 3^{m_3} \cdot 5^{m_5} \cdots$$

It follows that

$$n = \frac{k^2}{m^2} = \frac{2^{2k_2} \cdot 3^{2k_3} \cdot 5^{2k_5} \cdots}{2^{2m_2} \cdot 3^{2m_3} \cdot 5^{2m_5} \cdots}.$$

Since n is a natural number, m^2 must be a divisor of k^2, which is possible only if $2k_p \geq 2m_p$ for each prime number p. But then $k_p \geq m_p$ for each p, which means that m is a divisor of k. But if m is a divisor of k, then $\sqrt{n}$ is a natural number. ∎

1.4 Euclidean space

An ordered sequence of real numbers $(a_1, \ldots, a_n)$ is a *vector* of length n. Vectors of length n can be used to describe a point in n-dimensional space. In particular, vectors of length two describe points in an infinite plane and vectors of length three describe points in the three-dimensional world around us. Vectors of length four can be used to describe events, where the first three numbers describe the position, and the fourth number is the time when the even happens. When we use a vector $(a_1, \ldots, a_n)$ of length n to describe a point in n-dimensional space, the individual numbers $a_1, \ldots, a_n$ are called the *coordinates* of the vector.

Mathematicians like to have a short way of writing down a vector, and for that reason (when it is clear from the context what the dimension n is) they often abbreviate $(a_1, \ldots, a_n)$ as $\vec{a}$. Often, they even simply write a instead of $\vec{a}$ but for the moment being we will keep the arrow on top to remind ourselves that this is a vector and not a number. Instead of thinking of vectors as points in space, we can also think of them as arrows that have a length and a direction. For example, in two-dimensional space:

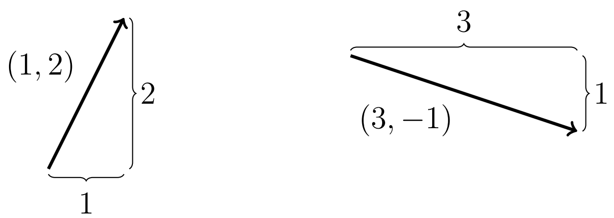

We denote the length of a vector by

$$\|\vec{a}\| := \sqrt{a_1^2 + \cdots + a_n^2}.$$

If the dimension n is two, then *Pythagoras' theorem* tells us that this is indeed the lenght of the arrow in the usual sense of the word. This is in fact also true in three dimensions, a fact that we do not prove here. In dimension one, vectors are just real numbers. In this case, we use the notation

$$|a| := \sqrt{a^2}$$

and we call $|a|$ the *absolute value* of a. Note that this is the same as the length of a, viewed as a vector of length one.

We can add vectors (of the same length) according to the definition:

$$\vec{a} + \vec{b} := \big(a_1 + b_1, \ldots, a_n + b_n\big),$$

i.e., the coordinates of $\vec{a} + \vec{b}$ are obtained by adding the corresponding coordinate of $\vec{a}$ and $\vec{b}$. In a picture, this amounts to the displacement we get by placing the vectors behind each other:

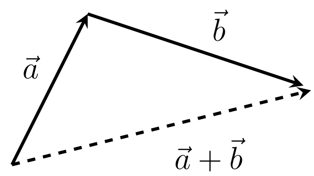

We can also multiply vectors by real numbers according to the definition:

$$a\vec{b} = \big(ab_1, \ldots, ab_n\big),$$

i.e., the coordinates of $a\vec{b}$ are obtained by multiplying each coordinate of $\vec{b}$ by a. If $a > 0$, then in a picture, this means that we multiply the length of $\vec{b}$ by a while keeping the direction as it was. If $a < 0$, then this means that we multiply the length of $\vec{b}$ by $|a|$ and reverse the direction.

It is well-known that the *surface* of a rectangle with sides a and b is equal to the product ab. Similarly, the *volume* of a cuboid with sides a, b, and c is equal to the product abc.

Chapter 2

Sets and functions

2.1 Naive set theory

A *set* is a collection of objects, that are called the *elements* of the set. Examples of sets are:

$$\begin{aligned}
\mathbb{R} &:= \text{the set of real numbers,}\\
\mathbb{Q} &:= \text{the set of rational numbers,}\\
\mathbb{Z} &:= \text{the set of integers,}\\
\mathbb{N} &:= \text{the set of natural numbers,}\\
\mathbb{N}_+ &:= \text{the set of positive natural numbers,}\\
\emptyset &:= \text{the empty set, which has no elements.}
\end{aligned}$$

The notation $a \in A$ (repectively, $\notin$) means that a is (repectively, is not) an element of the set A. So instead of writing "a is a real number" we can write more shortly "$a \in \mathbb{R}$". A set cannot contain an element more than once. In other words an object is either an element of a set, or not an element of a set, but there is no such concept as "being twice an element" of a set.

Two sets A and B are equal if each element of A is also an element of B and vice versa. As a result, all empty sets are equal, or to put it differently, there is only one empty set. For reasons that will become clear later in this chapter, a set can never be an element of itself. We say that A is a *subset* of B if each element of A is also an element of B. We denote this as $A \subset B$. Thus $A = B$ is the same as $A \subset B$ and $A \supset B$. If B is a set, and some of the elements of B have a property, and others not, then we use the notation

$$A := \{a \in \mathbb{R} : a \text{ has the property}\}$$

to denote the subset of B consisting of all elements of B that have that property. We use this notation even if all elements of B have the given property, in which case $A = B$, or if none of them have it, in which case $A = \emptyset$. For example:

$$\begin{aligned}
\mathbb{N} &:= \{n \in \mathbb{Z} : n \geq 0\}, \\
\mathbb{Q} &:= \{q \in \mathbb{R} : \text{there exist } n, m \in \mathbb{Z} \text{ with } m \neq 0 \text{ such that } q = n/m\}, \\
\emptyset &:= \{x \in \mathbb{R} : x < 0 \text{ and } x > 1\}.
\end{aligned}$$

If A and B are sets, then we let $A \cap B$ denote the *intersection* of A and B, we let $A \cup B$ denote the *union* of A and B, and we let $A \backslash B$ denote the *difference* of A and B, which are defined as

$$\begin{aligned}
A \cap B &:= \{x : x \in A \text{ and } x \in B\}, \\
A \cup B &:= \{x : x \in A \text{ or } x \in B\}, \\
A \backslash B &:= \{x : x \in A \text{ and } x \notin B\}.
\end{aligned}$$

Here we use the word "or" in an inclusive way. So, when we wite "or" between two properties, we mean that either only the first property holds, or only the second, *or both*. Note that the notation here is a bit different from our previous formulas. We could also have written

$$\begin{aligned}
A \cap B &:= \{x \in A : x \in B\}, \\
A \backslash B &:= \{x \in A : x \notin B\},
\end{aligned}$$

but there is no such formula for $A \cup B$, because a priori, before we have defined $A \cup B$, we do not have at our disposition a set that $A \cup B$ is a subset of. For *finite sets*, i.e., sets that have only a finite number of elements, we also write down a set simply by listing all its elements. For example:

$$\begin{aligned}
\{n \in \mathbb{Z} : 0 < n < 5\} &:= \{1, 2, 3, 4\}, \\
\{a \in \mathbb{R} : a = -a\} &:= \{0\}, \\
\{a \in \mathbb{R} : a^2 = a\} &:= \{0, 1\}.
\end{aligned}$$

Note that in such a list, each element can only occur once. Also, the order in which we list the elements does not matter, so

$$\{1, 2, 3, 4\} = \{3, 1, 2, 4\}.$$

In an informal text, we sometimes use this sort of notation that simply lists all elements also for infinite sets, when it is clear how the sequence continues:

$$\begin{aligned}
\{n \in \mathbb{N}_+ : 2 \text{ is a divisor of } n\} &= \{2, 4, 6, \ldots\}, \\
\{n \in \mathbb{N}_+ : n \text{ is a square}\} &= \{1, 4, 9, 16, 25, \ldots\}.
\end{aligned}$$

Intervals are subsets $I \subset \mathbb{R}$ of the real line with the property that if $a \in I$ and $b \in I$, then all numbers that lie between a and b are also elements of I. Our notation for intervals of finite length is as follows:

$$\begin{aligned}
[a, b] &= \{x \in \mathbb{R} : a \leq x \leq b\}, \\
(a, b] &= \{x \in \mathbb{R} : a < x \leq b\}, \\
[a, b) &= \{x \in \mathbb{R} : a \leq x < b\}, \\
(a, b) &= \{x \in \mathbb{R} : a < x < b\}.
\end{aligned}$$

We also write:

$$\begin{aligned}
[a, \infty) &= \{x \in \mathbb{R} : a \leq x\}, \\
(a, \infty) &= \{x \in \mathbb{R} : a < x\}, \\
(\infty, a] &= \{x \in \mathbb{R} : x \leq a\}, \\
(\infty, a) &= \{x \in \mathbb{R} : x < a\}.
\end{aligned}$$

Note that according to our definition, $\emptyset$ and $\mathbb{R}$ are also intervals.

The *Carthesian product* of two sets A and B is the set of all ordered pairs (a, b) with $a \in A$ and $b \in B$. We denote the Carthesian product of A and B as $A \times B$. So

$$A \times B := \big\{(a, b) : a \in A,\ b \in B\big\}.$$

The Carthesian product of three or more sets is defined similarly. We write

$$A^n := \underbrace{A \times A \times \cdots \times A}_{n \text{ times}}.$$

In particular,

$$\mathbb{R}^2 = \mathbb{R} \times \mathbb{R} \quad \text{and} \quad \mathbb{R}^3 = \mathbb{R} \times \mathbb{R} \times \mathbb{R}$$

are the two- and three dimensional euclidean spaces discussed in Section 1.4. Also,

$$[0, 1]^2 = [0, 1] \times [0, 1] = \big\{(a_1, a_2) : 0 \leq a_1 \leq 1 \text{ and } 0 \leq a_2 \leq 1\big\}$$

is a square.

It is even possible to make the Carthesian product of infinitely many sets. In particular, we let

$$A^{\mathbb{N}_+} := \big\{(a_1, a_2, \ldots) : a_k \in A \text{ for all } k \in \mathbb{N}_+\big\}$$

denote the set whose elements are infinite sequences of elements of A.

If A is a finite set, then we let $|A|$ denote the number of its elements, which is a natural number. Sums and products of natural numbers have a natural interpretation in set theory. In particular,

- $A \cap B = \emptyset$ implies that $|A \cup B| = |A| + |B|$,
- $|A \times B| = |A| \cdot |B|$.

