101 Illustrated Real Analysis Bedtime Stories

DRAFT August 14, 2016 First three chapters only



Contents

Preface

Ι	Tł	ne basics	1							
1	Sets, functions, numbers, and infinities									
	1	Paradoxes of the smallest infinity	3							
	2	Uncountability	13							
	3	Cantor's infinite paradise of infinities	18							
	4	The real deal	22							
	5	Building real numbers (Dedekind cuts)	28							
	6	The Cantor set	29							
2	2 Discontinuity									
	7	Guessing function values	33							
	8	The Dirichlet function	39							
	9	Conway's base-13 function	42							
	10	Continuity is uncommon	46							
	11	Thomae's function	47							
	12	Discontinuities of monotone functions	48							
	13	Discontinuities of indicator functions	52							
	14	Sets of discontinuities	53							
	15	The Baire category theorem	57							
3	Series									
	16	Stacking books	61							
	17	Inserting parentheses and rearranging series	68							
	18	A Taylor series that converges to the wrong function	73							
	19	Misshapen series	75							
	20	If you torture a series enough, it will converge	77							
4	Acknowledgments									

i

	inted Details	01
1	Adding books to the top of a stack	87
2	Analytic continuation	88

4

Preface

There's no ulterior practical purpose here. I'm just *playing*. That's what math is – wondering, playing, amusing yourself with your imagination.

Paul Lockhart [2]

In the genre of *entertainment through mathematics education*, content is scarce. Hopefully this book helps to remedy that. In each of 101 sections, we present a fun idea from real analysis. The word "story" in the title is meant loosely. Don't expect too many once-upon-a-times. (See Figure 1.)

This is *not* a real analysis textbook. We de-emphasize math's role as a *tool* in favor of its role as an *art form* and more generally as a *mode of expression*. The things you'll learn from this book might help you solve actual problems, but they're more likely to help you converse with a suitable set of guests at a cocktail party.

Like pants, the book is divided into two parts. You should find Part I gentle and accessible, as long as you've learned basic calculus. Maybe you're a curious high school student, or you work in a technical area outside of pure math, or you're currently learning undergraduate analysis. Or maybe you're familiar with the material but you want to read some bedtime stories! We include tidbits of math, history, and philosophy that are absent in standard presentations; hopefully even advanced students will find Part 1 worth reading.

Part II will be harder to digest. If you've taken a thorough undergraduate analysis course or two (or three), you should be able to follow along fine. Otherwise, well, you'll understand when you're older. But you should still get *something* out of every section.

As a reader, you'll have plenty of food for thought, but we hope you won't have to think so hard that reading this book becomes a chore. Many proofs are omitted or just outlined, since who wants a technical proof in the middle of a bedtime story? We try to provide references for the proofs though. Most can be found in standard undergraduate or graduate texts on real and functional analysis.

The title of this book is inspired by "101 Illustrated Analysis Bedtime Sto-



Figure 1: The approximate composition of this book.

ries" [1], a brilliant work of fiction. This book is both nonfictional and about \mathbb{R} , hence the name. We didn't include any complex analysis, because we feel that topics in complex analysis are more like *magic tricks* than bedtime stories. ("Can I get a bounded entire volunteer from the audience? Abracadabra, hocus pocus, tada! You're constant.") Happy reading!

Some students Earth, 2016

References

- [1] S. Duvois and C. Macdonald. 101 Illustrated Analysis Bedtime Stories. 2001. URL: http://people.maths.ox.ac.uk/macdonald/errh/.
- [2] Paul Lockhart. A mathematician's lament. Bellevue literary press New York, 2009.

Part I The basics

Chapter 1

Sets, functions, numbers, and infinities

OK, dude. Math is the foundation of all human thought, and set theory – countable, uncountable, etc. – that's the foundation of math. So even if this class was about Sanskrit literature, it should still probably start with set theory.

Scott Aaronson [1]

This chapter isn't exactly about real analysis, but it's fun stuff that you need to understand anyway. To appreciate the real analysis stories, you need to know something about the world in which they take place.

1 Paradoxes of the smallest infinity

Are there more even integers or odd integers? How about integers vs. rational numbers? Our goal for this section is to make sense of and answer questions like these by explaining how to *compare infinities*. (See Figure 1.1.)

A set is just a collection of objects, called the *elements* of the set.¹ A set has neither order nor multiplicity, e.g. $\{1, 2, 3\} = \{3, 2, 1, 1\}$. We write $x \in X$ (read "x in X" or "x is in X") to say that x is an element of the set X.

To compare two *finite* sets, you can simply *count* the elements in each set. Counting to infinity takes too long though. As you know from playing musical

 $^{^{1}}$ Are you unsatisfied with this definition? Strictly speaking, we're taking sets as our most primitive objects, so rather than *defining* them, we should just give some axioms about them which we will assume. We'll be assuming "the ZFC axioms." Look them up if you're curious. It shouldn't matter.

4



Figure 1.1: How can infinities be compared?



Figure 1.2: The centipede can compare two finite sets even though it doesn't know how to count. The centipede puts a clean sock on each foot until it either runs out of socks or runs out of feet. If it runs out of socks with some feet still bare, it can conclude that it has more feet than clean socks, so it's time to do laundry.

chairs, to compare two finite sets X and Y, you can avoid counting or numbers and instead use an even more primitive concept: *matching*. Pair off elements of X with elements of Y one by one. The sets are the same size if and only if you end up with no leftover elements in either set. (See Figure 1.2.)

Armed with this observation, way back in 1638, Galileo declared that infinities cannot be compared. Here's his reasoning. (See Figure 1.3.) Suppose we're interested in comparing the set of natural numbers $\mathbb{N} = \{1, 2, 3, ...\}$ with the set of perfect squares $S = \{1, 4, 9, ...\}$. On the one hand, "obviously," \mathbb{N} is **bigger** than S, because S is a *proper subset* of \mathbb{N} , i.e. every perfect square is a natural number but not vice versa. If we list the natural numbers on the left and the perfect squares on the right, we can match each perfect square n^2 on the right with the copy of that same number n^2 on the left, leaving a lot of lonely unmatched natural numbers.

But on the other hand, instead, we could match each natural number $n \in \mathbb{N}$ with the perfect square $n^2 \in S$. That would leave no leftovers on either side, suggesting that \mathbb{N} and S are the **same size**! We get two different answers based on two different matching rules. It's as if we play musical chairs twice, with the same set of people and the same set of chairs both times. In the first game, the chairs all fill up, with infinitely many losers still standing. But in the rematch, everybody finds a chair to sit in! Galileo concluded that this is all just nonsense [5]:

So far as I see we can only infer that the totality of all numbers is infinite, that the number of squares is infinite, and that the number of their roots is infinite; neither is the number of squares less than the totality of all the numbers, nor the latter greater than the former; and finally the attributes "equal," "greater," and "less," are not applicable to infinite, but only to finite, quantities.

Galileo was on the right track, but he didn't get it quite right. The main takeaway from Galileo's paradox is that we really do need a *definition* in order



Figure 1.3: Galileo's paradox. With infinite sets, different matching rules can lead to different outcomes.

to compare infinite sets. In the 1800s, Georg Ferdinand Ludwig Philipp Cantor provided a good one, declaring that X has the same cardinality as Y if there is some way to pair off the elements of X with elements of Y, leaving no leftovers in either set.² So the definition is "biased" in favor of declaring sets to be the same size. Cantor says, the appropriate way to handle Galileo's paradox is to say yeah, there really are just as many natural numbers in total as there are perfect squares. Infinity's weird like that.

To explore cardinality properly, we need to be more precise. Galileo's paradox involved two different ways of associating elements of \mathbb{N} with elements of S: two different *binary relations* between \mathbb{N} and S.

Definition 1. A binary relation consists of a set X (the domain), a set Y (the codomain), and a set G of ordered pairs (x, y) where $x \in X$ and $y \in Y$. (G is called the graph of the relation.)

Definition 2. A function $f : X \to Y$ (read "f from X to Y") is a binary relation with domain X and codomain Y whose graph G satisfies the following: For every $x \in X$, there is exactly one $y \in Y$ such that $(x, y) \in G$. We write f(x) = y instead of $(x, y) \in G$. Functions are also called *maps*.

It's sometimes useful to think of f as a machine, which is given the input x and produces the output f(x). But this idea doesn't always make too much sense, because there isn't necessarily a formula or algorithm for figuring out f(x). You're probably most familiar with functions $\mathbb{R} \to \mathbb{R}$, where \mathbb{R} is the set

 $^{^2 {\}rm Since \ this \ was \ such \ a \ brilliant \ insight \ of \ Cantor's, \ philosophers "honor" \ him \ by \ referring to it as \ Hume's \ principle.$



Figure 1.4: A binary relation. The domain is \mathbb{N} , the codomain is \mathbb{N} , and the graph is the set of pairs (n, n + k) where $n, k \in \mathbb{N}$. Of course, this is just the familiar < relation.



Figure 1.5: Let OBAMA be the binary relation with domain \mathbb{R} and codomain \mathbb{R} whose graph G is depicted above. OBAMA is *not* a function, for two reasons. First, for some values of x, there are *multiple* y so that $(x, y) \in G$. (OBAMA fails the "vertical line test.") Second, for some x, there *does not exist* a y so that $(x, y) \in G$. (OBAMA is not "entire.")



Figure 1.6: A function $f : \{1, 2, 3, 4\} \to \{a, b, c\}$, with e.g. f(4) = c.



Figure 1.7: The function $f : \mathbb{N} \to \mathbb{N}$ defined by f(x) = x + 1 is injective, because no two arrows point to the same number. In contrast, the function depicted in Figure 1.6 is not injective, because 1 and 2 collide.

of real numbers $(\frac{3}{2}, -10^9, \sqrt{2}, \pi \cdot e, \text{ etc.}^3)$ These real-valued functions of a real argument are going to be the main characters in most of our stories.

A collision of a function $f: X \to Y$ is a pair of distinct inputs $x_1, x_2 \in X$ such that $f(x_1) = f(x_2)$. A function is *injective* if it has no collisions. An injective function is *lossless*: you can recover the input from the output. "An *injection* preserves *information*." [3]. To put it another way, if X is a set of people and Y is a set of chairs, an injection $X \to Y$ is a seating arrangement where each person gets her own chair, possibly leaving some chairs empty.

You should think of the codomain Y as the set of "allowed" outputs of f. The *image* of f, denoted f(X), is the set of *actual* outputs of the function, i.e. f(X) is the set of all $f(x) \in Y$ as x ranges over X.⁴ E.g. the image of the function depicted in Figure 1.6 is $\{b, c\}$. We say that f is *surjective* if f(X) = Y.

 $^{^{3}}$ Unsatisfied by this "definition" as well? We'll discuss what real numbers really are in Section 4. For now, just think of points on a number line, or decimal expansions.

 $^{^4}$ You might have heard the term "range" before. The word "range" is ambiguous. Don't use it. When people say "range", sometimes they mean codomain, and other times they mean image.



Figure 1.8: "Sir Jective hits everything with his sword." –Kevin [12]. See also [11].

In other words, f is surjective if for every $y \in Y$, there exists an $x \in X$ so that f(x) = y. A surjection is a seating arrangement which fills every chair, possibly with many people sharing a single chair.

A function is *bijective* if it is injective and surjective. A bijection is also called a one-to-one correspondence:⁵ it is the notion of "matching with no leftovers" that we were looking for. A bijection is a seating arrangement in which every person is assigned her own chair and every chair is filled. Here's the official version of Cantor's definition.

Definition 3. Suppose X and Y are sets. We say that X has the same cardinality as Y if there exists a bijection $X \to Y$. We write |X| = |Y| in this case.

Example 1. The set of even integers has the same cardinality as the set of odd integers, because f(2k) = 2k + 1 is a bijection between these two sets. (See Figure 1.10.) This should be intuitive, since even and odd seem to be "on equal footing."

⁵Warning: some mathematicians use this phrase "one-to-one correspondence" to mean bijection, and then in the same breath use the term "one-to-one" to mean injection (note the omission of the word "correspondence.") Some say "f maps X onto Y" to say that f is surjective, while the subtly different "f maps X into Y" merely means that X and Y are the domain and codomain of f! It's a terminological disaster. Much better to stick with the injective/surjective/bijective terms, invented by the group of mathematicians known pseudonymously as Bourbaki.



Figure 1.9: The function $f(x) = x^2$ is not surjective when thought of as a function $\mathbb{R} \to \mathbb{R}$, because negative numbers are not part of its image. However, it is surjective if we think of it as a function $\mathbb{R} \to [0, \infty)$.



Figure 1.10: The bijection from the set of even integers to the set of odd integers.

1. PARADOXES OF THE SMALLEST INFINITY

Example 2. The set of *all* integers (positive, negative, and 0) has the same cardinality as \mathbb{N} (the set of positive integers). To see why, observe that we can *reorder* the integers as follows (see also Figure 1.11):

$$0, -1, 1, -2, 2, -3, 3, \ldots$$

The function f(n) which gives the *n*th element in the list is a bijection from \mathbb{N} to the set of integers. We denote the set of all integers by \mathbb{Z} (which stands for "Zahlen", the German word for number). So what we've just shown is that $|\mathbb{Z}| = |\mathbb{N}|$. This is counterintuitive: it feels like there are about "twice as many" integers as positive integers.



Figure 1.11: To count the integers, we fold \mathbb{Z} in half.



Figure 1.12: The proof that $|\mathbb{Q}| = |\mathbb{N}|$. We make an infinite table of fractions, with the row index being the denominator and the column index being the numerator. We circle all of the *reduced* fractions, and then we can make a list of all the rational numbers in the zig-zag order indicated by the arrows. A small simplification made in this illustration is that it omits the nonpositive rational numbers.

Example 3. The set \mathbb{Q} of all *rational* numbers (i.e. fractions of integers) has the same cardinality as \mathbb{N} ! This seems horribly wrong, because there are *infinitely* many rational numbers between every two integers. It's sufficiently surprising that at least one high school textbook [2] boldly asserts that \mathbb{Q} and \mathbb{N} have different cardinalities. But in fact, we can enumerate the rational numbers as follows. Every rational number can be written as a reduced fraction $\pm \frac{p}{q}$, where p is a nonnegative integer and q is a positive integer. First, we list all rational numbers with p + q = 1 (there's just one: zero.)

$$\frac{0}{1}$$

Then, we list all rational numbers with p + q = 2:

$$\frac{0}{1}, \frac{1}{1}, -\frac{1}{1}$$

Then all rational numbers with p + q = 3:

$$\frac{0}{1}, \frac{1}{1}, -\frac{1}{1}, \frac{1}{2}, -\frac{1}{2}, \frac{2}{1}, -\frac{2}{1}$$

Etc. etc. Every rational number will eventually be listed. Just like in the case of \mathbb{Z} , this reordering immediately gives a bijection between \mathbb{Q} and \mathbb{N} , showing that $|\mathbb{Q}| = |\mathbb{N}|$. (See Figure 1.12.)

We call a set *countably infinite* if it has the same cardinality as \mathbb{N} . (If you carefully count the elements in a countably infinite set, it's false that you

12



Figure 1.13: Hilbert's paradox of the Grand Hotel. There are countably infinitely many rooms, all of which are occupied, yet the hotel is still accepting more guests. When a new guest arrives, the hotel asks the patron in room n to move to room n+1, with the net effect being that room 1 is freed up for the new arrival. Do you see how the hotel can deal with countably infinitely many guests who all arrive simultaneously? (Credit for the "No vacancy, guests welcome" sign: [8])

will eventually have counted every element, but it's true that for each element, you will eventually have counted that element.) Countable infinity may be the smallest infinity, but it's got teeth. (See Figure 1.13.)

2 Uncountability

A simple observation: a set is finite if and only if you can write down the entire set, after giving each element a name. There's a similar characterization of countable sets. (A set is *countable* if it is either finite or countably infinite.) If X is countable, maybe you can't write down the entire set, but at least you can write down an arbitrary element of X.

Proposition 1. A set X is countable if and only if each element of X can be written down. More precisely, X is countable if and only if there is some finite alphabet Σ and an injection from X to the set Σ^* of finite strings of symbols from Σ .

Geology rocks.							
Vacuuming sucks.							
Don't drink and derive.							
Two wrongs can make a riot.							
Why is the letter before Z?							
Statisticians say mean things.							
Your calendar's days are numbered.							
A plateau is the highest form of flattery.							
÷							
Two fish are in a tank. One says to the other, "Do							
you know how to drive this thing?''							
Bobby Fischer got bored of playing chess with							
Russians. He asked the association to fix his next							

Figure 1.14: The set of all puns is countable, because every pun can be written down, and hence the puns can be enumerated: we start with the shortest, then move on to longer and longer puns. (We deserve no credit for the puns listed.)

match with some other Europeans, writing, "How about

a Czech mate?''

Before the proof, some examples: We can write down an arbitrary element of \mathbb{N} using the alphabet $\Sigma = \{0, 1, 2, 3, 4, 5, 6, 7, 8, 9\}$ and standard decimal notation. Similarly, to write down elements of \mathbb{Z} and \mathbb{Q} , just throw in two more symbols, - and /. (That was a much easier proof that \mathbb{Q} is countable than the zig-zag argument we did before!)

Proof of Proposition 1. If X is countable, then we can index each element of X with a natural number, which we can think of as its name. Writing down natural numbers is easy enough.

For the converse, we'll show that Σ^* is countable. To enumerate Σ^* , first list all the length-0 strings (there's only one: the empty string.) Then list all the length-1 strings, then the length-2 strings, etc. There are only finitely many strings of each length, so this gives a bijection $\mathbb{N} \to \Sigma^*$. (See Figure 1.14.)

Proposition 1 reveals tons of countable sets: the set of all finite subsets of \mathbb{Z} , the set of all polynomials with integer coefficients, the set of all possible computer viruses, the set of all possible recipes describing yummy food, the set of all love notes which can ever be written, the set of all theorems, the set of all proofs, the set of all stories, the set of all finite mazes, the set of all vague

philosophical questions, the set of all possible digital photographs, the set of all physical laws that we have any hope of making sense of...

Are there any sets that are *uncountable* – even bigger than \mathbb{N} ? Of course, by the that's-why-the-word-countable-was-invented principle. Where do we find one of these super-infinite sets, despised by Count von Count? Proposition 1 gives a hint: it ought to require an "infinite amount of information" to specify an element of the set. A *sequence* is a function with domain \mathbb{N} . If we are thinking of A as a sequence, we write A_n instead of the functional notation A(n).

Theorem 1. Let $2^{\mathbb{N}}$ denote the set of all sequences of zeroes and ones. Then $2^{\mathbb{N}}$ is uncountable.

The proof, due to Cantor, is unquestionably one of the greatest proofs of all time. Remember that the definition of cardinality was biased in favor of sets having the same cardinality, which makes it especially tricky to prove that two sets have different cardinalities. We have to prove that there *does not exist* a bijection $\mathbb{N} \to 2^{\mathbb{N}}$. Lots of people like to say that "you can't prove a negative" [4, 14, 10]. But we're about to do exactly that.

Proof. Consider any arbitrary function $f : \mathbb{N} \to 2^{\mathbb{N}}$; we will show that f is not a surjection.

