‘A brilliant gourmet feast of what maths is really about’ Ian Stewart

HOW TO BAKE π

Easy recipes for understanding complex maths

EUGENIA CHENG

‘Quirky recipes, personal anecdotes and a large dollop of equations are the key ingredients in this alternative guide to maths and the scientific process. You should find it as easy as cooking a pie.’ *Observer*

First published in Great Britain in 2015 by
PROFILE BOOKS LTD
3 Holford Yard
Bevin Way
London
WC1X 9HD
www.profilebooks.com

A CIP catalogue record for this book is available from the British Library.

eISBN 978 1 78283 082 5

CONTENTS

Prologue

part one | MATHEMATICS

1 What Is Maths?
2 Abstraction
3 Principles
4 Process
5 Generalisation
6 Internal vs. External
7 Axiomatisation
8 What Maths Is

part two | CATEGORY THEORY

9 What Is Category Theory?
10 Context
11 Relationships
12 Structure
13 Sameness
14 Universal Properties
15 What Category Theory Is

Acknowledgements
Index

To
my parents
and Martin Hyland

In memory of
Christine Pembridge

They say mathematics is a glorious garden. I know I would certainly lose my way in it without your guidance. Thank you for walking us through the most beautiful entrance pathway.

From a student's letter to the author
University of Chicago, June 2014

PROLOGUE

Here is a recipe for clotted cream.

Ingredients

Cream

Method

1. Pour the cream into a rice cooker.
2. Leave it on ‘warm’ with the lid slightly open, for about 8 hours.
3. Cool it in the fridge for about 8 hours.
4. Scoop the top part off: that’s the clotted cream.

What on earth does this have to do with maths?

Maths myths

Maths is all about numbers.

You might think that rice cookers are for cooking rice. This is true, but this same piece of equipment can be used for other things as well: making clotted cream, cooking vegetables, steaming a chicken. Likewise, maths is about numbers, but it's about many other things as well.

Maths is all about getting the right answer.

Cooking is about ways of putting ingredients together to make delicious food. Sometimes it's more about the method than the ingredients, just as in the recipe for clotted cream, which only has one ingredient – the entire recipe is just a method. Maths is about ways of putting ideas together, to make exciting new ideas. And sometimes it's more about the method than the 'ingredients'.

Maths is all either right or wrong.

Cooking can go wrong – your custard can curdle, your soufflé can collapse, your chicken can be undercooked and give everyone food poisoning. But even if it doesn't poison you, some food tastes better than other food. And sometimes when cooking goes 'wrong' you have actually accidentally invented a delicious new recipe. Fallen chocolate soufflé is deliciously dark and squidgy. If you forget to melt the chocolate for your cookies, you get chocolate chip cookies. Maths is like this too. At school if you write 10 + 4 = 2 you will be told that is wrong, but actually that's correct in some circumstances, such as telling the time – four hours later than 10 o'clock is indeed 2 o'clock. The world of maths is more weird and wonderful than some people want to tell you…

You're a mathematician? You must be really clever.

Much as I like the idea that I am very clever, this popular myth shows that people think maths is hard. The little-understood truth is that the aim of maths is to make things easier. Herein lies the problem – if you need to make things easier it gives the impression that they were hard in the first place. Maths *is* hard, but it makes hard things easier. In fact, since maths is a hard thing, maths also makes maths easier.

Many people are either afraid of maths, or baffled by it, or both. Or they were completely turned off it by their lessons at school. I understand this – I was completely turned off sport by my lessons at school, and have never really recovered. I was so bad at sport at school, my teachers were incredulous that anybody so bad at sport could exist. And yet I'm quite fit now, and have even run the New York Marathon. At least I now appreciate physical exercise, but I still have a horror of any kind of team sport.

How can you do research in maths? You can't just discover a new number.

This book is my answer to that question. It's hard to answer it quickly at a cocktail party, without sounding trite, or taking up too much of someone's time, or shocking the gathered company. Yes, one way to shock people at a polite party is to talk about maths.

It's true, you can't just discover a new number. So what can we discover that's new in maths? In order to explain what this 'new maths' could possibly be about, I need to clear up some misunderstandings about what maths is in the first place. Indeed, not only is maths not just about numbers, but the branch of maths I'm going to describe is actually not about numbers at all. It's called *category theory* and it can be thought of as the 'mathematics of mathematics'. It's about relationships, contexts, processes, principles, structures, cakes, custard.

