

INTRODUCTION TO LOGIC & PROOF

Chris Johnson

Fall 2023

Contents

Contents	ii
Introduction to the Course	vi
Course format	vii
Assignments	vii
Homework	vii
Pop quizzes	viii
Proof portfolio	viii
Midterm exams	viii
Final exam	ix
Make-up assignments	ix
Extra credit	x
Expectations	x
Online notes	xi
L ^A T _E X	xi
How to study in this course	xi
Some personal information	xiii
1 Motivation for Logic and Proof	1
1.1 What are proofs?	1
1.2 What is logic?	5
1.3 Why should we learn about logic and proofs?	6
1.4 Applied math and pure math	6
2 Basic Proof Techniques	8
2.1 Direct proofs and counterexamples	8
2.1.1 First examples of direct proofs	8
2.1.2 Counterexamples	10
2.1.3 More direct proof examples and exercises	11
2.1.4 Theorems, conjectures, lemmas, corollaries, etc.	14
2.1.5 Some practice in deductive reasoning	17

2.1.6	Divisibility	21
2.2	Proofs by induction	25
2.2.1	Warmup: The quicksort algorithm	25
2.2.2	Mathematical induction in proofs	28
2.2.3	Weak induction	29
2.2.4	Strong Induction	38
2.2.5	Prime numbers and the fundamental theorem of arithmetic	41
2.3	Proofs by contradiction	43
2.3.1	Basic examples	43
2.3.2	Irrational numbers	46
2.3.3	The infinitude of primes	47
3	Symbolic logic	49
3.1	Propositions and Predicates	49
3.1.1	Propositions	49
3.1.2	Variables and predicates/open sentences	50
3.1.3	The universal and existential quantifiers	50
3.2	Logical operations and truth tables	51
3.2.1	Conjunction (and)	52
3.2.2	Disjunction (or)	52
3.2.3	Negation	53
3.2.4	Negating quantifiers	54
3.2.5	Implication	54
3.3	Converses, Equivalences, and Contrapositives	56
3.3.1	Converses	57
3.3.2	Equivalences	57
3.3.3	Contrapositives	58
4	Sets	60
4.1	Basic ideas and definitions	60
4.1.1	Set-builder notation	63
4.1.2	Subsets and supersets	66
4.1.3	Equality	68
4.1.4	The empty set	69
4.1.5	Real numbers	70
4.2	Operations on sets	71
4.2.1	Unions	71
4.2.2	Intersections	72
4.2.3	Products	76
4.2.4	Complements	80

4.2.5	Difference	82
4.2.6	De Morgan's laws	84
4.3	Collections of sets	87
4.3.1	Power set	87
4.3.2	Indexing sets	87
4.3.3	Unions and intersections with indices	88
4.4	Maps between sets	89
4.4.1	Definitions and examples	89
4.4.2	Representing maps	92
4.4.3	Special types of maps	93
4.4.4	Images and preimages	97
4.5	Compositions of Maps	99
4.5.1	Definitions and basic examples	99
4.5.2	Composing three or more functions; associativity	100
4.5.3	Inverse maps	102
4.5.4	Identity maps	103
5	Relations	105
5.1	Basic definitions and examples	105
5.2	Special properties of relations	109
5.2.1	Symmetry	109
5.2.2	Antisymmetry	109
5.2.3	Reflexivity	110
5.2.4	Transitivity	111
5.3	Orderings	111
5.3.1	Total orders	111
5.3.2	Partial orders	112
5.4	Equivalence relations	114
5.4.1	Examples of equivalence relations	114
5.4.2	Equivalence classes	116
6	Binary Operators	120
6.1	Definitions and examples	120
6.1.1	Familiar examples	121
6.1.2	Composition of maps $X \rightarrow X$	122
6.1.3	Addition and multiplication of 2×2 matrices	123
6.2	Properties binary operators may have	125
6.2.1	Associativity	125
6.2.2	Commutativity	127
6.2.3	Identity	128
6.2.4	Inverses	130

6.2.5	Groups	132
6.3	Arithmetic of congruence classes	133
6.3.1	Defining addition and multiplication	134
6.3.2	Addition and multiplication are well-defined	135
6.3.3	Distribution	137
7	Odds and Ends	138
7.1	The Pigeonhole Principle	138
	Index	141

Introduction to the Course

The result of the mathematician's creative work is demonstrative reasoning, a proof; but the proof is discovered by plausible reasoning...

GEORGE POLYA

Welcome to Math 250, *Introduction to Logic and Proof*. The main goal of this course is to help students become proficient at mathematical proofs. By the end of the course, students should be comfortable with reading proofs, should appreciate the necessity of proofs, and be able to construct their own proofs. Throughout the course we will see the most commonly used proof techniques, starting informally by discussing the main ideas, and later becoming more precise after we have introduced basic logic. As we will see, “logic” here means something a bit more formal than the colloquial use of the word that you’re probably familiar with. For us, logic will refer to certain types of “sentences” built from symbols that have a specific meaning in mathematics, and rules for building more complicated sentences from simpler ones. (In a way this is similar to a computer program where certain basic keywords are combined according to specific rules to make the computer do something particular.)

In order to get practice proving theorems, we will need to introduce some different contexts where we will have interesting statements to prove. Our main contexts for proofs will be some basic number theory (arithmetic of integers) and set theory. The number theory and set theory we will introduce is only the tip of the iceberg for these topics, and either topic could easily be an entire semester-long course by itself. We will spend a bit of time on set theory since it is of fundamental importance for modern mathematics. As you will see if you continue taking more advanced courses, many mathematical objects of interests – groups and rings in abstract algebra, spaces in topology, manifolds in geometry, vector spaces in linear algebra, and many other examples – are defined in terms of sets. Thus knowledge of sets is important for any student of mathematics. (Sets also clarify many of the topics you know in less formal ways from previous mathematics courses.)

Course format

This semester I'm teaching this course as a discussion-based lecture, meaning that most days I will introduce a topic in class, and show you various examples of the topic, but I encourage you to stop me and ask questions as they arise. I will also periodically stop and ask if everyone has any questions, and check that everyone's still following the material. I will also sometimes ask for you to work on your own proofs in class in small, informal groups while I walk around the room answering individual questions as they arise. Sometimes I'll ask for someone to volunteer to put their proof on the board, and afterwards we'll discuss the proof, using any mistakes as an opportunity for more discussion to clarify any misconceptions or confusion.

The material in this course is likely very different from the material you've seen in earlier courses, and you should anticipate that you will make mistakes sometimes – and that's okay, it's all part of the learning process. I hope the format and atmosphere of the class will be casual enough that you feel comfortable making mistakes and using those as learning opportunities.

Assignments

Your grade in this class is determined primarily by four types of assignments: weekly homework, sporadic pop quizzes, a proof portfolio, and exams (three midterm exams, and a cumulative final exam).

Homework

Most weeks, homework will be posted to Canvas by 5pm on Friday afternoon and due by 11:59pm the following Thursday. Homework assignments will always be graded out of ten points, with about five points coming from effort and the remaining points coming from accuracy. The effort points you receive simply by making an "honest attempt" at each problem assigned, regardless of whether your solution is correct or not. The accuracy points come from one or two problems that I select to thoroughly grade. There are a few reasons for grading the homework this way. This will be a course where regular feedback on your work is very important, and so I want to be able to get your homework graded and back to you quickly. However, I also think it's important that you work on several different problems to help solidify your understanding of the

material. These goals (several problems to practice with, and thorough feedback given quickly) stand in contrast to one another: the more problems I grade, the slower the grading will go. Grading homework half for effort and half for accuracy is a compromise that helps me get you feedback quickly, while still giving you the opportunity to work on several different problems each week.

Pop quizzes

With the exception of the very beginning of the semester, I won't take attendance in class. However, in order to encourage you to come to class (and this is a class where coming regularly will be very important for your understanding), I will give pop quizzes at the start of class on occasion. These will be very short (sometimes consisting of a single problem), will only take five or ten minutes of class time, and are only graded for effort. That is, as long as you are present in class you will receive full credit on the pop quiz, even if your solutions to the problems on the quiz are incorrect. This is meant to make the pop quizzes a low stress, low stakes situation: you just need to be sure you're in class and you get credit. These will also serve as a place to start our discussion, as after the pop quizzes are collected we'll go over the solutions to the problems given and can discuss any possible confusion or concerns about the problems.

Proof portfolio

The proof portfolio is a semester-long project where you will have an opportunity to re-work some of the homework problems earlier in the semester that you might have missed in order to show that you've learned the material by the end of the course. This portfolio will consist of typed solutions to earlier homework problems and will be due by the last day of class. You will need to submit two drafts of the portfolio earlier in the semester, with the first draft being due on Friday October 6, and the second draft being due on Friday November 10. The final draft is due by 11:59pm on Friday, December 8.

Midterm exams

Our class will have three midterm exams during the semester. These will be closed book, closed note, individual exams. Midterm exams will always be graded out of 100 points, with each problem graded for accuracy. Partial credit is awarded when students begin solving a problem

correctly but make mistakes or simply stop solving the problem. However, students *must* begin solving the problem correctly to receive partial credit. Students will not receive partial credit for completely erroneous or illogical work, or for solving a problem different from what is asked on the exam.

The work you present on a midterm exam is expected to be written legibly and easy to follow.

All students are expected to take each midterm exam. As discussed in the *Make-up exams and homework* section below, make-up exams will only be allowed in a few specific circumstances. Students should always prepare to take the midterm at the date and time announced in class. ***Test anxiety is not a legitimate reason to delay an exam.*** Students who miss an exam for an unexcused reason will receive a grade of zero.

Tentatively, our midterm exams will take place on Wednesday September 13, Wednesday October 11, and Wednesday November 15. These dates are subject to change, however.

Final exam

We will have a cumulative final exam at the end of the semester, the exact date and time of which will be determined by the registrar. The structure and format of the final exam is very similar to that of the midterm exams, though the final will be somewhat longer and counts for a larger portion of the student's grade.

Make-up assignments

Generally speaking, no late work is accepted and no make-ups for missed assignments are allowed. Of course, there are exceptions to this. For example, if you are seriously ill or suddenly injured, then we will work together to find a reasonable solution to a missed assignment. Or, if you are student-athlete that will miss class because you are traveling with your team to a university-sanctioned event *and you notify me before you leave with documentation from your coach*, then we will find a reasonable solution to what you have missed. However, if you happen to miss class the day of an in-class assignment or when a written homework is due because you overslept, are hungover, or simply too anxious or feel unprepared, you **will not** be allowed to make up any missed assignments and ***you will receive a grade of zero on that assignment!***

As homeworks are taken up on Canvas, I do not plan to grant extensions or make-ups for these (even for university-sanctioned travel),

except in very extreme circumstances. Similarly, missed pop quizzes will have a recorded grade of zero, even if you have an “excused” absence. To compensate for this, I will drop some to-be-determined number of pop quizzes for everyone at the end of the semester.

Extra credit

There is no extra credit of any form in this class.

Expectations

This will be class where details are very important, and you will be held to a high standard when it comes to grading your work, both in terms of the content of your work and how you present it. Since part of the goal of this course is to help you transition to more advanced mathematics, it’s important that you learn how mathematics is communicated. This includes learning the standard ways in which mathematics is formatted when it is written, and you will have points deducted from an assignment’s grade if you don’t format your work in the standard way.

Students in this class are expected to be mature and conduct themselves in a professional manner. In terms of this classroom this means

- students are expected to come to class each day;
- be in class prepared with pencil and paper at the start of class
- students should have completed the assigned reading before coming to class;
- pay active attention during class and have any computers, phones, or tablets put away (students *may* take notes on a tablet, however);
- and be ready to participate in class by asking questions about examples from the previous lecture, problems from homework assignments, or any concepts discussed in class or the assigned reading.

Students are expected to spend a *minimum* of eight hours per week working on material for Math 250 (working on homework, reading the textbook, studying notes, etc.). Keep in mind eight hours is the minimum: each additional hour spent working outside of class will have been well-invested come exam time.

Students are strongly encouraged to take advantage of the various studying resources provided by the university and the mathematics department, such as the MTC.

Online notes

In addition to the textbook, I will be typing up my lecture notes for the course and posting them online in Canvas. Students are expected to read both the online lecture notes as well as the Bond & Keane textbook. The readings for each week will be posted to Canvas.

L^AT_EX

Virtually all serious mathematics (and really all scientific material) today is created using LaTeX (pronounced “lay-tech” – the ‘X’ at the end is supposed to be like the Greek letter chi, χ) which is a typesetting language that makes it relatively easy to “type math.”

While you will not be required to use LaTeX this semester, it would be a good idea if you took some time to go ahead and learn it. With this in mind, I’ll be posting some resources for LaTeX to Canvas each week to help you get started if you choose to use LaTeX. The goal will be that by the end of the semester, you’ll have everything you need to have typed up all of your homework assignments and your proof portfolio in LaTeX.

Just to emphasize, you *do not* have to use LaTeX if you don’t want to; you can use Microsoft Word or whatever other software you’d like to type your proof portfolio. LaTeX does have a learning curve and can seem daunting at first, but once you get over the initial “hump” of getting started with LaTeX, I think you’ll find it to be much nicer for writing complicated mathematical expressions than something like Word.

How to study in this course

There is no denying that this will be a difficult course: the ideas we will see will be very different from what you’ve seen in previous courses, and it will take time to become accustomed to the method of thinking necessary to prove theorems. You should anticipate that there will be times that you find the material in this course frustrating. I don’t mention that to scare you off from the course, but just to help prepare you for when you feel stuck and can’t seem to wrap your head around something

we discussed in class, because this will happen. The only real remedy for this is to be persistent. The absolute worst thing you could do in this course is to give up. Everyone (really, everyone) finds at least some of this material difficult the first time they encounter it, but everyone that doesn't give up eventually figures it out. Your brain is going to get stretched and exercised in this course, and like any exercise it can be uncomfortable and awkward when you first start, but if you stay with it you'll get better and better; your brain will get stronger and stronger.

In terms of concrete tips for studying for this course, I'd suggest a few specific things:

- Read the assigned sections of the lecture notes and the textbook. Make a point to read them regularly. Get in the habit of finding a quiet place where you can go and just sit down and read without distractions. This means turning off (not just silencing, but completely shutting down) your phone or anything else that can interrupt your thoughts as you read.
- While you read, try to be incredulous. Always ask yourself why we're doing what we're doing. You should be asking yourself things like "What was the point of some particular assumption we made?" or "Why was *that* particular manipulation the right thing to do in the proof?" As you get more comfortable with proofs, ask yourself more questions while you read and study. "What would happen if *this* particular hypothesis was changed to *that*?" If you keep pushing with questions like this, you eventually get into uncharted territory, asking questions no one else knows the answer to, and sometimes finding the answers yourself – and that's where math really starts to get fun and exciting.
- Take notes, not just in class (which you definitely should do), but also while you read. The process of just writing things down sometimes helps strengthen the neural pathways in your brain and can help you to recall information and understand things better. (I'm by no means an expert in these things, but my understanding is that there's real science from the fields of neuroscience and education behind this idea.) When discussing this with a collaborator once, he agreed and said something that's stuck with me: "You don't read math with your eyes, you read it with your hands." When I was a student I had two notebooks for my math classes. In one I wrote "quick notes," things I jotted down quickly during lecture or things I'd write down while studying the textbook (things like definitions

and statements of theorems). In the other notebook I tried to very carefully write down an organized version of the quick notes I'd taken earlier, trying to fill in all of the details, explaining why a certain topic was useful or interesting to myself, etc. It was kind of like writing my own version of a textbook while I was learning the topic. While this was slow and time consuming, it was also very effective for studying, and probably why I type up lecture notes like this now.

- Start on assignments early and work on them regularly. You will be setting yourself up for failure if you wait until the day an assignment's due to start it. The homework in this class will be hard and sometimes it will take a serious investment in time in order to complete it. You should just be aware of this early on and plan for it. Once an assignment is available, plan to start on it the next day and try to do a problem or two each day. This will save you from a lot of stress compared to trying to do the whole assignment in a short amount of time.
- Ask questions. Ask lots of questions. Ask questions of me during class, outside of class, through email, etc. Feel free to ask questions about "the bigger picture," and not just the specifics of a particular homework problem. Sometimes seeing the big picture and connections between topics we do in class to other things can help you stay motivated when things seem hard.
- Take advantage of the Math Tutoring Center. The tutors in the MTC know how hard a class like this can be: they've been there, and they have come out the other side. They can help you with homework problems, with conceptual questions about the "why" of what we're doing, tell you where they used some of these proof techniques in other classes, etc. Sometimes I think some students are "too proud" and don't want to get help from the tutors at the MTC, but this is a bit silly. I'd encourage you to regularly go to the MTC, even if just to have a place to sit and work and study.

Some personal information

I am very excited to be teaching Math 250 this semester at Western Carolina, partly because it was a class like this that first made me interested in math. When I first came to college (at WCU, in fact), I wasn't a math

major but was instead a computer science major. As a CS major I'd taken algebra, calculus, and statistics and while I did well in those classes, I didn't particularly love them. To me, math had seemed like a tool that might be useful for other people that wanted to be scientists or engineers, but I wanted to be a computer programmer and didn't find my previous math courses very interesting. Like a lot of students, I did what I felt I needed to do to get a good grade in the class, but then promptly forgot everything once the semester had ended, because I didn't think I'd actually need to know that material for anything else.

It was in Benjamin Schultz's course (which I think was called *Logic and Proof for Computer Science*, but I might be wrong; this was a CS version of this class that isn't taught anymore) that I first started to really appreciate the idea of actually *proving* something. I was really amazed by the idea that we could prove something was true in math (and computer science), not just collect evidence that it was likely true, like you do in other sciences. This intrigued me enough to add math as a minor, and then at some point I decided to add math as a double major. My intention was still to be a computer programmer, but I thought having the math double major would at least look good on a résumé.

While in college I did a couple of summer internships for small regional companies, and became a little bit disillusioned with the job of a computer programmer. I enjoyed writing code, because I viewed each little problem to solve as a riddle to figure out and I thought that was fun, but the typical day-to-day of a 9am-5pm programming job seemed less fun. I didn't really know what to do, so my senior year I applied to the masters program in math at WCU (which unfortunately no longer exists), and that was when I started to *really* enjoy math. As an undergraduate you think you know what math is about, but then in graduate school you start learning *a lot* more stuff and see a whole mathematical world you didn't even know existed.

After the master's program at WCU, I knew I wanted to get a PhD in math and keep learning. I wound up going to Clemson for my PhD, and learned more math I didn't know even existed beforehand. After finishing at Clemson in 2014, I had teaching positions at Wake Forest University, Indiana University, and Bucknell University, before a tenure-track job at Western opened up. I was happy to be hired by WCU and have been here since Fall 2020.

Today I'm still interested in the topics I studied in my dissertation, and do research in related problems. Roughly, this has to do with the dynamics (the study of how things changes) on certain surfaces that have a very special kind of geometry. I like this stuff because it combines a

lot of different kinds of math: topology, geometry, abstract algebra, complex analysis, and dynamical systems are all things that pop up in what I study. I also still get to use my computer programming skills to help me in my research by writing simulations and conducting little experiments, or having code that looks for examples of things that are too difficult to do “by hand.” I really enjoy this stuff and would be happy to talk about it with students, or even work on a undergrad research project if any students are interested. (Or if students wanted to do an independent study about something that’s at least tangentially related to my interest, I’d be happy to do that as well.)

Chris Johnson
Fall 2023

Motivation for Logic and Proof

1.1 What are proofs?

The content of this course is likely quite different from the mathematics you've previously learned. Most likely the classes you've had up to this point have been focused on computation, and how to interpret and apply the values you've computed. For example, in calculus you learned how to compute a derivative by following a certain algorithmic procedure, and you learned that this quantity has interpretations as the slope of a tangent line, or the velocity of a moving particle. If you've taken a course in probability, you've likely learned how to compute the expected value of a random variable and how to interpret this value as something like the long-term average of samples of that random variable.

The topics in this course are more fundamental and foundational than the topics in these other courses. That is, instead of learning *how* to perform specific computations, we are more interested in questions about the *why* of the theory that underlies those computations. For example, instead of learning how to use the chain rule to compute a derivative, we are interested in why the chain rule works: is it just a "made up" rule that we all believe, or is there a logical rationale for explaining where the chain rule comes from and takes the form that it takes?

As a simpler example, in your earlier algebra courses you how to use the quadratic formula to determine the value of x that solves the equation

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

In those courses you were likely simply told that the x that satisfies this equation is given by the formula

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

and most likely only took this on faith. But where does this formula come from? Why does this formula work to give us the x that satisfies the quadratic equation above? To justify *why* this formula works, we need a proof.

Informally, a "proof" can be thought of as an argument that justifies beyond the shadow of a doubt why a statement is true. One important thing to realize, which students often struggle with at first, is that a correct proof doesn't simply provide evidence that a statement is likely true, but it unquestionably establishes that a statement must be true.

To illustrate the point, suppose you wanted to justify that the quadratic formula above was correct that it always gave the value of x solving the equation $ax^2 + bx + c = 0$. You might be tempted to look at several examples, making various choices of a , b , and c , and for each one use the quadratic formula to compute x , plug back into the original equation and see if it's satisfied or not. For example, maybe you choose $a = 2$, $b = -9$ and $c = 7$ and so consider the quadratic equation $2x^2 - 9x + 7 = 0$. Then you apply the quadratic formula to compute

$$x = \frac{-(-9) \pm \sqrt{(-9)^2 - 4 \cdot 2 \cdot 7}}{2 \cdot 2} = \frac{9 \pm \sqrt{81 - 56}}{4} = \frac{9 \pm 5}{4} = 1 \text{ or } \frac{7}{2}.$$

Finally we might plug these two values, $x = 1$ and $x = 7/2$, back into the equation to see that the equation is satisfied:

$$\begin{aligned} 2 \cdot 1^2 - 9 \cdot 1 + 7 &= 2 - 9 + 7 = 0 \\ 2 \cdot \left(\frac{7}{2}\right)^2 - 9 \cdot \frac{7}{2} + 7 &= \frac{49}{2} - \frac{63}{2} + \frac{14}{2} = 0 \end{aligned}$$

At this point we believe the quadratic formula worked *for this particular equation*. But what if we chose different values of a , b , and c ? You might want to construct another example and see if the formula holds or not, or maybe even write a simple computer program that verifies if the quadratic formula works for several randomly chosen examples. Even if we were to look at tens of millions of examples and the quadratic formula gave the correct solutions to the equation for each one, you have not yet *proven* that the quadratic formula is correct; all we've done is collect evidence that it probably is correct. No matter how many millions, billions, trillions, ... of examples we consider, there will always be infinitely many examples we haven't yet considered. The power of a (correct) proof of the quadratic formula is that it takes care of infinitely-many examples at once.

So, how do we actually prove the quadratic formula, or any other mathematical statement for that matter? In the case of the quadratic formula, what we want is an argument that shows that no matter what a , b , and c are, the x 's that solve $ax^2 + bx + c = 0$ are given by $x = \frac{1}{2}(-b \pm \sqrt{b^2 - 4ac})$. There are conceivably many different ways such an argument might work, but the proof below is a pretty standard one.

Remark.

You can very safely skip over the proof of the quadratic formula that we present below. We will actually start our study of proofs with simpler things, so don't let the proof below scare you off!

Theorem 1.1 (The quadratic formula). *The values of x which satisfy the quadratic equation*

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

where the coefficients a , b , and c are real (or complex) numbers are given by

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

Proof. We will suppose that x solves the equation, and show that it must have the form described. To do this, ultimately we would like to solve for x in the equation $ax^2 + bx + c = 0$. Let us first move the c to the right-hand side of the equation to obtain

$$ax^2 + bx = -c.$$

We may now divide both sides of the equation by a to obtain

$$x^2 + \frac{b}{a}x = \frac{-c}{a}.$$

If we add the quantity $\frac{b^2}{4a^2}$ to both sides of the equation, we then have

$$x^2 + \frac{b}{a}x + \frac{b^2}{4a^2} = \frac{-c}{a} + \frac{b^2}{4a^2}.$$

The purpose of adding this term, $\frac{b^2}{4a^2}$, to both sides of the equation is that we can now factor the left-hand side. We can easily verify, simply by "FOILing," that $(x + \frac{b}{2a})^2 = x^2 + \frac{b}{a}x + \frac{b^2}{4a^2}$. Thus we may rewrite our above equation as

$$\left(x + \frac{b}{2a}\right)^2 = \frac{-c}{a} + \frac{b^2}{4a^2}.$$

Before we go any further, let's add the fractions on the right-hand side together. This will require we get a common denominator, and so we

must multiply and divide the first term, $\frac{-c}{a}$ by $4a$ giving us $\frac{-c}{a} \cdot \frac{4a}{4a} = \frac{-4ac}{4a^2}$. Adding this the second term on the right-hand side gives us

$$\left(x + \frac{b}{2a}\right)^2 = \frac{-4ac + b^2}{4a^2}.$$

We now take the square root of both sides. Notice that as squaring removes negatives, when we take the square root we must compensate by considering both the positive and negative square root. This leads us to

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

At this point we may finally solve for x to obtain

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

Simplifying by taking the square root of the denominator and then adding our fractions together, we obtain the familiar quadratic formula

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

□

Do not worry if there are some details in the proof above that you don't fully understand, or a step that doesn't completely "click" the first time you read the proof. The key thing to recognize is that we have some precise way of verifying the quadratic formula is true for all, infinitely-many coefficients a , b , and c . Throughout the course we will see many, many examples of proofs and will actually start off with simpler proofs than this one, but we want to go ahead and acknowledge the power of a mathematical proof: it establishes beyond any doubt that a statement must be true.

It's worth taking a minute to dwell on the significance of the last sentence in the paragraph above, as this is what distinguishes mathematics from virtually every other discipline. In fields such as physics, chemistry, or economics, we may develop ideas (theories) to explain phenomena we see, but these are only models for those phenomena and we can never truly know if those models are actually correct. In a field such as physics, the best we can ever do is to collect lots of evidence to support the notion that the theories may be correct, but we can never know for sure.

For instance, Isaac Newton developed a theory of gravity that seems to be correct in most typical “day-to-day” scenarios, but it was later shown that the theory is not perfect and it had to be replaced by a different, more accurate theory of gravitation (Einstein’s theory of general relativity).

Instances such as the above do not happen in mathematics: once we provide a proof to establish that a statement is true and declare it to be a theorem, it is impossible to be disproven later.¹

1.2 What is logic?

Mathematical logic is a very big subject, and we will only barely scratch the surface of logic in this class. Our real focus is on proofs, but some proof techniques really require a bit of logic in order to justify their use. Logic in math is a very formal topic where we consider quantities that have either a true or false value, and various ways we can build complicated statements from these quantities and determine if those statements are true or false. While mathematical logic is interesting, we will only touch on the parts of it that are relevant for us in this class. However, after we’ve introduced logic, you would be in a good position to learn more about mathematical logic if you were so inclined.

It’s worth pointing out that mathematicians like to think that all of math is essentially a consequence of logic. We like to believe that everything we do starts with very simple statements we all agree are true (called *axioms*), and by combining basic rules about manipulating true statements to get another true statement, we eventually arrive at interesting statements we care about. Almost no one actually does math this way, though. With the exception of some very basic (or very specialized) cases, the way we think about constructing proofs and the practice of doing mathematics is often pretty far removed from the axioms and the mechanistic manipulation of statements. In the late 19th and early 20th century there was a push to make the math that mathematicians actually do more like formal logic, led by people like Alfred Whitehead and Bertrand Russell. Russell and Whitehead tried to carefully, methodically lay out the foundations of modern mathematics in terms of logic. This was a huge endeavour, and while it’s important in the history of math

¹It could be, however, that our theorem requires assumptions that are not explicitly stated – or perhaps not even known – at first! For example, the parallel postulate from your high-school geometry class is true in the context of Euclidean geometry, but is not true in the hyperbolic geometry that was discovered by Gauss, Bolyai, and Lobachevsky in the 19th century.

and logic, it shows that actually boiling everything down to the axioms is really unwieldy for modern mathematics.

We will take the view in this course that the basics of logic are worth being aware of, but are secondary to our primary goal of getting comfortable reading and writing proofs.

1.3 Why should we learn about logic and proofs?

Although it may not feel like it based on your course work up to this point, mathematics is really all about proof, and mathematicians are in the business of proving theorems. Even “users” of mathematics who may not directly prove their own theorems will rely on theorems proven by other people, and it’s good to be able to read proofs on your own to be able to understand them instead of treating them like a black box that just somehow magically works. Most students won’t become professional mathematicians, but they will become users of mathematics in various forms.

It should also be pointed out that constructing a proof is really an exercise in deductive reasoning, and this type of reasoning transcends writing proofs or even doing math in general. Writing a computer program is also an exercise in deductive reasoning, and I think the two (writing code and constructing proofs) complement each other well. In both cases, you learn how to think very carefully and how to be very detail-oriented; if you’re not, then your proofs are incorrect and your code is full of bugs. The more experience you get with this kind of careful, deliberate reasoning (regardless of the form it takes), the easier all other exercises in deductive reasoning become.

1.4 Applied math and pure math

As mentioned earlier, this course will be quite different from most of the other math courses you’ve taken up to this point. In particular, this is probably your first experience with a course in “pure math.” Most mathematicians think of math as roughly dividing into two categories, usually called applied math and pure math. *Applied math* is the mathematics that’s concerned with solving problems that arise in concrete applications, such as problems from physics or engineering. *Pure math*, however, is mathematics that’s not necessarily developed with a particular

application in mind. Sometimes doing pure math is a bit like exploring an uncharted territory just for the sake of exploring it. This is not to say that pure math doesn't have applications, it's just that applications aren't the driving force behind the math.

One example of where math that was traditionally considered pure later became applied has to do with number theory and cryptography. One of the things that interest people in number theory has to do with prime numbers, integers like 3 or 19 that aren't divisible by anything but 1 and themselves. We'll say a bit about prime numbers later, but one thing that's interesting is that all integers are basically built out of prime numbers: every integer has a *prime factorization*, which is a way of expressing that number as a product of prime numbers. For example, the prime factorization of 28 is $2 \times 2 \times 7$. It turns out that constructing prime factorization is pretty difficult, even on a computer. For very large numbers, we don't have any very efficient ways of computing the prime factorization. This all seems very "pure math," very theoretical, and people thought that's all it was for hundreds of years. In later part of the 20th century however, people realized that they could use this fact that prime factorizations are difficult to compute to construct cryptographic algorithms: ways of encoding information that makes it very difficult for someone to decipher a message we don't want them to read. Today you use these algorithms all the time, even if you don't realize it. Every time you purchase something on the Amazon, for example, information such as your credit card number is encrypted before it's sent to Amazon. (Lots and lots of other information you send and receive electronically is encrypted as well, but credit card numbers are one thing we can all relate too – you *really* don't want some random person to be able to determine your credit card numbers.)

