Introduction to Pure Mathematics

Sets, Functions and Relations

This introductory course is intended for first-year students, to provide a good grounding in the essential language of pure mathematics, establish the notation used and point out common misconceptions. Course information and further study skills documentation are available from the Faculty website.

Dr Gareth Wilkes

Revised August 29, 2024

Introduction

One of the biggest jumps from school mathematics to university mathematics is in the level of abstraction required, particularly in 'pure' mathematics. Part of the difficulty of this transition is that the language and style of that mathematics change at the same time as you are trying to learn new material, like group theory or number theory.

This short introductory course is intended to assist with that transition. We will discuss the fundamental terminology in terms of which all other mathematics is phrased – sets, functions and equivalence relations. It is, in a sense, a course about 'concepts shorn of content'.

Sets and functions are the foundations of all branches of mathematics. Once you are familiar with the basic manipulations and language of these objects, you can go on to explore more interesting mathematical concepts. That is, once you have the 'words' of the language of mathematics, you can begin to combine them into ever more interesting and meaningful mathematical 'sentences'.

Much of the material covered in this course will also be covered at some point during your 'Numbers and Sets' course. However I, in collaboration with other experienced supervisors and Directors of Studies in the University, feel that it will be beneficial to meet these concepts in detail in a dedicated course at the beginning of your university education. You will then be better prepared for all your courses, and will have a better understanding when these concepts arise in other courses like Numbers and Sets.

One final comment: on the use of mathematical symbols as abbreviations for logical terms. The symbols '$\forall$' and '$\exists$' for 'for all' and 'there exists' often cause more confusion for new students than they are worth. Once you are more used to writing formal mathematical sentences, you can begin to abbreviate them using symbols. However for this introductory course I have chosen to avoid using these abbreviations. This makes the mathematics a little wordy at times, and you could consider which points the use of these symbols might make things look cleaner. I have written more fully about the use and interpretation of the symbols $\forall$ and $\exists$ elsewhere[1].

As for other logical symbols, '$\Leftrightarrow$' and '$\Rightarrow$' are rather less troublesome – provided they are used correctly – but I have still written them in words as 'if and only if' and 'implies' in this guide.

The formal logic symbols '$\wedge$', '$\vee$' and '$\neg$' for 'and', 'or' and 'not' seldom do anything but make your mathematics more difficult to read – and remember, mathematics is intended to be read by humans, not computers. I never use

[1] 'A Brief Guide to Mathematical Writing', available here.

$\wedge$, $\vee$ or $\neg$ in my own writing, however advanced the level.

How to use these notes

These notes, like all lecture notes, are intended to accompany the lectures rather than replace them. There is no expectation that you will have read these notes before the lectures. Some students find it beneficial to annotate a copy of the notes during the lectures rather than making their own notes; but for most, these typed notes will be most used after the lectures, as a means to supplement their own notes or as a reference for the future.

Throughout these notes there are various exercises. They are placed at their logical positions amid the material, but are also gathered onto a separate worksheet for easier reference. No solutions will be provided; if you find an exercise challenging, discuss it with another mathematician – either your fellow students or a supervisor. This is a much better way to understand something confusing than being given the solution on a piece of paper.

In common with many things at university-level study, the amount of effort you put into these exercises is up to you. Ignore them, and you cannot benefit from them; attempt them and you may learn something. This is, after all, your degree and no one else's.

> Scattered throughout this document are comments and exercises in boxes like this one. These are intended to be more advanced, subtle or obscure points of mathematics. Read them if you are comfortable with the main material and are interested; otherwise they may be safely ignored.

Contents

1 Sets

What is a set? We could try to say something like ‘a set is a collection of objects’ as a definition, but that only begs the question of ‘what is a collection?’. In fact, mathematics begins with two things that cannot be defined – because there are no concepts more basic to define them in terms of. We have an idea called a ‘set’, and a concept called ‘being an element of a set’. The sentence ‘x is an element of A’ is written $x \in A$. Everything else in mathematics can be defined beginning from these two concepts.

Specifying a set

The simplest way to write a set is to list all its elements, between a pair of braces. So $A = \{1, 2, 3\}$ is the set whose elements are 1, 2 and 3 (and nothing else). The set $\{\}$ with no elements at all is called the *empty set*, written $\emptyset$.

Of course, writing a set so explicitly does seem to require that you know all its elements in advance. More commonly you may see a set in which each element is given a ‘label’. For example, if you knew that a set A had four elements, you might very naturally label them with the natural numbers from 1 to 4 and write

$$A = \{a_1, a_2, a_3, a_4\}.$$

If I had said ‘117’ instead of ‘4’, obviously I wouldn’t want to write out 117 symbols. We would instead describe the elements of our set as ‘the elements a_i where i ranges from 1 to 117’. This kind of expression is what may be called an *indexed set*. If I is the set of labels used (the ‘indexing set’) then we may write

$$A = \{a_i \mid i \in I\}$$

to say that we have labelled the elements of A as a_i, where we have one element of A labelled with each element i of I. It is not guaranteed that these elements are all different; perhaps we do not yet know enough about A to list its elements with no repetitions, so we may have given some elements several labels.

Very commonly, as above, I is the set of natural numbers, or a subset of it. For the sake of laziness we often use an ellipsis ... to specify a range of integers. Our first example above may be written

$$A = \{a_1, a_2, a_3, a_4\} \text{ or } \{a_1, \ldots, a_4\} \text{ or } \{a_i \mid i \in \{1, \ldots, 4\}\}.$$

One might also see[2]

$$A = \{a_i \mid 1 \leq i \leq 4\}.$$

Similarly, for a whole sequence of elements labelled by the natural numbers[3], we may write

$$A = \{a_0, a_1, a_2, \ldots\} \text{ instead of } A = \{a_i \mid i \in \mathbb{N}\}.$$

However there is no reason for I to be limited to a collection of natural numbers. Perhaps A has its elements indexed by the real numbers; for example, the unit circle in the complex plane may be written

$$S = \{e^{i\theta} \mid \theta \in \mathbb{R}\}.$$

In this case, every element has received infinitely many labels: for instance, the element $1 = e^{i0} = e^{2\pi i}$ is labelled by $\theta = 0$ and $\theta = 2\pi$, among others.

Indexed sets are generally used as a way of labelling the elements of a pre-existing set. Very often, of course, we will want to construct new sets from old ones. Perhaps the most common way to do this is to define a set using a *property*. Such a construction starts with some 'known' set and builds a new set by taking only those elements which satisfy a certain property. The general format is $\{x \in X \mid x \text{ is} \ldots\}$.

For example, here are three ways of specifying the set of even natural numbers.

$$\begin{aligned} E &= \{0, 2, 4, 6, \ldots\} && \text{(explicit listing)} \\ &= \{a_i \mid i \in \mathbb{N}\}, \text{ where } a_i = 2i && \text{(indexed set)} \\ &= \{n \in \mathbb{N} \mid 2 \text{ is a factor of } n\} && \text{(property).} \end{aligned}$$

Notation. The same symbol $\mid$ is used for both an indexed set and a set defined by a property. This symbol can be read 'where' or 'such that' depending on context (and what sounds nicest). Many mathematicians use a colon : instead of a pipe |, or use both symbols interchangeably, or try to reserve one for indexed sets and one for sets defined by a property. I prefer the pipe for my part, but I may be in the minority.