If $A \cap B = \emptyset$, then we say that the sets A and B are *disjoint*.

Two mathematical symbols often come in handy when discussing sets. These are the symbols

$$\forall \text{ “for all”} \quad \text{and} \quad \exists \text{ “there exists”}.$$

So, for example, if $A_1, \ldots, A_n$ are sets, then

$$\begin{aligned} A_1 \cap \cdots \cap A_n &= \big\{a : a \in A_k \ \forall k \in \{1, \ldots, n\}\big\}, \\ A_1 \cup \cdots \cup A_n &= \big\{a : \exists k \in \{1, \ldots, n\} \text{ such that } a \in A_k\big\}. \end{aligned}$$

Sometimes we abbreviate even more and write “:” instead of “such that”. So $\{a : a \in A \text{ and } a \in B\}$ can be read as “the set of all a such that a is an element of A and a is an element of B”.

2.2 Functions

Let A and B be sets and let $F \subset A \times B$ be a subset of their Carthesian product with the property that

$$\text{for all } a \in A \text{ there is a unique } b \in B \text{ such that } (a, b) \in F.$$

In such a situation, it is convenient to have notation that tells us what b is, if we know a. Indeed, we can define

$$f(a) := \text{the unique } b \in B \text{ such that } (a, b) \in F.$$

Mathematicians say in such a situation that f is a *function* from the set A into the set B. The set F is called the *graph* of f. On a very formal level, a function and its graph are sort of the same thing, but in practice, we think about them differently. When we think about a graph, we really imagine a subset of the Carthesian product space $A \times B$. When we think about a function, we think about an object that takes an element the set A set, and turns it into an element of the set B. We often write

$$\text{Let } f : A \to B \text{ be a function}$$

to indicate that f goes “from A to B”, so to say. We have already seen many examples of functions:

- The function $f : \mathbb{R} \to \mathbb{R}$ defined as $f(a) := -a$ takes a real number and turns it into minus that number.
- The function $f : \mathbb{R}^n \to \mathbb{R}$ defined as $f(\vec{a}) := \|\vec{a}\|$ takes a vector gives out its length.
- The function $f : (\mathbb{R}\backslash\{0\}) \to \mathbb{R}$ defined as $f(a) := 1/a$ takes a nonzero real number and turns it into its inverse.
- The function $f : [0, \infty) \to [0, \infty)$ defined as $f(a) := \sqrt{a}$ takes a non-negative real number and gives its root.

A function f is:

- *surjective* if $\forall b \in B\ \exists a \in A$ such that $f(a) = b$.
- *injective* (also called *one-to-one*) if $\forall b \in B$, there exists at most one $a \in A$ such that $f(a) = b$.
- a *bijection* if it is both surjective and injective.

If $f : A \to B$ is a bijection, then the set

$$F^{-1} := \big\{(b, a) : (a, b) \in F\big\}$$

is the graph of a function $f^{-1} : B \to A$ that satisfies

$$f^{-1}\big(f(a)\big) = a \quad \forall a \in A \quad \text{and} \quad f\big(f^{-1}(b)\big) = b \quad \forall b \in B.$$

The function f^{-1} is called the *inverse* of the function f. Note that f^{-1} is also a bijection and f is the inverse of f^{-1}. Here are some examples of inverse functions that we have already seen.

- For each real number x, the function $f : \mathbb{R} \to \mathbb{R}$ defined as $f(x) := x+a$ is a bijection, and $f^{-1}(x) = x - a$ is its inverse.
- For each real number $a \neq 0$, the function $f : \mathbb{R} \to \mathbb{R}$ defined as $f(x) := ax$ is a bijection, and $f^{-1}(x) = x/a$ is its inverse.
- For each integer $n \geq 2$, the function $f : [0, \infty) \to [0, \infty)$ defined as $f(x) := x^n$ is a bijection, and $f^{-1}(x) = \sqrt[n]{x}$ is its inverse.

As we can see from these examples, if $f : X \to Y$ is a bijection, then for each $x \in X$, the *equation* $f(x) = y$ has a unique solution, which is given by $x = f^{-1}(y)$. Subtraction was invented to solve equations of the form

$x+a=y$. Division was invented to solve equations of the form $ax=y$. And the n-th root was invented to solve equations of the form $x^n=y$.

If X, Y, and Z are sets and $f : X \to Y$ and $g : Y \to Z$ are function, then we can define a function $h : X \to Z$ by $h(x) := g(f(x))$. The function h is called the *composition* of f and g and denoted by $h = g \circ f$. Thus:

$$g \circ f(x) := g\big(f(x)\big).$$

If $f : X \to Y$ is a bijection, then $f^{-1}(f(x)) = x$ for all $x \in X$, and $f(f^{-1}(y)) = y$ for all $y \in Y$, i.e., the function $f^{-1} \circ f : X \to X$ and $f \circ f^{-1} : Y \to Y$ are the *identity function* on X and Y, respectively, which map each element of X, respectively Y, into itself.

2.3 Relations

Let A be a set and let $R \subset A \times A$ be any subset of the Carthesian product of A with itself. Then R represents a property that an ordered pair (a, b) of elements of A can have: either $(a, b) \in R$ or $(a, b) \notin R$. Often, it is nice to use a bit different notation for this. For example, we can write

$$a \vdash b \text{ if } (a, b) \in R \quad \text{and} \quad a \not\vdash b \text{ if } (a, b) \notin R.$$

In this case, we call $\vdash$ a *relation*. Examples of relations that we have already seen are

$$< \quad \leq \quad > \quad \geq \quad =$$

Also ,if A is a set and $S := \{B : B \subset A\}$ is the set of all subsets of A, then $\subset$ is a relation on S. It is tempting to view also $\in$ as a relation, but if we do this, then it is not completely clear on which set this relation is defined.

Let $\vdash$ be a relation on a set A. Then we say that:

- $\vdash$ is *reflexive* if $a \vdash a$ for all $a \in A$,
- $\vdash$ is *symmetric* if $a \vdash b$ implies $b \vdash a$,
- $\vdash$ is *transitive* if $a \vdash b$ and $b \vdash c$ imply $a \vdash c$.

For example, $\leq$ and $=$ are reflexive, but $<$ is not. The relation $=$ is symmetric but $<$ and $\leq$ are not. Finally, all of the relations $<$, $\leq$, and $=$ are transitive.

An *equivalence relation* is relation that is reflexive, symmetric, and transitive. If $\sim$ is an equivalence relation on a set A and a is an element of A, then the subset of A defined as

$$[a] := \{b \in A : b \sim a\}$$

is called the *equivalence class* of a. Note that $a \in [a]$ by reflexivity. If a and b are elements of A such that $a \sim b$, then using symmetry and transitivity, it is easy to see that $[a] = [b]$. On the other hand, if $a \not\sim b$, then $[a] \cap [b] = \emptyset$, i.e., $[a]$ and $[b]$ are disjoint.

Equivalence relations are used a lot in mathematics. Often, we start with a set A whose elements a contain "too much information". By forgetting some of this information, one naturally arives at equivalence classes. In the next section, we demonstrate this on the example of calculation modulo a natural number.

2.4 Calculation modulo a natural number

Let $q \geq 2$ be a natural number. Then we can define an equivalence relation on $\mathbb{Z}$ by setting

$$n \sim m \text{ if and only if } \exists k \in \mathbb{Z} \text{ such that } k = m + kq.$$

(The reader should check that this relation is indeed reflexive, symmetric, and transitive.) Then

$$[k] = \{\ldots, k - 2q, k - q, k, k + q, k + 2q, \ldots\}.$$

The set of all equivalence classes is

$$\mathbb{Z}/q := \big\{[0], [1], \ldots, [q-1]\big\}.$$

Note that the set on the right indeed lists all equivalence classes, since $[q] = [0]$, $[q+1] = [1]$, etcetera, and on the other side, $[-1] = [q-1]$, $[-2] = [q-2]$ etcetera. The interesting thing about the equivalence relation $\sim$ is that (check!)

- $n \sim n'$ and $m \sim m'$ implies $n + m \sim n' + m'$,
- $n \sim n'$ and $m \sim m'$ implies $nm \sim n'm'$.

As a result of this, we can define the sum and product of elements of $\mathbb{Z}/q$ as follows:

$$[n] + [m] := [n + m] \quad \text{and} \quad [n] \cdot [m] := [nm] \qquad \forall n, m \in \mathbb{Z}.$$

Note that a priori, it is not clear that these are valid definitions. A priori, it is conceivable that there exist integers n, n', m, and m' such that $[n'] = [n]$ and $[m'] = [m]$, but $[n + m] \neq [n' + m']$. If that were the case, then according to our definition, $[n]+[m]$ would have to be equal to both $[n+m]$ and $[n'+m']$, which is impossible. But because of the special properties of $\sim$ listed above, this problem does not occur and our definition is valid.

The following theorem shows that if q is a prime number, then $\mathbb{Z}/q$ is a *field*, as defined in Section 1.2.

Theorem 2.1 (Calculation modulo q) *The sum and product on $\mathbb{Z}/q$ satisfy the properties (i)–(viii) of Section 1.2, where $[0]$ plays the role of 0 and $[1]$ plays the role of 1. If q is a prime number, then property (ix) holds too.*

Instead of $n \sim m$, one usually writes:

$$n = m \mod(q).$$

This is pronounced as "n and m are equal modulo q". Theorem 2.1 says that we can calculate modulo q more or less in the same way as we calculate normally. If q is a prime number, then we can even divide two numbers modulo q.

2.5 Cardinality

Recall that a function $f : A \to B$ is *injective* if $\forall b \in B$, there exists at most one $a \in A$ such that $f(a) = b$. If there exists an injection from A to B, then B must, in a sense, be "at least as large" as A. To indicate this, for two sets A and B, let us write $A \prec B$ if there exists an injection $f : A \to B$. Since the identity map is an injection, and since the composition of two injections is again an injection, the relation $\prec$ is reflexive and transitive.

Recall that a function $f : A \to B$ is a *bijection* if it is both surjective and injective, i.e., for each $b \in B$ there exists precisely one $a \in A$ such that $f(a) = b$. If there exists a bijection from A to B, then B must, in a sense, be "equally large" as A. To indicate this, for two sets A and B, let us write $A \sim B$ if there exists a bijection $f : A \to B$.

The following theorem was first proved by Dedekind. In line with Stigler's law,[1] the theorem is known as the Schroeder-Bernstein theorem.