We can represent f as a table, like the example in Figure 1.15. Let A be the diagonal sequence, defined by $A_n = f(n)_n$ – that is, the *n*th term of A is the *n*th term of the *n*th sequence. Let B be the opposite of A:

$$B_n = \begin{cases} 0 & \text{if } A_n = 1\\ 1 & \text{if } A_n = 0. \end{cases}$$
(1.1)

By construction, for every $n \in \mathbb{N}$, f(n) differs from B in its nth term. Thus, B is not in the image of f, so f is not surjective!

Here's a more familiar uncountable set:

Theorem 2. The set \mathbb{R} of all real numbers is uncountable.

Proof. Define $f: 2^{\mathbb{N}} \to \mathbb{R}$ by

 $f(a_1, a_2, \dots)$ = the real number represented by $0.a_1a_2...$ in base 10.

Then f is injective, and hence f is a bijection between $2^{\mathbb{N}}$ and $f(2^{\mathbb{N}}) \subsetneq (-1, 1)$. This shows that some *subset* of \mathbb{R} is uncountable, which implies that \mathbb{R} is uncountable.

The diagonalization argument in the proof of Theorem 1 is extremely clever, and it took Cantor a long time to figure it out. After struggling for many years to figure out whether $|\mathbb{R}| \stackrel{?}{=} |\mathbb{N}|$, he asked Richard Dedekind for help in 1873, but Dedekind couldn't solve the problem either. Cantor eventually published a proof that \mathbb{R} is uncountable in 1874. This early proof, which we'll see in Section 15,

n	f(n)										
1	0	0	0	0	0	0	0	0	0	0	
2	1	1	1	1	1	1	1	1	1	1	
3	0	1	0	1	0	1	0	1	0	1	
4	0	1	0	0	1	0	0	0	1	0	
5	1	1	1	1	0	1	0	1	0	0	
6	0	1	0	0	0	1	1	0	1	0	
7	1	0	0	1	0	0	1	0	0	1	
8	0	1	1	1	1	1	1	1	1	1	
9	1	0	1	1	1	1	1	1	1	1	
10	0	0	0	1	1	1	0	0	0	1	
÷	÷	÷	÷	÷	÷	÷	÷	÷	÷	÷	·
\overline{A}	0	1	0	0	0	1	1	1	1	1	
B	1	0	1	1	1	0	0	0	0	0	

Figure 1.15: How the proof of Theorem 1 works for one example function f. The sequence B cannot be in the image of f, because for every n, B and f(n) disagree at their nth position.

was less slick than the proof that we just saw. Cantor finally published his diagonalization argument in 1891 [6].

Theorem 2 is profound. Obviously some numbers, like π , have infinite decimal expansions. We still manage to write down such numbers, by using special notation, like the symbol π . But Theorem 2 tells us that no matter how much notation we make up, there will still be some numbers which cannot be written down! As Shakespeare said,

There are more things in heaven and earth, Horatio, than are dreamt of in your philosophy.

For example, there must exist *noncomputable numbers* – numbers for which there is no algorithm for listing the digits of the number. Numbers which turn up "in the wild" tend to be computable (π , e, $\sqrt{2}$, etc.) But the noncomputable ones are out there!

Notice that in the proof of Theorem 2, we actually showed that the interval (-1, 1) is already uncountable! Intuition suggests that \mathbb{R} has a greater cardinality than a puny little interval like (-1, 1), but you've probably learned by now that your intuition can be misleading in this business:

Proposition 2. For any real numbers a < b, $|(a,b)| = |\mathbb{R}|$.

Proof sketch. The function $f(x) = \tan(x)$ is a bijection $\left(-\frac{\pi}{2}, \frac{\pi}{2}\right) \to \mathbb{R}$. (See Figure 1.17.) By translating and scaling like in Figure 1.16, you can get a bijection $(a, b) \to \mathbb{R}$.

Proposition 2 is bizarre, because we like to think of intervals as having "different sizes," e.g. (0,2) should be twice as big as (0,1). We'll address that idea in depth in Chapter ??.

16



Figure 1.16: f(x) = 2x is a bijection $(0, 1) \to (0, 2)$.



Figure 1.17: Comparing (0, 1) and \mathbb{R} .

To illustrate the care that must be taken to show that one set is bigger than another, we conclude this section with some philosophical nonsense.

SIMPLICIO: I've discovered a proof that there are more bad experiences than good experiences. Take any good experience, and imagine altering it by setting yourself on fire. Now it's a bad experience! So there are at least as many bad experiences as good experiences, and of course there are some bad experiences in which you're not on fire, so the inequality is strict.

SALVIATI: No, that won't do. You've provided a map from the set of good experiences to the set of bad experiences (the "set-yourself-on-fire" map) which is injective, but not surjective. Your conclusion that there are more bad experiences than good experiences would only be justified if we were dealing with finite sets. (After all, the map $f : \mathbb{N} \to \mathbb{N}$ defined by f(x) = x + 1 is injective but not surjective! You don't think that \mathbb{N} is bigger than *itself*, do you?) But in actual fact, there are infinitely many experiences. Just consider the experience of holding n marbles, for $n \in \mathbb{N}$. There's a different experience for each n.

SIMPLICIO: No no, you've misunderstood what I mean by "experience." You thought that I meant a *situation*, which the subject

of the situation would judge to be good or bad. But I meant the perception of that situation. That is, an experience is a brain state, or rather a sequence of brain states, bounded in length by the human lifespan. I claim that there are only finitely many experiences. For example, if n and m are sufficiently large, then holding n marbles is indistinguishible from holding m marbles, and hence they are the same experience. My proof is salvaged.

SALVIATI: Ah, but if "experience" means sequence of brain states, then the set-yourself-on-fire map is *not injective*! Consider two good experiences in which you are watching a sunset. In one experience, a squirrel runs by at some distance from you. In the other experience, there is no squirrel. When you set yourself on fire, these cease to be distinct experiences, because you wouldn't notice the squirrel if you were on fire!

3 Cantor's infinite paradise of infinities

Definition 4. For sets X, Y, we say that $|X| \leq |Y|$ if there exists an injection $X \to Y$.

You'll be happy to know that every two sets can be compared in this way.

Theorem 3. For any two sets X and Y, $|X| \leq |Y|$ or $|Y| \leq |X|$.

(See [9] for a proof.) You'll also be happy to know that if $|X| \leq |Y|$ and $|Y| \leq |X|$, then |X| = |Y|.

Theorem 4 (Cantor-Bernstein-Schröder). Suppose X and Y are sets. If there is an injection $f: X \to Y$ and another injection $g: Y \to X$, then there exists a bijection $h: X \to Y$.

Despite the theorem's name, Dedekind was the first one to prove it, and Cantor never gave a proof for it (though he was the first to state it.) The proof of Theorem 4 is notoriously difficult, but in principle, it requires no deep mathematics education to understand. If you want to confuse yourself, read one of these proofs: [TODO references]

We've seen that there are at least two different kinds of infinity (countable and uncountable.) Obvious question: is there a *biggest* infinity?

Definition 5. For a set X, the power set of X (denoted $\mathcal{P}(X)$) is the set of all subsets of X. For example, if $X = \{1, 2, 3\}$, then

$$\mathcal{P}(X) = \left\{ \varnothing, \{1\}, \{2\}, \{3\}, \{1,2\}, \{1,3\}, \{2,3\}, \{1,2,3\} \right\}.$$

(The symbol \varnothing denotes the *empty set*, the set with no elements.)

Our next theorem implies that no matter how huge a set you come up with, there is always an even huger set. The proof is just a slightly more abstract version of the diagonalization argument that revealed uncountability.

3. CANTOR'S INFINITE PARADISE OF INFINITIES

Theorem 5 (Cantor's theorem). For every set X, $|X| < |\mathcal{P}(X)|$.

Proof. Consider an arbitrary function $f: X \to \mathcal{P}(X)$. Let A be the diagonal set, i.e.

$$A = \{ x \in X : x \in f(x) \}.$$

The above bit of notation is read "the set of all x in X such that x is in f(x)," and it means exactly what it sounds like. Let

$$B = X \setminus A = \{x \in X : x \notin A\}$$

Fix any $x \in X$; we'll show that $f(x) \neq B$. If $x \in f(x)$, then $x \in A$, so $x \notin B$. If $x \notin f(x)$, then $x \notin A$, so $x \in B$. Either way, f(x) disagrees with B on x. Hence, B is not in the image of f, so f is not surjective.

("Wait," you ask, "isn't $X = \emptyset$ a counterexample to Cantor's theorem?" No, because $\mathcal{P}(\emptyset) = \{\emptyset\}$, which has one element, whereas \emptyset has zero elements.)

Cantor's theorem uncovers a rabbit hole to the wonderland of set theory. ("No one shall expel us from the Paradise that Cantor has created." –David Hilbert [7]) This book is supposed to be about real analysis, so we're not going to explore the rich landscape of infinities in depth, but we'll visit it as tourists and see some sights, to give you a bit more intuition about infinity.

We've defined expressions like $|X| \leq |Y|$, but we haven't actually defined the object |X| by itself. If X is finite, |X| is just the number of elements in the set X. For X infinite, it's a bit trickier to give a suitable definition; suffice it to say that one can be given. These objects |X| are called *cardinal numbers*. Despite what you may have heard, infinity *is* a number, or rather many numbers.⁶ It's just not a *real number*, the kind of number with which you are most familiar.

So what's to be done with all these numbers? Arithmetic!

3.1 Cardinal addition

When you were a newborn baby learning arithmetic of natural numbers, you were taught that n + m is the number of apples you have in total if you combine a pile of n apples with a pile of m apples. Notice that there's a hidden technical assumption, which is that the two piles of apples are *disjoint*, i.e. they don't share any apples in common! (See Figure 1.18.)

Even with infinitely many apples, the definition still stands. For two sets X and Y, the union $X \cup Y$ is the set of all x such that $x \in X$ or $x \in Y$. The *intersection* $X \cap Y$ is the set of all x such that $x \in X$ and $x \in Y$. (See Figure 1.19.) If X and Y are disjoint (i.e. $X \cap Y = \emptyset$), we define

$$|X| + |Y| = |X \cup Y|.$$

If X and Y are not disjoint, just rename the elements of each set, giving new sets X' and Y' which are disjoint satisfying |X| = |X'| and |Y| = |Y'|.

 $^{^6{\}rm This}$ is actually one of many senses in which infinity is a number. See also ordinal numbers, hyperreal numbers, surreal numbers.



Figure 1.18: Despite what these two piles of apples may suggest, $5 + 4 \neq 7$.



Figure 1.19: The Boolean set operations: union, intersection, and set difference.

3. CANTOR'S INFINITE PARADISE OF INFINITIES

Our "fold in half" proof that $|\mathbb{N}| = |\mathbb{Z}|$ can easily be adapted to show that

$$|\mathbb{N}| + |\mathbb{N}| = |\mathbb{N}|.$$

We also have $|\mathbb{R}| + |\mathbb{R}| = |(0,1)| + |(0,1)| \le |(0,2)| \le |\mathbb{R}|$, and hence

$$\mathbb{R}|+|\mathbb{R}|=|\mathbb{R}|.$$

These two calculations are not coincidences: it turns out that for any infinite set X,

$$|X| + |X| = |X|.$$

Adding an infinity to itself doesn't do anything!

3.2 Cardinal multiplication

As an infant, you were taught that $n \cdot m$ is the number of apples in an $n \times m$ grid of apples. In general, for two sets X and Y, the *Cartesian product* $X \times Y$ is the set of all ordered pairs (x, y) where $x \in X$ and $y \in Y$. For example, $\mathbb{R} \times \mathbb{R}$ is the real plane. We define

$$|X| \cdot |Y| = |X \times Y|.$$

Our "zig-zag" proof that $|\mathbb{N}| = |\mathbb{Q}|$ can easily be adapted to show that

$$|\mathbb{N}| \cdot |\mathbb{N}| = |\mathbb{N}|.$$

As we'll discuss in Section ??,

$$\mathbb{R}|\cdot|\mathbb{R}| = |\mathbb{R}|.$$

Again, these are not coincidences: for any infinite set X,

1

 $|X| \cdot |X| = |X|.$

Even *multiplying* an infinity by itself doesn't do anything!

3.3 Cardinal exponentiation

As an infant, you were taught that n^m is the repeated product $n \cdot n \cdot n \cdots n$, with n appearing m times. Combinatorially, this is the number of different ways to fill out an m-question multiple choice exam where each question has n options. In other words, if we fix a set N with n elements and a set M with m elements, then n^m is the number of functions $M \to N$.

In general, for two sets X and Y, we define Y^X to be the set of functions $X \to Y$. For example, earlier, we denoted the set of all binary sequences by $2^{\mathbb{N}}$. If you identify the number 2 with the two-element set $\{0, 1\}$, our notation makes good sense. We define

$$|Y|^{|X|} = \left|Y^X\right|.$$

We saw that addition and multiplication are pretty boring for infinite cardinal numbers. Is exponentiation similarly boring? You can specify a subset $A \subseteq X$ by giving its *indicator function* $\chi_A : X \to \{0, 1\}$ defined by

$$\chi_A(x) = \begin{cases} 0 & \text{if } x \notin A \\ 1 & \text{if } x \in A. \end{cases}$$

 $(\chi_A(x) \text{ "indicates" whether } x \text{ is in } A.)$ Hence, Cantor's theorem implies that $|X| < 2^{|X|}$ for every set X. So cardinal exponentiation is not boring – unlike addition and multiplication, exponentiation can actually get you somewhere. You'll be happy to know that standard exponent rules, like $(a^b)^c = a^{b \cdot c}$, hold for cardinal numbers.

We use the notation \beth_0 (pronounced "bet nought") to denote $|\mathbb{N}|$, the cardinality of the countably infinite. (Here \beth is the second⁷ letter of the Hebrew alphabet.) Then we define $\beth_1 = 2^{\beth_0}$, and more generally $\beth_{n+1} = 2^{\beth_n}$, giving a whole sequence of increasingly enormous infinite cardinal numbers. It turns out that $|\mathbb{R}| = \beth_1$. In "normal mathematics" (i.e. outside set theory), the only infinite cardinalities you're likely to encounter are \beth_0 , \beth_1 , \beth_2 , and maybe \beth_3 .

That's all we're going to say⁸ about sets and functions, the basic foundations of math that they ought to teach in middle school. It's time to discuss the foundations of real analysis in particular.

4 The real deal

Real⁹ numbers seem like awfully familiar friends, but evidence suggests that *different people have different concepts in mind* when they talk about numbers. Back in the 5th century B.C., the Pythagoreans had a confused conception of number. On the one hand, they thought that all numbers were either integers or fractions (i.e. rational numbers.) But on the other hand, the Pythagoreans wanted to use numbers to describe Euclidean geometry, including, of course, the famous Pythagorean theorem. This led to trouble:

Proposition 3. $\sqrt{2}$ is irrational. That is, there is no rational number $\frac{p}{q}$ so that $\left(\frac{p}{q}\right)^2 = 2$.

Proof. Assume for a contradiction that $p^2 = 2q^2$, where $\frac{p}{q}$ is a *reduced* fraction. Then p^2 is even, which implies that p is even, and hence p^2 is divisible by 4.

22

⁷The cardinality of \mathbb{N} is more often denoted \aleph_0 ; here \aleph is the first letter of the Hebrew alphabet. We prefer the \beth notation, because the other \aleph numbers are much more confusing. Look up the continuum hypothesis and the generalized continuum hypothesis if you're curious.

⁸If you'd like a more detailed exposition of this sort of "background material" that most math books assume you already know, you might try Chapter 1 of James R. Munkres' book "Topology" [13]. The rest of that book is quite good too, if you're interested in learning about topology!

⁹The word "real" here doesn't actually mean anything. (It's Descartes' fault.) Imaginary numbers and real numbers are equally nonfictional. Real numbers should've been called "continuum numbers" or "line numbers" or something. Too late now.



Figure 1.20: Uh oh, rational numbers and triangles are not friends.

Hence $2q^2$ is divisible by 4, so q^2 is even, which implies that q is even. But that contradicts the fact that $\frac{p}{q}$ is reduced.

Legend has it that the irrationality of $\sqrt{2}$ was discovered by Hippasus while he was at sea, and his fellow Pythagoreans were so enraged that they threw him into the ocean, where he drowned. Unfortunately, as far as we can tell, the legend is just made up.

In modern times, many people claim that $0.999... \neq 1$ when asked. (Over 80% in one small study [15].) Mathematicians, on the other hand, are all confident that 0.999... = 1. (More on this in Section 6.) Are non-mathematicians just not thinking clearly when they say that 0.999... is "infinitesimally smaller" than 1? It's more reasonable to suggest that they simply were never told clearly what numbers are and how they are represented, so they came up with their own mental model which doesn't match the standard definitions used by mathematicians. Let's clear up these definitional issues now.

Definition 6 (Real numbers). A *real number system* is a 4-tuple¹⁰ ($\mathbb{R}, +, \cdot, \leq$), where \mathbb{R} is a set, + and \cdot are binary operations on \mathbb{R} , and \leq is a binary relation on \mathbb{R} , satisfying the "real number axioms."

Well that wasn't a very good definition. We'd better tell you what the real number axioms are, eh? Most of them are pretty boring. You should just skim them to get the flavor, except for Axiom 8 which is important. The first four axioms say that arithmetic works like it ought to.

Axiom 1. Addition is commutative (x+y=y+x) and associative ((x+y)+z = x + (y+z)). There is an additive identity 0 (x + 0 = x) and every number x has an additive inverse -x (x + -x = 0).

Axiom 2. Multiplication is commutative and associative, there is a multiplicative identity 1, and every nonzero number x has a multiplicative inverse $\frac{1}{x}$.

Axiom 3. Multiplication distributes over addition $(x \cdot (y + z) = x \cdot y + x \cdot z)$

Axiom 4 (Everyone's favorite axiom). $1 \neq 0$.



Figure 1.21: Number systems satisfying Axioms 1 through 4 (with no order structure) are called *fields*. There are a lot of bizarre fields which are nothing like \mathbb{R} . For example, $\mathbb{Z}/7\mathbb{Z}$ is the field you get by coiling \mathbb{Z} up into a circle, pretending that n and n + 7 are the same number for every n. Division in this field is pretty weird, e.g. $\frac{1}{3} = 5$, since $5 \cdot 3 = 15 = 1$. The point is, Axioms 5 through 8 are important.

The last four axioms deal with the order structure of \mathbb{R} .

Axiom 5. The order is transitive $(x \le y \text{ and } y \le z \text{ implies } x \le z)$, antisymmetric $(x \le y \text{ and } y \le x \text{ implies } x = y)$, and total (for every x, y, either $x \le y$ or $y \le x$.)

Axiom 6. If $x \leq y$, then $x + z \leq y + z$.

24

Axiom 7. If $x \ge 0$ and $y \ge 0$, then $x \cdot y \ge 0$.

So far, \mathbb{Q} has satisfied all of these axioms. From the axioms we've listed so far, you can derive familiar facts like "zero times anything is zero" and "a negative times a negative is a positive." Yawn. There's one more axiom, and it's the fun one. When is the last moment of Sunday? Midnight? No, that's Monday already...

Definition 7. Fix $X \subseteq \mathbb{R}$. We say $a \in \mathbb{R}$ is an *upper bound* for X if $a \ge x$ for all $x \in X$.