We mention this distinction between pure math and applied math because it might be easy to lose sight of our goals in this class, since some of the things we'll be doing won't have easy-to-describe, immediate applications. In a more applied class like ordinary differential equations, you're always seeing applications and that can help keep the material interesting. In this class the applications are often less transparent and harder to appreciate. For this reason it might be best to view our assignments in the class as exercises to build up your deductive reasoning abilities, than as something that has an immediate, easy-to-see application.

Basic Proof Techniques

Pure mathematics is, in its way, the poetry of logical ideas.

ALBERT EINSTEIN

New York Times obituary of Emmy Noether

May 4, 1935

In this chapter we begin our study of proofs in earnest, discussing a few of the most common proof techniques through several basic examples. Along the way we discuss some of the common terminology that you will encounter in mathematics (e.g., the distinction between a theorem, a lemma, and a corollary), explain the standard way theorems and proofs are formatted, as well as introduce some typical symbols that you will often see in mathematics texts.

This chapter has a few exercises *given without solutions* and you are strongly encouraged to at least attempt each of the exercises as the best way to learn math is not simply by reading, but by actively working on problems. If a problem seems too difficult, frustrating, or time-consuming that's okay; you should at least attempt the exercises, however, and ask the instructor if you run into trouble or are confused by anything.

2.1 Direct proofs and counterexamples

2.1.1 First examples of direct proofs

We begin our discussion of proofs by viewing several examples of a technique often called “direct proof.” A *direct proof* is a proof that follows a very linear order where we begin by stating any assumptions being made (the hypothesis of the theorem), and proceeds by showing immediate consequences of those assumptions, then consequences of those consequences, etc., until we arrive at the desired conclusion.

To get started, let's see how to prove that the sum of two even numbers is again an even number. What you might be tempted to do here is look at a few examples to see if the claim is reasonable. Perhaps you haphazardly pick a few examples and compute $2 + 4 = 6$, $8 + 12 = 20$, and $4 + 28 = 32$. In each instance we added two even numbers and saw the result was even. Thus we might believe the claim that the sum of

two even numbers is even is reasonable, but this isn't a proof. There are infinitely-many pairs of even numbers we need to consider, and we've only checked three, so there are infinitely-many examples we haven't yet considered. How can we be sure that amongst those infinitely-many pairs of even numbers there isn't some choice which summed to an odd number?

What we need is some general way of expressing each even number. If we could do that, then perhaps we can show the sum of our two numbers must have the same kind of expression. Put another way, we need to be very precise about the definition of an even number. So, let's recall that an *even number* is an integer (whole number) which is a multiple of the number 2. That is, an even number is a number that can be written as $2n$ for some integer n . For example, 4 equals $2 \cdot 2$, 28 equals $2 \cdot 14$, and 10 equals $2 \cdot 5$.

The fact we used " n " in our definition above, instead of some other letter or symbol, doesn't really matter. Thus if we want to consider two different even numbers, we might write one of them as $2n$ and the other as $2m$.

With all of this in mind, let's now prove that the sum of two even numbers is even.

Theorem 2.1. *The sum of two even numbers is an even number.*

Proof. Let x and y be two even numbers. Thus we may express x and y as $x = 2n$ and $y = 2m$ for some integers m and n . Their sum then equals

$$x + y = 2m + 2n = 2(m + n).$$

As the sum $x + y$ is a multiple of 2, being $2(m + n)$, it is an even number. \square

Before we discuss the structure of the proof, let's talk about how it is presented. Observe that our theorem is clearly set apart in the text above so that when you read it, you can easily see the claim being asserted. Immediately below the statement of the theorem, we have the proof that appears and is also set apart from the text. We have the word 'Proof' which indicates we are about to start the proof, and at the end of the proof there is a small block, \square . This block indicates the end of the proof, making it easier for the reader to know where the proof has stopped and the main text resumes. This symbol, \square , is called a *Halmos tombstone*, and was introduced by Paul Halmos. The story goes that Halmos thought of unproven theorems as enemies to be vanquished, and the proof was what killed the enemy, so he marked their death with a tombstone.

Remark.

Almost any mathematical proof you read today will end with a Halmos tombstone, but occasionally you'll see the letters *QED* appear at the end of a proof instead. This is an acronym for the Latin expression *quod erat demonstrandum* which means "which was to be demonstrated." That is, the author is claiming they have shown what they set out to show and are now done.

Let's now carefully walk through the structure of proof of Theorem 2.1. Our goal is to show that if any two even numbers are added together, the result is an even number. So, we first suppose that x and y are any two even numbers. Notice that we are not choosing two particular even numbers: we *are not* just picking an example of an x or a y . Instead, x and y are two "placeholders" that could be any of the infinitely-many choices of even numbers. Since x and y are even numbers, we know we can write them as $2m$ and $2n$ for some appropriate choice of m and n . (It's easy to see that m is $\frac{x}{2}$ and n is $\frac{y}{2}$, but the actual values of m and n don't really matter.) Now, when we add these two even numbers together, we can write their sum $x + y$ as $2m + 2n$. As each of those terms has a factor of 2, we can factor it out and write the sum as $2(m + n)$. This means the sum $x + y$ is a multiple of 2, which is exactly what it means to be an even number! We have thus shown that adding *any* two even numbers together always results in an even number.

2.1.2 Counterexamples

Notice again that our proof applies for all of the infinitely-many choices of even numbers, whereas explicitly testing specific pairs of numbers will only show us the statement is true for those pairs of numbers. However, we can *disprove* a false claim by finding any single example where the claim is not true.

For example, suppose someone were to erroneously claim that the sum of two odd numbers is always odd. If we can find a single pair of odd numbers which do not add up to an odd number, then we have disproven the claim (or, if you'd prefer, we've proven the claim is false). Just a second of thought easily produces an example, such as $1 + 3 = 4$. Here we have two odd numbers, 1 and 3, whose sum is an even number.

Thus the claim that “the sum of two odd numbers is always odd” must be false.

In general, a single example that disproves a claim is called a counterexample. Perhaps right now counterexamples seem a little bit silly and a bit contrary to what we want to do: aren’t we trying to prove things, not disprove them? Counterexamples are often useful for us to test our understanding and help us to be sure we aren’t going down a rabbit hole trying to prove something that isn’t actually true. That is, we may make some conjecture based on our experience or intuition that we hope will be true. Instead of immediately trying to prove the statement we hope is true must be true, we can instead try to construct a counterexample to see if working on a proof will be futile. This may still seem strange at the moment, but we’ll come to appreciate counterexamples more as we’ve seen more examples.

2.1.3 More direct proof examples and exercises

As we stated above, a “direct proof” is one which follows a very linear chain of reasoning where we start from a given hypothesis and proceed step by step to produce a chain of simple consequences until we reach our desired conclusion. Often these kinds of proofs are derivations where our starting point is the definition of some quantity and we perform basic algebraic manipulations one at a time to get to our result. This essentially what we did in our proof that the sum of two even numbers is always even. Let’s see a couple more examples of this kind of proof where we are ultimately “unwinding” the definition.

Some examples you might have seen before in calculus involve justifying the various derivative rules, such as proving the derivative of $f(x) + g(x)$ is $f'(x) + g'(x)$. We’ll walk through this proof in just a minute, but let’s first spend some time recalling the definition of the derivative.

There are several ways different ways you can think about what the derivative really represents¹, but for simplicity we’ll use the simple interpretation that the derivative $f'(a)$ represents the slope of the line tangent to the graph $y = f(x)$ at the point $(a, f(a))$. This quantity seems difficult to work with directly, but the key idea of calculus is to approximate dif-

¹There’s a well-known article by William Thurston, one of the greatest mathematicians of the 20th century, called *On Proof and Progress in Mathematics* where he discusses how people think about mathematics. At one point he mentions *seven* different ways to think about the derivative of a function! That article is available online at <https://arxiv.org/pdf/math/9404236.pdf> and is recommended reading for anyone interested in becoming a mathematician.

difficult quantities with easier to compute ones. Approximating the slope of the tangent line at $(a, f(a))$ with the slope of the secant line through $(a, f(a))$ and another nearby point, say $(a + h, f(a + h))$, leads to the following definition:

$$f'(a) = \lim_{h \rightarrow 0} \frac{f(a + h) - f(a)}{h}.$$

(There is one caveat here: this limit may not exist! If the limit does exist, then we say the function is *differentiable at a* . If this limit exists for all choices of $x = a$ in the domain of the function, then we just say the function is *differentiable*.)

With the definition of the derivative at our disposal, we're now ready to prove some basic properties of the derivative. In the examples below the key idea will be to start from the definition of the derivative and manipulate it a little bit at a time until we reach our claim.

Theorem 2.2. *If $f(x)$ and $g(x)$ are two functions which are both differentiable at a , then their sum $f(x) + g(x)$ is also differentiable at a and the derivative equals $f'(a) + g'(a)$.*

Before we “formally” write down the proof of this theorem, let's talk through how such a proof might work.

Let's begin by simply writing out what the definition of the derivative of $f(x) + g(x)$ at a should be:

$$\lim_{h \rightarrow 0} \frac{f(a + h) + g(a + h) - (f(a) + g(a))}{h}$$

Our goal is to manipulate this a little bit at a time until we ultimately arrive at the definition of $f'(a)$ plus $g'(a)$. One simple thing we can do to get started is to distribute the minus that appears in the numerator:

$$\lim_{h \rightarrow 0} \frac{f(a + h) + g(a + h) - f(a) - g(a)}{h}.$$

Now we can rearrange the terms in the numerator, grouping the “ f terms” together and similarly for the “ g terms”:

$$\lim_{h \rightarrow 0} \frac{f(a + h) - f(a) + g(a + h) - g(a)}{h}.$$

We may now break our fraction up into the sum of two fractions:

$$\lim_{h \rightarrow 0} \left(\frac{f(a + h) - f(a)}{h} + \frac{g(a + h) - g(a)}{h} \right).$$

At this point we should see the light at the end of the tunnel. We simply use one of the basic properties of limits to split our one limit above into a sum of two limits:

$$\lim_{h \rightarrow 0} \frac{f(a+h) - f(a)}{h} + \lim_{h \rightarrow 0} \frac{g(a+h) - g(a)}{h}.$$

We finally observe that these two terms are, by definition, the derivatives of $f'(a)$ and $g'(a)$, thus our limit definition of the derivative of $f(x) + g(x)$ at a can be rewritten as $f'(a) + g'(a)$, and the theorem is proven.

Typically a proof you read in a book won't be as verbose as what we have above. If you were to look up the proof of this theorem in a calculus textbook, it might more succinctly give you something like the following:

Proof of Theorem 2.2.

$$\begin{aligned} & \lim_{h \rightarrow 0} \frac{f(a+h) + g(a+h) - (f(a) + g(a))}{h} \\ &= \lim_{h \rightarrow 0} \frac{f(a+h) + g(a+h) - f(a) - g(a)}{h} \\ &= \lim_{h \rightarrow 0} \frac{f(a+h) - f(a) + g(a+h) - g(a)}{h} \\ &= \lim_{h \rightarrow 0} \left(\frac{f(a+h) - f(a)}{h} + \frac{g(a+h) - g(a)}{h} \right) \\ &= \lim_{h \rightarrow 0} \frac{f(a+h) - f(a)}{h} + \lim_{h \rightarrow 0} \frac{g(a+h) - g(a)}{h} \\ &= f'(a) + g'(a) \end{aligned}$$

□

However, being succinct in proofs *is not* always a good thing, especially when you're just learning to write proofs on your own. Instead, you should feel free to include as much detail as you feel is necessary to convincingly describe each step in your proofs. More details help you to be sure you're not inadvertently making incorrect steps in your proofs as you construct the proof, and they also make it easier for someone reading your proofs to follow your reasoning. (There is nothing more annoying than a proof that does not adequately explain each of its steps! It's kind of like reading code to a computer program that doesn't have enough comments to clearly describe how the code works.)

Remark.

Typically when you see a proof of a theorem in a text it will occur immediately after the statement of the theorem. However, sometimes authors will delay the proof and spend some time first discussing the logic and rationale behind the idea of the proof before jumping into technical details, or they might talk about applications of the theorem before giving its proof. When this happens the proof should reference what theorem is being proven, just as our proof above was preceded by *Proof of Theorem 2.2*.

Exercise 2.1.

Test your understanding of the ideas of proof introduced so far by proving each of the following statements. Format your answers to these exercises as a theorem followed by a proof.

- (a) Prove that the sum of two odd numbers is always even. (Hint: It might be helpful to first try to carefully define exactly what an odd number is.)
- (b) Prove that the product of two even numbers is always even.
- (c) Prove that the product of two odd numbers is always odd.

2.1.4 Theorems, conjectures, lemmas, corollaries, etc.

We have used the word “theorem” a few times already, and certainly you’ve seen other things called “theorems” in previous mathematics courses (e.g., the Pythagorean theorem), but we haven’t properly defined exactly what that term means just yet. The word “theorem” is itself a bit odd, since it’s not something most people typically use in their everyday speech. So, what exactly is this strange word and where does it come from?

The origins of the word “theorem” seem to go back to the Greek word $\theta\epsilon\omega\rho\eta\mu\alpha$ (transliterating from Greek to Latin this becomes *theorema*)

which means something like *to be contemplated*. The use of this word in the mathematical sense appears to go back to Euclid who used this word to describe the results he was explaining in his book *Elements* which surveyed the geometry and number theory known to the ancient Greeks at that time.

Mathematicians have continued the tradition started by Euclid by referring to their results as “theorems.” That is, the word *theorem* just means a statement that has been proven. In some technical settings the word *theorem* is used to literally mean any true statement. For example, $2 + 2 = 4$ is technically a theorem. However, most of the time when we use the word “theorem” we are talking about a particularly important statement (such as the Pythagorean theorem, or the fundamental theorem of calculus).

Statements that have been proven to be true, but aren’t quite as significant are often referred to as a *proposition*. The exact distinction between what gets called a “theorem” and what gets called a “proposition” is a little bit hand-wavy, vague, and subject to the personal opinions of whoever is authoring what you’re reading. Ultimately they’re the same thing (a true statement), but people usually like to reserve “theorem” for especially important statements, and use “proposition” for statements that are not quite as impressive.

Some other terms you’ll often encounter in mathematics texts are “lemmas” and “corollaries.” These too really just mean a true statement (just like propositions and theorems), but are used in specific situations. Lemmas are usually even less impressive than propositions: they are in some ways the most minor of results. However, usually lemmas appear as stepping stones or building blocks to some other larger proposition or theorem. That is, a *lemma* often refers to a statement which has been proven, and while perhaps not so interesting in and of itself, is used in proving some other result. You can think of lemmas more as convenience to someone reading a proof: instead of having one very long and tedious proof, it can be easier to digest and follow if some smaller pieces of a proof are pulled out, and we call those smaller pieces lemmas.

Remark.

The word “lemma” is again another term that sounds strange because it’s not something we really use in everyday speech. It comes from the Greek $\lambda\epsilon\mu\mu\alpha$ (transliterated as “lemma”) which means

something to peel off, like the rind of a fruit. This might seem like a very strange choice of word at first glance, but in some ways it makes sense. Lemmas are usually stepping stones to more interesting theorems, and so they are sort of peeling away some of the difficulty in the theorem's proof.

A **corollary** is again a true statement (something whose proof has been, or can easily be, established), but which is often an immediate consequence of some other theorem. As an easy example, since part (c) of Exercise 2.1 showed that the product of two odd numbers is odd, an immediate consequence is that the square of an odd number is odd.

Corollary 2.3 (Corollary of Exercise 2.1, part (c)). *The square of an odd number is odd.*

As a corollary is supposed to be an immediate consequence of a theorem, it's not uncommon that proofs of corollaries don't appear in texts. (Sometimes they do, sometimes they don't – it's another thing that comes down to an author's personal preference.)

Remark.

Even if the proof of a corollary isn't provided in a text, it's good practice to try to prove the corollary to yourself. After you've gotten some practice with proofs, this becomes an easy exercise you can do to make sure you're understanding what you've been reading. If you can't prove a corollary, that's often a hint that you may not understand something as well as you thought you did.

There's one more word we'll go ahead and introduce, though it's one whose meaning you might be able to guess: *conjecture*. A **conjecture** is a statement which is believed to be true, but which has not yet been proven. Often mathematicians make lots and lots of conjectures in the process of doing research; in a certain way, mathematics is a continual loop of making conjectures and then trying to upgrade the conjecture to a theorem by proving it. You could think of a conjecture as an unsolved problem, and a mathematician's goal is to solve the problem by providing a proof the conjecture is true, or a counterexample to show it is false.

One famous example of a conjecture (unsolved problem) is the *Collatz conjecture*, which has to do with the following process. Suppose you are given a positive integer n . You produce a new integer, let's call it m , according to the following rules:

- if n is even, then m is $n/2$; and
- if n is odd, then m is $3n + 1$.

We then repeat the process using the new number m in place of n . That is, we start off with some given number and then start producing a list of numbers according to the procedure described above.

If we started with $n = 17$ for example, we would next compute $3 \cdot 17 + 1 = 52$; we then compute $52/2 = 26$, and then $26/2 = 13$, then $3 \cdot 13 + 1 = 40$, and so on. We keep repeating this process until we produce the number 1. The Collatz conjecture states that this process will always terminate (i.e., we will always eventually produce the number 1) regardless of what number start with. In the case of 17, the complete list of numbers we produce will be

$$17, 52, 26, 13, 40, 20, 10, 5, 16, 8, 4, 2, 1.$$

As simple and innocuous as this problem seems, no one actually knows if the Collatz conjecture is true or false, and the problem has been around since the 1930's. People have used computers to check if the Collatz conjecture holds for numbers up to 2^{68} which is approximately 2.95×10^{20} or

$$2,950,000,000,000,000,000,000.$$

Despite checking for such an incredibly large number of values, no one has been able to prove the Collatz conjecture will hold for all starting points: we have only checked finitely many numbers, so there are still infinitely many left to worry about.

2.1.5 Some practice in deductive reasoning

Typically when you create your own proofs of theorems, you won't be able to follow an algorithm procedure where you follow a preordained set of steps and arrive at the result. Proof writing is an exercise in deductive reasoning, where we start with some hypotheses (which in this case means assumptions we're making, such as that two numbers x and y are even integers), and try to deduce some desired conclusion (such as the sum of x and y must also be even). However, there often isn't a simple

path in getting from “I know these assumptions” to “I can deduce this conclusion.” Often you’ll have to spend some time thinking about what you know and deriving simpler conclusions that are (hopefully) a little bit closer to your ultimate desired conclusion. It takes some practice to get used to this, so let’s spend some time working with some concrete problems to help us get more practice with deductive reasoning.

Suppose we are given a typical 8×8 chess board, and a collection of dominos which are 2×1 . (That is, if each square of the chess board was one inch by one inch, our dominoes would be two inches by one inch.) We are interested in covering the chess board by these dominoes in such a way that the following properties are satisfied:

1. every square of the chess board is covered by dominoes;
2. no dominoes overlap; and
3. every domino is completely contained on the board (no dominoes are hanging off the edge of the board).


We’ll say the chess board is *perfectly covered* if there is some way of placing dominoes on the chess board so these conditions hold.

A first question that comes to mind is the following: Can a standard 8×8 chess board be perfectly covered? This statement is either true or it’s false, and we want to establish which one it is. To prove the statement is true, we only need to exhibit a perfect covering. To prove the statement is false, we’d need to have some kind of reason to explain why it is false: failing to find a perfect covering through trial and error does not show the statement is false, it only shows we don’t know how to build a perfect covering.

If you have access to a chess board and dominoes, it’s worthwhile to sit down and see if you can build a perfect covering. It turns out the 8×8 chessboard can be perfectly covered, and in fact it can be covered in lots of different ways. Simply constructing an example of such a perfect covering is not too difficult, and just presenting such an example proves that an 8×8 chessboard can be perfectly covered by 2×1 dominoes. Let’s present this as a theorem.

Theorem 2.4. *An 8×8 chess board can be perfectly covered by 2×1 dominoes.*

Proof. The configuration of dominoes presented below with four rows of dominoes placed vertically is an example of such a covering.



□

Let's make one little observation about the proof of Theorem 2.4 above. The theorem is the statement that some particular "thing" (in this case a perfect covering of an 8×8 chess board) must exist. One completely reasonable way to prove such a theorem is to simply give an example of the "thing" we're claiming exists. However, sometimes these existence proofs are less concrete; there may be some abstract reasoning that justifies the "thing" in question must exist, without actually stating what it is. The proof of the intermediate value theorem from calculus² is an example of such a thing.

Question.

Here's a question that's worth pondering, but you shouldn't be too worried if you can't answer it: it's not an easy question, but something you might enjoy thinking about. Now that we know perfect coverings of an 8×8 chessboard exist, how many are there? How would you even start to determine the answer to such a question? Questions like this that ask how many ways something can be done are part of a branch of mathematics called *combinatorics*. Combi-

²The intermediate value theorem states that if a function $f(x)$ is continuous on $[a, b]$, then for each value of y_0 between $f(a)$ and $f(b)$, there must exist an x_0 in the interval such that $f(x_0) = y_0$.

natorics is often introduced together with graph theory in a course called *discrete mathematics* that math and computer science majors are both required to take.

That problem was easy, so let's consider a similar problem that's a little bit harder:

Exercise 2.2.

Suppose you covered up the right-most column and top row of the chess board with strips of paper, effectively leaving a 7×7 chess board. Is it possible to perfectly cover this remaining 7×7 board?

Often in mathematics we can start off with concrete examples of problems we care about, solve those concrete problems, and then look to see if we can generalize our solutions to get a theorem that applies to lots of situations. For example, is there anything really special about the number 7 in Exercise 2.2? If so, what's special and what other numbers have that special property? Can we ask the same question about another $n \times n$ board where n has the same "special" property that 7 had?

Let's now consider a more difficult problem involving chessboards: Suppose we take an 8×8 chess board and cover up two opposite corners (e.g., the upper left-hand and lower right-hand corners). Can we perfectly cover the remaining 62 squares with 2×1 dominoes? You should spend some time thinking on this problem first before reading the solution below.

In order to solve this problem we will need to make an observation that is perhaps not immediately obvious, but which will supply us with the necessary tools to show the 8×8 board with opposite corners blocked *can not* be perfectly covered by 2×1 dominoes. These types of observations are often the key steps in proving more involved theorems and sometimes only come after spending a fair deal of time pondering the problem before the "Ah-ha!" moment occurs.

Notice that every 2×1 domino covers exactly one black and one white square, whether it's oriented horizontally or vertically. This means if we have any hope of having a perfect covering, then the remaining squares on our board must have just as many black squares as white squares. However, the opposite corners of a chessboard always have the same

color: two opposite corners will be black, and the other pair of opposite corners will be white. As a consequence, our chessboard with opposite corners blocked will have a different number of black and white squares, either 30 black and 32 white or vice versa. As a consequence, it's impossible to perfectly cover an 8×8 chessboard with opposite corners removed by 2×1 dominoes!

2.1.6 Divisibility

To get some more practice with direct proofs, let's look at a few more examples concerning divisibility. We say that an integer a *divides* an integer b if there exists an integer m so that $b = am$, and we sometimes denote this by writing $a|b$. We also say that b is a *multiple* of a when $a|b$. For instance, 3 divides 21 since $21 = 3 \cdot 7$, so we might write $3|21$ and say 21 is a multiple of 3. However, 3 does not divide 16: there is no integer m that solves $16 = 3m$. (Notice we are explicitly talking about integers when we talk about divisibility like this. There is a real number m solving $16 = 3m$, namely $m = \frac{16}{3} \approx 5.3333$, but this is not an integer.) If a does not divide b , then we write $a \nmid b$. Since 3 does not divide 16, $3 \nmid 16$.

Let's prove one important property of divisibility, namely that it is *transitive*.

Theorem 2.5 (Transitivity of divisibility). *If a , b , and c are integers with $a|b$ and $b|c$, then $a|c$.*

Before just jumping into the proof of Theorem 2.5, let's try to informally reason our way through how the proof might have to work. We are assuming $a|b$ and $b|c$, so that must mean there are some integers m and n so that $b = am$ and $c = bn$. But, since b equals am , we can replace the b that appears in our equation $c = bn$ to get $c = amn$! So, our number that we multiply a by to get c as a multiple of a is just mn . If we wanted, we could call this number p – i.e., we just define p to be mn – then we'd have $c = ap$ and we've shown that $a|c$. This is basically the proof of Theorem 2.5, but let's write it down properly.

Proof of Theorem 2.5. As $a|b$, there exists an integer m so that $b = am$. Similarly, since $b|c$ there is an integer n so that $c = bn$. Notice that as b equals am , we may express the equation $c = bn$ as $c = amn$. Defining p to be the product mn , we have $c = ap$, and so a divides c . \square

Exercise 2.3.

Suppose a , b , and c are integers with $a|b$ and $a|c$. Show that $a|(b+c)$.

Sometimes it can be helpful to measure how far one number is from dividing another. For example, even though 16 is not divisible by 3, it's only one greater than a number that is divisible by 3. In particular, 3 divides 15 and 16 is only one greater than 15. As another example, 9 does not divide 25, and is 7 greater than a number is a multiple of 9: $25 = 18 + 7 = 9 \cdot 2 + 7$. You could interpret this number, 7, as being how much is left over (or how much remains) if you were to try to divide 9 into 25.

In general, whenever we have two integers a and b with $a \neq 0$, we can write $b = am + r$ where r is one of $0, 1, 2, \dots, |a| - 1$.

Theorem 2.6 (The Division Algorithm). *Let a and b be integers with $a > 0$. There are unique integers m and r with $0 \leq r < a$ so that $b = am + r$.*

The division algorithm promises us two things which might be a little bit subtle the first time you see a theorem like this: it promises the existence of the integers m and r , and it also promises us that these two integers m and r are unique. That is, not only does it tell us that there is some choice of m and r so that we can write $b = am + r$ with $0 \leq r < a$, but it also promises us there is exactly one choice of m and one choice of r that will make this hold. Our proof of the division algorithm will thus need to contain two distinct parts: proving existence of some m and r so our equation holds, and proving uniqueness of the m and r .

Let's momentarily suppose b is positive to get the idea of the existence part of the theorem figured out. The idea will basically be to take the largest possible multiple of a (this is the am part of the expression $am + r$) that keeps $am \leq b$. Once that's done, we need to see how much more we need to tack on (this is the $+r$ part of $am + r$) to get from am up to b . This number we tack on could be zero (if am happened to equal b), but it can't be any bigger than $a - 1$ (we'll leave this part as an exercise for you to think about).

Let's work through a concrete example to see the idea in action. Let's suppose we wanted to write $b = 77$ as $am + r$ where $a = 12$. We consider the multiples of 12 which remain less than 73. These are 12, 24, 36, 48, 60, and 72. The largest of these was $72 = 12 \cdot 6$ (i.e., $m = 6$). Now, how much

do we need to add onto $72 = 12 \cdot 6$ to bump it up to 77? Well, we need to add 5 on (i.e., $r = 5$). This means we can write $77 = 12 \cdot 6 + 5$.

Exercise 2.4.

If a and b are positive integers and m is the largest positive integer with $am \leq b$, show that $b - am$ is no larger than $a - 1$. (Hint: The fact m is the *largest* integer with this property is important.)

To see that our m and r will be unique, we employ what will turn out to be a very common trick for these kinds of proofs (i.e., you'll see this over and over if you keep taking proof based math courses): we'll suppose there are actually two different choices of m and r (call them m_1, m_2 and r_1, r_2) satisfying our equation (i.e., $b = am_1 + r_1 = am_2 + r_2$ with $0 \leq r_1, r_2 < a$), and then see if they actually have to be the same thing or not. If it turns out they have to be the same (that is, if it is forced upon us that $m_1 = m_2$ and $r_1 = r_2$), then the m and r are unique.

In the case of our example $77 = 12 \cdot 6 + 5$ above, we're saying $m_1 = 6$ and $r_1 = 5$, but supposing there was some other choice of m_2 and r_2 so that $77 = 12m_2 + r_2$ with $0 \leq r_2 < 12$. Both of these expressions equal 77, so they equal each other:

$$12 \cdot 6 + 5 = 12m_2 + r_2$$

Let's perform a bit of arithmetic now, by subtracting $12m_2$ from the right-hand side over to the left-hand side:

$$12 \cdot 6 - 12m_2 + 5 = r_2.$$

Factoring the 12 out, we can write this as

$$12 \cdot (6 - m_2) + 5 = r_2.$$

Now we make an observation: we claim $6 - m_2$ will need to be zero. To see why that must be the case, let's momentarily write N for $6 - m_2$, so our equation above is $12N + 5 = r_2$. If $N \neq 0$, it's either positive or negative. Let's suppose for a second it's positive. Then when we write $r_2 = 12N + 5$, we must have that $12N \geq 12$, so r_2 is then at least 17. But r_2 was supposed to be between 0 and 11, so this can't happen. Similarly, if $N \leq -1$, then $12N \leq -12$, so $r_2 = 12N + 5 \leq -7$. This is also a problem since r_2 was

supposed to be non-negative. Thus the only possibility, if we're really going to have $0 \leq r_2 < 12$ is that $N = 0$. But that means $6 - m_2 = 0$, so $m_2 = 6$. But then we have the expression $r_2 = 12 \cdot 0 + 5 = 0 + 5 = 5$, so $r_2 = 5$.

It's a little bit to digest the first time you see this idea, but try to carefully work through the example above, or (even better) make up your own concrete example and work through the same kind of reasoning with the numbers you come up with. All we're doing in the proof below is repeating the same kind of reasoning, just leaving the quantities a , b , m and r as variables instead of making a choice of a particular number for each one.

Proof of Theorem 2.6. Suppose momentarily that $b \geq 0$. Let m be the largest integer such that $am \leq b$. Now define the number r to be $b - am$. Since $b \geq am$, the number $r = b - am$ we defined is non-negative. Now notice

$$am + r = am + b - am = b,$$

and so we have found integers m and r so that $b = am + r$.

We still need to show that r satisfies $0 \leq r < a$, and that m and r are the unique integers satisfying $b = am + r$ with $0 \leq r < a$. As observed above, r is non-negative so $r \geq 0$. To see that $r < a$, notice that if r was greater than or equal to a , then we could write $r = a + s$ for some $s \geq 0$. But then

$$b = am + r = am + a + s = (m + 1)a + s$$

Since $s \geq 0$, though, $(m + 1)a = b - s \leq b$. But this is impossible since m is chosen to be the largest number with $am \leq b$. (It's impossible because $m + 1$ is necessarily bigger than m .)

To show m and r are the unique pair of integers satisfying $b = am + r$ with $0 \leq r < a$, suppose there was another such pair of integers, call them m_2 and r_2 . Since $b = am + r$ and $b = am_2 + r_2$, we must have $am + r = am_2 + r_2$, which we may rewrite as $am - am_2 + r = r_2$, or $r_2 = a(m - m_2) + r$. Since r and r_2 are both between 0 and a (not including a), we claim that we must have $m - m_2 = 0$. If this were not the case, then either $m - m_2 \geq 1$ or $m - m_2 \leq -1$.

