

A Gentle Introduction to the Art of
Mathematics

Version 3.1 N

Joe Fields

Southern Connecticut State University

Copyright © 2013 Joseph E. Fields. Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.3 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled “GNU Free Documentation License”.

The latest version of this book is available (without charge) in portable document format at

<http://www.southernct.edu/~fields/>

Acknowledgments

This is version 3.1 of *A Gentle Introduction to the Art of Mathematics*. Earlier versions were used and classroom tested by several colleagues: Robert Vaden-Goad, John Kavanagh, Ross Gingrich, Aaron Clark. I thank you all. A particular debt of gratitude is owed to Len Brin whose keen eyes caught a number of errors and inconsistencies, and who contributed many new exercises – thanks, Len.

A word about options

There are currently 4 optional versions of this text. The different options are distinguished with flags following the version number. Thus you may see a plain “version 3.1” (no flags) or any of the following: “version 3.1S”, “version 3.1N”, or “version 3.1SN”. The S flag is for those who prefer the symbol \sim to represent logical negation. (The default is to use the symbol \neg for logical negation.) The N flag distinguishes a version in which the convention that 1 is the smallest natural number is followed throughout. The absence of the N flag indicates the original version of the text in which the naturals do not contain 0 in chapters 1 through 4 and (after the Peano axioms are introduced) that convention is changed to the more modern rule.

Contents

1	Introduction and notation	1
1.1	Basic sets	1
1.2	Definitions: Prime numbers	12
1.3	More scary notation	20
1.4	Definitions of elementary number theory	23
1.4.1	Even and odd	23
1.4.2	Decimal and base- n notation	23
1.4.3	Divisibility	25
1.4.4	Floor and ceiling	26
1.4.5	Div and mod	27
1.4.6	Binomial coefficients	29
1.5	Some algorithms	36
1.6	Rational and irrational numbers	45
1.7	Relations	50
2	Logic and quantifiers	55
2.1	Predicates and Logical Connectives	55
2.2	Implication	68
2.3	Logical equivalences	74
2.4	Two-column proofs	88
2.5	Quantified statements	92
2.6	Deductive reasoning and Argument forms	101

2.7	Validity of arguments and common errors	110
3	Proof techniques I	119
3.1	Direct proofs of universal statements	119
3.2	More direct proofs	132
3.3	Contradiction and contraposition	137
3.4	Disproofs	143
3.5	By cases and By exhaustion	148
3.6	Existential statements	157
4	Sets	165
4.1	Basic notions of set theory	165
4.2	Containment	172
4.3	Set operations	177
4.4	Venn diagrams	188
4.5	Russell's Paradox	197
5	Proof techniques II — Induction	201
5.1	The principle of mathematical induction	201
5.2	Formulas for sums and products	211
5.3	Other proofs using PMI	222
5.4	The strong form of mathematical induction	230
6	Relations and functions	233
6.1	Relations	233
6.2	Properties of relations	243
6.3	Equivalence relations	251
6.4	Ordering relations	260
6.5	Functions	270
6.6	Special functions	283

7 Proof techniques III — Combinatorics	293
7.1 Counting	293
7.2 Parity and Counting arguments	309
7.3 The pigeonhole principle	323
7.4 The algebra of combinations	328
8 Cardinality	339
8.1 Equivalent sets	339
8.2 Examples of set equivalence	345
8.3 Cantor's theorem	356
8.4 Dominance	365
8.5 CH and GCH	374
9 Proof techniques IV — Magic	381
9.1 Morley's miracle	383
9.2 Five steps into the void	390
9.3 Monge's circle theorem	402
References	410
Index	410

List of Figures

1.1	The sieve of Eratosthenes.	15
1.2	Pascal's triangle.	30
1.3	A small example in pseudocode and as a flowchart	38
1.4	The division algorithm in flowchart form.	40
1.5	The Euclidean algorithm in flowchart form.	42
2.1	A schematic representation of a transistor.	60
2.2	Series connections implement <i>and</i>	61
2.3	Parallel connections implement <i>or</i>	61
2.4	Parenthesizations expressed as digital logic circuits.	63
2.5	Disjunctive normal form.	65
3.1	A four-color map.	149
3.2	Graph pebbling.	152
3.3	Graph pebbling move.	153
3.4	A \mathbb{Z} -module.	161
6.1	An example of a relation.	234
6.2	An example of the “divides” relation.	235
6.3	The graph of the “less than” relation.	239
6.4	The graph of the divisibility relation.	239
6.5	Some simple Hasse diagrams.	263
6.6	Hasse diagram for $(\mathcal{P}(\{1, 2, 3\}), \subseteq)$	264