Definition 8 (Supremum). Fix $X \subseteq \mathbb{R}$. A number $b \in \mathbb{R}$ is called the *supremum* of X if b is the *least upper bound* of X. That is, b is an upper bound for X, and if a is another upper bound for X, then $b \leq a$.

The supremum of X is denoted sup X; the abbreviation sup is pronounced like "soup." If a set X has a maximum, then $\sup X = \max X$. Some sets, like (0,1), have no maximum, but still have a supremum; $\sup(0,1) = 1$. There is no last moment of Sunday; the set of moments which are on Sunday does not have a maximum. But it does have a supremum: midnight. You might say that midnight is the *sup du jour!* (Ba dum tss.)

¹⁰An *n*-tuple is just an ordered list of *n* objects. So all we're saying is that a real number system has four parts: \mathbb{R} , +, \cdot , and \leq .



Figure 1.22: \mathbb{Q} is like Swiss cheese: it's riddled with holes. \mathbb{R} is like cheddar cheese: it tastes good grated over scrambled eggs.



Figure 1.23: The countable set $S = \{\frac{1}{n} : n \in \mathbb{N}\}$. This set has a maximum, $\max S = \sup S = 1$. It has no minimum, but inf S = 0.

Axiom 8 (Supremum axiom). If $X \subseteq \mathbb{R}$ is nonempty and X has an upper bound, then X has a supremum.

The supremum axiom is crucial in real analysis. It is the reason that $(\mathbb{Q}, +, \cdot, \leq)$ is not a real number system. To see why, consider the set $S = \{r \in \mathbb{Q} : r^2 < 2\}$. The supremum of S ought to be $\sqrt{2}$. But if \mathbb{Q} is our whole world, then there is no square root of 2! When S is considered to be a subset of \mathbb{Q} , it has no supremum: for every rational number a with $a^2 \geq 2$, there is a smaller rational number b which also satisfies $b^2 \geq 2$. This is why we prefer \mathbb{R} to \mathbb{Q} . In some sense, \mathbb{Q} is riddled with holes. The supremum axiom asserts that all holes are filled in. (See Figure 1.22.)

The *infimum* of a set X (denoted inf X) is the greatest lower bound of X. Just like the concept of supremum is a generalization of maximum, the concept of infimum generalizes that of minimum. It follows from the real number axioms that any nonempty set with a lower bound has an infimum. (See Figure 1.23.)

So that concludes the definition¹¹ of what a real number system is. But we're not done yet. Usually, people speak of *the* real number system, and we need to justify that terminology. Does there even *exist* a real number system? (If some of the real number axioms *contradict* each other, we are in serious trouble!)

Theorem 6. Thankfully, there does exist a real number system.

We'll sketch a proof of Theorem 6 in Section 5. So there's a real number system, but is it unique? Not quite. Let * denote some fixed object, e.g. the

 $^{^{11}\}mathrm{Mathematicians}$ summarize the whole definition by saying that a real number system is a "complete ordered field."

empty set, or Abraham Lincoln, or radical freedom. Given one real number system $(\mathbb{R}, +, \cdot, \leq)$, we can build another real number system. Our new set of real numbers is $\mathbb{R} \times \{*\}$, i.e. the set of all pairs (x, *) where $x \in \mathbb{R}$. Arithmetic is defined by (x, *) + (y, *) = (x + y, *) and $(x, *) \cdot (y, *) = (x \cdot y, *)$, and the order is defined by saying that $(x, *) \leq (y, *)$ if and only if $x \leq y$.

But that's dumb. All we did is *rename* each number x to (x, *), which shouldn't count as building a whole new real number system. "A rose by any other name would smell as sweet." This renaming silliness is the only thing that goes wrong; any two real number systems are *isomorphic*, i.e. each can be obtained from the other by renaming the elements. Precisely:

Theorem 7. The real number system is unique up to ordered field isomorphism. That is, if $(\mathbb{R}_1, +_1, \cdot_1, \leq_1)$ and $(\mathbb{R}_2, +_2, \cdot_2, \leq_2)$ are two real number systems, then there exists a bijection $f : \mathbb{R}_1 \to \mathbb{R}_2$ so that

- For all $x, y \in \mathbb{R}_1$, f(x + y) = f(x) + f(y).
- For all $x, y \in \mathbb{R}_1$, $f(x \cdot_1 y) = f(x) \cdot_2 f(y)$.
- For all $x, y \in \mathbb{R}_1$, $x \leq_1 y$ if and only if $f(x) \leq_2 f(y)$.

(For example, f(x) = (x, *) is an isomorphism $\mathbb{R} \to \mathbb{R} \times \{*\}$.) The proof of Theorem 7 is not too hard but somewhat tedious, so we'll omit it.

So now we can speak of *the* real number system \mathbb{R} , with the slight caveat that you're only allowed to ask questions which can be phrased in terms of $+, \cdot,$ and \leq . The reals are only *defined* up to ordered field isomorphism, so questions like "Is $\sqrt{2} \subseteq \pi$?" aren't meaningful. You might be thinking, "Duh, $\sqrt{2}$ and π aren't sets," but it's not that simple. In any particular real number system, $\sqrt{2}$ and π are sets! Under the hood, everything's a set. The question's not meaningful because in some real number systems (e.g. "Dedekind cuts") $\sqrt{2} \subseteq \pi$, but in others (e.g. "Cauchy sequences") $\sqrt{2} \not\subseteq \pi$. On the other hand, questions like "Does every real number have a real square root?" make perfect sense, because the answer is the same ("no") in every real number system.



Figure 1.24: In ancient times (circa 1970), engineers used *slide rules* to quickly multiply and divide numbers. A simple "circular slide rule" is depicted; the gray portion rotates relative to the white portion. The depicted position corresponds to multiplication/division by 2. Slide rules work because the exponential function $f(x) = e^x$ is an *isomorphism* between the additive structure of \mathbb{R} and the multiplicative structure of $(0, \infty)$, because of the standard exponent rule $e^{x+y} = e^x \cdot e^y$. (This is a slightly simpler kind of isomorphism than the ordered field isomorphism of Theorem 7, because here we're just preserving the structure of one operation, whereas in Theorem 7 we preserved the structure of two operations and a relation.)

5 Building real numbers (Dedekind cuts)

Let's start from \mathbb{Q} and build a real number system, thereby sketching a proof of Theorem 6. Dedekind observed that between any two real numbers, there is a rational number. Therefore, to specify a point X on a real number line, it suffices to specify the set of rational numbers less than X. So we can just *define* the real numbers to be the appropriate sets of rational numbers.

Definition 9. A *Dedekind cut* is a set $X \subseteq \mathbb{Q}$ with the following properties:

- 1. (X is closed downward) If a < b and $b \in X$, then $a \in X$.
- 2. (X has no maximum) If $a \in X$, there is some $b \in X$ with a < b.
- 3. (X is nontrivial) $\emptyset \neq X \neq \mathbb{Q}$.

For example, $\{x \in \mathbb{Q} : x < 0 \text{ or } x^2 < 2\}$ is a Dedekind cut, which will be the real number $\sqrt{2}$. (See Figure 1.25.) To officially define our real number system, we define \mathbb{R} to be the set of all Dedekind cuts. We identify each rational number x with the real number $X = \{y \in \mathbb{Q} : y < x\}$. If X and Y are Dedekind cuts, then we set

$$X + Y = \{x + y : x \in X, y \in Y\}.$$

We say that $X \leq Y$ if $X \subseteq Y$. Multiplication is a little more annoying because of minus signs. If $X, Y \geq 0$, then we define

$$X\cdot Y=\{x\cdot y: x\in X, y\in Y, x\geq 0, y\geq 0\}\cup\{x\in \mathbb{Q}: x<0\}.$$

We define -X by

$$-X = \{x - y : x < 0, y \notin X\}$$

And now we can extend our definition of multiplication to all reals by setting $(-X) \cdot Y = X \cdot (-Y) = -(X \cdot Y)$ and $(-X) \cdot (-Y) = X \cdot Y$. It's tedious, but it can be verified that these definitions make $(\mathbb{R}, +, \cdot, \leq)$ a real number system.

There are lots of alternative constructions of \mathbb{R} . For example, the *Cauchy* sequence construction, discovered by Cantor [TODO cite], is more in the spirit of real analysis. Cantor's idea is based on the fact that every real number is a limit of a sequence of rational numbers. This suggests defining real numbers to be convergent sequences of rational numbers. The trouble is that this is circular – "convergent" here means converging to a real number! So instead, Cantor defined real numbers to be those sequences of rational numbers which "deserve" to converge to something. You might enjoy reading about this construction in more detail [TODO reference], but you should read Section 7 first.



Figure 1.25: The Dedekind cut identified with $\sqrt{2}$ is the set of shaded rational numbers.

6 The Cantor set

Time for our first legitimate *real analysis* problem. A set $E \subseteq \mathbb{R}$ is *dense* if it intersects every open interval $J \subseteq \mathbb{R}$. For example, \mathbb{Q} is dense (this is why it's surprising that \mathbb{Q} is countable, and this is why the Dedekind cut construction works.) More generally, if $I \subseteq \mathbb{R}$ is an open interval, we say that E is *dense* in I if E intersects every subinterval $J \subseteq I$. For example, $\mathbb{Q} \cap [0, 1]$ is dense in I = (0, 1), but not in I' = (0, 2).

A set *E* is nowhere dense if there is no interval *I* in which *E* is dense. A nowhere dense set is just like your friend's arguments against your favorite political positions: no matter which part you zoom in on, you can see a gaping hole. For example, \mathbb{Z} is nowhere dense. For another example, the set $\{\frac{1}{n} : n \in \mathbb{N}\}$ is nowhere dense. (See Figure 1.26.) In some sense, a nowhere dense set is *small*. How does this notion of size interact with our earlier notion, cardinality? \mathbb{Q} shows that countable does not imply nowhere dense. Does nowhere dense imply countable?



Figure 1.26: The set $S = \{\frac{1}{n} : n \in \mathbb{N}\}$ is nowhere dense. Given any interval (such as the blue interval), there is a subinterval (in red) which completely misses S.

Nope! The simplest counterexample, discovered by Henry Smith in 1874, is denoted Δ and called the Cantor set. (Cantor popularized it.) To construct Δ , we start with the interval [0, 1], and remove a bunch of open intervals. In the first iteration, we remove the middle third interval $(\frac{1}{3}, \frac{2}{3})$. This leaves us with the two intervals $[0, \frac{1}{3}]$ and $[\frac{2}{3}, 1]$. Next, we remove the middle third interval from each remaining interval, so that we're left with four intervals. We continue the process of removing the middle third of each remaining interval ad infinitum, and then Δ is everything left over. That is, if we let Δ_n be the set that we have after *n* iterations of this process, then $\Delta = \bigcap_n \Delta_n$, the set of points which are in every Δ_n . The first few iterations are shown in Figure 1.27.



Figure 1.27: Construction of the Cantor set.

It's easy to see that Δ is nowhere dense: for any open interval I that intersects [0, 1], there is a sufficiently large n so that the nth step of constructing Δ involves removing a subinterval of I. In fact, after removing all those intervals, how much of [0, 1] is left over? The sum of the lengths of the intervals that make up Δ_n is $(\frac{2}{3})^n$, so if we take a limit as $n \to \infty$, we see that the "total length" of Δ is 0. (We'll come back to this calculation in Chapter ??.) So Δ must be empty... right? Wrong! For example, $0, 1 \in \Delta$. In fact, Δ has *infinitely* many points: all the endpoints, which include for example any number of the form $\frac{1}{3^n}$, are in Δ .

But to really understand $|\Delta|$, we need to take a detour. So far, we've talked about real numbers in the abstract. When you met \mathbb{R} as a child, real numbers were presented to you in the guise of *decimal expansions*. A decimal expansion is a *string*, something like 3.14159265..., which (by definition) represents¹² the real number

$$3 + \frac{1}{10} + \frac{4}{100} + \frac{1}{1000} + \frac{5}{10000} + \cdots$$

(We haven't talked about limits yet, but since all the terms are nonnegative, you can interpret this infinite sum as the supremum of the set of partial sums.)

Proposition 4. Every real number has a decimal expansion.

¹²You might complain that we haven't explained which real number is referred to by strings like 3, 10, 4, 100, etc.! Well, you understand which *integers* are referred to by such strings, right? And the real number 1 is part of the axioms. So identify the positive integer n with the real number $1 + 1 + \cdots + 1$ (with n ones.)


Figure 1.28: 0.999... apples are depicted. Or maybe "0.999... apple is depicted"?

(We'll skip the proof.) How about uniqueness? Annoyingly, some real numbers have *two different* decimal expansions. The real number 1 can also be represented as 0.999..., where there are an infinite number of nines after the decimal point. Do you doubt it? Let's prove it. Certainly 1 is an upper bound on $\{0.9, 0.99, 0.999, ...\}$. If there were a smaller upper bound, say $1 - \varepsilon$, then ε would be *infinitesimal*: greater than zero, but smaller than $\frac{1}{n}$ for every natural number *n*. Such numbers do not exist:

Theorem 8 (Archimedean Property). For any real $\varepsilon > 0$, there exists a natural number n > 0 so that $\varepsilon > \frac{1}{n}$.

Proof. Let $S = \{n \in \mathbb{N} : n \leq \frac{1}{\varepsilon}\}$. Our goal is to show that $S \neq \mathbb{N}$. If S is empty, we're done. Otherwise, by the supremum axiom, S has a least upper bound $\sup S$. By the minimality of $\sup S$, there exists $s \in S$ with $s > (\sup S) - 1$. Then $s+1 > \sup S$, so $\sup S$ is not an upper bound on \mathbb{N} . Therefore, $S \neq \mathbb{N}$. \Box

If you're still in doubt, maybe you'd be convinced by tripling both sides of the equation $\frac{1}{3} = 0.333...$ If you're *still* uncomfortable, maybe it helps to keep in mind that decimal expansions are *just strings*, not the numbers themselves.

So is 1 the only two-faced scoundrel in \mathbb{R} ? Nope, e.g. 97.842 = 97.841999...Every real number with a *finite* decimal expansion has a second decimal expansions. But that's the only thing that goes wrong.

Proposition 5. Every real number has at most two decimal expansions. A real number has two decimal expansions if and only if it has a finite decimal expansion.¹³

Now we can finally understand $|\Delta|$, by representing numbers in *ternary*, i.e. base 3. (All of our discussion of decimal expansions applies mutatis mutandis for any integer base $b \ge 2$, or even weirder bases like base 2i where i is the imaginary unit.) The interval $(\frac{1}{3}, \frac{2}{3})$ that we remove in the first iteration of the construction of Δ consists of all those real numbers $x \in [0, 1]$ whose first ternary

 $^{^{13}}$ Note that for this proposition, we count 3 and 3.0 and +003 as all being "the same" decimal expansion. If you were trying to be careful, you might disallow leading/trailing zeroes in decimal expansions.

digit (after the decimal point¹⁴) is 1. More precisely, $(\frac{1}{3}, \frac{2}{3})$ consists of those real numbers $x \in [0, 1]$ such that in *every* ternary representation of x, the first digit is 1. Similarly, in the *n*th step, we remove those real numbers x such that in every ternary representation of x, the nth digit is 1. So what we're left with, Δ , is the set of real numbers in [0,1] which can be represented in ternary without using the digit 1. But of course there are uncountably many such real numbers, because every sequence of 0s and 2s represents a distinct such real number!

So on the one hand, Δ is "big:" it is an uncountable set. But on the other hand, Δ is "small:" it is nowhere dense, and it has total "length" zero. We'll meet Δ again many times, when these odd properties make it useful.

References

- Scott Aaronson. Quantum computing since Democritus. Cambridge: Cam-[1]bridge University Press, 2013. ISBN: 978-0521199568.
- Algebra 2, Study Guide and Intervention Workbook. McGraw-Hill Educa-[2]tion, 2014.
- [3] Darius Bacon. Comment on post in "Shtetl-Optimized" blog. URL: http: //www.scottaaronson.com/blog/?p=391#comment-13569.
- [4] Can't Prove a Negative. URL: https://www.youtube.com/watch?v= qWJTUAezxAI.
- [5] Galileo Galilei. Discourses and Mathematical Demonstrations Relating to Two New Sciences. Italy, 1638.
- Fernando Q Gouvêa. "Was Cantor Surprised?" In: American Mathemati-[6] cal Monthly 118.3 (2011), pp. 198–209.
- [7] D. Hilbert. ger. In: Mathematische Annalen 95 (1926), pp. 161–190. URL: http://eudml.org/doc/159124.
- [8] Hilbert's Infinite Hotel - 60-Second Adventures in Thought. URL: https: //www.youtube.com/watch?v=faQBrAQ8714.
- [9] Chaim Samuel H{önig. "Proof of the well-ordering of cardinal numbers". In: Proceedings of the American Mathematical Society 5.2 (Feb. 1954), p. 312. DOI: 10.1090/s0002-9939-1954-0060558-3. URL: http://dx. doi.org/10.1090/S0002-9939-1954-0060558-3.
- [10]S. Jacoby S. Roell. An Interview with Susan Jacoby on Athiesm. URL: http://fivebooks.com/interviews/susan-jacoby-on-atheism.
- Jette. The Adventures of Sir Jective. URL: http://twistedpencil.com/ [11] posting/196.
- [12]Kevin. Comment on post in "Shtetl-Optimized" blog. URL: http://www. scottaaronson.com/blog/?p=391#comment-13542.
- James Munkres. Topology. 2nd ed. Prentice Hall, 2000. [13]

32

¹⁴It would really be more appropriate to call it a *radix point*, but whatever.

- [14] Nelson L. Price. Is There A God. URL: http://www.nelsonprice.com/ is-there-a-god/.
- [15] D. Tall. "Cognitive development in advanced mathematics using technology". In: *Mathematics Education Research Journal* 12 (Dec. 2000), pp. 196–218. DOI: 10.1007/BF03217085.

Chapter 2

Discontinuity

For half a century we have seen a mass of bizarre functions which appear to be forced to resemble as little as possible honest functions which serve some purpose.

Henri Poincaré [7]

7 Guessing function values

The heroes of this chapter are functions $f : \mathbb{R} \to \mathbb{R}$, i.e. functions which eat a number and spit a number back out. You met these functions in school and drew their graphs. (See Figure 2.1.) Roughly speaking, we say that f is *continuous* if you can draw its graph without ever picking up your pencil. (Euler "defined" continuity by saying that a f is continuous if the graph of f can be "described by freely leading the hand." [TODO cite]) So in Figure 2.1, f is continuous but g isn't.

Consider the function $f : \mathbb{R} \to \mathbb{R}$ defined by

$$f(x) = \sin(1/x),$$

shown in Figure 2.2. (For this section, we adopt the convention sin(1/0) = 0.) Is f(x) continuous? How about the function

$$g(x) = x\sin(1/x),$$

shown in Figure 2.3?

Evidently, Euler's "pencil-never-leaves-the-paper" nonsense is not a precise enough definition of continuity! The idea of the true definition, first given by Bolzano [1], is that f is continuous at x if f(x) is exactly what you'd expect it to be, based on how f behaves near x. Continuous functions are predictable. These "predictions" are, of course, limits.



Figure 2.1: Some boring functions $\mathbb{R} \to \mathbb{R}$.



Figure 2.2: The topologist's sine curve, $f(x) = \sin(1/x)$.



Figure 2.3: The topologist's sine curve after a pliers accident, $g(x) = x \sin(1/x)$.