Yes, even custard. Because mathematics is about drawing analogies, and I'm going to be drawing analogies with all sorts of things to explain how maths works. Including custard, cake, pie, pastry, doughnuts, bagels, mayonnaise, yoghurt, lasagne, sushi.

Whatever you think maths is … let go of it now.
This is going to be different.

part one

MATHEMATICS

1 WHAT IS MATHS?

Gluten-free chocolate brownies

Ingredients

115 g butter
125 g dark chocolate
150 g caster sugar
80 g potato flour
2 medium eggs

Method

1. Melt the butter and chocolate, stir together and allow to cool a little.
2. Whisk the eggs and the sugar together until fluffy.
3. Beat the chocolate into the egg mixture slowly.
4. Fold in the potato flour.
5. Bake in very small individual cases at 180°C for about 10 minutes, or until they're as cooked as you want them.

Maths, like recipes, has both ingredients and method. And just as a recipe would be a bit useless if it omitted the method, we can't understand what maths is unless we talk about the *way it is done*, not just the *things it studies*. Incidentally the method in the above recipe is quite important – these don't cook very well in a large tray. In maths the method is perhaps even more important than the ingredients. Maths probably isn't whatever you studied at school in lessons called 'maths'. Yet somehow I always knew that maths was more than what we did at school. So what *is* maths?

Recipe books

What if we organised recipes by equipment?

Cooking often proceeds a bit like this: you decide what you want to cook, you buy the ingredients, and then you cook it. Sometimes it might work the other way round: you go wandering around the shops, or maybe a market. You see what ingredients look good, and you feel inspired by them to conjure up your meal. Perhaps there's some particularly fresh fish, or a type of mushroom you've never seen before, and you go home and look up what to do with it afterwards.

Occasionally something completely different happens: you buy a new piece of equipment, and suddenly you want to try making all sorts of different things with that equipment. Perhaps you bought a blender, and suddenly you make soup, smoothies, ice cream. You try making mashed potatoes in it, and it goes horribly wrong (it looks like glue). Maybe you bought a slow cooker. Or a steamer. Or a rice cooker. Perhaps you learn a new technique, like separating eggs or clarifying butter, and suddenly you want to make as many things as possible involving your new technique.

So we might approach cooking in two ways, and one seems much more practical than the other. Most recipe books are divided up according to parts of the meal rather than by techniques. There's a chapter on starters, a chapter on soup, a chapter on fish, a chapter on meat, a chapter on dessert, and so on. There might be a whole chapter on an ingredient, say a chapter on chocolate recipes or vegetable recipes. Sometimes there are whole chapters on particular meals, say a chapter on Christmas lunch. But it would be quite odd to have a chapter on 'recipes that use a rubber spatula' or 'recipes that use a balloon whisk'. Having said that, kitchen gadgets often come with handy books of recipes you can make with your new equipment. A blender will come with blender recipes; likewise a slow cooker or an ice cream maker.

Something similar is true of subjects of research. Usually when you say what a subject is, you describe it according to the thing that you're studying. Maybe you study birds, or plants, or food, or cooking, or how to cut hair, or what happened in the past, or how society works. Once you've decided what you're going to study, you learn the techniques for studying it, or you invent new techniques for studying it, just like learning how to whisk egg whites or clarify butter.

In maths, however, the things we study are also determined by the techniques we use. This is similar to buying a blender and then going round seeing what you can make with it. This is more or less backwards from other subjects. Usually the techniques we use are determined by the things we're studying; usually we decide what we want for dinner, and then get out the equipment for making it. But when we're really excited about our new blender, we go round trying to make all our dinners in it for a while. (At least, I've seen people do this.)

It's a bit of a chicken-and-egg question, but I am going to argue that maths is defined by the techniques it uses to study things, and that the things it studies are determined by those techniques.