[2] This is, strictly speaking, ambiguous; is the indexing set here just $\{1, 2, 3, 4\}$ or does it include every real number between 1 and 4?

[3] Whether the set $\mathbb{N}$ includes the number 0 depends a bit on to whom you are talking and what branch of mathematics you are studying. For pure set theory, 0 is generally included in $\mathbb{N}$. In contexts where sequences are being studied, like Analysis, we usually start $\mathbb{N}$ with the element 1 so that a sequence reads $(a_1, a_2, a_3, \ldots)$.

There is a good reason for insisting that sets defined by a property should be subsets of an already existing set – to avoid paradoxes. One famous example is Russell's Paradox. Suppose you try to define[a] the 'set'

$$A = \{x \mid x \notin x\},$$

the 'set of all sets that are not elements of themselves'. The fatal question is whether A is an element of itself. Each of $A \in A$ and $A \notin A$ implies the other, an impossibility which means that the set A cannot exist. This also means we cannot have a 'set of all sets', as such a thing would contain A as a subset.

[a]The slash through the membership symbol turns it into a negative, $\notin$ 'is not an element of'.

Equality of sets

Throughout the above discussion we have been using the symbol $=$ to say that two sets are equal. It is worth pausing to dwell on what exactly equality of sets means.

Definition. Two sets A and B are *equal*, written $A = B$, if they have the same elements; that is,

$$x \in A \text{ if and only if } x \in B.$$

The set A is a *subset* of B, written $A \subseteq B$, if every element of A is an element of B:

$$x \in A \text{ implies } x \in B.$$

Thus $A = B$ if and only if $A \subseteq B$ and $B \subseteq A$.

This last statement is no mere tautology: it is very often how you actually prove two sets are equal. Take an arbitrary element of A, and prove it belongs to B. Then take an arbitrary element of B and prove it belongs to A.

Observe that set equality only cares about which elements belong to a set, not 'how many times they were included'. The sets

$$A = \{1, 2\} \text{ and } B = \{1, 1, 2, 2, 2, 1\}$$

are the same: every element of A is an element of B and vice versa. A variant notion of 'set' where we care about the possibility of including an element more than once is called a 'multiset', but that is a different branch of mathematics.

Note that only which elements belong to a set matter, not the way in which that set was defined. Very often new students will write monstruous manipulations of entire sets: trying to turn the conditions or properties which defined A into the conditions or properties that defined B. It is usually far better to take the elements one at a time.

Conversely, to show that two sets are different, you only need to find a single element of one set that does not belong in the other. Saying that the properties defining the two sets are different is not enough; for example, 'being an odd prime' and 'not being a power of 2' are very different properties in general, but the two sets

$$A = \{x \in \{1, \dots, 5\} \mid x \text{ is an odd prime}\}$$
$$B = \{x \in \{1, \dots, 5\} \mid x \text{ is not a power of } 2\}$$

are equal – they have the same elements 3 and 5.

Meanwhile, to show that the sets

$$A = \{x \in \{1, \dots, 7\} \mid x \text{ is an odd prime}\}$$
$$B = \{x \in \{1, \dots, 7\} \mid x \text{ is not a power of } 2\}$$

are not equal, no grandiose argument about why being prime and not being a power of 2 are different is needed. Simply point out that $6 \in B$ but $6 \notin A$.

More set constructions

So far, we have built our sets by taking subsets of things we already know about. There are other ways to build new sets from old ones. You probably have seen these before, but it is worth recapping them to establish all the notation properly.

Unions

Take a set of sets $\mathcal{A}$ – that is, each element $A \in \mathcal{A}$ will be a set. Let us write $\mathcal{A}$ as an indexed set $\mathcal{A} = \{A_i \mid i \in I\}$. We can form the *union* of the sets A_i – the set of all elements which belong to *at least one* A_i. The notation for the union is $\bigcup_{i \in I} A_i$, so we may express this definition by writing:

$$x \in \bigcup_{i \in I} A_i \text{ if and only if there exists } i \in I \text{ such that } x \in A_i.$$

If we haven't chosen an indexing for the set $\mathcal{A}$, one may see the union written simply[4] as

$$\bigcup_{A \in \mathcal{A}} A.$$

There are variant notations in the case when I is a subset of the natural numbers, with which you may be more familiar. If $I = \{1, 2\}$, so that $\mathcal{A}$ contains just two sets A_1 and A_2, we write $A_1 \cup A_2$ for the union. For $I = \{1, \ldots, n\}$ we may write any of

$$\bigcup_{i \in I} A_i \quad \text{or} \quad \bigcup_{i=1}^{n} A_i \quad \text{or} \quad A_1 \cup \cdots \cup A_n;$$

all of these mean exactly the same thing.

When $I = \{1, 2, 3, \ldots\}$ we can imitate the notation for an infinite sum by writing

$$\bigcup_{i \in I} A_i = \bigcup_{i=1}^{\infty} A_i.$$

I advise against writing '$A_1 \cup A_2 \cup \cdots$'. This makes it look like there is some sort of limiting process happening, which there isn't – the union happens 'all in one go'.

Intersections

Similarly, take a non-empty set of sets $\mathcal{A} = \{A_i \mid i \in I\}$. We can form the *intersection* of the sets A_i – the set of all elements which belong to *every* A_i. The intersection is written with the symbol $\bigcap$, so now we have the definition

$$x \in \bigcap_{i \in I} A_i \text{ if and only if } x \in A_i \text{ for all } i \in I.$$

Exercise. Why did I insist that $\mathcal{A}$ was non-empty? What happens if I is the empty set $\emptyset$? What would happen for a union?

Once again, there are several variant notations. We may write:

- $\bigcap_{A \in \mathcal{A}} A$ if $\mathcal{A}$ has not been indexed; or
- $\bigcap_{i=1}^{n} A_i$ or $A_1 \cap \cdots \cap A_n$ if $I = \{1, \ldots, n\}$; or
- $\bigcap_{i=1}^{\infty} A_i$ if $I = \{1, 2, 3 \ldots\}$.

[4] Or even more briefly as $\bigcup \mathcal{A}$, which I personally don't use – I think it is so terse as to be confusing.

Complements

In school you may have been taught the pernicious notation A^c for the complement of a set A – the set of 'everything not in A'. This is *bad notation* and should be avoided in most cases.

Why do I say this is bad notation? Consider the set E of even natural numbers. What is the 'complement' E^c? The 'set of odd numbers' is a sensible guess. But what about 1/2? Shouldn't 1/2 be in there? What about $\sqrt{-1}$? What about a fluffy unicorn? All of these are 'things that are not even numbers', so shouldn't they be in the complement E^c?

The point I am trying to make is that the notation A^c is too vague, and generates a 'set' that is 'too large' (see the box on Russell's paradox above). The phrase 'the set of everything not in A' needs to tell us what sort of 'thing' we're considering: that is, a complement is an operation of *two* sets, not one.