6.7	Hasse diagram of divisors of 72.	265
6.8	The sets related to an arbitrary function.	271
7.1	Full houses in Yahtzee.	296
7.2	Königsberg, Prussia.	313
7.3	Königsberg, Prussia as a graph.	314
7.4	A desk with pigeonholes.	323
8.1	Cantor's snake.	349
8.2	Equivalent intervals.	352
8.3	An interval is equivalent to a semi-circle.	353
8.4	Binary representations in the unit interval.	358
8.5	Setup for proving the C-B-S theorem.	369
8.6	An A -stopper in the proof of C-B-S.	371
9.1	The setup for Morley's Miracle.	384
9.2	The first Morley triangle.	385
9.3	Conway's puzzle proof.	388
9.4	Scaling in Conway's puzzle proof.	388
9.5	An infinite army in the lower half-plane.	391
9.6	Moving one step into the void is trivial.	393
9.7	Moving two steps into the void is more difficult.	394
9.8	Moving three steps into the void takes 8 men.	395
9.9	The taxicab distance to $(0, 5)$	397
9.10	Finding r	398
9.11	Setup for Monge's circle theorem.	403
9.12	Example of Monge's circle theorem.	405
9.13	Four triangles bounded by 6 line segments	407

List of Tables

2.1	Converse, inverse and contrapositive.	71
2.2	Basic logical equivalences.	83
2.3	The rules of inference.	106
3.1	The definitions of elementary number theory restated.	124
4.1	Basic set theoretic equalities.	180
6.1	Properties of relations.	244

To the student

You are at the right place in your mathematical career to be reading this book if you liked Trigonometry and Calculus, were able to solve all the problems, but felt mildly annoyed with the text when it put in these verbose, incomprehensible things called “proofs.” Those things probably bugged you because a whole lot of verbiage (not to mention a sprinkling of epsilons and deltas) was wasted on showing that a thing was true, which was *obviously true!* Your physical intuition is sufficient to convince you that a statement like the Intermediate Value Theorem just *has* to be true – how can a function move from one value at a to a different value at b without passing through all the values in between?

Mathematicians discovered something fundamental hundreds of years before other scientists – physical intuition is worthless in certain extreme situations. Probably you’ve heard of some of the odd behavior of particles in Quantum Mechanics or General Relativity. Physicists have learned, the hard way, not to trust their intuitions. At least, not until those intuitions have been retrained to fit reality! Go back to your Calculus textbook and look up the Intermediate Value Theorem. You’ll probably be surprised to find that it doesn’t say anything about *all* functions, only those that are continuous. So what, you say, aren’t most functions continuous? Actually, the number of functions that aren’t continuous represents an infinity so huge that it outweighs the infinity of the real numbers!

The point of this book is to help you with the transition from doing math

at an elementary level (which is concerned mostly with solving problems) to doing math at an advanced level (which is much more concerned with axiomatic systems and proving statements within those systems).

As you begin your study of advanced mathematics, we hope you will keep the following themes in mind:

1. Mathematics is alive! Math is not just something to be studied from ancient tomes. A mathematician must have a sense of playfulness. One needs to “monkey around” with numbers and other mathematical structures, make discoveries and conjectures and uncover truths.
2. Math is not scary! There is an incredibly terse and compact language that is used in mathematics – on first sight it looks like hieroglyphics. That language is actually easy to master, and once mastered, the power that one gains by expressing ideas rigorously with those symbols is truly astonishing.
3. Good proofs are everything! No matter how important a fact one discovers, if others don’t become convinced of the truth of the statement it does not become a part of the edifice of human knowledge. It’s been said that a proof is simply an argument that convinces. In mathematics, one “convinces” by using one of a handful of argument forms and developing one’s argument in a clear, step-by-step fashion. Within those constraints there is actually quite a lot of room for individual style – there is no one right way to write a proof.
4. You have two cerebral hemispheres – use them both! In perhaps no other field is the left/right-brain dichotomy more evident than in math. Some believe that mathematical thought, deductive reasoning, is synonymous with left-brain function. In truth, doing mathematics is often a creative, organic, visual, right-brain sort of process – however, in

communicating one's results one must find that linear, deductive, step-by-step, left-brain argument. You must use your whole mind to master advanced mathematics.