Cubism

When the style affects the choice of content

Characterising maths by the techniques it uses is similar to defining styles of art, like cubism or pointillism or impressionism, where the genre is defined by the techniques rather than the subject matter. Or ballet and opera, where the art form is defined by the methods, and the subject matter is duly restricted. Ballet is very powerful at expressing emotion, but not so good at expressing dialogue, or making demands for political change. Cubism is not that effective for depicting insects. Symphonies are good at expressing tragedy and joy, but not very good at saying ‘Please pass the salt.’

In maths the technique we use is *logic*. We only want to use sheer logical reasoning. Not experiments, not physical evidence, not blind faith or hope or democracy or violence. Just logic. So what are the things we study? We study *anything that obeys the rules of logic*.

> **Mathematics is the study of anything that obeys the rules of logic, using the rules of logic.**

I will admit immediately that this is a somewhat simplistic definition. But I hope that after reading some more you’ll see why this is accurate as far as it goes, not as circular as it sounds at first, and just the sort of thing a category theorist would say.

The prime minister

Characterising something by what it does

Imagine if someone asked you ‘Who’s the prime minister?’ and you answered ‘He’s the head of the government.’ This would be correct but annoying, and not really answering the right question: you’ve characterised the prime minister without telling us who it is. Likewise, my ‘definition’ of mathematics has *characterised* maths rather than telling you what it is. This is a little unhelpful, or at least incomplete – but it’s just the start.

Instead of describing what maths is *like*, can we say what maths *is*? What does maths actually study? It definitely studies numbers, but also other things like shapes, graphs and patterns, and then things that you can’t see – logical ideas. And more than that: things we don’t even know about yet. One of the reasons maths keeps growing is that once you have a technique, you can always find more things to study with it, and then you can find more techniques to use to study those things, and then you can find more things to study with the new techniques, and so on, a bit like chickens laying eggs that hatch chickens that lay eggs that hatch chickens…

Mountains

Conquering one enables you to see the higher ones

Do you know that feeling of climbing to the top of a hill, only to find that you can now see all the higher hills beyond it? Maths is like that too. The more it progresses, the more things it comes up with to study. There are, broadly, two ways this can happen.

First there's the process of 'abstraction'. We work out how to think logically about something that logic otherwise couldn't handle. For example, you previously only made rice in your rice cooker, and then you work out that you can use it to make cake, it's just a bit different from cake made the normal way in an oven. We take something that wasn't really maths before, and look at it differently to turn it into maths. This is the reason that *x*'s and *y*'s start appearing – we start by thinking about numbers, but then realise that the things we do with numbers can be done with other things as well. This will be the subject of the next chapter.

Secondly there's the process of 'generalisation': we work out how to build more complicated things out of the things we've already understood. This is like making a cake in your blender, and making the icing in your blender, and then piling it all up.[1] In maths this is how we get things like polynomials and matrices, complicated shapes, four-dimensional space, and so on, out of simpler things like numbers, triangles and our everyday world. We'll look into this in Chapter 5.

These two processes, abstraction and generalisation, will be the subject of the next few chapters, but first I want to draw your attention to something weird and wonderful about how maths does these two things.

Birds

They are not the same as the study of birds

Imagine for a second that you study birds. You study their behaviour, what they eat, how they mate, how they look after their young, how they digest food, and so on. However, you will never be able to build a new bird out of simpler birds – that just isn't how birds are made. So you can't do generalisation, at least not in the way that maths does it.

Another thing you can't do is take something that isn't a bird, and miraculously turn it into a bird. That also isn't how birds are made. So you can't do abstraction either. Sometimes we realise we've made a mistake of classification – for example the brontosaurus 'became' a form of apatosaurus. However, we didn't turn the brontosaurus into an apatosaurus – we merely realised it had been one all along. We're not magicians, so we can't change something into something it isn't. But in maths we can, because maths studies ideas of things, rather than real things, so all we have to do to change the thing we're studying is to change the idea in our head. Often, this means changing the way we think about something, changing our point of view or changing how we express it.

A mathematical example is knots.