Given two sets X and A, the *complement* $X \smallsetminus A$ is the set of all elements of X which are not in A:

$$X \smallsetminus A = \{x \in X \mid x \notin A\}.$$

Note that it is not necessary for A to be a subset of X in this definition, although it often will be. As an example, it is perfectly sensible to write

$$\{\text{prime numbers}\} \smallsetminus \{\text{odd numbers}\} = \{2\}.$$

Not every odd number is prime, but this is irrelevant; all we are doing is taking the set of prime numbers and removing all those prime numbers which happen to be odd.

Exercise. Find an example to show that

$$X \smallsetminus (A \smallsetminus B) = (X \smallsetminus A) \cup B$$

need not always be true. Think of a condition that guarantees the above equation must hold.

Notation. There are several other symbols used for complement; as well as $X \smallsetminus A$ you will also meet the notations $X - A$ and $X \setminus A$. The minus sign $-$ is fine but occasionally looks a little misleading; sometimes students expand $X-(A-B)$ to something like $(X-A)\cup B$ because it looks like a subtraction, but this equality is false in general.

The more steeply tilted backslash $X \setminus A$ is also used in some branches of mathematics – in particular Group Theory – for an operation of 'taking a quotient on the left'. You may well see this used for a set of cosets. It is usually clear from context which use is meant, but it is good to be aware of the conflict. For this reason I generally write my set subtractions with a 45 degree slash $X \setminus A$.

The notation X/A – a forward slash – is *definitely* wrong.

In spite of my criticisms, one may still see the notation A^c in use; most commonly in Probability and related topics, where the overarching set X is fairly obvious.

In common with all mathematical notation however, it is ultimately yours to use however you like, *provided it is clearly defined*. If you are about to answer a question where you know you'll have to write $X \setminus A$ a hundred times, it will save considerable ink to say 'in this question A^c means $X \setminus A$' at the start. Once you have defined the notation, you are free to use it.

Products

An *ordered pair* is a pair of objects (a, b), regarded as a single 'symbol' in its own right. Two ordered pairs (a, b) and (a', b') are *equal* if $a = a'$ and $b = b'$.

Given two sets A and B, the (Cartesian or direct) *product* $A \times B$ of A and B is the set of all ordered pairs (a, b) where $a \in A$ and $b \in B$:

$$A \times B = \{(a, b) \mid a \in A \text{ and } b \in B\}.$$

Notation. Observe that the same notation is used for an ordered pair (a, b) and for an open interval. If I just write $(0, 1)$ with no context at all, it is unclear whether I mean the point $(0, 1) \in \mathbb{R} \times \mathbb{R}$ or the open interval

$$(0, 1) = \{x \in \mathbb{R} \mid 0 < x < 1\}.$$

This is why we must be careful with our definitions in mathematics – we do not have enough symbols to give everything a unique meaning. Including extra words in your mathematics can help – if I say 'the interval $(0, 1)$' then I remind the reader what sort of object $(0, 1)$ is.

I claimed earlier that everything in mathematics can be defined in terms of sets and the relation $\in$; but I seem to have broken that promise by introducing the new symbol (a, b). The key word was '*can*' be defined. In the most formal set logic courses, the ordered pair (a, b) is defined as the set

$$(a, b) = \{\{a\}, \{a, b\}\}.$$

You can check for yourself that this construction has the claimed property: $(a, b) = (a', b')$ if and only if $a = a'$ and $b = b'$.
So an ordered pair *can* be defined in terms of set operations; but no one in their right mind would actually think about an ordered pair in this way.

Similarly, if one has n sets $A_1, \ldots, A_n$ then one can form the product of all the A_i to be the set of 'ordered n-tuples':

$$\prod_{i=1}^{n} A_i = A_1 \times \cdots \times A_n = \{(a_1, \ldots, a_n) \mid a_i \in A_i\}.$$

We do not usually worry much about the bracketing here ('associativity of $\times$'): $(A \times B) \times C$, $A \times (B \times C)$ and $A \times B \times C$ are usually regarded as essentially the same[5], even though the elements of each look slightly different: $((a, b), c)$ versus $(a, (b, c))$ versus (a, b, c). Sometimes you want to leave the brackets in, if you're treating the different sets in different ways.

Note however that $A \times B$ and $B \times A$ are usually different sets, as they have different elements: the ordered pair $(a, b) \in A \times B$ is not an element of $B \times A$ unless $a \in B$ and $b \in A$.

When the sets in a product are equal, we often write $A \times A = A^2$, $A \times A \times A = A^3$, and so on.

Exercise. Work out what the word 'usually' means in the sentence above. That is, when is $A \times B = B \times A$ true? Be careful...

One can also try to define products of infinite families of sets. For instance, if we have one set A_i for each $i = 1, 2, 3\ldots$ then $\prod_{i=1}^{\infty} A_i$ is the set of all sequences $(a_1, a_2, a_3, \ldots)$ where $a_i \in A_i$. Infinite products can run into some weird issues of foundational logic, however, so I'm leaving them out of the 'introductory' course.

[5] In language we will meet later, there are 'canonical bijections' between them.

Power sets

If A is a set, then we can form the *power set* $\mathcal{P}(A)$ of A to be the set of all subsets of A:

$$B \in \mathcal{P}(A) \text{ if and only if } B \subseteq A.$$

Note that the empty set $\emptyset$ is always an element of $\mathcal{P}(A)$.

2 Functions

One often reads definitions of the word 'function' along the lines of:

Definition. A function $f \colon X \to Y$ is a rule that assigns to each element x of X an element $f(x)$ of Y.

This will do for now, but I dislike the word 'rule' here. It would be better to have a nondescript word like 'object' instead of 'rule' – this sounds more vague, but is less misleading. I will explain the issue with the word 'rule' shortly.

Whatever word we use, the key point is that to define a function f we need to specify *three* things:

- a set X, called the *domain* of f;
- a set Y, called the *codomain* of f; and
- for every element $x \in X$, a single element $f(x) \in Y$, called the *value of f at x*.

All three of these things are necessary; simply writing 'the function $f(x) = x^2$' is not a sufficient definition. The sets are an important part of the definition, especially the domain. The domain tells us the set of values that can be 'fed to' the function f as an input. The codomain tells us the set of possible values $f(x)$ could conceivably take, but there is no claim that all these values are actually given by $f(x)$ for some x.

Note also that the function is called 'f', not '$f(x)$'. The symbol $f(x)$ refers to a single value of the function f evaluated at the point x. This may seem like pedantry, but becomes very important if, for example, Y is a collection of functions.

Again, having claimed that 'set' and 'element' should be the only undefined terms in mathematics, I seem to be breaking the rules when I use a word like 'rule' or worse, 'object'. There is, in fact, a way to say what a function is in terms of set theory. A function $f\colon X \to Y$ may be defined as a subset $f \subseteq X \times Y$ such that for every $x \in X$ there is exactly one $y \in Y$ such that $(x, y) \in f$; that element y is then called $f(x)$. As with the formal definition of 'ordered pair', it is almost never useful to work with this definition – although it does justify the notion of 'equality of functions' I will define below.

The two most common ways of writing a function definition are via an equation '$f(x) = \ldots$', viz.:

> Define $f\colon \mathbb{R} \to \mathbb{R}$ by $f(x) = x^2 - 1$.

or using the symbol $\mapsto$, read as 'which maps … to …':

> Define $f\colon \mathbb{R} \to \mathbb{R}$, $x \mapsto x^2 - 1$.

These sentences have exactly the same meaning.

