

Living Mathematics

Playing With Patterns

Gaurav Shah
gaurav.shah.contact@gmail.com

December 15, 2023

Contents

1	Preface	1
2	Introduction	3
3	The Number that Danced	5
3.1	The journey begins	5
3.2	A fiendishly difficult puzzle	8
3.3	Are you kidding me?	9
3.4	Not so trivial	9
3.5	The rules of the game	10
3.6	Beauty and power in mathematics	11
3.7	How about multiplication?	12
3.8	Whatever shall we do?	13
3.9	Shakespeare as an underrated mathematician	14
3.10	As promised, back to 142857	17
3.11	What if we start with a number other than 3?	18
3.12	Are there more numbers like this?	19
3.13	We come, at long last, to the space robots	20
3.14	A hypothetical scenario of the utmost importance	21
3.15	Back to the outer reaches of the Solar System	22

3.16	Here comes <i>another</i> diversion, but it does get to the point eventually	23
3.17	The next steps	25
3.18	A Perfectly Reasonable Explanation	26
3.19	You probably already know what the answer's going to be	28
3.20	Here's why	28
3.21	Okay, do we give up now?	29
3.22	What went wrong?	29
3.23	Doing it right	29
3.24	The actual proof, as promised	32
4	Space and Time are Child's Play	37
4.1	A Very Good Place to Start	37
4.2	Moving on from numbers	39
4.3	We get a little bit Zen	40
4.4	We don't have to pay Mr. Cayley royalties every time we use his table	40
4.5	Why Transformations form a Group	41
4.6	A Smooth Operator	42
4.7	Other ways to describe this group	44
4.8	The Simplest Trigonometry	45
4.9	Definition of Sin and Cos	45
4.10	Relation between Sin and Cos	46
4.11	Sum of Two Angles	46
4.12	A Mathematical Description of Rotations	47
4.12.1	Step 1: Find R and A	49
4.12.2	Step 2: Find the new R' and A'	49
4.12.3	Step 3: Find x' and y'	49
4.13	Some simple results	50

4.14	How to think like a mathematician	52
4.15	Jeez	53
4.16	Just a Minor Change	53
4.17	A Whole New Set of Function	53
4.18	The New Rotation, mathematically	55
4.19	Case 1: $x^2 > y^2$	55
4.19.1	Step 1: Find R and A	55
4.19.2	Step 2: Find the new R' and A'	56
4.19.3	Step 3: Find x' and y'	56
4.20	Case 2: $y^2 > x^2$	57
4.21	Case 3: $x = y$	57
4.22	What do the new “rotations” do?	58
4.22.1	The simplest possible case	58
4.22.2	A generalization	59
4.22.3	Pairs of points	61
4.23	Is This Real Life, or Just Fantasy?	62
4.23.1	This Entire Section Is Just A Blatant Setup For The Big Reveal That Comes Next	62
4.23.2	As Promised, Here’s The Completely Unexpected Big Reveal	63
4.23.3	I Feel That That Requires Some Elaboration	63
4.23.4	Let’s Back Up A Bit	63
4.23.5	What this means mathematically	64
4.23.6	What Are You Getting At?	65
4.24	How To Find “B”	67
4.25	You’ll never look at the Universe in the same way again	70
4.26	You got this	71
4.27	Get Lost!	71

5	A World of Information	75
5.1	Ice cream	75
5.2	Even more discussion about the ordering of ice cream	77
5.3	Information and Entropy. Also, Ice Cream.	78
5.4	How Do We Know How Much We Know?	79
5.5	Starting to get more challenging interesting	80
5.6	The Strawberry Strikes Back!	82
5.7	Uneven probabilities	82
5.8	Entropy in other situations	84
5.8.1	85
5.8.2	85
5.9	Entropy and atoms	87
5.10	The Boltzmann Equation	95
5.10.1	Uses of the Boltzmann Equation	95
5.11	Tweedledum and Tweedledee	97
5.11.1	I'm glad you asked.	97
5.11.2	A very small difference	98
5.11.3	Flipping coins	98
5.11.4	A cold, cold rain	99

Chapter 1

Preface

This book is intended for two types of people: those who love mathematics, and those who don't. If you belong to either one of these groups, I hope you will read on.

The mathematics that most people encounter (high school, typically) isn't a reflection of what higher math is like. This means that people decide whether or not they like mathematics, without ever knowing what mathematics is like. This book is a journey along the paths of mathematics as it can be. It's aimed at someone with an approximately high school level of knowledge,¹ with no knowledge of calculus, matrices, or complex numbers assumed.

At the same time, I wanted this to be a book *of* mathematics, rather than a book *about* mathematics. What does that mean? You're not going to be reading about concepts and proofs that other people did, you're going to be actually taking the steps yourself. And yet, we're going to actually reach some utterly amazing destinations.

To use mountaineering as a metaphor: this book is a hike to the top, as opposed to a technical rock-face climb. It's easier and doesn't need as much training, but you're still going to take the steps *yourself* to get to the top. It's going to be my responsibility as your guide to ensure that each of these steps is a reasonable ask for an amateur hiker.²

Importantly, you get to enjoy the view at the end that you arrived at by

¹Or you could be a really motivated middle schooler, that's okay too.

²If I don't do this right, I'm pretty sure you'll let me hear about it.

yourself, rather than look at a photo that someone else took.

And *most* importantly, you'll enjoy the journey.

Chapter 2

Introduction

“Something’s a little bit off”

Here’s a very simple question to start us off: what’s 142857×2 ?

I’ll even give you the answer:

$$142857 \times 2 = 285714$$

Did you notice that something about the answer is a little unusual? You probably did: the answer is simply the original number, except that it just happens to have two digits chopped off from the front, and moved to the back.

Hmmm.

It’s not much. It’s just . . . a little bit funny. But this is how adventures begin, mathematical adventures too. You find something that’s mildly curious, like the small sharp glint of metal in the dirt, and then you start digging:

*What if I multiply by numbers other than 2? Does it still work?
Are there any other numbers that behave like this?*

Why does it behave like this?

Are there other places I can apply what I’ve learned?

You examine the glinting metal a little more closely, and it turns out to be a locket. You open the locket, and it contains a clue: you solve the clue and you find a treasure map. The treasure map sends you halfway around the world. Before you know it, you're in a new country, scuba diving in a coral reef, surrounded by bright fish and sharks and coral. You're tired, and you're overwhelmed, and you're very very happy that you noticed that tiny sparkle of metal in the grass.

For example, we're going to dig into this particular little bit of funny business, the number that twirled. And when we're done digging, we'll be talking to robots beyond our solar system.

This kind of unexpected outcome isn't unusual. Later in this book, we'll look at how toddlers play with shapes, and end up with an understanding of Einstein's Theory of Relativity, and a better GPS system. Or we'll try to buy some ice cream, and somehow explain superconductivity.

But don't forget that it's not just the end results that matter, as amazing as they are. We're going to learn that mathematics isn't only useful, it's stunningly beautiful. But it's not just useful and beautiful: most of all, it's fun.

Chapter 3

The Number that Danced

3.1 The journey begins

We noticed in the Introduction that

$$142857 \times 2 = 285714,$$

where the answer is just the original number with the digits rotated. Does that make you wonder if this keeps happening when you multiply it with numbers other than 2? If you do, then congratulations: whether you know it or not, you're thinking like a mathematician.

Let's do the calculations and see:

$$142857 \times 2 = 285714$$

$$142857 \times 3 = 428571$$

$$142857 \times 4 = 571428$$

$$142857 \times 5 = 714285$$

$$142857 \times 6 = 857142$$

$$142857 \times 7 = 999999$$

Aha! We see that yes, the rule does indeed work for numbers other than 2. But then we got a bonus. The last multiplication, the one that *breaks* the

pattern, is a hint: the next clue in the treasure hunt. It tells us that 142857×7 is very nearly 1,000,000 or, alternatively, that 0.142857 is very close to $1/7$. In fact, if we look at the exact calculation, we see that $\frac{1}{7} = 0.14285714285714\dots$ where the digits 142857 keep repeating themselves over and over. There is no way that this could be a coincidence. We have found ourselves our first clue!

Let's take a closer look at what happens when divide 1 by 7. But this time, instead of concentrating on the *result*, we're going to focus on the *remainders* at each step.

$$\begin{array}{r}
 0 . 1 4 2 8 5 7 1 4 2 \dots \\
 \hline
 7 \) 1 . 0 0 0 0 0 0 0 0 0 0 \\
 \quad - 7 \\
 \hline
 \quad \color{green}{\boxed{3}} 0 \\
 \quad \quad - 2 8 \\
 \hline
 \quad \quad \color{green}{\boxed{2}} 0 \\
 \quad \quad \quad - 1 4 \\
 \hline
 \quad \quad \quad \color{green}{\boxed{6}} 0 \\
 \quad \quad \quad \quad - 5 6 \\
 \hline
 \quad \quad \quad \quad \color{green}{\boxed{4}} 0 \\
 \quad \quad \quad \quad \quad - 3 5 \\
 \hline
 \quad \quad \quad \quad \quad \color{green}{\boxed{5}} 0 \\
 \quad \quad \quad \quad \quad \quad - 4 9 \\
 \hline
 \quad \quad \quad \quad \quad \quad \color{green}{\boxed{1}} 0 \\
 \quad \quad \quad \quad \quad \quad \quad - 7 \\
 \hline
 \quad \quad \quad \quad \quad \quad \quad \color{green}{\boxed{3}} 0 \\
 \quad \quad \quad \quad \quad \quad \quad \quad - 2 8 \\
 \hline
 \quad \quad \quad \quad \quad \quad \quad \quad \color{green}{\boxed{2}} 0 \\
 \quad \quad \quad \quad \quad \quad \quad \quad \quad - 1 4 \\
 \hline
 \quad \quad \quad \quad \quad \quad \quad \quad \quad \color{green}{\boxed{6}} \dots
 \end{array}$$

Table 3.1: dividing 1 by 7


See the remainders? There are four simple facts about them that I want to point out that we'll need for later.

1. The remainder at each stage will always be an integer less than 7, because we're dividing by 7.

2. The first six of them are 3, 2, 6, 4, 5 and 1; and then they start repeating, just like the first six digits of the answer keep repeating.
3. If you know the remainder at any stage, you know what the next digit in the decimal answer is going to be. For example, if the remainder at some stage is 5, then you're going to "pull down the zero" to get 50 (which is of course the same as multiplying it by 10). The next digit in the answer is going to be what you get when you divide 50 by 7, which is 7.
4. If you know the remainder at any stage, you know what the *next* remainder is, too. For example, assume that the remainder at some stage is 3. We then "pull down the zero" to get 30 (in other words, multiply it by 10). Then divide 30 by 7, which gives us 4, and a remainder of 2. So if the remainder at one step is 3, the remainder at the next step has to be 2.¹

Like in any good treasure hunt, we're now going to go off in a completely different direction for a bit. Don't forget where we are now, because we'll have to come back here, but for the next step on our treasure hunt we must first solve a fiendishly difficult puzzle.

3.2 A fiendishly difficult puzzle



If it's 10 o'clock now, what time will it be in three hours?

¹If you say a word too many times, then you sometimes find that it seems to lose its meaning. This phenomenon is known as "semantic satiation". *Remainder, remainder, remainder!*

3.3 Are you kidding me?

Okay. It turns out it wasn't *that* difficult a question. If you live in a place with twelve hour clocks, three hours past 10 is 1 o'clock.

In the world of clock times, there are twelve possible hours: 1 through 12 (we're going to ignore the minutes in this discussion). The rules to add or subtract time are simple: if the answer is more than 12, you subtract 12. If it's less than zero, you add 12². The final number is always between 1 and 12. Six hours before 3 o'clock is 9 o'clock.

Mathematicians have a fancy way of describing this: this is called arithmetic modulo 12. And yes, they also call it "clock-number" arithmetic, because mathematicians are human beings too.

If you have a twelve hour clock, then "thirteen o'clock" is the same as one o'clock; the mathematical way of saying this is $13 \equiv 1 \pmod{12}$. This just means that the difference between 13 and 1 is divisible by 12, so the two numbers are equivalent if you're just looking at the time.

3.4 Not so trivial

This seems to be pretty trivial, but don't scoff at it. It's actually the starting point of something more profound.

What we've done here is we've moved on from ordinary numbers. We needed to describe the mathematical operations that we want to do on clock time ("What time will it be in three hours?"), and we created a whole new class of objects to do this. They look like integers, but they don't always act exactly like them. For one thing, there are only twelve of them, from 1 through 12. And addition is slightly different. When the mathematicians look at the numbers modulo 12, however, they make one more tweak: instead of 1 through 12, they use the numbers 0 through 11.

What can we say about this extremely simple system? Let's look at the table for addition, but this time we'll use modulo 6 instead of modulo 12 (table 3.5).

²Have you noticed that a superscript for footnote number can be easily confused with a superscript for raising to a power? An unscrupulous author could easily use this fact to drive his readers crazy.

If you remember, these tables are called Cayley's tables; we'll be using them a *lot* in this book.

+	0	1	2	3	4	5
0	0	1	2	3	4	5
1	1	2	3	4	5	0
2	2	3	4	5	0	1
3	3	4	5	0	1	2
4	4	5	0	1	2	3
5	5	0	1	2	3	4

Table 3.2: Addition modulo 6

3.5 The rules of the game

What can we say about adding modulo 6? Are there any rules that it follows? There are a few special aspects that may be obvious, but I want to point them out anyway, because they're so important.

1. The sum of any two "numbers" mod 6 is another number mod 6. (For example, $5 + 4 = 3$).
2. $0 +$ any number is just that number again. That is, $0 + x = x$.
3. Every row, and every column, in the table has the number 0 somewhere in it, exactly once.
4. $(x + y) + z = x + (y + z)$

When mathematicians look at these facts, they describe them in their own way. Here are the same statements, but translated into mathematical language.

1. The set is closed under addition modulo 6.
2. There is an identity element, e . (e is 0, in this particular case of addition).
3. Each element has an inverse. (That is, for every number x , we can find another number y so that $x + y = e$.)
4. The addition is associative.

With little fuss and no fanfare,³ we've reached someplace important. Any set of objects that obeys these rules is called a "group". (This particular group has a name, the Cyclic Group of order 6.) It's difficult to emphasize enough how powerful the concept of a group is in mathematics, or how central it is to so many different fields. We will be coming across many of them in our travels together and learning more about them at each stage.

3.6 Beauty and power in mathematics

But seriously. Given how important they are, couldn't mathematicians have found a more evocative name for this amazing structure, than the completely generic sounding name "groups"? A little bit of Hollywood razmatazz would probably have been helpful. I'm going to take a few minutes to talk about why groups are so wonderful.

What do we look for in mathematics? What concepts are powerful, or beautiful? There are a few considerations that come to mind. They actually contradict each other a little bit, making it difficult to satisfy all the considerations at the same time, which is actually a *good* thing, because if they didn't then life would be too easy, and what's the fun in that?

Here are two reasons why groups are such an amazing concept.

³Honestly, I've never actually heard a fanfare at a *single one* of the moments that genuinely called for one. I don't think I'm the only one with this complaint.

Groups have a very simple definition. The rules above are quite natural, and there aren't too many of them. This means that as you wander around in the fields of mathematics, many of the structures you find are likely to follow these rules. It's almost spooky how often you bump into them, in completely unexpected places. In this book, try to keep count of how many times we mention a new group.

Groups are powerful. Once you know that something is a group, you learn a lot about it. Groups are a very deep concept. There are many many theorems that describe their behaviour, and in fact, we're still learning more about them.

What these two reasons imply is that when it comes to groups, **we put in a little bit** (the definition of a group) **and get back a lot** (they're very common, and have interesting, unexpected behavior).

This is power in mathematics, this is beauty. Groups have it!

3.7 How about multiplication?

Where there's addition, you often see multiplication, too. Does multiplication modulo 6 form a group? Here in table 3.3 is a Cayley's table for that.

×	0	1	2	3	4	5
0	0	0	0	0	0	0
1	0	1	2	3	4	5
2	0	2	4	0	2	4
3	0	3	0	3	0	3
4	0	4	2	0	4	2
5	0	5	4	3	2	1

Table 3.3: Multiplication modulo 6

How does this work? For example, $4 \times 2 = 8$, and $8 \equiv 2 \pmod{6}$, because $8 - 2 = 6$.

The first law of groups is obeyed. When you multiply two numbers in the set, the answer is another number in the set.

The second law of groups is obeyed, but a little differently than before. The second law says that there has to be an identity; that is, an element “ e ” so that $e \times x = x$ whatever x is.

This is still valid, but the number that fulfils this role is different for multiplication: the identity was 0 for addition, but for multiplication, of course, the identity is 1.

The third law...whoops. The third law says that every row has to have the identity in it. To put it another way, for any x , you should be able to find y so that $x \times y = 1$.

We don’t see that here. The row for multiplying by 0 does not have a 1 in it; neither do the rows for 2, 3, or 4. In other words, 0, 2, 3 and 4 do not have inverses.

We struck out. These numbers do *not* form a group under multiplication.

3.8 Whatever shall we do?

We *could* mope about this, I guess. We could weep, and bemoan our fate. Or we could try to fix it.

We know that $0 \times x = 0$, always. So 0 is *never* going to have an inverse. If 0 doesn’t want to play along nicely with our group, let’s just go ahead and drop it, shall we?

The other numbers in this table that don’t have inverses are 2, 3, and 4. They have something in common: they share a common factor with 6. This makes sense: for example, any multiple of 2 modulo 6 will still be an even number, which means 1 cannot be a multiple of 2 modulo 6.

So in this case, it’s the modulus, 6, that’s the culprit. We can do better: let’s choose to look at numbers modulo 7 instead. As 7 is a prime number, none of the numbers less than it will have any common factor with it.