If $m - m_2 \geq 1$, then $r_2 = a(m - m_2) + r \geq a + r \geq a$, but this is impossible since $0 \leq r_2 < a$. If instead $m - m_2 \leq -1$, then $r_2 = a(m - m_2) + r \leq -a + r < 0$, and this is also impossible since $0 \leq r_2 < a$. Together these mean that $m - m_2 = 0$, so $m = m_2$.

Once we know that $m = m_2$, our expression $am + r = am_2 + r_2$ from above becomes $am + r = am + r_2$, or simply $r = r_2$. \square

Exercise 2.5.

Suppose a is an integer and $0|a$. Show that $a = 0$.

2.2 Proofs by induction

The next proof technique we will discuss is referred to as *induction* and the idea of a proof by induction will be familiar to anyone who has used “recursion” in computer programming. Before describing induction in the general mathematical sense, we will see some recursive ideas from computer science to help us warm up to mathematical induction.

2.2.1 Warmup: The quicksort algorithm

One common problem in computer science has to do with sorting a list of values. The idea is simply that we have some collection of data that is “out of order” and we wish to put it in an order from the least elements to the greatest elements.

As a simple example, suppose we have a website that keeps track of certain users which have a first and last name. This could be a course management system used by a university, such as Canvas or Blackboard, where the users are students in a course. Students may add or drop the course at various points and the list of users for the site might not be in alphabetical order: there’s no particular reason why students would register for a course in alphabetical order. The instructor of the course might wish to have an alphabetized list of students, however, and so our list of users needs to be sorted. This begs the question of how do we sort this data: how do we instruct a computer to take unordered data and put it in order?

There are several different *algorithms* (specific lists of instructions for accomplishing some general task) that computer scientists have developed for sorting data, and different algorithms have different strengths and weaknesses. Some algorithms may sort data very quickly (which is usually desirable), but might need to use a lot of memory in the process (which is usually undesirable); other algorithms might use less memory, but be slower. There are lots and lots of sorting algorithms out there, but to illustrate the idea of recursion (and induction) we will consider the *quicksort algorithm*.

The idea with quicksort is that we have some unordered collection of data which we will essentially break into two halves: one half will be less than some particular element in the list (let's say the first element), and the other half is greater than that element. We will then order each half, and put the halves together. If we can do this, we've ordered the list.

The key thing in the description of the quicksort algorithm above is that we take our list, break it into smaller lists *and then order those smaller lists*. Now, how are we going to order the smaller lists? Well, we could just repeat the process on those smaller lists! In general, we take our lists and break them into half, then break those "sublists" into smaller lists, and break those into smaller lists, etc. Eventually we get to a point where our sublist consists of only one element, which is already sorted. Once we get down to this "base case" we work our way back up to the original list, putting everything back in order as we go.

To be concrete, let's try to apply this process to order the following list of numbers:

5, 8, 3, 2, 7, 4

We first pull off the first element of the list, 5, and look at the remaining list of numbers: 8, 3, 2, 7, 4. Now we take these lists and look through each entry and compare it to 5, storing all the elements less than 5 in one list (this would be 3, 2, 4) and the elements greater than 5 in another list (8, 7). If we could sort these lists, giving us 2, 3, 4 and 7, 8, we could then easily put everything back together to get our ordered list 2, 3, 4, 5, 7, 8. But now we have to sort these smaller lists, such as 3, 2, 4. How do we do that? Let's "rinse and repeat," applying the same procedure to 3, 2, 4.

We'll pull off the first entry, 3, and then divide the remaining entries into lists which are smaller than 3 (just 2) and greater than 3 (just 4). Now notice that each of these little lists of a single element is necessarily already sorted! So, we take the list of smaller entries, concatenate 3 to it, then add the list of larger entries. This gives us 2, 3, 4 which is sorted.

We can do the same thing to our list 8, 7 to turn it into 7, 8. Now we take our smaller list, 2, 3, 4, attach 5 to it, then attach the larger list 7, 8, to get the ordered list of numbers:

2, 3, 4, 5, 7, 8

For the sake of concreteness, here is the quicksort algorithm implemented in Python:

```
def quicksort(data):
    if len(data) <= 1:
        return data

    pivot = data[0]
    lessThan = [datum for datum in data[1:] if datum <= pivot]
    greaterThan = [datum for datum in data[1:] if datum > pivot]

    lessThan = quicksort(lessThan)
    greaterThan = quicksort(greaterThan)

    return lessThan + [pivot] + greaterThan
```

Running this code does indeed organize the list of data we give it:

```
>>> quicksort([5, 8, 3, 2, 7, 4])
[2, 3, 4, 5, 7, 8]
```

Remark.

It's completely fine if you're not familiar with Python, or writing code in general. This is just meant to be an example of how the idea of induction can be applied. You should make an effort to understand the idea of the quicksort algorithm, but you certainly don't need to feel like you need to understand the code above if you don't want to.

There are two crucial properties of the quicksort algorithm we've presented:

- we apply the same algorithm to a smaller list of data than what we are originally given; and
- the algorithm immediately returns its value, instead of running again, if it's given a list of one (or zero) elements.

That is, we continually apply the same algorithm to smaller and smaller data sets, and the algorithm stops once we reach the smallest possible data set. This last step is extremely important: without a condition like this, the algorithm might try to repeat forever and ever. In general, for recursive functions it's important to have a "base case" where the function stops calling itself, otherwise we could have an "infinite recursion" that never ends.

2.2.2 Mathematical induction in proofs

Proof by induction is a technique where we prove infinitely-many statements by taking one statement, and expressing in terms of a simpler statement. The idea is that we're basically breaking down something complicated into simpler, simpler, simpler pieces until we get all the way down to the simplest situation which we can prove. This simplest situation is called the *base case*, and it's something that we'll have to verify using some other technique (such as a direct proof). We then have an *inductive step* that tells us how to relate a more complicated statement to a simpler one. One crucial thing is that the inductive step needs to be moving us a little bit closer to the base case each time.

There are two flavors of mathematical induction that we'll discuss, called *weak induction* and *strong induction*. The names are a little misleading because they make it seem like "strong induction" would somehow be more powerful than "weak induction," but that's not the case. The difference between them is what kind of assumption we're going to make in the induction step; or, put another way, how much we're breaking down our complicated statement.

In the case of weak induction we're basically relating our current state in the proof to just the previous states, whereas in strong induction we're relating the current state to all the previous states. The distinction will make more sense after we see some examples.

Remark.

Often when people discuss induction they don't distinguish between "strong induction" and "weak induction" and just use the word "induction" to mean whichever technique they're using.

2.2.3 Weak induction

Let's try to be a little bit more clear about precisely what (weak) induction is before we start jumping into examples. Let's say that we have infinitely-many related statements we wish to prove are true, and let's say we have some way of ordering these statements. That is, there is a first statement, a second statement, a third statement and so on. Whatever the statements happen to be, let's denote them by S_1 for the first statement, S_2 for the second statement, etc.

For example, suppose we wanted to show that for all positive integer n we had that $n \leq 2^n$. The infinitely-many statements here are $1 \leq 2^1$, $2 \leq 2^2$, $3 \leq 2^3$, and so forth. In term of our notation, we're writing S_1 for the statement that $1 \leq 2^1$, and S_2 for the statement that $2 \leq 2^2$, and S_3 is our short-hand for $3 \leq 2^3$, etc. Here we have infinitely-many statements, that have a natural "order" to them: S_n is the statement $n \leq 2^n$ for each positive integer n .

Now, how can we prove all of these infinitely-many statements? For any *finite* number of cases we can verify each individual statement, but that doesn't prove all infinitely-many statements. It's conceivable that the first 500 million statements are true, but the next one is false; verifying a finite number of statements always falls short of proving infinitely-many.

Induction gives us a tool for proving these infinitely-many statements. In particular, the principle of *weak induction* states that all of the statements S_n will be true provided two things hold: S_1 is true, and we can show that if S_{n-1} is true then S_n must be true for each $n \geq 2$.

Remark.

Informally, induction is like climbing a ladder. Is it true that you can climb a ladder as high as you'd like? Well, if you can get on the ladder, and you know how to go from one rung of the ladder to the next, then you should be able to climb as high as you'd like. Here getting on the ladder is like the base case, and going from one rung of the ladder to the next is the inductive step.

We should think of the first statement, S_1 (e.g., $1 \leq 2^1$) as being the simplest case, and it's something we can often verify directly. This is

called the *base case*. The second part, showing that if S_{n-1} is true, then S_n must be true is called the *inductive step*.

So, in the case of our claim that $n \leq 2^n$ above, the base case that S_1 is true (i.e., $1 \leq 2^1$) is something we directly check, since 1 is less than 2. For the inductive step, we need to justify that if $n - 1 \leq 2^{n-1}$ for $n \geq 2$, then it must be true that $n \leq 2^n$ as well. So, we will suppose that n is at least 2 $n - 1 \leq 2^{n-1}$ is true; we assume this is something we have already shown. (This statement that we're assuming $n - 1 \leq 2^{n-1}$ – or more generally, that we assume S_{n-1} is true in proving S_n – is referred to as the *inductive hypothesis*.)

Now we need to use this assumption to show that n must be less than 2^n . To do this, we need to try to manipulate the expression we have assumed is true (that $n - 1 \leq 2^{n-1}$) to get the expression we want to show is true (that $n \leq 2^n$). To do this, let's do something really simple: let's just add 1 to each side of the expression:

$$\begin{aligned} n - 1 &\leq 2^{n-1} \\ \implies n - 1 + 1 &\leq 2^{n-1} + 1 \\ \implies n &\leq 2^{n-1} + 1 \end{aligned}$$

Above we used the symbol \implies for the first time, but we'll see it several more times in this course, and you may have seen it in previous courses. This symbol means "implies," and when we write something like $A \implies B$ that means that B is a logical consequence of A : if A is true, then B must be true. For example, above we had $n - 1 \leq 2^{n-1} \implies n - 1 + 1 \leq 2^{n-1} + 1$. This means that if $n - 1 \leq 2^{n-1}$ is true, then a logical consequence is that $n - 1 + 1 \leq 2^{n-1} + 1$ is true. We'll be a bit more precise about the idea of implication when we discuss logic later, but for now we're meaning this in an intuitive kind of way.

Anyway, back to our proof. We now know that $n - 1 \leq 2^{n-1}$ is true, then a consequence is that $n \leq 2^{n-1} + 1$. Let's notice that since $n \geq 2$ we have $1 \leq n - 1$, so we can write

$$n \leq 2^{n-1} + 1 \leq 2^{n-1} + n - 1$$

But now we can use our inductive hypothesis again! Since we've assumed $n - 1 < 2^{n-1}$ we have

$$2^{n-1} + n - 1 \leq 2^{n-1} + 2^{n-1}$$

and now we can do some arithmetic. We may rewrite $2^{n-1} + 2^{n-1}$ as $2 \cdot 2^{n-1}$ which equals 2^n . Putting all of this together, we have that $n \leq 2^n$ provided $n - 1 \leq 2^{n-1}$. Let's rewrite this as a theorem and its proof.

Theorem 2.7. *For every positive integer n , the inequality $n \leq 2^n$ is satisfied.*

Proof. We will induct on n . The base case, $1 \leq 2^1$ is verified as $1 \leq 2$. Now suppose that for $n \geq 2$, we have $n - 1 \leq 2^{n-1}$. We then have

$$\begin{aligned} n - 1 &\leq 2^{n-1} \\ \implies n - 1 + 1 &\leq 2^{n-1} + 1 \\ \implies n &\leq 2^{n-1} + 1 \end{aligned}$$

We now observe that since $n \geq 2$, we have $n - 1 \geq 1$ and so have the following inequalities:

$$\begin{aligned} n &\leq 2^{n-1} + 1 \\ &\leq 2^{n-1} + n - 1 \end{aligned}$$

We now apply our inductive hypothesis that $n - 1 \leq 2^{n-1}$ to extend our string of inequalities above,

$$\begin{aligned} n &\leq 2^{n-1} + 1 \\ &\leq 2^{n-1} + n - 1 \\ &\leq 2^{n-1} + 2^{n-1} \\ &= 2 \cdot 2^{n-1} \\ &= 2^n. \end{aligned}$$

The theorem has been established. □

Remark.

In the proof of Theorem 2.7 above, we started off by saying “We will induct on n .” You’ll often see this phrase at the start of a proof by induction, and this does two things: it tells us we’re about to see a proof by induction; and it also tells us that n is the variable that indexes all of the statements we have. Recall in our description of the principle of weak induction above, we had infinitely many statements S_n for $n = 1, n = 2, n = 3$, and so on. Sometimes in other contexts we may want our variable to be something different like k instead of n . All we’re doing by saying we’re going to “induct on n ” is informing the reader that n is the variable in our statements. Sometimes you’ll see this at the start of a proof by induction, but sometimes you won’t: it just depends on the preference of whoever

wrote the proof you're reading.

Let's walk through a proof by (weak) induction to get some familiar formulas for summations. Recall that the notation

$$\sum_{i=1}^n f(i)$$

is short-hand for the sum

$$f(1) + f(2) + f(3) + \cdots + f(n),$$

where $f(i)$ is some expression involving i . We can start this sum at some place besides $i = 1$. If we wanted to start it at $i = 0$, we'd have

$$\sum_{i=0}^n = f(0) + f(1) + f(2) + \cdots + f(n).$$

For example,

$$\sum_{i=1}^n i = 1 + 2 + 3 + \cdots + n$$

and

$$\sum_{i=0}^n 2^i = 2^0 + 2^1 + 2^2 + \cdots + 2^n,$$

We would like to have a way to easily compute what this quantity is without manually adding up things 2^0 plus 2^1 plus 2^2 and so on, which can take a long time.

The first of these formulas you may have seen before (for example, it comes up in calculus when evaluating a Riemann sum), and so you may know that this first expression has a nice formula:

$$\sum_{i=1}^n = \frac{n(n+1)}{2}.$$

Let's try to prove that this formula holds for every n . Using induction isn't the only way to do this, but since we want an example of an inductive proof, we'll use induction right now.

Let's observe that the formula holds for the simplest possible case of $n = 1$. In that case we have

$$\sum_{i=1}^1 i = 1 = \frac{1 \cdot 2}{2}.$$

That is, if we replace all the n 's on both sides of the equation above by 1, we can directly verify the formula holds. We could also directly verify the formula for $n = 2$, or $n = 3$, or $n = 4$, and so on. But we want to verify the formula *for all* n , and we can't do that directly since there are infinitely-many choices of n . The idea with induction though, is that we can cook up an argument that says if we know the formula's true for one n , then we can show it must be true for the next n as well. Thus once we know the formula holds for $n = 1$, we'll have an argument to show it holds for $n = 2$. Once it's established for $n = 2$, the same argument will show the formula holds for $n = 3$, and so on. Thus we get to prove the formula holds for infinitely-many n at once.

So, let's say we want to verify the formula for some choice of n , but have already verified it for the previous value, $n - 1$. That is, we are assuming we know

$$\sum_{i=1}^{n-1} i = \frac{(n-1)n}{2}$$

(Here we're replacing all the n 's on the right-hand side of our formula by $n - 1$.) Now, how do we "upgrade" the formula we've assumed holds for $n - 1$ to hold for n ? Well, let's try to rewrite our summation in terms of $n - 1$:

$$\sum_{i=1}^n i = 1 + 2 + 3 + \cdots + (n-1) + n$$

Notice the first n terms are exactly $\sum_{i=1}^{n-1} i$, and so we can write

$$\sum_{i=1}^n i = \sum_{i=1}^{n-1} i + n.$$

But we have assumed our formula holds for the sup up to $n - 1$. Plugging that in we have

$$\sum_{i=1}^n i = \frac{(n-1)n}{2} + n.$$

Now let's add these quantities on the right-hand side together:

$$\begin{aligned}
 \sum_{i=1}^n i &= \frac{(n-1)n}{2} + n \\
 &= \frac{(n-1)n}{2} + \frac{2n}{2} \\
 &= \frac{n^2 - n}{2} + \frac{2n}{2} \\
 &= \frac{n^2 - n + 2n}{2} \\
 &= \frac{n^2 + n}{2} \\
 &= \frac{n(n+1)}{2}
 \end{aligned}$$

Finally, just for the sake of presentation, let's state our formula above as a theorem and then write down its proof.

Theorem 2.8. *For every positive integer n ,*

$$\sum_{i=1}^n i = \frac{n(n+1)}{2}.$$

Proof. We will prove this by induction on n . The base case, $n = 1$, is verified by direct computation:

$$\sum_{i=1}^1 i = 1 = \frac{1 \cdot (1+1)}{2}.$$

For the inductive step, suppose the theorem has been proven for $n - 1$ with $n \geq 1$. That is, we suppose

$$\sum_{i=1}^{n-1} i = \frac{(n-1) \cdot (n-1+1)}{2} = \frac{(n-1)n}{2}.$$

Now to establish the proof for n , we rewrite our summation as

$$\sum_{i=1}^n i = n + \sum_{i=1}^{n-1} i.$$

The second term on the right-hand side can be rewritten using our formula to obtain

$$\sum_{i=1}^n i = n + \frac{(n-1)n}{2}.$$

Finally, we simply add n and $\frac{(n-1)n}{2}$ by getting a common denominator and combining like-terms to obtain

$$n + \frac{(n-1)n}{2} = \frac{2n}{2} + \frac{n^2 - n}{2} = \frac{2n + n^2 - n}{2} = \frac{n^2 + n}{2}.$$


Factoring out an n in the numerator on the last line above completes our proof,

$$\sum_{i=1}^n i = \frac{n(n+1)}{2}.$$

□

Remark.

The numbers that arise as $\sum_{i=1}^n i$ for increasing choices of n (i.e., $n = 1$, then $n = 2$, then $n = 3$, ...) form a sequence of numbers called the *triangular numbers*. These numbers (such as 1, 3, 6, 10, ...) are called “triangular numbers” because if you were to create a triangle where you had 1 object (say a bowling pin, or red Solo cup) on at the top of the triangle, then two objects on the row beneath that, and three on the row beneath that, and so on down to n rows, the n -th triangular number corresponds to how many objects (bowling pins, red Solo cups, ...) you will need to construct your triangle. For instance, you need ten objects to create a triangle with four rows, because 10 is the fourth triangular number, $\sum_{i=1}^4 i = \frac{4 \cdot 5}{2} = 10$.



In the previous example the formula for our summation $\sum_{i=1}^n i$ was given to us, and it was our job to prove the formula was correct. But what if the formula had not been handed to us? What if we were just told to figure out what the formula was supposed to be ourselves, and then prove our conjectured formula was actually correct? Let's see what to do in that situation by considering the other example we had mentioned, $\sum_{i=0}^n 2^i$.

Finding the correct formula for our summation is a two-step process: first we need to make a conjecture about what a formula for this expression should be, and then we need to actually prove our formula is correct.

Making a conjecture about what the formula should be can sometimes be difficult, but let's try to look at a few simple examples and see if we notice any patterns. Let's just write down the first few cases of the sum above.

- If $n = 0$, then our expression is just 2^0 which just equals 1.
- If $n = 1$, our summation is $2^0 + 2^1 = 1 + 2 = 3$.
- If $n = 2$, our summation is $2^0 + 2^1 + 2^2 = 1 + 2 + 4 = 7$.
- If $n = 3$, our summation is $2^0 + 2^1 + 2^2 + 2^3 = 1 + 2 + 4 + 8 = 15$.

Have any "obvious" patterns emerged? After looking at these numbers for a minute, you might realize that they look like one less than a power of 2:

$$1 = 2 - 1$$

$$3 = 4 - 1$$

$$7 = 8 - 1$$

$$15 = 16 - 1$$

So, it might be reasonable to conjecture that the sum $\sum_{i=0}^n 2^i$ equals $2^{n+1} - 1$, as this formula agrees with the first four values we've computed above. But how do we *prove* this formula is actually correct? This is a case where we're naturally set up for induction since there's a "simplest" case (when $n = 0$), and each other situation has a simpler case before it.

To be more precise, in our base case we can easily see the formula holds just by direct computation:

$$\sum_{i=0}^0 2^i = 2^0 = 1 = 2^{0+1} - 1.$$

Now to justify the formula holds for larger n , suppose we've already proven the formula is true for $n - 1$. That is, suppose we have already shown that $\sum_{i=0}^{n-1} 2^i = 2^n - 1$. Our goal is to rewrite the summation for n in such a way we can take advantage of this formula. The "obvious" thing to do is to pull off the largest term, and apply the formula on the remaining summation:

$$\sum_{i=0}^n 2^i = 2^n + \sum_{i=0}^{n-1} 2^i = 2^n + 2^n - 1.$$

Now let's just add $2^n + 2^n$ together. To do this we could factor 2^n from each term leaving us with $2^n(1+1)$, and of course $1+1 = 2$ so this becomes $2^n \cdot 2$ or simply 2^{n+1} . And this is basically our proof.

Theorem 2.9. *For each integer $n \geq 0$, we have*

$$\sum_{i=0}^n 2^i = 2^{n+1} - 1.$$

Proof. We will prove this by induction on n . The base case, $n = 0$, is verified by direct computation:

$$\sum_{i=0}^0 2^i = 2^0 = 1 = 2^{0+1} - 1.$$

Now suppose the formula has been proven for $n - 1$. That is, suppose we have shown

$$\sum_{i=0}^{n-1} 2^i = 2^{(n-1)+1} - 1 = 2^n - 1$$

for $n \geq 0$. To verify the formula for n we compute

$$\begin{aligned} \sum_{i=0}^n 2^i &= 2^n + \sum_{i=0}^{n-1} 2^i \\ &= 2^n + 2^n - 1 \\ &= 2^n(1 + 1) - 1 \\ &= 2^n \cdot 2 - 1 \\ &= 2^{n+1} - 1 \end{aligned}$$

□

Exercise 2.6.

Here are some exercises to consider to help you get used to induction problems.

(a) Show that for every positive integer n , 5 divides $6^n - 1$.

(b) Show that for every non-negative integer n the following equality holds:

$$\sum_{k=0}^n 3^k = \frac{3^{n+1} - 1}{2}.$$

(c) Show that for every integer satisfying $n \geq 4$ we have $2^n < n!$.

2.2.4 Strong Induction

Strong induction is very similar to weak induction: the idea is still that we have a base case we can verify, and then want to show that the statement is true if “simpler versions” of that statement are true. However, instead of just supposing the preceding statement is true, we suppose *all* of the previous statements were true.

To be more precise, the principle of *strong induction* says that statements S_n , for integers $n \geq 1$, will be true if the first k statements S_1, S_2, \dots, S_k are true (the base cases – notice there may be more than one base case), and if S_1, S_2, \dots, S_{n-1} are true, then S_n is also true. This will be a little clearer after a few examples.

One place where strong induction is useful is in finding solutions to “recurrence relations.” A *recurrence relation* is a way of defining a sequence of numbers such that the “next” element in the sequence depends on previous elements. One famous example that you may have seen before is the sequence of *Fibonacci numbers*. The Fibonacci numbers are numbers F_n defined as follows. We define F_1 and F_2 to both be 1, and then each subsequent number F_n is defined to be the sum of the previous two Fibonacci numbers, $F_n = F_{n-1} + F_{n-2}$.

For example, the first few Fibonacci numbers are

$$1, 1, 2, 3, 5, 8, 13, 21, 34, 55, \dots$$

and our notation is that F_n is the n -th Fibonacci number. E.g., $F_7 = 13$.

As another example of a recurrence relation, consider the sequence of numbers a_n defined by

$$\begin{aligned} a_1 &= 1 \\ a_2 &= 3 \\ a_n &= 2a_{n-1} - a_{n-2} \text{ for } n \geq 3 \end{aligned}$$

Using this expression for a_n , we can compute the next few entries in the sequence are

$$\begin{aligned} a_3 &= 2a_2 - a_1 = 6 - 1 = 5 \\ a_4 &= 2a_3 - a_2 = 10 - 3 = 7 \\ a_5 &= 2a_4 - a_3 = 14 - 5 = 9 \end{aligned}$$

At this point you likely see a pattern and may conjecture that $a_n = 2n - 1$. Let's see if we can justify this using induction.

If we were to use weak induction, we would only be supposing $a_{n-1} = 2(n-1) + 1$, so when proving the inductive step we would have

$$\begin{aligned} a_n &= 2a_{n-1} - a_{n-2} \\ &= 2(2(n-1) + 1) - a_{n-2} \\ &= 2(2n - 2 + 1) - a_{n-2} \\ &= 2(2n - 1) - a_{n-2} \\ &= 4n - 2 - a_{n-2} \end{aligned}$$

and now we're a little bit stuck because we were only making an assumption about a_{n-1} instead of a_{n-2} . This should be a hint that strong induction might be the more useful tool.

If we were to use strong induction, we would suppose $a_k = 2k - 1$ for each $k = 1, 2, 3, \dots, n-2, n-1$. Now when manipulating a_n , we could write

$$\begin{aligned} a_n &= 2a_{n-1} - a_{n-2} \\ &= 2(2(n-1) - 1) - (2(n-2) - 1) \\ &= 2(2n - 2 - 1) - (2n - 4 - 1) \\ &= 2(2n - 3) - (2n - 5) \\ &= 4n - 6 - 2n + 5 \\ &= 2n - 1 \end{aligned}$$

which is what we wanted to show. Let's now write this down precisely.

Theorem 2.10. *In the sequence of numbers defined by the recurrence relation*

$$\begin{aligned}a_1 &= 1 \\a_2 &= 3 \\a_n &= 2a_{n-1} - a_{n-2} \text{ for } n \geq 3\end{aligned}$$

the n -th term, a_n , is given by $2n - 1$.

Proof. We will use the principle of strong induction. For our base cases, we observe

$$a_1 = 1 = 2 \cdot 1 - 1 \quad \text{and} \quad a_2 = 3 = 2 \cdot 2 - 1$$

and so the formula holds for $n = 1$ and $n = 2$.

Now we will suppose the formula $a_k = 2k - 1$ holds for each k from 1 through $n - 1$. Then we may rewrite the n -th term in the sequence as

$$\begin{aligned}a_n &= 2a_{n-1} - a_{n-2} \\&= 2(2(n-1) - 1) - (2(n-2) - 1) \\&= 2(2n - 2 - 1) - (2n - 4 - 1) \\&= 2(2n - 3) - (2n - 5) \\&= 4n - 6 - 2n + 5 \\&= 2n - 1\end{aligned}$$

and the theorem is proven. □

Remark.

It will turn out that weak and strong induction are actually equivalent: anything you can prove using one, you can also prove with the other. However, it can sometimes be easier to think through the proof and the reasoning will be “cleaner” if we use strong induction. If you wanted you could always just use strong induction or always just use weak induction, but the proofs might be unnecessarily cumbersome for some types of problems.

Exercise 2.7.

- Show that the sum of the first n Fibonacci numbers equals $F_{n+2} - 1$.
- Show that the sum of the first n odd-indexed Fibonacci numbers,

$$\sum_{k=1}^n F_{2k-1} = F_1 + F_3 + F_5 + \cdots + F_{2n-1},$$

always equals the n -th even-indexed Fibonacci number, F_{2n} .

2.2.5 Prime numbers and the fundamental theorem of arithmetic

Recall that we say an integer a divides an integer b , denoted $a|b$, if $b = am$ for some integer m . We also denote this relationship by a and b by saying b is a multiple of a , or that a is a divisor of b . Notice that for every integer n , n always divides itself and 1 always divides n . If the only positive divisors of an integer are 1 and the integer itself, we say that integer is a prime number. For example, the first few prime numbers are

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

(By convention we do not consider 1 to be a prime number. This may seem like an odd convention, but we'll see why this convention is adopted in just a little bit.)

If a positive integer is not prime, we call it a composite number. For example, 4, 6, 8, 9, 12, and 15 are composite numbers since they have non-trivial divisors; that is, we can write composite numbers as products

of other numbers:

$$4 = 2 \cdot 2$$

$$6 = 2 \cdot 3$$

$$8 = 2 \cdot 4$$

$$9 = 3 \cdot 3$$

$$12 = 3 \cdot 4$$

$$15 = 3 \cdot 5$$

In fact, we can write composite numbers as products of prime numbers, though we may need to multiple the same prime by itself multiple times. For example, 8 equals 2^3 and 12 equals $2^2 \cdot 3$. The collection of prime numbers which multiply together to give us our composite numbers is called the *prime factorization* of the number, and it turns out every positive integer has a unique prime factorization. That is, not only can composite numbers be written as a product of primes, but there's exactly one way to do this. This fact is known as the *fundamental theorem of arithmetic*, and can be proven using strong induction.

Theorem 2.11 (The fundamental theorem of arithmetic). *Every positive integer $n \geq 2$ has a unique prime factorization.*

Proof. Let $n \geq 2$ be a positive integer. We will use the principle of strong induction to show n has a unique prime factorization. The base case is simply when $n = 2$. In this case 2 is prime and the prime factorization is simply 2. For the inductive hypothesis we will suppose that for each $k = 2, 3, 4, \dots, n - 1$, the number k has a unique prime factorization.

There are two cases to consider: whether n is prime or composite. If n is prime, the its prime factorization is simply that number itself. This is unique because if it were not, the number could be written as a product of other numbers, but then would not be prime.

If n is composite, then by definition it can be written as a product of other numbers: say $n = ab$ where neither a nor b equals 1. (Again, the fact n is composite implies we must be able to write n as a product like this.) Notice that a and b are less than n and greater than 1, and so by the inductive hypothesis they each have a unique prime factorization. This product of a and b thus provides a prime factorization of n . To see that the factorization is unique, suppose two different factorizations existed: say

$$n = p_1^{e_1} p_2^{e_2} \cdots p_i^{e_i} = q_1^{f_1} q_2^{f_2} \cdots q_j^{f_j}$$

where the p 's and q 's above are distinct prime numbers. □

2.3 Proofs by contradiction

The last proof technique that we'll mention for now is called *proof by contradiction*, and it gives us a way to prove something must be true by essentially showing it's impossible for it to be false. This is quite a bit different from our previous proof techniques, and it can seem a little weird and it can take some time to wrap your head around it.

The way that proofs by contradiction work is by doing something that seems extremely counterintuitive at first. The very first step in a proof by contradiction is to assume that what you want to show is true is actually false. You then show that if the statement were false, some other statement which you know must be true would have to be false as well – *this* is the contradiction.

As we'll discuss more when we talk about logic in the next chapter, all of our statements are either true or false. They must take on exactly one of these values: our statements can not be both true and false. So, if we had some chain of reasoning that told us a statement we already know must be true would have to be false, then there must be a faulty argument with our chain of reasoning. The idea is that if our first step is assuming some statement were false, and all the intermediate steps are logically consistent and we arrive at saying some true statement is false, then the faulty step must be our initial assumption. Which would mean the original statement, which we assumed was false, is actually true. Again, this kind of reasoning is pretty weird the first time you see it, but it will make more sense after we see some examples.