Equality of functions

Just as we did with sets, we must define what it means for two functions to be 'the same'.

Definition. Two functions f and g are *equal* if they have the *same domain* and the *same value at every point* of that domain.

You may possibly have been expecting the codomain to appear here too. It is part of the definition of a function, but is not usually important for the statement $f = g$. The codomain is most important for knowing when a composition is valid, and in questions of surjectivity (see below for both these cases).

This definition explains why the word 'rule' is misleading in the definition of the word 'function'. Consider the following three functions.

$$f\colon \{0,1\} \to \mathbb{R} \qquad g\colon \{0,1\} \to \mathbb{C} \qquad h\colon \mathbb{R} \to \mathbb{R}$$
$$x \mapsto x \qquad\qquad x \mapsto x^2 \qquad\qquad x \mapsto x^2$$

The functions f and g are certainly different 'rules' or 'formulas'; you are doing different things when computing values of f and g. However f and g have the same domain $\{0, 1\}$ and take the same value at every point of that domain, since $0^2 = 0$ and $1^2 = 1$. So f and g are the same function – there is no way to distinguish the functions f and g. *A function is not the same thing as a formula.*

Meanwhile g and h appear to be defined by the same 'rule'. They are nevertheless different functions, because they have different domains of definition: $h(2) = 4$ is a perfectly sensible value of h, but $g(2)$ is not a permissible thing to write: it is not defined.

The fact that I happened to write $\mathbb{C}$ rather than $\mathbb{R}$ for the codomain of g is not really relevant here; since all the values of g are in $\mathbb{R}$ I could have written $g\colon \{0,1\} \to \mathbb{R}$ without changing g in any important way.

As with sets, this definition of equality should be borne in mind when assessing equality of two functions with the same domain: they are equal if you can prove that they take the same value at every point of that domain. This may be in the form of a string of equation manipulations; or it may be a case-by-case breakdown.

To show that two functions on the same domain are not equal, there is no need for a complicated argument to say why the two formulae you happen to have used to define f and g are different – all you need to do is find a *single* element x *of the correct domain* such that $f(x) \neq g(x)$.

In the example above, there is clearly some relationship between the functions g and h. This relationship is expressed through the language of a 'restriction'.

Definition. Let $f\colon X \to Y$ be a function and let A be a subset of X. The *restriction* $f|_A$ *of* f *to* A is the function

$$f|_A\colon A \to Y$$

defined by $f|_A(a) = f(a)$ for all $a \in A$.

If $g\colon A \to Y$ is a function and $g = f|_A$, we may call f *an extension of* g *to* X. Note the indefinite article here; there is no such thing as 'the' extension of g to X. If you want to extend a function to a larger domain, you have to define it!

So, in our example above, we have

$$f = g = h|_{\{0,1\}}.$$

If we also included the function $k\colon \mathbb{R} \to \mathbb{R}$, $x \mapsto x$, then $k|_{\{0,1\}} = f = h|_{\{0,1\}}$ but $k \neq h$ since $k(2) \neq h(2)$.

Composition of functions

If $f\colon X \to Y$ and $g\colon Z \to W$ are functions, and the codomain Y of f is contained in the domain Z of g, one may form the *composition* $g \circ f$ (also written simply gf); this is the function $g \circ f \colon X \to W$ defined by

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

Note that it is important for the codomain of f to be a subset of the domain of g – otherwise we do not know that the value $f(x)$ is a valid input to feed to the function g.

Identity and inclusion maps

Given a set X, the *identity map*, or *identity function*, on X is the function

$$\mathrm{id}_X \colon X \to X, \quad \mathrm{id}_X(x) = x.$$

If A is a subset of X then the *inclusion map* of A into X is the function

$$\iota_A \colon A \to X, \quad \iota_A(a) = a.$$

The notation id_X is pretty common; the notation ι_A is less standard and it is best to explicitly give an inclusion map a name if you want to use it.

Inclusion maps can be useful for changing the domain and codomain of a function. Given a function $f \colon X \to Y$ and a subset A of X, we may 'shrink the domain of X to A' by forming the composition $f \circ \iota_A \colon A \to Y$. Observe that this is exactly the same thing as the restriction of f to A:

$$f|_A = f \circ \iota_A.$$

Indeed, the inclusion map itself is simply a restriction of the identity map: $\iota_A = \mathrm{id}_X |_A$.

> We can similarly 'enlarge the codomain' of f: if Y is a subset of Z, and $\iota_Y \colon Y \to Z$ is the inclusion map, then we have a function $\iota_Y \circ f \colon X \to Z$. This can be a useful thing to do if you want separate names for, say, a function $f \colon X \to \mathbb{R}$ and the function $g \colon X \to \mathbb{C}$ defined by $g(x) = f(x)$; one could write $g = \iota_{\mathbb{R}} \circ f$. Note that, unlike with a restriction, f and g have the same domain and the same values on that domain, so formally they satisfy our definition of 'equality of functions'.

Exercise. Let X and Y be non-empty sets such that X has n elements and Y has m elements (written $|X| = n$, $|Y| = m$). How many functions are there from X to Y?

Exercise. The set of all functions from X to Y is sometimes denoted Y^X. Does this make sense, given your answer to the last exercise? What is Y^X when:

1. Y is non-empty but $X = \emptyset$?
2. X is non-empty but $Y = \emptyset$?
3. X and Y are both empty?

What does this tell you about the value we 'should' assign to the quantities 1^0, 0^1 and 0^0?

'Well-defined' functions

You will often meet the phrase 'this function is well-defined' in pure mathematics. This is rather an odd sentence; there is no such thing as a function which is not 'well-defined'. This phrase is supposed to mean: "We have written down something that looks like the definition of a function. It is *well-defined* in the sense that it actually *is* a function: every element of the domain has been assigned exactly one value in the codomain".

Often it is fairly obvious that a function is well-defined, and no further comment is needed. Whenever you define a function, take a moment to consider whether it is actually well-defined. It may be that you do not need to say anything; sentences like '$x = y$ implies $f(x) = f(y)$ so f is single-valued' are a bit strange and unnecessary. But it may be that there is an issue that you need to address.

It may perhaps be best to show this concept by means of examples.

Example 2.1. Which of the following 'functions' is well-defined?

$$f\colon [0,1] \to \mathbb{R} \qquad\qquad g\colon [0,1] \to \mathbb{R}$$

$$f(x) = \begin{cases} x & 0 \le x \le 1/2 \\ 2 - x & 1/2 \le x \le 1 \end{cases} \qquad\qquad g(x) = \begin{cases} x & 0 \le x \le 1/2 \\ 1 - x & 1/2 \le x \le 1 \end{cases}$$

Here f is not well-defined – that is, it is not really a function. I have given '$f(1/2)$' two different values when it should only have one unique value. For g, we have again defined $g(1/2)$ twice, by different formulae, but since $1/2 = 1 - 1/2$ we still have the unique value $g(1/2) = 1/2$. So g is actually a function, and we may say 'g is well-defined'.