Figure 2.4: This guy (Bernard Bolzano) invented continuity in 1817. Don't let anyone try to tell you that it was Cauchy's idea.

Definition 10 (Limit of a sequence). Suppose x_1, x_2, \ldots is a sequence of real numbers, and L is a real number. We say that x_n converges to L if for all $\varepsilon > 0$, there exists an N such that for all n > N,

$$|x_n - L| < \varepsilon.$$

(See Figure 2.5.) In this situation, we write $\lim_{n \to \infty} x_n = L$, or just $x_n \to L$.



Figure 2.5: The definition of the limit of a sequence. For any error margin $\varepsilon > 0$, for all sufficiently large n, x_n is within ε of L.

Traditionally, real analysis students find the epsilontics involved in the defi-

nition of a limit to be confusing.¹ Maybe a real-life example would help clarify. You are the pilot of a helicopter carrying secret agents. For their secret spy mission, it's important that you hover L feet off the ground. Let x_n be the altitude of the helicopter after you've made n adjustments. (It's a digital helicopter.) Then $x_n \to L$ means that no matter what tolerance $\varepsilon > 0$ your crazy boss demands of you, by making enough careful adjustments, you can eventually guarantee that the helicopter is within ε of L and always will be in the future. Meeting higher standards takes more time, of course: if ε is very small, then Nmight have to be very big.

Now let's move on to defining continuity. You and your spouse want to go on a trip to the Moon. Your spouse has been obsessively watching the fluctuating rocket ticket prices, trying to get the best possible deal. "I finally bought the tickets just now at time t," your spouse says.

"How much did they end up costing?" you ask.

"You don't wanna know S" your spouse replies. But you really do wanna know, so you ask, "Well how much did they cost at time t - 100?"

"Only \$200! We should've bought them then!"

"How about at time t - 10?"

"They shot up to \$1000, which scared me."

"And at time t - 1?"

"Down to \$600. I thought I'd better grab them soon."

"What about at time t - 0.1?"

"\$580." Having learned the prices at times near t, you can *extrapolate* to guess the price at t, but you'd have to assume that the price doesn't fluctuate too wildly. You keep needling your spouse, learning the prices at times t - 0.01, t - 0.0001, t - 0.0001... You gain more and more confidence in your extrapolations, because you have to assume less and less about the behavior of the price. After infinitely many questions, you've learned the price at a sequence of times t_n with $t_n \to t$, so you just have to extrapolate *infinitesimally* to infer the price at t. All you're assuming now is that the price function is *continuous at* t.

Definition 11 (Continuity). We say that $f : \mathbb{R} \to \mathbb{R}$ is continuous at $x \in \mathbb{R}$ if for every sequence of inputs x_1, x_2, \ldots converging to x, the corresponding sequence of outputs $f(x_1), f(x_2), \ldots$ converges to f(x).

To put it another way, f is discontinuous at x if there is some "misleading" sequence $x_n \to x$ with $f(x_n) \not\to f(x)$. So $f(x) = \sin(1/x)$ is discontinuous at 0, because f(0) = 0, yet there is a sequence $x_n \to 0$ so that $f(x_n) = 1$ for every n. (Figure 2.6.) Remember, though, we declared f(0) = 0 by fiat. Maybe that was a mistake. Would it be better to choose f(0) = 1? Nope: there's another sequence $y_n \to 0$ with $f(y_n) = -1$. (Figure 2.7.) So there's no value of f(0) that would make f continuous at 0.

¹Steven Krantz reports [8] that when asked to give the ε - δ definition of continuity on a quiz, one student responded: "For every $\varepsilon > 0$ there is a $\delta > 0$ such that you can draw the graph without lifting your pencil from the paper."

7. GUESSING FUNCTION VALUES

On the other hand, $g(x) = x \sin(1/x)$ is continuous everywhere. (The only worrisome spot is x = 0, but observe that $|g(x)| \le |x|$.) Let's see you draw the graph of *that*, Euler! If a function is continuous everywhere, we just say that it is *continuous*. In other words, a continuous function is one which *commutes with limits*, i.e.

$$\lim_{n \to \infty} f(x_n) = f\left(\lim_{n \to \infty} x_n\right)$$



Figure 2.6: A sequence (in red) showing that $f(x) = \sin(1/x)$ is discontinuous at 0. The sequence suggests that f(0) = 1, but in actuality f(0) = 0.



Figure 2.7: Another sequence showing that $f(x) = \sin(1/x)$ is discontinuous at 0. This time, the sequence suggests that f(0) = -1.

We'll end this section with a ridiculous theorem about infinitesimal extrapolation even in the face of discontinuity, from [10]. Let's play a game. We choose a function $f : \mathbb{R} \to \mathbb{R}$. Then a point $x_* \in \mathbb{R}$ is randomly chosen (drawn from, say, a standard normal distribution, or whatever.) We reveal to you the restriction of f to $\mathbb{R} \setminus \{x_*\}$. (I.e. you get to know f(x) for every $x \neq x_*$.) Then you have to guess what $f(x_*)$ is. You win if you get it right; we win if you get it wrong.

You're probably thinking, "I'll take a limit!" You could find a sequence $x_n \to x_*$ with $x_n \neq x_*$, and evaluate $\lim_{n\to\infty} f(x_n)$. If that limit exists, it seems like the obvious guess. If we choose a continuous f and you follow this strategy, you're guaranteed to win.

But we're not going to make it that easy. We don't make any promises at all about f. Can you still force a guaranteed win? Nah, you're doomed to



Figure 2.8: This is the sort of picture that you have to deal with in our guessing game. Every value of the function is revealed except one mysterious point.

occasionally give wrong answers. But, absurdly, you can force an *almost sure* win:

Theorem 9. There is a strategy you can follow which ensures that for any function f we choose, there are only finitely many values of x_* which lead you to lose. In particular, no matter which f we choose, your probability of winning is 100%.

Proof. Define a binary relation \sim on the set of all functions $\mathbb{R} \to \mathbb{R}$ by declaring that $f \sim g$ if f and g agree on all but finitely many points. This relation \sim is an equivalence relation, i.e. it is reflexive $(f \sim f)$, symmetric $(f \sim g \implies g \sim f)$ and transitive $(f \sim g, g \sim h \implies f \sim h)$ Therefore, \sim partitions the set of all functions $\mathbb{R} \to \mathbb{R}$ up into equivalence classes – maximal sets of functions any two of which agree on all but finitely many points. For each equivalence class \mathcal{C} , choose one representative function $f_{\mathcal{C}} \in \mathcal{C}$.

When you're presented with f with its value at x_* hidden, figure out which equivalence class f belongs to (call it C.) Then guess that $f(x_*) = f_C(x_*)$. For any f, there are only finitely many x_* causing you to lose, because $f \sim f_C!$

That proof was our first² encounter with the Axiom of Choice (AC), which is the axiom of set theory which allows the step where we defined $f_{\mathcal{C}}$.

Axiom 9 (Axiom of Choice). Suppose U is a set, $\mathcal{F} \subseteq \mathcal{P}(U)$, and $\emptyset \notin \mathcal{F}$. Then there exists $f : \mathcal{F} \to U$ such that for every $X \in \mathcal{F}$, $f(X) \in X$. (The function f is called a *choice function*.)

AC frustrates many people, because it allows for very *nonconstructive* proofs. Notice that our proof of Theorem 9 doesn't actually explain how you should play, in practice. It just shows that there exists, in the abstract, a strategy with the desired properties. Some mathematicians prefer to avoid AC when possible, but sometimes it is unavoidable. AC will be a recurring character in this book. See [5] for similar, even more ridiculous theorems, also relying heavily on the axiom of choice.

 $^{^{2}}$ Actually, several of the results that we stated without proof in Chapter 1 rely on AC.



Figure 2.9: Countably infinitely many discontinuities.

8 The Dirichlet function

Can a function be discontinuous in infinitely many places? Sure, easy peasy: a "step function" with infinitely many steps, like the floor function. (See Figure 2.9.) How about a function which has *uncountably* many discontinuities? We'll do even better. Johann Peter Gustav Lejeune Dirichlet (pronounced "deerish lay") discovered a function which is discontinuous *everywhere*.

Dirichlet realized he could exploit the fact that in every open interval $(a, b) \subseteq \mathbb{R}$, there are both rational and irrational numbers. (\mathbb{Q} and $\mathbb{R} \setminus \mathbb{Q}$ are both dense.) The Dirichlet function is another name for $\chi_{\mathbb{Q}}$, the indicator function of the rationals. As a reminder, the definition is

$$\chi_{\mathbb{Q}}(x) = \begin{cases} 1 & \text{if } x \in \mathbb{Q} \\ 0 & \text{if } x \notin \mathbb{Q}. \end{cases}$$

(See Figure 2.10.)

Proposition 6. For every $x \in \mathbb{R}$, the Dirichlet function is discontinuous at x.

Proof. There's a sequence of rational numbers x_1, x_2, \ldots converging to x, and there's another sequence of irrational numbers y_1, y_2, \ldots converging to x. By definition, $\chi_{\mathbb{Q}}(x_n) = 1$ and $\chi_{\mathbb{Q}}(y_n) = 0$, so $\chi_{\mathbb{Q}}(x_n)$ and $\chi_{\mathbb{Q}}(y_n)$ cannot both converge to $\chi_{\mathbb{Q}}(x)$.

An oddity of the Dirichlet function is that it is periodic and nonconstant, yet it has no *smallest* period: for any rational number r and any real number x, $\chi_{\mathbb{Q}}(x+r) = \chi_{\mathbb{Q}}(x)$, so $\chi_{\mathbb{Q}}$ is periodic with period r.

Dirichlet's idea spawns more monstrosities. Here's a function which is continuous at one point, but discontinuous everywhere else:

$$f(x) = \begin{cases} x & \text{if } x \text{ is rational} \\ -x & \text{if } x \text{ is irrational.} \end{cases}$$



Figure 2.10: The Dirichlet function.

(See Figure 2.11.) Even better, here's a function which is *differentiable* at one point, but discontinuous everywhere else:

$$f(x) = \begin{cases} x^2 & \text{if } x \text{ is rational} \\ 0 & \text{if } x \text{ is irrational.} \end{cases}$$

(See Figure 2.12.)

Part of the reason Dirichlet gets his name attached to $\chi_{\mathbb{Q}}$ is that it played a role in clarifying the concept of a function. Mathematicians were churning out functions way back in the 1600s in the course of developing calculus. But shockingly, it seems that the now-standard definition of function that we gave in Section 1 first appeared in a 1954 book [2]! So how did mathematicians get by in the intervening several hundred years? Well, they played around with many different notions of function, of varying degrees of rigor. For the first couple hundred years, it was popular to think of functions in terms of "formulas" or "analytic expressions," whatever that means. E.g. in 1748, Euler gave a "definition" [4]:

A function of a variable quantity is an analytic expression composed in any way whatsoever of the variable quantity and numbers or constant quantities.

In 1829 [3], Dirichlet gave $\chi_{\mathbb{Q}}$ as an example of a function with no integral (see Section ??). Since $\chi_{\mathbb{Q}}$ is not really defined by a formula, some infer that Dirichlet had internalized the modern concept of a function, for which they therefore give him credit. But Lakatos correctly points out [9, p 151] that the credit is undeserved. Dirichlet never gave any such definition.



Figure 2.11: Continuous at precisely one point.



Figure 2.12: Differentiable at 0, but discontinuous everywhere else.

9 Conway's base-13 function

You might remember the Intermediate Value Theorem from calculus class, which says that if your position is a continuous function of time, then you can't teleport. (See Figures 2.13 and 2.14.)

Theorem 10 (Intermediate Value Theorem). Suppose f is continuous and a < b. Then for any y between f(a) and f(b), there is an $x \in (a, b)$ so that f(x) = y.

The IVT seems pretty obvious, because "no teleporting" sounds almost what we *meant* by being continuous! At least, it sounds awfully similar to Euler's pencil-never-leaves-the-paper idea of continuity... but remember, that wasn't the actual definition of continuity.

Hey, maybe now we can argue that Bolzano's formal definition of continuity successfully captures Euler's intuitive idea! A function which satisfies the *conclusion* of the IVT is called a Darboux function. That is, f is a Darboux function if for every a < b and every y between f(a) and f(b), there is an $x \in (a, b)$ so that f(x) = y. Maybe that's a definition of continuity that Euler could get behind! The IVT says that every continuous function is Darboux, so now we just have to prove that every Darboux function is continuous.

There's one small hitch: that last statement is *extremely false*! The function $f(x) = \sin(1/x)$ is Darboux, but discontinuous at one point. It gets worse. We'll give a function $f : \mathbb{R} \to \mathbb{R}$ such that for every open interval (a, b), we have $f((a, b)) = \mathbb{R}$. That is, for every open interval (a, b) and every $y \in \mathbb{R}$, there exists $x \in (a, b)$ so that f(x) = y. So f is certainly Darboux, but f is not even remotely close to continuous. In fact, it's discontinuous at every point (like the Dirichlet function, but much crazier.)

Figure 2.15 is a little misleading. The graph of f isn't *all* of \mathbb{R}^2 (it's a function, after all!) But every disc in \mathbb{R}^2 contains a point in the graph of f. In other words, the graph of f is a dense subset of \mathbb{R}^2 .

So what function has this bizarre property? One example is by British mathematician John Horton Conway, who (as of 2016) is still alive, unlike the other mathematicians we've encountered. His idea is to represent numbers in base 13, with these symbols:



Every real number has a unique base-13 expansion with no trailing (\cdot) symbols (recall Section 6.) Conway's base-13 function f is defined with respect to this expansion as follows.

- For the interesting case, suppose the base-13 expansion of x is of the form AB, where removing all the circles from the symbols in B yields a sensible base-10 expansion for a real number y. Then set f(x) = y.
- Otherwise, just set f(x) = 0.



Figure 2.13: The intermediate value theorem: in order for a continuous function to get from one value to another, it must pass through every value in between.



Figure 2.14: When you drive in a car, your distance from Wellington varies continuously. Every point on the Earth's surface which is 550 miles away from Wellington is at sea. So by the IVT, if you want to drive from New Zealand to Australia, you're going to have to build a car that can drive through water. Or a bridge or something.



Figure 2.15: A truncated graph of Conway's base-13 function (in black).

Figure 2.16: A truncated graph of Conway's base-13 function (in white).

For example, let x be the real number with base-13 expansion

$$x = \underbrace{-(+)}_{A} \underbrace{(-)}_{B} \underbrace{(-$$

Notice that if we start at the - digit, then if we removed the circles, we would get a string y = -3.1415926..., which is a base-10 expansion for the real number $-\pi$. So we set $f(x) = -\pi$. For another example, let x be a real number with infinitely many + symbols in its base-13 expansion. Then f(x) = 0. Note that in the definition of f, we require base-10 expansions to start with either + or -; this ensures that B is unique and hence f is well defined.

Proposition 7. Let f denote Conway's base-13 function. Then for every a < b and every y, there is some x such that a < x < b and f(x) = y.

Proof. Start with the base-13 expansion for the midpoint $\frac{1}{2}(a+b)$. If we go out

$$\frac{1}{2}(a+b) = \underbrace{0} \cdot \underbrace{5} - \underbrace{-1} \cdot \underbrace{8} \underbrace{1} \cdot \underbrace{0} \cdot \underbrace{5} \cdot 9 \dots \quad \text{(base 13)}$$

L

$$y = +3.1415926...$$
 (base 10)

$$x = 0.5 - 1.8 + 3.14...$$
 (base 13)

Figure 2.17: The proof of Proposition 7. The location of the vertical bar in the base-13 expansion of $\frac{1}{2}(a+b)$ is chosen based on how big b-a is, to make sure that $x \in (a, b)$.

far enough in this base-13 expansion, we can change anything we want and we'll still have a number in (a, b). So in particular, we can replace the sequence of subsequent digits with the circled base-10 expansion of y, to obtain an $x \in (a, b)$ such that f(x) = y. (See Figure 2.17.)

So Darboux functions are a lot more complicated than continuous functions. In fact, Darboux functions are absurdly "expressive":

Theorem 11 (Sierpinski). For every function $f : \mathbb{R} \to \mathbb{R}$, there are two Darboux functions g, h so that f = g + h.

Proof. Define an equivalence relation \sim on \mathbb{R} by declaring that $x \sim y$ if $x - y \in \mathbb{Q}$. Let E be the set of equivalence classes. Observe that

$$|\mathbb{R}| = |E \times \mathbb{Q}| \le |E \times E| = |E| \le |\mathbb{R}|$$

so $|E| = |\mathbb{R}|$. Partition E up into two disjoint sets $E = E_1 \cup E_2$ so that $|E_1| = |E_2| = |\mathbb{R}|$. There are bijections $\alpha_1 : E_1 \to \mathbb{R}$ and $\alpha_2 : E_2 \to \mathbb{R}$. Define

$$g(x) = \begin{cases} \alpha_1([x]) & \text{if } [x] \in E_1 \\ f(x) - \alpha_2([x]) & \text{if } [x] \in E_2; \end{cases}$$
(2.1)

$$h(x) = \begin{cases} f(x) - \alpha_1([x]) & \text{if } [x] \in E_1\\ \alpha_2([x]) & \text{if } [x] \in E_2. \end{cases}$$
(2.2)

(Here, [x] denotes the equivalence class to which x belongs.) By construction, f = g + h. To show that g and h are Darboux, we'll show that even better, they (like Conway's base-13 function) map every open interval surjectively onto \mathbb{R} . Fix a < b and y. Let $[x] = \alpha_1^{-1}(y)$. Since \mathbb{Q} is dense in \mathbb{R} , we can find $x' \sim x$ so that a < x' < b. Then $\alpha(x') = y$, showing that $g((a, b)) = \mathbb{R}$. The same argument works for h.

What's the moral of this story? Is there something wrong with Bolzano's definition of continuity? Nah. Euler would probably agree that Conway's base-13 function does not deserve to be called continuous. The notion of a Darboux

function is *not* a reasonable definition of continuity. It's hard to say what it means for the graph of a function to be "described by freely leading the hand," but it really ought to be more *conservative* than continuity, not more liberal.

10 Continuity is uncommon

In the previous couple of sections, we saw some really nasty functions with tons of discontinuities. But in "everyday life," it seems like we only run into continuous functions. You might be tempted to infer that *most* functions are continuous. But in truth, in the sense of cardinality, the *vast majority* of functions are discontinuous!

Proposition 8. Let $C(\mathbb{R}, \mathbb{R})$ be the set of all continuous functions $\mathbb{R} \to \mathbb{R}$. Then $|C(\mathbb{R}, \mathbb{R})| = |\mathbb{R}| = \beth_1$. (In contrast, note that the set $\mathbb{R}^{\mathbb{R}}$ of all functions $\mathbb{R} \to \mathbb{R}$ has cardinality \beth_2 .)

Proof. Since \mathbb{Q} is dense, to specify a continuous function $f : \mathbb{R} \to \mathbb{R}$, it suffices to give the restriction of f to \mathbb{Q} . (The value of f at any point x can be recovered from its values on \mathbb{Q} , because there's a sequence x_1, x_2, \ldots of rational numbers converging to x, and $f(x) = \lim_{n \to \infty} f(x_n)$. See Figure 2.18.) Therefore,

$$|C(\mathbb{R},\mathbb{R})| \le |\mathbb{R}|^{|\mathbb{Q}|} = (2^{\beth_0})^{\beth_0} = 2^{\beth_0 \cdot \beth_0} = 2^{\beth_0} = \beth_1.$$

Constant functions establish the reverse inequality $|C(\mathbb{R},\mathbb{R})| \geq \beth_1$.