2.3.1 Basic examples

Let's begin by showing something that would be pretty hard to prove directly. We will show that for each integer n , if $n^3 + 5$ is odd, then n must be even. To do this, let's suppose the opposite. For the sake of providing a contradiction, we will suppose that if $n^3 + 5$ is odd, then n is also odd. This would mean $n = 2m + 1$, and so

$$\begin{aligned}n^3 + 5 &= (2m + 1)^3 + 5 \\&= (2m)^3 + 3 \cdot (2m)^2 + 2 \cdot 2m + 1 + 5 \\&= 8m^3 + 12m^2 + 4m + 6 \\&= 2(4m^3 + 6m^2 + 2m + 3)\end{aligned}$$

Notice this means $n^3 + 5$ would be even, but this contradicts our original assumption that $n^3 + 5$ was odd. Thus our conclusion must be that n is

necessarily even. Let's now write this up properly as a theorem and its proof.

Theorem 2.12. *For each integer n , if $n^3 + 5$ is odd, then n is even.*

Proof. Let n be an integer and suppose that $n^3 + 5$ is odd. Suppose for the sake of contradiction that n was also odd, and so $n = 2m + 1$ for some integer m . We may then rewrite $n^3 + 5$ as

$$\begin{aligned} n^3 + 5 &= (2m + 1)^3 + 5 \\ &= (2m)^3 + 3 \cdot (2m)^2 + 2 \cdot 2m + 1 + 5 \\ &= 8m^3 + 12m^2 + 4m + 6 \\ &= 2(4m^2 + 6m^2 + 2m + 3). \end{aligned}$$

However, this is an even number and so contradicts the earlier assumption that $n^3 + 5$ was odd. Hence n must have been an even number. \square

As another example, we will show that if x and y are two odd numbers, then $x^2 + y^2$ can never be a perfect square. To do this, we will suppose that instead $x^2 + y^2$ was a perfect square, we would have to have a contradiction. Since x and y are odd, we can write $x = 2m + 1$ and $y = 2n + 1$. Then $x^2 + y^2$ is

$$(2m + 1)^2 + (2n + 1)^2 = 4m^2 + 4m + 1 + 4n^2 + 4n + 1 = 4(m^2 + n^2 + m + n) + 2$$

Let's observe that this number is even, but *is not* divisible by 4 since the remainder when dividing by 4 would be 2. Notice though, that if a perfect square z^2 is even, then z is necessarily even. In fact, since $z = 2k$ we have $z^2 = 4k^2$, so an even number squared is divisible by 4. But now we have a contradiction as our calculation above shows that $x^2 + y^2$ would be a perfect even square which is not divisible by 4. Hence if x and y are both odd, we must conclude that $x^2 + y^2$ is not a perfect square.

To write all of this up as a theorem and a proof, it might be convenient to first prove a little lemma about our observation that perfect even squares are divisible by 4.

Lemma 2.13. *Every even perfect square is divisible by 4.*

Proof. Let n be an integer which is an even perfect square. As n is a perfect square, $n = m^2$ for some integer m . As an odd number squared is necessarily odd, we must have that m is an even number and can write $m = 2k$ for some integer k . We then have $n = m^2 = (2k)^2 = 4k^2$, and so 4 divides n . \square

Theorem 2.14. *If x and y are both odd numbers, then $x^2 + y^2$ is never a perfect square.*

Proof. Let x and y be odd numbers. We may write $x = 2m + 1$ and $y = 2n + 1$ for some integers m and n . Suppose for the sake of contradiction that $x^2 + y^2$ is a perfect square. That is, we are assuming $x^2 + y^2 = z^2$ for some integer z . Observe we may rewrite $x^2 + y^2$ as

$$\begin{aligned} x^2 + y^2 &= (2m + 1)^2 + (2n + 1)^2 \\ &= 4m^2 + 4m + 1 + 4n^2 + 4n + 1 \\ &= 4(m^2 + n^2 + m + n) + 2. \end{aligned}$$

Notice this is an even number. As this quantity also equals z^2 , we then have z^2 is an even number. However, by Lemma 2.13, z^2 must be divisible by 4. Our expression above contradicts this, though, as it expresses $z^2 = 4k + 2$ for some integer k . Hence we must conclude that $x^2 + y^2$ is not a perfect square. \square

Let us do one more basic example of a proof by contradiction before we turn to more interesting examples. As you know from experience, when solving equations in algebra, solutions to seemingly “simple” equations may not be “simple” numbers. However, in some problems we may only be interested in particular simple solutions. For example, if we have a problem an expression where the variables represent quantities like people, or cars, or the number of objects produced by a factory, we may only be interested in whole number solutions to our equations. One interesting question to consider, then, is whether a given equation has any integer solutions.

Let’s show that there are no integer solutions to the equation $18x + 6y = 1$. To do this, we will suppose that some integer solution (meaning both x and y are integers) exists, and derive a contradiction. To do this, we observe that if $18x + 6y = 1$ then we may rewrite the equation as $6(3x + y) = 1$ which we can further write as $3x + y = \frac{1}{6}$. And herein lies our contradiction: if x and y are integers, then any product or sum with other integers always yields another integer, not a fraction.

Theorem 2.15. *There are no integers x and y that solve the equation $18x + 6y = 1$.*

Proof. Suppose for the sake of contradiction that x and y were two integers satisfying $18x + 6y = 1$. By dividing both sides of the equation by 6, this would imply that x and y are integers solving $3x + y = \frac{1}{6}$, yet this is

impossible as sums and products of integers are integers. In particular, $3x + y$ is an integer, whereas $\frac{1}{6}$ is not. This contradiction establishes that $18x + 6y = 1$ has no integer solutions. \square

2.3.2 Irrational numbers

We will now give a famous example of a proof by contradiction to establish that certain numbers can not be written as ratios of whole numbers. In particular, we will show that $\sqrt{2}$ can not be written as a ratio of integers by supposing it can be, and arriving at a contradiction.

Theorem 2.16. *The number $\sqrt{2}$ can not be written as a ratio of integers.*

Proof. Suppose for the sake of contradiction that $\sqrt{2}$ was equal to p/q where p and q are integers. Without loss of generality, we may suppose that p and q do not have any common factors. (If p and q *did* have common factors, then we could cancel those factors and replace p and q by two numbers without common factors.) We would then have $p^2/q^2 = 2$, which we could rewrite as $p^2 = 2q^2$. This would imply p^2 is an even number, and so p would also be an even number. As p is even, we may write $p = 2k$ for some integer k . Our equation $p^2 = 2q^2$ above would then become $4k^2 = 2q^2$. Dividing both sides of this last equation by 2, we would have $2k^2 = q^2$ showing that q^2 is even, and so q must also be even. We have now contradicted the fact that p and q were taken to not have any common factors, as both numbers are even and so have a common factor of 2. \square

Theorem 2.16 shows that some numbers can not be written as a ratio of integers. In general, if a number *can* be written as a ratio of integers, we call that number ***rational***. For example, $3/7$, $-2/3$, and 7 are all rational numbers. Numbers that can not be written as a ratio of two integers are called ***irrational***. Our proof above only shows that one particular irrational number exists, but the proof can be extended to show many other square roots are also irrational.

Exercise 2.8.

Show that if an integer n is not a perfect square, then its square root is irrational.

It will turn out that “most” real numbers are actually irrational, but the proof of this fact will have to wait until we’ve established some set theory. Despite this, we can go ahead and prove lots of interesting properties of rational and irrational numbers.

Exercise 2.9.

- (a) Show that the sum and product of two rational numbers is always rational.
- (b) Is it true that sums and products of irrational numbers are always irrational? If so, give a proof. If not, find a counterexample.
- (c) Show that the sum of a rational and irrational number is always irrational.
- (d) Show that the product of a non-zero rational and irrational number is always irrational.

2.3.3 The infinitude of primes

We give one more important example of a proof by contradiction, showing that there are infinitely-many prime numbers. Recall that a prime number is a positive integer whose only divisors are 1 and itself, such as 2 or 19. The proof below is attributed to Euclid and uses the fact that every composite number can be written as a product of primes to show that any finite list of prime numbers is necessarily incomplete.

Theorem 2.17. *There are infinitely-many prime numbers.*

Proof. Suppose for the sake of contradiction that there were only finitely-many prime numbers, say p_1, p_2, \dots, p_n . Consider the number

$$n = p_1 \cdot p_2 \cdot \dots \cdot p_n + 1.$$

Notice this number is not divisible by any of p_1, p_2, \dots, p_n as dividing n by any p_i will have a remainder of 1. However, n is either prime or

composite. If n is prime, then our list of primes was incomplete as n is not an element of the list. If n is composite, then our list of primes is incomplete since n has a prime factorization, yet none of the prime numbers p_1 through p_n appear in its factorization. In either case, we have contradicted the claim that the finite list p_1 through p_n contains all prime numbers. \square

Symbolic logic

In this chapter we begin our formal study of logic.

3.1 Propositions and Predicates

3.1.1 Propositions

We had previously mentioned that a *statement* in mathematics is a claim that has either a true or a false value. One technical point here is that a statement must definitively have a true or false value and that value can not change. For example, “*Albert Einstein won a Nobel prize on November 9, 1922*” is a statement because it has a definitive true/false value (true); and “*Mitt Romney defeated Barack Obama in the United States’ 2012 presidential election*” is also a statement since it has a definitive true/false value (false). These true/false sentences are sometimes also referred to as *propositions*.

In symbolic logic, our goal is often to determine ways of combining statements to get other statements in such a way that we can determine the true/false value of a complicated statement from the true/false value of the simpler statements that are used to build the complicated statement. The actual content of the statement is irrelevant in this case, and only its true/false value matters.

For example, consider the following string of statements:

All men die.
Socrates is a man.
Thus Socrates will die.

or these statements:

All lizards have scales.
An iguana is a lizard.
Thus iguanas have scales.

Both of these examples can be rephrased as

All x have property P .
 y is an x .
Thus y has property P .

The line of reasoning in both examples above is the same; the actual content of what property P is or what x or y is doesn't matter. All that matters is that we have one statement that can be used to infer another, and we can substitute in various values for the variables involved later and we will have a valid line of reasoning.

3.1.2 Variables and predicates/open sentences

Often our sentence will have *variables* in them, symbols that represent an as-yet-undetermined quantity. This could be a place holder for a certain type of number (such as an integer, or maybe more specifically an even integer, or a prime number), or for some more abstract mathematical object like a set or an element of a set. For example, one sentence with a variable is $x^2 - 4 = 0$. A declaration such as this, $x^2 - 4 = 0$, is not true or false until we make a choice of what the variable is. These declarations with unassigned variables are referred to as *open sentences* or *predicates*. We will often refer to a predicate involving a variable x as $P(x)$ or $Q(x)$. If $P(x)$ is the proposition $x^2 - 4 = 0$, then $P(2)$ and $P(-2)$ are true, but $P(3)$ is false. Of course, our predicates may have multiple variables and we may write things such as $P(x, y)$ or $Q(a, b, c)$ for predicates with two or three variables.

3.1.3 The universal and existential quantifiers

The truth value of a predicate depends on the value of variables we give it, but we can often "upgrade" a predicate to a statement by specifying the statement must hold for all choices of the variable. For example, if $P(x)$ is the predicate " $x^2 + x$ is even", then $P(x)$ will be true for all integers x . We get a statement by *quantifying* our predicate. A *quantifier* is a symbol that indicates that the predicate that follows the *quantifier* should hold for certain choices of variables. We have two quantifiers: the *universal quantifier*, denoted \forall means that the predicate should hold for all choices of variable. For example, if $P(x)$ is a predicate, the statement $\forall x P(x)$ is read "for all x the predicate $P(x)$ holds." Sometimes we further restrict our variables based on what type of variable we allowed to plug in for x . For example, the predicate " $P(x): x^2 + x$ is even" holds for all integers, so we may write \forall integers $x, P(x)$ to indicate we are only claiming $P(x)$ is true for all integers x .

The other quantifier we will use is the *existential quantifier*, denoted \exists and is used to mean "there exists." For instance, consider the claim that $3x > 2$. This is true for some values of x , but not all x . Thus the statement

$\forall x, 3x > 2$ would be false since it does not hold for all x . However, since it holds for at least one x (such as $x = 1$), the statement $\exists x, 3x > 2$ is true. This is read “There exists an x so that $3x > 2$.”

When a variable in a predicate is attached to a quantifier, such as $\forall x, P(x)$ or $\exists y, Q(y)$, then we say the variable is a ***bound variable***. If a variable is not bound, it is called a ***free variable***. Notice that predicates only get upgraded to statements when all of their variables are bound.

Predicates that involve several variables may have different quantifiers attached to the variables. For example, if $P(x, y)$ is a predicate involving both x and y , then $\forall x \exists y, P(x, y)$ is read “for every x there exists a y so that $P(x, y)$.” As an example, for every real number x there exists a real number y so that $y + x > 0$ and so the statement $\forall x \exists y, y + x > 0$ is true.

In general the order in which quantifiers appear is very important and changing the order of the quantifiers can affect whether the statement is true or false. For example, suppose $P(x, y)$ is the predicate $y = x^2$. The statement \forall real numbers $x \exists$ real number $y, P(x, y)$, which would be read “for every real x there exists a real y so that y equals x^2 ”, is true. However, we swap the quantifiers to be $\exists y \forall x P(x, y)$ our statement becomes “there exists a real y so that for every real number $x, y = x^2$ ” and this statement is false. (No matter what real you take y to be, there will be a number whose square is not y .)

3.2 Logical operations and truth tables

Whenever we have a statement, we can try to combine it with other statements to form a new statement, and try to determine the truth or falsity of the new statement based on the truth or falsity of the initial statements. In a certain sense, all of mathematics boils down to starting from very simple statements we all agree on (called *axioms*) and combining them in various ways to deduce more interesting statements (*theorems*). We’ll have more to say about this point of view of mathematics later, but for the moment we just want to get acquainted with the basic ways we can combine mathematical statements together.

As we introduce the various logical operations, we will start seeing *truth tables*, which give us a concise way of describing whether a complicated statement built from simpler statements is true or false, based on whether those simpler statements are true or false.

3.2.1 Conjunction (and)

One of the simplest operations we can perform between two logical statements P and Q is to consider their *conjunction*, which is also referred to as a *logical and*. That is, given two statements P and Q , we construct a new statement $P \wedge Q$ (read as P and Q) which will be true precisely when both P and Q are true, and false if at least one (or both) of P or Q is false. We summarize this with the following table, which is our first example of a truth table.

P	Q	$P \wedge Q$
T	T	T
T	F	F
F	T	F
F	F	F

To be precise about what the table above means, a *truth table* contains one column per simpler statement (the P and Q) used in building a more complicated statement ($P \wedge Q$), and then we have one row per possible truth value of each of our simpler statements. Here since we have two simple statements, P and Q , and each one takes on one of two truth values, there are a total of four rows.

As an intuitive example of the logical and, suppose P is the statement *It will rain today*, and Q is the statement *There will be a rainbow in the sky today*. Their conjunction would be the statement *It will rain today and there will be a rainbow in the sky today*. The only way for this longer statement to be true is if both statements *It will rain today* and *There will be a rainbow in the sky today*. If it either doesn't rain (whether there's a rainbow or not), or there's not a rainbow (where it rained or not) is false, then the conjunction $P \wedge Q$ is false.

3.2.2 Disjunction (or)

The *disjunction* of two statements, denoted $P \vee Q$ and read P or Q , is true if either P or Q (or both) is true. The corresponding truth table is

P	Q	$P \vee Q$
T	T	T
T	F	T
F	T	T
F	F	F

So, for instance, if P was the statement *The interior angles of a triangle add to 180°* and Q was the statement *Every four-sided figure is a rectangle*, then the statement $P \vee Q$ would be true because P is true, even though Q is false. The statement $P \wedge Q$, however, would be false.

3.2.3 Negation

Every logical statement P has a negative, denoted $\neg P$ and read *not P*, which simply has the opposite truth value as P . The corresponding truth table is rather boring:

P	$\neg P$
T	F
F	T

For example, the negation of *It will rain today* is *It will not rain today*. Notice that exactly one of P or $\neg P$ is true, and the other is false.

We can combine the negation operation with our conjunction and disjunction operations to negate a conjunction or a disjunction. That is, we can consider $\neg(P \wedge Q)$ and $\neg(P \vee Q)$. Of course, the truth tables for these statements are easily determined from the truth tables of $P \wedge Q$ and $P \vee Q$:

P	Q	$P \wedge Q$	$\neg(P \wedge Q)$	$P \vee Q$	$\neg(P \vee Q)$
T	T	T	F	T	F
T	F	F	T	T	F
F	T	F	T	T	F
F	F	F	T	F	T

It is sometimes helpful to observe that these negations can be rewritten as follows: The negative of $P \wedge Q$, $\neg(P \wedge Q)$, can be written as $\neg P \vee \neg Q$. That is, the “opposite” of P and Q is *not P or not Q*. Notice that if we compare the truth tables of $\neg P \vee \neg Q$ and $\neg(P \wedge Q)$, they are identical:

P	Q	$\neg P$	$\neg Q$	$P \wedge Q$	$\neg(P \wedge Q)$	$\neg P \vee \neg Q$
T	T	F	F	T	F	F
T	F	F	T	F	T	T
F	T	T	F	F	T	T
F	F	T	T	F	T	T

Similarly, the negation of $P \vee Q$, $\neg(P \vee Q)$, is $\neg P \wedge \neg Q$: the negative of P or Q is *not P and not Q*.

Exercise 3.1.

Verify that the truth table of $\neg(P \vee Q)$ agrees with the truth table of $\neg P \wedge \neg Q$.

The take-away from this is that we can “distribute” a negative across a statement, but conjunctions and disjunctions are exchanged: *ands* become *ors*, and *ors* become *ands*.

Exercise 3.2.

Suppose P , Q , R , and S are logical (true/false) statements.

- (a) Construct a truth table to verify that the negation of $P \wedge (Q \vee \neg R)$ is $\neg P \vee (\neg Q \wedge R)$.
- (b) Determine the negation of $(P \wedge \neg Q) \vee (\neg R \vee S)$, then construct a truth table to verify that you have the correct negation.

3.2.4 Negating quantifiers

Suppose that our statement comes from a proposition with a quantifier. For example, consider the statements $\forall x, P(x)$ and $\exists x, Q(x)$. What would the negations of these statement be?

If the statement $\forall x, P(x)$ is not true, then it *is not* the case that $P(x)$ is true for all x . That must mean there is at least one x where $P(x)$ is false. That is,

$$\neg(\forall x, P(x)) \text{ is equivalent to } \exists x, \neg P(x).$$

Similarly, if $\exists x, Q(x)$ is not true, then it is not the case that there is some x making $Q(x)$ true and so for every x , $Q(x)$ is false:

$$\neg(\exists x, Q(x)) \text{ is equivalent to } \forall x, \neg Q(x).$$

3.2.5 Implication

Perhaps the most important of our logical connectives (these operations that combine simpler statements to get more complicated statements) is

implication. **Implication** expresses a relationship between two statements that if one statement is true, it must be the case that the other statement is true. You should think of this as an “if-then” kind of statement. Notice we have seen several examples of these kinds of statements when we discussed theorems and proofs in the last chapter. For instance, we had theorems such as “If x is an integer, then $x^2 + x$ is an even integer.” Here we have “if-then” type statement where if the first part of the statement (x is an integer) is true, then it follows that the second part ($x^2 + x$ is an even integer) is also true.

In general, for two statements P and Q , we write $P \Rightarrow Q$ (read “ P implies Q ”) to mean that if P is true, then Q is necessarily true. For example, if P is the statement x is an integer and Q is the statement $x^2 + x$ is an even integer, then $P \Rightarrow Q$ is a true statement.

We can’t just arbitrarily string together statements P and Q with $P \Rightarrow Q$ and get a true statement, though. For example, if we continued to let P denote x is an integer but change Q to be $x^2 + x$ is an odd integer, then $P \Rightarrow Q$ is false: it does not follow that $x^2 + x$ is odd because x is an integer.

So, what are the various ways in which $P \Rightarrow Q$ can be true or false? Or, put another way, what is the truth table of $P \Rightarrow Q$? Our examples above should tell us that if P and Q are both true, the $P \Rightarrow Q$ is true as well. If P is true and Q is false, however, then $P \Rightarrow Q$ must be false. But what about the other two possibilities, where P is false and Q is either true or false?

It confuses everyone the first time they hear it, but “false implies true” and “false implies false” are both true statements! The intuition for this is something like the following: the statement $P \Rightarrow Q$ should mean that if P is true, then Q must be true as well. However, it does not inherently mean that Q must be false (or true) if P is false. To help this make sense, here is an example shamelessly stolen from the Mathematics Stackexchange¹:

Suppose a parent tells their child that if they eat their vegetables at dinner, then they will get a dessert.

(So, P is the statement “child eats their vegetables” and Q is the statement “child gets a dessert”.)

Our question is when is the implication $P \Rightarrow Q$ true or false, or put another way, when did the parent lie to the child by

¹<https://math.stackexchange.com/questions/431639/intuition-of-implication-in-propositional-logic>

making this statement?

The only way the parent lies is if the child eats their vegetables, but does not get a dessert. That is, if P is true but Q is false, then $P \Rightarrow Q$ would be false.

If the child eats their vegetables and gets a dessert, then the parent did not lie. That is, if P is true and Q is also true, then $P \Rightarrow Q$ is true.

In the remaining situations, the child does not eat their vegetables (i.e., P is false). Did the parent lie if the child does not get a dessert (Q is false)? No, the parent was consistent with what they claimed: eating vegetables gets a dessert, and so the child shouldn't be surprised to not get a dessert if they don't eat their veggies.

But what if the child does not eat their vegetables (P is false), yet still gets a dessert (Q is true)? Did the parent lie to the child in that situation? No, the parent only promised a dessert if the child ate their veggies: eating vegetables guarantees a dessert. But there could be another reason the child received a dessert besides eating their veggies (maybe they did all their chores early, or completed all of their homework).

Putting all of this together, our truth table for implication is

P	Q	$P \Rightarrow Q$
T	T	T
T	F	F
F	T	T
F	F	T

3.3 Converses, Equivalences, and Contrapositives

In the last section we learned about implication, which is a way for us to describe when one statement implies another: $P \Rightarrow Q$ means that if P is true, then Q must be true as well. Sometimes this is stated by saying that P is a sufficient condition for Q , or that Q is a necessary condition for P .

3.3.1 Converses

Any time we have an implication $P \Rightarrow Q$ we may consider the statement $Q \Rightarrow P$ which is known as the *converse* of $P \Rightarrow Q$. In general, the validity of $P \Rightarrow Q$ tells us nothing about the validity of its converse, $Q \Rightarrow P$. For example, consider the statements

P : the integer n is even

Q : the integer $n^2 + n$ is even

Notice that in this case $P \Rightarrow Q$ is a true statement as it is correct that if n is even, then $n^2 + n$ is even. But what about $Q \Rightarrow P$? This would be the statement *If $n^2 + n$ is even, then n is even*. However this implication is **not** true. In particular, $n^2 + n$ is always even regardless of whether n is or not. Sometimes the converse of a valid statement is valid, but sometimes it is not: it just depends on exactly what the statement is.

3.3.2 Equivalences

We say that two logical statements P and Q are equivalent if they always have the same truth values. For example, if A , B , and C are statements and we let P denote the statement $\neg(A \wedge (B \vee \neg C))$ and Q denotes the statement $\neg A \vee (\neg B \wedge C)$, then the truth tables of these statements will verify that these two statements are equivalent.

A	B	C	$\neg(A \wedge (B \vee \neg C))$	$\neg A \vee (\neg B \wedge C)$
T	T	T	F	F
T	T	F	F	F
T	F	T	T	T
T	F	F	F	F
F	T	T	T	T
F	T	F	T	T
F	F	T	T	T
F	F	F	T	T

As these statements are equivalent, you can use one in place of the other in any logical sentence. This will be very important for establishing certain types of proofs as we'll see later: sometimes we can replace a seemingly difficult to prove statement with a logically equivalent, but easier to think about, statement that we can prove.

Notice that two statements P and Q are equivalent, then it will necessarily be the case that $P \Rightarrow Q$. That is, if P is true exactly when Q is

true (and P is false exactly when Q is false), then $P \Rightarrow Q$ will be a true statement (in the case P and Q are both true we have true implies true, which we know is true; and if P and Q are both false, we have false implies false which is also false). In this particular case the order of P or Q doesn't matter: if P and Q are equivalent, then it is also the case that $Q \Rightarrow P$. Thus the equivalence of two statements is expressed symbolically by the statement $(P \Rightarrow Q) \wedge (Q \Rightarrow P)$. This is sometimes abbreviated as $P \Leftrightarrow Q$, which you can think of as the implication arrows going both ways: from P to Q and also from Q to P . This is pronounced *P equivalent to Q* or *P if and only if Q*, and you will sometimes see *if and only if* abbreviated as *iff*.

3.3.3 Contrapositives

We now mention one particularly important equivalence. As noted above, the validity of $P \Rightarrow Q$ does not necessarily mean anything about the validity of $Q \Rightarrow P$. However, there is a statement that is *always* equivalent to $P \Rightarrow Q$, which is the statement $\neg Q \Rightarrow \neg P$ and called the *contrapositive* of $P \Rightarrow Q$. We can verify that these are equivalent statements by looking at their truth tables:

P	Q	$P \Rightarrow Q$	P	Q	$\neg Q$	$\neg P$	$\neg Q \Rightarrow \neg P$
T	T	T	T	T	F	F	T
T	F	F	T	F	T	F	F
F	T	T	F	T	F	T	T
F	F	T	F	F	T	T	T

This will turn out to be an extremely useful trick in proving theorems because it means we can establish that a theorem $P \Rightarrow Q$ is true by proving its contrapositive, $\neg Q \Rightarrow \neg P$, as these are equivalent statements. It's hard to appreciate this the first time you learn about it, but it is sometimes the case that a contrapositive is easier to think about than the original statement.

As an example of establishing the proof of a theorem by proving the contrapositive, let's consider proving the following statement: if n is an odd integer, then $\sqrt{2n}$ is not an integer. Proving this directly is not so clear, as the only obvious strategy is to write $n = 2m + 1$ (as n is odd), and then we must show that $\sqrt{2(m+1)} = \sqrt{2m+2}$ is not an integer. But how do we do that? If we try to prove the contrapositive instead, then our goal is prove that if $\sqrt{2n}$ is an integer, then n is an even number. This seems a bit more straight-forward as if $\sqrt{2n}$ is an integer then we

may write it as $m = \sqrt{2n}$ for some integer m , and by squaring both sides we have that $m^2 = 2n$, but this means m^2 is an even number, and we had previously established that even numbers square to even numbers and odd numbers square to odd numbers, so m is also an even number and we can write $m = 2k$. But as $m = \sqrt{2n}$ we can write $2k = \sqrt{2n}$, or $4k^2 = 2n$ which we can further rewrite as $n = 2k^2$, showing that n is an even integer.

4

Sets

A set is a Many that allows itself to be thought of as a One.

GEORG CANTOR

“Sets” are abstract tools used for collecting information, and provide a sort of universal language for expressing mathematics¹. Even if you’ve never seen or heard of sets before, they’ve been lurking in the background of most of the mathematics you’ve done in other courses. It can be convenient sometimes to know a little bit of the basic language of sets, as this will allow us to easily and compactly describe certain mathematical objects that we’ll encounter.

The goal of this chapter is to introduce the basic ideas, language, and notation of the theory of sets. Some of this may seem strange and abstract if you’ve never seen it before, so lots of examples are included to hopefully elucidate anything that initially seems unintuitive. The main thing you should take away from this chapter is to have a general idea of what a set is; to learn some of the common notation; and obtain a basic understanding how certain simple geometric objects (such as lines and circles) can be represented as sets.

4.1 Basic ideas and definitions

A *set* is an unordered collection of objects. These objects could be numbers, points in space, functions, words, symbols, other sets, or (almost) anything else. Most of mathematics is described in terms of sets, even though this isn’t always made explicit.

We sometimes describe a set by explicitly writing out everything in the set, separated by commas, and surrounded by curly braces. For example, the set containing the first few positive, even numbers is

$$\{2, 4, 6, 8, 10\}.$$

¹This idea of sets being a “universal language” for mathematics isn’t *really* correct, but it’s a convenient way of thinking about sets. “Most” mathematicians most of the time think of most mathematics as being described in terms of sets, but there are exceptions.

The only thing that matters when we talk about a set is what is in the set. The order in which an object occurs in a set does not matter, so the following sets are all the same:

$$\{2, 4, 6, 8, 10\} = \{10, 8, 6, 4, 2\} = \{8, 2, 4, 10, 6\}.$$

The number of times we write an object in the set also does not matter (as long as it occurs at least once):

$$\{2, 4, 6, 8, 10\} = \{2, 2, 2, 4, 4, 6, 8, 10, 10, 10, 10, 10\}.$$

We use the symbol \in to denote that something is an element of a set, and \notin to denote that something is not an element of a set:

$$\begin{aligned} 2 &\in \{2, 4, 6, 8, 10\} \\ 3 &\notin \{2, 4, 6, 8, 10\}. \end{aligned}$$

Many times a set will be too big for us to write out all of the elements, and in that situation we need some other notation to describe the set. One common notation is to list a few elements in a set and then write “...” to mean “continue the pattern.” For example,

$$\{2, 4, 6, 8, 10, 12, 14, \dots\}$$

denotes the set of all positive even numbers; while

$$\{5, 10, 15, 20, 25, 30, \dots\}$$

denotes the set of all positive multiples of 5.

Exercise 4.1.

Write down a set which contains all positive integers that satisfy the following conditions: each number is a multiple of 4, a multiple of 6, and is less than 50. We begin with the set of all positive multiples of 4 less than 50,

$$\{4, 8, 12, 16, 20, 24, 28, 32, 36, 40, 44, 48\},$$

then we remove everything which is not a multiple of 6, leaving

$$\{12, 24, 36, 48\}.$$

Of course, it can get tedious to write sets down in this way every time we want to refer to a set. To save ourselves some writing if we are going to refer to a set multiple times, we will often assign the set a name. For example, if we write

$$E = \{2, 4, 6, 8, 10, 12, 14, \dots\}$$

then we are saying we want to use the symbol E to refer to the set of all positive even numbers. We are then justified in writing things like $28 \in E$, $17 \notin E$, and $-2 \notin E$.

It will sometimes be convenient to say that several things are or are not in a given set. In this case we list all of those things separated by commas and followed by \in or \notin :

$$8, 32, 96, 384 \in E$$

$$3, 347, -10 \notin E$$

Many times the sets we will be interested in will be “special,” and we will only be interested in those sets for a little while – e.g., while we’re solving a particular problem. So, we might use E to denote one set now and then later use the same symbol again to denote a different set. For instance, in solving one problem we may let E denote the set $\{1, 2, 3\}$, and let we’ll use E to denote the set $\{-3, 7, 8\}$. It will usually be clear from context which set a given symbol refers to.