In table 3.6 we have a Cayley’s table of multiplication modulo 7, and we’ve dropped the 0 as it didn’t want to play nice. That leaves the numbers 1 through 6.

x	1	2	3	4	5	6
1	1	2	3	4	5	6
2	2	4	6	1	3	5
3	3	6	2	5	1	4
4	4	1	5	2	6	3
5	5	3	1	6	4	2
6	6	5	4	3	2	1

Table 3.4: Multiplication modulo 7

And we're off! You can check that this actually follows all the rules that qualify it to be a group.^{4,5}

Also, note that the numbers 0-7 form a group under addition (in the same way that the numbers 0-6 did), and the numbers 1-7 (dropping the "0") form a group under multiplication. There's a special name for this kind of structure, namely, a *field*. We will be looking at fields for the rest of this chapter, but we won't be talking too much about them or using any particularly complicated properties of fields.

3.9 Shakespeare as an underrated mathematician

Do you notice anything unusual about table 3.5 and table 3.6? I'd be surprised if you did. The aspect I'm referring to is completely unexpected, subtle, and also quite fundamental.

Here's the answer: *they're exactly the same table*. How could something as simple as that pass us by? It's because the players in the game changed their names. They're exactly the same table, but you need a translation table to see the correspondence between the two of them.

To make it clearer, I'm going to repeat both of the tables right here, even though we just saw them not that long ago. This is because, let's face it, we're humans, and humans have sadly short attention spans.

⁴Really. You *can* check it.

⁵Have you checked it yet?

+	0	1	2	3	4	5
0	0	1	2	3	4	5
1	1	2	3	4	5	0
2	2	3	4	5	0	1
3	3	4	5	0	1	2
4	4	5	0	1	2	3
5	5	0	1	2	3	4

Table 3.5: Addition modulo 6

×	1	2	3	4	5	6
1	1	2	3	4	5	6
2	2	4	6	1	3	5
3	3	6	2	5	1	4
4	4	1	5	2	6	3
5	5	3	1	6	4	2
6	6	5	4	3	2	1

Table 3.6: Multiplication modulo 7

The first part of the translation is obvious. The first table has addition, the second has multiplication, so $+$ \iff \times .

The next correspondence is easy to guess at, too. The identity in the addition table is 0, while the identity in the multiplication table is 1, so that is our next correspondence: $0 \iff 1$.

The other numbers are trickier, however, as they don't line up either. You have to look carefully to figure out what to do with them. Here's the full translation table, table 3.7.

+	×
0	1
1	3
2	2
3	6
4	4
5	5

Table 3.7: Correspondences between adding mod 6 and multiplying mod 7

For example, $3 + 4 \equiv 1 \pmod{6}$ is something we can read from table 3.5. Reading from the correspondences in table 3.7, that statement becomes $6 \times 4 \equiv 3 \pmod{7}$, which we can check in table 3.6 to see if it's true or not: and it is!

To emphasize this, we're going to combine both of these Cayley's tables into one. The parts in purple are addition, while the parts in blue are multiplication, and each pair in the same cell is connected by the correspondence table 3.7.

+,×	0,1	1,3	2,2	3,6	4,4	5,5
0,1	0,1	1,3	2,2	3,6	4,4	5,5
1,3	1,3	2,2	3,6	4,4	5,5	0,1
2,2	2,2	3,6	4,4	5,5	0,1	1,3
3,6	3,6	4,4	5,5	0,1	1,3	2,2
4,4	4,4	5,5	0,1	1,3	2,2	3,6
5,5	5,5	0,1	1,3	2,2	3,6	4,4

Table 3.8: Addition mod 6, multiplication mod 7

In "Romeo and Juliet" (Act II, Scene II), Shakespeare says:

What's in a name? that which we call a rose
By any other name would smell as sweet.

He's right, and we should probably listen to him more often. When we go from addition modulo 6 to multiplication modulo 7, the names get changed, but the deeper underlying structure doesn't, and that is what is most important. I'm not going to add the proof here, but this is in fact true for every prime number

p : The multiplicative group of numbers from 1 to $p - 1$ (modulo p) is actually the same as the additive group of numbers from 0 to $p - 2$ (modulo $p-1$), no matter how hard it tries to disguise itself and look completely different.⁶

This is one of the ways mathematics makes progress

You solve one problem, you understand it, and then you find it pop up in a hundred other places it was hiding. It turns out that it was just wearing a wig or a false moustache, or using different names. And the results you got from the work you put into solving it in the first place, can now get used for free in these new situations.

In particular, it's obvious that we can start with 0, and keep adding 1, and thereby go through all the numbers modulo 6: 0, 1, 2, 3, 4, 5, 0, 1, 2, But we also know that addition modulo 6 is structurally the same as multiplication modulo 7.

That tells us that we can start with 1, and keep multiplying by 3, and thereby go through all the numbers modulo 7 (except 0, of course): 1, 3, 2, 6, 4, 5, 1, 3, 2, ...⁷.

Any element of the group that creates *all* the rest of the elements of the group in this way is called a generator. (Another term for this is a "primitive element").

So 1 is a generator of the addition group mod 6, and 3 is a generator of the multiplication group mod 7.

You can see that 2 is *not* a generator of the addition group mod 6. If you start from 0 and keep adding 2 (modulo 6), you get: 0, 2, 4, 0, You never get to see the elements 1, 3 or 5.

3.10 As promised, back to 142857

Here's where we were before we dove into group theory: The digits 142857 are found repeating when we divide 1 by 7. In this process of division, the remainders

⁶Here's a technical term for it: "isomorphism", which is a fancy way of saying that the groups are identical other than possibly having different names for the elements.

⁷Don't let me stop you from checking this yourself. Really, be my guest.

at each step keep repeating, too: 3, 2, 6, 4, 5, 1, 3, 2, 6, We can now look at this again in the light of what we've just learned about group theory.

Here's how we go from one remainder to the next in this series, 3, 2, 6, 4, 5, 1, ..., described in two different ways.

Ordinary Division	Group Theory
Start with the remainder at any stage	Start with any number modulo 7
Multiply it by ten ("Bring down the zero")	Multiply it by 3 (because $10 \equiv 3 \pmod{7}$)
Find the remainder when divided by 7	Convert to a number modulo 7

What this means, simply:

if the remainder at one step is a ,
*then the remainder at the next step will be $3 * a$,*

where "*" means "multiplication in the modulo 7 multiplication group".

We've already seen that the number 3 is a generator of the group, so if we keep multiplying by 3, we will cover *all* the members of the group (in some order), and then start repeating the sequence.⁸

3.11 What if we start with a number other than 3?

Recall that the multiplicative group modulo 7 is the same group as the additive group modulo 6. So to ask the same question another way, what if we started at, say, 4, and kept adding 1, modulo 6? That's not a very difficult question. It would go 4, 5, 0, 1, 2, 3, 4, and we've started all over again. In other words, the same numbers, in the same order, just starting at a different value.

And, of course, the behavior will be exactly the same when we keep multiplying by 3: you'd still go over all the elements of the group, one by one in some

⁸Anyone feel like confirming that the sequence 3, 2, 6, 4, 5, 1, ... follows the rule "each element is the previous element multiplied by three"?

order, and then keep repeating the sequence: 6, 4, 5, 1, 3, 2, 6, 4, 5, ... In other words, it's the same numbers in the same order, with a different starting point.

This, of course, is exactly what happens when we divide, say, 2 by 7!

The first remainder is 2 ($2 = 0 \times 7 + 2$), and then we cycle through all the remainders in exactly the same order, just starting with 2 instead of starting with 1.

And if you know what the remainder is, you know what the next digit in the answer is. **So if the remainders are going to keep repeating, the answer is going to be the original answer, but with the digits rotated.** Sound familiar?

We've solved the puzzle of the twirling numbers, which was our initial motivation! But no need to stop here. On the way, we learned enough to keep us going further.

3.12 Are there more numbers like this?

We've seen why 142857 is what we call a cyclic number, so we can look for more of them. There are basically two criteria: 7 is a prime number, and 10 is a generator of the multiplicative group modulo 7 (actually 3 is, in this case, but $10 \equiv 3 \pmod{7}$, so it's the same thing). So all we need to do is find a number other than 7 that fulfills these criteria.

Let's choose another prime number, and see if 10 is a generator in that multiplicative group. How about 11?

$$10^0 \equiv 1 \pmod{11}$$

$$10^1 \equiv 10 \pmod{11}$$

$$10^2 \equiv 1 \pmod{11}$$

That was quick. If we start from 1 and keep multiplying by 10, we get only two different elements of the multiplicative group modulo 11. As mathematicians put it: 10 has order 2 in the group. If 10 is going to be the generator of the group, it has to have order $11 - 1 = 10$ (the "-1" is because it has to generate every number modulo 11 *except* 0).

It turns out 10 is not a generator of the multiplicative group modulo 13, either.⁹ But the order of 10 in the multiplicative group modulo 17 is indeed 16, so it *is* a generator of the group. The digits in $1/17$ that repeat are “0588235294117647” (just look at $1/17$ in a calculator). Since we’re working with 17 rather than 7, it now works for all factors from 1 through 16 instead of merely 1 through 6. Quite frankly, this is awe-inspiring, so much that I’m actually going to take the trouble to list them all, with the ‘0’ in red to make it easier to convince yourself of the truth of this marvel.

$$\begin{aligned}
 0588235294117647 &\times 1 = 0588235294117647 \\
 0588235294117647 &\times 2 = 1176470588235294 \\
 0588235294117647 &\times 3 = 1764705882352941 \\
 0588235294117647 &\times 4 = 2352941176470588 \\
 0588235294117647 &\times 5 = 2941176470588235 \\
 0588235294117647 &\times 6 = 3529411764705882 \\
 0588235294117647 &\times 7 = 4117647058823529 \\
 0588235294117647 &\times 8 = 4705882352941176 \\
 0588235294117647 &\times 9 = 5294117647058823 \\
 0588235294117647 &\times 10 = 5882352941176470 \\
 0588235294117647 &\times 11 = 6470588235294117 \\
 0588235294117647 &\times 12 = 7058823529411764 \\
 0588235294117647 &\times 13 = 7647058823529411 \\
 0588235294117647 &\times 14 = 8235294117647058 \\
 0588235294117647 &\times 15 = 8823529411764705 \\
 0588235294117647 &\times 16 = 9411764705882352 \\
 0588235294117647 &\times 17 = 9999999999999999
 \end{aligned}$$

Do *not* try to pretend you’re not impressed.

3.13 We come, at long last, to the space robots

On September 5th, 1977, Voyager 1 was launched from Earth, on a mission to explore the outer Solar System.¹⁰ On August 25, 1989, it passed the orbit of

⁹Go ahead and verify this yourself. You *know* you want to!

¹⁰Rather strangely, Voyager 2 was launched *before* Voyager 1. It made sense at the time, I guess. **The 70s were a strange, strange decade.**

3.14. A HYPOTHETICAL SCENARIO OF THE UTMOST IMPORTANCE 21

Neptune, *the farthest of our eight planets*.¹¹ It's on a course that will send it by the star system Gliese 445, 17.6 light years from earth, in 40,000 years.

Amazingly, it's still functional, *after more than forty years in vacuum and temperatures near absolute zero*. Right now, it is about 21 billion kilometers from the sun, and in the time it took you to read this sentence it just travelled about *150 kilometers*¹² further. At that distance, the signals from Voyager moving toward us at the speed of light take *20 hours*¹³ to reach our antennas.

The power supply for this space probe came from the heat generated by radioactive isotopes. Right now, the power generated is a little more than 200 Watts, or the power that is needed for two old-style incandescent lightbulbs. For the whole space probe.

Out of the 200 Watts, the power allocated to communications is about 20 Watts, which is a little bit more than most night lights, but not much more. And it uses this power supply to communicate all the way to ground control, on Earth! My car radio gets staticky just because I drive underneath a bridge; can you imagine how much static noise there has to be in the signal all the way from outside the solar system?

So how do they do it? How do they get the correct information from the plucky little probe? Let's talk about that. This is a pretty deep and intricate field of study, but we can go quite a ways into it, and when we do, we'll meet an old friend.

3.14 A hypothetical scenario of the utmost importance

You're at home. Someone from the other room shouts out and asks you if you want pineapple on your pizza. What kind of question is that? You're a right-thinking person – you know that pineapple-flavored pizza is an abomination. *How do you react?* Think fast, buddy.

You could just say "Nah". The problem is, "Nah" sounds a lot like "Yah",

¹¹We can argue about this all day. Pluto is not a planet. You have to *let it go*.

¹²When writing about Voyager, I find myself using *lots of italics*. It's just that cool.

¹³See?

and if the person in the other room thinks that you said this, you will end up with tropical fruit on your cheese and tomatoes. And the chances of being misheard are even worse if there's a lot of noise in the house: maybe from the traffic outside, or perhaps someone's playing some music.¹⁴

Here's one thing you can try: repeat your answer multiple times: "NO NO NO NO NO". The chances are that this will communicate your feelings adequately to whoever's ordering.^{15,16}

3.15 Back to the outer reaches of the Solar System

This gives us one possible way to help Voyager: we could program it to send the data multiple times. For example, if Voyager wanted to send the bit "0", it would instead send the zero repeated five times: **00000**.

Many of those bits would get corrupted on the way to Earth, but hopefully not *most* of them. The information as received on Earth, in this scenario, would look something like this: **01001**.

Since there are two ones and three zeros in this set, we could figure out that the bit that the probe wanted to send was actually a zero, by simple majority.

The problem with this solution is that it's exceptionally inefficient. If the space probe were to repeat every bit of data five (or more) times, then that means the stream of information coming in to us would be five times slower. That's not a good thing.

Our strategy was an emergency measure we came up with on the spur of the moment, designed to save us from getting pineapple on our pizza. Given time and the efforts of some brilliant scientists and mathematicians, there are better options for us when we design the trip to the star Gliese 445. What we're going to do right now is describe a system of error-correcting communication called Reed-Solomon codes, which are actually what NASA chose to use for this purpose.

¹⁴such as "Never Gonna Give You Up" by Rick Astley

¹⁵You'll probably just end up with anchovies instead. That's life.

¹⁶Is it obvious that I'm overcompensating to fit in because of social pressure? The truth is that I actually like pineapple on my pizza. Seriously, what's wrong it?

3.16 Here comes another diversion, but it does get to the point eventually

One of the great books of our times is the classic tome, “How to Avoid Huge Ships” by Capt. John W. Trimmer. It’s exactly what the title says, a useful compendium of various techniques to introduce into your life in order to not get run over by an ocean liner or oil tanker. I’ve scrupulously followed the recommendations found therein, and can truthfully claim that I’ve not even once been swamped by a boat and forced to live the rest of my life on an uninhabited desert island.

If you’re lucky enough to own a copy of this book, open it up. You will find, among the first few pages, the ISBN¹⁷ number of the book: 0881000191¹⁸. That’s a lot of digits!

If you wanted to order the book by providing the ISBN number, it would be very easy to make a mistake. If you did make a mistake, you’d go to your local bookstore where you thought you ordered the book, but instead of pulling out your copy of “How to Avoid Huge Ships”, they’d make a mistake and give you instead, for example, “The Stray Shopping Carts of Eastern North America: A Guide to Field Identification”, by Julian Montague. Now, “The Stray Shopping Carts of Eastern North America: A Guide to Field Identification”, by Julian Montague, is a terrific book, but I mean no disrespect when I say it’s certainly no “How to Avoid Huge Ships”.

You’d be disappointed to see this, wouldn’t you? The local bookstore owner would certainly notice that you were disappointed and then *they* would feel bad too (local bookstore owners are notoriously empathetic).¹⁹

This is a terrible scenario, I know you’ll agree. And it’s to avoid situations like this that the International ISBN Agency²⁰ has given us a solution. All ISBN numbers are constructed in such a way as to be impervious to single digit mistakes. What does this mean?

¹⁷International Standard Book Number

¹⁸For one particular edition. Also, later ISBN numbers were 13 digits rather than 10. We’re going to focus only on 10 digit ISBN numbers for convenience.

¹⁹Yes, they’re real books. They were among the winners of the Bookseller/Diagram Prize For Oddest Title Of The Year, and I think you’ll agree that they were both worthy winners

²⁰No, I’m not kidding, it really exists. I guess the members of the International ISBN Agency get to call themselves International ISBN Agents?

Not all ten digit numbers are allowed to be valid ISBN numbers. The valid ones have been chosen in such a way that that if you make a mistake when writing an ISBN number, by getting *any* one digit wrong, the computer would be able to look at the new number and realize that an error has been made.

For example, if 0881000191 is a legitimate ISBN number, then **6**881000191 is not allowed for any other books, because it differs from 0881000191 in just one digit. 088100**5**191 wouldn't be allowed either. **6**88100**5**191, however, *would* be allowed, since it differs from 0881000191 in *two* places.

That's simple, of course. The trickier part is figuring out which numbers are valid, out of all possible ten digit numbers.

A quick note: we will also be referring to valid numbers as "codes". This doesn't mean a secret message that needs to get deciphered, it's just a way of referring to these strings of digits which are a subset of all possible strings of digits. What we're looking at here, in particular, is called an "error correcting code".

How does this work?

To make a 10 digit ISBN number, start with the first 9 digits, which have the actual information about the book, and can be anything you want. Then add the tenth digit in a special way: so that

$$10*(\text{first digit}) + 9*(\text{second digit}) + 8*(\text{third digit}) \dots + 1*(\text{tenth digit})$$

is a multiple of 11.²¹ This is called a check digit. This sum is a simple example of a what is called a *checksum*.