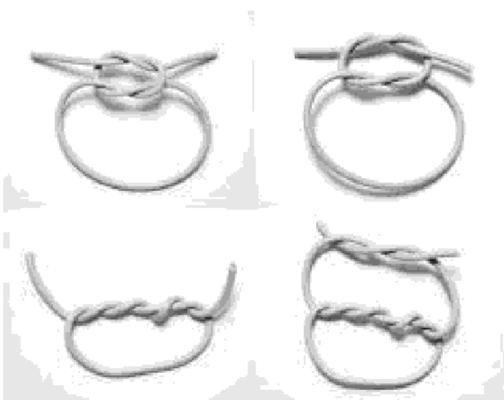

In the eighteenth and nineteenth centuries, Vandermonde, Gauss and others worked out how to think of knots mathematically, so that they could be studied using the rules of logic. The idea is to imagine sticking together the two ends of the piece of string so that it has become a closed loop. This makes the knots impossible to make without glue, but much easier to reason with mathematically. Each one can be expressed as a circle that has been mapped to three-dimensional space. There are many techniques for studying this kind of thing in the field of *topology*, which we'll come back to later. We can then deduce things not only about real knots in string, but also about the apparently impossible ones that arise in nature in molecular structures.

Geometrical shapes are another, much older example of this process of turning something from the 'real' world into something in the 'mathematical' world.

We can think of maths as developing in the following stages:

1. It started as the study of numbers.
2. Techniques were developed to study those numbers.
3. People started realising that those techniques could be used to study other things.
4. People went round looking for other things that could be studied like this.

Actually there's a step 0, before the study of numbers: someone had to come up with the idea of numbers in the first place. We think of them as the most basic things you can study in maths, but there was a time before numbers. Perhaps the invention of numbers was the first ever process of *abstraction*.

The story I'm going to tell is about abstract mathematics. I'm going to argue that its power and beauty lie not in the answers it provides or the problems it solves, but in the *light* that it sheds. That light enables us to see clearly, and that is the first step to understanding the world around us.

Footnote

[1] Mathematical generalisation isn’t the same as the kind where you go round making sweeping statements about things, but we’ll come to that later.

2 ABSTRACTION

Mayonnaise or hollandaise sauce

Ingredients

2 egg yolks
300 ml olive oil
Seasoning

Method

1. Whisk the egg yolks and seasoning using a hand whisk or immersion blender.
2. Drip the olive oil in very slowly, while continuing to whisk. For hollandaise sauce, use 100 g melted butter instead of the olive oil.

At some level mayonnaise and hollandaise sauce are the same – they use the same method, but with a different type of fat incorporated into the egg yolk. In both cases, the amazing near-magic properties of egg yolks create something rich and unctuous. It looks so much like magic, I never tire of watching it happen.

The similarity between mayonnaise and hollandaise sauce is the sort of thing that mathematics goes round looking for – situations where things are somehow the same apart from some small detail. This is a way of saving effort, so that you can understand how to do both things at once. Books might tell you that hollandaise sauce needs to be done differently, but I ignore them to make my life simpler. Maths is also there to make things simpler, by finding things that look the same if you ignore some small details.

Pie

Abstractions as blueprints

Cottage pie, shepherd's pie and fisherman's pie are all more or less the same – the only difference is the filling that is sitting underneath the mashed potato topping. Crumble is also very similar – you don't really need a different recipe for different types of crumble, you just need to know how to make the crumble part. Then you put the fruit of your choice in a dish, and put the crumble on top, and bake it.

Another favourite of mine is upside-down cake. You put the fruit in the bottom of the cake tin, pour the cake mix on top, and after baking it you turn it out upside down so that the fruit is on top. For extra effect you can put melted butter and brown sugar on the bottom of the cake tin first, to caramelise the fruit a bit. Of course, this works better with some fruit than others: bananas, apples, pears and plums work well. Grapes less well. Watermelon would be terrible. The same is true for crumble. Watermelon crumble? Probably not.

Savoury tarts and quiches also follow a general pattern. You bake an empy pastry case, put in some filling of your choice, and then top it up with a mixture of egg and milk or cream, before baking it again. The filling could be bacon and cheese, or fish, or vegetables – whatever you feel like.

In all these cases the 'recipe' is not a full recipe, but a blueprint. You can insert your own choice of fruit, or meat, or fillings to make your own variations, within reason.

This is also how maths works. The idea of maths is to look for similarities between things so that you only need one 'recipe' for many different situations. The key is that when you ignore some details, the situations become easier to understand, and you can fill in the variables later. This is the process of abstraction.