Theorem 2.2 (Schroeder-Bernstein) *$A \prec B$ and $B \prec A$ imply $A \sim B$.*

Clearly,[2] $\sim$ is an equivalence relation. We will be interested in the equivalence classes. Let us first look at finite sets. If A and B are finite sets,

[1]Stigler's law states that no scientific discovery is named after its original discoverer. True to its name, Stigler is not the original discoverer of this "law".

[2]Well, not so clearly. It is true that $\sim$ is reflexive, symmetric and transitive. But it is not so clear *on which set* $\sim$ is defined. Naively, one would say, "on the set of all sets", but as we will later see, there is no such thing. We ignore this problem for the moment but will get back to it later.

say:

$$A = \{a_1, \dots, a_n\} \quad \text{and} \quad B = \{b_1, \dots, b_m\},$$

then one has $A \sim B$ if and only if $n = m$, i.e., if A and B have the same number of elements. Thus, each natural number n corresponds to an equivalence class, which consists of all sets that have precisely n elements. Note that $\emptyset$ is the only set with zero elements. As we have already seen in Section 2.1, the sum and product on $\mathbb{N}$ have a natural interpretation in terms of sets, since $A \cap B = \emptyset$ implies that $|A \cup B| = |A| + |B|$ and $|A \times B| = |A| \cdot |B|$.

Now let us look at infinite sets. If A is a finite set and $b \notin A$, then $|A \cup \{b\}| = |A| + 1$ and hence $A \cup \{b\} \not\sim A$. But for infinite sets, this is no longer true. Indeed, $\mathbb{N}$ contains precisely one element that is not contained in $\mathbb{N}_+$, yet the function $f : \mathbb{N}_+ \to \mathbb{N}$ defined as $f(n) := n - 1$ is a bijection. It is even possible to define a bijection from the natural numbers to the even natural numbers, by setting $f(n) := 2n$. This phenomenon is informally known as "Hilbert's hotel".

In view of this, one might be tempted to think that all infinite sets are "equally large" in the sense of $\sim$, but this is *not true*. Recall from Section 2.1 that

$$\{0,1\}^{\mathbb{N}_+} = \big\{(x_1, x_2, \dots) : x_k \in \{0,1\} \text{ for all } k \in \mathbb{N}_+\big\}$$

denotes the set of all infinite sequences of zeros and ones. The next theorem says that this set is strictly larger than $\mathbb{N}_+$.

Theorem 2.3 (An uncountable set) *One has $\mathbb{N}_+ \prec \{0,1\}^{\mathbb{N}_+}$. On the other hand, $\mathbb{N}_+ \not\sim \{0,1\}^{\mathbb{N}_+}$.*

Proof We can define an injection $f : \mathbb{N}_+ \to \{0,1\}^{\mathbb{N}_+}$ by setting

$$f(k) := (0, \dots, 0, 1, 0, \dots).$$
$$\uparrow \ k\text{-th coordinate}$$

On the other hand, we claim that a function $f : \mathbb{N}_+ \to \{0,1\}^{\mathbb{N}_+}$ can never be surjective, and as a result, there exist no bijections between $\mathbb{N}_+$ and $\{0,1\}^{\mathbb{N}_+}$. To see this, let $f : \mathbb{N}_+ \to \{0,1\}^{\mathbb{N}_+}$ be a function. Then we can write

$$f(k) = \big(f_1(k), f_2(k), f_3(k), \dots\big),$$

where for each $n \in \mathbb{N}_+$, $f_n : \mathbb{N}_+ \to \{0,1\}$ is a function. Now we can define an infinite sequence $(x_1, x_2, \dots)$ of zeros and ones by setting

$$x_k := \begin{cases} 1 & \text{if } f_k(k) = 0, \\ 0 & \text{if } f_k(k) = 1. \end{cases}$$

Then $x_k \neq f_k(k)$ and as a result

$$\big(x_1, x_2, x_3, \dots\big) \neq \big(f_1(k), f_2(k), f_3(k), \dots\big) \qquad \forall k \in \mathbb{N}_+.$$

This shows that f is not surjective. ∎

If a set A satisfies $A \sim \mathbb{N}_+$, then we say that A is *countably infinite*. We have already seen that $\mathbb{N}$ is countably infinite. The same is true for $\mathbb{Z}$, since we can define a bijection $f : \mathbb{N}_+ \to \mathbb{Z}$ by setting:

$$f(1) := 0, \quad f(1) := 1, \quad f(2) := -1, \quad f(3) := 2, \quad f(4) := -2, \quad f(5) := 3,$$

etcetera. We claim that the Carthesian product $\mathbb{N} \times \mathbb{N}$ is also countably infinite. Indeed, we can define a bijection $f : \mathbb{N}_+ \to \mathbb{N} \times \mathbb{N}$ by

$$\begin{array}{llll} f(1) := (0,0) & f(2) := (0,1) & f(3) := (1,0) & f(4) := (0,2) \\ f(5) := (1,1) & f(6) := (2,0) & f(7) := (0,3) & f(8) := (1,2) \\ f(9) := (2,1) & f(10) := (3,0) & f(11) := (0,4) & f(12) := (1,3) \end{array}$$

etcetera, which goes through $\mathbb{N} \times \mathbb{N}$ as follows:

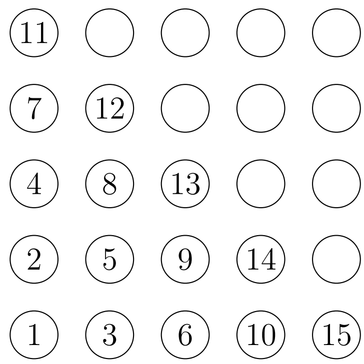

We moreover note that:

- Each infinite subset $A \subset \mathbb{N}_+$ satisfies $A \sim \mathbb{N}_+$.

As a result, there are no infinite sets that are "smaller" than $\mathbb{N}$. A set is called *countable* if it is either finite or countably infinite. Sets that are not countable are called *uncountable*. Note that by what we just said, such sets are "larger" than $\mathbb{N}_+$. Theorem 2.3 shows that the set $\{0,1\}^{\mathbb{N}}$ consisting of infinite sequences of zeros and ones is uncountable.

Theorem 2.4 (Reals are uncountable) *The set $\mathbb{R}$ is uncountable.*

Proof By Theorem 2.3, the set $\{0,1\}^{\mathbb{N}}$ of infinite sequences of zeros and ones is uncountable. Let S denote the set of real numbers between zero and one whose decimals only consist of zeros and ones. For example, an element of S could look like this: $0.110100011010\ldots$. Then clearly, there exists a bijection between $\{0,1\}^{\mathbb{N}}$ and S, so these sets are "of the same size". Thus, S must be uncountable. Since $S \subset \mathbb{R}$, this seems to say that $\mathbb{R}$ must be even larger.

Indeed, if $\mathbb{R}$ would be countably infinite, then there would exist a bijection $f : \mathbb{N}_+ \to \mathbb{R}$. Let $A := \{n \in \mathbb{N}_+ : f(n) \in S\}$. Then the restriction of f to A is a bijection from A to S. But since $A \subset \mathbb{N}_+$, the set A is countable, so this contradicts the fact that S is uncountable. ∎

2.6 Axiomatic set theory

Mathematics is first and foremost the art of proving statements. Statements cannot be proved out of nothing. Indeed, a mathematical proof is always an argument that shows that if certain statements are true, then certain other statements must also be true. Therefore, every mathematician needs a set of elementary statements to start with, from which all other statements can then be derived. These elementary statements are called *axioms*.

The first set of axioms to be widely used are the axioms for planar geometry formulated by Euclid in his book Elements, written approximately 300 BC in Alexandria. He started his book with 5 "common notions", which sound like definitions of the main objects of interest such as points and lines, and 5 "postulates", which sound more like statements about properties of these objects. In what follows, he derives a large number of nontrivial statements from these common notions and postulates, although critics have pointed out several places where he seems to use "obvious" facts that nevertheless are not part of his original $5+5$ assumptions.

Nevertheless, the Elements were hugely influential. In the 17th century, when modern calculus was being developed, different "sorts" of mathematics (based on different sets of elementary assumptions) started to rise to prominence. Nevertheless, Euclid's geometry was still considered the golden standard of mathematical rigor, and mathematicians like Newton made an effort to show their new differential calculus could be derived from Euclid's axioms. The Elements remained prominent till the late nineteenth century and influenced school books well into the twentieth century.

Depending on what one is interested in, one can use various systems of axioms. For example, there exist axioms for the real numbers, from which all properties of the real numbers[3] can be deduced. These axioms include the properties (i)–(ix) from Section 1.2. As we have seen in Section 2.4, the axioms (i)–(ix) from Section 1.2 are not enough to characterize the real numbers uniquely, as they are also satisfied by $\mathbb{Z}/q$ when q is a prime number. Nevertheless, by adding some additional axioms, one can uniquely characterize the real numbers.

Obviously, one can not just write down any collection of axioms. In particular, one would like to know that a given system of axioms, such as the one for the real numbers, is *consistent*. With this we mean that if with a system of axioms it is possible to prove a certain statement, then it should not be possible to also prove the converse of that statement. Unfortunately, for most[4] systems of axioms, it is impossible to be sure they are consistent. In particular, Gödel proved in 1930 that the consistency of a system of axioms cannot be proved within their own system. In view of this, mathematicians like to work with well-known systems of axioms that they trust.

Today, almost all of mathematics is based on the axioms of set theory, which were developed in the early twentieth century. The Zermelo-Fraenkel axioms are best known but certain other systems of axioms, such as those developed by Von Neumann, Bernays, and Gödel, also enjoy wide popularity. The latter system is a bit stronger than the Zermelo-Fraenkel axioms and more suitable for category theory, but in practice, for most branches of mathematics, the differences between the axiomatic systems play no role in daily life.

Early, "naive" set theory defined sets as objects of the form

$$\{a : a \text{ has property } \phi\},$$

which can be read as "the set of all a that have property ϕ". In 1901, Bertrand Russell noted that if one allows this sort of definitions without any restriction, then it would seem perfectly allowed to define a set R by:

$$R := \{A : A \notin A\}$$

i.e., R is the set of all sets that do not contain themselves as an element. One can now ask whether R contains itself as an element. If it does, then that implies it does not, but if it does not, then that implies that it does. Thus, the statement $R \in R$ cannot be true, but it cannot be false either. This is known as *Russel's paradox*. As a result, naive set theory is *inconsistent*.

Modern set theory, such as the theory of Zermelo and Fraenkel, carefully avoids Russel's paradox by restricting which objects can be called "sets". If

[3] At least, almost all properties that one is usually interested in.