Let's assume that the first 9 digits are 088100019. Now,

$$10 \times 0 + 9 \times 8 + 8 \times 8 + 7 \times 1 + 6 \times 0 + 5 \times 0 + 4 \times 0 + 3 \times 1 + 2 \times 9 = 164 \equiv 10 \pmod{11}$$

So simply choose the last, tenth, digit to be 1. Then, for 0881000191:

$$10 \times 0 + 9 \times 8 + 8 \times 8 + 7 \times 1 + 6 \times 0 + 5 \times 0 + 4 \times 0 + 3 \times 1 + 2 \times 9 + 1 \times 1 = 165 \equiv 0 \pmod{11}$$

Let's look at a concrete example: maybe one of the numeral "1"s looks like a

²¹Sometimes, none of the digits from 0 through 9 will make this work, and the correct value required is 10. In this case, they use an "X" as the last digit.

“7”, so the ISBN is read as 0881000791. When we enter that mistaken number into the computer, the first thing it will do is to calculate the sum:

$$10 \times 0 + 9 \times 8 + 8 \times 8 + 7 \times 1 + 6 \times 0 + 5 \times 0 + 4 \times 0 + 3 \times 7 + 2 \times 9 + 1 \times 1 = 183 \equiv 7 \pmod{11}$$

When it sees that the sum is equivalent to 7 (mod 11) rather than zero, it will realize that an error has been made, and let us know.

And it's not just this case: if we've made an error in one of the digits, the new sum

$$10*(\text{first digit}) + 9*(\text{second digit}) + 8*(\text{third digit}) \dots + 1*(\text{tenth digit})$$

can *never* be zero. For example, in the case we considered, when we thought the “1” was a “7”, the sum increased by 6×3 . The “6” comes from $7-1$; the “3” is the coefficient. This increase will never be equal to 0 modulo 11, because it's the product of two numbers that are less than 11, and 11 is a prime number.

Since the checksum was 0 for the original number, and the increase is not equal to 0, the new checksum will obviously not be 0.

Note that this works because 11 is a prime, and is related to the fact that we discussed earlier – that multiplication modulo 11 forms a group only because 11 is a prime number.

If phone numbers worked this way, if you made a mistake in one of the digits, you'd *never* get a wrong number, you'd only get a message saying that you dialed a number that doesn't belong to anyone. It's sad that this system wasn't used when phone numbers were first assigned.

In other words, we've increased the length of the message by only one digit – 11% – and we still are able to discover if single digit errors have been introduced into the message. If we had done a crude repetition code, we would have doubled each bit to achieve this same performance – an increase of 100

3.17 The next steps

This is a great first step, and works great for the numbering of books, but it's not yet enough for our space mission. Given how weak the signals are, it can easily happen that we get *two* errors in a message stretch of 10 digits. In that

case, two messages that are different at two digits can be confused with each other.

To be able handle this situation, we need to be able to create lists of messages of length 10 so that any two messages have at least 3 digits different from each other. The good news is that once we figure out how to do this, we can extend it easily – the same method that lets us choose a coding scheme so that any two messages have at least 3 digits different from each other can be extended as far as we like, until all 10 digits are different.²²

To make life easier for us, we’re going to work in base 7. This means that the length of the final word can be only 6 digits, including any check digits. With this change, to find the check sum, the new formula is going to be

$$6 * d_5 + 5 * d_4 + 4 * d_3 + 3 * d_2 + 2 * d_1 + 1 * d_0 \equiv 0 \pmod{7}.$$

And here’s the next step: instead of having the coefficients be 6, 5, 4, 3, 2, 1 in order, we’re going to scramble them. It should be clear that doing this doesn’t change anything critical. The new series is going to be 5, 4, 6, 2, 3, 1. That is, a set of 6 digits is a valid, acceptable code sequence if

$$5 * d_5 + 4 * d_4 + 6 * d_3 + 2 * d_2 + 3 * d_1 + 1 * d_0 \equiv 0 \pmod{7}.$$

This works just as well. But why did we do that? A hint: you’ve seen this sequence earlier in this chapter.

3.18 A Perfectly Reasonable Explanation

You got it!²³ This sequence, 5, 4, 6, 2, 3, 1 is just the reverse of 1, 3, 2, 6, 4, 5, which is the series of consecutive powers of 3, in the multiplication group modulo 7. We saw this when we were discussing the 142857 sequence; in particular, this was the sequence of remainders when dividing 1 by 7.

²²This extreme case, of course, is very easy. We only have two possible messages to choose from: 0000000000 and 1111111111. This is an example of the “NO NO NO NO NO” code we cleverly came up with a while ago.

²³Actually, I have no idea whether you actually got it or not.

So we let the first *five* numbers be arbitrary, but choose the sixth number (d_0) such that

$$3^5 * d_5 + 3^4 * d_4 + 3^3 * d_3 + 3^2 * d_2 + 3^1 * d_1 + 3^0 * d_0 \equiv 0 \pmod{7} \quad (3.1)$$

(In case you didn't know, $3^0 = 1$. In fact, any positive number raised to zero is equal to 1^{24} .)

Just to be clear, so far we have *not* made any progress past what we were already able to do with the ISBN numbers. We've just changed the order of the coefficients, a very minor change that will be helpful later on.

Here's another way to look at equation 3.1. For every sequence $\overline{d_5 d_4 d_3 d_2 d_1 d_0}$, define a polynomial:

$$p(x) \triangleq d_5 x^5 + d_4 x^4 + d_3 x^3 + d_2 x^2 + d_1 x + d_0 \quad (3.2)$$

[By the way, the symbol \triangleq means it's a *definition* of $p(x)$.]

For example, $532103 \longleftrightarrow 5x^5 + 3x^4 + 2x^3 + 1x^2 + 0x + 3$. For every sequence, there's a polynomial; for every polynomial, there's a sequence. We don't even bother to distinguish between them in most conversation.

You can see that the left side of 3.1 is what we get from 3.2 when we set $x = 3$. So what eq. 3.1 is saying is simply that, if $\overline{d_5 d_4 d_3 d_2 d_1 d_0}$ is a valid codeword, and $p(x)$ is its related polynomial, then

$$p(3) \equiv 0 \pmod{7} \quad (3.3)$$

This becomes significant: because for any polynomial $p(x)$, if $p(a) = 0$, then $p(x)$ is divisible by $(x - a)$. So $d_5 x^5 + d_4 x^4 + d_3 x^3 + d_2 x^2 + d_1 x + d_0 = q(x)(x - 3)$.

In other words, the valid codewords are codewords whose polynomials are divisible by $(x - 3)$.

²⁴Have you noticed how easy it is to confuse a superscript that means "footnote" with a superscript that means "raise to a power"? An unscrupulous author could exploit this by scattering these instances throughout a book that he wrote, and driving his readers crazy.

This also makes the working very obvious. If two numbers differ in only one digit, say the third digit from the left, then the difference between them is cx^3 , where c is the difference between the correct and the incorrect versions of the digit. And cx^3 cannot be a multiple of $x - 3$.

But this gives us a way to generalize! To create a set of numbers that have at least *three* digits different between each member of the set, choose another number beside 3 – say, 4. The first four digits can be arbitrary, but now we add *two* check digits. Calculate the fifth and sixth numbers so that the resulting polynomial is a multiple of $(x - 3)(x - 4)$:

$$d_5x^5 + d_4x^4 + d_3x^3 + d_2x^2 + d_1x + d_0 = q(x)(x - 3)(x - 4)$$

Does this work?

3.19 You probably already know what the answer's going to be

No. It doesn't.

3.20 Here's why

Consider some simple examples. I'll do the calculations for you.

If the first four digits are 0000, the fifth and sixth digits are 0 and 0, so the sequence is **000000**.

If the first four digits are 0001, the fifth and sixth digits are 0 and 5, so the sequence is **000105**.

If the first four digits are 0002, the fifth and sixth digits are 0 and 3, so the sequence is **000203**.

These sequences differ in *two* places, not three.

3.21 Okay, do we give up now?

Never.

3.22 What went wrong?

Look at $(x - 3)(x - 4)$:

$$\begin{array}{rcl}
 (x - 3)(x - 4) & = & x^2 - (3 + 4)x + 12 \\
 & = & x^2 - 7x + 12 \\
 & \equiv & x^2 + 5 \pmod{7} \\
 & \longleftrightarrow & 0000105
 \end{array}$$

So $(x - 3)(x - 4)$ itself only has two non-zero coefficients, instead of the three non-zero coefficients that most quadratic polynomials have. In other words, the distance from $(x - 3)(x - 4)$ to 0 is just two digits, instead of three. So we failed coming right out of the gate.

Our first factor was $(x - 3)$, because 3 is a generator of the multiplication group modulo 7. That was a good choice.²⁵

We chose to use $(x - 4)$ as the second factor in the polynomial to generate our code words, but it turns out not all numbers are created equal. In particular, $3 + 4 \equiv 0 \pmod{7}$, so one of the coefficients in $(x - 3)(x - 4)$ is zero. Which means that the codewords are not as far apart from each other as we want.

3.23 Doing it right

Simply fixing this problem is easy: just choose a number other than 4, so that the sum is no longer 0.

But this becomes more difficult as get into larger code words and more error correction. If we want to generate codewords that are 4, 5, 6 or more digits

²⁵Yay us.

apart from each other, how do we choose the right numbers?²⁶ If we multiply $(x - 3)(x - 4)(x - 1)(x - 2)(x - 7)(x - 8)(x - 12)(x - 13)(x - 15)$ (modulo 17), that will be a polynomial of order 9, which has 10 coefficients. How can we be sure that *all* of them are not zero, other than trial and error?

In fact, for the actual Voyager mission, there were 32 check digits, and they were effectively working in base 256. You can imagine this would be very difficult to design for without some further insight.

Luckily,²⁷ there's a very clever trick. Start with any generator (in our example, it was 3), so the first factor is $(x - 3)$. Then, for the next factor, we choose $(x - 3^2)$; then $(x - 3^3)$, $(x - 3^4)$, etc – that is, consecutive powers of the generator.

Why does this work? I'm glad you asked.

We're going to prove that this works. In order to not get lost in the depths of abstraction, we're going to make some assumptions for the values that we'll be using:

- we'll be working modulo 7
- we're going to try to create code words that are 6 digits long
- we will require that any two legitimate code words differ from each other in at least 4 digits.

It's important to realize that none of these values that we've chosen (7, 6 and 4) are critical: our proof is actually a general proof.²⁸

Also remember that we will be working modulo 7 for all of the proof. And we will have no compunctions in switching back and forth between codes as sequences of digits, and codes as polynomials, with no warning whatsoever. Which representation we mean should be clear from context.

Let's begin! Here's some setup first.

²⁶To go further out, of course, we can no longer work modulo 7, because base 7 codewords can only be 6 digits long. That's not a problem; you can work in larger systems. We can actually use any prime, or any power of a prime.

²⁷Luck actually has nothing to do with this. A lot of thought went into this.

²⁸For whatever it's worth: these numbers, 764, also do **not** form the combination that opens a locker at Grand Central Station containing the mythical long-lost lantern of the legendary Cynical Greek philosopher Diogenes. So don't waste your time trying.

1. We start by finding a generator of the multiplicative group modulo 7. We know one of these generators: it's 3. But it's actually easier to not choose a specific value. Instead we'll call it α . You can imagine a 3 everywhere you see α if you want to.
2. Our goal is to create a set of code words, such that any two members of the set differ in at least *four* digits.
3. We will say that allowable code words are only those whose polynomials are multiples of $(x - \alpha)(x - \alpha^2)(x - \alpha^3)$. Call this polynomial $g(x)$; it's often called the "generator polynomial" of the code. (Please don't confuse this with the generator α of the multiplicative group modulo 7.)

In other words, every polynomial of the form $h(x)g(x)$ is a legitimate code word.

4. If $q_1 = h_1(x)g(x)$ and $q_2 = h_2(x)g(x)$ are legitimate code words, then

$$\begin{aligned} q_1 - q_2 &= h_1(x)g(x) - h_2(x)g(x) \\ &= [h_1(x) - h_2(x)]g(x) \\ &= h_3(x)g(x) \end{aligned}$$

where $h_3(x) = h_2(x) - h_1(x)$ is also a polynomial. So if q_1 and q_2 are legitimate code words, then $q_1 - q_2$ is also a legitimate code word, because it's also a polynomial multiplied by the generator polynomial. [And remember, we find $q_1 - q_2$ by doing a subtraction at each digit, modulo 7. For example, $3 - 6 = 4$, because $6 + 4 = 3$, modulo 7.]

5. In other words, if the difference between any two code words has at least 4 non-zero digits, then every code word has at least 4 non-zero digits, and vice versa. Because the difference between any two code words is another code word.
6. Assume that there is a code word q that has *less* than 4 non-zero digits. If you write the code word as a polynomial, that means that the polynomial has less than 4 non-zero coefficients. To make things as specific as possible, assume that polynomial for the code word q is $p_0 + p_1x + p_2x^2 + p_3x^3 + \dots$. And of these coefficients, less than 4 are non-zero. So let's assume, for

the sake of concreteness, that all the coefficients except p_0 , p_2 and p_5 are zero.

Then, viewed as a polynomial, $q = p_0 + p_2x^2 + p_5x^5$.

Spoiler: what we're going to show is that if the other coefficients are zero, then p_0 , p_2 and p_5 are zero too, so q is all zeros. Which means that if $q = q_1 - q_2$, then $q_1 = q_2$. Conversely, if $q_1 \neq q_2$, there have to be at least four digits different between them.

The preliminaries now over, let's start with the actual proof.

3.24 The actual proof, as promised

Step 1: Recap of what we just mentioned

1. α is a generator of the multiplication group modulo 7 (it could be 3, for example, but we choose to leave it as a symbol).
2. The generator polynomial is:

$$g(x) = (x - \alpha)(x - \alpha^2)(x - \alpha^3) \quad (3.4)$$

3. q is a code word, so its polynomial is a multiple of the generator polynomial.

$$q(x) = h(x)g(x) \quad (3.5)$$

(And conversely, any multiple of the generator polynomial is a legal code word.)

4. Every coefficient of the polynomial for q is zero, except for p_0 , p_2 and p_5 .

$$q(x) = p_0 + p_2x^2 + p_5x^5 \quad (3.6)$$

We're going to eventually prove that under these conditions, p_0 , p_2 and p_5 must *also* be zero.

Step 2: $q(x)$ is zero for $\alpha, \alpha^2, \alpha^3$

This is pretty obvious. If you look at eq. 3.4, you can see that $g(x)$ is zero for α, α^2 and α^3 . And eq. 3.5 says that $q(x)$ is a multiple of $g(x)$.

So let's write that out. We're going to use eq. 3.6 to show us what $q(x)$ is, and set that to zero when $x = \alpha, \alpha^2, \alpha^3$.

$$p_0 + p_2(\alpha)^2 + p_5(\alpha)^5 = 0 \quad (3.7)$$

$$p_0 + p_2(\alpha^2)^2 + p_5(\alpha^2)^5 = 0 \quad (3.8)$$

$$p_0 + p_2(\alpha^3)^2 + p_5(\alpha^3)^5 = 0 \quad (3.9)$$

Step 3: Swap the exponents

We know that $(a^b)^c = a^{b \times c} = (a^c)^b$. So, for example, in eq. 3.8, $(\alpha^2)^5 = (\alpha^5)^2$. Let's do that throughout the three equations.

$$p_0 + p_2\alpha^2 + p_5\alpha^5 = 0 \quad (3.10)$$

$$p_0 + p_2(\alpha^2)^2 + p_5(\alpha^5)^2 = 0 \quad (3.11)$$

$$p_0 + p_2(\alpha^2)^3 + p_5(\alpha^5)^3 = 0 \quad (3.12)$$

Let's simplify this set of equations by setting $\beta = \alpha^2$, and $\gamma = \alpha^5$.

$$p_0 + p_2\beta + p_5\gamma = 0 \quad (3.13)$$

$$p_0 + p_2\beta^2 + p_5\gamma^2 = 0 \quad (3.14)$$

$$p_0 + p_2\beta^3 + p_5\gamma^3 = 0 \quad (3.15)$$

Step 4: Take a deviously clever linear combination of these equations

First, define a (slightly mysteriously motivated) new polynomial

$$r(x) \triangleq (x-1)(x-\beta) \quad (3.16)$$

We know that $r(x)$ is a quadratic polynomial (that is, the highest power is 2). So we can write

$$r(x) = a + bx + cx^2 \quad (3.17)$$

and we don't even have to bother to calculate the exact values of a , b and c for now – we just have to remember that this is how a , b and c are defined.

Now take ($a \times$ eq. 3.13) + ($b \times$ eq. 3.14) + ($c \times$ eq. 3.15). And then regroup the terms so terms in p_0 are together, terms in p_2 are together, terms in p_5 are together:

$$a(p_0 + p_2\beta + p_5\gamma) + b(p_0 + p_2\beta^2 + p_5\gamma^2) + c(p_0 + p_2\beta^3 + p_5\gamma^3) = 0 \quad (3.18)$$

$$\Rightarrow p_0(a + b + c) + p_2(a\beta + b\beta^2 + c\beta^3) + p_5(a\gamma + b\gamma^2 + c\gamma^3) = 0 \quad (3.19)$$

$$\Rightarrow p_0(a + b + c) + p_2\beta(a + b\beta + c\beta^2) + p_5\gamma(a + b\gamma + c\gamma^2) = 0 \quad (3.20)$$

Step 5: The part where the magic happens

From eq. 3.16 and eq. 3.17:

$$a + bx + cx^2 = (x - 1)(x - \beta) \quad (3.21)$$

In eq. 3.21 set, firstly, $x = 1$, giving:

$$a + b + c = (1 - 1)(1 - \beta) = 0 \quad (3.22)$$

and now do it again, but set $x = \beta$ instead of $x = 1$, giving:

$$a + b\beta + c\beta^2 = (\beta - 1)(\beta - \beta) = 0 \quad (3.23)$$

(These equations were the motivation for defining $r(x)$ as we did.)