Notice that the same basic argument actually shows that the vast majority of functions have *uncountably many* discontinuities! (To specify a function f with countably many discontinuities, just give $f \upharpoonright_{\mathbb{Q}}$ along with $f \upharpoonright_{D}$ where D is the set of x values at which f is discontinuous.)

For whatever reason, this is a recurrent phenomenon in mathematics. Usually, the vast majority of cases are pathological (in appropriate senses of "vast majority" and "pathological.")



Figure 2.18: To recover the full graph of f given the values of f on \mathbb{Q} , just connect the dots.

11 Thomae's function

We've seen some very discontinuous functions. But bigger is not always better. Maybe you're especially fond of some set $D \subseteq \mathbb{R}$. Like discontinuity connoisseurs, we can look for a function which is discontinuous exactly at the x values in D. For now, let's consider the case $D = \mathbb{Q}$. In the 19th century, the German mathematician Carl Johannes Thomae devised his namesake function:

$$f(x) = \begin{cases} 0 & \text{if } x \text{ is irrational} \\ \frac{1}{q} & \text{if } x = \frac{p}{q}, \text{ with } \frac{p}{q} \text{ reduced and } q > 0. \end{cases}$$
(2.3)

(See Figure 2.19.)

Proposition 9. Thomae's function is continuous at irrational x and discontinuous at rational x.

Proof. First, suppose x is rational, so that f(x) > 0. There's a sequence of irrational numbers x_1, x_2, \ldots converging to x, but $f(x_n) = 0 \neq f(x)$, so f is discontinuous at x.

Conversely, for the harder direction, suppose x is irrational. The intuition here is to think about rational approximations to x, and notice that *close approximations must have large denominators*. So as x' gets very close to x, f(x')really will get very close to 0. Now for the proof:

Consider an arbitrary sequence x_1, x_2, \ldots converging to x, and fix an arbitrary $\varepsilon > 0$. There are only finitely many rational numbers within distance 1 of x with denominator no more than $\frac{1}{\varepsilon}$, so one of them (call it y) is closest to x. If n is sufficiently large, the sequence x_n is even closer to x than y is, and hence $f(x_n) < \varepsilon$. Since ε was arbitrary, $f(x_n) \to 0$, showing that f is continuous at x.



Figure 2.19: Thomae's function (in black).

12 Discontinuities of monotone functions

A function f is monotone increasing if $x \leq x'$ implies $f(x) \leq f(x')$. Monotone decreasing is defined in the obvious way, and monotone just means either monotone increasing or monotone decreasing. (See Figure 2.20.) The pathological functions we've seen so far have *not* been monotone. The following theorem, due to Darboux despite its name, gives an excuse:

Theorem 12 (Froda's theorem). Suppose f is monotone. Then f has only countably many discontinuities.

The key to proving Theorem 12 is a recharacterization of continuity.

Definition 12. Fix $f : \mathbb{R} \to \mathbb{R}$ and $c \in \mathbb{R}$. We write $\lim_{x \to c} f(x) = L$ to mean that for every $\varepsilon > 0$, there exists $\delta > 0$ such that

$$0 < |x - c| < \delta \implies |f(x) - L| < \varepsilon.$$

(It's basically like the definition of the limit of a sequence, with x playing the role of n and δ playing the role of N.) If you just check, you'll see that f is continuous at c if and only if $\lim_{x\to c} f(x) = f(c)$. So now we can divide the crime of discontinuity into three tiers, depending on how badly $\lim_{x\to c} f(x)$ fails to equal f(c) (see Figure 2.21):

1. Suppose $\lim_{x\to c} f(x)$ exists, but it doesn't equal f(c). Then f is charged with having a *removable discontinuity* at c. For this minor infraction, f is required to enroll in a 12-step program, where it learns how to change its value at c and thereby become continuous.



Figure 2.20: A monotone increasing function on the left and a monotone decreasing function on the right.



Figure 2.21: Removable discontinuity, jump discontinuity, and essential discontinuity.

- 2. The left limit, denoted $\lim_{x\to c^-} f(x)$ or f(c-), is defined just like $\lim_{x\to c} f(x)$, except we only pay attention to x < c. Similarly for right limits. Suppose f(c-) and f(c+) both exist, but they're not equal, and hence $\lim_{x\to c} f(x)$ doesn't exist. Then f is charged with having a *jump discontinuity* at c. For this misdemeanor, f is incarcerated in a correctional facility, where professionals attempt to decrease f(x) for all x on one side of c, thereby eliminating the jump and restoring continuity.
- 3. Finally, suppose either f(c-) or f(c+) does not exist. Then f is charged with having an *essential discontinuity* at c, which is a felony. Making f continuous at c would require fundamentally altering f's character. So f is just sentenced to life imprisonment, to protect society from its incorrigible, deviant behavior.

Proof of Froda's theorem. Without loss of generality, assume f is monotone *increasing*. Suppose f is discontinuous at $c \in \mathbb{R}$. Monotonicity implies that it's a jump discontinuity. Since f(c-) < f(c+), there is some rational number q_c with $f(c-) < q_c < f(c+)$. The map $c \mapsto q_c$ is an injection from the set of discontinuities of f to \mathbb{Q} .

Time for a fun application of Froda's theorem in the form of a game. We choose two distinct real numbers $x_1 < x_2$ and put each in an unmarked envelope.



Figure 2.22: You open the right envelope and see 10^6 . Do you guess that $x_1 = 10^6$ or $x_2 = 10^6$? Does 10^6 seem like a small number, or a big number? What a dumb question. Surely, all you can do is toss a coin and hope for the best... right? Nope!



Figure 2.23: The strategy which gives you a win probability greater than 50%. The area of the green region is the probability that y really does fall between x_1 and x_2 , in which case you win. If the blue or yellow event occurs, you'll win if and only if you open the envelope containing x_2 or x_1 , respectively.

We shuffle the envelopes and give them to you. You choose an envelope and open it, learning the real number inside. You then guess whether you're looking at x_1 or x_2 . If you're right, you win. If you're wrong, we win. (See Figure 2.22.)

You can trivially achieve a win probability of 50% by just opening a random envelope and saying " x_1 ." Bizarrely, you can *beat* 50%. Here's what you do: Pick your own third number y randomly, from (say) a standard normal distribution. Choose a random envelope and open it. Assume that y falls between x_1 and x_2 , and guess accordingly.

Here's why it works: Whatever values x_1, x_2 we choose, there's a positive probability that y falls between them. In that case, you'll win. And in the other case, it all depends on which envelope you open, so you've still got a 50-50 shot. So overall, your probability of winning is

$$Pr(win) = Pr(y \in (x_1, x_2)) \cdot 1 + Pr(y \notin (x_1, x_2)) \cdot 0.5$$

= 0.5 + 0.5 \cdot Pr(y \in (x_1, x_2)) > 0.5.

(See Figure 2.23.)

Admittedly, it's a bit anticlimatic. The strategy beats 50%, but only by ε , where we can force ε to be as small as we want by choosing x_1 and x_2 very close together. Maybe there's a cleverer strategy which guarantees you probability of success at least p where p > 0.5?

Nope! Here's why. Fix an arbitrary strategy. Let f(x) be the probability that you guess " x_2 " given that you observed the number x in the envelope you opened. Then your probability of success is

$$Pr(win) = 0.5 \cdot f(x_2) + 0.5 \cdot (1 - f(x_1)) = 0.5 + 0.5[f(x_2) - f(x_1)].$$

If f is not monotone increasing, we can choose $x_1 < x_2$ so that $f(x_1) > f(x_2)$, putting your win probability below 50%. If f is monotone, then by Froda's theorem, it has a point of continuity, so we can force $f(x_2) - f(x_1)$ to be smaller than whatever $\varepsilon > 0$ we choose.

How about the converse to Froda's theorem? Yep, countability *characterizes* the sets of discontinuities of monotone functions!

Theorem 13. Suppose $D \subseteq \mathbb{R}$ is countable. Then there is some monotone function $f : \mathbb{R} \to \mathbb{R}$ which is discontinuous precisely at points in D.

Proof. Say $D = \{d_1, d_2, ...\}$. Define

$$f(x) = \sum_{d_i \le x} 2^{-i}.$$
 (2.4)

(See Figure 2.24.) The sum makes sense, because the terms are all nonnegative, so the order of summation doesn't matter. The sum converges to a finite number between 0 and 1, since $\sum_{i=1}^{\infty} 2^{-i} = 1$. It's immediate that f is monotone increasing; as x gets bigger, we add up more and more things. And of course f is discontinuous at $d_i \in D$, because the value jumps up by 2^{-i} there.

Finally, fix $x \notin D$; we must show that $\lim_{x'\to x} f(x') = f(x)$. Consider an arbitrary $\varepsilon > 0$. Let N be large enough that $\sum_{i=N}^{\infty} 2^{-i} < \varepsilon$. Let δ be small enough that the interval $[x - \delta, x + \delta]$ doesn't contain any of the first N elements of D. Then while traversing this interval $[x - \delta, x + \delta]$, the value of f changes by at most $\sum_{i=N}^{\infty} 2^{-i} < \varepsilon$ as desired.

Notice that this provides another example of a function which is discontinuous at exactly the rationals (like Thomae's function.) But this time, it's *monotone*! We'll revisit the construction in the proof of Theorem 13 in Section ?? after developing measure theory, and hopefully it will seem more natural then.



Figure 2.24: The function f used to prove Theorem 13 in the case $D = \mathbb{N}$.

13 Discontinuities of indicator functions

For a function f, let D(f) be the set of x values such that f is discontinuous at x. So far, in every example we've seen, D(f) has either been countable or else has contained an interval. Can D(f) be an uncountable nowhere dense set? E.g., is there a function with $D(f) = \Delta$, where Δ is the Cantor set from Section 6? Yep! Oddly enough, the indicator function χ_{Δ} is an example! This is in contrast to the situation with \mathbb{Q} , whose indicator function is discontinuous everywhere.

Proposition 10. χ_{Δ} is discontinuous precisely at Δ .

Proof. First, suppose $x \notin \Delta$. Remember that to construct Δ , we just removed a bunch of open intervals from [0, 1], so there is some open interval I such that $x \in I$ and $I \cap \Delta = \emptyset$. Then χ_{Δ} is 0 on all of I, so it is continuous at x.

Conversely, suppose $x \in \Delta$. Remember that Δ is nowhere dense, so in particular, Δ does not contain any intervals. Therefore, there are points arbitrarily close to x which are not in Δ , where χ_{Δ} is 0. Therefore, χ_{Δ} is discontinuous at x.

Let's generalize, so we can understand what just happened. It's time to introduce you to topology. The definitions are a bit more intuitive in \mathbb{R}^n . For a point $x \in \mathbb{R}^n$ and a radius r > 0, let $B_r(x)$ denote the open ball of radius r centered at x.

Definition 13 (Interior, exterior, boundary). Fix a set $E \subseteq \mathbb{R}^n$ and a point $x \in \mathbb{R}^n$.

• If there's some $\varepsilon > 0$ so that $B_{\varepsilon}(x) \subseteq E$, we say that x is an *interior point* of E.



Figure 2.25: Let E denote the gray region. Then x is an interior point of E, y is an exterior point of E, and z is a boundary point of E.

- If there's some $\varepsilon > 0$ so that $B_{\varepsilon}(x) \subseteq E^c$, we say that x is an *exterior* point of E.
- If x is neither an interior point nor an exterior point of E, we say that x is a *boundary point* of E.

(See Figure 2.25.)

The *interior* of E, denoted int(E), is the set of interior points of E. The *exterior* of E is denoted ext(E), and the *boundary* of E is denoted ∂E . For example, if I is an interval from 0 to 1, then regardless of which endpoints are included, we have int(I) = (0, 1), $\partial I = \{0, 1\}$, and $ext(I) = (-\infty, 0) \cup (1, \infty)$. A couple other examples of boundaries: $\partial \mathbb{R} = \emptyset$, $\partial \mathbb{Q} = \mathbb{R}$, $\partial \mathbb{Z} = \mathbb{Z}$, and $\partial \Delta = \Delta$.

Proposition 11. For any set $E \subseteq \mathbb{R}$, χ_E is discontinuous precisely on ∂E .

Proof. A point x is in ∂E if and only if there points arbitrarily close to x in E and points arbitrarily close to x in E^c .

14 Sets of discontinuities

Does Thomae's function have a twin? That is, does there exist a function which is continuous at rational points and discontinuous at irrational points?

As in the last section, D(f) is the set of points where f is discontinuous. We've seen examples of messed up functions with $D(f) = \mathbb{R}$, $D(f) = \mathbb{Q}$, $D(f) = \Delta$, etc. We have not, however, seen any hints about how you might *rule out* the possibility of a function f with some given discontinuity set.

In Section 12, we saw a satisfying theorem: There exists a *monotone* function f such that D(f) = D if and only if D is countable. In this section, we'll prove an analogous theorem without the "monotone" qualifier.

Definition 14. Fix a set $E \subseteq \mathbb{R}^n$. We say that E is *closed* if $\partial E \subseteq E$. We say that E is *open* if $\partial E \subseteq E^c$.



Figure 2.26: Adolf Hitler does not appreciate the terms "open" and "closed" [6].

For example, thankfully, open intervals are open and closed intervals are closed. An open set is one where each point has some wiggle room. "Fuzzy set" probably would have been a better term for open sets. The term "closed set" is more reasonable, because a set E is (topologically) closed if and only if it is closed under the operation of *taking limits*. That is, E is closed if and only if whenever x_n is a convergent sequence of points in E, we have $\lim x_n \in E$. Warning: some sets, like [0, 1), are neither open nor closed, and other sets, like \emptyset , are both open and closed. (See Figure 2.26.)

If $E = \partial E$, like the case $E = \Delta$, then there's a function f with D(f) = E, namely $f = \chi_E$. By adapting Dirichlet's simple trick, we can handle all closed sets, even the ones with nonempty interiors.

Proposition 12. Suppose $E \subseteq \mathbb{R}$ is closed. Then there is some function f with D(f) = E.

Proof. Define

 $f(x) = \begin{cases} 1 & \text{if } x \in E \cap \mathbb{Q} \\ -1 & \text{if } x \in E \setminus \mathbb{Q} \\ 0 & \text{if } x \notin E. \end{cases}$

(See Figure 2.27.) This is obviously continuous on $x \notin E$, because there's a neighborhood around x on which f is constant. Conversely, suppose $x \in E$, so that $f(x) \neq 0$. Let x_n, y_n be sequences converging to x with $x_n \in \mathbb{Q}, y_n \notin \mathbb{Q}$. Then $f(x_n)$ is nonnegative and $f(y_n)$ is nonpositive, so they can't both converge to f(x).

(An alternative way to prove Proposition 12 is to show that every closed subset of \mathbb{R} is the boundary of some set.) How about the converse? Do we have our characterization – are discontinuity sets precisely closed sets? Nah, that hypothesis has already been falsified. For example, \mathbb{Q} is not closed, but it's the discontinuity set of Thomae's function. The real criterion is slightly more complicated.



Figure 2.27: The function used to prove Proposition 12 in the case E = [-1, 1].

Definition 15. A set $E \subseteq \mathbb{R}$ is F_{σ} if it can be written as a countable union of closed sets.

(The term F_{σ} comes from the French words "fermé" and "somme," meaning "closed" and "union.") For example, any countable set, like \mathbb{Q} , is F_{σ} , because singleton sets are closed. Any closed set, like Δ or \mathbb{R} , is trivially F_{σ} . The set $\mathbb{R} \setminus \{0\}$ is F_{σ} , because

$$\mathbb{R} \setminus \{0\} = \bigcup_{n \in \mathbb{N}} \left(-\infty, -\frac{1}{n}\right] \cup \left[\frac{1}{n}, \infty\right).$$

Notice that every set of discontinuities that we've encountered so far is F_{σ} ! This is no coincidence. Using the basic idea behind Thomae's function, we can tweak the proof of Proposition 12 to handle arbitrary F_{σ} sets.

Theorem 14. Suppose E is F_{σ} . Then there exists a function $f : \mathbb{R} \to \mathbb{R}$ with D(f) = E.

Proof. Say $E = \bigcup_n E_n$, where each E_n is closed. Define

$$f(x) = \begin{cases} \max\{\frac{1}{n} : x \in E_n\} & \text{if } x \in E \cap \mathbb{Q} \\ -\max\{\frac{1}{n} : x \in E_n\} & \text{if } x \in E \setminus \mathbb{Q} \\ 0 & \text{if } x \notin E. \end{cases}$$

(See Figure 2.28.) First, suppose $x \in E$. The proof used for Proposition 12 still applies, showing that f is discontinuous at x. Conversely, suppose $x \notin E$, so f(x) = 0. Suppose $x_m \to x$. Since each E_n is closed, the sequence x_m must eventually escape E_n and never return. Once x_m has escaped E_1, \ldots, E_n , we have $|f(x_m)| \leq \frac{1}{n}$. So $f(x_m) \to 0$, and f is continuous at x.

And the converse is also true: F_{σ} -ness characterizes discontinuity sets.

Theorem 15. For any function $f : \mathbb{R} \to \mathbb{R}$, the set D(f) is F_{σ} .

Remember how the key to Froda's theorem was to classify discontinuities as more or less severe? That's true here too, in a slightly different sense. The



Figure 2.28: The function used to prove Theorem 14 in the case $E_n = [\frac{1}{n}, 3 - \frac{1}{n}]$, which is discontinuous precisely on E = (0, 3).



Figure 2.29: The oscillation of f in E is the height of the smallest box that contains the graph of the restriction of f to E. For example, the oscillation of $\sin(1/x)$ in any interval containing 0 is 2.

diameter of a set $E \subseteq \mathbb{R}$ is defined by

$$\operatorname{diam}(E) = \sup_{x,y \in E} |x - y|.$$

For a function $f : \mathbb{R} \to \mathbb{R}$ and a set $E \subseteq \mathbb{R}$, the *oscillation* of f in E is defined by $\omega_f(E) = \text{diam}(f(E))$. (See Figure 2.29.) The oscillation of f at a *point* x is defined by

$$\omega_f(x) = \lim_{\varepsilon \to 0} \omega_f(B_\varepsilon(x)).$$

The oscillation of f at x measures the extent to which f is discontinuous at x. For example, if f has a removable discontinuity at x, then $\omega_f(x)$ is the distance from the actual value f(x) to the better value $\lim_{x'\to x} f(x')$. In particular, $\omega_f(x) = 0$ if and only if f is continuous at x. Proof sketch of Theorem 15. We can write

$$D(f) = \bigcup_{n \in \mathbb{N}} \left\{ x \in \mathbb{R} : \omega_f(x) \ge \frac{1}{n} \right\}.$$
 (2.5)

If you just check, you'll see that $\{x : \omega_f(x) \ge \varepsilon\}$ is a closed set.

Theorems 14 and 15 help a lot toward understanding which sets are discontinuity sets. For example, the vast majority of sets are *not* discontinuity sets.