[4] In particular, all systems that are "rich" enough to define the natural numbers.

A is a set and ϕ is a property that elements of A can have, then we can always define a subset of A as $\{a \in A : a \text{ has property } \phi\}$. However, in modern set theory, there exists no "set of all sets", so the definition of Russel's set R is not allowed. Nevertheless, some versions of modern set theory allow for other objects called *categories* that are more general than sets. In such systems, there is a category of all sets.

Axiomatic set theory is quite complicated. Luckily, in practice, we do not need it very much as long as we make sure that we always start from some well-known sets and define new sets in terms of these old sets by well-known legal operations, such as the Carthesian product or defining a subset of an existing set in terms of some property of its elements.

However, the fact that there exists no "set of all sets" means that some of our discussion in Section 2.5 was a bit too simplistic. In particular, the relation $\sim$ is defined on the category of all sets, which is not a set, and the resulting equivalence classes, which are all possible cardinal numbers, also form a catagory but not a set. Nevertheless, we can identify the cardinal numbers of finite sets with the natural numbers, which form a set.

A well-known axiom that is usually included in the axioms of set-theory is the *axiom of choice*:

- If X is a set whose elements are nonempty sets, and $Y := \{x : x \in A \text{ for some } A \in X\}$ is the union of all sets $A \in X$, then there exists a function $f : X \to Y$ such that $f(A) \in A$.

Informally, this says that for given a collection of nonempty sets, it is possible to choose one element out of each set. Another equivalent formulation is that the Carthesian product of (maybe infinitly many) nonempty sets is nonempty.

Here is a simple example of an application of the axiom of choice. Recall that in Section 2.5 we wrote $A \prec B$ if there exists a injection $f : A \to B$.

Theorem 2.5 (Injections and surjections) *Let A and B be nonempty sets. Then there exists an injection $f : A \to B$ if and only if there exists a surjection $g : B \to A$.*

Proof Assume that there exists an injection $f : A \to B$. Let $B' := \{b \in B : \exists a \in A \text{ such that } f(a) = b\}$ denote the *image* of A under the map f. Then $f : A \to B$ is a bijection. Since A is nonempty, there exists some $a \in A$. Now we can define a surjection $g : B \to A$ by setting $g(b) := f^{-1}(b)$ if $b \in B'$ and $g(b) := a$ otherwise.

Conversely, assume that there exists a surjection $g : B \to A$. Then for each $a \in A$, the set $\{b \in B : g(b) = a\}$ is nonempty. By the axiom of choice, for each $a \in A$, we can choose an element $f(a) \in \{b \in B : g(b) = a\}$. Then $f : A \to B$ is an injection. ∎

Although the axiom of choice sounds innocent and has nice consequences such as Theorem 2.5, we will later see that it can also have very counterintuitive consequences. These counterintuitive results only occur when the axiom is applied to a uncountable collections of sets. In view of this, mathematicians sometimes work with alternative systems of axioms, where the axiom of choice is valid only for countable collections of sets. It is interesting that Theorem 2.2 can be proved without the axiom of choice.

We already mentioned that Gödel proved that for any "decent"[5] system of axioms, it is impossible to prove that it is consistent. He also proved that any such system must be *incomplete* in the sense that there are statements that are *undecidable* in the sense that they cannot be proved, but their converse can also not be proved. Recall from Section 2.5 that $\mathbb{R}$ is larger than $\mathbb{N}$ in the sense that $\mathbb{N} \prec \mathbb{R}$ but $\mathbb{N} \not\sim \mathbb{R}$. One may ask if there exist sets A such that $\mathbb{N} \prec A \prec \mathbb{R}$ but $\mathbb{N} \not\sim A \not\sim \mathbb{R}$. Cantor conjectured in 1878 that such sets do not exist. This conjecture became known as the *continuum hypothesis*. In 1963, Paul Cohen proved that the question whether such sets exist is undecidable with the usual axioms for set theory. More precisely, he showed that assuming that the usual axioms for set theory are consistent, one can add the continuum hypothesis and still have a consistent set of axioms, but one can also postulate the converse of the continuum hypothesis and again have a consistent set of axioms.

2.7 Construction of the rational numbers

As we have already seen in Sections 2.1 and 2.5, there is a one-to-one correspondence between the cardinals of finite sets and the natural numbers $\mathbb{N}$. Moreover, addition and multiplication have natural interpretations in terms of the union of two disjoint sets and the Carthesian product. In view of this, one can easily prove that the sum and product satisfy properties (i)–(vii) from Section 1.2.

In order to also satisfy property (viii), we need to extend $\mathbb{N}$. One can prove that $\mathbb{Z}$ is the smallest possible extension of $\mathbb{N}$ that satisfies (i)–(viii). We will not give a full proof here but sketch the main line. Since $0 + 0 = 0$, (viii) can be satisfied for $a = 0$ by setting $-0 := 0$. For all elements $a \in \mathbb{N}_+$ we introduce a new "number" called $-a$. To extend the definition of the sum to $\mathbb{Z}$, we first prove that for all $a, b \in \mathbb{N}$ with $0 \leq a \leq b$, there exists a unique

[5] In the sense that it is "rich" enough to define the natural numbers.

element $c \in \mathbb{N}$ such that $b = a + c$. Denoting this element by $c := b - a$, we define for all $a, b \in \mathbb{N}$ with $0 \leq a \leq b$:

$$(-a) + b = b + (-a) := b - a, \quad a + (-b) = (-b) + a := -(b - a),$$
$$\text{and} \quad (-a) + (-b) = (-b) + (-a) := -(a + b).$$

Next, we extend the definition of the product by defining, for $a, b \in \mathbb{N}$:

$$a \cdot (-b) := -ab, \quad (-a) \cdot b := -ab, \quad \text{and} \quad (-a) \cdot (-b) := ab.$$

With these definitions, one can prove that the sum and product on $\mathbb{Z}$ satisfy properties (i)–(viii) from Section 1.2.

In order to also satisfy property (ix) from Section 1.2, we need to introduce rational numbers. We can view a rational number as a pair of numbers n/m with $n \in \mathbb{Z}$ and $m \in \mathbb{N}_+$. We consider two pairs to be equivalent, which we denote by writing $n/m = n'/m'$, if there exist numbers $k \in \mathbb{N}_+$ and $k' \in \mathbb{N}_+$ such that $kn = k'n'$ and $km = k'm'$. We claim that this defines an equivalence relation. Reflexivity and symmetry are easy. To see that transitivity holds, assume that $n/m = n'/m'$ and $n'/m' = n''/m''$. Then there exist k, k' such that $kn = k'n'$ and $km = k'm'$, and moreover, there exist l', l'' such that $l'n' = l''n''$ and $l'm' = l''m''$. It follows that

$$(kl')n = k'l'n' = (k'l'')n'' \quad \text{and} \quad (kl')m = k'l'm' = (k'l'')m'',$$

proving that $n/m = n''/m''$. Formally, rational numbers then correspond to equivalence classes with respect to this equivalence relation. One can prove that there is a unique way to extend the definition of the sum and product to $\mathbb{Q}$ so that properties (i)–(vii) remain valid, and that $\mathbb{Q}$ also satisfies (viii) and (ix).

In the next chapter, we will see how the real numbers can be formally defined as limits of natural numbers.

Chapter 3

Limits

3.1 Limits and continuity

Let $(x_k)_{k\geq 1} = (x_1, x_2, \ldots)$ be an infinite sequence of real numbers. By definition, we say that the sequence $(x_k)_{k\geq 1}$ *converges* to a *limit* $x \in \mathbb{R}$ if

$$\forall \varepsilon > 0 \ \exists m \in \mathbb{N}_+ \text{ such that } |x_n - x| \leq \varepsilon \ \forall n \geq m.$$

This definition looks a bit complicated at first, but it expresses a simple idea: *as n increases, the points x_n approximate the limit point x.* The definition is illustrated in the picture below: if we think of the horizontal axis as time, then for every $\varepsilon > 0$, there exists an m such that from time m onwards, the sequence never leaves the interval $[x - \varepsilon, x + \varepsilon]$.

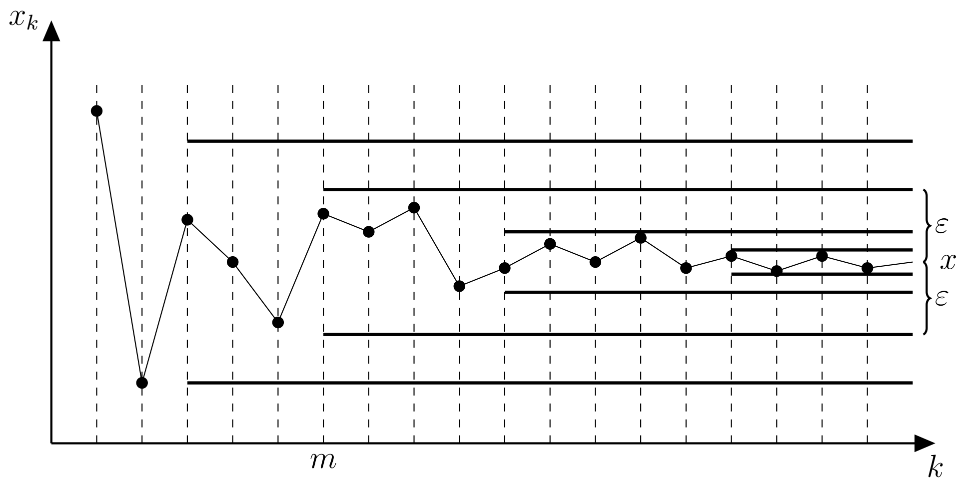

There exist a number of ways to write down that a sequence $(x_k)_{k\geq 1}$ converges to a limit x. The formulas

$$\lim_{k\to\infty} x_k = x, \qquad x_k \to x \text{ as } k \to \infty, \qquad x_k \underset{k\to\infty}{\longrightarrow} x,$$

all mean the same thing, namely, that $(x_k)_{k\geq 1}$ converges to the limit x. Note that not every sequence has a limit. For example, the seqences

$$(1,2,3,4,5,\ldots) \quad \text{and} \quad (1,0,1,0,1,\ldots)$$

do not converge to any limit. The following elementary properties of limits are easy to prove:

- A sequence can have only one limit, i.e., $x_k \underset{k\to\infty}{\longrightarrow} x$ and $x_k \underset{k\to\infty}{\longrightarrow} x'$ imply $x = x'$.
- $x_k \underset{k\to\infty}{\longrightarrow} x$ and $y_k \underset{k\to\infty}{\longrightarrow} y$ imply $x_k + y_k \underset{k\to\infty}{\longrightarrow} x + y$.
- $x_k \underset{k\to\infty}{\longrightarrow} x$ and $y_k \underset{k\to\infty}{\longrightarrow} y$ imply $x_k y_k \underset{k\to\infty}{\longrightarrow} xy$.