Now look at eq. 3.20, keeping eq. 3.22 and eq. 3.23 in mind.

$$p_0(a + b + c) + p_2\beta(a + b\beta + c\beta^2) + p_5\gamma(a + b\gamma + c\gamma^2) = 0 \quad (3.24)$$

$$\Rightarrow p_0 \times 0 + p_2 \times 0 + p_5 \times r(\gamma) = 0 \quad (3.25)$$

$$\Rightarrow p_5\gamma(\gamma - 1)(\gamma - \beta) = 0 \quad (3.26)$$

It is obvious, and yet profound, that $\gamma \neq 1$, and $\gamma \neq \beta$. The reason is that $\gamma = \alpha^5$, and $\beta = \alpha^2$. Remember that α is a generator of the multiplicative group, so the different powers of α will be different, right until you reach α^6 , which is back to 1. This – THIS! – is why we chose α to be a generator of the group lo all those many steps ago.²⁹

So eq. 3.26 says that p_5 multiplied by three non-zero numbers, gives an answer of zero. But 7 is a prime number, so the set of non-zero numbers under multiplication modulo 7 forms a group: the product of any two non-zero numbers cannot be zero.³⁰

So p_5 has to be zero. But we can use the same method to show that p_2 and p_0 are zero, too!

To show $p_2 = 0$, choose $r(x) = (x - 1)(x - \gamma)$, and find the coefficients a , b and c from that polynomial instead.

And to show that $p_0 = 0$, choose $r(x) = (x - \beta)(x - \gamma)$, and find the coefficients a , b and c from that polynomial instead.

Step 6: Put it all together

To sum up:

- A code word is one for which the polynomial is multiple of the generator polynomial similar to eq. 3.4. Call the order of that generator polynomial n .
- If that polynomial has less than $n + 1$ non-zero coefficients, then *all* of the coefficients are zero. It's impossible for exactly 1, 2, ... n coefficients to be non-zero.
- The difference between any two code words is another code word, because they are all multiples of the generator polynomial.
- So either two code words are different in at least $n + 1$ locations, or they're completely identical in *all* locations.

²⁹I really hope you're impressed.

³⁰Note that this would NOT be true for non-primes. For example, $2 \times 3 \equiv 0$, modulo 6 (remember the multiplicative table modulo 6 we tried earlier?).

- And *that* gives us an efficient way of talking in noisy environments, so that errors can be corrected for.
- And *that* lets us talk to our space probes, the extensions of the human race we've sent to go where we currently cannot, to explore the unknown, to expand our home.³¹
- And what took us here, what started our journey to the outer edges, was that we noticed that a certain number, when multiplied, gives . . . interesting results.

³¹Yes, I'm underselling the applications of error correcting codes. Space probe communication may not be the most useful application, but it's certainly up there for the coolest.


Chapter 4

Space and Time are Child's Play

4.1 A Very Good Place to Start


Many books claim that they will start at the beginning. How many really, actually, definitively, do? *This book does.* We're going back, back before middle school, back before kindergarten. Back to the very first toys you may have played with.

You may have seen this toy: a box with various shaped holes in it, and pieces with the same shapes that could fit into the right holes.





A toy for children ... or cutting edge mathematics?

Here's the triangle from that set. It has a smudge on one of its corners, which is hopefully just some jam or something. We can use this little smudge to keep track of which direction that particular vertex is pointing.



The triangle piece, with a little smudge.

Here's a simple question to start us on our way: in how many ways can this piece be fitted into the slot? We've already seen one orientation that works, here are two more:



Pretty simple: there are three ways of orienting the triangle so that the piece fits through the slot.

4.2 Moving on from numbers

We talked about groups earlier; our groups were Cyclic Groups, formed by addition or multiplication of the numbers modulo some base. These aren't the same as the numbers we're used to ($3 \times 4 = 5$ in the multiplication group modulo 7, not $3 \times 4 = 12$), but they're pretty similar, and clearly derived from numbers.

But groups can be more abstract. One of the most common ways groups are formed is from *transformations*. In particular, consider **the set of transformations of some object that leave some property of object unchanged**.

Did you notice how vague this description is, and how *exceptionally* unhelpful it is in terms of understanding? Let's look at a more concrete example.

We have above a piece for our puzzle, in the shape of a triangle. Right now, our smudge is to the upper right,¹ and it's ready to get through the slot. That's what we're going to take as our property: "being in the right position to pass through the opening". So what kind of transformation will keep that property unchanged?

Not a very difficult question: we can rotate it 120° clockwise (that's one third of a full circle), and it's ready to go through the opening again. This time the smudge is on the upper left, but we never chose the position of the smudge as a property we wanted to keep unchanged, so that's okay. So the first element of our group is the *transformation* – that is, a rotation of 120° clockwise. Don't confuse this with the *orientation* of the piece, that is, which direction the smudge is pointing. Only the change in orientation is relevant.

There's another transformation that keeps unchanged the property of being able to fit through the opening, a rotation of 120° counter-clockwise.

There's another word for a property of an object that remains unchanged under a transformation: that's a symmetry. So the rotations by 0° , 120° and 240° are the only rotations of the triangle that maintain the position of the equilateral triangle as able to fit through the opening. In ordinary conversation,

¹Really hoping the smudge is just jam.

symmetry is usually used to refer specifically to *mirror symmetry*, like the letters A, H, or M². In mathematics, the word is more general and means any property that can remain unchanged after some transformation.

4.3 We get a little bit Zen

Is there anything else that we can do that keeps this property unchanged?³ Yes: nothing.

That is, if the triangle can fit through the opening, then *keeping it in that position without making any changes* will quite obviously keep it in position to go through the opening. This counts as a transformation in much the same way that zero counts as a number. You could ask me, “How many keys do you have on your keychain?” and the answer would be 10. Or you could ask me, “How many porcupines do you own as pets?”, and the answer, truthfully, is zero.

4.4 We don't have to pay Mr. Cayley royalties every time we use his table

So these three actions (do nothing, rotate 120° clockwise, rotate 120° counter-clockwise) are the three things we can do to that will still keep the piece in position to fit through the opening. Let's give them easier names to call them. By tradition, the transformation *do nothing* is called e . Let's give the name a to rotate 120° clockwise, and b to rotate 120° counter-clockwise.

You can see that e , a , and b form a group, because if you rotate clockwise 120° and then rotate clockwise 120° again, you've done the same transformation as if you rotated counter-clockwise by 120°. In other words, $aa = b$. This is sometimes written as $a \circ a = b$, where $a \circ b$ simply means, “do the transformation b , and then the transformation a ”.

Similarly, $b \circ b = a$.

²Random question: what's the longest word you can think of that uses only mirror symmetric letters? For left-right symmetric letters, off the top of my head, I see MAMMOTH. For up-down symmetric letters, I see DECODED. I suspect you guys can do a lot better.

³“this particular property”, of course, refers to being able to fit through the opening.

○	e	a	b
e	e	a	b
a	a	b	e
b	b	e	a

Table 4.1: Rotation group for the triangle

Do you recognize this? You might. Here's an old friend, the addition group modulo 3:

+	0	1	2
0	0	1	2
1	1	2	0
2	2	0	1

Table 4.2: Addition group modulo 3

Yes, they're the same. Remember what we called this correspondence, *isomorphism*. Of course, it's not restricted to triangles. The group of rotations that respect the symmetry of the square is isomorphic to the addition group modulo 3, and etcetera.

4.5 Why Transformations form a Group

As a reminder, here's the definition of a group from before: it's a set along with an operation (such as addition or multiplication) that obeys the following rules⁴:

1. The set is closed under the operation.
2. There is an identity element
3. Each element has an inverse.

⁴This is from the previous chapter. You do *not* have to be embarrassed you don't remember them all.

4. The operation is associative.

Now, let a and b be two particular transformations of an object that respect some symmetry. Then here are some statements you can make:

1. If you apply transform b , then a , then the result is also a transformation that respects that symmetry.

[For example, for the letter "H", if you take the mirror image horizontally, the shape remains the same. And if you take the mirror image vertically, the shape remains the same. So if you take the mirror image vertically and then horizontally, the shape has to remain the same.]

2. You can choose to do nothing at all.


3. Anything you do, you can also undo (that is, going back to the starting position is a legitimate transformation, too).

4. If you apply a lot of transformations, only the order matters, not how you group them together.


These are pretty general statements, and they explain why transformations and groups are so closely related.

4.6 A Smooth Operator


This is a circle:



And this is the same circle after being transformed by a rotation of 120° clockwise:



Here's the circle after being rotated by 1° counter-clockwise:



Would you agree that these images are *seriously* lacking in any sort of drama? The conclusion will not be a shock: if you rotate the circle by any angle whatsoever, it stays the same. This is different from the symmetries that we've looked at earlier in this chapter, which were *discrete*: the groups that we looked at had a finite number of elements in them.

How do you describe this group? We can't make a Cayley table; there are an infinite number of elements in this group.

But instead, we can specify a general set of rules for the group. If the rotation is by A° , then let the element of the group be $R(A)$. That is, $R(A)$ is the transformation, "Rotate by A° clockwise". Then there are two rules that specify the group:

$$R(A_1)R(A_2) = R(A_1 + A_2) \quad (4.1)$$

$$R(360^\circ) = R(0^\circ) \quad (4.2)$$

None of this is surprising!

- The first rule just says that if you rotate clockwise by 20° , and then rotate clockwise by 15° degrees, then it's the equivalent of having rotated clockwise by 35° degrees.

- The second rule just means that if you rotate 360° , it's the equivalent of having not rotated at all.

4.7 Other ways to describe this group

As we just described, we can define the group of rotations by eq. 4.1 and eq. 4.2. But there's another way to define the group, and that other way is perhaps more fundamental, and easier to understand. That other way is in terms of symmetry.

As we *just* recently happened to have mentioned,⁵ the set of transformations that keep some property unchanged forms a group; we also say that these transformations "respect a symmetry". So is there a property that you can think of that remains unchanged during rotations?

Why yes, there is. When you rotate around a point, the distance from the center to any other point doesn't change. (Not just that, but if there are multiple points rotating around a center, the distance between any two points doesn't change either. We're not going to be pointing this out a lot, but it's useful to keep in mind.)

So: the group of rotations, can also be described as *the group of transformations that don't change distances*.⁶

This is nice! Thinking in terms of the symmetry is a very, very useful technique. It's a way of gaining a better understanding of a problem – as you might have realized, it's a simpler definition than eq. 4.1 and eq. 4.2 are.

We're going to care mostly about the distance from the center of the rotation, and there's another way of describing that, using coordinates. If a point starts out with coordinates x and y , then the distance from the center is $\sqrt{x^2 + y^2}$. And if there are two points, (x_1, y_1) and (x_2, y_2) , the distance between them is $\sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2}$.

So: the group of rotations, with a little impreciseness, can also be described as *the group of transformations that don't change $x^2 + y^2$* .

⁵What a coincidence.

⁶Mirror reflections don't change distances either, but we will ignore them going forward. Including or excluding them won't change what we learn significantly.

4.8 The Simplest Trigonometry

Here's a reminder of some *very simple* trigonometry that we will be using later. *Please don't stop reading this book.*⁷ We're just going to state some basic facts, for the benefit of anyone who's not familiar.

4.9 Definition of Sin and Cos

Consider a circle of radius R . Let P be a point on that circle, with coordinates (x, y) , and assume that \overline{OP} makes an angle of A with the positive x -axis.




Figure 4.1: $\sin A = 3/5$; $\cos A = 4/5$

Then here are some definitions for the sine, cosine and tan of the angle A ;

⁷Please?

and a reminder of what Pythagoras's theorem says for the triangle.

$$\cos A = x/R \quad (4.3)$$

$$\sin A = y/R \quad (4.4)$$

$$\tan A = y/x \quad (4.5)$$

$$A = \arctan(y/x) \quad (4.6)$$

$$R^2 = x^2 + y^2 \quad (4.7)$$

4.10 Relation between Sin and Cos

Let's take a further look at that last line, $x^2 + y^2 = R^2$ (Pythagoras's Theorem). That tells you that $x^2/R^2 + y^2/R^2 = 1$, or

$$\sin^2 A + \cos^2 A = 1$$

for all values of A .

Please remember this; we'll be using it, and some interesting variations of it, in the future.

4.11 Sum of Two Angles

I'm just going to give you the formula and not worry about proving this. Here it is:

If you know the sin and cos of two angles, A_1 and A_2 , you can find the sin and cos of their sum, $A_1 + A_2$:

$$\begin{aligned}\sin(A_1 + A_2) &= \sin(A_1) \cos(A_2) + \cos(A_1) \sin(A_2) \\ \cos(A_1 + A_2) &= \cos(A_1) \cos(A_2) - \sin(A_1) \sin(A_2)\end{aligned}$$

Enough preliminaries⁸!

4.12 A Mathematical Description of Rotations

In this section, we're going to answer this question: if we know where a point is (that is, we know its coordinates), and we rotate it around the origin by a certain angle, what are its new coordinates going to be?

⁸I'm glad you're still with us. I mean it.




Figure 4.2:

In other words, we know the coordinates of A , which are (x, y) . We know the angle θ that we're rotating A by. We want to find the coordinates of the new point A' , (x', y') .

Let's find a way to describe these rotations mathematically.

Here's the setup: point P has coordinates (x, y) . Line Segment OP (\overline{OP}) has length R , and makes an angle of A with the positive x-axis.

We rotate it by an angle B , and point P becomes point P' . The coordinates of P' are (x', y') . The question is, *how do you find x' and y' in terms of x, y , and B ?*

4.12.1 Step 1: Find R and A

We already know this!

$$R = \sqrt{x^2 + y^2} \quad (4.8)$$

$$A = \arctan y/x \quad (4.9)$$

$$x = R \cos A \quad (4.10)$$

$$y = R \sin A \quad (4.11)$$

4.12.2 Step 2: Find the new R' and A'

We see that rotations keep $x^2 + y^2$ constant, which means that R is constant. So to find the new point after a rotation, all we do is keep R the same, and change A . Let's change it by a certain amount, say, B . Then the new distance from the origin and angle, R' and A' , are given by

$$R' = R$$

$$A' = A + B$$

4.12.3 Step 3: Find x' and y'

$$\begin{aligned} x' &= R \cos A' \\ &= R \cos(A + B) && \text{Because } A' = A + B \\ &= R(\cos A \cos B - \sin A \sin B) && \text{Expansion of } \cos(A + B) \\ &= (R \cos A) \cos B - (R \sin A) \sin B && \text{Regrouping the terms} \\ \implies x' &= x \cos B - y \sin B && \text{eq. 4.10 and eq. 4.11} \end{aligned}$$

In the same way, if you do the calculations,⁹ you can see that

$$y' = x \sin B + y \cos B$$

⁹Hint, hint.

Wait.

We just reached an important conclusion, but we didn't . . . mark the occasion. Let's go ahead and say the same thing once more, shall we, but we'll *frame* it.¹⁰

$$\begin{aligned}x' &= x \cos B - y \sin B \\y' &= x \sin B + y \cos B\end{aligned}$$

Much better.

Here's the important thing to take from the derivation: we used two important sets of facts in deriving it. Spoiler, we will be coming back to look at this.

1.
 - $x = R \cos A$
 - $y = R \sin A$
2.
 - $\sin(A_1 + A_2) = \sin(A_1) \cos(A_2) + \cos(A_1) \sin(A_2)$
 - $\cos(A_1 + A_2) = \cos(A_1) \cos(A_2) - \sin(A_1) \sin(A_2)$

4.13 Some simple results

We're going to look at a couple of simple situations, and what can happen to the x- and y- coordinates when the observer gets rotated. If this seems way too obvious, please stick around anyway, we'll get to something more interesting very shortly!

Here's something to get used to: rotating a point by an angle of A clockwise, is mathematically the same thing as rotating the axes by an angle of A *counter-clockwise*. Sometimes we'll look at rotations in the first way, sometimes in the second.¹¹ Again, mathematically and conceptually, there's very little difference.

¹⁰This is the mathematics equivalent of popping open a bottle of champagne and spraying it everywhere.

¹¹We will switch from one way to another with very little warning, and even less remorse.

- Here are two points, P and Q , in fig. 4.3. They have the same x -coordinate, which is 3. And p' and q' are the x co-ordinates of P and Q as seen by another observer who has rotated compared to the first one.




Figure 4.3: P and Q have the *same* x -coordinate with respect to one set of coordinate system. But as seen by another observer (red coordinate system), they have *different* x coordinates.

No surprise: after rotation, they don't have the same x -coordinates any more.

- Another situation, in fig. 4.4: A has smaller x -coordinate than B .




Figure 4.4: $x_A < x_B$, but $x'_A > x'_B$

But after rotation, A' has a larger x-coordinate than B' does. Really, no shock here at all.¹²

4.14 How to think like a mathematician

Well, that was easy. Now, what can we do to make it more difficult?

¹²Did you notice that I rotated the axes in the previous diagram, and rotated the *points* themselves in this one? Part of the reason for that is that it made it more convenient for me to draw those particular diagrams. I'm not exactly Michaelangelo.

4.15 Jeez

We've done a good job studying simple rotations in two dimensions.¹³ There's always a next step, though. How can we move beyond transformations in two dimensions, that keep $x^2 + y^2$ constant?¹⁴

There are a couple of possibilities that come to mind. For example, what if we rotated something in three or more dimensions, instead of two? Or, what if we kept $x^2 - y^2$ constant instead of $x^2 + y^2$? As it turns out, both of these possibilities are extremely interesting and take us fun places.¹⁵ In this chapter, we look at transformations that keep $x^2 - y^2$ constant.