Also, there are amusing quotations at the start of every chapter.

Preface: for Instructors

At many universities and colleges in the United States a course which provides a transition from lower-level mathematics courses to those in the major has been adopted. Some may find it hard to believe that a course like Calculus II is considered “lower-level” so let’s drop the pejoratives and say what’s really going on. Courses for Math majors, and especially those one takes in the Junior and Senior years, focus on proofs — students are expected to learn *why* a given statement is true, and be able to come up with their own convincing arguments concerning such “why”s. Mathematics courses that precede these typically focus on “how.” How does one find the minimum value a continuous function takes on an interval? How does one determine the arclength along some curve. Et cetera. The essential *raison d’etre* of this text and others like it is to ease this transition from “how” courses to “why” courses. In other words, our purpose is to help students develop a certain facility with mathematical proof.

It should be noted that helping people to become good proof writers – the primary focus of this text – is, very nearly, an impossible task. Indeed, it can be argued that the best way to learn to write proofs is by writing a lot of proofs. Devising many different proofs, and doing so in various settings, definitely develops the facility we hope to engender in a so-called “transitions” course. Perhaps the pedagogical pendulum will swing back to the previous tradition of essentially throwing students to the wolves. That is, students might be expected to learn the art of proof writing while actually writing

proofs in courses like algebra and analysis¹. Judging from the feedback I receive from students who have completed our transitions course at Southern Connecticut State University, I think such a return to the methods of the past is unlikely. The benefits of these transitions courses are enormous, and even though the curriculum for undergraduate Mathematics majors is an extremely full one, the place of a transition course is, I think, assured.

What precisely are the benefits of these transitions courses? One of my pet theories is that the process one goes through in learning to write and understand proofs represents a fundamental reorganization of the brain. The only evidence for this stance, albeit rather indirect, are the almost universal reports of “weird dreams” from students in these courses. Our minds evolved in a setting where inductive reasoning is not only acceptable, but advisable in coping with the world. Imagine some Cro Magnon child touching a burning branch and being burned by it. S/He quite reasonably draws the conclusion that s/he should not touch *any* burning branches, or indeed anything that is on fire. A Mathematician has to train him or herself to think strictly by the rules of deductive reasoning – the above experience would only provide the lesson that at that particular instant of time, that particular burning branch caused a sensation of pain. Ideally, no further conclusions would be drawn – obviously this is an untenable method of reasoning for an animal driven by the desire to survive to adulthood, but it is the *only* way to think in the artificial world of Mathematics.

While a gentle introduction to the art of reading and writing proofs is the primary focus of this text, there are other subsidiary goals for a transitions course that we hope to address. Principal among these is the need for an introduction to the “culture” of Mathematics. There is a shared mythos and language common to all Mathematicians – although there are certainly some distinct dialects! Another goal that is of extraordinary importance is

¹At the University of Maryland, Baltimore County, where I did my undergraduate work, these courses were actually known as the “proofs” courses.

impressing the budding young Mathematics student with the importance of play. My thesis adviser² used to be famous for saying “Well, I don’t know! Why don’t you monkey around with it a little . . .” In the course of monkeying around – doing small examples by hand, trying bigger examples with the aid of a computer, changing some element of the problem to see how it affected the answer, and various other activities that can best be described as “play,” eventually patterns emerged, conjectures made themselves apparent, and possible proof techniques suggested themselves. In this text there are a great many open-ended problems, some with associated hints as to how to proceed (which the wise student will avoid until hair-thinning becomes evident), whose point is to introduce students to this process of mathematical discovery.

To recap, the goals of this text are: an introduction to reading and writing mathematical proofs, an introduction to mathematical culture, and an introduction to the process of discovery in Mathematics. Two pedagogical principles have been of foremost importance in determining *how* this material is organized and presented. One is the so-called “rule of three” which is probably familiar to most educators. Propounded by (among others) Gleason, Hughes-Hallett, *et al.* in their reform Calculus it states that, when possible, information should be delivered via three distinct mechanisms – symbolically, graphically and numerically. The other is also a “rule of three” of sorts, it is captured by the old speechwriter’s maxim – “Tell ’em what you’re gonna tell ’em. Tell ’em. Then tell ’em what you told ’em.” Important and/or difficult topics are revisited at least three times in this book. In marked contrast to the norm in Mathematics, the first treatment of a topic is *not* rigorous, precise definitions are often withheld. The intent is to provide a bit of intuition regarding the subject material. Another reason for providing a crude introduction to a topic before giving rigorous detail revolves around the way human memory works. Unlike computer memory, which (excluding

²Dr. Vera Pless, to whom I am indebted in more ways than I can express.

the effects of the occasional cosmic ray) is essentially perfect, animal memory is usually imperfect and mechanisms have evolved to ensure that data that are important to the individual are not lost. Repetition and rote learning are often derided these days, but the importance of multiple exposures to a concept in “anchoring” it in the mind should not be underestimated.