If $I \subset \mathbb{R}$ is an interval and $f : I \to \mathbb{R}$ is a function, then we say that f is *continuous* if:

$$x_k \underset{k\to\infty}{\longrightarrow} x \quad \text{implies} \quad f(x_k) \underset{k\to\infty}{\longrightarrow} f(x).$$

One can check that this is equivalent to:

$$\forall x \in I \text{ and } \varepsilon > 0\ \exists \delta > 0 \text{ such that } \big|f(y) - f(x)\big| \leq \varepsilon\ \forall y \in [x-\delta, x+\delta].$$

Informally, a real function is continuous if its graph can be drawn in a single stroke, without lifting the pen from the paper.

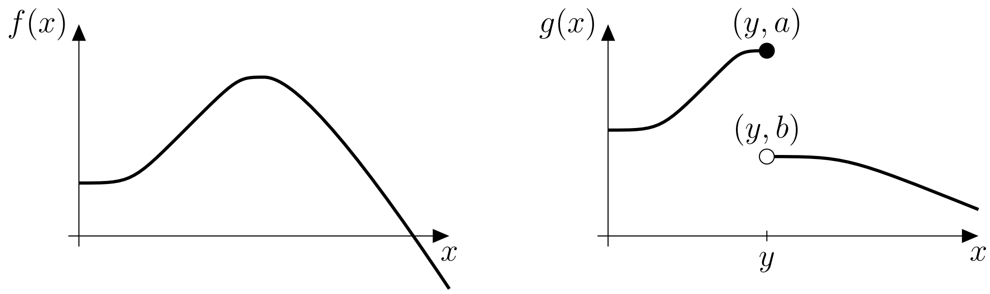

a continuous function a discontinuous function

The function g in the picture on the right is continuous everywhere except in the point y, where it jumps from the value a to the value b. The closed and open circle in the picture indicate that $g(y) = a$ and $g(y) \neq b$. If $(y_k)_{k\geq 1}$ is a sequence of real numbers such that $y_k \to y$, then we can distinguish the following three cases:

- If there exists an m such that $y_k \leq y$ for all $k \geq m$, then $g(y_k) \underset{k\to\infty}{\longrightarrow} a = g(y)$.

- If there exists an m such that $y_k > y$ for all $k \geq m$, then $g(y_k) \underset{k\to\infty}{\longrightarrow} b$.

- If both $y_k \leq y$ and $y_k > y$ occur for infinitely many values of k, then the limit $\lim\limits_{k\to\infty} g(y_k)$ does not exist.

Of course, discontinuous functions can be much more complicated than this. For example, we can define $h(x) := 1$ if x is a rational number and $h(x) := 0$ if x is irrational. We cannot draw the graph of such a function in an understandable way.

The following rules are easy to prove:

(i) The constant function $f(x) := 1$ is continuous on $\mathbb{R}$.

(ii) The function $f(x) := x$ is continuous on $\mathbb{R}$.

(iii) The function $f(x) := 1/x$ is continuous on the intervals $(-\infty, 0)$ and $(0, \infty)$.

(iv) If f and g are continuous, then $h(x) := f(x) + g(x)$ is continuous.

(v) If f and g are continuous, then $h(x) := f(x)g(x)$ is continuous.

(vi) If I and J are intervals and $f : I \to J$ and $g : J \to \mathbb{R}$ are continuous functions, then their concatenation $g \circ f : I \to \mathbb{R}$ is continuous.

Using these rules, we can prove for many functions that they are continuous. For example, the function $f(x) := x^2$ is continuous by (ii) and (v). Using also (i) and (iv), we see that $f(x) := 1 + x^2$ is continuous. Applying (iii) and (vi), we obtain that

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

is continuous.

3.2 Metric spaces

In Section 2.1, we introduced the notation $\mathbb{R}^n$ for the n-dimensional real space consisting of all n-dimensional vectors. In Section 1.4, we denotes such a vector by $\vec{x} = (x_1, \ldots, x_n)$ and defined its length as

$$\|\vec{x}\| := \sqrt{x_1^2 + \cdots + x_n^2}.$$

We call $\vec{0} := (0, \dots, 0)$ the *origin*. As already explained in Section 1.4, we can think of vectors as arrows, but we can also think about them as points in space. The distance between two points $\vec{x}$ and $\vec{y}$ is then equal to the length of the vector that connects them. We denote this distance by

$$d(\vec{x}, \vec{y}) := \|\vec{x} - \vec{y}\| = \sqrt{(x_1 - y_1)^2 + \cdots + (x_n - y_n)^2}. \tag{3.1}$$

In general, if A is a set and $d : A \times A \to \mathbb{R}$ is a function satisfying:

(i) $d(a, b) = 0$ if and only if $a = b$,

(ii) $d(a, b) = d(b, a)$,

(iii) $d(a, c) \leq d(a, b) + d(b, c)$,

then such a function d is called a *metric*. Property (iii) is called the *triangle inequality*. Note that properties (i)–(iii) imply $0 = d(a, a) \leq d(a, b) + d(b, a) = 2d(a, b)$ and hence $d(a, b) \geq 0$ for all a, b. One can check that the distance function in (3.1) is a metric. By definition, a *metric space* is a pair (A, d) where A is a set and d is a metric on A.

Let (A, d) be a metric space. By definition, a sequence $(a_k)_{k \geq 1}$ of elements of A converges to a limit $a \in A$ if

$$d(a_k, a) \underset{k \to \infty}{\longrightarrow} 0.$$

This is equivalent to

$$\forall \varepsilon > 0 \ \exists m \in \mathbb{N}_+ \text{ such that } d(a_n, a) \leq \varepsilon \ \forall n \geq m.$$

For each $r > 0$ and $a \in A$, we call

$$B_r(a) := \{b \in A : d(a, b) < r\}$$

the *open ball* of *radius* r around a. By definition, a set $B \subset A$ is *open* if

$$\forall b \in B \ \exists \varepsilon > 0 \text{ such that } B_\varepsilon(b) \subset B.$$

Note that the open ball $B_r(a)$ is indeed an open set, since $d(a, b) < r$ implies that $B_\varepsilon(b) \subset B_r(a)$ for all $\varepsilon < r - d(a, b)$. By definition, a set $B \subset A$ is *closed* if

$$b_k \in B \ \forall k \geq 1 \text{ and } b_k \underset{k \to \infty}{\longrightarrow} b \in A \quad \text{imply} \quad b \in B,$$

i.e., B contains all limits of sequences in B. For example, the *closed ball* of radius r around a, defined as

$$\overline{B}_r(a) := \{b \in A : d(a, b) \leq r\}$$

is a closed set.

In mathematics, it is custom to call only the most important results "theorems". A less important result is called a "proposition". Small results, often of a technical nature, are called "lemmas". We are ready to formulate our first lemma.

Lemma 3.1 (Open and closed sets) *Let (A,d) be a metric space and $B \subset A$. Then B is closed if and only if its complement $A\backslash B$ is open.*

Proof Assume that B is closed. We want to show that $A\backslash B$ is open, i.e., for all $a \in A\backslash B$ there exists an $\varepsilon > 0$ such that $B_\varepsilon(a) \subset A\backslash B$. Assume that this is not true. Then there exists an $a \in A\backslash B$ such that $B_{1/k}(a) \cap B \neq \emptyset$ for all $k \geq 1$. Now we can choose $b_k \in B_{1/k}(a) \cap B$. Then $d(b_k, a) < 1/k$ so $\lim_{k\to\infty} b_k = a$. Since $b_k \in B$ for each k this contradicts the assumption that B is closed.

Assume that $A\backslash B$ is open. We want to show that B is closed. Assume the converse. Then there exist $b_k \in B$ such that $d(b_k, a) \to 0$ for some $a \in A\backslash B$. This implies that for each $\varepsilon > 0$ we can find $b_k \in B$ such that $d(b_k, a) < \varepsilon$. In other words, $B_\varepsilon(a) \cap B \neq \emptyset$ for all $\varepsilon > 0$, which contradicts the assumption that $A\backslash B$ is open. ∎

If (A,d) is a metric space, then according to our definitions, the sets $\emptyset$ and A, viewed as subsets of A, are the same time both open and closed. By definition, a metric space is *connected* if these are the only subsets of A with this property, i.e., if

$$B \subset A \text{ is both open and closed} \qquad \text{implies} \qquad B = \emptyset \text{ or } B = A.$$

The set of real numbers, equipped with the metric $d(x,y) := |x-y|$, is a connected metric space. On the other hand, if we equip the space $A := [0,1] \cup [2,3]$ with the same metric, then A is not connected, since $B := [0,1]$ is both open and closed[1] as a subset of A. Also, if we equip the set of rational numbers $\mathbb{Q}$ with the metric $d(x,y) := |x-y|$, then $\mathbb{Q}$ is not connected, since the set $B := \{x \in \mathbb{Q} : x > \sqrt{2}\}$ is both open and closed as a subset of $\mathbb{Q}$. Indeed, it is easy to see that B is open, and the same is true for $\mathbb{Q}\backslash B = \{x \in \mathbb{Q} : x < \sqrt{2}\}$. Here we have used Theorem 1.3 which implies that $\sqrt{2} \notin \mathbb{Q}$.

The next lemma says that we can always "close" a set B by adding all limits of sequences in B. The set $\overline{B}$ is called the *closure* of B.

Lemma 3.2 (Closure of a set) *Let (A,d) be a metric space and $B \subset A$. Then*

$$\overline{B} := \{b \in A : \exists b_k \in B \text{ such that } b_k \xrightarrow[k\to\infty]{} b\}$$

[1] Indeed, $B := [0,1]$ and $A\backslash B = [2,3]$ are both closed.

is a closed subset of A.

Proof Let

$$C := \{c \in A : \exists \varepsilon > 0 \text{ such that } B_\varepsilon(c) \cap B = \emptyset\}.$$

We claim that C is an open set. To see this, let $c \in C$. Then there exists an $\varepsilon > 0$ such that $B_\varepsilon(c) \cap B = \emptyset$. Now by the triangle inequality, there cannot exist points $a \in A$ and $b \in B$ such that $d(c, a) < \varepsilon/3$ and $d(a, b) < \varepsilon/3$. As a result, for each $a \in B_{\varepsilon/3}(c)$ we must have $B_{\varepsilon/3}(a) \cap B = \emptyset$ and hence $a \in C$. This shows that $B_{\varepsilon/3}(c) \subset C$. Since c is arbitrary, we conclude that C is open.