Example 2.2. Which of these functions is well-defined?

$$f\colon \mathbb{R} \to \mathbb{R} \smallsetminus \{0\} \qquad g\colon \mathbb{R} \to \mathbb{R} \smallsetminus \{0\} \qquad h\colon \mathbb{C} \to \mathbb{C} \smallsetminus \{0\}$$
$$x \mapsto x^2 - 1 \qquad\qquad x \mapsto x^2 + 4 \qquad\qquad x \mapsto x^2 + 4$$

Here f is not well-defined. Even though each point of the domain has been assigned a unique value, some of these values are not in the correct codomain: '$f(1)$' is not an element of $\mathbb{R} \smallsetminus \{0\}$. You may protest that f is sensible enough, as a function from $\mathbb{R} \to \mathbb{R}$. Indeed it is, but that isn't what we have written. Having written $f\colon \mathbb{R} \to \mathbb{R} \smallsetminus \{0\}$, we must verify that all values of f actually lie in the given codomain[6].

However, g is well-defined: $x^2 + 4$ is always in $\mathbb{R} \smallsetminus \{0\}$ for $x \in \mathbb{R}$. If you defined this function yourself in a question you might add a sentence "note that $x^2 + 4 \neq 0$ for all $x \in \mathbb{R}$, so g is well-defined".

Just like f, the 'function' h is not well-defined, because it has 'values' outside the correct codomain: $h(2i)$ is not in $\mathbb{C} \smallsetminus \{0\}$. Even though it is 'defined' by the same formula as the function g, there is a problem with h. This reinforces the message that a formula is not the same thing as a function: the domain and codomain are important too.

Example 2.3. A common one: a plausible-looking formula may fail to define a function at every point of the domain. The classic example is something like

$$f\colon \mathbb{R} \to \mathbb{R}, \quad x \mapsto 1/x$$

which does not assign any value to $x = 0$.

What you should *not* say is something like '$1/0$ doesn't exist so $f(0)$ is undefined' – the failure of f to be well-defined is *your* fault for not defining it properly! If f is defined by

$$f\colon \mathbb{R} \to \mathbb{R}, \quad f(x) = \begin{cases} 1/x & (x \neq 0) \\ 117.4 & (x = 0) \end{cases}$$

then f is well-defined! No one is claiming $1/0 = 117.4$; again, a function is not the same thing as a formula!

Example 2.4. Possibly the most important type of problem where 'well-defined' is relevant is any problem where *choice* is involved. Consider the following attempt at a function definition.

[6]This issue comes up often in Group Theory. In some texts you will see an explicit axiom of 'closure' for a group operation $*$, stating $g * h \in G$ for $g, h \in G$. Other texts will leave this more implicit, by stating that $(g, h) \mapsto g * h$ defines a function $G \times G \to G$. The condition you must check in the two cases is the same!

Define $f\colon \mathbb{Q} \to \mathbb{N}$ by $f(x) = q$ where $x = p/q$ for integers p and q.

The problem here is that there are many choices for how to write the rational number x. What value does $f(1/2)$ have? 2? But $1/2 = 2/4$, so we also have $f(1/2) = 4$! Because we haven't specified which choice to make, f is not a well-defined function (that is, it isn't a function at all).

Watch out for situations where an element may have several different 'names', or 'labels', or 'representations'.

In this case, we may fix the function definition by choosing a *specific* representation to use; for example by saying $x = p/q$ where $q > 0$ and p and q have no common factors. This specifies a unique way of writing x and we are now comfortable defining $f(x) = q$.

Here it was possible to be very explicit. Sometimes we simply fix one choice arbitrarily and use that. We may say 'for each $x \in \mathbb{Q}$ choose $p_x, q_x \in \mathbb{Z}$ with $p_x/q_x = x$ and set $f(x) = q_x$'. The great disadvantage of this is that, unless we know which values we chose, we have no idea how to compute f. This kind of arbitrary choice is common in arguments about the 'countability' of infinite sets, which you will meet towards the end of your Numbers and Sets course.

The other common circumstance in which a function depending on a choice is well-defined is when it did not, in fact, matter which choice was made. For example, the function $g\colon \mathbb{Q} \to \mathbb{Q}$ given by $g(x) = (p-q)/q$ where $x = p/q$ is well-defined, since $(p-q)/q$ has the same value $(x-1)$ no matter which representative fraction we choose for x. We will meet many examples of this type later when we talk about equivalence relations.

Exercise. 'Define' two functions $f, g\colon \mathbb{C} \setminus \{0\} \to \mathbb{R}$ by $f(re^{i\theta}) = r$, $g(re^{i\theta}) = \theta$. Which of these functions are well-defined?

Injectivity and surjectivity

There are several important adjectives that are frequently applied to functions.

Definition. A function $f\colon X \to Y$ is *injective* if, for all elements $x_1 \neq x_2$ of X, we have $f(x_1) \neq f(x_2)$. Equivalently, f is injective if $f(x_1) = f(x_2)$ implies $x_1 = x_2$.

'Injective' is sometimes also called 'one-to-one'. For an injective function, each element $y \in Y$ is equal to $f(x)$ for *at most* one $x \in X$. Two different elements of X are never mapped to the same place by f.

Note that injectivity depends critically on the domain of f. The function $f\colon (0,\infty) \to \mathbb{R}$ defined by $f(x) = x^2$ is injective, but the function $g\colon \mathbb{R} \to \mathbb{R}$ defined by the same formula $g(x) = x^2$ is not injective because $g(-1) = g(1)$.

The adjective 'injective' also has a noun form – an *injection* is an injective function. The other adjectives defined in this section have similar noun forms.

Definition. A function $f\colon X \to Y$ *surjects* Y, or *is surjective*, if for every element $y \in Y$ there exists some $x \in X$ such that $f(x) = y$.

'Surjective' is also called 'onto'. Note that there could be many $x \in X$ which are mapped to each $y \in Y$; all 'surjective' tells you is that there is *at least* one such $x \in X$.

Note that while 'injective' only cared about the domain, 'surjective' depends crucially on *both* the domain *and* the codomain used to define f.

For example, the function $f\colon \mathbb{R} \to \mathbb{R}$ defined by $x \mapsto x^2$ is not surjective because there is no $x \in \mathbb{R}$ such that $f(x) = -1$. We aren't allowed to just say '$f(x) = -1$ for $x = i$' because i is not in the domain of f.

If instead we define the function $f\colon \mathbb{R} \to [0,\infty)$ by $x \mapsto x^2$, then f *is* surjective – the value -1 is no longer in the given codomain, so we no longer care about it.

> You may have noticed that the definition of 'surjective' causes problems with our definition of 'equality of functions'. The two functions called f in the last two paragraphs are equal, but one is surjective and one is not. This is because 'surjective' depends both on the function f and a specified choice of codomain for f. If this is ever an issue, the phrasing 'f surjects Y' may be preferable.

Definition. If $f\colon X \to Y$ is both injective and surjective, then f is called *bijective.*

When f is bijective, every $y \in Y$ is equal to $f(x)$ for *exactly one* $x \in X$. We may then define a function $g\colon Y \to X$ by declaring $g(y)$ to be the unique value $x \in X$ with $f(x) = y$. This function g is called the *inverse* of f, written $g = f^{-1}$.

Exercise. Prove that $f\colon X \to Y$ is bijective if and only if there is a unique function $g\colon Y \to X$ such that $g \circ f = \mathrm{id}_X$ and $f \circ g = \mathrm{id}_Y$, where id_X denotes the identity function on a set X.