As with the watermelon crumble, once you've made the abstract 'recipe' you will find that you won't be able to apply it to *everything*. But you are at least in a position to try, and sometimes surprising things turn out to work in the same recipe.

Think about the symmetry of an equilateral triangle:

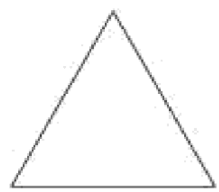

There's reflectional symmetry, and rotational symmetry. How can we describe the different symmetries without cutting out the triangle and folding it up or waving it around?

One way is that we could label the corners 1, 2 and 3,

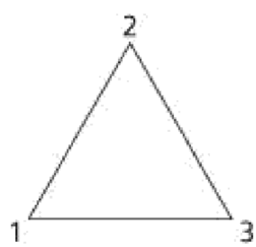

and then just talk about how the numbers get swapped around. For example, if we reflect the triangle in a vertical line, we will swap the numbers 1 and 3. Whereas if we rotate the triangle 120° clockwise we will send 1 to where 2 was, 2 to where 3 was, and 3 to where 1 was.

You can try checking that the six symmetries of the triangle correspond exactly to the six different ways of shuffling the numbers 1, 2 and 3. There are three lines of symmetry, and they correspond to swapping 1 and 3, or 1 and 2, or 2 and 3. There are three types of rotational symmetry: 120° clockwise, 240° clockwise, and the 'trivial' one where nothing moves.

This shows that the symmetry of an equilateral triangle is *abstractly* the same as the permutations of the numbers 1, 2, 3, and the two situations can be studied at the same time.

Kitchen clutter

Abstraction as tidying away the things you don't need

Abstraction is like preparing to cook something, and putting away the equipment and ingredients that you don't need for the recipe, so that your kitchen is less cluttered. It is the process of putting away the ideas you don't need for the present purposes, so that your *brain* is less cluttered.

Are you better at this in your kitchen, or in your brain? (I am definitely better at it in my brain.) Abstraction is the important first step of doing mathematics. It's also a step that can make you feel uneasy because you're stepping away from reality a little bit. I never put my food processor away because it's such a hassle to move it, and I want to know that I can use it any minute now without going through the rigmarole of getting it out of the cupboard. You might feel like that about abstraction in the brain as well.

Try the following problem:

I buy two stamps for 36p each. How much does it cost?

When children do this sort of thing at primary school it sometimes get called a 'word problem', because it has been stated in words, and they're told that the first step in solving this 'word problem' is to turn it into numbers and symbols:

$$36 \times 2 = ?$$

This is a process of abstraction. We have thrown away, or ignored, the fact that the thing we were buying was *stamps*, because it didn't make any difference to the answer. It could have been apples, bananas, monkeys, …, the sum would still be the same, and so the answer would still be the same: 72p.

What about this one:

My father is three times as old as I am now but in ten years' time he will be twice as old as me. How old am I?

Or this one:

> *I have a recipe for icing the top and sides of a 6-inch cake. How much icing do I need for the top and sides of an 8-inch cake?*

For the question about stamps you probably didn't need to write down a sum, because the answer was immediately obvious to you. However, for these last two questions, perhaps you would need to perform some abstraction to work out the answer, where you throw away the fact that you're talking about your father, or a cake and icing, and write down some sums, with numbers and symbols. We'll see what sums we get from these word problems a bit later in this chapter.

Sweets

How things that are too real don't obey mathematics

If you've ever tried teaching arithmetic to small children, you might have come up with the following problem. You try and get them to think about a real-life situation such as:

> *If grandma gives you five sweets and grandpa gives you five sweets, how many sweets will you have?*

And the child answers: 'None, because I'll eat them all!'

The trouble here is that sweets do not obey the rules of logic, so using maths to study them doesn't quite work. Can we force sweets to obey logic? We could impose an extra rule on the situation by adding '… and you're not allowed to eat the sweets'. If you're not allowed to eat them, what's the point of them being sweets? We could treat the sweets as just *things* rather than sweets. We lose some resemblance to reality, but we gain scope and with it efficiency. The point of numbers is that we can reason about 'things' without having to change the reasoning depending on what 'thing' we are thinking about. Once we know that 2 + 2 = 4 we know that two things and another two things make four things, whether they are sweets, monkeys, houses, or anything else. That is the process of abstraction: going from sweets, monkeys, houses, or whatever, to numbers.