4.16 Just a Minor Change

So let's look at transformations that keep $x^2 - y^2$ constant. Effectively, instead of defining the "distance" as $x^2 + y^2$, we're redefining it as $x^2 - y^2$. And again, we're looking for transformations that keep the distances between all points constant, except we're doing it with the *new* definition of distance.

As a preliminary, however, we will need a whole new set of functions.

4.17 A Whole New Set of Function

Similar to \sin and \cos are three new functions, the "hyperbolic sine", the "hyperbolic cosine", and the "hyperbolic tangent", or \sinh , \cosh and \tanh . We're going to define them here, but in case you're not familiar with exponentials, please don't worry. The exact definition isn't as important as some of the properties of the functions, that we'll describe below.

By the way, the \triangleq sign means "is defined as". In other words, we don't ask

¹³You have to admit it was pretty easy.

¹⁴Reminder: we're using this as simple shorthand for "keeps distances between all points constant, i.e. keeps all values such as $(x_A - x_B)^2 + (y_A - y_B)^2$ constant, for all points A and B ".

¹⁵Seriously, what did you expect. It was pretty obvious that we wouldn't have spent that much time studying something as simple as rotations in two dimensions if it wasn't leading somewhere.

why these following equations are true – it just says that the right hand side is the *definition* of what is on the left hand side.

And the e that is used in the following definitions is Euler's number, which is approximately 2.71828... It's an important number in mathematics, comparable to π , but if you're not familiar with it, you don't have to worry about it to follow what comes next.

$$\begin{aligned}\sinh A &\triangleq \frac{e^A - e^{-A}}{2} \\ \cosh A &\triangleq \frac{e^A + e^{-A}}{2} \\ \tanh A &\triangleq \frac{\sinh A}{\cosh A}\end{aligned}$$

If $\sinh A = x$, then we turn it around and also write $A = \operatorname{arcsinh} x$; and similarly for \cosh and \tanh .

What matters most about this definition is that $\cosh^2 A - \sinh^2 A = 1$. You can check it for yourself from the definition.¹⁶ Compare it to what we saw earlier, $\cos^2 A + \sin^2 A = 1$ – it's similar, but has a minus sign rather than a plus sign.

There was just one more characteristic of the \sin and \cos functions that we used in understanding rotations, which is how they behave for the sum of two angles. So let's detail that here for the hyperbolic functions, too.

$$\begin{aligned}\sinh(A_1 + A_2) &= \sinh A_1 \cosh A_2 + \cosh A_1 \sinh A_2 \\ \cosh(A_1 + A_2) &= \cosh A_1 \cosh A_2 + \sinh A_1 \sinh A_2\end{aligned}$$

As before – pretty similar to what we had with \sin and \cos but with a switch in sign.

¹⁶But only if you want to.

4.18 The New Rotation, mathematically

We take the same steps as we did earlier, with just a few minor changes.

As a recap, when we kept $x^2 + y^2$ constant, we chose two new quantities, R and A , so that

$$\begin{aligned} R^2 &= x^2 + y^2 \\ x &= R \cos A \\ y &= R \sin A \end{aligned}$$

This is possible because of the property: $\sin^2 A + \cos^2 A = 1$.

We're going to choose two new quantities now, too. The first is a pretty obvious choice, now we take $R^2 = x^2 - y^2$. Unlike $x^2 + y^2$, $x^2 - y^2$ can sometimes be negative, so just for now, we are going to just look at values of (x, y) where $x^2 > y^2$. We'll come back and look at $y^2 > x^2$, but it won't really make much of a difference.¹⁷ And then we'll deal with the case where $x^2 = y^2$.

4.19 Case 1: $x^2 > y^2$

4.19.1 Step 1: Find R and A

We're going to define the numbers R and A by two new equations. When we did this before, for $x^2 + y^2$, we understood R to be distance, and A to be angle. But for $x^2 - y^2$, there's no such obvious interpretation of what the numbers "mean" – but we can define them by these equations anyway.

$$R^2 \triangleq x^2 - y^2 \tag{4.12}$$

$$A \triangleq \operatorname{arctanh}(y/x) \tag{4.13}$$

¹⁷Spoiler alert: I'm not even going to do it, I'll just tell you you can work it out yourself if you want to. Be prepared. And if you're more comfortable with complex numbers, you don't really have to treat these two cases separately.

With these new quantities defined, we can see what x and y are:

$$x = R \cosh A \quad (4.14)$$

$$y = R \sinh A \quad (4.15)$$

It's actually easy to plug eq. 4.14 and eq. 4.15 back into 4.12 and 4.13 and see that those two equations are satisfied¹⁸.

4.19.2 Step 2: Find the new R' and A'

Our transformation is going to keep $x^2 - y^2$ constant, so it's not going to change R . If it changes A by the amount B ,

$$\begin{aligned} R' &= R \\ A' &= A + B \end{aligned}$$

4.19.3 Step 3: Find x' and y'

$$\begin{aligned} x' &= R \cosh A' \\ &= R \cosh(A + B) && A' = A + B \\ &= R(\cosh A \cosh B + \sinh A \sinh B) && \text{Expansion of } \cosh(A + B) \\ &= (R \cosh A) \cosh B + (R \sinh A) \sinh B && \text{Regrouping the terms} \\ \implies x' &= x \cosh B + y \sinh B && \text{eq. 4.14 and eq. 4.15} \end{aligned}$$

Solving for y' is very similar, and you can do it yourself.¹⁹

$$x' = x \cosh B + y \sinh B \quad (4.16)$$

$$y' = x \sinh B + y \cosh B \quad (4.17)$$

¹⁸Sigh. You already know what I'm going to say: check it yourself!

¹⁹...

Technically, we should also confirm that the distance between *any* two points remains the same with this transformation, not just the distance of one point from the origin. It's not conceptually any different; if you can do one, you can do the other, so we'll skip it.²⁰

4.20 Case 2: $y^2 > x^2$

Just switch x and y around so that $x = R \sinh A$ and $y = R \cosh A$. The steps are the same, and in fact, the result is still exactly the same. You can go through the steps yourself.²¹

4.21 Case 3: $x = y$

This is pretty similar, too. In this case, $R^2 = x^2 - y^2 = 0$, so *any* point (x', x') will also have a distance of 0 from the origin. So choose any other point P (say, $(1,0)$), and set up an equation that says the distance of (x, x) from P has to be the same as the distance of (x', x') from P – in other words, the distance between those two points cannot change after the transformation.

This takes no extra conceptual insight, so I'm not going to go through all the steps. The answer turns out to be the same as the other cases. So in *all three cases*,²²

$$\begin{aligned}x' &= x \cosh B + y \sinh B \\y' &= x \sinh B + y \cosh B\end{aligned}$$

²⁰To be precise, I'll skip it. You're welcome to dive right in.

²¹Or not. Whatever.

²²Here comes the frame.

4.22 What do the new “rotations” do?

We’re going to drop the quotation marks on “rotations”. They look sort of sarcastic, and you can’t keep mentally making quotation marks with your fingers every time you come across them. Just remember that the new distance is defined as $x^2 - y^2$, and in this section, when we say rotation, we mean a transformation that keeps the *new* distance constant.

Doing the mathematics was fun, but we also want to go beyond that and develop a feeling for what these transformations look like. To help us do that, let’s look at a few simple cases.


4.22.1 The simplest possible case

Here’s a point, P , starting off on the x-axis, with coordinates $(1,0)$. After a rotation by an “angle”²³ of B , it moves to point P' . To be specific, let’s choose B to have the suspiciously specific value of 0.6931.²⁴ For this value of B , $\cosh B = 1.25$, and $\sinh B = 0.75$, so $P' = (1.25, 0.75)$ ²⁵.

²³Oops. We did agree to avoid the quotation marks.

²⁴In case you are familiar with logarithms, this is simply the natural logarithm of two, but don’t worry about it if this is new to you.

²⁵Obviously, the reason we chose $B = 0.6931$ is that it produces nice clean values of $\sinh B$ and $\cosh B$.

Figure 4.5: P moves to P'

4.22.2 A generalization

Ordinary rotations keep $x^2 + y^2$ constant, so points move along paths where $x^2 + y^2$ is constant – which are circles.




Figure 4.6: After rotation, a point on any contour will move to another point on the same contour.

Our new rotations keep $x^2 - y^2$ constant, so points move along paths where $x^2 - y^2$ is constant. These are known, in general, as hyperbolas, and they look like this:




Figure 4.7: After the new transformations, a point on a contour will move to another point on the same contour.

Note that the contours look a little different in the different cases: $x^2 > y^2$, $x^2 < y^2$, $x^2 = y^2$ (blue, red, and green, respectively).

4.22.3 Pairs of points

The effects of the new rotations on pairs of points are pretty similar to what we looked at with the old rotations – but the differences are interesting!

We'll just look at a simple example. Here are two points, P and Q . They have the same x -coordinates. And P' and Q' are what we get after P and Q are

rotated by the value of $B = 0.6931$ – or, as seen by another observer who has rotated compared to the first one. Remember, these two situation are the same.




Figure 4.8: P' and Q' no longer have the same x-coordinate after transformation.

Shock, shock: P and Q no longer have the same x-coordinate. The same thing would happen if they had started with the same y-coordinate.

4.23 Is This Real Life, or Just Fantasy?

4.23.1 This Entire Section Is Just A Blatant Setup For The Big Reveal That Comes Next

Wow, wouldn't it be weird if real life worked like this?

4.23.2 As Promised, Here's The Completely Unexpected Big Reveal

It does.

4.23.3 I Feel That That Requires Some Elaboration

These new, weird rotations do in fact describe the world that we live in. But to see how, we're going to have to look at them in a very different, slightly disorienting way: the x and y coordinates are no longer going to describe the x -coordinate and the y -coordinate. Instead, they describe the x -coordinate and *time*.

4.23.4 Let's Back Up A Bit

In 1905, a young patent clerk in Switzerland built on several years of experiments and hypotheses and proposed something radically simple: that the speed of light is a constant, and that anyone who measured it would always come up with the same answer. Albert Einstein turned out to be right; in hindsight, he was probably not achieving his full potential working as a patent clerk.

This proposal, of the constancy of the speed of light, sounds obvious, but it's not. To understand why this is so radical, you can contrast it with the speed of sound. If you are traveling in the direction of a sound wave, it would seem to you to be traveling slower with respect to you. If you travel fast enough, you could, in fact, catch up with the sound – that's called breaking the sound barrier, and is celebrated with a loud noise called a sonic boom.²⁶

If the speed of light, c , is a constant as Einstein proposed, then that could never happen. You could never catch up with a light beam, because it would *always* be moving faster than you, by precisely the speed of light, c . No matter how fast you were moving²⁷.

²⁶...that *really* annoys anyone who happens to live right below you while you're making the sound.

²⁷In fact, one of the consequences of this fact is that the question "how fast were you moving?" becomes *meaningless* – only velocity relative to someone else can be measured.

4.23.5 What this means mathematically

Let's say you're at the origin, and you send out a sudden flash of light forwards along the x-axis. If x is the position of this flash of light at time t , then $x = ct$. That's what it means to say that the speed of light is c ! In other words, $x^2 - c^2t^2 = 0$.²⁸ We are in frame territory here.

$$(\text{The speed of light is } c) \implies (x^2 - c^2t^2 = 0)$$

But this has a very simple consequence: if anybody *else* measures the position x' and time t' , then they have to measure the same value for the speed of light. That is, $x'^2 - c^2t'^2 = 0$, too. This is critical. In fact, I feel another frame coming on.

Any transformation that keeps $x^2 - c^2t^2$ constant, will keep the speed of light constant for all reference frames.

Does that look familiar? It's the same as the "new" rotations²⁹ that we looked at. The only difference is that instead of $x^2 - y^2$, we have $x^2 - c^2t^2$. Since c is a constant, ct is just proportional to t . It's very similar to simply changing the units, from, say, inches to centimeters. The proportionality constant for that conversion is 2.54 – most importantly, it's *always* 2.54. The conversion would be tricky if the ratio of inches to centimeters kept changing in different places, but it's *constant*. It's the same for converting from t to ct .³⁰

²⁸If we had sent it *backwards* along the negative x-axis, then we would have $x = -ct$. This way of writing it includes both cases, and is also easier to generalize when we talk about beams of light along directions other than the x-axis.

²⁹Those quotation marks again.

³⁰The Mars Climate Observer was a spacecraft that crashed into Mars, instead of orbiting the planet, precisely because someone forgot to convert units from feet-pounds-seconds to metric. If there had been Martians watching, it would have seemed to them that we were bombarding their planet with interplanetary cannons. They might have attacked us back with giant space lasers, so please do be careful with units.

4.23.6 What Are You Getting At?

I'm saying something very simple. The Special Theory of Relativity is the same thing as the new rotations we studied, under a different name.

Transformations that keep $x^2 - y^2$ constant	Relativistic transformations of space and time
The x-y plane in two dimensions	Just one space dimension (x) and one time dimension (t)
A point in the x-y plane	An "event" that takes place at one point, that lasts for an instant. Think, for example, of the flash of a firefly.
A rotation changes the coordinates of a point, but keeps $x^2 - y^2$ constant	A new observer who is moving with respect to the original observer will see different values for the location (x-coordinate) and the time (t-coordinate) of an event.
$x' = x \cosh B + y \sinh B$ $y' = x \sinh B + y \cosh B$	$x' = x \cosh B + ct \sinh B \quad (4.18)$ $ct' = x \sinh B + ct \cosh B \quad (4.19)$
Two points can start with the same x coordinates, but different y coordinates. They will have different x coordinates after a rotation.	Two events can have appear to be at the same location (that is, have the same x-coordinate), but at different times, according to one observer. But another observer moving with respect to the first one would see them as being in different locations.
Two points can start with the same y coordinates, but will have different y coordinates after the transformation.	Two events can have appear to be simultaneous (that is, have the same t-coordinate) according to one observer. But another observer moving with respect to the first one would see them as being at different times.
The difference in y coordinates between two points can change after a transformation.	Two observers, moving at different speeds, would get different measurements for the time difference between two events. For example, if a firefly flashes twice, they would have different measurements for the time between flashes.

4.24 How To Find “B”

We claim that these equations describe the real world through Einstein’s Special Theory of Relativity, but they involve the quantity “B”. That’s a little weird, because if you look around you in the real world, you find lengths, velocities, masses, temperatures, but none of these mysterious B s³¹. In the case of ordinary rotations, “B” was simply the angle of rotation, which is something we understand. Now it’s our job now to relate this B parameter to something that we actually observe, too.

It’s actually easy to see. We’re going to go back to the situation we talked about: an observer (“Abeni”) sitting stationary according to her own coordinate system, and another observer (“Bakul”) flying past at velocity v . According to “Bakul”, *he’s* the one who’s stationary, and Abeni is going *backwards* with velocity $-v$ ³².




Figure 4.9: If someone is flying by you at a significant fraction of the speed of light, **do not attempt to high-five them!** Bad idea, Abeni and Bakul!

The notation for x and t in the different coordinate systems is important here, and we need to keep track of it.

1. Abeni’s coordinates for her own position and time are x_a and t_a .

³¹If you were expecting a joke about Hymenoptera, I’m going to disappoint you here.

³²Coordinate systems are just a mathematical expression of narcissism.

That is, x_a is where she thinks she is, and t_a is the time on her wristwatch.³³

2. **Similarly, Bakul's coordinates for his own position and time are x'_b and t'_b .**
3. **Abeni's coordinates for Bakul are x_b and t_b .** That is, when Abeni says that the time is t_b on *her* watch, then she says that Bakul is at position x_b with respect to her.
4. **And Bakul's coordinates for Abeni's position and time are x'_a and t'_a .** That is, when Bakul says that the time is t'_a , then he says that Abeni is at position x'_a .

In other words, the underscript "a" or "b" tells you whose position and time are being recorded (Abeni or Bakul), and the normal letters or the presence of an apostrophe³⁴ tell you who is doing the recording, Abeni or Bakul.

So let's see how this plays out!

1. $x_a = 0$, for all values of t_a . Simple. In her own coordinate system, Abeni is always motionless, at the origin.
2. $x'_b = 0$, for all values of t'_b . The same for Bakul. In his own coordinate system, Bakul feels that he is motionless, and at the origin.
3. $x_b = v \cdot t_b$. Abeni sees Bakul moving past her with speed v along the x axis.
4. $x'_a = -v \cdot t'_a$. Bakul sees Abeni moving past him with speed v along the negative x axis.

None of this is exactly rocket science, although it is quite useful *in* rocket science. But it's useful to take stock of what we've achieved. Equations 4.18 and 4.19 tell us how x_a and t_a are related to x'_a and t'_a , since those two pairs of quantities are ways of measuring the same things in different coordinate systems, and they do it in terms of the rather mysterious quantity, B .

³³Okay, I don't wear a wristwatch either. It's what people used to have in the time period after grandfather clocks, and before mobile phones.

³⁴Please do not write to me and tell me that it's technically not an apostrophe. But if you do, please also include your favorite joke.

And $x_a = 0$ and $x'_a = -v \cdot t'_a$ also give you a way of comparing those same two pairs, but in terms of v , which is something we understand.

So comparing those two relations will tell us how B and v are related. Let us proceed.