There are some sets that are used over and over, again and again, and those sets have special names and symbols that are reserved only for those particular sets. One such set is the set of *natural numbers*, which is the set of all positive whole numbers and is denoted by a capital N , but written in what is often called “blackboard bold” and looks like \mathbb{N} :

$$\mathbb{N} = \{1, 2, 3, 4, 5, 6, \dots\}.$$

Remark.

In older textbooks this \mathbb{N} was originally written as a bold N . It’s difficult to write bold letters on paper or a blackboard, however, and so people started writing an extra line in the letter to denote the letter was bold. This way of writing bold letters eventually became popular enough that it made its way into typed works as a special

typeface.

The set of all whole numbers (positive, negative, and zero) is called the set of integers and is denoted by a blackboard bold \mathbb{Z} :

$$\mathbb{Z} = \{\dots, -3, -2, -1, 0, 1, 2, 3, \dots\}.$$

Remark.

Using the letter \mathbb{Z} might seem like a weird choice for integers, but it's only weird if you're an English speaker. Many influential mathematicians of the past, including Georg Cantor who is considered the father of set theory, were German and so they of course used the German equivalent of these words and used the first letter of those German words. The German word for numbers is *die Zahlen* (*die* is the feminine definite article in German, like *la* in French or Spanish), hence the \mathbb{Z} .

Conveniently, German and English have some commonalities and so some German words are very similar to their English counterparts, so most of these blackboard bold letters are actually what you would guess using the English words. For example, *the natural numbers* in German is *die natürliche Zahlen*, so \mathbb{N} makes sense in both German and English.

The number of distinct elements in a set A is called the cardinality of the set and is denoted by either $\#A$ or $|A|$. For example $\#\{7, 8, 0, 4, 3\} = 5$ while $\#\{3, 6, 9, 12, \dots, 84, 87, 90\} = 30$. The cardinality can be infinite as well; both \mathbb{N} and \mathbb{Z} have infinite cardinality.

4.1.1 Set-builder notation

Unfortunately, there are times when the ... notation mentioned above can be ambiguous. For example,

$$\{2, 4, \dots\}$$

could mean the set of all even numbers, or it could be all the powers of 2: both of the following sets match the pattern

$$\{2, 4, 6, 8, 10, \dots\}$$

$$\{2, 4, 8, 16, 32, \dots\}.$$

To get around this ambiguity we sometimes use *set builder notation*. In this notation we write two curly braces, like normal, but separated into two parts by a vertical bar. On the left-hand side of the bar we write a variable (or sometimes a collection of variables) that give us some pattern that all of the elements in the set follow, and on the right-hand side we given a condition (usually in the form of an equation or inequality, but sometimes written in words) that the variable must satisfy in order to be an element of the set. The collection of all positive even integers, for example may be written in set builder notation as

$$E = \{x \mid x = 2n \text{ for some } n \in \mathbb{N}\}.$$

That is, we start off by considering the natural numbers, but to be an element of E , a given natural number x has to be two times some other natural number. (A number is even if and only if it is divisible by two.)

We could define the set of all positive odd numbers as

$$\mathcal{O} = \{x \mid x = 2n - 1 \text{ for some } n \in \mathbb{N}\}.$$

Exercise 4.2.

- (a) Write the set of all positive, whole number multiples of 5 in set builder notation.
- (b) Write the set of all whole number multiples (including negatives) of 5 in set builder notation.

(a)

$$\{x \mid x = 5n \text{ for some } n \in \mathbb{N}\}$$

(b)

$$\{x \mid x = 5n \text{ for some } n \in \mathbb{Z}\}$$

Another common set of numbers is the set of **rational numbers**, which are ratios of integers where the denominator is not zero. These are quotients², so the set of all rational numbers is denoted \mathbb{Q} . In set builder notation we can express \mathbb{Q} as

$$\mathbb{Q} = \left\{ \frac{p}{q} \mid p, q \in \mathbb{Z} \text{ and } q \neq 0 \right\}.$$

In the examples thus far we have only considered sets of numbers, but there is nothing special about numbers: the elements of a set can be any type of object. They could be names of people,


{William, Charles, Percy, Fred, George, Ron, Ginny},

or abstract symbols,

{♥, ♣, ♦, ♠},

or points in space,

$$\{(x, y) \mid -2x + 8y = 10\}.$$



You can even have sets that contain other sets:

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

Sets are ubiquitous in mathematics: the vast majority of things you work with are, or are defined in terms of, sets. This point may not have been made clear to you before in earlier mathematics courses because it may not have been needed, but for our purposes in this class we will need to deal with sets on a regular basis, so it's important that we have a good understanding of them.

²Conveniently, the German word for *the quotient* is *der Quotient*, and so the \mathbb{Q} makes sense for English speakers too!

4.1.2 Subsets and supersets

We say that a set A is a subset of a set B if every element of A is also an element of B . When this happens we write $A \subset B$.

Example 4.1.

Every natural number is an integer, so the set of natural numbers is a subset of the set of integers: $\mathbb{N} \subset \mathbb{Z}$. Every integer is also a rational number (e.g., $3 = 3/1$), so $\mathbb{Z} \subset \mathbb{Q}$.

Example 4.2.

Suppose that A is the set of all the multiples of 3, and B is the set of all multiples of 12:

$$A = \{x \mid x = 3n \text{ for some } n \in \mathbb{N}\},$$

$$B = \{x \mid x = 12n \text{ for some } n \in \mathbb{N}\}.$$

Since every multiple of 12 is also a multiple of 3 (because 3 divides 12), B is a subset of A : $B \subset A$.

When A is a subset of B we say that B is a superset of A . That is, when we write $A \subset B$ the set on the left is a subset of the set on the right; and the set on the right is a superset of the set on the left. The superset is the “larger” set, and the subset is the “smaller” set. Sometimes it will be convenient for the symbol \subset to be written in the other direction: for example $B \supset A$. Here B is still the larger superset, and A is the smaller subset. (Compare this to writing $3 < 4$ and $4 > 3$.)

In our mind’s eye we often picture the relationship between a set and a any subsets or supersets as shown in Figure 4.1.

We will use pictures like this, which are called Venn diagrams, many times when describing sets, even if the sets we’re talking about don’t really look like the two-dimensional shapes we’ll draw: though the pictures aren’t technically accurate (e.g., A and B may not be actually be the set of points making up two ovals in the plane), it’s often very helpful to




Figure 4.1: If $A \subset B$, then we imagine A as sitting inside of B .

use these kinds of abstract pictures because they provide us with some intuition about how different sets are related to one another.

Notice again that we say A is a subset of B if every element of A is also an element of B . This means, in particular, that for every set A , A is a subset of itself: every element of A is also an element of A . We are thus justified in writing $A \subset A$. If we want to explicitly exclude this possibility, we use the symbol \subsetneq : writing $A \subsetneq B$ means that A is a subset of B and A is not all of B .

Exercise 4.3.

Suppose that A and B are two sets and $A \subsetneq B$. Show that this means there must exist at least one element of B which is not an element of A . Suppose $A \subsetneq B$. That is, $A \subseteq B$ but $A \neq B$. This means $B \not\subseteq A$. Since B is not a subset of A , it is not the case that every element of B is also an element of A ; in other words, there exists at least one element of B (possibly many more, but at least one) which is not an element of A .

When $A \subsetneq B$ we call A a **proper subset** of B . For example, the natural numbers are a proper subset of the integers, and the integers are a proper subset of the rational numbers.

Remark.

There is a little bit of ambiguity that can occur with the symbol \subset : some authors use \subset to mean \subsetneq , and use \subseteq to mean \subset . That is, some people will use $A \subseteq B$ to mean that A is a subset of B , possibly all

of B , and $A \subset B$ to mean that A is a subset of B but not all of B . This is reminiscent to using \leq and $<$ in comparing numbers, but it's not completely standard.

To avoid any potential ambiguity we will typically use $A \subsetneq B$ to mean that A is a proper subset of B , and $A \subseteq B$ to mean that A is a subset of B but could potentially be all of B .

4.1.3 Equality

We say that two sets A and B are *equal* if they have precisely the same elements: that is, if $x \in A$ then $x \in B$ and if $y \in B$, then $y \in A$ as well. This is exactly the same thing as saying $A \subseteq B$ and $B \subseteq A$. When this happens we, unsurprisingly, write $A = B$.

Example 4.3.

Let A and B be the sets described below:

$$A = \{x \in \mathbb{Z} \mid x = 2n \text{ for some } n \in \mathbb{Z}, \text{ and } x = 3m \text{ for some } m \in \mathbb{Z}\},$$

$$B = \{y \in \mathbb{Z} \mid y = 6n \text{ for some } n \in \mathbb{Z}\}.$$

Show that A and B are equal.

Here, A is the set of all integers which are simultaneously multiples of 2 and 3, while B is the set of all integers which are multiples of 6. If you start writing down a few elements of A , then you'll probably be convinced pretty quickly that, sure enough, everything in A is a multiple of 6, but let's actually prove this.

We first want to show that $A \subseteq B$: i.e., every integer which is a multiple of both 2 and 3 must be a multiple of 6. So suppose $x \in A$, we want to show that $x \in B$ as well. If $x \in A$ then $x = 2n = 3m$ for some pair of integers m and n . This equation means, in particular, that 2 divides $3m$. Since 2 is a prime number it must divide either 3 or m (this is basically the definition of a prime number; see the Wikipedia page about prime numbers for more information). Since 2 does not divide 3, it must divide m . Thus $m = 2k$ for some k . This means $x = 3m = 3 \cdot 2k = 6k$, and so x must be a multiple of 6. Hence if $x \in A$, then $x \in B$ as well, so $A \subseteq B$.

We also need to show that $B \subseteq A$. Suppose that $y \in B$, so $y = 6k$ for some k . But then $y = 3 \cdot 2 \cdot k$, and so y is simultaneously a multiple of 2 (take $n = 3k$ in the definition of A) and a multiple of 3 (let $m = 2k$). Thus $B \subseteq A$.

As $A \subseteq B$ and $B \subseteq A$, $A = B$.

Remark.

Just a reminder that it's okay if you don't understand an example when you first read it in these notes. The important thing is to make an effort to try to understand it. Usually just making an effort, even if you don't feel comfortable that you understood what you just read, still helps to get your brain thinking about the idea. You may find that if you read something you don't understand, then step away from it for a while (a few hours, maybe a day or two) and then re-read it, it might make sense on the second reading. If you still don't understand the example on a second reading, don't beat yourself up about it. Feel free to ask questions about the idea through email, office hours, or in class if you're still unable to understand what's going on. The most important thing is to keep trying and not let one thing you don't understand discourage you from trying anything else.

4.1.4 The empty set

There is one special set in mathematics called *the empty set* which is the only set that contains no elements; it is the only set of cardinality zero and is denoted \emptyset .

A set without anything in it might sound uninteresting, but there is at least one surprising thing about the empty set: the empty set is a subset of every other set. That is, for any set A , $\emptyset \subseteq A$. Why is this the case? We should only write $\emptyset \subseteq A$ if every element of \emptyset is also an element of A . Since \emptyset has no elements, however, it immediately satisfies this definition! All the elements of \emptyset (all zero of them) are also elements of A !

Exercise 4.4.

If the idea that the empty set is a subset of every other set sounds a little bit odd, re-read the above paragraph and think about the logic behind the last sentence until it makes sense. The solution to this exercise in the appendices gives another way to think about this if the first explanation above simply won't "click" for you. Here's another way to think about subsets that might make $\emptyset \subseteq A$ a little easier to digest. By definition, $B \subseteq A$ if every element of B is also an element of A . Think of this like a test: you hand me an element of B and I tell you *Pass* or *Fail*, where I say *Pass* if the element is an element in A , and *Fail* if it's not. To see if $B \subseteq A$ or not, we'll subject *every* element of B to this test. If any element of B fails the test, then B is not a subset of A . However, if no element fails the test, then B is a subset of A .

Now, for any set A let's try to apply this test to \emptyset . So, for every element of \emptyset we apply our test, and if nothing fails, then $\emptyset \subseteq A$. There are no elements of \emptyset , however, so there's nothing to fail. There's no failure, so $\emptyset \subseteq A$.

4.1.5 Real numbers

So far we have described three different sets of numbers: the natural numbers \mathbb{N} , the integers \mathbb{Z} , and the rational numbers \mathbb{Q} . We now describe one more set of numbers which we will use in this class: the real numbers.

To define the real numbers rigorously would take us very far afield, and so we will be a little bit hand-wavy in the definition. A ***real number*** is simply the coordinate of a point on the real line; equivalently, it is the collection of all numbers that we can write down with a (possibly infinite) decimal expansion. All of the numbers described thus far (natural numbers, integers, and rational numbers) are real numbers: we can write 6 as 6.000...; we can write -3 as -3.000...; we can write $\frac{22}{7}$ as 3.142857142857142857...

The set of all real numbers is denoted \mathbb{R} . Notice we have the following string of subsets:

$$\mathbb{N} \subsetneq \mathbb{Z} \subsetneq \mathbb{Q} \subsetneq \mathbb{R}.$$

Notice that the examples of real numbers we wrote down above all have a decimal expansion which is eventually repeating. However there

are numbers that can't be written in this way. One simple example is $\sqrt{2}$. We can write $\sqrt{2}$ as an infinite decimal expansion $\sqrt{2} = 1.414213562\dots$, this expansion never repeats no matter how many digits you write down. We won't do it, but it can be shown that every rational number has an eventually repeating decimal expansion. So, another way to say that the decimal expansion of $\sqrt{2}$ never repeats, is to say that it is impossible to write $\sqrt{2}$ as a ratio of two integers. That is, $\sqrt{2}$ is not a rational number. A real number that is not rational is called an *irrational number*.

4.2 Operations on sets

4.2.1 Unions

Given a collection of sets there are many different ways we can combine the sets together to get new sets. Here we discuss the three most important such operations: unions, intersections, and products.

Given two sets A and B , their *union* is the "smallest" set which contains every element of A as well as every element of B , and is denoted $A \cup B$. You should think of the union as gluing two sets together to get a bigger set.

Example 4.4.

Let $A = \{2, 4, 6, 8, 10\}$ be the set of all even integers between 1 and 10, and let $B = \{1, 3, 5, 7, 9\}$ be the set of all odd integers between 1 and 10. Then their union $A \cup B$ is the set of all integers between 1 and 10:

$$A \cup B = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}.$$

Exercise 4.5.

Let A and B be any two sets. Show that $A \subseteq A \cup B$ and $B \subseteq A \cup B$. To show $A \subseteq A \cup B$, we need to show that every element of A is also an element of $A \cup B$. Let $x \in A$ be any element of A ; we need to show $x \in A \cup B$ as well. Notice, however, that $A \cup B$ consists of all elements which are in A or B . Since x is in A , it is certainly in A or B , and so $x \in A \cup B$. Thus $A \subseteq A \cup B$.

The argument that $B \subseteq A \cup B$ is exactly the same, but with B 's where A 's appeared above.

Given any two sets A and B , there are going to be *lots* of other sets that contain A and B as subsets. In the example above, for instance, the set

$$\{1, 2, \dots, 10, 11\}$$

contains both A and B as a subset, as does \mathbb{N} and \mathbb{Z} . The union $A \cup B$ is the *smallest* set containing both A and B as subsets in the following sense: If $A \subseteq C$ and $B \subseteq C$, then $A \cup B \subseteq C$.

4.2.2 Intersections

Another operation we can perform on two sets is to intersect them. The *intersection* of two sets A and B , denoted $A \cap B$, consists precisely of all of the elements which are in both A and B . That is, $x \in A \cap B$ if and only if $x \in A$ and $x \in B$.

Example 4.5.

Let A be the set of all multiples of 6, and B the set of all multiples of 10,

$$A = \{\dots - 18, -12, -6, 0, 6, 12, 18, \dots\},$$

$$B = \{\dots - 30, -20, -10, 0, 10, 20, 30, \dots\}.$$

Then $A \cap B$ is the set of all the numbers which are both multiples of 6 and 10.

$$A \cap B = \{\dots, -90, -60, -30, 0, 30, 60, 90, \dots\}.$$

Example 4.6.

Suppose that S is the set of all characters that have ever appeared in

a Star Wars film,

$$S = \{\text{Luke Skywalker, Obi-Wan Kenobi, Kylo Ren, } \dots\},$$

that R is the set of all droids from the Star Wars films,

$$R = \{\text{R2D2, C3P0, BB-8, } \dots\},$$

D is the set of all characters corrupted by the dark side of The Force,

$$D = \{\text{Darth Vader, Kylo Ren, Emperor Palpadine, } \dots\}$$

and V is the set of all characters which appeared in *Star Wars V: The Empire Strikes Back*,

$$V = \{\text{Luke Skywalker, Lando Calrissian, Boba Fett, } \dots\}.$$

Then the set of all droids that appeared in *The Empire Strikes Back* is the intersection of the set of all droids and the set of all characters that were in that movie:


$$R \cap V = \{\text{C3P0, R2D2}\}.$$

The set of all characters which were corrupted by the dark side of The Force and were in *The Empire Strikes Back* is the intersection of all characters corrupted by the dark side of The Force and the set of all characters in *The Empire Strikes Back*:

$$D \cap V = \{\text{Darth Vader, Emperor Palpadine}\}.$$


Example 4.7.

Suppose that A is the set of all points in the plane (all (x, y) -pairs) that are at most distance 1 from the origin.




$$A = \{(x, y) \mid x^2 + y^2 \leq 1\}$$

And suppose that B is the line with slope 1 through the point $(0, -0.25)$



$$B = \{(x, y) \mid y = x - 0.25\}$$

Then the intersection $A \cap B$ is then the portion of the line B that remains inside the disc A . This is the dark purple line segment in the figure below. (The original disc and line are drawn in very lightly just for comparison; they are not part of $A \cap B$.)



It may happen that two sets have nothing in common: for example, the set $A = \{1, 2, 3\}$ and the set $B = \{4, 5, 6\}$ have no common elements. In a situation such as the intersection of the two sets is empty, $A \cap B = \emptyset$, and we say that A and B are disjoint.

Exercise 4.6.

Let A and B be any two sets. Show that $A \cap B$ is a subset of A and also a subset of B . By definition, $A \cap B$ contains everything that is in both A and in B . Thus every element of $A \cap B$ is an element of A , and this is exactly what it means to say $A \cap B \subseteq A$. By the same token, $A \cap B \subseteq B$.

Exercise 4.7.

Show that if $A \subseteq B$, then $A \cap B = A$. By assumption $A \subseteq B$, and so every element of A is also an element of B . Since $A \cap B$ consists of all the elements of both A and B , and everything in A is already an element of B , we see $A \cap B$ doesn't "remove" anything from A .

Just as the union $A \cup B$ was the smallest set containing both A and B as subsets, the intersection $A \cap B$ is the largest subset of both A and B in the following sense: If $C \subset A$ and $C \subseteq B$, then $C \subseteq A \cap B$.

Anytime you have several operations defined on some collection of objects (e.g., unions and intersections defined for sets), you might be interested in how those operations interact with one another. For unions and intersections this interaction is similar distributive law for normal numbers (e.g., that $x \cdot (y + z) = x \cdot y + x \cdot z$).

Proposition 4.1. *For any sets A , B , and C we have the following two distributive laws:*

$$\begin{aligned} A \cap (B \cup C) &= [A \cap B] \cup [A \cap C] \\ A \cup (B \cap C) &= [A \cup B] \cap [A \cup C] \end{aligned}$$

Proof. We will only prove the first distributive law; the proof of the second one is almost identical.

Notice that elements of $A \cap (B \cup C)$ are elements of A which are also elements of either B or C . The elements of $A \cap B$ are elements of both A and B ; the elements of $A \cap C$ are elements of both A and C . Unioning $A \cap B$ and $A \cap C$ together, we have exactly the elements of A which are also in either B or C . \square

4.2.3 Products

One last operation we will mention is the Cartesian product, which we will usually refer to simply as the “product.” Given two sets, A and B , their (Cartesian) product is a set denoted $A \times B$ and which consists of all ordered pairs (a, b) where $a \in A$ and $b \in B$:

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


Example 4.8.

Let $A = \{x, y, z\}$ and $B = \{u, v, x\}$. Then

$$\begin{aligned} A \times B &= \{(x, u), (x, v), (x, x), \\ &\quad (y, u), (y, v), (y, x), \\ &\quad (z, u), (z, v), (z, x)\} \end{aligned}$$

Example 4.9.

Let A be the interval $[1, 4]$ and B the interval $[2, 3]$. Then the product $A \times B$ consists of all pairs of numbers (i.e., all (x, y) pairs in the plane) where the first coordinate is between 1 and 4, and the second coordinate is between 2 and 3:



$$A \times B = \{(x, y) \mid 1 \leq x \leq 4 \text{ and } 2 \leq y \leq 3\}$$

It is fairly often that we will want to consider the product of a set with itself, $A \times A$. In such a situation we will usually simply write A^2 to mean $A \times A$.

The three operations we described above can be defined for more than two sets. For example, it makes sense to talk about the union, intersection, or product of three sets. It is completely reasonable, for example, to say that the union $A \cup B \cup C$ should be the smallest set containing all the elements of A , all the elements of B , and all the elements of C . The intersection $A \cap B \cap C$ should contain only those elements that are in all three sets A , B , and C .

Example 4.10.

Consider the sets A , B , and C described below:

$$A = \{1, 2, 3, \dots, 10\}$$

$$B = \{2, 4, 6, \dots, 20\}$$

$$C = \{-12, -9, -6, \dots, 6, 9, 12\}.$$

The union of these sets is

$$A \cup B \cup C = \{-12, -9, -6, -3, 0, 1, 2, 3, \dots, 10, \\ 12, 14, 16, 18, 20\}.$$

The intersection is

$$A \cap B \cap C = \{6\}.$$

Of course, there's nothing magical about having two sets or three sets: we can define unions and intersections for any number of sets – even infinitely-many.

Example 4.11.

For each $n \in \mathbb{N}$ define the set A_n to be the interval $[-\frac{1}{2^n}, \frac{1}{2^n}]$. The first few intervals are thus

$$A_1 = [-1/2, 1/2]$$

$$A_2 = [-1/4, 1/4]$$

$$A_3 = [-1/8, 1/8]$$

$$A_4 = [-1/16, 1/16]$$

⋮

The intersection of all these intervals is usually written in one of two ways,

$$A_1 \cap A_2 \cap A_3 \cap \dots \quad \text{or} \quad \bigcap_{n=1}^{\infty} A_n,$$

and consists of all the elements which are in *every* A_n . In this case

the only such element is 0:

$$\bigcap_{n=1}^{\infty} A_n = \{0\}.$$

Exercise 4.8.

For each $n \in \mathbb{N}$, let B_n be the following interval:

$$B_n = \left[\frac{1}{2^n}, 1 - \frac{1}{2^n} \right].$$

What is the infinite union of all the B_n 's, $\bigcup_{n=1}^{\infty} B_n$? The union of the B_n 's is the open interval $(0, 1)$. To see this, let's let U denote the infinite union, $U = \bigcup_{n=1}^{\infty} B_n$. We want to show $U = (0, 1)$, which means we need to show $U \subseteq (0, 1)$ and $(0, 1) \subseteq U$. It is easy to see $U \subseteq (0, 1)$ since each $B_n \subseteq (0, 1)$. To see $(0, 1) \subseteq U$, let $x \in (0, 1)$ be any arbitrary element. Since $x > 0$, there exists some value of m_1 such that $x > \frac{1}{2^{m_1}}$ as $\frac{1}{2^n}$ decreases to 0 as n increases. Notice if $x > \frac{1}{2^{m_1}}$, then $x > \frac{1}{2^M}$ for any $M > m_1$. Similarly, since $x < 1$, there exists some m_2 such that $x < 1 - \frac{1}{2^{m_2}}$. Note also that if $M > m_2$, then $x < 1 - \frac{1}{2^M}$.

Now let M be the maximum of m_1 and m_2 , $M = \max\{m_1, m_2\}$. Then $x > \frac{1}{2^M}$ and $x < 1 - \frac{1}{2^M}$; i.e., $x \in B_M$. Since $B_M \subseteq U$, this shows $x \in U$.

Thus we have established that $(0, 1) = U$.

The product might be slightly, but not very, surprising. When we write a product of three sets we will mean the collection of ordered triples; a product of four sets is the collection of ordered quadruples. In general, the product of n sets is the set of all ordered n -tuples. (An n -*tuple* is an ordered list of n items. A 2-tuple is simply a pair; a 3-tuple is a triple; a 5-tuple has the form (a, b, c, d, e) .)

Example 4.12.

Let A , B , and C be the following sets:

$$A = \{1, 2, 3\}$$

$$B = \{\alpha, \beta\}$$

$$C = \{\#, b\}$$

Then $A \times B \times C$ is the following set

$$\begin{aligned} &\{(1, \alpha, \#), (1, \alpha, b), (1, \beta, \#), (1, \beta, b), \\ &\quad (2, \alpha, \#), (2, \alpha, b), (2, \beta, \#), (2, \beta, b), \\ &\quad (3, \alpha, \#), (3, \alpha, b), (3, \beta, \#), (3, \beta, b)\} \end{aligned}$$

It is very common to consider the Cartesian product of a set A with itself n times, so we usually denote this as A^n .

Example 4.13.

The set of all ordered triples of integers could be written \mathbb{Z}^3 :

$$\mathbb{Z}^3 = \{(x, y, z) \mid x, y, z \in \mathbb{Z}\}$$

We can also talk about products of infinitely-many sets, but for simplicity we will avoid that for the time being.

4.2.4 Complements

The last operation on sets we will describe is not defined for all sets, but only for subsets of some given set. That is, in some applications there will be some ambient set “in the background,” and all other sets we are interested in will be subsets of this ambient set. In such a situation, we sometimes call the ambient set the *universe* because it consists of everything we care about for the problem at hand. For example, in geometry the universe may be the set of all points in the plane, \mathbb{R}^2 – for some geometric problems everything you care about might take place in the plane, so that is your universe.

Once we have a universal set \mathcal{U} , we can define the *complement* of any subset $E \subseteq \mathcal{U}$, which you should think of as being the complete opposite of E . To be more precise, given any set E inside the universe \mathcal{U} , the complement of E , denoted E^c , is the set of all elements in \mathcal{U} which are not in E :

$$E^c = \{x \in \mathcal{U} \mid x \notin E\}.$$

Example 4.14.

Suppose the universe \mathcal{U} consists of all integers between 1 and 10,

$$\mathcal{U} = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}.$$

- If E is the set of all even numbers between 1 and 10, $E = \{2, 4, 6, 10\}$, then its complement consists of all the odd numbers, $E^c = \{1, 3, 5, 7, 9\}$.
- if E is the set of all numbers in \mathcal{U} greater than 7, $E = \{8, 9, 10\}$, then its complement is the set of all numbers less-than-or-equal-to 7, $E^c = \{1, 2, 3, 4, 5, 6, 7\}$.


Exercise 4.9.

Let \mathcal{U} be any universal set and $E \subseteq \mathcal{U}$ any subset. Show $(E^c)^c = E$. For notation convenience, let's write $D = E^c$ for the moment. Then D is made up of all the $x \in \mathcal{U}$ such that $x \notin E$. So what is D^c , aka $(E^c)^c$? By definition, D^c is the set of all $x \in \mathcal{U}$ such that $x \notin D$. But what does it mean if $x \notin D$? Since D consists of everything *not* in E , if $x \notin D$ that must mean $x \in E$. That is $D^c = E$.

Example 4.15.

Suppose the universe \mathcal{U} consists of all points in the plane, $\mathcal{U} = \mathbb{R}^2$. If E is the set of all points whose distance to the origin is at most 1

(so, E is the circular disc of radius 1 centered at the origin), then its complement E^c consists of all the points distance more than 1 from the origin (this would be the entire plane with a “hole” of radius 1 centered at the origin).



Exercise 4.10.

Given some universe \mathcal{U} , what is the complement of the empty set \emptyset ? What is the complement of \mathcal{U} ? The complement of the empty set, by definition, is the collection of all elements of \mathcal{U} which are not elements of the empty set. But since the empty set has no elements, nothing in \mathcal{U} is in the empty set, and so the complement of \emptyset is the entire universe \mathcal{U} .


The complement of \mathcal{U} is the set of all elements of \mathcal{U} which are not elements of \mathcal{U} – of course, there are no such elements (an element can not simultaneously be in \mathcal{U} and not in \mathcal{U} , and so the set of all such elements is empty. I.e., $\mathcal{U}^c = \emptyset$).

4.2.5 Difference


The difference between two sets E and F , denoted $E \setminus F$, is the set of all elements in E which are not also elements in F :

$$E \setminus F = \{x \in E \mid x \notin F\}.$$

To have a picture of this, imagine that E and F are the overlapping regions indicated below.



Then the set difference $E \setminus F$, is the shaded region below.

**Example 4.16.**

Let S be the set of all Star Wars movies,

$$S = \{\text{Star Wars, The Empire Strikes Back, Return of the Jedi,} \\ \text{The Phantom Menace, The Clone Wars, Revenge of the Sith,} \\ \text{The Force Awakens, Rogue One, The Last Jedi,} \\ \text{Solo}\}$$

and let D be the set of all movies produced by Disney,

$$D = \{\text{Snow White, Pinocchio, ..., Coco, The Force Awakens, ...}\}.$$

Then $S \setminus D$ would be the set of all Star Wars movies not produced by Disney,

$$S \setminus D = \{\text{Star Wars, The Empire Strikes Back, Return of the Jedi,} \\ \text{The Phantom Menace, The Clone Wars, Revenge of the Sith}\}$$

Exercise 4.11.

Show that $E \setminus F$ is equal to $E \setminus (F \cap E)$. To show that two sets are equal, we need to show that each is a subset of the other. That is, we must show $E \setminus F \subseteq E \setminus (F \cap E)$ and $E \setminus (F \cap E) \subseteq E \setminus F$.

First note that if $x \in E \setminus F$, that means x is in E but not in F . If x is not in F , then in particular it's not in $F \cap E$ (everything in $F \cap E$ is in F). Thus $x \in E \setminus (F \cap E)$, and so $E \setminus F \subseteq E \setminus (F \cap E)$.

Now suppose $x \in E \setminus (F \cap E)$. That is, $x \in E$ but x is not in $F \cap E$. This means in particular that $x \notin F$: we already know $x \in E$ so if $x \in F$ as well, we would have $x \in F \cap E$. Hence $x \in E$ but $x \notin F$, which precisely means $x \in E \setminus F$. Thus $E \setminus (F \cap E) \subseteq E \setminus F$.

Together these mean that the two sets are equal.

4.2.6 De Morgan's laws

It is very common in mathematics to have multiple possible operations you can perform on a given type of object, and then to ask how these operations interact with one another. For example, in arithmetic two basic operations are addition and multiplication, and these two operations "interact" via the distributive law $a \cdot (b + c) = a \cdot b + a \cdot c$.

At this point we have a few different operations we can perform on sets, and we want to know how they interact with each other. In particular, we have unions, intersections, and complements. These three operations are related by two rules called *de Morgan's laws*, which essentially say that unions turn into intersections (and intersections turn into unions) when we take complements.