To prove the lemma, it now suffices to prove that $\overline{B} = A\backslash C$. If $c \in C$, then there exists an $\varepsilon > 0$ such that $B_\varepsilon(c) \cap B = \emptyset$ i.e., $d(b, c) \geq \varepsilon$ for all $b \in B$. This clearly implies that there cannot exist $b_k \in B$ such that $b_k \to c$. On the other hand, if $a \notin C$, then $B_{1/k}(a) \cap B \neq \emptyset$ for all $k \geq 1$ so we can choose $b_k \in B$ such that $d(b_k, a) \to 0$, proving that $a \in \overline{B}$. ∎

In mathematics, a result that follows immediately from another result, or that follows immediately a side-result of a proof, is called a "corollary". In particular, our previous proof yields the following corollary. The set $\mathring{B}$ is called the *interior* of B.

Corollary 3.3 (Interior of a set) *Let (A, d) be a metric space and $D \subset A$. Then*

$$\mathring{D} := \{a \in A : \exists \varepsilon > 0 \textit{ such that } B_\varepsilon(a) \subset D\}.$$

is an open subset of A.

Proof Let $B = A\backslash D$. Then the set C defined in our proof of Lemma 3.2 is the same as $\mathring{D}$. We already proved that C is open, so we are done. ∎

For any set B, we call $\partial B := \overline{B}\backslash\mathring{B}$ the *boundary* of B. Then a set B is closed if and only if it includes its boundary, i.e., if $\partial B \subset B$, and open if and only if it does not include its boundary, i.e., if $\partial B \cap B = \emptyset$. For example, for a ball in $\mathbb{R}^n$ of radius r around $\vec{0}$, the boundary is the surface $\{\vec{x} \in \mathbb{R}^n : \|\vec{x}\| = r\}$. When we think about subsets of $\mathbb{R}^n$, we often imagine "nice" sets such as balls or retangles that have the properties:

$$\overline{\mathring{B}} = \overline{B} \quad \text{and} \quad \mathring{\overline{B}} = \mathring{B}.$$

There are, however, many sets that do not have these nice properties. For example, the interior of $\mathbb{Q}$, viewed as a subset of $\mathbb{R}$, is empty, while $\overline{\mathbb{Q}} = \mathbb{R}$. In general, if (A, d) is a metric space and $D \subset A$, then we say that D is *dense* in A if $\overline{D} = A$.

3.3 Completeness

In the previous section, we have defined the closure $\overline{B}$ of a subset $B \subset A$ of a metric space (A, d). For example, if we view the set $\mathbb{Q}$ of rational numbers as a subset of the space $(\mathbb{R}, d)$ with $d(x, y) := |x - y|$, then $\overline{\mathbb{Q}} = \mathbb{R}$. This suggests that it should be possible to define the real numbers as some sort of "closure" of the rational numbers, by adding the limits of all sequences that "should" converge, but whose limit is not in $\mathbb{Q}$.

Let (A, d) be a metric space. By definition, a *Cauchy sequence* is a sequence $(a_k)_{k\geq 1}$ of elements of A such that:

$$\forall \varepsilon > 0 \ \exists m \in \mathbb{N}_+ \text{ such that } d(a_k, a_n) \leq \varepsilon \ \forall k \geq m \text{ and } n \geq m.$$

Note that this looks a lot like the definition of convergence to a limit, but the limit point a is never mentioned. Instead, from the point m onwards, all elements of the sequence stay close to each other. It is easy to see that:

- Every convergent sequence is a Cauchy sequence, i.e., if $(a_k)_{k\geq 1}$ satisfies $a_k \to a$ for some $a \in A$, then $(a_k)_{k\geq 1}$ is a Cauchy sequence.

By definition, a metric space (A, d) is *complete* if the converse conclusion can be drawn, i.e., if every Cauchy sequence has a limit. More formally, (A, d) is complete if:

$$\forall \text{ Cauchy sequence } (a_k)_{k\geq 1} \ \exists a \in A \text{ such that } a_k \underset{k\to\infty}{\longrightarrow} a.$$

We state the following fact without proof:

- The space $\mathbb{R}$, equipped with the metric $d(x, y) := |x - y|$, is complete.

As we will see in Section 3.4 below, the completeness of $\mathbb{R}$ is more of less a direct conequence of the formal mathematical definition of $\mathbb{R}$, which we have not given yet.

The following theorem says that for any metric space A, we can take some sort of "closure" of A even if a priori we do not view A as a subset of some larger space.

Theorem 3.4 (Completion of a metric space) *Let (A, d) be a metric space. Then there exists a metric space $(\overline{A}, \overline{d})$ with the following properties:*

(i) *$(\overline{A}, \overline{d})$ is complete.*

(ii) *A is a dense subset of $\overline{A}$.*

(iii) *$d(a, b) = \overline{d}(a, b)$ for all $a \in A$ and $b \in A$.*

Proof (sketch) Let C be the set of all Cauchy sequences $(a_k)_{k\geq 1}$ in A. We call two Cauchy sequences $(a_k)_{k\geq 1}$ and $(a'_k)_{k\geq 1}$ equivalent, which we denote as $(a_k)_{k\geq 1} \sim (a'_k)_{k\geq 1}$, if

$$\forall \varepsilon > 0\ \exists m \text{ such that } d(a_k, a'_k) \leq \varepsilon\ \forall k \geq m.$$

We let $\overline{A}$ denote the set of equivalence classes for this equivalence relation. If $(a_k)_{k\geq 1}$ and $(b_k)_{k\geq 1}$ are Cauchy sequences, then one can check that $\big(d(a_k, b_k)\big)_{k\geq 1}$ is a Cauchy sequence in $\mathbb{R}$, and hence by the completeness of $\mathbb{R}$, the limit

$$\hat{d}\big((a_k)_{k\geq 1}, (b_k)_{k\geq 1}\big) := \lim_{k\to\infty} d(a_k, b_k)$$

exists. Moreover, one can check that

$$\begin{aligned} &(a_k)_{k\geq 1} \sim (a'_k)_{k\geq 1} \text{ and } (b_k)_{k\geq 1} \sim (b'_k)_{k\geq 1} \\ &\quad \text{imply} \quad \hat{d}\big((a_k)_{k\geq 1}, (b_k)_{k\geq 1}\big) = \hat{d}\big((a'_k)_{k\geq 1}, (b'_k)_{k\geq 1}\big). \end{aligned}$$

As a result, $\hat{d}\big((a_k)_{k\geq 1}, (b_k)_{k\geq 1}\big)$ depends only on the equivalence classes of $(a_k)_{k\geq 1}$ and $(b_k)_{k\geq 1}$, so there exists a function $\overline{d} : \overline{A} \times \overline{A} \to \mathbb{R}$ such that

$$\hat{d}\big((a_k)_{k\geq 1}, (b_k)_{k\geq 1}\big) = \overline{d}\big([(a_k)_{k\geq 1}], [(b_k)_{k\geq 1}]\big),$$

where $[(a_k)_{k\geq 1}]$ denotes the equivalence class containing $(a_k)_{k\geq 1}$. Now one can prove that $\overline{d}$ is a metric on $\overline{A}$, i.e., it satisfies properties (i)–(iii) of the definition of a metric.

If $(a_k)_{k\geq 1}$ and $(a'_k)_{k\geq 1}$ are equivalent Cauchy sequences, then it is not hard to prove that there are two possibilities. Either $(a_k)_{k\geq 1}$ and $(a'_k)_{k\geq 1}$ both converge to the same limit $a \in A$, of $(a_k)_{k\geq 1}$ and $(a'_k)_{k\geq 1}$ both do not have a limit in A. We use this to view A as a subset of $\overline{A}$. More precisely, we identify $a \in A$ with the equivalence class of sequences that converge[2] to a. If we view A as a subset of $\overline{A}$ in this way, then one can prove that $d(a, b) = \overline{d}(a, b)$ and moreover A is dense in $\overline{A}$. Finally, one can prove that the metric space $(\overline{A}, \overline{d})$ is complete. ∎

The proof of Theorem 3.4 is very formal, and, if one wants to fill in all the details, also quite long. Things are much easier if we already know that A is a subset of another metric space, which is complete. If (A, d) is a metric space and $B \subset A$ is a subset of A, then we can view B as a metric space equipped with the metric $d' : B \times B \to \mathbb{R}$ defined as $d'(a, b) := d(a, b)$, i.e., d' is the restriction of d to B. We say that B is *complete* (as a subset of A) if (B, d') is complete.

[2]Recall that every convergent sequence is automatically a Cauchy sequence.

Lemma 3.5 (Complete subsets) *Let (A, d) be a complete metric space. Then a subset $B \subset A$ is complete if and only if it is closed.*

Proof Assume that B is complete. If $b_k \in B$ satisfy $b_k \to a$ for some $a \in A$, then $(b_k)_{k\geq 1}$ must be a Cauchy sequence, and hence by the completeness of B must converge to a limit in B. It follows that $a \in B$. Since this holds for every sequence in B, the set B is closed.

Conversely, if B is closed and $(b_k)_{k\geq 1}$ is a Cauchy sequence in B, then by the completeness of A we must have $b_k \to a$ for some $a \in A$. Since B is closed, this implies $a \in B$. Thus, every Cauchy sequence in B has a limit in B proving that B is complete. ∎

In particular, if (A, d) be a complete metric space and $B \subset A$ is any subset of A, then $\overline{B}$ (equipped with the metric d restricted to $\overline{B}$) is a complete metric space that contains B as dense subset. Since moreover the metric on $\overline{B}$ agrees with the metric on B, this means that $\overline{B}$ has properties (i)–(iii) of the completion of B as stated in Theorem 3.4, so we can view $\overline{B}$ as the completion of B.

3.4 Construction of the real numbers

In Section 1.1, we defined real numbers as numbers that can be written in decimal notation, and in Section 1.2 we claimed without proof that addition and multiplication of such numbers can be defined in such a way that properties (i)–(ix) hold. In Section 1.3 we introduced rational numbers and showed that there are real numbers that are not rational.

In Chapter 2 and more specifically in Section 2.7 we indicated, without going into details, how one can use the axioms of set theory to construct the rational numbers and to prove that addition and multiplication of rational numbers can be defined in such a way that properties (i)–(ix) of Section 1.2 hold. In the present section, we will extend this construction to the real numbers. Finally, in Section 3.6 below, we will the prove that each real number has a unique decimal notation. This then (finally!) confirms that indeed, the real numbers as defined in Section 1.1 can be added and multiplied in the way we claimed in Section 1.2.