1. Write down equations 4.18 and 4.19, for Abeni's position and time, as described by Abeni herself and by Bakul:

$$\begin{aligned}x'_a &= x_a \cosh B + ct_a \sinh B \\ ct'_a &= x_a \sinh B + ct_a \cosh B\end{aligned}$$

2. Substitute $x_a = 0$, because Abeni thinks of herself as stationary:

$$\begin{aligned}x'_a &= ct_a \sinh B \\ ct'_a &= ct_a \cosh B\end{aligned}$$

3. Substitute $x'_a = -v \cdot t'_a$, because Bakul sees Abeni moving backwards with speed v .

$$\begin{aligned}vt'_a &= ct_a \sinh B \\ ct'_a &= ct_a \cosh B\end{aligned}$$

4. Divide these two equations, and that gives us B in terms of v :

$$\frac{v}{c} = \tanh B$$

Technically, this solves the problem we wanted to answer, of what B is – B is simply $\operatorname{arctanh}(v/c)$. If we know v , we could calculate B , and plug that into the equations 4.18 and 4.19. But there's a way to simplify it further, based on the fact that $\cosh^2 A - \sinh^2 A = 1$ that we noticed earlier. (By the way, $\cosh^2 A$ is another way of writing $(\cosh A)^2$ – it might be rather confusing notation.)

$$\begin{aligned}
& \tanh B &= -v/c \\
\implies & \sinh B / \cosh B &= -v/c \\
\implies & \sinh^2 B / \cosh^2 B &= v^2/c^2 \\
\implies & (\cosh^2 B - 1) / \cosh^2 B &= v^2/c^2 \\
\implies & \cosh B &= 1/\sqrt{1 - v^2/c^2} \\
\implies & \sinh B &= (-v/c)/\sqrt{1 - v^2/c^2}
\end{aligned}$$

(The last line comes from $\cosh^2 B - \sinh^2 B = 1$.)

Let's plug these values back into eq 4.18 and eq 4.19! And if anything we've done deserves a frame, these two equations do.

$$x' = \frac{x}{\sqrt{1 - v^2/c^2}} - \frac{vt}{\sqrt{1 - v^2/c^2}} \quad (4.20)$$

$$ct' = \frac{-vx/c}{\sqrt{1 - v^2/c^2}} + \frac{ct}{\sqrt{1 - v^2/c^2}} \quad (4.21)$$

These are the famous Lorentz transformations that describe Einstein's Special Theory of Relativity.

4.25 You'll never look at the Universe in the same way again

There are many other ways of deriving these equations, but in my opinion, thinking in terms of symmetry gives us the best understanding of what's going on.

We derived the fact that $x^2 - c^2 \cdot t^2$ is invariant from the constancy of the speed of light. But thinking in terms of symmetry allows us to invert this process. We can *start* from the fact that $x^2 - c^2 \cdot t^2$ is an invariant that does not change no matter how fast an observer is moving, and derive everything else from that.

By doing so, we realize that there's nothing central about electromagnetic radiation! **The equations of relativity aren't really about light (and its speed). Instead, they are about the invariant properties of space and time.** Which is a much more fundamental, interesting, and fruitful way of thinking about them.

4.26 You got this

But don't forget – if you understand ordinary rotations in two dimensions, then you understand our weird rotations that keep $x^2 - y^2$ constant. And if you understand *those* rotations, you understand special relativity. Points moving that preserve a symmetry – it really is that simple.

4.27 Get Lost!

The Global Positioning System³⁵ (GPS) is a system of 31 satellites that circle the globe at a height of more than 20,000 km over the earth's surface. We don't have the time³⁶ to describe how the GPS system works in detail, but here's a brief overview.

There are always at least four satellites visible in the sky at any position on the globe, and each satellite broadcasts a signal to the GPS receiver. This signal tells the receiver the position of the satellite, and also includes a clock signal that marks the time the signal was transmitted from the satellite.

There are a few different ways that the receiver is able to use this information to estimate the position on the planet and successfully guide the user to the nearest coffee shop. We're going to discuss the simplest method, rather than the most commonly used method, and we'll do some order-of-magnitude calculations³⁷.

³⁵There are multiple GPS systems in the world, but we're going to discuss the Navstar GPS, as that is the one that is by far the most used.

³⁶Or the space, depending on how fast we want to proceed. (That was a joke about space, time and relativity. I shouldn't need to point that out, folks.)

³⁷The other methods may make more sense in practice, but they won't change the conclusions that we arrive at.

1. The receiver has its own internal clock. All of the satellites have their own internal clocks, too, and they continually broadcast a radio message that does nothing but say what time it is on their own clock.³⁸
2. The receiver compares the time on its own clock with the time on the signal from each of the satellites, and finds the difference in time.
3. This difference in time is just the time taken for the signal to get from the satellite to the receiver.
4. The distance to each of the satellites is then just $c\Delta t$, where c is the speed of light (approximately 300,000 km/sec), Δt is the difference in time between the signal and the receiver's internal clock.
5. Since we know the *position* of each satellite,³⁹ and the *distance* from each of the satellites, we know the position of the receiver: it's at the point of intersection of the spheres with known centers and known radiuses.
6. Since we know our position, we can then figure out where the nearest coffee shop is, in order to get there before we collapse.

Do you notice a tricky part in this calculation? It's in the second step. We're comparing the clocks in the satellites with the clocks in the receiver. Let's assume that the clocks in the receiver and the satellite are calibrated and working correctly at the start of the day. This means that at the start of the day, $t = t' = 0$. Let's answer the question: how accurate will the GPS be by the end of the day?

We know that the GPS satellites are at a height of 22,000 km above the earth's surface. Satellites at this height travel at a speed of 4000 m/s. We look at the Lorentz equations and plug in $x = 0$ (because the receiver always stays at the center of her own coordinate system), $t = 86,400$ (because there are 86,400 seconds in a day), and look at the value of t' . Here's what we get:

³⁸And it's still not the most boring radio show out there.

³⁹That was clearly me glossing over an important point; I'm just not going to get into details of how we know the positions, since it's not important for this purpose.

$$\begin{aligned}
 t' &= \frac{t}{\sqrt{1 - v^2/c^2}} \\
 &= \frac{86,400}{\sqrt{1 - (4,000/300,000,000)^2}} \\
 &= 86,399.999993
 \end{aligned}$$

Does that seem a little underwhelming? Instead of 86,400 seconds, we are confronted with a value of 86,399.999993 seconds. Stop the presses?

Well, actually, yes. The difference is 0.000007 seconds in terms of *time*, but how much of a mistake will that make in terms of *position*? The answer is, approximately, $c \times 0.000007$, which is $300,000,000 \text{ m/sec} \times 0.000007 \text{ sec}$. Which is more than 2 km, and if you're off by that much, you're going to take a *really* long time to get your cup of coffee. Especially if the GPS drives you off a cliff first. The *only* way GPS works even remotely well, is by including relativistic analysis into every single calculation, which is pretty cool when you think about it.⁴⁰

Maybe you can use the GPS to guide you to a store to pick up one of those toddlers' toys, while you're at it.

⁴⁰And every day, the uncertainty gets worse by 2 more kilometers. There's also an addendum to this analysis: we're only considering the *Special* Theory of Relativity. The General Theory of Relativity (which deals with the effects of gravitation), actually adds an even larger error every day.

Chapter 5

A World of Information

This is going to be a funny chapter.¹

The crux of the chapter, the essence of what we will learn, is what we *don't* know. What we can *never* know. And ice cream.

5.1 Ice cream

Alice, Bob, Carol and Dan want to order ice cream. Of course, as everyone knows, there are only two flavors of ice cream, chocolate and vanilla.^{2,3} Also, as everyone knows, there's a fifty-fifty chance of any particular person liking chocolate or vanilla.

What do you think the final ice cream order is likely to be? Let's make a table and see, checking all the different choices each person can possibly make.

¹*Funny* as in *strange*, not *funny* as in *a bunch of rather strained jokes that are clearly a cry for help*.

²If they were the kind of people to order avocado-and-liquorice flavored ice cream, they'd be in somebody else's book.

³Okay, I'll grant you strawberry. Let's just pretend strawberry doesn't exist for now, please.

	Alice	Bob	Carol	Dan	the ice cream order
1	chocolate	chocolate	chocolate	chocolate	0 vanilla, 4 chocolate
2	chocolate	chocolate	chocolate	vanilla	1 vanilla, 3 chocolate
3	chocolate	chocolate	vanilla	chocolate	1 vanilla, 3 chocolate
4	chocolate	chocolate	vanilla	vanilla	2 vanilla, 2 chocolate
5	chocolate	vanilla	chocolate	chocolate	1 vanilla, 3 chocolate
6	chocolate	vanilla	chocolate	vanilla	2 vanilla, 2 chocolate
7	chocolate	vanilla	vanilla	chocolate	2 vanilla, 2 chocolate
8	chocolate	vanilla	vanilla	vanilla	3 vanilla, 1 chocolate
9	vanilla	chocolate	chocolate	chocolate	1 vanilla, 3 chocolate
10	vanilla	chocolate	chocolate	vanilla	2 vanilla, 2 chocolate
11	vanilla	chocolate	vanilla	chocolate	2 vanilla, 2 chocolate
12	vanilla	chocolate	vanilla	vanilla	3 vanilla, 1 chocolate
13	vanilla	vanilla	chocolate	chocolate	2 vanilla, 2 chocolate
14	vanilla	vanilla	chocolate	vanilla	3 vanilla, 1 chocolate
15	vanilla	vanilla	vanilla	chocolate	3 vanilla, 1 chocolate
16	vanilla	vanilla	vanilla	vanilla	4 vanilla, 0 chocolate

Table 5.1: The different ways ice cream can be ordered

All of these possibilities sort of make you appreciate the splendor and diversity of humanity, don't they? It's important to realize that since each person is equally likely to order chocolate and vanilla, each row of this table is equally likely to happen.

But the ice cream counter person doesn't need to know who wants which type of ice cream, of course. They only need to know how many of each type of ice cream are needed, in total. That information is found in the last column of the table.

5.2. EVEN MORE DISCUSSION ABOUT THE ORDERING OF ICE CREAM⁷

the order	how many times it occurs in the previous table
0 vanilla, 4 chocolate	1 time
1 vanilla, 3 chocolate	4 times
2 vanilla, 2 chocolate	6 times
3 vanilla, 1 chocolate	4 times
4 vanilla, 0 chocolate	1 time

Table 5.2: How many times each order occurs

That was a lot of work to get to something really obvious: the most likely ice cream order for four people is evenly divided, 2 chocolate and 2 vanilla ice creams. Are we seriously going anywhere with this? Stay tuned!⁴

5.2 Even more discussion about the ordering of ice cream

We're going to re-state that result again, using different language. Let's call each of the different rows in table 5.1 a "micro-state", while the total order we'll call a "macro-state". For example, as an example of a micro-state, we have "Alice orders chocolate; Bob orders vanilla; Carol orders vanilla; Dan orders vanilla".

The person taking the order at the counter doesn't need all this information, of course: all she hears is "Three vanilla, one chocolate". That's an example of what we're calling a macro-state.

So the statement, "Two chocolate and two vanilla occurs the most times in table 5.1" can be re-phrased: "The macro-state that corresponds to the most micro-states is *two chocolate, two vanilla*".

And if each possibility is equally likely – if each person has an equal probability of ordering chocolate or vanilla – then the most likely macro-state is the one that has the most micro-states within it.

⁴Yes. Yes, we are.

5.3 Information and Entropy. Also, Ice Cream.

Let's talk about information now. Here's a rough, non-technical working definition of "information": *something we didn't know before*. Or, in other words, a surprise. Of course, there are many other ways of defining information, but let's use this particular one for this chapter.

If someone comes up to me and tells me that there's a piece of spinach stuck in my teeth, that's information, because I didn't already know they were going to tell me this.

On the other hand, if the neighbour comes up and tells you that the local sports team (the Fighting Earwigs) are going to do well next year, that's not a surprise, because that's what that neighbour says every single year, again and again and again and again. Since we *already know* that he's going to say that, there's zero information content. The Earwigs could be good, they could be bad – since our neighbour always says that they're going to do well, this conversation does not actually give us any information about how good they are.

Some more examples:

- When we study information, one of the most common models is a stream of zeros and ones. For example, in an earlier chapter, we talked about the Voyager spacecraft sending us data back from the outer reaches of the Solar System⁵. Each individual bit that we get, zero or one, is information.
- If we know the macro-state of some system (such as ordering ice-cream), then we have a range of possible micro-states. To know which particular micro-state we actually have from those possibilities, is a piece of information.

And if we know only the macro-state and not which micro-state is chosen (like knowing the ice cream order, but not knowing which person gets which flavor), that is a *lack* of information that we can still talk about.

When we quantify our lack of information, such as by knowing the macro-state but not the exact micro-state, then **the amount of information that's hidden in the macro-state is called the entropy of that state**. We will

⁵For the moment, let's ignore the noise that corrupts the data, and imagine that the ones and zeros are received just as they're sent.

be using that term later. It's still information, but it's a way of describing how much information we're missing rather than how much information we have.

Please be aware: entropy is a concept that is used widely in many different contexts, and a lot of people spend a lot of time arguing about the exact definition. It seems to make them happy for some reason, like ice cream or a hot-air balloon ride. That's a game we're not going to play: we'll use the definition we just saw.

5.4 How Do We Know How Much We Know?

If we want to really study the theory of information, then our first requirement will be to *quantify* information. Only once we've done this, can we dive in further.

What should we use to measure information? The answer might seem a little circular: by using other information. But if you think about it, that's how we measure most things. Length is measured in terms of a standard length (one centimeter, for example); weight in terms of a standard weight (such as one kilogram). So information can be measured in units of a standard amount of information.

Here's one unit that we can use, the simplest one: the amount of information, or surprise, in randomly choosing a zero or a one. This is called one "bit" of information.

We see single bits of information all over the place. If we toss a coin, it's either heads or tails; and these outcomes can (quite obviously) be put in a one-to-one correspondence with a zero and a one. So the random toss of a coin choosing heads or tails, has the same amount of information as randomly choosing a zero or a one: one bit.

Some more examples of single bits of information:

- In Morse code, every character is a dash or a dot
- Every location in a computer memory is a zero or a one
- A light switch can be on or off
- When we get a signal from Voyager 1, each character is a zero or a one.
- An ice cream order can be chocolate or vanilla

It's been really easy so far. It will get more challenging soon.^{6,7}

5.5 Starting to get more challenging interesting

So one bit of information allows us to distinguish between two equally likely states. What about two bits?

Obviously, two bits allow you four combinations: 00, 01, 10, 11. Three bits allow you eight, four bits allow you sixteen. And it doesn't take much to generalize that: n bits of information allow us to specify between 2^n micro-states.

Now that we've established that, going the other direction is pretty clear. 2^n micro-states contain n bits of information. More generally, if we have k possible micro-states, then knowing which one of them is chosen gives us $\log_2 k$ bits of information. In the same way, if we have a signal coming in that can be in one of k different states, it provides $\log_2 k$ bits of information.

Here's where it starts getting powerful: we can use this formula even when k is not a power of 2. For example, when we roll a die, there are 6 equally likely possibilities. Using this formula tells us that this is $\log_2 6$ bits of information, or approximately 2.58 bits. So one roll of a die is the equivalent of 2.58 coin tosses.

And of course, this works for entropy, too. If the order (aka the macro-state) is "one chocolate, three vanilla", then table 5.4 tells us that there are four micro-states that are possible. So the entropy of this macro-state is $\log_2 4$, or 2 bits.

And if the order is 2 chocolate, 2 vanilla: we already saw that there are 6 micro-states that correspond to that order. So the entropy of the "two chocolate, two vanilla" ice cream order is 2.58 bits.

The fact that the most likely order is "two chocolate, two vanilla", then, can be re-phrased: that order has the highest entropy of all orders. These are two different ways of saying the same thing.

I don't know if any of these steps seemed particularly hard, but this equation is a Really Big Deal. We celebrate that in the traditional way, by putting it inside a nice little frame:

⁶For the record, that's meant as a promise, not a threat.

⁷I meant *interesting*! I meant to say that it will get more *interesting* soon.

$$S = \log \Omega$$

where:

S is the entropy,

Ω is the number of micro-states.

Did I mention this equation is a Really Big Deal? Ludwig Boltzmann, one of the great physicists of the Nineteenth century, was so proud of it that he had it carved on his gravestone.⁸



⁸In his version of the equation, k is called Boltzmann's constant. It's a constant, so multiplying by it is simply a change of units, like measuring in meters instead of centimeters, or furlongs instead of nautical miles. Don't worry about it, it doesn't change anything important.

5.6 The Strawberry Strikes Back!

For this formula $S = \log \Omega$ to be true, there's an implicit assumption, that all possibilities are equally likely. For example, when we calculated that the entropy of one roll of a die was 2.58 bits, we assumed that it was a *fair* die, that all options from 1 through 6 were equally likely. For that matter, even a toss of the coin doesn't deliver one bit of information if heads and tails aren't equally likely.

The case where all possibilities are equally likely is technically called a *uniform distribution*. But *non-uniform* distributions are found all over the place.⁹

As an example of an "unfair", non-uniform distribution, consider asking people that you meet on the street how many years old they are.¹⁰ The answers you get are numbers between 0 and 120, but they're certainly not equally likely for all numbers between 0 and 120! You're far more likely to meet someone who's 20 than someone who's 120.

When that is the case, the entropy can't just be the log of the number of micro-states, because each micro-state is no longer equally likely. So the number of micro-states is no longer an indicator of how likely a given macro-state is.

5.7 Uneven probabilities

In this section, we are going to find the information content in a distribution that doesn't have equal probabilities. Also, we will order more ice cream.

Alice, Bob, Carol and Dave have a friend named Elizabeth who likes ice cream too. But in addition to chocolate and vanilla, she also likes strawberry ice cream. She can't decide which one she likes best, in fact, so when asked what she wants, she chooses randomly from all three, with equal probability (so a probability of $1/3$ for each of them). Here's a question: how much information content is packed into her choice?

This is a question we know how to answer: it's $\log_2 3$ bits, or approximately 1.58 bits. We are now going to extract those same 1.58 bits in another way, in two steps instead of one.

⁹Because the world isn't fair.

¹⁰Do not actually do this.