A theme that has recurred over and over in my own thinking about the transitions course is that the “transition” is that from inductive to deductive mental processes. Yet, often, we the instructors of these courses are ourselves so thoroughly ingrained with the deductive approach that the mode of instruction presupposes the very transition we hope to facilitate! In this book I have, to a certain extent, taken the approach of teaching deductive methods using inductive ones. The first time a concept is encountered should only be viewed as providing evidence that lends credence to some mathematical truth. Most concepts that are introduced in this intuitive fashion are eventually explicated in a rigorous manner – there are exceptions though, ideas whose scope is beyond that of the present work which are nonetheless presented here with very little concern for precision. It should not be forgotten that a good transition ought to blend seamlessly into whatever follows. The courses that follow this material should be proof-intensive courses in geometry, number theory, analysis and/or algebra. The introduction of some material from these courses without the usual rigor is intentional.

Please resist the temptation to fill in the missing “proper” definitions and terminology when some concept is introduced and is missing those, uhhh, missing things. Give your students the chance to ruminate, to “chew”³ on these new concepts for a while *on their own!* Later we’ll make sure they get the same standard definitions that we all know and cherish. As a practical matter, if you spend more than 3 weeks in Chapter 1, you are probably filling in too much of that missing detail – so stop it. It really won’t hurt them to

³Why is it that most of the metaphorical ways to refer to “thinking” actually seem to refer to “eating”?

think in an imprecise way (at first) about something so long as we get them to be rigorous by the end of the day.

Finally, it will probably be necessary to point out to your students that they should actually *read* the text. I don't mean to be as snide as that probably sounds. . . Their experiences with math texts up to this point have probably impressed them with the futility of reading — just see what kind of problems are assigned and skim 'til you find an example that shows you “how to do one like that.” Clearly such an approach is far less fruitful in advanced study than it is in courses which emphasize learning calculational techniques. I find that giving expressed reading assignments and quizzing them on the material that they are supposed to have read helps. There are “exercises” given within most sections (as opposed to the “Exercises” that appear at the end of the sections) these make good fodder for quizzes and/or probing questions from the professor. The book is written in an expansive, friendly style with whimsical touches here and there. Some students have reported that they actually enjoyed reading it!⁴

Earlier versions of this text included a systematic error regarding how the natural numbers are defined. In the first chapter they began with 1 and somewhat later the convention shifted so that the smallest natural was 0. I received quite a bit of more or less good natured ribbing about this “continuity error.” In the book you are holding I have endeavored to keep to the convention that 0 is not a natural number.

⁴Although it should be added that they were making that report to someone from whom they wanted a good grade.

Chapter 1

Introduction and notation

Wisdom is the quality that keeps you from getting into situations where you need it. –Doug Larson

1.1 Basic sets

It has been said¹ that “God invented the integers, all else is the work of Man.” This is a mistranslation. The term “integers” should actually be “whole numbers.” The concepts of zero and negative values seem (to many people) to be unnatural constructs. Indeed, otherwise intelligent people are still known to rail against the concept of a negative quantity – “How can you have negative three apples?” The concept of zero is also somewhat profound.

Probably most people will agree that the natural numbers *are* a natural construct – they are the numbers we use to count things. Traditionally, the natural numbers are denoted \mathbb{N} .

At this point in time there seems to be no general agreement about the status of zero (0) as a natural number. Are there collections that we might possibly count that have *no* members? Well, yes – I’d invite you to consider

¹Usually attributed to Kronecker – “Die ganze Zahl schuf der liebe Gott, alles Übrige ist Menschenwerk.”

the collection of gold bars that I keep in my basement. . .

The traditional view seems to be that

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

i.e. that the naturals don't include 0. My personal preference would be to make the other choice (i.e. to include 0 in the natural numbers), but for the moment, let's be traditionalists.

Be advised that this is a choice. We are adopting a convention. If in some other course, or other mathematical setting you find that the other convention is preferred, well, it's good to learn flexibility. . .