Proposition 13. Let \mathcal{D} denote the set of all F_{σ} subsets of \mathbb{R} . Then $|\mathcal{D}| = |\mathbb{R}|$ (which is smaller than $|\mathcal{P}(\mathbb{R})|$ by Cantor's theorem.)

Proof sketch. It turns out that every open set $U \subseteq \mathbb{R}$ can be written as a countable union of disjoint open intervals. A closed set is just a complement of an open set, so a closed set can be specified by a sequence of real numbers. Hence, an arbitrary element of \mathcal{D} is specified by a sequence of sequences of reals. Therefore,

$$|\mathcal{D}| \le (|\mathbb{R}|^{|\mathbb{N}|})^{|\mathbb{N}|} = |\mathbb{R}|^{|\mathbb{N}|} = |\mathbb{R}|.$$

But the story so far isn't entirely satisfying, because it's not obvious how to identify examples of sets which are not F_{σ} . Can the set $\mathbb{R} \setminus \mathbb{Q}$ be written as a countable union of closed sets? It's difficult to say! (That's the thing about characterization theorems. You're never really sure when you're done.) Stay tuned, we'll answer this question in Section 15.

15 The Baire category theorem

Our goal in this section is to prove that Thomae's function does *not* have a twin. That is, $\mathbb{R} \setminus \mathbb{Q}$ is not F_{σ} . On the way, we'll meet the *meager sets*. Meagerness might seem like a technical, awkward concept. At the very least, it's a useful tool. And meager sets are actually pretty fun to hang out with, once you get to know them.

In Section 2, we saw Cantor's famous 1891 diagonal argument, which proved that \mathbb{R} is uncountable. Diagonalization is a great trick to have up your sleeve; we saw in Section ?? that it can be used to prove that $|\mathcal{P}(S)| > |S|$ for every set S. Historically, diagonalization was not the first technique used to prove that \mathbb{R} is uncountable. Let's take a look at (a slight variant of) Cantor's *original* proof that \mathbb{R} is uncountable, from 1874. The older proof is actually more real-analysisish than the slick diagonalization trick, and if you understand the proof, you'll be ready to meet meager sets. A set $E \subseteq \mathbb{R}$ is *bounded* if diam $(E) < \infty$.

Theorem 16 (Cantor's intersection theorem). Suppose $E_1 \supseteq E_2 \supseteq \ldots$ is a nested sequence of closed, bounded, nonempty sets. Then $\cap_n E_n \neq \emptyset$.

Proof. Closedness implies that each E_n has a minimum $x_n = \min E_n$. Then x_n is a bounded, monotone increasing sequence, so it has a finite limit $x = \sup_n x_n$. For any E_m , the sequence x_n is eventually in E_m , so closedness implies that $x \in E_m$. (See Figure 2.31.)



Figure 2.30: The hypotheses of Cantor's intersection theorem are important. On the left, the nested sequence of closed, nonempty, *unbounded* sets $E_n = [n, \infty)$ has empty intersection. On the right, the nested sequence of *open*, nonempty, bounded sets $E_n = (0, \frac{1}{n})$ has empty intersection.



Figure 2.31: The proof of Cantor's intersection theorem in the case $E_n = [-\frac{1}{n}, \frac{1}{n}]$. The left endpoints limit to 0, the sole element of $\bigcap_n E_n$.

Theorem 17. \mathbb{R} is uncountable.

Proof from 1874. Let x_1, x_2, \ldots be an arbitrary sequence. Inductively define closed, bounded intervals $I_1 \supseteq I_2 \supseteq \ldots$, with $x_n \notin I_n$. (See Figure 2.32.) By Cantor's intersection theorem, there is some $x \in \bigcap_n I_n$. Then $x \neq x_n$ for every n. Since the sequence was arbitrary, we can conclude that no sequence exhausts all of \mathbb{R} .



Figure 2.32: Cantor's original proof that \mathbb{R} is uncountable. Having already defined I_{n-1} (in black), we can find a subinterval I_n (in blue) which misses the single point x_n (in red.)

Uncountability is a sort of *bigness*. In 1899, René-Louis Baire realized³ that by tweaking Cantor's proof, we can show that \mathbb{R} is "big" in a stronger sense

³Dunno if this was actually Baire's thought process. But it's a reasonable guess.

than mere uncountability. Recall from Section 6 that a set $E \subseteq \mathbb{R}$ is nowhere dense if for every open interval $I \subseteq \mathbb{R}$, there is an open subinterval $J \subseteq I$ so that $E \cap J = \emptyset$. In the proof that \mathbb{R} is uncountable, we avoided the sequence of points x_1, x_2, \ldots , but the argument actually allows us to avoid a sequence of sets E_1, E_2, \ldots , as long as each E_n is nowhere dense. This idea led Baire to classify subsets of \mathbb{R} as falling into two "categories."

Definition 16. A set $E \subseteq \mathbb{R}$ is *meager*, or *first category*, if it can be written as a countable union of nowhere dense sets.

You should think of "meager" as meaning "small" (though this is a more relaxed sense of smallness than nowhere dense or countable.) Sometimes, people describe meager sets as "thin." Some examples: Any countable set (like \mathbb{Q}) is meager, because a singleton set $\{x\}$ is nowhere dense. Any nowhere dense set (like Δ) is meager. Let $\Delta + \mathbb{Q} = \{\delta + q : \delta \in \Delta, q \in \mathbb{Q}\}$. (In words, we put a copy of Δ at every rational number. This is called the *Minkowski sum* of Δ and \mathbb{Q} .) Then $\Delta + \mathbb{Q}$ is meager, even though it's uncountable and dense.

Definition 17. A set $E \subseteq \mathbb{R}$ is *nonmeager*, or *second category*, if it isn't meager.



Figure 2.33: Baire categories.

Theorem 18 (The Baire category theorem). \mathbb{R} is nonmeager.

Proof. Just repeat the proof that \mathbb{R} is uncountable, replacing the sequence x_1, x_2, \ldots of real numbers with a sequence E_1, E_2, \ldots of nowhere dense sets. \Box

The Baire category theorem opens the door to a host of nonmeager sets. A *comeager* set is the *complement* of a meager set.⁴ Since the union of two meager sets is meager, Baire's category theorem implies that every comeager set, such as $\mathbb{R} \setminus \mathbb{Q}$, is nonmeager.

We promised that all this Baire category stuff was going to help us to show that $\mathbb{R} \setminus \mathbb{Q}$ is not F_{σ} . Maybe F_{σ} sets are always meager? Nah, that's not true. There's really no reasonable sense in which F_{σ} sets are "small," because \mathbb{R} itself is F_{σ} ! The true connection is a little subtler: F_{σ} sets are *either* small, or big, but never medium! Precisely:

Proposition 14. Suppose E is F_{σ} . Then either E is meager ("E is small") or else E contains an interval ("E is big.") In particular, $\mathbb{R} \setminus \mathbb{Q}$ is not F_{σ} .

Proof. Say $E = \bigcup_n E_n$, where every E_n is closed. If E doesn't contain an interval, then each E_n is nowhere dense (if it were dense in I, it would contain I by closedness.) So E is meager.

A couple other examples: The set of transcendental numbers is not F_{σ} . The set of noncomputable numbers is not F_{σ} . $(\Delta + \mathbb{Q})^c$ is not F_{σ} .

References

- Bernard Bolzano. Rein analytischer Beweis des Lehrsatzes daß zwischen je zwey Werthen, die ein entgegengesetzetes Resultat gewähren, wenigstens eine reelle Wurzel der Gleichung liege. Gottlieb Haase, 1817.
- [2] Nicolas Bourbaki. Berlin: Springer, 2006. ISBN: 3540340343.
- [3] P. G. L. Dirichlet. "On the convergence of trigonometric series which serve to represent an arbitrary function between two given limits". In: 4 (1829), pp. 157–169.
- [4] L. Euler. Introductio in analysin infinitorum. Vol. 1. 1748.
- [5] Christopher S. Hardin and Alan D. Taylor. "A peculiar connection between the Axiom of Choice and predicting the future". In: American Mathematical Monthly (2008), pp. 91–96.
- [6] Hitler Learns Topology. URL: https://www.youtube.com/watch?v= SyD4p8_y8Kw.
- [7] Israel Kleiner. "Evolution of the function concept: A brief survey". In: *The College Mathematics Journal* 20.4 (1989), pp. 282–300.

 $^{^{4}}$ The "co-" prefix convention for complements is especially popular in mathematical logic. It's useful for constructing low-quality jokes. E.g., a coconut should have just been called a nut.

- [8] Steven G Krantz. Mathematical apocrypha: stories and anecdotes of mathematicians and the mathematical. MAA, 2002.
- [9] Imre Lakatos. Proofs and refutations : the logic of mathematical discovery. Cambridge New York: Cambridge University Press, 1976. ISBN: 0521290384.
- [10] Set Theory and Weather Prediction. URL: http://xorshammer.com/ 2008/08/23/set-theory-and-weather-prediction/.

Chapter 3

Series

On the whole, divergent series are the work of the devil, and it's a shame that one dares base any demonstration upon them.

> Niels Henrik Abel [6], possibly mistranslated

16 Stacking books

How far over the edge of a table can a stack of books protrude without toppling? (Figure 3.1)



Figure 3.1: The book stacking problem with N = 4. We are interested in maximizing the overhang d.

To be more precise, we have N identical unit-length books, and in our stack, no two books may have the same vertical position. How large can the horizontal distance be between the edge of the table and the right edge of the rightmost book? For N = 1, you can achieve an overhang of $d = \frac{1}{2}$, but if you push the book any farther it will fall off the table. (Figure 3.2)

What about N > 1? So far, this is a physics question, but we can turn it into a math question by trusting Newton: assume that the stack falls if and only if for some $n \leq N$, the center of mass (COM) of the top n books is *not* above the surface on which those n books rest. (Figure 3.3)



Figure 3.2: An optimal "stack" of 1 book.

If you're in the mood to solve this puzzle yourself, close this book now and ponder. Otherwise, read on for the solution.



Figure 3.3: The stack on the left is unbalanced and will topple over. The COM of the entire stack (marked \times) is over the table like it should be, but the COM of the top two books (marked \circ) is to the right of the third book. The top two books will pivot about the top right corner of the third book, as shown on the right. Note: We assume that the books have uniform density, so the COM of a set of books is just the average of their spatial centers.

Definition 18. The *harmonic stack* is defined inductively as follows. To build a harmonic stack of N books, first place a book on the table poking over the edge a distance of $\frac{1}{2N}$. Then build a harmonic stack of N - 1 books, treating that first book as if it were the table. For example, the harmonic stack of 4 books is depicted in Figure 3.4.



Figure 3.4: A harmonic stack of 4 books.

Proposition 15. Harmonic stacks do not topple over.

Proof. We proceed by induction¹ on N, the number of books in the stack. The case N = 1 is trivial. Now consider N > 1. By induction, we can assume that if

¹Never seen a proof by induction before? Fear not, this is a great introductory example. Mathematical induction (not to be confused with "inductive reasoning") is a technique for proving that for every $N \in \mathbb{N}$, blab blab blab. Here's the idea: Start by showing that your theorem is true for N = 1 (this is the "base case.") Then show that your theorem being true
you held the bottom book steady, the stack wouldn't fall over. So we just need to verify that the COM of the whole stack is over the table. Put the origin at the lower right corner of the bottom book, so that the horizontal COM of the top N-1 books is at most 0 (by induction.) Hence, the horizontal COM of all the books is at most the contribution from the bottom book, namely $-\frac{1}{2} \cdot \frac{1}{N}$. By our choice of coordinate system, that's precisely the location of the right edge of the table. (Figure 3.5)



Figure 3.5: The proof of Proposition 15 in the case N = 4. The unshaded books form a harmonic stack of 3, so we can assume we've already proven that they won't fall. So their COM (marked \times) is not to the right of the *y* axis. The shaded book is sufficiently far to the left that this implies that the COM of the entire stack (marked \circ) is not to the right of the edge of the table.

The overhang achieved by the harmonic stack of N books is

$$d = \frac{1}{2} \sum_{n=1}^{N} \frac{1}{n}$$

You should recognize the famous harmonic series. (See Figure 3.6.)



Figure 3.6: The harmonic series is not to be confused with the harmonica series.

A series is an expression² of the form $\sum_{n=1}^{\infty} a_n$, where a_1, a_2, \ldots is a sequence of real numbers (the *terms* of the series.) The sequence of partial sums of the series is the sequence S_1, S_2, \ldots where $S_N = \sum_{n=1}^{N} a_n$. We say that the series

for some value of N makes it also true for N + 1 (this is the "inductive step.") And then it's QED o'clock! The theorem being true for N = 1 makes it true for N = 2, which makes it true for N = 3, which makes it true for N = 4, etc. Basically, you make your theorem prove itself – it's a mathematical bootstrapping maneuver.

For N = 3, which makes it the for N = 4, etc. Dasteally, you make your theorem prove itself - it's a mathematical bootstrapping maneuver. ²Notice that strictly speaking, two series $\sum_{n=1}^{\infty} a_n$ and $\sum_{n=1}^{\infty} b_n$ are equal only if they are equal termwise, i.e. only if $a_n = b_n$ for every n. But typically, when $\sum_{n=1}^{\infty} a_n$ appears in a mathematical expression, it stands for the value of the series, $\lim_{N\to\infty} \sum_{n=1}^{N} a_n$, rather than for the series itself. So for example, even though $\sum 2^{-n}$ and $\sum 2 \cdot 3^{-n}$ are two different series, we still write $\sum 2^{-n} = \sum 2 \cdot 3^{-n} = 1$.

converges/diverges if the sequence of partial sums converges/diverges. Series can diverge because the limit is infinite, e.g. $1+1+1+\cdots$, or because the limit does not exist, e.g. $1-1+1-1+\cdots$.

Theorem 19. The harmonic series $\sum_{n=1}^{\infty} \frac{1}{n}$ diverges, i.e.

$$\lim_{N \to \infty} \sum_{n=1}^{N} \frac{1}{n} = \infty.$$



Figure 3.7: Other than his proof that the harmonic series diverges, Oresme's main contribution to the world may have been the invention of bar charts.

Proof. This proof was discovered by the philosopher Nicole Oresme in the 1300s. (Figure 3.7) We'll make the series a little smaller, and show that it *still* diverges. Replace each term with the next power of two to appear:

$$\sum_{n=1}^{\infty} \frac{1}{n} = \frac{1}{1} + \frac{1}{2} + \frac{1}{3} + \frac{1}{4} + \frac{1}{5} + \frac{1}{6} + \frac{1}{7} + \frac{1}{8} + \frac{1}{9} + \cdots$$
$$\geq \frac{1}{1} + \frac{1}{2} + \underbrace{\frac{1}{4} + \frac{1}{4}}_{1/2} + \underbrace{\frac{1}{8} + \frac{1}{8} + \frac{1}{8} + \frac{1}{8} + \frac{1}{8} + \frac{1}{16} + \frac{1}{16} + \cdots$$

Grouping together like powers of two gives

$$\sum_{n=1}^{\infty} \frac{1}{n} \ge \frac{1}{1} + \frac{1}{2} + 2 \cdot \frac{1}{4} + 4 \cdot \frac{1}{8} + 8 \cdot \frac{1}{16} + \dots$$
$$= \frac{1}{1} + \frac{1}{2} + \frac{1}{2} + \frac{1}{2} + \frac{1}{2} + \frac{1}{2} + \dots$$
$$= \infty.$$

(See Figure 3.8.)



Figure 3.8: The proof that $\frac{1}{2} + \frac{1}{3} + \frac{1}{4} + \ldots$ diverges. We divide the infinitely many terms of the series into blocks, and alternatingly color the blocks gray and red. Each block has only finitely many terms (twice as many as the previous block) yet each block has a total width of at least $\frac{1}{2}$.

The implications for book stacking are astounding. The overhang of the harmonic stack of N books limits to ∞ as $N \to \infty$! So for any distance d, no matter how large – a mile, a million miles, $10^{10^{10^{10}}}$ miles, anything – you could, in principle, build a precariously balanced stack of books which hangs a distance d over the edge of the table! (Figures 3.9, 3.10)

However, the harmonic series diverges *slowly*. Oresme's proof that the harmonic series diverges suggests that $\sum_{n=1}^{N} \frac{1}{n}$ scales like log N. In fact, it turns out that there is a constant $\gamma \approx 0.58$ called the Euler-Mascheroni constant such that $\sum_{n=1}^{N} \frac{1}{n} \approx \gamma + \ln N$ in the sense that

$$\lim_{N \to \infty} \left(\left[\sum_{n=1}^{N} \frac{1}{n} \right] - \ln N \right) = \gamma.$$

(Figures 3.11, 3.12).

To paraphrase³ the number theorist Daniel Shanks, $\ln N$ goes to infinity with great dignity. Turning things around, the number of books you'd need to achieve an overhang of d using a harmonic stack grows very rapidly with d; it scales like e^{2d} . Even for smallish distances like d = 30, you would need far more books than can be found on Earth.

So the harmonic stack isn't as exciting as it seemed. Unfortunately, the harmonic stack is optimal:

Proposition 16. The maximum overhang that can be achieved by a stack of N books is that achieved by the harmonic stack of N books.

(The proof, which simply amounts to moving all the centers of mass as far right as possible, is omitted.) One way to get around this annoyance is to relax the model by allowing multiple books at each vertical position, side by side. (Figure 3.13.) It turns out that in this new model, the number of books needed to reach a distance d scales like d^3 instead of like e^{2d} [4]. Much more practical.

Finally, we'll address two misconceptions about book stacking. Misconception one: Some people mistakenly summarize our discussion of harmonic stacks

³The original quote: " $\log \log \log x$ goes to infinity with great dignity."



Figure 3.9: A harmonic stack of 52 books, which achieves an overhang of about 2.27.



Figure 3.10: You can get near the theoretical optimal overhang with a deck of 52 playing cards.



Figure 3.11: Euler discovers the Euler-Mascheroni constant.



Figure 3.12: The harmonic stack is shaped like the exponential function (or the natural log function if your head is sideways.) The distance marked *a* is approximately $\gamma/2$, where γ is the Euler-Mascheroni constant.



Figure 3.13: When you allow books to be side by side (unlike our original problem), new possibilities open up. A harmonic stack of 9 books achieves an overhang of $d \approx 1.41$, but this simple "diamond" stack of 9 books achieves a superior overhang of d = 1.5.

by saying, "You can build a stack of books that reaches *infinitely far* away from the table." But "infinitely far" is much different than "arbitrarily far". (What physics is even supposed to apply to an *infinite stack* of books?)

Misconception two: Some people mistakenly believe that you can add books to the top of an ever-growing stack, one by one, in such a way that the overhang goes to ∞ as time progresses. Our discussion of harmonic stacks did not prove this claim; notice that to get from a harmonic stack of N books to a harmonic stack of N + 1 books, you have to add another book to the *bottom* of the stack! And in fact, in the model where no two books can have the same vertical position, the claim is false. Since this point is a little bit subtle, and it isn't discussed anywhere outside this book to the best of our knowledge, we give a fairly detailed statement and proof in Appendix 1.