Sometimes writing down such a function g is easier than directly checking 'injective' and 'surjective'.

Exercise. Which of the following functions are injective? Which are surjective?

1. $f\colon \mathbb{C} \to \mathbb{C}$, $x \mapsto x^3$.
2. $f\colon \mathbb{R} \to \mathbb{R}^2$, $x \mapsto (x^2, x^3)$.
3. $f\colon \mathbb{R} \to (0, \infty)$, $x \mapsto e^x$.
4. $f\colon (0, \infty) \times \mathbb{R} \to \mathbb{C}$, $(x, y) \mapsto xe^{iy}$.

   [Observe the conflict here between the uses of the notation (a, b) that we mentioned before. Here $(0, \infty)$ can only reasonably mean a set – the open interval $\{x \in \mathbb{R} \mid x > 0\}$. Meanwhile (x, y) must mean an element of $(0, \infty) \times \mathbb{R}$ – an ordered pairs with $x \in (0, \infty)$ and $y \in \mathbb{R}$.]
5. $f\colon \mathbb{R} \to \mathbb{R}$ defined by

$$f(x) = \begin{cases} x & (x \geq 0), \\ -x^3 & (x \leq 0). \end{cases}$$

We are now in a position to explain what was said earlier about the sets $(A \times B) \times C$ and $A \times (B \times C)$ being 'essentially the same': the function

$$f\colon (A \times B) \times C \to A \times (B \times C)$$
$$((a, b), c) \mapsto (a, (b, c)).$$

defines a *bijection* between them.

This function f is an example of what mathematicians call a *canonical* bijection. The word 'canonical' is another word that mathematicians use a lot but defies an easy definition. Loosely, it means that no choices are made in the definition of f, which looks the same no matter which sets we are dealing with. One may also think of the word 'canonical' as meaning 'the obvious reasonable thing to write down' – if you asked two reasonable mathematicians to give you definitions of a canonical object, they will write down the same thing.

For instance, if you ask two mathematicians to give you 'the canonical injection $f\colon \mathbb{Q} \to \mathbb{R}$', they will both write down the function $x \mapsto x$. There are plenty of other injections from $\mathbb{Q}$ to $\mathbb{R}$, but a reasonable mathematician will never call a weird function like, say, $x \mapsto x\sqrt{2}/\pi^7$ the 'canonical' injection $\mathbb{Q} \to \mathbb{R}$.

Similarly, if A, B and C are sets, there may be many many bijections from $(A \times B) \times C$ to $A \times (B \times C)$ involving all manner of shuffling of elements. There is only one 'canonical' bijection though: the only reasonable, obvious map to write down which looks the same no matter what A, B and C are is the canonical bijection f defined above.

Images and pre-images

The concepts of injectivity and surjectivity are closely linked to set operations called the *image* and *pre-image*.

Definition. Let $f\colon X \to Y$ be a function.

For any subset $A \subseteq X$, the *image of* A *under* f is the set of all elements of Y which equal $f(a)$ for some $a \in A$:

$$\begin{aligned}\mathrm{im}_f(A) &= \{f(a) \mid a \in A\} \\ &= \{y \in Y \mid \text{there exists } a \in A \text{ with } y = f(a)\}.\end{aligned}$$

For a subset $B \subseteq Y$, the *pre-image of* B *under* f is the set of all elements x of X such that $f(x)$ lies in B:

$$\mathrm{preim}_f(B) = \{x \in X \mid f(x) \in B\}.$$

> Note how similar the image $\{f(a) \mid a \in A\}$ looks to an indexed set. In fact they are essentially the same thing: an indexed set $A = \{a_i \mid i \in I\}$ is the same thing as a set A with a choice of surjective function $f\colon I \to A$, $f(i) = a_i$.

Notation. The notations I have given above are very precise, and cannot reasonably be confused with anything else. They take the form of an 'image function'

$$\mathrm{im}_f\colon \mathcal{P}(X) \to \mathcal{P}(Y)$$

and a 'pre-image function'

$$\mathrm{preim}_f\colon \mathcal{P}(Y) \to \mathcal{P}(X).$$

Unfortunately, no-one really uses these notations. Instead, we 'overload' or 'abuse' the notation f: we write $f(A) = \mathrm{im}_f(A)$. Since A is a subset of X, not an element, we should not really write this, but we do[7]. That is, we use the same letter f to denote *both* the original function f *and* the image function im_f. Occasionally people will use square brackets to denote the set operation: $f[A] = \mathrm{im}_f(A)$.

Meanwhile, the pre-image function preim_f gets the notation f^{-1}:

$$f^{-1}(B) = \mathrm{preim}_f(B).$$

[7] Things get even worse if A is both an element of X and a subset... what if $X = \{x, \{x\}\}$ and $A = \{x\}$?

Writing this *does not mean* that there actually is an inverse function f^{-1} to f; we have not assumed that f is bijective. We are simply using the same symbol for the pre-image function, which always exists, as for the inverse function, which only exists if f is a bijection.

If f actually is bijective, so that the inverse function $f^{-1}: Y \to X$ exists, the two possible meanings of '$f^{-1}(B)$' agree:

$$\operatorname{im}_{f^{-1}}(B) = \operatorname{preim}_f(B) = f^{-1}(B).$$

Definition. The *image*[8] of the function $f: X \to Y$ is the set

$$\operatorname{im}(f) = f(X) = \{f(x) \mid x \in X\}.$$

In a sense, the image of f is 'the smallest sensible codomain for f'. A function always surjects its own image.

A function $f: X \to Y$ is then surjective if and only if $f(X) = Y$. Alternatively, f is surjective if and only if[9] $f^{-1}(\{y\})$ is non-empty for every $y \in Y$.

For $y \in Y$, the set $f^{-1}(\{y\})$ is the set of all elements $x \in X$ such that $f(x) = y$. So injectivity can also be expressed using pre-images: a function is injective if and only if $f^{-1}(\{y\})$ has zero or one elements for every $y \in Y$.

A function $f: X \to Y$ is bijective if and only if $f^{-1}(\{y\})$ contains exactly one element for every $y \in Y$ – the element $f^{-1}(y)$.

Exercise. As practice in proving set (in)equalities, establish the following facts about images and pre-images. Let $A, A' \subseteq X$ and $B, B' \subseteq Y$.

1. $f^{-1}(B \cap B') = f^{-1}(B) \cap f^{-1}(B')$
2. $f^{-1}(B \cup B') = f^{-1}(B) \cup f^{-1}(B')$
3. $f^{-1}(Y \setminus B) = X \setminus f^{-1}(B)$
4. $f(A \cap A') \subseteq f(A) \cap f(A')$. Show that equality is guaranteed when f is *injective*.
5. $f(A \cup A') = f(A) \cup f(A')$.
6. $f(X \setminus A) \supseteq f(X) \setminus f(A)$. Show that equality is guaranteed when f is *injective*.

[8] Some people use the word 'range' for the image. Since other people use 'range' to mean 'codomain', however, it is best not to use it for either.

[9] Some people get lazy here and write $f^{-1}(y)$ instead of $f^{-1}(\{y\})$. I urge you not to do this unless you really know what you're doing.

7. $f^{-1}(f(A)) \supseteq A$. Show that equality is guaranteed when f is *injective*.