Numbers are so fundamental, it's difficult to imagine life without them, and difficult to imagine the process of inventing them. We don't even notice that we're making a leap of abstraction when we count things. It's much more noticeable if you watch small children struggling to do it, because they're not yet used to making that leap.

Eeny meeny miny moe

Numbers as an abstraction

I remember a wonderfully feisty mother at a primary school I was helping at. She also helped there, and remarked on how frustrating it was when other mothers competitively declared that *their* child could count up to 20 or 30. 'My son can count up to three,' she said defiantly. 'But he knows what three *is*.'

And she had a point.

When a child first 'learns to count to ten' they aren't really doing more than learning to recite a little poem, like 'Incy wincy spider climbed up the spout…' It just so happens that the little 'poem' goes:

'*One, two, three, four, five, six, …*'

Then they learn that this has something to do with pointing at things, so they start pointing while reciting the 'poem', a bit haphazardly.

Next they learn that they're supposed to point at one thing per word in the poem, but they have trouble making sure they have only pointed at each thing once, so they will get rather variable answers if you ask them 'How many ducks are in this picture?' Or they might latch on to a particular number – say, six – and somehow manage to count everything as being six, no matter how many ducks there really are.

Finally they'll get the idea that they're supposed to match get to the heart, you have to strip away clothes and skin and flesh and bone.

Road signs

Abstraction as the study of ideal versions of things

Road signs are a form of abstraction. They don't precisely depict what is going on in the road, but represent some idealised form of it, where just the essence is captured. Not every humpbacked bridge looks exactly like this:[1]

but this captures the essence of humpbacked-bridge-ness. Similarly, not all children crossing the road look exactly like this:

Nevertheless the benefits of this system are clear. It's much quicker to take in a symbol than read some words while you are driving. Also it's much easier for foreigners to understand. The disadvantage is that when you first start driving you have to learn what all these funny symbols mean. Some of them, for example

are much closer to reality than others, for example

This 'No Entry' sign is entirely abstract: it doesn't look like the thing it is representing at all. (What does 'No Entry' look like?) But it's also more important – you will probably encounter more of those in your driving life than the one warning you there might be deer crossing the road.

One side effect of the abstraction of maths is that a variety of funny symbols get used as well, for the same sorts of reasons: once you know what they mean, the symbols are much quicker to take in, and you can reserve your mathematical brain power for the more complicated parts of the maths you're supposed to be focusing on. It also makes the maths easier to understand across different languages – it’s surprisingly easy to read a maths book in a language you don’t know.

> The most basic ‘funny symbols’ used in maths are the ones for normal arithmetic: +, –, ×, ÷, =. Once you’re comfortable with these symbols, it’s much quicker and easier to read
>
> $$2 + 2 = 4$$
>
> than ‘two plus two equals four’. As maths gets more and more complicated, the symbols get more and more complicated as well, with things like
>
> $$\sum, \int, \oint, \otimes, \Leftrightarrow, \models, \ldots$$

I’m not going to explain what the more complex symbols mean here – this is just to give an idea of some of the symbols that get used. As with road signs, it makes maths look a bit incomprehensible at first, but it makes it easier in the long run.

Google Maps

The difficulty of relating the map to the reality

What is difficult about reading a map? It’s not the actual reading of the map that’s hard, but matching that up with reality in order to put the map to practical use. A map is an abstraction of reality. It depicts certain aspects of reality that are supposed to help you find your way around. The difficulty, in practice, is in translating between the abstraction and the reality. That is, making the link between the map and the place you’re actually wandering around.

Google Maps gives us a brilliant way of moving from the abstract to the concrete, via Google Street View and GPS. Often the hardest part about using a map is working out

a. where you are in the first place, and
b. which way you’re facing.

Those are the crucial pivot points between the map and the reality. GPS has sorted out the business of working out where you are and Google Street View has sorted out the business of which way you’re facing, by giving us a very realistic representation of reality in the form of an actual picture of it.