17 Inserting parentheses and rearranging series

In Oresme's proof that the harmonic series diverges, there was a step where we grouped together like powers of two:

 $\frac{1}{1} + \frac{1}{2} + \frac{1}{4} + \frac{1}{4} + \frac{1}{8} + \frac{1}{8} + \frac{1}{8} + \frac{1}{8} + \frac{1}{8} + \cdots = \frac{1}{1} + \frac{1}{2} + 2 \cdot \frac{1}{4} + 4 \cdot \frac{1}{8} + \cdots$

Seems pretty true. But how can we legitimately justify this move? Effectively, we are *inserting parentheses* into the series, so that e.g. the two terms $\frac{1}{4}, \frac{1}{4}$ in the original series are replaced with a single term $(\frac{1}{4} + \frac{1}{4})$ in the new series.

This smells like the familiar *associative law* for addition, which says that we can insert and remove parentheses in *finite* sums without changing the value, e.g. a + (b + c) = (a + b) + c. Does associativity still hold for infinite sums (series)?

Nope! For an easy counterexample, let's look at Grandi's series,

$$1 - 1 + 1 - 1 + 1 - 1 + \cdots, (3.1)$$

which diverges since its partial sums form the divergent sequence 1, 0, 1, 0, ...Now insert some parentheses to help it along:

$$(1-1) + (1-1) + (1-1) + \cdots$$
 (3.2)

73



Figure 3.14: Guido Grandi experiences an *identity* crisis. Actually, the paradox didn't bother Grandi at all. He found it theologically illuminating: "By putting parentheses into the expression 1 - 1 + 1 - 1 + ... in different ways, I can, if I want, obtain 0 or 1. But then the idea of the creation ex nihilo is perfectly plausible." [1]

This series is just $0 + 0 + 0 + 0 + \cdots$, which converges to 0. We can even insert parentheses a different way and evaluate

$$1 + (-1 + 1) + (-1 + 1) + \cdots, \qquad (3.3)$$

which then converges to 1. (Figure 3.14.) Evidently, we can't get associativity for infinite sums in general. Uh oh. Does Oresme's proof have a gaping hole in it? It seemed so convincing!

No, not a gaping hole, just a tiny technicality to address. It is true that if $\sum a_n$ converges, then we can insert parenthesis wherever we want and it will still converge to the same thing. Proof: Inserting parenthesis amounts to looking at a subsequence of the sequence of partial sums. (Figure 3.15.) If the sequence of partial sums converges to begin with, then all subsequences also converge to the same thing, so we can add parentheses to the series willy-nilly and the sum won't change. Adding parentheses can only help the series converge. So Oresme's proof works, because⁴ if the harmonic series did converge, then the series $1 + \frac{1}{2} + \frac{1}{2} + \frac{1}{2} + \ldots$ would have to converge to something smaller.

Now that we've seen that associativity does not generalize to infinite series, we'll look at commutativity (a + b = b + a). Can we rearrange terms of a series without affecting the sum?

The answer is, again, no in general. As an example, let's rearrange the al-

⁴Another way to justify Oresme's argument: Every term of the harmonic series is nonnegative, so the sequence of partial sums is monotone. Every subsequence of a monotone sequence has the same convergence behavior as the original sequence.



Figure 3.15: The sequence of partial sums for Grandi's series oscillates and diverges. But the subsequence consisting of just the blue dots converges to 1, and the subsequence consisting of just the black dots converges to 0.

ternating harmonic series, $\sum_{n=1}^{\infty} \frac{(-1)^{n+1}}{n}$. Using Taylor series, we can evaluate:

$$\sum_{n=1}^{\infty} \frac{(-1)^{n+1}}{n} = 1 - \frac{1}{2} + \frac{1}{3} - \frac{1}{4} + \dots = \log 2.$$

By the way, that's a *natural* logarithm.⁵⁶ Now let's rearrange the series like this:

$$S = 1 - \frac{1}{2} - \frac{1}{4} + \frac{1}{3} - \frac{1}{6} - \frac{1}{8} + \frac{1}{5} - \frac{1}{10} - \frac{1}{12} + \cdots,$$

which is a pattern of an odd denominator followed by two consecutive even denominators. This new series does *not* converge to $\log 2$. If it did, we could insert parentheses without altering the sum, but:

$$\begin{pmatrix} 1 - \frac{1}{2} \end{pmatrix} - \frac{1}{4} + \begin{pmatrix} \frac{1}{3} - \frac{1}{6} \end{pmatrix} - \frac{1}{8} + \begin{pmatrix} \frac{1}{5} - \frac{1}{10} \end{pmatrix} - \frac{1}{12} + \cdots$$

$$= \frac{1}{2} - \frac{1}{4} + \frac{1}{6} - \frac{1}{8} + \frac{1}{10} - \frac{1}{12} + \cdots$$

$$= \frac{1}{2} \left[1 - \frac{1}{2} + \frac{1}{3} - \frac{1}{4} + \frac{1}{5} - \frac{1}{6} + \cdots \right]$$

$$= \frac{1}{2} \log 2.$$

(Figure **3.16**.)

How far can we push this madness? Which series have sums which depend on the order of summation? And which values can such a series be made to sum to?

We'd better clarify what it means to "rearrange" the terms of a series. Intuitively, we just want to add up the terms in a different order. But of course,

$$1 + \frac{1}{2} + \frac{1}{4} + \frac{1}{8} + \frac{1}{16} + \dots$$

 $^{^5 {\}rm In}$ analysis, when the base of a logarithm isn't specified, you should assume it's base e. This is in contrast to e.g. computer science, where logs are base 2 by default.

 $^{^{6}\}mathrm{Here's}$ a joke: What do analysts and number theorists throw into the fireplace? Answer: Natural logs!



Figure 3.16: You can think about a series such as $\sum \frac{(-1)^{n+1}}{n}$ financially. The positive terms of the series are paychecks and the negative terms are bills. When you get a paycheck, you immediately deposit it, and when you get a bill, you immediately pay it off. The series converges to log 2, which means that as time progresses, your bank account balance will converge to log 2. The paychecks sum to infinite wealth, and the bills sum to infinite debt, so your bank account balance converging is the result of a careful balancing act. Each paycheck puts your bank account balance a little above log 2, and each bill puts your bank account balance a little below log 2. It should make sense that if you start getting two bills for every paycheck, you won't be able to maintain such a high bank account balance.

should not count as a "rearrangement" of the harmonic series, because some terms of the harmonic series will never appear. We want every term of the original series to appear exactly once in the new series. To make this precise, we start with the definition of a *permutation* of a set S as a bijection $\sigma: S \to S$.

Definition 19. A *rearrangement* of the series $\sum_{n=1}^{\infty} a_n$ is a series of the form $\sum_{n=1}^{\infty} a_{\sigma(n)}$, where σ is a permutation of \mathbb{N} .

Recall that a convergent series $\sum_{n=1}^{\infty} a_n$ is conditionally convergent if $\sum_{n=1}^{\infty} |a_n| = \infty$. For example, the alternating harmonic series is conditionally convergent.

Theorem 20 (Riemann's rearrangement theorem). Let $\sum a_n$ be a conditionally convergent series. Then for any $L \in \mathbb{R} \cup \{\pm \infty\}$, there is a permutation $\sigma(n)$ so that $\sum a_{\sigma(n)} = L$.

Apparently, conditionally convergent series are so weak-willed that they can be persuaded to converge to *anything at all* by permuting the terms! Let's get started with the proof.

Lemma 1. Suppose $\sum_{n=1}^{\infty} a_n$ is conditionally convergent. Then the sum of the positive terms diverges $(to +\infty)$ and the sum of the negative terms diverges $(to -\infty)$.



Figure 3.17: The rearrangement of the alternating harmonic series that the proof of Theorem 20 constructs for the target sum L = 1.2.

Proof. Let $a_n^+ := \max\{a_n, 0\}$ and $a_n^- := \min\{a_n, 0\}$, so that $\sum a_n^+$ is the sum of the positive terms and $\sum a_n^-$ is the sum of the negative terms. Since $\sum a_n$ converges, either $\sum a_n^+$ and $\sum a_n^-$ both converge, or they both diverge. But they can't both converge, because that would imply that $\sum |a_n| = \sum a_n^+ - \sum a_n^-$ converges.⁷

Proof sketch of Theorem 20. First suppose $L \in \mathbb{R}$. Without loss of generality, assume $L \geq 0$. By the lemma, our positive terms are worth ∞ and our negative terms are worth $-\infty$, so let's use them! Add a bunch of positive terms until our partial sum exceeds L. Then throw in some negative terms until we drop below L, then back to positive terms, etc. We switch to adding terms of the other sign as soon as we pass L. In this way, we "use up" all the terms in the series, and the error between our partial sum and L goes to 0 as time progresses, since $a_n \to 0$ as $n \to \infty$. (Figure 3.17)

Now suppose $L = +\infty$. Then we simply add up a lot of positive terms, then a negative term, then a lot of positive terms, then a negative term, etc. By the lemma, we can always add enough positive terms to more than make up for the negative term. The $L = -\infty$ case is symmetric.

If $\sum_{n=1}^{\infty} |a_n|$ converges, we say that $\sum_{n=1}^{\infty} a_n$ is absolutely convergent. Here's

76

 $^{^7\}mathrm{Here}$ we're using the easy fact that rearranging a series with nonnegative terms does not affect the sum.

a converse to Riemann's rearrangement theorem. Dirichlet showed that rearranging an absolutely convergent series never changes the sum:

Theorem 21 (Dirichlet). Any rearrangement of an absolutely convergent series converges to the same thing as the original series.

Proof. Consider a rearrangement $\sum_{n=1}^{\infty} a_{\sigma(n)}$ of the absolutely convergent series $\sum_{n=1}^{\infty} a_n = L$, and fix $\varepsilon > 0$. The key fact is that by absolute convergence, there is some N so that $\sum_{n=N+1}^{\infty} |a_n| \le \varepsilon$. The main weight of the sum is in the first N terms, so wait for $T \in \mathbb{N}$ large enough that $\{1, \ldots, N\} \subseteq \{\sigma(1), \ldots, \sigma(T)\}$. Now apply the triangle inequality a bunch of times:

$$\left| L - \sum_{n=1}^{T} a_{\sigma(n)} \right| \leq \left| L - \sum_{n=1}^{N} a_n \right| + \left| \sum_{\substack{n \leq T \\ \sigma(n) > N}} a_{\sigma(n)} \right|$$
$$\leq \left| \sum_{n=N+1}^{\infty} a_n \right| + \sum_{\substack{n \leq T \\ \sigma(n) > N}} |a_{\sigma(n)}|$$
$$\leq \varepsilon + \varepsilon.$$

1

18 A Taylor series that converges to the wrong function

Remember Maclaurin series from calculus? Here's the formula:

$$f(x) = \sum_{n=0}^{\infty} \frac{f^{(n)}(0)}{n!} x^n.$$
(3.4)

For example, $e^x = \sum_{n=0}^{\infty} \frac{x^n}{n!}$. Here's a joke: Why does the Maclaurin series fit f so well? Because it's Taylor-made! (Recall that a Maclaurin series is a special kind of Taylor series.)

There are some caveats to Equation (3.4), though. Obviously the formula only makes sense if f is infinitely differentiable at 0. And sometimes the Taylor series *diverges* for some values of x, e.g. $\frac{1}{1-x} = \sum_{n=0}^{\infty} x^n$ only works for |x| < 1.

Are there any other caveats? Suppose the Taylor series makes sense and converges everywhere; then is Equation (3.4) true? Surprisingly, not even close! There exists an infinitely differentiable $f : \mathbb{R} \to \mathbb{R}$ whose Taylor series converges everywhere – to a different function! Such a function f is admittedly a counterexample to our joke, but it's just a joke. Here's an example of such a function (Figure 3.18):

$$f(x) = \begin{cases} e^{-1/x}, & x > 0\\ 0, & x \le 0 \end{cases}$$



Figure 3.18: Plot of $e^{-1/x}$ near zero. It is quite flat at zero but then ever so slowly, just like a turtle, makes its way up away from the x-axis.

We'll outline a proof that f is infinitely differentiable at 0, but its Maclaurin series is just the zero function $T(x) \equiv 0$, which disagrees with f for all x > 0. So there isn't even *any* neighborhood of 0 in which the Taylor series gives the right answer!

Since f(x) = 0 for x < 0, the left-derivatives $f_{-}^{(n)}(0)$ are all clearly 0. To compute the right hand limit at zero, we use the following lemma, which can be proved using induction.

Lemma 2. For x > 0,

$$f^{(n)}(x) = p_n(1/x)e^{-1/x},$$

where p_n is a polynomial of degree at most 2n.

Using this, the definition of the derivative, and the fact that exponentials dominate polynomials, one can show that $f_{+}^{(n)}(0) = 0$, so that $f^{(n)}(0) = 0$. Then the Taylor series for f centered at x = 0 is simply

$$T(x) = \sum_{n=0}^{\infty} \frac{f^{(n)}(0)}{n!} x^n = \sum_{n=0}^{\infty} 0 = 0,$$

which is certainly convergent for all $x \in \mathbb{R}$.

Remark 1. f has a relative

$$g(x) = \begin{cases} e^{-1/x^2}, & x \neq 0\\ 0, & x = 0 \end{cases}$$

which is even more mischevious: the Maclaurin series of g disagrees with g everywhere except 0! (Figure 3.19.)

Remark 2. We say h is *real analytic* if for every $a \in \mathbb{R}$, there is some neighborhood of 0 on which the Maclaurin series of h(x - a) converges to h(x - a). So f and g from this section are examples of functions that are infinitely differentiable, but not real analytic. Maybe you hope these non real analytic functions



Figure 3.19: The "seagull function" $f(x) = e^{-1/x^2}$ is infinitely differentiable, but all of its derivatives at 0 are 0, so its Taylor series converges to the function $g(x) \equiv 0$.

are "rare". Too bad! Let's say we start with a function h that is real analytic. Then add f to it, giving a function h + f which is not real analytic. This shows that of all infinitely differentiable functions, at least as many are not real analytic as are (in the sense of cardinality). In fact, even the set of infinitely differentiable but nowhere analytic functions on \mathbb{R} is nonmeager in $\mathcal{C}^{\infty}(\mathbb{R})$! (See [3].)

Remark 3. This nonsense disappears in magical complex analysis land. Complex differentiability is equivalent to (complex) analyticity and infinite differentiability. See Section 2.

19 Misshapen series

So far, we've investigated *standard* series, of the form $\sum_{n=1}^{\infty} a_n$. But standards are for chumps. How about a two-sided series? E.g.

$$\sum_{n=-\infty}^{\infty} 2^{-|n|} = \dots + 2^{-2} + 2^{-1} + 2^{-0} + 2^{-1} + 2^{-2} + \dots$$

It seems pretty clear that this series should converge to $1 + 2 \sum_{n=1}^{\infty} 2^{-n} = 3$. (Figure 3.20.)



Figure 3.20: Summing a two-sided series.

How about a two-dimensional series?

$$2^{-1} + 2^{-2} + 2^{-3} + 2^{-4} + 2^{-5} + 2^{-6} + \cdots + 2^{-2} + 2^{-3} + 2^{-4} + 2^{-5} + 2^{-6} + \cdots + 2^{-3} + 2^{-4} + 2^{-5} + 2^{-6} + \cdots + 2^{-4} + 2^{-5} + 2^{-6} + \cdots + 2^{-5} + 2^{-6} + \cdots + 2^{-6} + \cdots \vdots$$

This one ought to converge to $\sum_{n=1}^{\infty} n2^{-n} = 2$. (Figure 3.21.)



Figure 3.21: Summing a two-dimensional series in your flower garden.

More generally, if we have a countable index set I, we can make sense of the series via a bijection $\mathbb{N} \to I$. As we saw in Section 17, the value of the sum might depend on which bijection we choose. But there are no discrepancies if all the terms are positive.

What if we want to sum up *uncountably many* terms? You might remember that back in Section 12, we actually found it *useful for a proof* to use a series with terms that were not indexed by \mathbb{N} . The definition we used there generalizes nicely to the uncountable case: Suppose I is some arbitrary index set, and for $i \in I$, a_i is a *nonnegative* real number. Then we define

$$\sum_{i \in I} a_i = \sup \left\{ \sum_{i \in J} a_i : J \text{ is a finite subset of } I \right\}.$$

For example, if $I = \mathbb{N}$, this definition matches the standard notion of convergence for series. By taking $I = \mathbb{R}$, we can add up all the values of some nonnegative function $\mathbb{R} \to \mathbb{R}$! But uncountable sums are not as exciting as you might hope:

80

Proposition 17. Suppose that for uncountably many $i \in I$, $a_i > 0$. Then $\sum_{i \in I} a_i = \infty$.

Proof. For $n \in \mathbb{N}$, let $E_n = \{i \in I : a_i > \frac{1}{n}\}$, so that $\cup_n E_n$ is uncountable. A countable union of finite sets is countable, so some E_n must be infinite. Therefore,

$$\sum_{i \in I} a_i \ge \frac{1}{n} + \frac{1}{n} + \frac{1}{n} + \dots = \infty.$$

What if our series has negative terms?

$$-3 - 2 - 1 - 0 + 1 + 2 + 3 + \dots = ???$$

Sometimes we can say something about such a series. We can separate our series $S = \sum_{i \in I} a_i$ into the positive part and the negative part:

$$S^+ = \sum_{\substack{i \in I \\ a_i \ge 0}} a_i, \qquad S^- = \sum_{\substack{i \in I \\ a_i \le 0}} a_i.$$

Both of these series make sense by our earlier definition, and if at most one of S^+ and S^- is infinite, then we can define $S = S^+ + S^-$. But if $S^+ = \infty$ and $S^- = -\infty$, we just leave S undefined. All of these ideas are generalized by measure theory and the Lebesgue integral. But that's a story for Chapter ??.

20 If you torture a series enough, it will converge

Earlier, we saw Grandi's series $S = 1 - 1 + 1 - 1 + \cdots$, which can be made to sum to 0 or 1 by judiciously inserting parentheses. Mathematicians are sane, clearthinking folk, so *of course*, all the great mathematicians of history understood that Grandi's series obviously *diverges*, and thus it simply doesn't have a sum: it's not 0, it's not 1, and it's *certainly* not anything else... right?

(Leibniz, 1674 [5]) $\frac{1}{1+1} = \frac{1}{1} - \frac{1}{1+1}$. Ergo $\frac{1}{1+1} = 1 - 1 + 1 - 1 + 1 - 1$ etc.

(Leibniz, 1713 [7]) ...And now since from that one [Gerolamo Cardano] who wrote of the values of the gambling games, it had been shown that when the average between two even quantities is found by calculation, the arithmetic mean ought to be found, which is onehalf of the sum, and in such a way this nature of things attends to the same law of righteousness; hence although 1 - 1 + 1 - 1 + 1 - 1 +etc is 0 in the case with an finite even number of elements, in the case with a finite odd number of elements it is equal to 1; it follows that in the case with both sides vanishing into multitude of infinite elements, where the law is confounded by the presence of both evens and odds, and there is such a great sum on both sides, that $\frac{0+1}{2} = \frac{1}{2}$ emerges, which is what has been proposed. "Gambling games"? "Law of righteousness"? What was Leibniz smoking? But he's in good company!

(Euler, 1760 [2]) For if in a calculation I arrive at this series 1 - 1 + 1 - 1 + 1 - 1 etc. and if in its place I substitute 1/2, no one will rightly impute to me an error, which however everyone would do had I put some other number in the place of this series. Whence no doubt can remain that in fact the series 1 - 1 + 1 - 1 + 1 - 1 + 1 - 1 + 1 - 1 + 1 - 1 etc. and the fraction 1/2 are equivalent quantities and that it is always permitted to substitute one for the other without error.