8. $f(f^{-1}(B)) \subseteq B$. Show that equality is guaranteed when f is *surjective*.

For those facts needing an extra condition to ensure equality, find an example of a function and sets where the fact fails if f is not injective/surjective. Don't overcomplicate matters – always look for 'small' examples first.

Similar facts are true when we have arbitrary families $\{A_i \mid i \in I\}$ and $\{B_j \mid j \in J\}$ of subsets of X and Y instead of just two. If you found the above exercise easy, you may wish to think about the more general case.

3 Equivalence relations

Equivalence relations give a formal way of breaking down a set into several smaller 'pieces' or 'boxes'. Usually we would like to regard all the elements in a 'box' as sharing some property or looking in some way similar, though this is not necessary; any way of breaking down a set into 'pieces' will yield an equivalence relation, as we will soon see. Before we define equivalence relations properly, let us begin with a more mathematical formulation of the expression 'breaking down a set into pieces'.

Definition. Let X be a set. A *partition* of X is a collection of non-empty subsets $\mathcal{A} \subseteq \mathcal{P}(X) \setminus \{\emptyset\}$ such that:

- $\bigcup_{A\in\mathcal{A}} A = X$ ('every element of X is in a box'), and
- if $A, B \in \mathcal{A}$ then $A \cap B = \emptyset$ unless $A = B$ ('no element is in more than one box').

Notation. This is sometimes phrased as 'X is the *disjoint union* of the sets $A \in \mathcal{A}$', written

$$X = \bigsqcup_{A\in\mathcal{A}} A \text{ or } X = \dot{\bigcup_{A\in\mathcal{A}}} A.$$

> Note that there is now a well-defined function
>
> $$f_{\mathcal{A}} \colon X \to \mathcal{A}$$
>
> defined by $f(x) = A$, where A is the unique element of $\mathcal{A}$ with $x \in A$.

For example, the following are partitions of the set of natural numbers $\mathbb{N}$:

- $\mathcal{A} = \{\{\text{even numbers}\}, \{\text{odd numbers}\}\}$.

- $\mathcal{A} = \{\{\text{primes}\}, \{\text{composites}\}, \{0\}, \{1\}\}$.
- $\mathcal{A} = \{\{\text{numbers ending in } 0\}, \ldots, \{\text{numbers ending in } 9\}\}$.

The set of all Cambridge undergraduate students may be partitioned into {{Jesus students}, {Selwyn students}, ...}.

Writing down all the sets of a partition like this, and listing all the elements of each set, can be a lengthy process and often it may be easier instead to give a criterion for deciding when two elements of X belong in the same box. The last example of a partition of $\mathbb{N}$ given above would be much better written by saying 'two numbers go in the same box if they have the same last digit'. For the Cambridge college example, instead of listing all the Colleges and listing all the students in each college, it would be better to say 'two students are in the same set if they are in the same college'.

This leads us to the idea of an *equivalence relation* – a partition of X allows us to say, roughly, 'x and y are related if they're in the same subset $A \in \mathcal{A}$'. What properties would this notion of 'related' have? Certainly each x is in the same box as itself; if x and y are in the same box, then y and x are in the same box; and if x and y are in the same box, and y and z are in the same box, then x and z are in the same box. These three properties show up as conditions (R), (S) and (T) in the formal definition below.

Definition. A *relation between a set X and a set Y* is a subset $\mathcal{R} \subseteq X \times Y$. A *relation on a set X* is a relation between X and itself: a subset $\mathcal{R} \subseteq X \times X$. We may write $x\mathcal{R}y$ instead of $(x, y) \in \mathcal{R}$.

A relation $\mathcal{R}$ on a set X is an *equivalence relation* on X if $\mathcal{R}$ satisfies the following three conditions.

(R) 'reflexive': $x\mathcal{R}x$ for every $x \in X$.

(S) 'symmetric': if $x\mathcal{R}y$ then $y\mathcal{R}x$.

(T) 'transitive': if $x\mathcal{R}y$ and $y\mathcal{R}z$ then $x\mathcal{R}z$.

If $\mathcal{A}$ is a partition of X then the relation '$x\mathcal{R}y$ if and only if x and y are in the same set $A \in \mathcal{A}$' is an equivalence relation.

Exercise. For each of the eight possible combinations of (R), (S) and (T), find a set X and a relation $\mathcal{R}$ on X which satisfies those properties but not the others. Try to keep your set X as small as possible.

Exercise. Consider the following argument.

> Suppose $\mathcal{R}$ satisfies (S) and (T). Let $x \in X$ and take some $y \in X$ with $x\mathcal{R}y$. Then by (S) we have $y\mathcal{R}x$, and by (T) these two conditions force $x\mathcal{R}x$. So (S) plus (T) implies (R).

What do you think of this argument?

Exercise. Let X be the set of all triangles in the plane. Which of the following relations on X are equivalence relations?

1. $x\mathcal{R}y$ if x and y have one equal side.
2. $x\mathcal{R}y$ if x and y have the same set of side lengths.
3. $x\mathcal{R}y$ if x and y are congruent.
4. $x\mathcal{R}y$ if x and y have the same area.
5. $x\mathcal{R}y$ if x and y share a vertex.

Do 2 and 3 define the same relation?

We have already seen that partitions give rise to equivalence relations; in fact the converse is also true.

Definition. Let X be a set and let $\mathcal{R}$ be an equivalence relation on X. For each $x \in X$, the *equivalence class of* x is the set

$$[x]_{\mathcal{R}} = \{y \in X \mid y\mathcal{R}x\}.$$

The *set of equivalence classes* of X *modulo* $\mathcal{R}$ (more rarely 'the quotient of X by $\mathcal{R}$') is the set

$$X/\mathcal{R} = \{[x]_{\mathcal{R}} \mid x \in X\}.$$

Claim. *$X/\mathcal{R}$ is a partition of X.*

Proof. If $x \in X$ then, by property (R), we know $x \in [x]_{\mathcal{R}}$. So every element belongs to some equivalence class, and X is the union of the equivalence classes.

Now suppose that $[x]_{\mathcal{R}} \cap [z]_{\mathcal{R}} \neq \emptyset$. We are required to prove that $[x]_{\mathcal{R}} = [z]_{\mathcal{R}}$. Let $y \in [x]_{\mathcal{R}} \cap [z]_{\mathcal{R}}$, so that $y\mathcal{R}x$ and $y\mathcal{R}z$. By (S) we also have $x\mathcal{R}y$ so by (T) we have $x\mathcal{R}z$.

For any $w \in [x]_{\mathcal{R}}$ we have $w\mathcal{R}x$. Applying (T) again we have $w\mathcal{R}z$. So $w \in [z]_{\mathcal{R}}$ also. We conclude $[x]_{\mathcal{R}} \subseteq [z]_{\mathcal{R}}$. Similarly, $[x]_{\mathcal{R}} \supseteq [z]_{\mathcal{R}}$ and the two equivalence classes are equal as required. □

Note once again that the sets $[x]_{\mathcal{R}}$ and $[z]_{\mathcal{R}}$ are equal *because they have the same elements*. The elements x and z serve as different 'labels' for this one set.