As already suggested in Section 3.3, we want to construct $\mathbb{R}$ in such a way that it is the completion of $\mathbb{Q}$ with respect to the metric $d(x, y) := |x - y|$. We can, however, not simply apply Theorem 3.4 to define $\mathbb{R}$. Indeed, the proof of Theorem 3.4 already uses the fact that $\mathbb{R}$ is complete, so such an argument would be circular. Nevertheless, by being a bit more careful, we can use the same idea as in the proof of Theorem 3.4.

Let $\mathbb{Q}_+ := \{x \in \mathbb{Q} : x > 0\}$. We start by noting that $|x - y| \in \mathbb{Q}_+$ for all $x, y \in \mathbb{Q}$ and that a sequence $(x_k)_{k\geq 1}$ of rational numbers is a Cauchy sequence if and ony if

$$\forall \varepsilon \in \mathbb{Q}_+ \ \exists m \in \mathbb{N}_+ \text{ such that } |x_k - x_n| \leq \varepsilon \ \forall k \geq m \text{ and } n \geq m.$$

Thus, we do not need the real numbers to say what a Cauchy sequence of rational numbers is. Similar to what we did in the proof of Theorem 3.4, we call two such Cauchy sequences $(x_k)_{k\geq 1}$ and $(x'_k)_{k\geq 1}$ equivalent if and only if

$$\forall \varepsilon \in \mathbb{Q}_+ \ \exists m \text{ such that } |x_k - x'_k| \leq \varepsilon \ \forall k \geq m,$$

and we define $\mathbb{R}$ to be the corresponding set of equivalence classes. Now one can show that

$$\begin{aligned} &(x_k)_{k\geq 1} \sim (x'_k)_{k\geq 1} \text{ and } (y_k)_{k\geq 1} \sim (y'_k)_{k\geq 1} \\ &\quad \text{imply} \quad (x_k + y_k)_{k\geq 1} \sim (x'_k + y'_k)_{k\geq 1} \text{ and } (x_k y_k)_{k\geq 1} \sim (x'_k y'_k)_{k\geq 1}, \end{aligned}$$

which allows us to unambiguously define the sum and product of two equivalence classes. Next, one can prove that the sum and product on $\mathbb{R}$, defined in this way, satisfy properties (i)–(ix) from Section 1.2.

We note that if two sequences $(x_k)_{k\geq 1}$ and $(y_k)_{k\geq 1}$ satisfy

$$\exists m \in \mathbb{N}_+ \text{ such that } x_k < y_k \ \forall k \geq m, \tag{3.2}$$

and $(x_k)_{k\geq 1} \sim (x'_k)_{k\geq 1}$ and $(y_k)_{k\geq 1} \sim (y'_k)_{k\geq 1}$, then (3.2) also holds with x_k and y_k replaced by x'_k and y'_k. We can use this to unambiguously define the relation $<$ on $\mathbb{R}$. We can use this to define $|x|$ and prove that $\mathbb{R}$, equipped with the metric $d(x, y) := |x - y|$, is a complete metric space. This then implies that $\mathbb{R}$ is the completion of $\mathbb{Q}$, as we originally wanted. Filling in all the details is quite a long job, but not very difficult, so we skip the boring details.

3.5 Calculating with infinity

Let us define a function $\psi : \mathbb{R} \to \mathbb{R}$ by

$$\psi(x) := \frac{x}{1 + |x|}.$$

This function is continuous, satisfies $\psi(x) < \psi(y)$ for all $x < y$, as well as $\lim_{k\to\infty} \psi(-k) = 1$ and $\lim_{k\to\infty} \psi(k) = 1$.

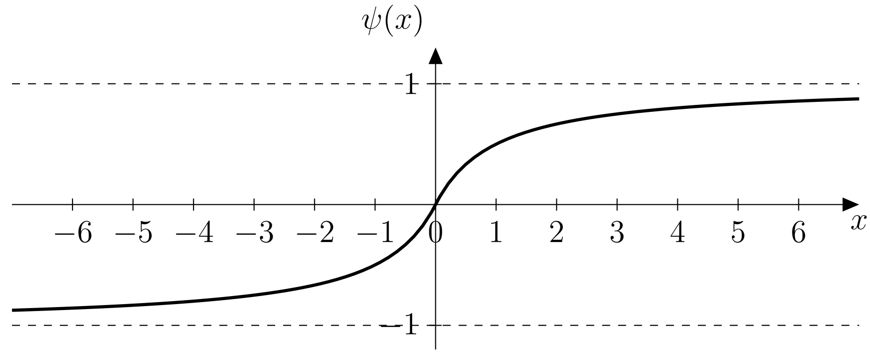

In Section 1.1, we drew the real numbers on an infinite straight line. The function ψ is a bijection from $\mathbb{R}$ to the interval $(-1, 1)$. By drawing a real number x at the position $\psi(x)$ instead of x, we get a picture of the real line that looks like this:

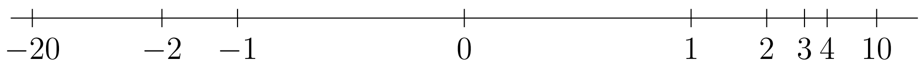

It is now very natural to give names to the endpoints of this line segment. We call the left endpoint $-\infty$ "minus infinity" and the right endpoint ∞ "infinity", also called "plus infinity". The picture then looks like this:

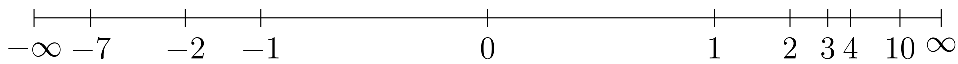

We call the set $\overline{\mathbb{R}} := \mathbb{R} \cup \{-\infty, \infty\}$ the *extended real numbers*. Another way of looking at $\overline{\mathbb{R}}$ is as follows. Let us define $\psi(-\infty) := -1$ and $\psi(\infty) := 1$ and let us define a function $\overline{d} : \overline{\mathbb{R}} \times \overline{\mathbb{R}} \to \mathbb{R}$ by

$$\overline{d}(x, y) := \big|\psi(x) - \psi(y)\big|,$$

i.e., $\overline{d}(x, y)$ is the distance between x and y in our last two pictures, where we have used the map ψ to transform the infinite real line into a finite line segment. Then it is easy to see that d is a metric on $\overline{\mathbb{R}}$ (i.e., satisfies the properties (i)–(iii) of the definition of a metric). Moreover, if $(x_k)_{k\geq 1}$ is a sequence of real numbers and $x \in \mathbb{R}$, then

$$|x_k - x| \underset{k\to\infty}{\longrightarrow} 0 \text{ if and only if } \overline{d}(x_k, x) \underset{k\to\infty}{\longrightarrow} 0,$$

so a sequence $(x_k)_{k\geq 1}$ of real numbers converges to a limit $x \in \mathbb{R}$ with respect to the usual distance if and only if it converges with respect to the new distance function $\overline{d}$. However, certain sequences that did not have a limit before now have a limit. For example,

$$\overline{d}(k^2, \infty) \underset{k\to\infty}{\longrightarrow} 0 \quad \text{and} \quad \overline{d}(-k, -\infty) \underset{k\to\infty}{\longrightarrow} 0$$

which we can write down differently as

$$\lim_{k\to\infty} k^2 = \infty \quad \text{and} \quad \lim_{k\to\infty} (-k) = -\infty.$$

A slightly different way of looking at this is as follows. Let $d : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ denote the restriction of $\overline{d}$ to $\mathbb{R} \times \mathbb{R}$. Then $(\mathbb{R}, d)$ is a metric space and $(\overline{\mathbb{R}}, \overline{d})$ is the completion of $(\mathbb{R}, d)$.

For intervals in $\overline{\mathbb{R}}$ the usual notation applies. In particular $[-\infty, \infty] = \overline{\mathbb{R}}$ and $(-\infty, \infty) = \mathbb{R}$. By definition, $-\infty < x < \infty$ for all $x \in \mathbb{R}$. In many ways, it is possible to calculate with ∞ as if it were a real number. The usual conventions are that

$$x + \infty := \infty \ \forall x \in \overline{\mathbb{R}}, \quad x \cdot \infty := \infty \ \forall x > 0, \quad \text{and} \quad 0 \cdot \infty := 0.$$

Then the set of extended nonnegative numbers $[0, \infty]$ satisfies the rules (i)–(vii) from Section 1.2 without an exception. However, we cannot extend $[0, \infty]$ in a way so that also rules (viii) and (ix) hold. We could try to satisfy (viii) by defining $\infty - \infty := 0$, but this would lead to problems. For example, this would violate rule (ii) since such a definition would have the result that

$$3 + (\infty - \infty) = 3 + 0 = 3 \neq 0 = \infty - \infty = (3 + \infty) - \infty.$$

In view of this, the usual convention is that $\infty - \infty$ is *undefined*. As soon as this occurs somewhere in a formula, the result is undefined. However, as long as we never have to subtract infinity from infinity, rules (i)–(vii) remain valid.

It is also not possible to satisfy rule (ix) for $a = \infty$. One sometimes defines $1/\infty := 0$, but then using our earlier rules we have $\infty \cdot (1/\infty) = \infty \cdot 0 = 0 \neq 1$, in violation of rule (ix).

The extended real numbers have some pleasant properties that the real numbers do not have. We say that a sequence $(x_k)_{k\geq 0}$ is *increasing* if $x_k \leq x_{k+1}$ for all $k \geq 1$. If $x_k < x_{k+1}$ for all $k \geq 1$, then we say that the sequence is *strictly increasing*. Similarly, we say that a sequence is *decreasing* (respectively, *strictly decreasing*) if $x_k \geq x_{k+1}$ (respectively, $x_k > x_{k+1}$) for all $k \geq 1$. Instead of *increasing* one sometimes says *nondecreasing* and instead of *decreasing* one sometimes says *nonincreasing*, to stress the fact that one does not mean strictly increasing/decreasing. Some pleasant properties of the extended real numbers are:

- Each increasing sequence has a limit in $\overline{\mathbb{R}}$.
- Each decreasing sequence has a limit in $\overline{\mathbb{R}}$.
- For each set $A \subset B$, there exists a number $y \in \overline{\mathbb{R}}$ such that $\{x \in \overline{\mathbb{R}} : x \geq a \text{ for all } a \in A\} = [y, \infty]$. This number y is called the *supremum* of A and denoted as $\sup A := y$.
- For each set $A \subset B$, there exists a number $y \in \overline{\mathbb{R}}$ such that $\{x \in \overline{\mathbb{R}} : x \leq a \text{ for all } a \in A\} = [-\infty, y]$. This number y is called the *infimum* of A and denoted as $\inf A := y$.