Other giants of mathematics (e.g. Bernoulli, de Morgan) expressed similar confidence that Grandi's series sums to 1/2. Were they all crazy or stupid? It's more reasonable to be charitable. These mathematicians must have had some *different definitions* in mind for how infinite sums work. According to the standard definitions, it's not true that Grandi's series sums to 1/2... but let's make it true!

Definition 20. Suppose a series $\sum_{n=1}^{\infty} a_n$ has partial sums S_1, S_2, \ldots The *Cesàro sum* of the series is the limit of the *arithmetic mean of the first m* partial sums:

$$C = \lim_{m \to \infty} \frac{\sum_{N=1}^{m} S_N}{m}.$$

Proposition 18. If $\sum_{n=1}^{\infty} a_n = L \in \mathbb{R}$, then the Cesàro sum of $\sum_{n=1}^{\infty} a_n$ is L. Proof. Fix an arbitrary $\varepsilon > 0$. Choose N_0 large enough so that for every $N > N_0$, $|S_N - L| < \varepsilon$. Then apply the triangle inequality a few times:

$$\begin{split} \left| \frac{\sum_{N=1}^{m} S_{N}}{m} - L \right| &\leq \left| \frac{N_{0}}{m} L - \frac{\sum_{N=1}^{N_{0}} S_{N}}{m} \right| + \left| \frac{m - N_{0}}{m} L - \frac{\sum_{N=N_{0}+1}^{m} S_{N}}{m} \right| \\ &\leq \frac{1}{m} \cdot (\text{no } m \text{ dependence}) + \frac{m - N_{0}}{m} \varepsilon \\ &\leq 2\varepsilon \quad \text{for } m \text{ sufficiently large.} \end{split}$$

So Cesàro provides the right answer when you give him a convergent series. But sometimes, he even gives an answer if you give him a divergent series! The partial sums of Grandi's series are 1, 0, 1, 0, 1, 0, ..., hence the Cesàro sum of the series is $\frac{1}{2}$. (See Figure 3.22 for another example.)

Does *every* series have a Cesàro sum? Nah, that's too good to be true. Consider the series

$$1 - 2 + 3 - 4 + 5 - 6 + 7 - 8 + \dots$$

Obviously this series diverges, but Euler claimed that this series sums to $\frac{1}{4}$ for some reason. Unfortunately, the arithmetic mean of the first *m* partial sums is

$$\begin{cases} \frac{\lceil m/2 \rceil}{m} & \text{if } m \text{ is odd} \\ 0 & \text{otherwise.} \end{cases}$$

82



Figure 3.22: The series $\sum_{n=0}^{\infty} \sin(n)$. The top graph shows the *terms* of the series. The middle graph shows the *partial sums* of the series. The bottom graph shows the *arithmetic means* of the partial sums of the series. As the images suggest, the sequence of partial sums oscillates and hence the series diverges, but the Cesàro sum of the series is $\frac{1}{2} \cot \frac{1}{2} \approx 0.915$.

When m is large, this oscillates between 0 and approximately 1/2, so it does not converge. Hence, the series has no Cesàro sum. We'll need to be cleverer to justify Euler's claim.

The oscillation behavior we're trying to deal with is suspiciously similar to the behavior of the partial sums of Grandi's series. It seems like we just need to somehow "Cesàro-ify" the series *again*! To see how to properly do that, define

$$e_k = \begin{cases} 1 & \text{if } k = 1\\ 0 & \text{otherwise.} \end{cases}$$

Then for a series $\sum_{k=1}^{\infty} a_k$, the *n*th partial sum is

$$S_n = \frac{\sum_{k=1}^n a_k}{\sum_{k=1}^n e_k},$$

and the arithmetic mean of the first m partial sums is

$$M_m = \frac{\sum_{n=1}^m \sum_{k=1}^n a_k}{\sum_{n=1}^m \sum_{k=1}^n e_k}$$

Hopefully it's clear how to generalize. Fix an integer $p \ge 0$ which tells how many times to Cesàro-ify, and fix a series $\sum_{k=1}^{\infty} a_k$. Define

$$C_n^p = \frac{\sum_{n_0=1}^n \sum_{n_1=1}^{n_0} \sum_{n_2=1}^{n_1} \cdots \sum_{n_p=1}^{n_{p-1}} a_{n_p}}{\sum_{n_0=1}^n \sum_{n_1=1}^{n_0} \sum_{n_2=1}^{n_1} \cdots \sum_{n_p=1}^{n_{p-1}} e_{n_p}}.$$

For example, C_n^0 is the *n*th partial sum, and C_n^1 is the mean of the first *n* partial sums. We define the *p*-Cesàro sum of the series to be $\lim_{n\to\infty} C_n^p$. For example, the 0-Cesàro sum is the standard sum, and the 1-Cesàro sum is the original Cesàro sum.

Now we can generalize our earlier claim that Cesàro gives the right answer for convergent series. As p gets larger, Cesàro never goes back on his word:

Theorem 22. If the p-Cesàro sum of a series is L, then the (p + 1)-Cesàro sum of the series is also L.

(We omit the proof.) If you're bored, you can check that $1 - 2 + 3 - 4 + 5 - 6 + \cdots$ has a 2-Cesàro sum, namely $\frac{1}{4}$. (Figure 3.23) This isn't exactly the calculation that Euler performed to arrive at the sum $\frac{1}{4}$, but it's also not a coincidence.

You might hope that by considering sufficiently large p, we can force *every* series to converge. But that's still too good to be true. For example, Theorem 22 works even for $L \in \{\pm \infty\}$, so Cesàro is worthless for dealing with series that diverge to $\pm \infty$.

In light of this defect, let's look at another (fishier) way to force divergent series to converge. What is $1 + 2 + 4 + 8 + 16 + \cdots$? Obviously ∞ , but that answer is unacceptable. Recall the formula for geometric series: if |x| < 1, then

$$1 + x + x^2 + x^3 + \dots = \frac{1}{1 - x}.$$



Figure 3.23: The series $1 - 2 + 3 - 4 + 5 - 6 + \cdots$. The top graph shows the terms of the series. The next three graphs show C_n^0 , C_n^1 , and C_n^2 .

However, the right side of the equation, $\frac{1}{1-x}$, makes sense for any $x \neq 1$. By plugging in x = 2, we arrive at the fun claim that

$$1 + 2 + 4 + 8 + 16 + \dots = -1.$$

And by plugging in x = -1, we recover our earlier calculation that Grandi's series sums to 1/2. For another example of this fishy calculation technique, let's look at the following "sum" that many people try to assert:



Again, keep in mind that with respect to *standard* summation, this formula is *just plain false*. But we can make the formula seem true, just like we did for $1+2+4+8+\cdots = -1$. We know that $1+2+3+4+\cdots$ in reality diverges, but we can multiply each term by a *damping factor* to make a convergent series, and then maybe try to take a limit. So instead of $\sum_{n=1}^{\infty} n$, we take $S(x) := \sum_{n=1}^{\infty} nx^{n-1}$, where we imagine that x is small so x^{n-1} gets smaller and smaller as n increases. Eventually, the plan is to plug in x = 1 to get the value of $\sum_{n=1}^{\infty} n$. Observe

$$S(x) = 1 + 2x + 3x^{2} + 4x^{3} + \dots = \frac{d}{dx} \left[x + x^{2} + x^{3} + \dots \right]$$
$$= \frac{d}{dx} \frac{x}{1 - x} = \frac{1}{(1 - x)^{2}}.$$

We can't get a meaningful expression by evaluating x = 1 here unfortunately. But we can clear some denominators and verify the identity

$$\frac{1}{(1+x)^2} = S(x) - 4xS(x^2),$$

at least for x where S(x) and $S(x^2)$ are defined. But if we use this equation with x = 1, we obtain

$$S(1) = 1 + 2 + 3 + 4 + \dots = -\frac{1}{12}.$$

Yay! So this fishy calculation technique, where you find a formula that works for some numbers and then just boldly plug in other numbers, seems to be very powerful. You have to be careful though, because you can get conflicting answers. For example, Callet noticed that for |x| < 1, we have

$$1 - x^{2} + x^{3} - x^{5} + x^{6} - x^{8} + x^{9} - x^{11} + x^{12} - \dots = \frac{1 + x}{1 + x + x^{2}}$$

Again, the right-hand side makes sense for other values of x, so we can plug in x = 1 to find this time that Grandi's series sums to 2/3!

There's a way to actually make rigorous mathematical sense of this fishy calculation technique. But it involves complex numbers, so we put it in Appendix 2.

References

- [1] Giorgio T Bagni. Appunti di didattica della matematica. 2005.
- [2] Edward J Barbeau and PJ Leah. "Euler's 1760 paper on divergent series". In: *Historia Mathematica* 3.2 (1976), pp. 141–160.
- [3] R. B. Darst. "Most infinitely differentiable functions are nowhere analytic". In: *Canad. Math. Bull.* 16 (Jan. 1973), pp. 597–598. DOI: 10.4153/cmb-1973-098-3. URL: http://dx.doi.org/10.4153/CMB-1973-098-3.
- [4] Mike Paterson et al. "Maximum overhang". In: American Mathematical Monthly 116.9 (2009), pp. 763–787.
- [5] Siegmund Probst. Differenzen, Folgen und Reihen bei Leibniz (1672– 1676). na, 2006.
- [6] Arild Stubhaug and Richard H Daly. *Niels Henrik Abel and his times:* called too soon by flames afar. Springer, 2000.
- [7] Robert Woodhouse. The principles of analytical calculation. 1803.

TO BE CONTINUED...

Chapter 4

Acknowledgments

Many Bothans died to bring us this information.

Mon Mothma [2]

We thank many people for their inspirations, suggestions, and contributions to this project. Here's a list of some of them, in no particular order: Sunshine Dobois and Colin Macdonald [1], William Hoza, Laura Shou, Oriel Humes, Adam Jermyn, Eric Bobrow, Pairs Mielke, George Washington, Xander Rudelis, Shival Dasu, Kathleen Hoza, Nancy Zhang, Algae Elbaum, Molly Weasley, Alex Bourzutschky, Paul Zhang, Zachary Chase, Helen Xue, Matthew Weidner, Tiffany Jager.

References

- [1] S. Duvois and C. Macdonald. 101 Illustrated Analysis Bedtime Stories. 2001. URL: http://people.maths.ox.ac.uk/macdonald/errh/.
- [2] George Lucas. Star Wars Episode VI: The Return of the Jedi. 1983.

Appendix A

Omitted Details

1 Adding books to the top of a stack

This section elaborates on the last paragraph of Section ??.

Definition 21. A *plan* is a sequence x_1, x_2, \ldots of real numbers. (Think of x_n as the horizontal position of the center of book n, where the origin is at the upper right corner of the table. Remember, no two books can have the same vertical position.)

Definition 22. A plan x_1, x_2, \ldots is *sound* if for every $N \in \mathbb{N}$, if you were to place N books at positions x_1, \ldots, x_N (with book 1 on the bottom and book N on the top), then that stack of N books would not topple. (In other words, if you were to build an ever-growing stack by adding books to the top one by one in positions x_1, x_2, \ldots , then the stack would never topple.) See Figure A.2.

Proposition 19. Suppose x_1, x_2, \ldots is a sound plan. Then for every $n, x_n \leq \frac{1}{2}$. So the overhang of the stack always satisfies $d \leq 1$.



Figure A.1: The plan $0, -\frac{1}{2}, 0, 0, 0, 0, 0, \dots$ is sound. It achieves an overhang of $d = \frac{1}{2}$.



Figure A.2: The plan $-\frac{3}{8}$, $-\frac{5}{24}$, 0, $\frac{1}{2}$, 0, 0, 0, 0, ... is sound. It achieves an overhang of d = 1, showing that Proposition 19 is tight.

Proof. Suppose x_1, x_2, \ldots is a plan with $x_n = \frac{1}{2} + \varepsilon$ for some $n \in \mathbb{N}$ and some $\varepsilon > 0$. We'll show that the plan is not sound. For each N > n, define

$$a_N = \frac{1}{N} \sum_{i=1}^N x_i$$
$$b_N = \frac{1}{N-n} \sum_{i=n+1}^N x_i.$$

(So a_N is the COM of books $1, \ldots, N$, and b_N is the COM of books $n+1, \ldots, N$.) Then we have

$$a_{N} = \frac{(N-n)b_{N} + \sum_{i=1}^{n} x_{i}}{N}$$
$$= b_{N} - \frac{n}{N}b_{N} + \frac{1}{N}\sum_{i=1}^{n} x_{i}$$

First, suppose that $|b_N - x_n| > \frac{1}{2}$ for some N. Then when the stack has N books, it will topple over, pivoting about one of the top corners of book n.

Therefore, assume instead that $|b_N - x_n| \leq \frac{1}{2}$ for all N. Then $\frac{n}{N}b_N \to 0$ as $N \to \infty$, and of course $\frac{1}{N} \sum_{i=1}^n x_i \to 0$ as $N \to \infty$, so $|a_N - b_N| \to 0$ as $N \to \infty$. Choose N large enough that $|a_N - b_N| < \varepsilon$. Then by the triangle inequality, $|a_N - x_n| < \frac{1}{2} + \varepsilon$, so $a_N > 0$. Therefore, when the stack has N books, it will topple, pivoting about the upper right corner of the table.

2 Analytic continuation

This section elaborates on the last paragraph of Section 20. In particular, we will investigate another way to argue that $1 + 2 + 3 + \cdots = -\frac{1}{12}$.



2. ANALYTIC CONTINUATION

We are going to need analytic continuation from complex analysis. If you're already familiar with this, feel free to skip a few paragraphs down to the part about the Riemann zeta function. The functions we will be thinking about in this section eat complex numbers and spit out more complex numbers. In other words, they map $\Omega \to \mathbb{C}$, where $\Omega \subseteq \mathbb{C}$ is an open subset of \mathbb{C} . For example, the functions f(z) = z or $g(z) = z^2$ are fine examples $\mathbb{C} \to \mathbb{C}$. The notion of (complex-)differentiability looks basically the same as in the real case: we say f is complex-differentiable or holomorphic at $z_0 \in \Omega$ if

$$f'(z_0) = \lim_{h \to 0} \frac{f(z_0 + h) - f(z_0)}{h}$$

exists. It looks the same as the real case, except that $h \neq 0$ is allowed to take complex values.¹ We'll find out later, though, that complex-differentiability seems a much stronger condition than the usual real-valued differentiability.

The idea of analytic continuation is to take a (complex-)differentiable function defined on some set $\Omega \subset \mathbb{C}$ and try to extend it nicely to a larger set Ω' . For example, $f(z) := \sum_{n=0}^{\infty} z^n$ is equal to $\frac{1}{1-z}$ for |z| < 1 but is not defined for |z| > 1. But the honest function $z \mapsto \frac{1}{1-z}$, which agrees with f on the unit disk |z| < 1, is defined and complex-differentiable on the larger set $\mathbb{C} \setminus \{1\}$. We say that $z \mapsto \frac{1}{1-z}$ is an *analytic continuation* of f.

Remark 4. You may recall that an analytic function is one that is locally given by a convergent power series. In real analysis, there are examples of infinitely differentiable functions that are not (real-)analytic (Section 18). But in complex analysis, it turns out that one-time differentiable, infinitely differentiable, and analytic are all equivalent.



Figure A.3: The rightward arrows are true in complex analysis, but not in real analysis. Complex differentiability is so strong that it implies infinite differentiability and analyticity.

Analytic continuation is incredibly useful because of the *uniqueness of analytic functions*:

Theorem 23 (uniqueness of analytic functions). If $f, g : \Omega \to \mathbb{C}$ are complexdifferentiable on a connected open set $\Omega \subseteq \mathbb{C}$ and f(z) = g(z) for all z in a sequence of distinct points with an accumulation point, then f(z) = g(z) on all of Ω .

¹We do also require that $z_0 + h \in \Omega$.

This is an amazing result! It says there is only *one* analytic continuation: if f and g agree on some smaller set S with an accumulation point (for example, any non-empty open set), then they must be equal on the *entire* set Ω . So we can talk about *the* analytic continuation. Complex analysis is of course magic. For more magic, see a complex analysis book like [2].

The Riemann zeta function $\zeta(s) := \sum_{n=1}^{\infty} \frac{1}{n^s}$ is especially important in analytic number theory. This particular expression, as an infinite sum, is defined for $\Re \mathfrak{e} s > 1$, which is where it converges. However, we would like to analytically continue it to all of $\mathbb{C} \setminus \{1\}$. While we won't prove the analytic continuation here, we'll try to provide some ideas from a proof.

The Riemann zeta function has a friend called the *Gamma function*,

$$\Gamma(s):=\int_0^\infty e^{-t}t^{s-1}\,dt,\quad \mathfrak{Re}\,s>0.$$

Gamma may look a bit scary, but by integrating by parts, we can verify the functional equation $\Gamma(s+1) = s\Gamma(s)$ and conclude that $\Gamma(n+1) = n!$ for $n \in \mathbb{N}$. (And yes, that is a factorial, not just an exclamation point.) Although the integral expression for $\Gamma(s)$ is only useful for $\Re \mathfrak{e} s > 0$, we can use the functional equation to analytically continue it. We simply copy-paste everything to the left one unit at a time. The only ugly part is the singularity at s = 0, which gets translated to all the negative integers.



Figure A.4: A plot of $\Gamma(s)$ along the real line. It matches the factorial function (shifted over by one) for $n \in \mathbb{N}$. The pole at zero is copied over to all the negative integers.

Gamma plays nicely with Riemann zeta. One way to prove the analytic continuation of Riemann zeta is to form the auxiliary function

$$\xi(s) := \pi^{-s/2} \Gamma(s/2) \zeta(s),$$

REFERENCES

and then prove analytic continuation of ξ . This involves the *Poisson summation* formula and Theta series $\theta(t) := \sum_{k \in \mathbb{Z}} e^{-t\pi k^2}$ along with some sums, integrals, and computations. The end result is an integral formula for ξ that works for $s \in \mathbb{C}, s \neq 0, 1$. Then we get a formula for ζ by dividing by Γ .

Theorem 24 (analytic continuation of Riemann zeta). $\zeta(s) = \sum_{n=1}^{\infty} \frac{1}{n^s}$, $\Re \mathfrak{e}(s) > 1$, extends to a meromorphic function on \mathbb{C} with only a simple pole at s = 1.

We three in the word "meromorphic", but that just means analytic except for some isolated poles, like the one at s = 1. With analytic continuation of ζ , we can have fun assigning values to sums like $\zeta(-1) = 1 + 2 + 3 + \cdots$. By comparing poles and residues of Γ and ζ , it turns out

$$1 + 2 + 3 + \dots = \zeta(-1) = -\frac{1}{12}$$

and $1 + 1 + 1 + \dots = \zeta(0) = -\frac{1}{2}$.

Tada! Complex analysis is magic.

References: complex analysis [2], analytic continuation of the Riemann zeta function [1]

References

- [1] Anto Deitmar. A First Course in Harmonic Analysis. Springer, 2005.
- [2] Elias Stein and Rami Shakarchi. *Complex Analysis*. Vol. 3. Princeton Lectures in Analysis. Princeton University Press, 2003.