We see that equivalence relations and partitions are two sides of the same coin. Equivalence relations are *often* a more powerful mathematical tool: writing a partition needs us to decide exactly which elements belong to each set, but writing an equivalence relation only needs us to give rules to decide when two elements belong in the same set.

Of course, this is not *always* true; if I partition the set $\{1, \ldots, 7\}$ into

$$\mathcal{A} = \{\{1, 2, 6\}, \{4, 5, 7\}, \{3\}\},$$

there is no particularly neat equivalence relation describing this partition. This example is better specified as a partition rather than an equivalence relation.

A class of equivalence relations which appears all over the first-year syllabus is the following example. Take $X = \mathbb{Z}$ and let n be a positive integer. Define an equivalence relation $\sim_n$ on $\mathbb{Z}$ by

$$x \sim_n y \text{ if and only if } n \text{ divides } x - y.$$

There are n equivalence classes here; one for each possible 'remainder modulo n'. We may choose to label these equivalence classes by the numbers $0, \ldots, n-1$, and we often do, but these are only a choice of labels. The equivalence class $[1]_{\sim_n}$ could equally well be labelled with $n+1$, $2n+1$, $1-n$, and so on.

There are many variant notations for this equivalence relation. Perhaps the most common is

$$x \sim_n y \text{ if and only if } x \equiv y \text{ modulo } n,$$

but there are others as well. In Group Theory you will encounter such notations as $x + n\mathbb{Z} = y + n\mathbb{Z}$ with exactly the same meaning.

The set of equivalence classes $\mathbb{Z}/{\sim_n}$ is often denoted $\mathbb{Z}_n$ or, especially in Group Theory and its relatives, $\mathbb{Z}/n\mathbb{Z}$. Note that $\mathbb{Z}/{\sim_n}$ has exactly n elements. Each of these elements is an equivalence class; and each equivalence class is an infinite set of integers.

Equivalence classes and well-defined functions

A very common situation in pure mathematics is the need to define a function out of a set of equivalence classes. Given sets X and Y and an equivalence relation $\mathcal{R}$ on X, how might we define a function $f\colon X/\mathcal{R} \to Y$?

The natural thing to do is to take an equivalence class and label it by one of its elements: $[x]_{\mathcal{R}}$. We might then write some expression depending on x; for instance, a function like

$$\begin{aligned} f\colon \mathbb{Z}/{\sim_n} &\to \mathbb{Z}/{\sim_n} \\ [x]_{\sim_n} &\mapsto [x+1]_{\sim_n}. \end{aligned}$$

However, we have made a *choice* here. We do not yet know that this function is well-defined; what if we assigned different values to $f([x]_{\mathcal{R}})$ depending on which label x we chose?

For such a situation, there are essentially two resolutions. One is to fix a specific label for each equivalence class and work with that. For example, for the equivalence relation $\sim_n$ on $\mathbb{Z}$, each equivalence class $[x]_{\sim_n}$ contains exactly one element $r_{[x]}$ such that $0 \leq r < n$. This gives us an unambiguous 'label' for each equivalence class, which can be used to define functions. For instance, the function

$$f\colon \mathbb{Z}/{\sim_n} \to \mathbb{Z}, \quad [x]_{\sim_n} \mapsto r_{[x]}$$

is well-defined; we have said exactly where to send each equivalence class, and there is no ambiguity in which label we chose.

The other resolution – by far the more common one – is to show that the value assigned to $f([x]_{\mathcal{R}})$ is actually independent of which choice of label we made. For example, the function $[x]_{\sim_n} \mapsto [x+1]_{\sim_n}$ given above is well-defined. If $[x]_{\sim_n} = [y]_{\sim_n}$ then n divides $x - y$. But then n also divides $(x+1)-(y+1)$ so $[x+1]_{\sim_n} = [y+1]_{\sim_n}$. Changing the label from x to y did not change the value assigned to $f([x]_{\sim_n})$, so this function is well-defined.

On the contrary, something like

$$f\colon \mathbb{Z}/{\sim_7} \to \mathbb{Z}/{\sim_4}, \quad [x]_{\sim_7} \mapsto [x]_{\sim_4}$$

is *not* well-defined. The equivalence class $[1]_{\sim_7}$ could be labelled by, among other things, 1 and 8; but $[1]_{\sim_4}$ is *not* equal to $[8]_{\sim_4}$. So in this case the output of f depends on a choice of label, and the function is not well-defined (that is, it is not really a function).

Exercise. For which pairs of integers n and m is the 'function'

$$f\colon \mathbb{Z}/{\sim_n} \to \mathbb{Z}/{\sim_m}, \quad [x]_{\sim_n} \mapsto [x]_{\sim_m}$$

well-defined?

In several of the above examples, the codomain Y also happened to be a set of equivalence classes. This is common but by no means necessary.

You may have noticed that what we really did in the above examples was to write down a function $\hat{f}\colon X \to Y$ and attempt to define a function $f\colon X/\mathcal{R} \to Y$ by the formula

$$f([x]_{\mathcal{R}}) = \hat{f}(x).$$

In order for such an f to be a well-defined function, we need its output to not depend on the choice of label x for an equivalence class. That is, we need to know that $x\mathcal{R}y$ implies $\hat{f}(x) = \hat{f}(y)$. If this is the case, we say 'f respects the equivalence relation $\mathcal{R}$'. Sometimes it helps clean the writing up to give a name to $\hat{f}$, but often we just leave it implicit as in the above examples.

Indeed, every function out of $X/\mathcal{R}$ can be written in this way. If $\mathcal{R}$ is an equivalence relation on X, denote by $q\colon X \to X/\mathcal{R}$ the 'quotient map' or 'quotient function' defined by

$$q(x) = [x]_{\mathcal{R}}.$$

Then for each function $f\colon X/\mathcal{R} \to Y$ the composition $\hat{f} = f \circ q$ is a function from X to Y which respects $\mathcal{R}$.
We may summarise this discussion as a formal lemma.

Lemma. *If $f\colon X/\mathcal{R} \to Y$ is a (well-defined) function then $\hat{f} = f \circ q$ respects the equivalence relation $\mathcal{R}$. Conversely, if $\hat{f}\colon X \to Y$ respects the equivalence relation $\mathcal{R}$ then the function*

$$f\colon X/\mathcal{R} \to Y, \quad f([x]_{\mathcal{R}}) = \hat{f}(x)$$

is well-defined and $f \circ q = \hat{f}$.

This kind of lemma, with various added adjectives, will appear again and again in Group Theory, Linear Algebra, Topology...

Exercise. Consider the following relation on the set $X = \mathbb{Z} \times (\mathbb{Z} \smallsetminus \{0\})$:

$$(a, b) \sim (c, d) \text{ if and only if } ad = bc.$$

1. Show that $\sim$ is an equivalence relation. What happens if you try to define this relation on $\mathbb{Z} \times \mathbb{Z}$?

2. Show that the function

$$f\colon (X/\sim) \times (X/\sim) \to X/\sim$$
$$f([(a,b)]_\sim, [(c,d)]_\sim) = [(ad+bc, bd)]_\sim$$

is well-defined.

3. Construct a bijection $g\colon X/\sim \to \mathbb{Q}$.

Note: we have here a function of *two* equivalence classes; you need to check what happens when you relabel *both*. Is there a more efficient way to do that?