More precisely, if E and F are two subsets of some universe \mathcal{U} (recall we always need a "universe" to discuss complements), then we have the following:

$$\begin{aligned}(E \cup F)^c &= E^c \cap F^c \\ (E \cap F)^c &= E^c \cup F^c\end{aligned}$$

That is, we can intentionally turn unions into intersections and vice versa, but we also have to take the complement of the sets involved. Right now it might be hard to appreciate why this is something we'd like to do, but we'll see later that when calculating probabilities we will have special rules for calculating probabilities of unions and intersections. In some

types of problems we use de Morgan's laws to turn a complicated problem involving probabilities of unions into a simpler problem involving probabilities of intersections. (This is a little ways down the road from where we are now, but that's where we're heading.)

Example 4.17.

Suppose the universal set is $\mathcal{U} = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}$ and let $E = \{1, 2, 3\}$ and $F = \{3, 4, 5\}$. Verify directly that de Morgan's laws are satisfied.

Here we just want to compute the four sets stated in de Morgan's laws above (two sets per equation) and see if the equalities that are claimed to be true are in fact satisfied.

First note $E \cup F = \{1, 2, 3, 4, 5\}$. Hence $(E \cup F)^c = \{6, 7, 8, 9, 10\}$. Now note $E^c = \{4, 5, 6, 7, 8, 9, 10\}$ and $F^c = \{1, 2, 6, 7, 8, 9, 10\}$. Their intersection is $E^c \cap F^c = \{6, 7, 8, 9, 10\}$. So the first equation in de Morgan's laws is satisfied.

For the second equation we note $E \cap F = \{3\}$, so $(E \cap F)^c = \{1, 2, 4, 5, 6, 7, 8, 9, 10\}$. Now E^c and F^c we computed above, and their union is $E^c \cup F^c = \{1, 2, 4, 5, 6, 7, 8, 9, 10\}$, and so the second equation in de Morgan's laws is satisfied.

The proof of de Morgan's laws essentially has to do with working out what each side of each equation means. We will simply prove the first law, leaving the second one as an exercise.

Proof of de Morgan's first law. We wish to show that $(E \cup F)^c = E^c \cap F^c$. To show two sets are equal we must show each one is a subset of the other: i.e., we must show $(E \cup F)^c \subseteq E^c \cap F^c$ and also that $E^c \cap F^c \subseteq (E \cup F)^c$.

Let $x \in (E \cup F)^c$. That is, x is an element of \mathcal{U} which is in neither E nor F . Since $x \notin E$ and $x \notin F$, we have $x \in E^c$ and $x \in F^c$, so $x \in E^c \cap F^c$. This shows $(E \cup F)^c \subseteq E^c \cap F^c$.

Now to show the other inclusion, let $x \in E^c \cap F^c$. Thus x is in both E^c and F^c . This means x is in neither E nor F , and hence $x \notin E \cup F$. By the definition of the complement, that means $x \in (E \cup F)^c$. Hence $E^c \cap F^c \subseteq (E \cup F)^c$. \square

Exercise 4.12.

Prove the second law of de Morgan. That is, if E and F are subsets of a universal set \mathcal{U} , then $(E \cap F)^c = E^c \cup F^c$.

There's nothing really special about our using two sets in the statements of de Morgan's laws above instead of three or four or five or ... In general, given any collection of subsets E_1, E_2, \dots, E_n of some universal set \mathcal{U} , de Morgan's laws extend to

$$\begin{aligned}(E_1 \cup E_2 \cup \dots \cup E_n)^c &= E_1^c \cap E_2^c \cap \dots \cap E_n^c \\ (E_1 \cap E_2 \cap \dots \cap E_n)^c &= E_1^c \cup E_2^c \cup \dots \cup E_n^c.\end{aligned}$$

If you believe the proof of de Morgan's laws for two sets, then it's easy to see how to get de Morgan's laws for more than two sets. Let's consider the case when there are three sets, and let's just call them E, F , and G . The first law says that

$$(E \cup F \cup G)^c = E^c \cap F^c \cap G^c.$$

How can we get this if we know only have de Morgan's laws for two sets? We'll just cheat and rewrite the above as two sets. If we write $H = E \cup F$, then $E \cup F \cup G$ can be written as $H \cup G$. De Morgan's laws on two sets then tell us

$$(E \cup F \cup G)^c = (H \cup G)^c = H^c \cap G^c.$$

Now let's figure out what H^c is: since $H = E \cup F$, we must have $H^c = (E \cup F)^c$. But now de Morgan's laws for two sets tell us $H^c = (E \cup F)^c = E^c \cap F^c$. Plugging this in for H^c on the right-hand side above we have

$$(E \cup F \cup G)^c = (H \cup G)^c = H^c \cap G^c = E^c \cap F^c \cap G^c.$$

The same trick works for de Morgan's second law for three sets.

Now that we know de Morgan's laws for three sets, it's easy to extend it to de Morgan's law for four sets; once we have de Morgan's laws for four sets, we can easily extend to five sets; etc. We just keep taking a "complicated" de Morgan's law with lots of sets and rewriting it in terms of de Morgan's law with fewer sets. Repeating this process several times we can always boil everything back down to de Morgan's law with two sets which we already know.

4.3 Collections of sets

In this section we discuss some of the ways that we can collect multiple sets together. As we had previously seen, sets can contain other sets as elements. One particularly important example is the “power set” of a set, which is the collection of *all* subsets of a given set. Other important families of sets of sets are “index sets,” which give us a way of specifying a set for each element in another set. We will also discuss the idea of a partition of a set which is simply a way of breaking a given set up into disjoint pieces.

4.3.1 Power set

Given any set X , we define its *power set*, denoted $\mathcal{P}(X)$ or $\mathbb{P}(X)$ or 2^X , is the set of all sets which are subsets of X :

$$\mathcal{P}(X) = \{Y \mid Y \subseteq X\}$$

So, for example, if $X = \{1, 2, 3\}$, then its power set is

$$\mathcal{P}(X) = \{\emptyset, \{1\}, \{2\}, \{3\}, \{1, 2\}, \{1, 3\}, \{2, 3\}, \{1, 2, 3\}\}.$$

As another example, the power set of the empty set is the set that contains the empty set:

$$\mathcal{P}(\emptyset) = \{\emptyset\}.$$

4.3.2 Indexing sets

Sometimes it is convenient to be able to describe very large collections of sets, and an “index set” is one way of doing this. To be more precise, we refer to a set I as an *index set* if we have a collection of sets S_i , one for each $i \in I$. This idea is hard to appreciate the first time you learn it, but it turns out to be of fundamental importance in defining certain mathematical objects.

Although in practice we usually use index sets to describing very large collections of related sets, our index sets are not required to be large. For example, suppose I was the set $\{1, 2, 3\}$. This forms an index set for any collection of three sets, say

$$S_1 = \{0, 2, 5\}, \quad S_2 = \mathbb{Z}, \quad \text{and } S_3 = \{2, 3, 5, 7, 11, 13, \dots\}$$

Our index set (in this case $\{1, 2, 3\}$) is just a convenient way of recording the allowable indices, the i 's, the parametrize our collection of three sets.

As a more involved example, an index set could be the collection of all natural numbers, \mathbb{N} . A family of sets indexed by \mathbb{N} is just any infinite collection of sets where there's a first set (corresponding to 1), a second set (corresponding to 2), and third set, a fourth set, and so on.

For instance, we might consider the sets


$$X_n = \{m \in \mathbb{N} \mid n|m\} \text{ for each } n \in \mathbb{N}.$$

Here X_1 would be the set containing all natural numbers divisible by 1 (so X_1 would just be all natural numbers), X_2 would be the set containing all natural numbers divisible 2 (all the even numbers), X_3 would be the set of all natural numbers divisible by 3, and so on.

As one more example, let's suppose our indexing set was the interval $(0, 1)$ – the set of all real numbers x satisfying $0 < x < 1$. Now suppose for each $x \in (0, 1)$ we considered the set

$$Y_x = (-x - 1, x^3) = \{y \in \mathbb{R} \mid -x - 1 < y < x^3\}.$$

So, for example, $Y_{1/2}$ would be the interval $(-3/2, 1/8)$, and $Y_{3/4}$ would be the interval $(-7/4, 27/64)$. Graphically, you can picture the Y_x 's as the consisting of the y -coordinates that live on the vertical line segments connecting points on the graph $y = -x - 1$ to points on the graph $y = x^3$.



4.3.3 Unions and intersections with indices

Once we have an indexed family of sets (i.e., a family of sets S_i , one per each i in an index set I), we can consider operations like unions and intersections for all the sets in our collection. The union of all these S_i

would consist of the elements that are in at least one S_i and is denoted $\bigcup_{i \in I} S_i$:

$$\bigcup_{i \in I} S_i = \{x \mid \exists i \in I, x \in S_i\}.$$

For example, using our examples S_i , X_n , and Y_x sets mentioned in the previous section, we would have

$$\bigcup_{i \in I} S_i = \mathbb{Z}, \quad \bigcup_{n \in \mathbb{N}} X_n = \mathbb{N}, \quad \text{and} \quad \bigcup_{x \in (0,1)} Y_x = (-2, 1).$$

The intersection of an indexed family of sets, denoted $\bigcap_{i \in I} S_i$, consists of all the x that are elements of every S_i :

$$\bigcap_{i \in I} S_i = \{x \mid \forall i \in I, x \in S_i\}.$$

Using our three families of sets from before, their intersections are

$$\bigcap_{i \in I} S_i = \{2\}, \quad \bigcap_{n \in \mathbb{N}} X_n = \emptyset, \quad \bigcap_{x \in (0,1)} Y_x = (-1, 0).$$

4.4 Maps between sets


The material in this chapter will be crucial when we begin to discuss random variables later, however we will not need this material until then. You may want to only skim over this chapter for now, then return to it once we begin discussing random variables in class.

4.4.1 Definitions and examples

Given two sets A and B , a **map** from A to B (also called a **function** from A to B) is a rule which associates to each element of A an element of B . Sometimes we will represent maps pictorially by drawing A on the left, B on the right, and then having arrows going from elements of A to elements of B .

Example 4.18.

Suppose $A = \{a, b, c, d\}$ and $B = \{1, 2, \dots, 6\}$. One possible map from A to B is pictured below:



The map in Example 4.18 associates 1 to a ; associates 2 to b ; 2 is also associated to c ; and finally d gets associated to 6.

It is convenient to give a map a name so that we can refer to it without drawing pictures like this all of the time. Let's refer to the map from Example 4.18 as f . To say that f takes elements of A and associates an element of B to them we write $f : A \rightarrow B$. We then call A the *domain* of f and B is called the *codomain* of f . The *range* of f is the subset of B which actually get associated to an element of A . For the map in Example 4.18 the range is $\{1, 2, 6\}$.

There are several different notations that are used to describe which elements of B a map associates to elements of A . Some commonly used ones are $f(a) = b$ and $a \mapsto b$. The first one you've probably seen before, but the second one might be new. We pronounce $a \mapsto b$ as " a maps to b ."

Example 4.19.

Considering the map f shown in Example 4.18 we have

$$f(a) = 1$$

$$f(b) = 2$$

$$f(c) = 2$$

$$f(d) = 6$$

Using the other notation we would write

$$a \mapsto 1$$

$$b \mapsto 2$$


$$c \mapsto 2$$

$$d \mapsto 6$$

It is important to realize that a map $f : A \rightarrow B$ can only associate one element of B to a given element of A (even though there could be several elements of A associated to a given $b \in B$). A map $f : A \rightarrow B$ must also associate *every* element of A to something in B , even though not every element of B necessarily have something associated to it. (The range of $f : A \rightarrow B$ is by definition the set of all elements in B which have an element of A associated to them.)


Example 4.20.

The following is not a map:



Example 4.21.

The following is not a map:



4.4.2 Representing maps

It is common to represent a map by a formula, for example consider the map $f : \mathbb{Z} \rightarrow \mathbb{Z}$ which takes a given number and squares it. It's not really reasonable to represent this map pictorially since \mathbb{Z} has infinitely-many elements, so we instead describe the map by an algebraic rule and write $f(x) = x^2$ or $x \mapsto x^2$.


Another way to represent a map is to consider its graph. In general, the **graph** of a map $f : A \rightarrow B$, which we will denote $\text{Graph}(f)$, is a subset of $A \times B$ which consists of pairs of the form $(a, f(a))$. That is,

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

When we have a map from the set of real numbers \mathbb{R} (defined below) to itself, it is common to actually draw these points in the plane \mathbb{R}^2 . That is, given a map $f(x)$ we plot all of the pairs (x, y) where $y = f(x)$.

Example 4.22.

The graph of the map $f : \mathbb{R} \rightarrow \mathbb{R}$ given by $x \mapsto x^2 - 3x$ is



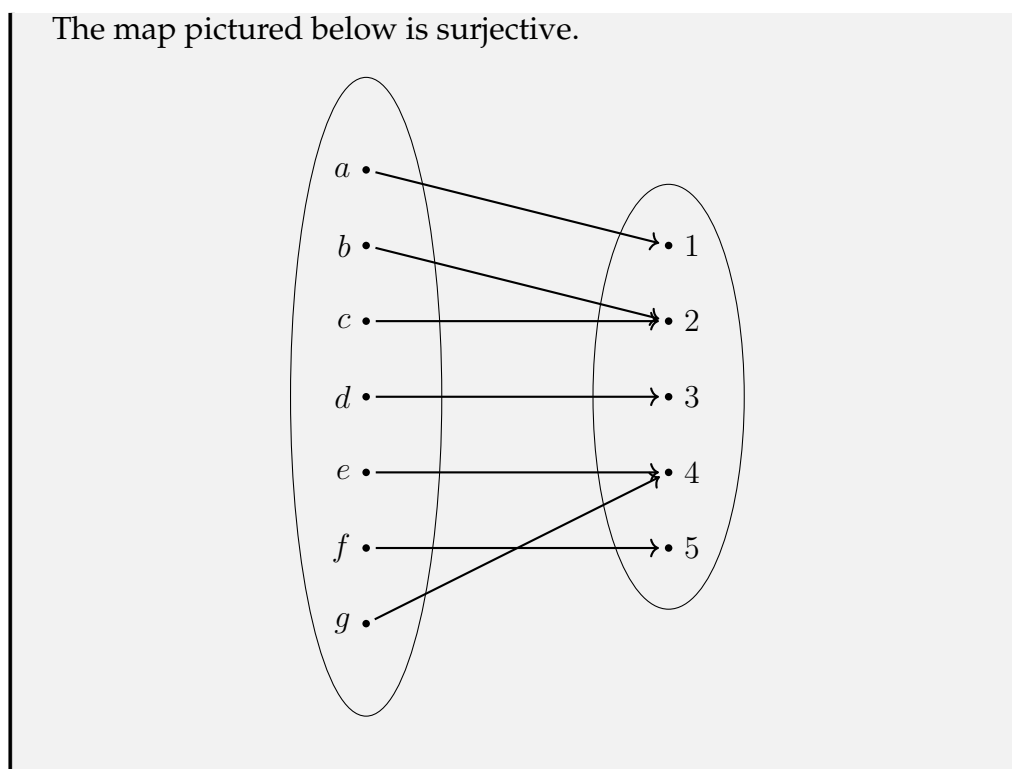
4.4.3 Special types of maps

As mentioned above, a map $f : A \rightarrow B$ must associate every element of A to some element of B (i.e., for every $a \in A$, $f(a)$ is defined), but not every element of B must have an element of A associated to it (there may be some $b \in B$ such that for every $a \in A$, $f(a) \neq b$). In the special case where every element of B *does* have an element of a associated to it, we say the map f is *surjective* or *onto*. Equivalently, a map is surjective when its codomain and range are the same.

Example 4.23.

The map $f : \mathbb{Z} \rightarrow \mathbb{Z}$ defined by $f(x) = x + 3$ is surjective. Every y in the codomain \mathbb{Z} gets associated an x from the domain, namely $x = y - 3$.

Example 4.24.

**Remark.**


The terms *surjective* and *onto* are completely synonymous, and which one a person uses is largely a matter of personal preference.

Notice that the map pictured in Example 4.24 has the property that multiple elements of the domain get associated to the same element in the codomain: both b and c get associated to 2, while both e and g are associated to 4. When this *does not* happen, we give the map a special name.

We say that a map $f : A \rightarrow B$ is *injective* or *one-to-one* (commonly denoted **1-1**) if each element of A is associated to a unique of B . That is, if a_1 and a_2 are distinct elements of A , then $f(a_1) \neq f(a_2)$.

Example 4.25.

The following map is injective.


**Example 4.26.**

The map $f(x) = x + 3$ from Example 4.23 is injective: if $x_1 \neq x_2$, then $f(x_1) = x_1 + 3 \neq x_2 + 3 = f(x_2)$.

When a map is both injective and surjective, we say the map is *bijjective*. Bijective maps play a special role in most areas of mathematics because having a bijection between two sets means those two sets are “the same.” That is, you may label the elements of the sets differently and think of them in different ways, but each element in one set has exactly one element in the other set associated to it: we can pair the elements of the sets together one by one.

Example 4.27.

The following map is injective.

**Example 4.28.**

The map $f(x) = x + 3$ from Example 4.23 is bijective as it is both surjective and injective.

Remark.

If we know that a given map $f : A \rightarrow B$ is injective, surjective, or bijective, then we also instantly know how the cardinalities of A and B are related. If f is injective, then $\#A \leq \#B$. If f is surjective, then $\#A \geq \#B$. If f is bijective, then $\#A = \#B$. This holds even when A and B have infinitely-many elements! These ideas can be used to make sense of when one “type” of infinity is bigger than another type of infinity.


The notion of different sizes of infinity was very controversial when first proposed by Georg Cantor in the late 19th century, but today is a commonly accepted and understood part of mathematics. For a very easy and brief introduction to the idea of different sizes of infinity, watch the short short TED-Ed video *How big is infinity?*,

<https://youtu.be/UPA3bwVVzGI>.

When a map $f : A \rightarrow B$ is bijective, there is always a map $g : B \rightarrow A$ which “undoes” f in the following sense: for every $a \in A$, $g(f(a)) = a$, and for every $b \in B$, $f(g(b)) = b$. We call the map g the *inverse* of f and usually denote it by f^{-1} . (Notice that f^{-1} is *not* f raised to the negative first power! This is simply a common, if unfortunate, notation for the inverse.)

Example 4.29.

The bijective map f is denoted in black in the image below, while its inverse f^{-1} is given in red.



4.4.4 Images and preimages

Just as a map $f : A \rightarrow B$ associates elements of B to elements of A , it also associates subsets of B to subsets of A by applying f to every element of a subset of A .


Suppose that $f : A \rightarrow B$ is any map and that $X \subseteq A$ is any subset of A . We can define a subset of B , which we'll denote $f(X)$, in the following way:

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

This set $f(X)$ is called the image of X under f .

Example 4.30.

Let $A = \{a, b, c, d\}$, $B = \{1, 2, \dots, 6\}$ and let f be the map from Example 4.18. If $X = \{a, b, c\}$, then its image $f(X)$ is $\{1, 2\}$.



Given any $Y \subseteq B$, the preimage of Y is the set of all elements in A which get mapped to an element of Y . The preimage is often denoted $f^{-1}(Y)$, even if f is not bijective.

$$f^{-1}(Y) = \{x \in A \mid f(x) \in Y\}$$

Example 4.31.

Let A, B , and f be as in Example 4.30. If $Y = \{1, 2\}$, then $f^{-1}(Y) =$

$\{a, b, c\}.$
Exercise 4.13.

Let A , B , and f be as in Example 4.30. What is the preimage of $\{3, 4, 5\}$? The preimage of $\{3, 4, 5\}$ is the empty set, \emptyset .

4.5 Compositions of Maps

4.5.1 Definitions and basic examples

We can sometimes take a pair of maps and combine them together to get a new map, called the “composition” of the original maps. In particular, we would like to take the output of one map and use it as the input to another map. For this to be defined, however, we need to be careful about the domains and codomains of the maps involved. In particular, let’s suppose that $f : A \rightarrow B$ is a map with domain A and codomain B , and that $g : B \rightarrow C$ is a map with domain B and codomain C . In this situation we can take an element $a \in A$, apply f to it to obtain an element $f(a) \in B$. Since $f(a)$ is in B , though, and g takes inputs from the set B , we can now apply g to $f(a)$ to get an element $g(f(a)) \in C$. This gives us a new map called the *composition* of f and g , sometimes denoted $g \circ f$, which is a map with domain A and codomain C ,

$$g \circ f : A \rightarrow C \text{ defined by } a \mapsto g(f(a)).$$

Notice that this notation $g \circ f$ tells you the functions being applied from the right to the left. That is, in $g \circ f$ the map f is applied first, and is followed by g .

Compositions of maps are sometimes indicated in simple diagrams such as

$$A \xrightarrow{f} B \xrightarrow{g} C$$

where following the arrows of the diagram indicates the composition where we start with an element of A , apply f to obtain an element of B , then apply g to obtain an element of C .

Notice that compositions often can not be “reversed.” That is, if $f : A \rightarrow B$ is a map from A to B , and $g : B \rightarrow C$ is a map from B to C , then we can define the composition $g \circ f$, but *we cannot* define the composition $f \circ g$. We can’t define $f \circ g$ because g takes an element of B as input and produces an element of C as an output, whereas f requires its inputs come from A , and so this composition $f \circ g$ is not defined.

As a few simple concrete applications of compositions, consider the map $f : \mathbb{N} \rightarrow \mathbb{Q}$ given by $f(n) = \frac{1}{n}$ and $g : \mathbb{Q} \rightarrow \mathbb{R}$ given by $g(x) = \sqrt{x}$. The composition $g \circ f : \mathbb{N} \rightarrow \mathbb{R}$ is a map which takes initial inputs from the natural numbers and ultimately produces real numbers as an output,

$$\mathbb{N} \xrightarrow{f} \mathbb{Q} \xrightarrow{g} \mathbb{R}.$$

The actual value of $g \circ f$ at a natural number n is determined by simply applying the maps in sequence:

$$(g \circ f)(n) = g(f(n)) = g\left(\frac{1}{n}\right) = \sqrt{\frac{1}{n}}.$$

For example, $(g \circ f)(2)$ is $\sqrt{\frac{1}{2}}$ and $(g \circ f)(9)$ is $\sqrt{\frac{1}{9}} = \frac{1}{\sqrt{9}} = \frac{1}{3}$.

4.5.2 Composing three or more functions; associativity

We can extend our operation of composition by chaining multiple compositions together. For instance if we have maps $f : A \rightarrow B$, $g : B \rightarrow C$, and $h : C \rightarrow D$, then we can consider the composition $h \circ (g \circ f)$, which takes an element $a \in A$ and ultimately maps it to $h(g(f(a)))$.

$$A \xrightarrow{f} B \xrightarrow{g} C \xrightarrow{h} D$$

For example, letting $f : \mathbb{N} \rightarrow \mathbb{Q}$ and $g : \mathbb{Q} \rightarrow \mathbb{R}$ be the maps earlier, we could compose with the map $h : \mathbb{R} \rightarrow \{0, 1\}$ defined by

$$h(x) = \begin{cases} 0 & \text{if } \lfloor x \rfloor \text{ is even} \\ 1 & \text{if } \lfloor x \rfloor \text{ is odd} \end{cases}$$

This map produces a 0 or 1 based on whether the floor of its argument is even or odd: e.g., $h(\pi) = 1$ as $\lfloor \pi \rfloor = 3$ is odd; and $h(e) = 0$ as $\lfloor e \rfloor = 2$ is even.

The composition of all three maps, $h \circ (g \circ f)$,

$$\mathbb{N} \xrightarrow{f} \mathbb{Q} \xrightarrow{g} \mathbb{R} \xrightarrow{h} \{0, 1\}$$

is map from \mathbb{N} to $\{0, 1\}$ which could be expressed as

$$(h \circ (g \circ f))(n) = \begin{cases} 0 & \text{if } \left\lfloor \sqrt{\frac{1}{n}} \right\rfloor \text{ is even} \\ 1 & \text{if } \left\lfloor \sqrt{\frac{1}{n}} \right\rfloor \text{ is odd} \end{cases}$$

Notice that if we compose three maps,

$$A \xrightarrow{f} B \xrightarrow{g} C \xrightarrow{h} D$$

there are two conceivable ways we could get a composition: we could consider $(h \circ g) \circ f$ or $h \circ (g \circ f)$. The difference between these two options is which composition is “constructed” first before being composed with the remaining map. Conveniently, though, these are actually the same map. That is, the map from A to D given by $(h \circ g) \circ f$ will equal the map from A to D given by $h \circ (g \circ f)$. The reason for this is that for each $a \in A$, both expressions will take us to the same element of D :

$$\begin{aligned} ((h \circ g) \circ f)(a) &= (h \circ g)(f(a)) \\ &= h(g(f(a))) \\ &= (h \circ (g \circ f))(a) \end{aligned}$$

So we don’t actually need to bother with writing all these parentheses and can just write $h \circ g \circ f$. This property of sliding parentheses around (choosing different precedences for which maps are composed first) without changing the final result is called the ***associative property of composition***.

We can continue to consider compositions of several maps between sets, as long as we are careful that the codomain of a map corresponds to the domain of the “next” map in the composition. For instance if we have maps

$$A \xrightarrow{f} B \xrightarrow{g} C \xrightarrow{h} D \xrightarrow{i} E \xrightarrow{j} F \xrightarrow{k} G$$

then we can define the composition $k \circ j \circ i \circ h \circ g \circ f$ to define a map with domain A and codomain G .

Remark.

In more advanced branches of mathematics such as algebraic topology, algebraic geometry, and homological algebra, it is actually common to consider certain special compositions of infinitely-many maps. It seems unwieldy at first glance, but under certain circumstances you can use properties of some finite number of the compositions to get information about other parts of the infinite composition, and

this is one common “trick” used to make complex calculations much simpler in those areas of math.

Exercise 4.14.

Let $f : X \rightarrow Y$ and $g : Y \rightarrow Z$ be two maps between sets.

- (a) Show that if f and g are both injective, then the composition $g \circ f$ is injective.
- (b) Show that if f and g are both surjective, then the composition $g \circ f$ is surjective.

4.5.3 Inverse maps

When a map $f : A \rightarrow B$ is bijective, we can define a map called the *inverse map* of f , denoted f^{-1} , which is a map from B to A which “undoes” the map f . That is, for $b \in B$ we define $f^{-1}(b)$ to be the element $a \in A$ so that $f(a) = b$. Notice that such an a must exist since f is surjective, and that element a is unique because f is injective. Thus we have a well-defined map $f^{-1} : B \rightarrow A$.

Exercise 4.15.

Suppose $f : A \rightarrow B$ is bijective and $f^{-1} : B \rightarrow A$ is its inverse. Show that f^{-1} is also bijective.

Two key properties of bijective maps and their inverses are the following:

$$\forall a \in A, f^{-1}(f(a)) = a, \text{ and}$$

$$\forall b \in B, f(f^{-1}(b)) = b.$$

These two properties justify the idea that f and f^{-1} “undo” one another.

Example 4.32.

The following are examples of bijective maps with their inverses.

- The map $f : \mathbb{R} \rightarrow \mathbb{R}$ given by $f(x) = x^3$ is bijective with inverse $f^{-1}(x) = \sqrt[3]{x}$.
- The map $f : \mathbb{R} \rightarrow \mathbb{R}$ given by $f(x) = 2x + 3$ is bijective with inverse $f^{-1}(x) = \frac{1}{2}(x - 3)$.
- The map $f : \mathbb{Z} \rightarrow \mathbb{Z}$ given by $f(x) = x + c$, where c is any fixed, constant integer, is bijective with inverse $f^{-1}(x) = x - c$.

Exercise 4.16.

Suppose that $f : A \rightarrow B$ and $g : B \rightarrow C$ are both bijective with inverses $f^{-1} : B \rightarrow A$ and $g^{-1} : C \rightarrow B$. Show that the composition $g \circ f : A \rightarrow C$ is also bijective and its inverse equals $f^{-1} \circ g^{-1} : C \rightarrow A$.

4.5.4 Identity maps

Every set A has a special map associated to it called the identity map, and denoted $\text{id}_A : A \rightarrow A$. This is the map which simply sends each element of A to itself; it is the map that “does nothing.” For each $a \in A$, $\text{id}_A(a)$ is defined to be a .

Notice that our earlier properties for a bijective map and its inverse can be expressed in terms of identity maps as follows: if $f : A \rightarrow B$ is bijective and $f^{-1} : B \rightarrow A$ is its inverse, then

$$f^{-1} \circ f = \text{id}_A \text{ and } f \circ f^{-1} = \text{id}_B.$$

Exercise 4.17.

Let $f : A \rightarrow B$ be any map (no assumptions about injectivity or surjectivity). Show that $\text{id}_B \circ f = f$ and $f \circ \text{id}_A = f$.

Relations

5.1 Basic definitions and examples

We often want to compare mathematical objects to one another and establish various “relationships” between quantities. For example, when considering integers we may compare integers together by saying that one is less than another, or say that two integers are related if one divides the other. As we will see, maps between sets can also be thought of as a relationship between mathematical objects (elements in the domain and in the codomain of the function). There are also special types of relationships called “equivalence relations” that permeate mathematics and which you have seen examples of before, even if you didn’t realize it at the time. We will get started, though, by first simply defining what “relation” between two sets is, and then seeing several examples.

Given two sets X and Y , a **relation** between X and Y is simply a subset R of $X \times Y$. Simply defining a relation has no other restrictions: it’s just some chosen subset of ordered pairs (x, y) . (There are special types of relations that do have restrictions, but we’ll discuss those later.)

For example, if $X = \{1, 2, 3, 4\}$ and $Y = \{a, b, c\}$, then one possible relation would be

$$R = \{(1, a), (1, b), (2, b), (3, b), (3, a)\}.$$

If we have an element $(x, y) \in R$, we sometimes write xRy . For example, in the relation described above, we have $1Ra$ and $3Ra$.

Notice that a map $f : X \rightarrow Y$ determines a special type of relation: given any such map we can consider a relation where the ordered pairs of $X \times Y$ are of the form $(x, f(x))$ for each $x \in X$. This is a perfectly legitimate subset of $X \times Y$, and so it defines a relation. In general, this particular relation determined by a function is sometimes called the **graph** of the function and denoted $\text{Graph}(f)$,

$$\text{Graph}(f) = \{(x, y) \in X \times Y \mid y = f(x)\}.$$

This generalizes the familiar notion of a graph of a function you’re familiar with from algebra and calculus.