- For the first step, her friend Alice asks her, “Hey, Elizabeth, we’re ordering ice cream. Is the flavor you want either chocolate or vanilla?” Elizabeth decides on a flavor, randomly, between chocolate, vanilla, and strawberry. If the flavor is chocolate or vanilla (which is $2/3$ of the time), she answers “yes”; if it’s strawberry ($1/3$ the time), she answers “no”.

We don’t know the information content of this step. It’s a yes-or-no question, but not a 50-50 one, a $2/3$ to $1/3$ probability question.

- For the second step, **if** Elizabeth answered “yes” to the first question, then Alice asks her a second question: “Okay, then, which one, chocolate or vanilla?”. This question distinguishes between two equally likely alternatives, so it provides one bit of information, as we’ve seen before. But it only gets asked $2/3$ of time; the other $1/3$ of the time, zero bits of information are provided. So on the average, the amount of information transferred in this second step is $(2/3) \times 1 = 2/3$ bits of information.

What we want to find is the amount of information transferred in the *first* step. So we simply state the obvious: the information transferred by answering “chocolate, vanilla or strawberry?” directly ($\log_2 3$ bits) has to be equal to the information transferred in the two step process – because it’s the same information, after all.

If the information stored in “Is it either chocolate or vanilla?” is S , then:

$$\begin{aligned}
 S + 2/3 &= \log_2 3 \\
 \implies S + 2/3 \log_2 2 &= (2/3) \log_2 3 + (1/3) \log_2 3 && \log_x x = 1 \text{ for any } x \\
 \implies S &= (2/3)(\log_2 3 - \log_2 2) + (1/3) \log_2 3 && \text{subtracting } 2/3 \log_2 2 \text{ from both sides} \\
 \implies S &= (2/3) \log_2(3/2) + (1/3) \log_2 3 && \log b - \log a = \log b/a \\
 \implies S &= -(2/3) \log_2(2/3) - (1/3) \log_2(1/3) && \log a = -\log 1/a
 \end{aligned}$$

This can be generalized very easily, to more than just a simple case of two choices with $2/3$ and $1/3$ probability. I’m not going to go through all the steps, but all the key ideas are already present in the discussion we just went through.

If the signal consists of one of n different symbols ($n = 3$ for chocolate, vanilla, strawberry; $n = 2$ for zero or one; $n = 6$ for rolling a die; $n = 120$ or so if you ask people their age), and the probability of these symbols are p_1, p_2, \dots, p_n , then the information content (or, alternatively, the entropy) of the system is

$$S = -p_1 \log p_1 - p_2 \log p_2 - \dots - p_n \log p_n$$

This formula is *huge*. It has applications all through physics, digital signal processing, and information theory. We're going to look at some of these applications now, but before we do that, let's try to understand it a bit better. We'll make things easier on ourselves by looking at the case of $n = 2$. $n = 2$ means that every symbol transmitted is either 0 or 1 and nothing else, so $p_1 + p_2 = 1$. We'll call the probability of the symbol being 0 p ; then the probability that the symbol is 1 is $1 - p$.

Then, $S = -p \log p - (1 - p) \log(1 - p)$.

Let's look at some special cases.

- In particular, $S = 0$ if $p = 0$ or $p = 1$. If you think about it, this makes sense. $p = 0$ means that every single symbol received has to be a 1; $p = 1$ implies that every single symbol received has to be a 0. Since you can predict exactly what the next symbol is, there is no information transfer happening, so $S = 0$.
- If it's an equally likely distribution with n choices, then $p_1 = p_2 = p_3 \dots = p_n = 1/n$.

Then $S = -n \times 1/n \log(1/n) = \log n$ (since $\log(1/n) = -\log n$).

So we get back our formula that we already saw earlier. Not a surprise, but it's nice to know that our work is consistent.

5.8 Entropy in other situations

As usual, I'm going to present some problems that appear to have nothing to do with real life, just so I can later pull out a whole bunch of real-world applications, and make a big deal about it. We're going to start off with a simple question, then work step by step.

5.8.1

The first question is something we've already solved before:

If we have 6 numbers, each of which is 1, 2 or 3, what's the most likely distribution of 1s, 2s and 3s?

We looked at this when studying ice cream order distributions. The most likely distribution is two 1s, two 2s, and two 3s, the even distribution. Just as the most likely ice cream order for four people was two chocolate, two vanilla.

5.8.2

That was a little bit too easy, so here's the next level of complexity.

Just as before, let's say we have 6 numbers, each of which was either 1, 2, or 3. *But now we also require that the sum of the 6 numbers must equal 8.* So what's the most likely distribution of 1s, 2s and 3s?

It's always easier if we can look at all the alternatives. So here's a table. It's a list of all the different ways that six numbers, all of which are between 1 and 3, can add up to 8. Keep an eye on the last three columns, which tell you how many 1s, 2s and 3s are in each sum.

	sum	number of 1s	number of 2s	number of 3s
1	$1 + 1 + 1 + 1 + 1 + 3 = 8$	5	0	1
2	$1 + 1 + 1 + 1 + 2 + 2 = 8$	4	2	0
3	$1 + 1 + 1 + 1 + 3 + 1 = 8$	5	0	1
4	$1 + 1 + 1 + 2 + 1 + 2 = 8$	4	2	0
5	$1 + 1 + 1 + 2 + 2 + 1 = 8$	4	2	0
6	$1 + 1 + 1 + 3 + 1 + 1 = 8$	5	0	1
7	$1 + 1 + 2 + 1 + 1 + 2 = 8$	4	2	0
8	$1 + 1 + 2 + 1 + 2 + 1 = 8$	4	2	0
9	$1 + 1 + 2 + 2 + 1 + 1 = 8$	4	2	0
10	$1 + 1 + 3 + 1 + 1 + 1 = 8$	5	0	1
11	$1 + 2 + 1 + 1 + 1 + 2 = 8$	4	2	0
12	$1 + 2 + 1 + 1 + 2 + 1 = 8$	4	2	0
13	$1 + 2 + 1 + 2 + 1 + 1 = 8$	4	2	0
14	$1 + 2 + 2 + 1 + 1 + 1 = 8$	4	2	0
15	$1 + 3 + 1 + 1 + 1 + 1 = 8$	5	0	1
16	$2 + 1 + 1 + 1 + 1 + 2 = 8$	4	2	0
17	$2 + 1 + 1 + 1 + 2 + 1 = 8$	4	2	0
18	$2 + 1 + 1 + 2 + 1 + 1 = 8$	4	2	0
19	$2 + 1 + 2 + 1 + 1 + 1 = 8$	4	2	0
20	$2 + 2 + 1 + 1 + 1 + 1 = 8$	4	2	0
21	$3 + 1 + 1 + 1 + 1 + 1 = 8$	5	0	1

Table 5.3: Six numbers (between 1 and 3) that add up to 8

number of 1s	number of 2s	number of 3s	how many times
5	0	1	6
4	2	0	15

Table 5.4: How many times each combination of 1s, 2s and 3s occurs

That last table tells you the important thing: the most likely distribution is four 1s and two 2s (some arrangement of $1 + 1 + 1 + 1 + 2 + 2 = 8$). Instead

of the most common distribution being an even arrangement of 1s, 2s, and 3s, the most common distribution has more smaller numbers than larger ones.

Of course, if we fix the sum of the six numbers to be 8, then that's the same thing as saying as we're fixing the *average* of the six numbers to be $8/6$, or 1.33. And when we do this, in the most likely arrangement, 66.6% of the numbers (4 out of 6) are 1s, 33.3% (2 out of 6) are 2s, and 0% are 3s.

5.9 Entropy and atoms

That was pretty easy, and not particularly painful. So ask yourself this: are we the type of people to be satisfied to leave it at that?¹¹

So let's talk about gas.

To be specific, let's talk about the Helium atoms inside a Helium balloon. There are many, many atoms inside the balloon,¹² and each atom is moving around with a certain value for its speed, and hence its energy. We cannot know where each individual atom of helium is, how fast it's moving, and whether it prefers chocolate or vanilla ice cream. All of these are random and impossible to measure in any practical sense.

What we *can* actually measure are quantities such as the total mass and the temperature. So specifying a given mass of helium, and the temperature it's at, will be our macro-state.¹³ This is like the ice cream counter worker, who just knows how many chocolate and how many vanilla ice creams are ordered, but does not know which individual gets which flavor.

Knowing the mass and the temperature actually gives us some insight:

- The total mass is the number of particles, times the mass of each particle. So from the mass, we can find the number of particles.
- The temperature is proportional to the average kinetic energy of each particle, so we know the average kinetic energy.

¹¹No?

¹²Seriously. Like a bajillion.

¹³This is the actual original usage of macro and micro states. These concepts weren't really developed to describe the ordering of ice cream, you know.

- From the number of particles, and the average kinetic energy of each particle, we can find the total kinetic energy of all the particles put together.

Here is our goal for this section: given a macro-state (for example, a known mass of helium in the balloon, at a known temperature), what is the most likely distribution of kinetic energies for the particles?¹⁴

If you think about it, we've been building up to this for a while now. All the other problems have been pretty similar to this one – that is, we've tried to answer the question “What's the most likely distribution?” for a bunch of different scenarios.

- We looked at ordering ice cream. We discovered that in this case, the most likely distribution is the “flat” one. That is, chocolate and vanilla are equally likely; if there were more flavors, it would be equally distributed among all of those, too.
- We looked at the scenario in which we chose exactly 6 numbers (out of 1, 2 and 3), such that their sum was fixed to be 8. For this scenario, again, we found the most likely distribution, by counting the micro-states for each distribution. We discovered that the most likely distribution was a little less evenly distributed. In particular, in the most likely distribution, smaller numbers were more common than larger ones. In this case, the most likely distribution had the number “1” occurring 4 times, the number “2” occurring 2 times, and the number “3” occurring zero times.

It turns out that the second scenario is almost exactly the same as the current one! In the helium atom case, too, the total number of constituents is fixed (it was 6 in the second scenario). And their sum is fixed – because the total energy tells you the temperature.

However, there are a few differences between the problem we already solved, and the current one.

1. Instead of six numbers adding up to 8, there are about 1.5×10^{23} (per gram of Helium) different values for kinetic energy that add up to a fixed amount

¹⁴If we know the kinetic energies, we also know the speeds, so that is another equivalent way to look at the problem.

– which, for all practical purposes, is pretty much an infinite number of values.

2. The possible values themselves were limited to the discrete numbers 1, 2 and 3 in the example we looked at, but can be a continuous range of values for the kinetic energy of the atoms.

Neither of these is a major stumbling block, but they make it impossible to simply enumerate all the possible micro-states, see which macro-state each one belongs to, and thereby see which macro-state has the most micro-states. So instead we're going to have to *think* about it in order to solve it.¹⁵

We're going to make one simplification: instead of letting the energies take a continuous range of possible values, we will assume they take a discrete (but infinite) set of possible values: $E_0 < E_1 < E_2 < E_3 < \dots$. This allows us to focus on the reasoning instead of getting caught up in technical issues that are not relevant. The number of particles is going to be fixed, but we will assume that it is a very large number.¹⁶

I'm going to say this ahead of time: the problem we are about to solve is a pretty challenging one, and it's very impressive that we can actually solve it with the machinery that we've just built, after simply fiddling around with ice cream orders for a few pages!

For a fixed number of particles, and a fixed total energy, how many particles are expected to be in each energy level E_0, E_1, \dots ?

1. We'll start off with **the really cool equation that we recently discovered**: $S = -p_0 \log p_0 - p_1 \log p_1 - p_2 \log p_2 - \dots - p_k \log p_k - \dots$
 - As a reminder, p_0 is the fraction of the total number of particles that are in energy level E_0 , and similarly for p_1, p_2, \dots
 - Also as a reminder, the entropy ("S" on the left hand side of the equation) is a measure of the number of micro-states available in this distribution of particles among energy states.

¹⁵Which is a good thing, in case you were wondering.

¹⁶This is an example of a "micro-canonical ensemble". Don't worry too much about what that means. The other kinds of ensembles are the "canonical ensemble" and the "grand canonical ensemble". **Who can deny that "Grand Canonical Ensemble" would be a fantastic name for a band?**

- Obviously, this means that the most likely state is the one that has a higher entropy than any other state.
2. **The two constraints** are that the total number of particles is fixed to be N , and the total energy of all the particles is fixed to be E .

- **Total number of particles:** Let the number of particles in energy level E_0 be n_0 , and similarly for n_1, n_2, \dots

$$n_0 + n_1 + n_2 + \dots = N$$

We are going to assume N is a very large number, much greater than 1. Another way to write that is $N \gg 1$. In fact, we assume that the number of particles in each level is much greater than 1: $n_k \gg 1$ for each k .

Also, p_0 , is the ratio of the number of particles in energy level E_0 to the total number of particles (N):

$$p_k = \frac{n_k}{N}$$

These two equations give us a simple relationship that is no surprise:

$$p_0 + p_1 + p_2 + \dots = 1$$

- **Total Energy:** A particle in energy level E_k has energy E_k .¹⁷ If there are n_k particles in that energy level, they will together have energy $n_k \cdot E_k$. So the total energy (which we know is fixed to be E) is given by:

$$E = n_0 \cdot E_0 + n_1 \cdot E_1 + \dots + n_k \cdot E_k + \dots$$

3. The k^{th} term in the formula for entropy is $-p_k \log p_k$. **How does this change if we change the number of particles in the energy level by a small amount?**

- Here is a fairly elementary result from calculus, which can be shown in other ways, too. We're not going to prove it, we'll just take it as given. The formula is this:

If x is a fraction very close to zero ($x \ll 1$), then

$$\log(1 + x) \approx x$$

¹⁷Insert sarcasm emoji.

- We know that $p_k = \frac{n_k}{N}$, so

$$-p_k \log p_k = -\frac{n_k}{N} \log \frac{n_k}{N}.$$
- We want to see what happens to $-p_k \log p_k$ when n_k changes by a small amount.

Say the change in n_k is d_k ; that is, the new number of particles in the energy level is $n'_k = n_k + d_k$. Remember that the change is small, so d_k is much less than n_k ($d_k \ll n_k$).

The fraction of particles in energy level k is now changed from $p_k = \frac{n_k}{N}$ to

$$p'_k = \frac{n'_k}{N} = \frac{n_k + d_k}{N}$$

- The change in the quantity $-p_k \log p_k$ is the new value minus the old value. I want to emphasize that this calculation, as ugly as it looks, is just a little rearrangement of terms – you should feel free to skip to the end, since it's details rather than concepts.¹⁸

$$\begin{aligned} & (-p'_k \log p'_k) - (-p_k \log p_k) \\ &= \left(-\frac{n'_k}{N} \log \frac{n'_k}{N}\right) - \left(-\frac{n_k}{N} \log \frac{n_k}{N}\right) \\ &= \left(-\frac{n_k + d_k}{N} \log \frac{n_k + d_k}{N}\right) - \left(-\frac{n_k}{N} \log \frac{n_k}{N}\right) \\ &= \left(-\frac{n_k + d_k}{N} \log\left[\left(\frac{n_k}{N}\right)\left(1 + \frac{d_k}{n_k}\right)\right]\right) - \left(-\frac{n_k}{N} \log \frac{n_k}{N}\right) \\ &= \left(-\frac{n_k + d_k}{N} \left[\log\left(\frac{n_k}{N}\right) + \log\left(1 + \frac{d_k}{n_k}\right)\right]\right) - \left(-\frac{n_k}{N} \log \frac{n_k}{N}\right) \\ &\approx \left(-\frac{n_k + d_k}{N} \left[\log\left(\frac{n_k}{N}\right) + \frac{d_k}{n_k}\right]\right) - \left(-\frac{n_k}{N} \log \frac{n_k}{N}\right) \\ &= -\frac{d_k}{N} \left[\log \frac{n_k}{N} + 1 - \frac{d_k}{n_k}\right] \\ &\approx -\frac{d_k}{N} \left(\log \frac{n_k}{N} + 1\right) \\ &= -\frac{d_k}{N} (\log p_k + 1) \end{aligned}$$

Explanation:

$$p_k = \frac{n_k}{N}$$

$$n'_k = n_k + d_k$$

$$n_k + d_k = n_k \left(1 + \frac{d_k}{n_k}\right)$$

$$\log(ab) = \log a + \log b$$

$$\log(1+x) \approx x \text{ if } x \ll 1$$

multiply the terms in the brackets

$$d_k \ll n_k$$

$$p_k = n_k/N$$

That was easier than it looks!¹⁹ There were plenty of steps, but each step is quite simple. Look at what changed from each line to the

¹⁸Minor technicality: we assume that the total count N doesn't change, even though we're changing n_k . This will be explained later, please don't worry about it now.

¹⁹This is your cue to chime in and agree with me.

next, and use the hints if you need to.

The important part is that last line. And the most important part of the last line is that the change in the term is directly proportional to d_k (as long as d_k is small, of course). So if you add d_k particles to the k^{th} energy level, the entropy will go up by $-\frac{d_k}{N}(\log p_k + 1)$; and if you *subtract* d_k particles, then the entropy will go *down* by the same amount. If you double the change in the number of particles, you double the change in the entropy – as long as the approximation is still correct, that the change is small compared to the number of particles present.

4. **What does it mean for the entropy to be a maximum?** Assume that the configuration is such that the entropy is at a maximum. Now, shuffle a *small* number of particles around in such a way that after the shuffling, the constraints (total number of particles, and total energy of the system) are still satisfied. Here's my question: what happens to the entropy?

It's easy to say that the entropy has to go down; after all, we started with the entropy at the maximum! But if that were true, then we could *reverse* the process of shuffling – instead of, say, adding 2 particles to energy level E_k , we could subtract 2 particles from that energy level instead; and the same for every other energy level too.

We already saw that the change in entropy is linear. So if it went down in one set of transformations of all the p_k , then it would have to go up if we reversed the transformations – which contradicts our assumption that we started off at a maximum.