Maths has to go through these steps as well. First you have to turn the reality into an abstraction. Then you do your logical reasoning in the abstract world. Then finally you have to turn that back into reality again. Different people are good at different parts of this process. But really the key part is being able to move back and forth between the abstract and the real. Still, *someone had to draw the map*.

For example, suppose you have a recipe for an 8-inch square cake, but you want to make it round instead. What size of round cake tin should you use? First you perform an abstraction to turn this ‘real-life’ question into a piece of maths. We want to find a circle whose area is the same as the area of the given square, which is $8^2 = 64$. Now we have to remember that the area of a circle is πr^2, where r is the radius. If we write d for the diameter of the circle (because cake tins are measured by their diameter not their radius), this means we need

$$\pi\left(\frac{d}{2}\right)^2 = 64.$$

Now we actually do the logical reasoning, manipulating the algebra to find out what the diameter d needs to be. This is the only part that’s actually maths.

$$\left(\frac{d}{2}\right)^2 = \frac{64}{\pi}$$

$$\frac{d}{2} = \sqrt{\frac{64}{\pi}}$$

$$d = 2 \times \sqrt{\frac{64}{\pi}}$$

$$\approx \pm 9.027$$

Finally we take the context into account and turn this back into reality. First of all, we don't want the negative answer because we're talking about cake tins here, so the answer needs to be a positive number. Secondly, we don't need all those decimal places – cake tins are usually only measured to the nearest inch. So the answer in reality is that we need a 9-inch round tin for our cake.

The key in maths, and with maps, is to find the most appropriate level of abstraction for the given moment. Do you need little pictures of all the buildings on a street when you're looking at a street map? Do you need to know where there is grass and where there isn't? It depends what you're using the map for, and you'll need different maps for different situations. If you're driving, then you'll want to know which streets are one-way, but that's not very relevant if you're on foot. The same is true of maths. There are different levels of abstraction available for different situations.

What is the number 1? Here are two different ways of answering that question, at different levels of abstraction.

First answer: 1 is the basic building block of counting.

Second answer: 1 is the only number with the property that multiplying by it does nothing.

Each of these answers is useful in different contexts. The first is for when we are most interested in adding numbers up; in mathematics this characterises numbers as something called a 'group' – a world in which we can do addition. The second is for when we are also interested in multiplying; this characterises numbers as something called a 'ring' – a graph. Some of the squares will only be partial squares, so you will only get a truly accurate answer if you use infinitesimally small squares.

Rigorous calculus makes this argument into something logically watertight, but baffles people because it doesn't pin down an answer in the way that people are expecting. Instead it says something like: there's no such thing as graph paper with infinitesimally small squares, so we use progressively smaller and smaller squares and observe that the answer gets closer and closer to $\frac{8}{3}$ as the squares get smaller. Then we prove that no matter how close we wanted it to get to $\frac{8}{3}$, there is a size of square that would get us that close.

A level at which advanced mathematicians sometimes reach their abstract limit is category theory. They react in much the same way that teenagers do when they meet *x*'s and *y*'s – they say they don't see the point, and resist any further abstraction. I am always reminded of Prof. John Baez, who said the following during an argument about abstraction on the worldwide 'Category Theory emailing list':

> *If you do not like abstraction, why are you in mathematics? Perhaps you should be in finance, where all the numbers have dollar signs in front of them.*

I haven't yet met my abstract limit, but I do remember various key moments in my life where I was pushing a boundary and felt I had to make a conscious effort to get over the next bar.

From numbers to pictures

My mother taught me how you can draw a graph of x^2, like this:

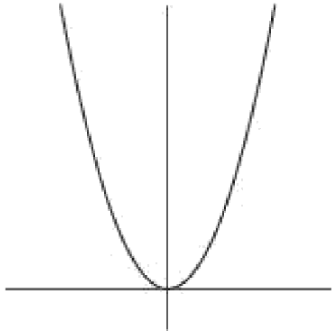

I distinctly remember my bafflement at the fact that you could turn the process of squaring numbers into a *picture* of a curve. I sat in our big green armchair at home thinking and thinking about this until my brain felt like it was popping out of my head. And in my memory this is the exact same feeling I've had every time I've thought about a difficult mathematical concept in my research.