As an example, consider the map $f : \mathbb{R} \rightarrow \mathbb{R}$ defined by $f(x) = x^2$. The graph of this function (in the algebra/calculus sense you’re familiar with) is a parabola in the plane, the (x, y) pairs that satisfy $y = x^2$. We can interpret this as a subset of points (x, y) in $\mathbb{R}^2 = \mathbb{R} \times \mathbb{R}$, namely those

(x, y) -pairs where the y value is equals to x^2 for each x . This is exactly the object $\text{Graph}(f)$ described above:

$$\text{Graph}(f) = \{(x, y) \in \mathbb{R} \times \mathbb{R} \mid y = x^2\}.$$

So, every map $f : X \rightarrow Y$ defines a relation we call the graph. Notice, though, that not every relation is necessarily the graph of a map. For example, the relation R between $X = \{1, 2, 3, 4\}$ and $Y = \{a, b, c\}$ above is not the graph of a map, for a few reasons:

1. The element $1 \in X$ appears in two different entries in R , $(1, a)$ and $(1, b)$.
2. The element $3 \in X$ appears in two different entries in R , $(3, a)$ and $(3, b)$.
3. The element $4 \in X$ does not appear in any entry in R .

In general, for a relation $R \subseteq X \times Y$ to be the graph of a map, the relation must have that each x appears as the first entry in one and one element of R .

Exercise 5.1.

Determine if each of the relations between $X = \{1, 2, 3, 4\}$ and $Y = \{a, b, c\}$ below is a graph of a map or not.

- (a) $R = \{(1, a), (2, b), (3, c), (4, a)\}$
- (b) $R = \{(1, a), (2, a), (3, a), (4, a)\}$
- (c) $R = \{(1, c), (2, b), (2, a), (3, c)\}$
- (d) $R = \{(1, a), (1, b), (1, c)\}$

Since some relations are graphs of functions, but not all are, you can think of a relation between sets as being a generalization of the idea of a function.

The sets X and Y that appear in our definition of a relation don't have to be distinct; that is, we could consider a subset R of $X \times X$, and we call this a relation on X . For instance, if $X = \mathbb{Z}$, a subset R of $\mathbb{Z} \times \mathbb{Z}$ is a relation

\mathbb{Z} . You actually already know some common relations on \mathbb{Z} , even if you never thought of them as subsets of $\mathbb{Z} \times \mathbb{Z}$.

Of course, any map $f : \mathbb{Z} \rightarrow \mathbb{Z}$ defines a relation (the graph), but some other relations you already know are the following:

Divisibility

If we consider the subset of $\mathbb{Z} \times \mathbb{Z}$ where the first entry divides the second entry,

$$\{(x, y) \in \mathbb{Z} \times \mathbb{Z} \mid \exists n \in \mathbb{Z}, y = nx\}$$

then this defines a relation. Let's momentarily refer to this relation as R . We would then have elements like $(2, 8) \in R$ and $(-3, 12) \in R$ and we might write this as $2R8$ and $-3R12$. Notice that if instead of using the letter R for this set we use the symbol $|$, this becomes our familiar notation for divisibility, $2|8$ and $-3|12$.

Inequality

Now consider the subset of $\mathbb{Z} \times \mathbb{Z}$ where the first entry is less-than-or-equal-to the second entry,

$$R = \{(x, y) \in \mathbb{Z} \times \mathbb{Z} \mid x \text{ is less than } y \vee x = y\}$$

then we have a relation where, for instance, $0R5$ and $-7R-1$. Letting the symbol \leq denote this set (i.e., $\leq = R$), we then have our familiar notation $0 \leq 5$ and $-7 \leq -1$.

Neither of the examples described above is the graph of a function, because for example $2|2$ and $2|4$; and $0 \leq 0$ and $0 \leq 1$. That is, there are "x-coordinates" or our ordered pairs which appear multiple times, and this can't happen for functions.

As another relation on the integers, we could consider ordered pairs (x, y) where we require that either both x and y are even, or both x and y are odd. This is sometimes stated more succinctly by saying that x and y have the same *parity*, which just means they're both even or they're both odd, whichever it happens to be. Writing this down "directly" in set-builder notation is a bit cumbersome, but let's notice that we could rephrase this saying that the difference between x and y is divisible by 2.

To see this, notice that if x and y were both even, say $x = 2m$ and $y = 2n$ for integers m and n , then

$$x - y = 2m - 2n = 2(m - n)$$

and so the difference of two even numbers is divisible by 2. If we had x and y were both odd, say $x = 2m + 1$ and $y = 2n + 1$ for integers m and n , then

$$x - y = 2m + 1 - (2n + 1) = 2m + 1 - 2n - 1 = 2m - 2n = 2(m - n),$$

and again the difference of two odd numbers is divisible by 2. Finally, if one of the numbers was even and one was odd, say $x = 2m + 1$ and $y = 2n$ for integers m and n , then

$$x - y = 2m + 1 - 2n = 2(m - n) + 1$$

which will not be divisible by 2, since we would have a remainder of 1.

Putting this all together, we can define relation R where xRy if x and y have the same parity (both are even or both are odd):

$$R = \{(x, y) \in \mathbb{Z} \times \mathbb{Z} \mid 2 \mid (x - y)\}$$

So, for example, $-2R14$ and $3R127$. This particular relation is sometimes expressed as saying x and y are ***congruent modulo 2*** if xRy where R is the relation above. This relation also has some special notation where we write \equiv_2 for R . I.e., $-2 \equiv_2 14$ and $3 \equiv_2 127$.

We can extend this notion of congruence for any integer. That is, for any integer m we can define a relation \equiv_m on $\mathbb{Z} \times \mathbb{Z}$ by saying $x \equiv_m y$ (i.e., the pair (x, y) is an element of the subset $equiv_m \subseteq \mathbb{Z} \times \mathbb{Z}$ we are about to define) if their difference is divisible by m , $m \mid (x - y)$,

$$\equiv_m = \{(x, y) \in \mathbb{Z} \times \mathbb{Z} \mid m \mid (x - y)\}.$$

So, for example, $5 \equiv_4 17$ since $17 - 5 = 12$ which is divisible by 4; and $10 \equiv_9 19$ since $19 - 10 = 9$ which is divisible by 9. This special relation we have defined, called ***congruence modulo m*** , is also sometimes written as

$$x = y \pmod{m}$$

For example,

$$5 = 17 \pmod{4} \text{ and } 10 = 19 \pmod{9}.$$

5.2 Special properties of relations

If you spent a lot of time thinking about the various common relations we could define on a set, you might begin to notice there are a few special properties that are common among many different relations. It might be convenient, then, to give some of these commonly recurring properties special names, so that you could easily state that a relation has a particular property without redefining it each time. Here we mention a few of the most common properties.

5.2.1 Symmetry

Many relations we will consider will be *symmetric* in the sense that xRy if and only if yRx . For example, in our congruence relation \equiv_m defined in the last section, it's easy to verify that $x \equiv_m y$ if and only if $y \equiv_m x$. To see this, suppose x and y are integers with $x \equiv_m y$. Then, by definition, this means $m|(x - y)$, and so $x - y = mn$ for some integer n . But then since $y - x$ is just the negative of $x - y$, $y - x = m(-n)$ and so $m|(y - x)$ as well and we have $y \equiv_m x$. (Since we're stating this as xRy "if and only if" yRx , we technically still need to show that if $y \equiv_m x$ then we must have $x \equiv_m y$. The proof, however, is exactly the same as the proof above, just with the roles of x and y reversed.)

Notice that not all relations will be symmetric. For example, $2|4$ but $4 \nmid 2$. Thus symmetry is a special property of some, but not all, relations.

5.2.2 Antisymmetry

Another common property many relations satisfy is called "antisymmetry." This is not simply the "opposite" of the symmetry defined above, but means something a bit more precise. We say a relation R on a set X (i.e., $R \subseteq X \times X$) is *antisymmetric* if when xRy and yRx , we must have $x = y$. For example, the relation \leq on the integers has this property: if $x \leq y$ and $y \leq x$, then $x = y$. The divisibility relation on the natural numbers also has this property:

Lemma 5.1. *If $x, y \in \mathbb{N}$ with $x|y$ and $y|x$, then $x = y$.*

Proof. Suppose x and y are integers where x divides y and y divides x . Since $x|y$ we can write $y = xm$ for some integer m , and similarly as $y|x$ we have $x = yn$ for some integer n . We may use this to write

$$x = ym = (xn)m = xnm.$$

We could subtract xnm from both sides of the equation to write $x - xnm = 0$ or $x(1 - nm) = 0$. This means either $x = 0$ or $1 - nm = 0$. However, as we working over the integers, we can not have $x = 0$. Thus we must consider $1 - nm = 0$ meaning $nm = 1$. The only positive integer solutions to this, though are $n = m = 1$ and so our equations $x = ym$ and $y = xn$ both simply become $y = x$. \square

Exercise 5.2.

Is the divisibility relation on \mathbb{Z} antisymmetric? That is, if x and y are integers (including zero and the negatives) with $x|y$ and $y|x$, is it true $x = y$?

Exercise 5.3.

Is it possible for a relation R on a set X to be both symmetric and antisymmetric? If so, give an example. If not, prove no such relation can exist on any set X .

Exercise 5.4.

Is it possible for a relation to be neither symmetric nor antisymmetric? If so, give an example. If not, prove that every relation must be symmetric or antisymmetric.

5.2.3 Reflexivity

We call a relation R on a set X reflexive if for every $x \in X$ we have xRx . Many of our examples of relations thus far are reflexive:

- The divisibility relation is reflexive since every integer divides itself.

- The \leq relation is reflexive as for every x we have $x \leq x$.
- The congruence relations on the integers are reflexive as for every x we have $x - x = 0$ and every integer divides 0, so $x \equiv_m x$ for every integer x and every integer m .

5.2.4 Transitivity

A relation R is called transitive if xRy and yRz implies xRz . For example, the inequality \leq is transitive as if $x \leq y$ and $y \leq z$, then we must have $x \leq z$ as well.

Exercise 5.5.

Show that the divisibility relation on the integers is transitive. That is, show that if $x, y, z \in \mathbb{Z}$ with $x|y$ and $y|z$, then we must have $x|z$ as well.

5.3 Orderings

Some of the most important relations on a set are “orderings,” which give us a way of comparing different elements in a set. For example, you’re familiar with the orders of $<$, \leq , $>$, and \geq for the real numbers (and subsets of the reals such as integers and rationals). We will define two different types of orderings called “total orderings” and “partial orderings” that generalize the essential properties of these relations.

5.3.1 Total orders

A **total order** (also sometimes called a **linear order**) on a set X is a relation R that has the following properties:

- R is reflexive (i.e., $\forall x \in X, xRx$);
- R is transitive (i.e., $\forall x, y, z \in X, (xRy \wedge yRz) \implies xRz$);
- R is antisymmetric (i.e., $\forall x, y \in X, (xRy \wedge yRx) \implies x = y$); and
- All elements of X are comparable under R (that is, for every $x, y \in X$ we have either xRy or yRx : symbolically, $\forall x, y \in X, xRy \vee yRx$).

The first three properties we have seen defined in the last section, but the last property is new. It simply means there every pair of elements can be compared by our relation R .

As mentioned above, the “obvious” examples of total orders are the orderings you already know and love on the real numbers. It will in fact turn out that every set can be given a total order (in fact, a very special type of total order called a “well-order”), but we will postpone discussing that for now since it requires a detour through a more sophisticated (and in some ways controversial) part of set theory called the “axiom of choice.”

Exercise 5.6.

Construct a total order on the set $\{\alpha, B, \odot, 13\}$. How many total possible total orders are there on this set? More generally, how many possible total orders are there on a finite set of cardinality n ?

5.3.2 Partial orders

While it will turn out that every set can be given a total order, many of the relations we will care about won't satisfy the fourth condition of a total order. That is, not all orders we will care about will allow us to compare every possible pair of elements in the set. These are referred to as “partial orders.”

To be more precise, a *partial order* on a set X is a relation R on X that is reflexive, transitive, and antisymmetric. Notice that every total order is necessarily also a partial order, just a very special type of partial order. Thus the “obvious” examples of total orders (the usual orderings on the real numbers and its subsets) are also examples of partial orders.

As for examples of partial orders that *are not* total orders, we have already seen a few:

Example 5.1.

The relation of divisibility is an ordering on the the natural numbers, \mathbb{N} . We can easily verify each of the three required properties:

Reflexivity The divisibility relation is reflexive since every natural number divides itself: $n|n$ as $n = 1 \cdot n$.

Transitivity Suppose $m, n,$ and p are natural numbers where $m|n$ and $n|p$. That is, $n = am$ and $p = bn$ for some integers a and b . We then have

$$p = bn = b(am)$$

and so $m|p$ as well.

Antisymmetry Suppose m and n are natural numbers with $m|n$ and $n|m$. Then $n = am$ and $m = bn$ for integers a and b . We may write $n = am = abn$ and so $n - abn = 0$, or $n(1 - ab) = 0$. Since we are only working with natural numbers, n can not be zero, and we must have $1 - ab = 0$. The only natural number solution to this, though is $a = b = 1$, and so $n = m$.

Thus divisibility is a partial order, but it is not a total order. To verify this we just need to show that there natural numbers m and n where $n \nmid m$ and $m \nmid n$. One easy example is $n = 2$ and $m = 3$.

Example 5.2.

Recall that given a set X , its powerset (denoted $\mathcal{P}(X)$ or 2^X) is the set of all subset of X . There is a partial ordering on the powerset of any set X given by the subset relation. That is, we consider the relation $A \subseteq B$ for subsets A and B of X .

We can check this satisfies the three necessary properties to be a partial order:

Reflexivity Every subset A of X (i.e., every element $A \in \mathcal{P}(X)$) is a subset of itself, so $A \subseteq A$.

Transitivity If A, B, C are subsets of X with $A \subseteq B$ and $B \subseteq C$, then $A \subseteq C$ as every element $a \in A$ is also an element of B , and every element of B (which includes those elements of A since $A \subseteq B$) is an element of C .

Antisymmetry If A and B are subsets of X where $A \subseteq B$ and $B \subseteq A$, then $A = B$.

Notice that the partial order \subseteq is not a total order if X contains at least two elements. In this case say the elements are $x_1 \neq x_2$, then the sets $\{x_1\}$ and $\{x_2\}$ are not comparable: neither is a subset of the other.

5.4 Equivalence relations

Another special type of relation on a set is an “equivalence relation,” which gives us a notion of two elements in a set being equivalent, and sets of equivalent elements form “equivalence classes” which are subsets of the set. It will turn out that equivalence relations are ubiquitous in mathematics: many objects you are familiar with are in fact equivalence classes. In fact, equivalence classes are necessary to even carefully define what a real number is in a precise way.

An *equivalence relation* on a set X is a relation which is reflexive, symmetric, and transitive. Recall this means that for every $x \in X$ we must have xRx (reflexivity); for every x and y in X we have that xRy if and only if yRx (symmetry); and for all $x, y, z \in X$ if xRy and yRz , then xRz as well (transitivity).

5.4.1 Examples of equivalence relations

One example of an equivalence relation is the congruence relationship between integers we defined earlier. Recall that for a fixed non-zero integer $m \in \mathbb{Z}$, we consider a relation denoted $x \equiv_m y$ defined by the condition that $m|(x - y)$.

Exercise 5.7.

Check that congruence modulo m is an equivalence relation on \mathbb{Z} .

We often use the symbol \sim to denote an equivalence relation. That is, we write $x \sim y$ if \sim is an equivalence relation on some set X containing x and y , and say that x is equivalent to y .

Example 5.3.

As another example of an equivalence relation on pairs of non-zero integers, $(a, b) \in (\mathbb{Z} \setminus \{0\})^2$, we may consider two pairs of non-zero integers, (a, b) and (c, d) to be equivalent if $ad = bc$. Let's first check that this is in fact an equivalence relation by verifying it is reflexive, transitive, and symmetric.

For notation convenience, let us write $(a, b) \sim (c, d)$ if $ad = bc$.

Reflexivity Notice that for any (a, b) we have $(a, b) \sim (a, b)$ as $ab = ba$.

Symmetry Suppose $(a, b) \sim (c, d)$, and so by assumption $ad = bc$. We need to show $(c, d) \sim (a, b)$ as well, meaning we need to verify $cb = da$. Of course, if $ad = bc$ then $cb = da$ since multiplication of integers is commutative.

Transitivity Suppose $(a, b) \sim (c, d)$ and $(c, d) \sim (e, f)$. We wish to show $(a, b) \sim (e, f)$ meaning we need to verify that $af = be$. By assumption we know $ad = bc$ and $cf = de$. Now observe

$$\begin{aligned} ad &= bc \\ \implies adef &= bcef \quad (\text{multiplying both sides by } ef) \\ \implies afde &= bec f \quad (\text{rearranging the factors}) \\ \implies afde &= bede \quad (\text{using } de = cf) \\ \implies af &= be \quad (\text{cancelling } de \text{ from each side}) \end{aligned}$$

Example 5.4.

One example of an equivalence relation that is important in a branch of mathematics called "algebraic geometry" is the following. We take the set X to be the set of all points in the plane \mathbb{R}^2 except for the origin: $X = \mathbb{R}^2 \setminus \{(0, 0)\}$. Now we declare two of these points, say (x_1, y_1) and (x_2, y_2) to be equivalent there is some $\lambda \neq 0$ so that $x_1 = \lambda x_2$ and $y_1 = \lambda y_2$. We can check that this is an equivalence relation:

Reflexivity $(x, y) \sim (x, y)$ by taking $\lambda = 1$.

Symmetry If $(x, y) \sim (x', y')$, then there is some non-zero λ so that $x = \lambda x'$ and $y = \lambda y'$. We can rewrite these equations as $x' = \frac{1}{\lambda}x$ and $y' = \frac{1}{\lambda}y$. This shows (x', y') is a multiple of (x, y) and so $(x', y') \sim (x, y)$.

Transitivity Suppose $(x, y) \sim (u, v)$ and $(u, v) \sim (s, t)$. We must show $(x, y) \sim (s, t)$. Notice that there are numbers λ and μ so that $x = \lambda u$, $y = \lambda v$, and $u = \mu s$, $v = \mu t$. But then we easily see $x = \lambda u$ becomes $x = \lambda \mu s$ and $y = \lambda \mu t$.

5.4.2 Equivalence classes

If \sim is an equivalence relation on X , then we can partition X into “equivalence classes.” An *equivalence class* for \sim is a subset of X consisting of elements which are all equivalent to one another, and containing all of the equivalent elements. For example, if \sim is the equivalence relation on \mathbb{Z} given by $x \sim y$ if x and y have the same parity (equivalently, $x \equiv_2 y$ or $2|(x - y)$), then there are two equivalence classes: the even integers and the odd integers.

As another example, if we consider the equivalence relation \equiv_3 on \mathbb{Z} (so $x \equiv_3 y$ if $3|(x - y)$) there are three equivalence classes, indicated below:

$$\{\dots, -9, -6, -3, 0, 3, 6, 9, \dots\}$$

$$\{\dots, -8, -5, -2, 1, 4, 7, 10, \dots\}$$

$$\{\dots, -7, -4, -1, 2, 5, 8, 11, \dots\}$$

Let us momentarily refer to these subsets of \mathbb{Z} as S_0 , S_1 , and S_2 (meaning the three sets described above which contain 0, 1, or 2, respectively). Notice that these partition \mathbb{Z} :

$$\mathbb{Z} = S_0 \cup S_1 \cup S_2.$$

Remark.

Because the sets S_0 , S_1 , and S_2 together form a partition of \mathbb{Z} , it's not simply that their union is \mathbb{Z} , but each pair of sets is disjoint: $S_i \cap S_j = \emptyset$ for any choice of $i, j \in \{0, 1, 2\}$ with $i \neq j$. This is sometimes indicated by replacing the symbol \cup with \sqcup which means *disjoint union*. That is, if we write $A = B \sqcup C$ that means not only is A the union of B and C , but it also means B and C are disjoint, $B \cap C = \emptyset$. Thus above we could write $\mathbb{Z} = S_0 \sqcup S_1 \sqcup S_2$.

In the case of the equivalence relation on $\mathbb{R}^2 \setminus \{(0, 0)\}$ described in Example 5.4, the equivalence classes are precisely lines in \mathbb{R}^2 which go through the origin (but with the origin removed).

We often want to consider the set of all equivalence classes of some equivalence relation \sim on a set X . This set is denoted X/\sim . As a simple toy example, let us consider the set

$$X = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}.$$

We will define an equivalence relation on X by declaring $1 \sim 2$, $2 \sim 4$, $3 \sim 5$, $3 \sim 6$, $4 \sim 8$, $7 \sim 7$, $9 \sim 10$. (Technically we haven't completely described the relation since we would also need to declare $1 \sim 4$, for example, and $6 \sim 3$. However, we can extend the list above uniquely to define an equivalence relation, and we are implicitly doing that.) The equivalence classes would then be

$$\begin{aligned} &\{1, 2, 4, 8\}, \\ &\{3, 5, 6\}, \\ &\{7\}, \text{ and} \\ &\{9, 10\}. \end{aligned}$$

And so X/\sim would be the set containing these equivalence classes,

$$X/\sim = \{\{1, 2, 4, 8\}, \{3, 5, 6\}, \{7\}, \{9, 10\}\}.$$

This becomes very cumbersome to write, so we adopt a common notation that makes it easier to describe the equivalence classes. Given any set X with equivalence relation \sim and any $x \in X$, we often refer to the equivalence class containing x as $[x]$. So, for example, in the example above we have $[1] = \{1, 2, 4, 8\}$ and $[3] = \{3, 5, 6\}$. Notice that we can

pick any element of the equivalence class and it would describe the same set. For instance, $[2] = \{1, 2, 4, 8\}$, $[4] = \{1, 2, 4, 8\}$, and $[8] = \{1, 2, 4, 8\}$ as well. Because these sets are all equal to each other, we are justified in writing $[1] = [2] = [4] = [8]$, as well as $[3] = [5] = [6]$ in our example above.

In the case of our equivalence relation \equiv_3 on \mathbb{Z} above, the set of equivalence classes can be described as

$$\mathbb{Z}/\equiv_3 = \{[0], [1], [2]\}.$$

Sometimes for certain commonly used equivalence relations, the notation for equivalence classes will be a little bit different. For example, when discussing congruence modulo m , it is common to see the equivalence class of a number $[x]$ to be written as \bar{x} . So, for example, the equivalence classes \mathbb{Z}/\equiv_3 are often denoted as $\bar{0}$, $\bar{1}$, and $\bar{2}$. This set of equivalence classes, which we've so far denoted \mathbb{Z}/\equiv_3 is often denoted \mathbb{Z}_3 . Of course, we can consider congruence modulo other values. Here are a few other collections of equivalence classes of \mathbb{Z} under congruence modulo a number:

$$\begin{aligned}\mathbb{Z}_2 &= \{\bar{0}, \bar{1}\} \\ \mathbb{Z}_3 &= \{\bar{0}, \bar{1}, \bar{2}\} \\ \mathbb{Z}_4 &= \{\bar{0}, \bar{1}, \bar{2}, \bar{3}\} \\ \mathbb{Z}_5 &= \{\bar{0}, \bar{1}, \bar{2}, \bar{3}, \bar{4}\}\end{aligned}$$

Remark.

The equivalence classes of \mathbb{Z} under congruence modulo m (i.e., elements of \mathbb{Z}_m) are often referred to as *congruence classes*.

The set of equivalence classes from the equivalence relation defined in Example 5.4 is often denoted \mathbb{RP}^1 and called the *real projective line*. The equivalence classes themselves, as we mentioned above, are lines through the origin in \mathbb{R}^2 but with the origin removed. Given a point (x, y) , we often refer to the equivalence class containing that point (i.e., the line through (x, y) which passes through the origin, sans the origin itself) as $[x : y]$, and these are called the *homogeneous coordinates* of that

point in \mathbb{RP}^1 . Notice that each such line determines a unique point on the upper half of the unit circle, we can think of each equivalence class as being a point on that upper half of the circle, and the 'ends' of that (corresponding to the points $(1, 0)$ and $(-1, 0)$ in the plane) are actually in the same equivalence class, and so are identified. This shows that we can think of \mathbb{RP}^1 as being a circle!

Remark.

This idea of identifying points that lie on the same line that goes through the origin is fundamental in algebraic geometry. It turns out that if we consider these kinds of equivalence classes, we can do geometry on the set of all equivalence classes and many statements in geometry are actually greatly simplified in this setting, even though the setup seems very abstract the first time you encounter it.

6

Binary Operators

Mathematics is the study of analogies between analogies. All science is. Scientists always want to show that things that don't look alike are really the same. That's one of their innermost Freudian motivations. In fact, that's what we mean by "understanding."

GIAN-CARLO ROTA
*Mathematics, Philosophy, and Artificial
Intelligence*

In this chapter we discuss the notion of a “binary operation” on a set, which is an abstract generalization of many familiar notions in mathematics. As you study more mathematics, you may observe that it is extremely common to combine two objects of the same mathematical type in some way to get another object of the same type. For example, we may add two integers to obtain a new integer; we might multiply two rational numbers to get another rational number; we could compose two functions to get another function; or we might multiply two square matrices to get another square matrix.

Here we will discuss the general idea of “combining two objects to get a third,” where the type of objects we combine come from some fixed set, be it the the set of integers, the rationals, or something else. After giving the basic definitions and some exmaples, we will spend some time discussing desirable properties that our binary operations may have, and finally spend some time discussing one particularly important family of examples of binary operators.

6.1 Definitions and examples

A **binary operator** on a set X is simply a map that takes two elements of X as inputs, and produces an element of X as an output. That is, it is simply a map from $X \times X$ to X . In general this can be any arbitrary map of the form $X \times X \rightarrow X$ without any rhyme or reason as to how the map combines two elements to get a third.

You have already dealt with binary operators many times before in your mathematics courses, but you may not have thought of them as

maps in this form. For example, consider the map $\mathbb{Z} \times \mathbb{Z} \rightarrow \mathbb{Z}$ given by $(x, y) \mapsto x + y$. This is of course just addition of integers, simply expressed as a map. As another familiar example, consider the map $\mathbb{Q} \times \mathbb{Q} \rightarrow \mathbb{Q}$ given by $(x, y) \mapsto xy$; this is just the usual multiplication of rational numbers that you know and love, simply expressed as a map.

Before we see any more examples of binary operators, it's worth pointing out that binary operators often have a special symbol associated with them that is used in *infix notation*. That is, suppose we have a map $f : X \times X \rightarrow X$, this gives us some binary operator. Instead of writing $f(x, y)$, people will often use notation such as $x f y$ as a short-hand for $f(x, y)$, using the name of the function inbetween its two arguments, similar to how we use $+$ inbetween two numbers such as $2 + 3$.

In these notes we will use \star as our generic binary operator, unless there is some other commonly used notation for that operation. (You will often see other textbooks use the dot \cdot or simply use juxtaposition to indicate a binary operator is used.)

For example, we might define a map $\star : \mathbb{Z} \times \mathbb{Z} \rightarrow \mathbb{Z}$ by $x \star y = 3x - 2y$. Applying this to the integers $x = 7$ and $y = 4$, for example, we would have $7 \star 4 = 3 \cdot 7 - 2 \cdot 4 = 13$, $6 \star 1 = 16$, and $-3 \star 5 = -19$.

Exercise 6.1.

Does division form a binary operator on the set of rational numbers, \mathbb{Q} ? If not, can we restrict to a subset of \mathbb{Q} to get a binary operator? Is there any subset of the integers, $X \subseteq \mathbb{Z}$, so that division forms a binary operator on X ?

6.1.1 Familiar examples

We already mentioned a few familiar examples of binary operators on \mathbb{Z} and \mathbb{Q} , but let's mention a few more familiar examples you're familiar with.

Addition, Subtraction, and Multiplication on \mathbb{R} The familiar arithmetic operations of addition, subtraction, and multiplication on the real numbers are also binary operators: they take two given real numbers x and y and produce a third real number, $x + y$, $x - y$, or xy .

Division on $\mathbb{R} \setminus \{0\}$ Division on the reals *does not* define a binary operation on all of the reals since division by zero is undefined. However, if we restrict from the reals to the subset of non-zero reals, $\mathbb{R} \setminus \{0\}$, then we do have a binary operator: dividing one non-zero real number by another non-zero real number will produce a non-zero real number.

Exponentiation on the natural integers On the set of natural numbers \mathbb{N} we can define an operation by sending the pair of natural numbers (x, y) to the number x^y . Since x and y these are both positive integers, the result of x^y will also be a positive integer.

Unions and intersections Suppose X is any fixed set, and consider its powerset $\mathcal{P}(X)$, the set of all subsets of X . On the powerset there are two familiar binary operations we have previously seen: unions and intersections. Given any two subsets of X , $A, B \subseteq X$, their union $A \cup B$ and their intersection $A \cap B$ are both subsets of X , and so elements of $\mathcal{P}(X)$.

Exercise 6.2.

Is exponentiation, $(x, y) \mapsto x^y$, a binary operation on the set of all integers, \mathbb{Z} ?

6.1.2 Composition of maps $X \rightarrow X$

As another example of a binary operation that might feel a little less familiar, we could consider composition of functions. That is, suppose we have any fixed set X and we let f and g both be functions with domain and codomain X . That is, $f : X \rightarrow X$ and $g : X \rightarrow X$ are two maps. We can construct a new map, also from X to X , by considering the composition $g \circ f$.

Just to have a concrete example, suppose X is the set $\{1, 2, 3, 4\}$ and consider the maps f and g indicated below:

$$\begin{array}{ll} f(1) = 1 & g(1) = 2 \\ f(2) = 4 & g(2) = 2 \\ f(3) = 3 & g(3) = 4 \\ f(4) = 2 & g(4) = 4 \end{array}$$

The composition $g \circ f$ gives us a new map from the set $\{1, 2, 3, 4\}$ to itself, which we can easily determine:

$$\begin{array}{l} (g \circ f)(1) = g(f(1)) = g(1) = 2 \\ (g \circ f)(2) = g(f(2)) = g(4) = 4 \\ (g \circ f)(3) = g(f(3)) = g(3) = 4 \\ (g \circ f)(4) = g(f(4)) = g(2) = 2 \end{array}$$

6.1.3 Addition and multiplication of 2×2 matrices

Two more extremely important examples of binary operators are addition and multiplication of matrices. In order to make the discussion as simple as possible, we will only consider addition and multiplication of 2×2 matrices, but these ideas can be defined for matrices of other sizes as well.

By a 2×2 *real matrix*, we simply mean a collection of four real numbers listed as a table with two rows and two columns. For example,

$$\begin{pmatrix} 1 & 0 \\ 3 & \pi \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$$

are two such matrices. The collection of all such two by two matrices is denoted $M_2(\mathbb{R})$:

$$M_2(\mathbb{R}) = \left\{ \begin{pmatrix} a & b \\ c & d \end{pmatrix} \mid a, b, c, d \in \mathbb{R} \right\}.$$

We define the sum of two such matrices “component-wise,” simply meaning that add the corresponding entries of the two matrices (e.g., the entries in the first row, first column; or the entries in the second row, first column; etc.) to produce a new matrix with these numbers in the corresponding positions. That is, we define the sum of two matrices as

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} + \begin{pmatrix} e & f \\ g & h \end{pmatrix} = \begin{pmatrix} a + e & b + f \\ c + g & d + h \end{pmatrix}$$

This is a nice binary operation on $M_2(\mathbb{R})$: we take two 2×2 matrices, and combine them together to get a third 2×2 matrix.