If $\sup A \in A$, then we call $\sup A$ the *maximum* of A. Similarly, if $\inf A \in A$, then we call $\inf A$ the *minimum* of A. Note, however, that in general a set need not contain its supremum or infimum. For example, $\sup(0,1) = 1$ and $\inf \mathbb{R} = -\infty$. We use the word "maximum" or "maximal element" only if $\sup A$ is an element of A. The same applies to the words "minimum" and "minimal element". For finite sets, there is of course no difference between the two concepts. We use the notation

$$x \vee y := \max\{x, y\} = \sup\{x, y\} \quad \text{and} \quad x \wedge y := \min\{x, y\} = \inf\{x, y\}.$$

3.6 Infinite sums

If $(x_k)_{k\geq 1}$ is a sequence of real numbers, then we let

$$\sum_{k=1}^{n} x_k := x_1 + x_2 + \cdots + x_n$$

denote the sum of its first n terms. Sometimes it is more natural to consider sequences indexed by $\mathbb{N}$ instead of $\mathbb{N}_+$ and to start the sum with x_0. In that case, we write $\sum_{k=0}^{n}$. In general,

$$\sum_{k=m}^{n} x_k := x_m + x_{m+1} + \cdots + x_n.$$

By definition, we write

$$\sum_{k=m}^{\infty} x_k := \lim_{n\to\infty} \sum_{k=m}^{n} x_k$$

if the limit exists as an extended real number (including $-\infty$ and ∞); otherwise, the infinite sum is not defined.

Lemma 3.6 (Finite geometric sum) *For any real number $a \neq 1$, one has*

$$\sum_{k=0}^{n-1} a^n = \frac{1-a^n}{1-a}.$$

Proof Recall that in Section 1.2 we defined $a^0 := 1$. Then

$$(1-a)\sum_{k=0}^{n-1} a^n = \sum_{k=0}^{n-1} a^n - \sum_{k=0}^{n-1} a \cdot a^n = \sum_{k=0}^{n-1} a^n - \sum_{k=1}^{n} a^n = 1 - a^n.$$

Using the assumption that $a \neq 1$, we can divide both sides of this equation by $1-a$ to arrive at the formula in the lemma. ∎

Remark A different way to write down the proof of Lemma 3.6, that does not use the new notation for sums, is as follows:

$$(1-a)\big(1+a^1+\cdots+a^{n-1}\big) = \big(1+a^1+\cdots+a^{n-1}\big) - \big(a^1+a^2+\cdots+a^n\big) = 1-a^n.$$

Note that when we write $(1-a)\sum_{k=0}^{n-1} a^n$, we mean $(1-a)\big(\sum_{k=0}^{n-1} a^n\big)$, even when we do not write the brackets.

Proposition 3.7 (Infinite geometric sum) *If a is a real number such that $|a| < 1$, then*

$$\sum_{k=0}^{\infty} a^k = \frac{1}{1-a}.$$

If $a \geq 1$, then $\sum_{k=0}^{\infty} a^k = \infty$, while for $a \leq -1$, the infinite sum is not defined.

Proof If $|a| < 1$, then $a^n \to 0$ as $n \to \infty$. Since the function $f(x) := (1-x)/(1-a)$ is continuous, it follows that

$$\lim_{n\to\infty} \frac{1-a^n}{1-a} = \frac{1}{1-a},$$

so the claim follows from Lemma 3.6. If $a = 1$, then $\sum_{k=0}^{n-1} a^k = n$, which tends to infinity. If $a > 1$, then $\sum_{k=0}^{n-1} a^k = (a^n - 1)/(a-1)$ which tends to infinity since a^n tends to infinity. If $a = -1$, then $\sum_{k=0}^{n-1} a^k = 1$ for odd n and $= 0$ for even n, so the limit as $n \to \infty$ does not exist. If $a < -1$, then $\lim_{n\to\infty}(1-a^{2n})/(1-a) = -\infty$ while $\lim_{n\to\infty}(1-a^{2n+1})/(1-a) = +\infty$, so $(a^n-1)/(a-1)$ does not converge as $n \to \infty$. ∎

Remark In particular,

$$1 + \frac{1}{n} + \frac{1}{n^2} + \frac{1}{n^3} + \cdots = \frac{1}{1-\frac{1}{n}} = \frac{n}{n-1} = 1 + \frac{1}{n-1}.$$

So, for example,

$$\begin{aligned}
&\frac{1}{2} + \frac{1}{2^2} + \frac{1}{2^3} + \cdots = \frac{1}{1} = 1, \\
&\frac{1}{3} + \frac{1}{3^2} + \frac{1}{3^3} + \cdots = \frac{1}{2}, \\
&\frac{1}{4} + \frac{1}{4^2} + \frac{1}{4^3} + \cdots = \frac{1}{3},
\end{aligned}$$

etcetera.

In many ways, it is possible to calculate with infinite sums as if they were finite sums. There is one point one has to be careful with. In Section 3.5, we have seen that $\infty - \infty$ is ill-defined. As a consequence, strange things can happen when the sum of all positive terms is $+\infty$ while the sum of all negative terms is $-\infty$. On the other hand, if all terms are nonnegative, then the following theorem shows that infinite sums are always well-defined (though the outcome can be $+\infty$) and that the outcome does not depend on the summation order.

Theorem 3.8 (Nonnegative sums) *Let I be a countable set and let $(x_k)_{k\in I}$ be numbers with $x_k \in [0,\infty]$. Then there exists a number $S \in [0,\infty]$ such that*

$$S = \lim_{n\to\infty} \sum_{k=1}^{n} x_{\phi(k)}$$

for every bijection $\phi : \mathbb{N}_+ \to I$. In particular, the outcome does not depend on the choice of the bijection ϕ.

Proof Since each increasing sequence has a limit in $\overline{\mathbb{R}}$, it is clear that the limit exists. We need to show that outcome does not depend on the choice of the bijection. Let $\phi : \mathbb{N}_+ \to S$ and $\psi : \mathbb{N}_+ \to S$ be bijections and let

$$S = \lim_{n\to\infty} \sum_{k=1}^{n} x_{\phi(k)} \quad \text{and} \quad T = \lim_{n\to\infty} \sum_{k=1}^{n} x_{\psi(k)}.$$

We observe that

$$\forall n \in \mathbb{N}_+ \ \exists m \in \mathbb{N}_+ \text{ such that } \{\psi(1),\ldots,\psi(m)\} \supset \{\phi(1),\ldots,\phi(n)\}.$$

As a consequence

$$\forall n \in \mathbb{N}_+ \ \exists m \in \mathbb{N}_+ \text{ such that } \sum_{k=1}^{m} x_{\psi(k)} \geq \sum_{k=1}^{n} x_{\phi(k)}.$$

Since $\sum_{k=1}^{m'} x_{\psi(k)} \geq \sum_{k=1}^{m} x_{\psi(k)}$ for $m' \geq m$, it follows that

$$T \geq \sum_{k=1}^{n} x_{\phi(k)} \qquad \forall n \in \mathbb{N}_+,$$

and hence $T \geq S$. Since ϕ and ψ play symmetric roles, the same argument shows that $S \geq T$. ∎

We denote the number $S \in [0, \infty]$ from Theorem 3.8 by $\sum_{k \in I} x_k$. More generally, we define

$$\sum_{k \in I} x_k := \sum_{x \in I_+} x_k - \sum_{x \in I_-} (-x_k)$$
$$\text{with} \quad I_+ := \{k \in I : x_k > 0\} \text{ and } I_- := \{k \in I : x_k < 0\},$$

provided this is not $\infty - \infty$, which is not defined. If the sum of the positive and negative terms yields $\infty - \infty$, then it will still be true that the limit

$$\lim_{n \to \infty} \sum_{k=1}^{n} x_{\phi(k)}$$

exists for *some* bijection $\phi : \mathbb{N}_+ \to I$, but the limit will depend on the choice of ϕ and for some bijections the limit does not exist at all. Here are some additional rules that hold for infinite sums of nonnegative numbers. In the first rule, $(x_{(n,m)})_{(n,m) \in \mathbb{N}_+ \times \mathbb{N}_+}$ is any collection of nonnegative extended numbers indexed by $\mathbb{N}_+ \times \mathbb{N}_+$, which as we have seen in Section 2.5 is countable.

- $\displaystyle\sum_{n=1}^{\infty} \sum_{m=1}^{\infty} x_{(n,m)} = \sum_{(n,m) \in \mathbb{N}_+ \times \mathbb{N}_+} x_{(n,m)}.$
- $\displaystyle\Big(\sum_{n=1}^{\infty} x_n\Big)\Big(\sum_{m=1}^{\infty} y_m\Big) = \sum_{(n,m) \in \mathbb{N}_+ \times \mathbb{N}_+} x_n y_m.$
- $\displaystyle\sum_{k=0}^{\infty} x_k + \sum_{k=0}^{\infty} y_k = \sum_{k=0}^{\infty} (x_k + y_k).$

A sequence $(x_k)_{k \in \mathbb{N}_+}$ is *absolutely summable* if

$$\sum_{k=1}^{\infty} |x_k| < \infty.$$

By our previous remarks, if $(x_k)_{k \in \mathbb{N}_+}$ is absolutely summable, then the infinite sum $\sum_{k=1}^{\infty} x_k$ is well-defined, the outcome is finite, and does not depend on the summation order.

Recall that at the end of Section 1.1, we defined

$$a^{-n} := \frac{1}{a^n} \qquad \forall a \neq 0,\ n \in \mathbb{Z}.$$

Let $n \geq 2$ be an integer. Then we claim that every nonnegative real number a can be written in the form

$$a = \sum_{k=m}^{\infty} x_k n^k \quad \text{with} \quad m \in \mathbb{Z} \text{ and } x_k \in \{0, \ldots, n-1\}\ \forall k.$$

For example, for $n = 10$, we have

$$\begin{aligned} \pi &= 3.14159265\ldots \\ &= 3 \cdot 10^0 + 1 \cdot 10^{-1} + 4 \cdot 10^{-2} + 1 \cdot 10^{-3} + 5 \cdot 10^{-4} + 9 \cdot 10^{-5} + \cdots, \end{aligned}$$

i.e., this is just the usual decimal expansion of a real number, that we started our informal discussion with in Section 1.1.