From numbers to letters

I was perfectly comfortable solving equations with *x*'s, say

$$2x + 3 = 7.$$

I knew this would turn into

$$\begin{aligned} 2x &= 7 - 3 \\ &= 4 \\ x &= \frac{4}{2} \\ &= 2. \end{aligned}$$

But then I met one with *a*'s, *b*'s and *c*'s instead of the numbers, something like

$$ax + b = c$$

and I vividly remember feeling completely at a loss as to how on earth to find out what *x* was in this case, without knowing *a*, *b* and *c*. I think I knew that I should start by subtracting *b* from both sides, but I had no idea what that would give on the right-hand side. I do remember that when someone explained to me that it would be *c* – *b* I felt extremely stupid. Why couldn't I have worked that out myself? The answer is then

$$x = \frac{c - b}{a}.$$

Well, as I say to my students – feeling stupid for not having understood something before just shows that you are *now cleverer* than you were then.

From numbers to relationships

This is the last big leap of abstraction I remember having to make, and it was when I was first learning category theory. For the sake of completeness and perhaps amusement value, I'll include here what it was: it was the idea that a *one-object category is exactly a monoid*. Laugh as much as you like; there it is. I sat for days thinking about it and feeling like my brain was popping out of my head, just like when I was a child and thinking about a graph for the first time in my life. And the fact that a one-object category is exactly a monoid is now so obvious to me that I know I am definitely cleverer now than I was then. It's a bit early to explain this example now, but I'll come back to it in the second part of the book.

> We will see that category theory studies relationships between objects. A *category* is a mathematical context for studying these relationships. A *monoid* is a mathematical context for studying something much more concrete: multiplication of things like numbers. The fact that a 'one-object category is a monoid' corresponds to viewing numbers as relationships between the world and itself. This sounds quite strange, but is remarkably powerful.

The goose that laid the golden eggs

Making machines for solving problems

It would be lovely to find a way of making golden eggs. But it would be even better to find a way of making a goose that lays golden eggs: a goose-that-lays-golden-eggs machine. But wouldn't it be even better to make a machine that makes these machines? A 'goose-that-lays-golden-eggs machine' machine. This is a form of abstraction. It's the idea of building a machine to do something, rather than directly doing the thing yourself. So really it's just a form of conservation of energy, or of reserving human brain power for the things machines can't do.

In order to build a machine to do something rather than doing it yourself, you have to understand that thing at a different level. It's like giving someone directions. When you walk somewhere you know well, you don't really think about exactly what streets you're walking on, or which way you're turning and when. You probably go somewhat instinctively. But when you're telling someone else how to get there you have to analyse how you do it more carefully, in order to explain it. You might have noticed that if you ask a local person where a certain street is, they will often not be very sure, as you don't really think about street names when you're wandering around your own town.

Something similar happens when learning a language. When you learn it yourself as your mother tongue, you don't really think about how it works – you pick it up from the adults around you instinctively. Then when you're an adult and a foreigner asks you to explain some aspect of the language that is confusing them, you have to go back and analyse how you speak in a completely different way.

If you're building a machine to make a cake, you'll have to analyse each step rather carefully in order to work out how to get a machine to do it. Even cracking an egg would require careful thought – how do we know how hard to tap the egg against the bowl?

The previous example of solving equations is an example of this type of machine. We start by understanding how to solve equations such as

$$2x + 3 = 7.$$

Then we make a 'machine' for solving all such equations, that is, we solve the equation

$$ax + b = c$$

because then *a, b* and *c* can be any numbers at all.

***

We can then try it for quadratic equations

$$ax^2 + bx + c = 0$$

and we learn that the 'machine' for solving these gives the famous solution

$$x = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}.$$

As a further level of building a machine that makes these machines, there is the *fundamental theorem of algebra* which tells us that every polynomial equation has at least one solution, as long as we allow complex numbers, which we'll come to later.

Cake cutting

An example of abstraction

I remember the first GCSE maths investigation I had to do at school. It was about cutting a cake into as many pieces as possible while making a fixed number of cuts. Obviously if you can only make one cut (in a straight line) you'll only get two pieces of cake, and if you can only make two cuts, you'll get at most four pieces. But what about three cuts? Four cuts? And so on?

The best answer for three cuts is: seven pieces of cake, like this.