The other operation we want to define gives us a way of multiplying two matrices. The first time you see this operation it seems strange and complicated, but there is an easy way to remember it that we will describe shortly.

We define the product of two 2×2 matrices to be the following:

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \cdot \begin{pmatrix} e & f \\ g & h \end{pmatrix} = \begin{pmatrix} ae + bg & af + bh \\ ce + dg & cf + dh \end{pmatrix}$$

For example,

$$\begin{pmatrix} 1 & 2 \\ 0 & 3 \end{pmatrix} \cdot \begin{pmatrix} 0 & 1 \\ -1 & 2 \end{pmatrix} = \begin{pmatrix} 1 \cdot 0 + 2 \cdot (-1) & 1 \cdot 1 + 2 \cdot 2 \\ 0 \cdot 0 + 3 \cdot (-1) & 0 \cdot 1 + 3 \cdot 2 \end{pmatrix} = \begin{pmatrix} -2 & 5 \\ -3 & 6 \end{pmatrix}.$$

One way that people often remember how to multiply matrices is by multiplying entries in from a row of the matrix on the left with entries from a column of the matrix on the right, and adding the result together. For example, when multiplying

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \cdot \begin{pmatrix} e & f \\ g & h \end{pmatrix}$$

we consider the first row of the left-hand matrix (with entries a and b), and the first column of the right-hand matrix (with entries e and g), multiply the corresponding elements of these rows and columns (giving ae and bg), and then adding the result together (getting $ae + bg$). This corresponds to the entry in the first row, first column of the resulting product matrix.

To get the entry in the first row and second column of the product, we consider the first row of the left-hand matrix (entries a and b), the second column of the right-hand matrix (entries f and h), multiply corresponding entries (af and bh), add the results together ($af + bh$), and places this in the first row, second column of the product. This pattern continues for the other entries in the product: the entry in the i -th row and j -th column comes from multiplying and adding the entries of the i -th row of the left-hand matrix and the j -th column of the right-hand matrix.

Though this operation seems weird and unwieldy when you first encounter, it is extremely important for more advanced mathematics courses, both in pure mathematics and in applications.

Remark.

We won't take the time to go through the details, but matrix multiplication is defined this way explicitly so that it corresponds to composition of certain types of functions. If you take linear algebra, you will learn there are special types of objects called "vectors," and collections of vectors form "vector spaces." Matrices give a way of representing maps with nice properties between vector spaces, called "linear transformations," and multiplication of matrices really corresponds to composition of these linear transformations.

6.2 Properties binary operators may have

A "binary operation" in itself is often too vague to be useful: there are just too many possibilities for what a map $X \times X \rightarrow X$ could possibly be for us to say much about binary operators in general. If, however, we restrict ourselves to binary operators that have some desirable properties, it will turn out we can actually start proving interesting theorems.

Remark.

We won't take the time in this course to say much about these various "special" binary operators with desirable properties, but this is essentially the beginning of a branch of mathematics called *abstract algebra*. As you will see if you take abstract algebra, we can actually prove quite a lot of interesting things if we assume the binary operator has enough properties.

6.2.1 Associativity

A binary operator \star on a set X is said to be *associative* if for all $x, y, z \in X$ we have $(x \star y) \star z = x \star (y \star z)$. Intuitively, being associative means that if we string binary operators together to perform the operation on several elements two at a time, the way we group the elements together does not change the result.

Most of the binary operators we have introduced will be associative, though binary operators in general need not be. The usual arithmetic operations of addition and multiplication on \mathbb{Z} , \mathbb{Q} , and \mathbb{R} will all be associative, but subtraction is not associative. This is easy to verify with a counterexample. Consider the case when $x = 2$, $y = 3$, and $z = 4$, then

$$(x - y) - z = (2 - 3) - 4 = -1 - 4 = -5$$

however

$$x - (y - z) = 2 - (3 - 4) = 2 - (-1) = 3,$$

and so subtraction is not associative. Notice that since $2, 3, 4 \in \mathbb{Z}$ and $\mathbb{Z} \subseteq \mathbb{Q} \subseteq \mathbb{R}$, this counterexample shows that subtraction is not associative in \mathbb{Z} , \mathbb{Q} , or \mathbb{R} .

One particularly important example of an associative binary operation is function composition, so we will state this as a theorem together with a proof.

Theorem 6.1. *Function composition is associative. In particular, if f , g , and h are functions from a set X to itself, then $(f \circ g) \circ h = f \circ (g \circ h)$.*

Before giving the proof of Theorem 6.1, let's walk through the idea. First notice that we're trying to show that two maps are the same. Recall that if two maps are equal, then they must map elements from the domain to the codomain in the same way. That is, if φ and ψ were both maps from a set A to a set B , then $\varphi = \psi$ means that for every $a \in A$ we must have $\varphi(a) = \psi(a)$. In our situation, this means we want to show that $((f \circ g) \circ h)(x)$ equals $(f \circ (g \circ h))(x)$ for every $x \in X$. We can do this by simply "unwinding" the definition of $f \circ g$ (and $g \circ h$, etc.) for both compositions, and see if we get the same element. We'll do this by having an arbitrary element x in X , and verifying that after unwinding the definition, both maps send x to the same place. Since our x was arbitrarily chosen, this means this argument applies to all choices of $x \in X$, and so the maps must be equal.

Proof of Theorem 6.1. Let $x \in X$ be arbitrary. Notice that applying $(f \circ g) \circ h$ to x yields

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

Similarly, applying $f \circ (g \circ h)$ to x gives

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

In both cases we have that the maps send x to $f(g(h(x)))$. This holds for all x since $x \in X$ was arbitrary, and so the maps are equal. \square

We can also show that matrix multiplication is associative by directly verifying that the products

$$\left(\begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} e & f \\ g & h \end{pmatrix} \right) \begin{pmatrix} i & j \\ k & \ell \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} a & b \\ c & d \end{pmatrix} \left(\begin{pmatrix} e & f \\ g & h \end{pmatrix} \begin{pmatrix} i & j \\ k & \ell \end{pmatrix} \right)$$

are equal. This is not difficult, but is tedious (and a good exercise in matrix multiplication), so we will leave it as an exercise.

Exercise 6.3.

Verify that multiplication of 2×2 matrices is associative.

Remark.

If we're willing to believe our earlier claim that matrix multiplication really corresponds to composition of some special types of maps, then we can sidestep the tedious computation of Exercise 6.3 by appealing to Theorem 6.1.

6.2.2 Commutativity

We say that a binary operator \star on a set X is *commutative* if for all $x, y \in X$ we have that $x \star y = y \star x$. Multiplication of integers/rationals/reals is commutative, but division (when defined) is not. For example, if we divide the rational numbers 1 and 3, we could get either $3/1 = 3$ or $1/3$ depending on the order in which we perform the division, and so division is not commutative.

Function composition is also typically not commutative. For example, if X is the set $\{1, 2, 3\}$ and f and g are the maps from X to itself given below,

$$\begin{array}{ll} f(1) = 1 & g(1) = 2 \\ f(2) = 1 & g(2) = 3 \\ f(3) = 2 & g(3) = 1 \end{array}$$

then $g \circ f$ applied to 1 would give us $g(f(1)) = g(2) = 3$, whereas $f \circ g$ applied to 1 gives $f(g(1)) = f(2) = 1$. Hence $g \circ f \neq f \circ g$. So we see that function composition is not going to be commutative in general.

Exercise 6.4.

Are there any restrictions on the set X that would make function composition of maps from X to itself commutative?

Matrix multiplication we can also verify is not commutative by giving a counterexample. Notice

$$\begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 0 & 3 \\ 4 & 0 \end{pmatrix} = \begin{pmatrix} 8 & 3 \\ 4 & 0 \end{pmatrix}$$

but if we reverse the order of the matrices we're multiplying above, we would have

$$\begin{pmatrix} 0 & 3 \\ 4 & 0 \end{pmatrix} \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 0 & 4 \\ 3 & 8 \end{pmatrix}.$$

Thus matrix multiplication is not commutative.

Exercise 6.5.

Consider the binary operator on the integers given by $x \star y = x^2 + xy + y^2$. For example, $3 \star 5 = 3^2 + 3 \cdot 5 + 5^2 = 49$. Show that this operation is commutative.

6.2.3 Identity

We had seen before that every set X has a special map from X to itself called the "identity map," id_X defined simply by $x \mapsto x$. When we mentioned the identity map before, we had stated that it had the special property that for *every* map $f : X \rightarrow X$, the compositions $f \circ \text{id}_X$ and $\text{id}_X \circ f$ simply gave us back f . That is, composing a map with the identity map doesn't change the map. Something similar happens if we

consider the binary operation of addition on the integers (or rationals or reals): adding zero to a number does not change the number. Similarly, multiplying an integer (or rational or real number) by 1 does not change anything. In each of these cases there's some special element of our set which does not affect other elements when we perform a binary operation. These are all examples of "identities" for the operation.

To be precise, if we have a binary operator \star on a set X , we say that an element $e \in X$ is an identity for \star if for every element $x \in X$ we have $e \star x = x = x \star e$.

Remark.

The letter e is often used for the identity element of a binary operator, and should not be confused with Euler's constant. The use of e seems to go back to an article by Heinrich Burkhardt from 1899, where he referred to the identity element by the German word *Einheitselement* and so used e to represent this special element.

We can verify with a little bit of computation that the follow matrix, called the identity matrix, is an identity for matrix multiplication:

$$\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

By simply multiplying this special matrix on both the left and the right of an arbitrary 2×2 matrix, we will see that it does not affect the matrix:

$$\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} a & b \\ c & d \end{pmatrix} = \begin{pmatrix} 1 \cdot a + 0 \cdot c & 1 \cdot b + 0 \cdot d \\ 0 \cdot a + 1 \cdot c & 0 \cdot b + 1 \cdot d \end{pmatrix} = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} a \cdot 1 + c \cdot 0 & b \cdot 1 + d \cdot 0 \\ a \cdot 0 + c \cdot 1 & b \cdot 0 + d \cdot 1 \end{pmatrix} = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

It's worth noting that not every binary operator necessarily has an identity. As a silly example, consider the binary operator of addition defined on the natural numbers (positive integers). This is a perfectly legitimate operation: if you add two natural numbers you get a natural number. However, there is no natural number that could play the role of the identity since any two non-zero positive integers added together will yield an integer that is larger than either of the terms. Since we have explicitly removed 0 from consideration, we don't have an identity.

Exercise 6.6.

Does the operator \star on \mathbb{Z} we defined earlier as $x \star y = x^2 + xy + y^2$ have an identity? That is, is there any integer e so that $e \star x = x = x \star e$ for all integers x ?

As we've defined it, we've left open the possibility that there could be multiple identity elements, but is this actually possible? It turns out that if a binary operator has an identity, the identity must be unique: if e_1 and e_2 were both identities, then e_1 and e_2 are actually the same element.

Exercise 6.7.

Show that if a binary operator has an identity, the identity is unique.

6.2.4 Inverses

Notice that for every non-zero rational number $x = \frac{p}{q}$ there exists a non-zero rational number $y = \frac{q}{p}$ so that $xy = 1$: the product of the two numbers is the identity element. Similarly, for every integer x there is a corresponding integer $-x$ so that their sum $x + (-x)$ is 0, the identity element for addition of integers.

In general, given a binary operator \star on a set X with an identity e , we say that an element $y \in X$ is an *inverse* of an element $x \in X$ if $x \star y = e = y \star x$. That is, if when we combine the two elements together with the binary operation, we get back the identity element.

Notice that even when our binary operator has an identity, it could be the case that some elements have inverses and some do not. For example, multiplication is a binary operator on all the rational numbers, but it's precisely the non-zero elements that have a multiplicative inverse: There is no rational number y so that $0 \cdot y = 1$, since multiplication by zero always gives zero.

As another example, consider function composition. As we mentioned when we were defining maps, a function $f : X \rightarrow X$ need not have an inverse. In fact, an inverse exists precisely when f is bijective; if f is not bijective, f does not have an inverse.

As for matrices, some matrices have inverses, but some do not. For example, we can check by direct computation that the matrices

$$\begin{pmatrix} 1 & 2 \\ 3 & 1 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} -1/5 & 2/5 \\ 3/5 & -1/5 \end{pmatrix}$$

are inverses:

$$\begin{pmatrix} 1 & 2 \\ 3 & 1 \end{pmatrix} \begin{pmatrix} -1/5 & 2/5 \\ 3/5 & -1/5 \end{pmatrix} = \begin{pmatrix} 1 \cdot (-1/5) + 2 \cdot 3/5 & 1 \cdot 2/5 + 2 \cdot (-1/5) \\ 3 \cdot (-1/5) + 1 \cdot 3/5 & 3 \cdot 2/5 + 1 \cdot (-1/5) \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

(The computation with the order of the product reversed is easily checked.)

The matrix $\begin{pmatrix} 1 & 2 \\ 0 & 0 \end{pmatrix}$ does not have an inverse, however. To see this, suppose for the sake of contradiction that there was some inverse matrix, $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$. This would mean

$$\begin{pmatrix} 1 & 2 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} a & b \\ c & d \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

If we actually multiply out the matrices on the left, however, we are left with

$$\begin{pmatrix} a + 2c & a + 2d \\ 0 & 0 \end{pmatrix}$$

Notice we must have 0 in the lower right-hand corner, but this precludes us from finding any choice of a , b , c , and d so that the product is the identity. Hence this matrix can not have an inverse.

It can be shown that a 2×2 matrix $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$ will have an inverse precisely when $ad - bc \neq 0$, and in this case the inverse matrix is

$$\begin{pmatrix} \frac{d}{ad-bc} & \frac{-b}{ad-bc} \\ \frac{-c}{ad-bc} & \frac{a}{ad-bc} \end{pmatrix}.$$

It is easy to check that this does indeed form an inverse of the original matrix:

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} \frac{d}{ad-bc} & \frac{-b}{ad-bc} \\ \frac{-c}{ad-bc} & \frac{a}{ad-bc} \end{pmatrix} = \begin{pmatrix} \frac{ad}{ad-bc} + \frac{-bc}{ad-bc} & \frac{-ab}{ad-bc} + \frac{ab}{ad-bc} \\ \frac{cd}{ad-bc} + \frac{-cd}{ad-bc} & \frac{-bc}{ad-bc} + \frac{ad}{ad-bc} \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

Combining this with the result of the exercise below shows that this is *the* inverse of our matrix.

Exercise 6.8.

Suppose \star is an associative binary operator with identity e on a set X , and suppose that $x \in X$ has inverse y : i.e., $x \star y = e = y \star x$. Show that the inverse is unique. That is, show that if there existed a $y' \in X$ so that $x \star y' = e = y' \star x$, then $y = y'$.

It is important to note that the associativeness above is required for us to guarantee that inverses are unique. It is possible to construct examples of non-associative binary operators where the inverse is not unique. For instance, if our set X is $\{a, b, c, e\}$ and we define an operator \star by

$$\begin{array}{llll} a \star a = a & a \star b = e & a \star c = e & a \star e = a \\ b \star a = e & b \star b = b & b \star c = b & b \star e = b \\ c \star a = e & c \star b = c & c \star c = b & c \star e = c \\ e \star a = a & e \star b = b & e \star c = c & e \star e = e \end{array}$$

then this will be a binary operator where e is the identity, and both b and c are inverses of a . Notice this operator will not be associative, though, as

$$(a \star b) \star c = e \star c = c \quad \text{yet} \quad a \star (b \star c) = a \star b = e.$$

In the special case of an associative operator, the exercise above shows us that inverses are unique, and so often adopt the notation that x^{-1} represents the inverse of x , if it exists.

6.2.5 Groups

It is very, very common for binary operators to be associative, have an identity, and for *every* element of the set to have an inverse. This situation is so common that we give sets with such binary operators a special name: we call them a *group*. That is, a **group** is a set X together with a binary operator \star that is associative, has an identity, and where every element has an inverse. We've actually already seen a few examples of groups: the integers together with addition form a group; the non-zero rationals with multiplication form a group; the set of all bijective maps from a set X to itself form a group. The collection of all matrices does not form a group since not every element has an inverse. We've seen above, though, that a 2×2 real matrix will have an inverse provided $ad - bc \neq 0$.

If we restrict ourselves to the set of matrices where this condition is satisfied, then we will have a group which is denoted $GL_2(\mathbb{R})$:

$$GL_2(\mathbb{R}) = \left\{ \begin{pmatrix} a & b \\ c & d \end{pmatrix} \mid a, b, c, d \in \mathbb{R} \text{ and } ad - bc \neq 0 \right\}.$$

We won't say anything more about groups for now, but if you were to take abstract algebra, learning about groups would be the starting point of that class.

6.3 Arithmetic of congruence classes

We had previously defined an equivalence relation \equiv_m on the set of integers \mathbb{Z} by declaring $x \sim y$ if $m \mid (x - y)$. The collection of equivalence classes under this equivalence relation is often denoted \mathbb{Z}_m . Recall that given an $x \in \mathbb{Z}$, we use the notation $[x]$ to represent its equivalence class, the set of all integers which are equivalent to $[x]$.

For example, if $m = 5$, then the set of equivalence classes \mathbb{Z}_5 is the following

$$\mathbb{Z}_5 = \{[0], [1], [2], [3], [4]\}$$

where

$$\begin{aligned} [0] &= \{\dots, -15, -10, -5, 0, 5, 10, 15, \dots\}, \\ [1] &= \{\dots, -14, -9, -4, 1, 6, 11, 16, \dots\}, \\ [2] &= \{\dots, -13, -8, -3, 2, 7, 12, 17, \dots\}, \\ [3] &= \{\dots, -12, -7, -2, 3, 8, 13, 18, \dots\}, \text{ and} \\ [4] &= \{\dots, -11, -6, -1, 4, 9, 14, 19, \dots\}. \end{aligned}$$

Notice that when we write $[0]$ or $[2]$ for an equivalence class, we're making a choice of which element in the class to use to represent the class. If we replace that element with any equivalence element, we still have the same equivalence class, and so we could reasonably refer to \mathbb{Z}_5 as

$$\mathbb{Z}_5 = \{[5], [-4], [7], [13], [-11]\}$$

since $[0] = [5]$, $[1] = [-4]$, $[2] = [7]$, $[3] = [13]$, and $[-11] = [4]$. However, usually when people discuss these equivalence classes (also called ***congruence classes***), they will use the smallest non-negative elements of the equivalence class as the representative.

6.3.1 Defining addition and multiplication

We can define binary operators of addition and multiplication on \mathbb{Z}_m by simply adding or multiplying the representatives, and then taking the equivalence class of the resulting sum or product. For example, in \mathbb{Z}_5 we could compute $[3] + [4]$ by adding $3 + 4$ to get 7, and then taking the equivalence class $[7]$ as the result. Since $[7] = [2]$, we might reasonably write $[3] + [4] = [2]$. Similarly, the product $[2] \cdot [3]$ would be given by the equivalence class represented by $2 \cdot 3 = 6$. Since $[6] = [1]$, we have $[2] \cdot [3] = [1]$.

As another example, consider addition and multiplication in \mathbb{Z}_{12} . Here $[8] + [6]$, for example, would give us $[2]$ since $8 + 6 = 14$ and $12 \mid (14 - 2)$. Notice also that something a bit odd happens with multiplication: it is possible to multiply two non-zero congruence classes together and get zero, something which does not happen with integers, rationals, or reals. For instance, $[8] \cdot [3] = [8 \cdot 3] = [24] = [0]$.

For small values of m , we can represent all the possibilities of adding and multiplying congruences in \mathbb{Z}_m by creating a *Cayley table*. These are $m \times m$ tables (one for addition, and one for multiplication) where we have a row and a column per congruence class, and then fill in the table with the values of $[r][c]$ where $[r]$ corresponds to the congruence class for a row, and $[c]$ is the congruence class for a column. For example, in \mathbb{Z}_6 the table for addition is

+	[0]	[1]	[2]	[3]	[4]	[5]
[0]	[0]	[1]	[2]	[3]	[4]	[5]
[1]	[1]	[2]	[3]	[4]	[5]	[0]
[2]	[2]	[3]	[4]	[5]	[0]	[1]
[3]	[3]	[4]	[5]	[0]	[1]	[2]
[4]	[4]	[5]	[0]	[1]	[2]	[3]
[5]	[5]	[0]	[1]	[2]	[3]	[4]

Whereas the table for multiplication in \mathbb{Z}_6 is

+	[0]	[1]	[2]	[3]	[4]	[5]
[0]	[0]	[0]	[0]	[0]	[0]	[0]
[1]	[0]	[1]	[2]	[3]	[4]	[5]
[2]	[0]	[2]	[4]	[0]	[2]	[4]
[3]	[0]	[3]	[0]	[3]	[0]	[3]
[4]	[0]	[4]	[2]	[0]	[4]	[2]
[5]	[0]	[5]	[4]	[3]	[2]	[1]

Exercise 6.9.

Construct the Cayley tables for multiplication and addition in \mathbb{Z}_7 and in \mathbb{Z}_8 .

6.3.2 Addition and multiplication are well-defined

Before we go too far in our study of arithmetic with congruence classes (often also referred to as *modular arithmetic*), we should point out there is one conceivable problem with the binary operators we defined above: are they actually well-defined? That is, if we chose different representatives of our equivalence classes, would we still get the same result? For instance, $[3] = [13]$ and $[4] = [-11]$, but is it true that $[3] + [4]$ will be the same as $[13] + [-11]$? Here it's easy to see these will be equal since we computed $[3] + [4] = [2]$ and since $13 + (-11) = 2$, but will this always happen? If our operator is to be "well-defined," we need to know that we will always compute quantities in the same, consistent way, regardless of the representatives of the equivalence classes that we choose to work with.

The following theorem tells us that these operations we want to define will actually be well-defined: they result of multiplying or adding congruence classes will be independent of the representatives we choose to work with.

Theorem 6.2. *Suppose $x \equiv x' \pmod{m}$ and $y \equiv y' \pmod{m}$. (That is, $[x] = [x']$ and $[y] = [y']$ in \mathbb{Z}_m .) Then $x + y \equiv x' + y' \pmod{m}$ and $xy \equiv x'y' \pmod{m}$. (Or, in terms of our congruence classes, $[x] + [y] = [x'] + [y']$ and $[x][y] = [x'][y']$.)*

Proof. Since $x \equiv x' \pmod{m}$ and $y \equiv y' \pmod{m}$, there are integers a and b so that $x - x' = am$ and $y - y' = bm$. Thus $x' = am + x$ and $y' = bm + y$. Now consider $x' + y'$. We may write this as

$$x' + y' = am + x + bm + y = (a + b)m + x + y$$

or

$$(x' + y') - (x + y) = (a + b)m$$

so $m \mid ((x' + y') - (x + y))$, and $x' + y' \equiv x + y \pmod{m}$.

Similarly,

$$x'y' = (am+x)(bm+y) = abm^2 + amy + bmx + xy = (abm + ax + by)m + xy$$

and so $x'y' \equiv xy \pmod{m}$. \square

Theorem 6.2 is exactly what we need to justify that the addition and multiplication of congruence classes defined above is independent of our choices of representatives of the classes involved in doing the computation:

Corollary 6.3 (Corollary of Theorem 6.2). *If $[x] = [x']$ and $[y] = [y']$ in \mathbb{Z}_m , then $[x] + [y] = [x'] + [y']$ and $[x][y] = [x'][y']$.*

Proof. To show that $[x] + [y]$ equals $[x'] + [y']$, notice that

$$[x] + [y] = [x + y] \quad \text{and} \quad [x'] + [y'] = [x' + y'].$$

Theorem 6.2 tells us that $x + y \equiv x' + y' \pmod{m}$, and so $[x + y] = [x' + y']$. The proof for multiplication is similar. \square

Exercise 6.10.

Show that for each positive integer m , \mathbb{Z}_m forms a group under addition. That is, show that addition in \mathbb{Z}_m is associative, has an identity, and each element has an additive inverse. Show that addition in \mathbb{Z}_m is commutative.

Exercise 6.11.

Show that multiplication in \mathbb{Z}_m is associative, commutative, and there exists an identity element. Does \mathbb{Z}_m necessarily form a group under multiplication? What if we consider $\mathbb{Z}_m \setminus \{0\}$? Are there any restrictions on m so that $\mathbb{Z}_m \setminus \{0\}$ forms a group under multiplication?

6.3.3 Distribution

Now that we have two binary operations on \mathbb{Z}_m , we may be interested in how they interact with one another. That is, can we say anything about expressions in \mathbb{Z}_m which involve both multiplication and addition? Do nice properties such as distribution hold? To be more precise, can we distribute multiplication across addition, just as we do with normal integers? I.e., is it true that $[x]([y] + [z]) = [x][y] + [x][z]$ for all $[x], [y], [z] \in \mathbb{Z}_m$?

You shouldn't be too surprised that we can distribute, but let's give a quick proof.

Theorem 6.4. *Multiplication distributes across addition in \mathbb{Z}_m . That is, for all $[x], [y]$, and $[z]$ in \mathbb{Z}_m we have*

$$\begin{aligned} [x]([y] + [z]) &= [x][y] + [x][z], \text{ and} \\ ([x] + [y])[z] &= [x][z] + [y][z]. \end{aligned}$$

Proof. We simply note that since multiplication of regular integers distributes, we can apply the distributive property to representatives of our congruence classes, and then take the corresponding congruence class. That is,

$$[x]([y] + [z]) = [x][y + z] = [x(y + z)] = [xy + xz] = [xy] + [xz] = [x][y] + [x][z].$$

The other distributive rule is proved similarly. \square

Odds and Ends

7.1 The Pigeonhole Principle

The “pigeonhole principle” is a simple idea that we can sometimes use to prove some surprisingly counterintuitive facts. The idea behind the pigeonhole principle is the following: if n pigeons are distributed amongst $m < n$ nests, then at least one nest must have two pigeons. For example, if there are four nests and five pigeons, then at least one nest has two pigeons. It could be that one nest has two pigeons and the others have exactly one, or it could be that one nest has all five pigeons and the other nests are empty, or some other configuration of pigeons in their nests, but at least one nest must have at least two pigeons.

A more mathematical-sounding version of the pigeonhole principle would be to say that there are no injective maps from a set X of cardinality n to a set Y of cardinality $m < n$.

As simple as the pigeonhole principle sounds, it can be surprisingly useful as the next few examples show.

Example 7.1.

Suppose $n \geq 2$ people are at a party. At least two people at the party must know the same number of people attending the party.

Before justifying why this is true, let’s be sure we understand what the claim is. We have a collection of people at a party, and perhaps some people at the party know several attendees, some people know very few, maybe someone knows everyone, maybe someone knows no one. For each person attending the party we assign a number, which is how many people at the party they know. If you know six people at the party, your number is six; if you know no one at the party, your number is zero; if you know twelve people at the party, your number is twelve. We claim there must be two people with the same number. Or, put another way, it’s impossible that each person at the party knows a different number of party attendees.

To justify our claim using the pigeonhole principle, notice that the number we assign to each person is between 0 and $n - 1$. Yet

there are n numbers to assign, so at least one number must be assigned to at least two people.

Remark.

A more mathematical sounding version of the example above is to say that every graph has at least two vertices of the same degree. We won't take the time to define what those terms are, but you will see them again when you take a course in discrete structures and learn about graph theory.

Example 7.2.

Any subset of eleven numbers from $\{1, 2, 3, \dots, 19\}$ will contain two numbers whose sum is 20. That is, if you choose eleven numbers from 1 through 19, you will necessarily have two that add to twenty.

To justify this, we imagine there being ten "boxes" labelled $(1, 19)$, $(2, 18)$, $(3, 17)$, $(4, 16)$, $(5, 15)$, $(6, 14)$, $(7, 13)$, $(8, 12)$, $(9, 11)$, and (10) . Now each element of our chosen set of eleven elements in $\{1, \dots, 19\}$ is placed in a corresponding box. Since there are eleven elements and ten boxes, at least one box must contain two elements, which means we will have two numbers that add to twenty.

Example 7.3.

If ten points are randomly placed in a 3×3 square, there are at least two points that are within distance $\sqrt{2}$ from one another.

Proof. Imagine the 3×3 square being made of nine 1×1 squares, and our random points are thus placed in one of these nine "subsquares." Since

we have ten points and only nine squares though, at least one square must contain at least two points. But the furthest apart two points in a 1×1 square can be is $\sqrt{2}$, corresponding to opposite corners of the square. \square

Index

- 2×2 real matrix, 123
- n -tuple, 79
- (Cartesian) product, 76
- 1-1, 94

- antisymmetric, 109
- associative, 125
- associative property of composition, 101

- base case, 30
- bijjective, 95
- binary operator, 120
- bound variable, 51

- cardinality, 63
- Cayley table, 134
- codomain, 90
- Collatz conjecture, 17
- commutative, 127
- composite number, 41
- composition, 99
- congruence classes, 118, 133
- congruence modulo m , 108
- congruent modulo 2, 108
- conjecture, 16
- conjunction, 52
- contrapositive, 58
- converse, 57
- corollary, 16
- counterexample, 11

- de Morgan's laws, 84
- differentiable, 12
- differentiable at a , 12
- disjoint, 75
- disjoint union, 117
- disjunction, 52
- divides, 21
- domain, 90

- equal, 68
- equivalence class, 116
- equivalence relation, 114
- equivalent, 114
- even number, 9
- existential quantifier, 50

- Fibonacci numbers, 38
- free variable, 51
- function, 89
- fundamental theorem of arithmetic, 42

- graph, 92, 105
- group, 132

- homogeneous coordinates, 118

- identity, 129
- identity map, 103
- identity matrix, 129
- if and only if, 58
- iff, 58
- image, 98
- Implication, 55
- index set, 87
- inductive hypothesis, 30
- inductive step, 28, 30
- injective, 94
- integers, 63
- intersection, 72
- inverse, 97, 130
- inverse map, 102
- irrational, 46
- irrational number, 71

- lemma, 15
- linear order, 111
- logical and, 52

- map, 89

- modular arithmetic, 135
- multiple, 21
- natural numbers, 62
- necessary condition, 56
- negative, 53
- one-to-one, 94
- onto, 93
- open sentences, 50
- parity, 107
- partial order, 112
- power set, 87
- predicates, 50
- preimage, 98
- prime factorization, 42
- prime number, 41
- proper subset, 67
- proposition, 15
- propositions, 49
- quantifier, 50
- quicksort algorithm, 25
- range, 90
- rational, 46
- rational numbers, 65
- real number, 70
- real projective line, 118
- recurrence relation, 38
- reflexive, 110
- relation, 105
- relation on X , 106
- set, 60
- set builder notation, 64
- strong induction, 38
- subset, 66
- sufficient condition, 56
- superset, 66
- surjective, 93
- symmetric, 109
- the empty set, 69
- theorem, 15
- total order, 111
- triangular numbers, 35
- truth table, 52
- union, 71
- universal quantifier, 50
- universe, 80
- variables, 50
- Venn diagrams, 66
- weak induction, 29