The only way out of this is that for *small* changes in each p_k that satisfy the constraints, **the entropy doesn't change.**

This is a big deal!

As a reminder, this is true only for small changes, for which the change in entropy is linearly proportional to the change in the number of particles.

5. **The payoff.**

- Choose three energy levels. For simplicity, we will choose one of them to be the first energy level E_0 ; the other two are arbitrary and will be written as E_i and E_j , where i and j are any two integers.

- Also for simplicity, we can choose to say that the energy of E_0 is zero. This doesn't affect any calculations, since only *differences* in energy between any two levels are relevant. If we didn't make this assumption, it would not make any changes in our result, but it would make the calculations a little bit more complicated.
- Now shuffle a small number of particles around between those levels, in a very specific way, so that our constraints on total number of particles and total energy are satisfied both before and after shuffling. The shuffling that we do is to add d_0 particles to energy level E_0 , d_i particles to energy level E_i , and d_j particles to energy level E_j . Obviously, not all of d_0 , d_i , and d_j can be positive, but that's not a problem. All the other energy levels are unchanged.²⁰

- Then the increase in the total number of particles has to be zero, so $d_0 + d_i + d_j = 0$

And the increase in the total energy has to be zero, so

$$(d_0 \cdot 0) + (d_i \cdot E_i) + (d_j \cdot E_j) = 0$$

We will assume that d_0 is a known quantity. So solving these two equations for d_i and d_j in terms of d_0 gives:

$$d_i = d_0 \cdot \frac{-E_j}{E_j - E_i}, \text{ and}$$

$$d_j = d_0 \cdot \frac{E_i}{E_j - E_i}$$

- If the entropy is already at a maximum for all configurations satisfying the constraints, we just talked about how making small changes that satisfy the constraints cannot change the entropy. We also just showed that if the number of particles in the k^{th} energy level changes by d_k , then the change in entropy has to be $-\frac{d_k}{N}(\log p_k + 1)$.

So in our case, the total change in entropy, which is the sum of the changes in entropy from each level, has to be zero.

$$\begin{aligned} & -\frac{d_0}{N}(\log p_0 + 1) - \frac{d_i}{N}(\log p_i + 1) - \frac{d_j}{N}(\log p_j + 1) = 0 \\ \implies & -\frac{d_0}{N}(\log p_0 + 1) - \frac{d_0 \cdot \frac{-E_j}{E_j - E_i}}{N}(\log p_i + 1) - \frac{d_0 \cdot \frac{E_i}{E_j - E_i}}{N}(\log p_j + 1) = 0 \end{aligned}$$

²⁰This is the justification I promised earlier for changing the number of particles in the k -th energy level, but not changing the total number of particles, N .

Some very simple rearrangement of terms gives us:

$$\begin{aligned} \implies \log p_0 + \frac{-E_j}{E_j - E_i} \cdot \log p_i + \frac{E_i}{E_j - E_i} \cdot \log p_j &= 0 \\ \implies (E_j - E_i) \cdot \log p_0 + (-E_j) \cdot \log p_i + E_i \cdot \log p_j &= 0 \quad \text{Multiply all terms by } (E_j - E_i) \\ \implies E_j \cdot (\log p_0 - \log p_i) - E_i \cdot (\log p_0 - \log p_j) &= 0 \quad \text{Rearrangement} \\ \implies \frac{1}{E_i} \cdot \log \frac{p_i}{p_0} = \frac{1}{E_j} \cdot \log \frac{p_j}{p_0} &\quad \log a - \log b = \log(a/b) \end{aligned}$$

Wait. That last line!²¹ Do you notice something funny about it?²²

The left hand side of the equation depends on the index i , and not j ; the right hand side of the equation depends on the index j , and not i . And i and j , if you remember, are *any* energy levels, not especially chosen ones.

That means that the quantity $\frac{1}{E_k} \cdot \log \frac{p_k}{p_0}$ is the same value for *any* value of k . In other words, it doesn't depend on k at all – it's a constant.

Let's call the value of the constant $-C$. The minus sign is a little uncalled for now, but will make more sense in a few minutes.

$$\begin{aligned} \frac{1}{E_k} \cdot \log \frac{p_k}{p_0} &= -C \\ \implies \frac{p_k}{p_0} &= p_0 \exp(-C \cdot E_k) \end{aligned}$$

Sorry, excuse me, *this* is what I meant to say:

$$p_k = p_0 \exp(-C \cdot E_k)$$

Because it's just that big of a deal. This is the famous Boltzmann distribution, and we find it turning up all over the place. We'll discuss it a bit further in a minute.

That was a rather long derivation, but the reason that it was so long was that we explained every small step in detail. I want to run through the same derivation again, but without all the details, so you don't get lost in the details.

²¹If I feel that a sentence deserves an exclamation mark, I will darn well go ahead and add an exclamation mark.

²²"Do you notice something funny about it?" is a recurring theme in this book, I'm beginning to realize.

The goal is this: *What distribution of particles, with a fixed total number and a fixed total energy, is most likely?*

Here's the recap of what we just did:

1. The most likely distribution is the one that has the highest entropy.
2. Find a way to shuffle a small fraction of the particles among three energy levels to keep the total number of particles and the total energy unchanged.
3. If the entropy is at the maximum, then this shuffling should also not change the entropy. Solve the equation that says "the entropy before this shuffling is the same as the entropy after the shuffling". That gives you the Boltzmann distribution.

A few comments about the parameter C : We chose a minus sign in front of it, so that the number of particles in each energy level would decrease, rather than increase. If it had a positive sign, there would need to be either an infinite number of particles, or a finite number of energy levels.

Also, C can be written as $1/kT$, where T is the temperature, and k is Boltzmann's constant. We can calculate this by finding the average energy of each particle, and relating it to the temperature. It's pretty easy but it's more trouble than is necessary for an introduction.

5.10 The Boltzmann Equation

5.10.1 Uses of the Boltzmann Equation

This set of equations is called the Boltzmann distribution, and has consequences *everywhere*.

- The most obvious is what we've already talked about: tell us how many atoms in a gas have a kinetic energy within a given range, at a given temperature. (With a little extra work, this also tells us how many have a given range of speed.)

- This can further help determine most of the thermodynamic properties of a substance, such as specific heat (how much heat is required to increase the temperature by a certain amount).
- In the same way, the Boltzmann distribution can also tell us as a first approximation how the atmosphere thins out as we climb above the surface – why the air is so thin over Kilimanjaro, for example. In this case, the energy is proportional to the height h , by the formula $E = mgh$.
- Chemical reaction rates are determined in part by how often two molecules collide with enough energy to react, which is determined by the Boltzmann distribution. This works also for calculating nuclear fusion rates, in stars and in trying to create controlled fusion for energy generation.
- The Boltzmann distribution is also used in machine learning. Let's say you have a neural net that takes as input, say, photographs of cats. We've somehow trained the neural net to try to distinguish between cute cats and cats that aren't cute. When we feed in a cat photo, however, it gives us as output two numbers: say "5" for cuteness, "2" for non-cuteness. How do we go from these rather cryptic numbers to a rating of whether this cat photo is likely to go viral or not? The answer that many people use, called the softmax method, is an attempt to find the most likely set of probabilities that is associated with these two numbers; and another word for *most likely* is *highest entropy*. So we can apply a measure of

$$\frac{e^5}{e^5 + e^2}$$

as a cuteness probability and

$$\frac{e^2}{e^5 + e^2}$$

as the non-cuteness probability. This gives us something easy to understand: cuteness probability now goes from 0 to 1, with 1 being a guaranteed meme monster.²³

²³For three days, after which you have to find something else to stare at.

5.11 Tweedledum and Tweedledee

You know what, getting to the Boltzmann equation is pretty impressive. That's not a bad location to end up at, maybe put up your feet. And yet . . . there's a lot more to play with.

What if there was a system exactly like the gas with the Maxwell-Boltzmann distribution, except for a difference in how we calculated the entropy? How would the behavior change? That would be a great showcase for how something simple like counting states changes the entire behavior of an actual physical system.

5.11.1 I'm glad you asked.

In the early 1920s, a young Indian physics lecturer in Dhaka (now in Bangladesh), made an embarrassing mistake.²⁴ He was trying to explain to his classroom of students why physics was *unable* to correctly predict, say, the glow emitted by a red hot piece of charcoal. Unfortunately, he made a simple mistake, right there in front of all his students . . . and he ended up with the *right* answer, the one no one else had been able to obtain, instead of the wrong one.²⁵

He was smart enough to realize that he was actually on to something. Here's the mistake he made: he assumed that photons, the particles that make up light, were indistinguishable. Not just hard to distinguish from each other, not just similar to each other. Two electrons, for example, always have the same mass, the same charge, and the same spin. There no way, even in principle, to take a permanent marker and put a red smudge on an electron.

So let's say you swapped a photon on Mars with one on Earth. According to this way of thinking, you haven't changed the universe in any way. The two states are actually just one state.

This scientist's name was Satyendranath Bose. His "mistake" is now called Bose-Einstein statistics, and the particles he described are known as bosons.²⁶

²⁴I haven't been able to track down a primary source for this story, but the story itself seems to be well known. Please see, for instance, <https://www.isical.ac.in/econophys/bose.html>.

²⁵Awkward.

²⁶There is another type of particle: the fermion (electrons, for example, are fermions). Fermions, like bosons, are indistinguishable particles. But they are different from bosons in one very important way: no two fermions are allowed to be in exactly the same state.

5.11.2 A very small difference

Bosons are very similar to distinguishable particles, in terms of interactions. The only difference is that because they're indistinguishable, the number of different ways they can be split up changes. In other words, their entropy.

This small difference leads to very strange behavior, as we'll examine.

5.11.3 Flipping coins

Let's say you flip two distinguishable coins. Each coin, of course, can be heads (H) or tails (T). There are four possibilities, or micro-states, which we can describe as HH, HT, TH, TT. Each of these possibilities has a probability of 1/4th of happening.

But if the coins are indistinguishable, there are only *three* possibilities: HH, HT, TT – because HT and TH are now the same. And each of these has probability of 1/3 happening.

This gets more extreme when there are more coins. If there are, say, 10 indistinguishable coins, then there are only 11 very simple possibilities, each of which has probability 1/11:

- HHHHHHHHHH
- HHHHHHHHHT
- HHHHHHHHTT
- HHHHHHHTTT
- HHHHHHTTTT
- HHHHHTTTTT
- HHHHTTTTTT
- HHHTTTTTTT

²⁷It's also useful to realize that every journal he tried to submit his work to rejected the paper, one after the other. Bose then wrote to Einstein and asked for help getting the paper published, and what do you know, it turns out that having a famous physicist give a recommendation improved the paper immensely.

- HHTTTTTTTTT
- HTTTTTTTTTT
- TTTTTTTTTTT

So all ten heads is just as likely as five heads, five tails. However, if these were normal, distinguishable coins instead, then having five heads and five tails would be 252 times more likely than ten heads.

To put it another way, for indistinguishable particles, the “all heads” state has the *exact same* entropy as the evenly distributed state. The phenomenon we saw earlier, where ordering equal numbers of chocolate and vanilla ice creams was the most likely case because that had the highest entropy, is no longer true. Bosons are much more likely to hang out together in the same state compared to distinguishable particles.

And if we are dealing with approximately 10^{26} particles or so, instead of just ten, the difference in behavior is even more extreme.

This gets . . . weird.

5.11.4 A cold, cold rain

As we just said²⁸ the new Bose-Einstein statistics can lead to some very weird behaviors. We’re going to look into them in this section.²⁹ The particular phenomenon we’ll look at is called Bose-Einstein condensation, and it is responsible for the bizarre behaviors of many superconductors, and also for the even more bizarre behavior of liquid Helium.

Obviously, we can’t go into a full explanation of this phenomenon, because it’s too complicated and challenging.³⁰ Instead, we’re going to play with a toy model that exhibits very similar behavior – a simplified version of the actual system. The objective is to get an understanding of the phenomenon, not be able to grasp all of the details.³¹

²⁸Literally, the previous sentence.

²⁹The method we use is taken from the Wikipedia article on the Bose-Einstein condensate.

³⁰Not going to lie: I’d be thrilled if this insult enrages you to look up other explanations and dive right into them.

³¹I count approximately 18 physicists who have received Nobel Prizes relating to work in superconductivity, superfluidity, and Bose-Einstein condensation of atoms. So I think it’s okay to leave a *little* bit out of this account of the phenomenon and not quite explain *all* of it.

Imagine a system of N particles, which can be in either of two energy levels, E_0 or E_1 . The energy in E_0 is 0, the energy in E_1 is Δ . The temperature of the system is T .

Our task now is to calculate what fraction of particles are in state E_0 , and what fraction are in state E_1 . This is similar to what we did earlier in finding the Boltzmann distribution, except that these particles are Bosons and are indistinguishable from each other.

In particular: we cannot find the answer by finding the macro-state with the highest entropy, as we did in calculating the Boltzmann distribution case. This is because we discovered that for Bosons, each macro-state has only one micro-state, so they all have the same entropy.

We are going to get around this by loosening one of the assumptions we made earlier. Instead of assuming that the total energy is a constant, we will assume that the temperature is a constant. At a fixed temperature, you can still have a distribution of possible energies.

Instead of the energy of each particle, we will look at the energy of the system. Since the particles are indistinguishable, there are only $N + 1$ possibilities: the number of particles in the higher energy state can be a number, k , between 0 and N .

Then the number of particles in the lower energy state will be $N - k$; and the energy of this configuration is $k \cdot \Delta$. Each of these configurations has the same entropy, because swapping any two particles doesn't lead to a new state – that's what it means to say that the particles are indistinguishable.

Here's the critical part: The *particles* are indistinguishable, but the *configurations* have different numbers of particles in each level, so the configurations themselves are distinguishable from each other.

So the Boltzmann distribution statistics is valid, but not for the distribution of particles, but for the distribution of configurations. That is, the probability of having k particles in the higher energy level will be

$$p(k) = C \exp(-k\Delta/T)$$

For convenience, define $p = \exp(-\Delta/T)$. Then,

$$p(k) = C \exp(-k\Delta/T) = C \exp(-\Delta/T)^k = Cp^k.$$

Obviously, the sum of the probabilities of being in each of the configurations

must be 1. That is,

$$p(0) + p(1) + p(2) + \dots = Cp^0 + Cp^1 + Cp^2 + \dots Cp^N = 1$$

Let's make the assumption that $N \gg 1$, so we can approximately consider the series to be an infinite series.

$$\begin{aligned} Cp^0 + Cp^1 + Cp^2 + \dots &= 1 \\ C \cdot (p^0 + p^1 + p^2 + \dots) &= 1 \end{aligned}$$

$p^0 + p^1 + p^2 + \dots$ is known as a geometric series. The sum of the series is well known to be $\frac{1}{1-p}$, assuming $0 < p < 1$. That gives us:

$$\begin{aligned} C \cdot \frac{1}{1-p} &= 1 \\ C &= (1-p) \end{aligned}$$

Here's the question we will be trying to answer: **what is the average number of particles in the higher energy level?**

This is how we calculate that average number:

1. For a given configuration, find the probability of being in that configuration
2. Multiply that by the number of particles in the higher energy level for that configuration.
3. Then add up for all configurations.

[For example: if you were flipping a coin with a "1" on one side and a "2" on the other side, and a 25% chance of the "1" coming up. What is the average of the numbers that come up, if you flipped the coin many times? It's $(1/4 \times 1) + (3/4 \times 2) = 7/4$.]

So, the average number of particles in the higher energy level over all configurations:

$$\begin{aligned}
&= (C \cdot p^0 \times 0) + (C \cdot p^1 \times 1) + (C \cdot p^2 \times 2) + \dots \\
&= C \cdot (1 \cdot p + 2 \cdot p^2 + 3 \cdot p^3 + \dots) \\
&= (1 - p) \cdot (1 \cdot p + 2 \cdot p^2 + 3 \cdot p^3 + \dots)
\end{aligned}$$

The series, $(1 \cdot p + 2 \cdot p^2 + 3 \cdot p^3 + \dots)$ is another well known series, called the “arithmetico-geometric series”.³² It’s not difficult to find the sum³³: it’s $\frac{p}{(1 - p)^2}$.

Putting it all together, this means that the average number of particles in the higher energy level, over all configurations is $\frac{p}{1 - p}$. Do you notice something funny about this?TM

This number of particles does *not* depend on N ! That means that as you increase the total number of particles, all of them except $\frac{p}{1 - p}$ of them will be in the lower energy level, on the average. As N grows, close to 100% of the particles will be in the lower energy level, even though the temperature is not zero.

That is, even though the temperature gives the particles enough energy to be in the higher energy level, they congregate in the lower energy level, because of the weird statistics that happens when the particles are indistinguishable.³⁴

This is different from what we see in “normal” particles, where the Boltzmann distribution tells us that at any non-zero temperature, there’s a non-zero probability of any particular particle being in the higher energy level.

This phenomenon – in real life, not in this extremely simplified model – is known as Bose-Einstein condensation. It’s responsible for, among other things, superconductivity in some materials. To put it very crudely: electrical resistance occurs when the atoms of the material slows down the progress of an electron that is conducting a current. Critically, the atoms of the material may be able to slow down a single electron, but they can’t slow down a zombie horde of electrons that have paired up, and are all in the same state and completely indistinguishable from each other.

³²Okay, honestly, not nearly as well-known as the geometric series.

³³or to look up the sum

³⁴It’s like the zombie apocalypse of physics.

And I want to say that it's insanely amazing that we could get here, just starting from counting different ways to order ice cream. Get an ice cream for yourself, please, any flavor you want, because you certainly deserve it.³⁵

³⁵And thank you for reading the footnotes too.