Perhaps the best way of saying what a set is, is to do as we have above. List all the elements. Of course, if a set has an infinite number of things in it, this is a difficult task – so we satisfy ourselves by listing enough of the elements that the pattern becomes clear.

Taking \mathbb{N} for granted, what is meant by the “all else” that humankind is responsible for? The basic sets of different types of “numbers” that every mathematics student should know are: \mathbb{N} , \mathbb{Z} , \mathbb{Q} , \mathbb{R} and \mathbb{C} . Respectively: the naturals, the integers, the rationals, the reals, and the complex numbers. The use of \mathbb{N} , \mathbb{R} and \mathbb{C} is probably clear to an English speaker. The integers are denoted with a \mathbb{Z} because of the German word *zählen* which means “to count.” The rational numbers are probably denoted using \mathbb{Q} , for “quotients.” Etymology aside, is it possible for us to provide precise descriptions of these remaining sets?

The integers (\mathbb{Z}) are just the set of natural numbers together with the negatives of naturals and zero. We can use a doubly infinite list to denote this set.

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

To describe the rational numbers precisely we'll have to wait until Section 1.6. In the interim, we can use an intuitively appealing, but somewhat

imprecise definition for the set of rationals. A rational number is a fraction built out of integers. This also provides us with a chance to give an example of using the main other way of describing the contents of a set – so-called set-builder notation.

$$\mathbb{Q} = \left\{ \frac{a}{b} \mid a \in \mathbb{Z} \text{ and } b \in \mathbb{Z} \text{ and } b \neq 0 \right\}$$

This is a good time to start building a “glossary” – a translation lexicon between the symbols of mathematics and plain language. In the line above we are defining the set \mathbb{Q} of rational numbers, so the first symbols that appear are “ $\mathbb{Q} =$.” It is interesting to note that the equals sign has two subtly different meanings: assignment and equality testing, in the mathematical sentence above we are making an assignment – that is, we are declaring that from now on the set \mathbb{Q} will *be* the set defined on the remainder of the line.² Let’s dissect the rest of that line now. There are only 4 characters whose meaning may be in doubt, $\{$, $\}$, \in and $|$. The curly braces (a.k.a. *french braces*) are almost universally reserved to denote sets, anything appearing between curly braces is meant to define a set. In translating from “math” to English, replace the initial brace with the phrase “the set of all.” The next arcane symbol to appear is the vertical bar. As we will see in Section 1.4.3 this symbol has (at least) two meanings – it will always be clear from context which is meant. In the sentence we are analyzing, it stands for the words “such that.” The last bit of arcana to be deciphered is the symbol \in , it stands for the English word “in” or, more formally, “is an element of.”

Let’s parse the entire mathematical sentence we’ve been discussing with an English translation in parallel.

²Some Mathematicians contend that only the “equality test” meaning of the equals sign is real, that by writing the mathematical sentence above we are asserting the truth of the equality test. This may be technically correct but it isn’t how most people think of things.

\mathbb{Q}	=	{
The rational numbers	are defined to be	the set of all
$\frac{a}{b}$		
fractions of the form a over b	such that	
$a \in \mathbb{Z}$	and	$b \in \mathbb{Z}$
a is an element of the integers	and	b is an element of the integers
and	$b \neq 0$	}
and b is nonzero.	(the final curly brace is silent)	

It is quite apparent that the mathematical notation represents a huge improvement as regards brevity.

As mentioned previously, this definition is slightly flawed. We will have to wait 'til later to get a truly precise definition of the rationals, but we invite the reader to mull over what's wrong with this one. Hint: think about the issue of whether a fraction is in lowest terms.

Let's proceed with our menagerie of sets of numbers. The next set we'll consider is \mathbb{R} , the set of real numbers. To someone who has completed Calculus, the reals are perhaps the most obvious and natural notion of what is meant by "number." It may be surprising to learn that the actual definition of what is meant by a real number is extremely difficult. In fact, the first reasonable formulation of a precise definition of the reals came around 1858, more than 180 years after the development of the Calculus³. A precise definition for the set \mathbb{R} of real numbers is beyond the scope of this book, for the moment consider the following intuitive description. A real number is a

³Although it was not published until 1736, Newton's book (*De Methodis Serierum et Fluxionum*) describing both differential and integral Calculus was written in 1671.

number that measures some physical quantity. For example, if a circle has diameter 1 then its circumference is π , thus π is a real number. The points $(0, 0)$ and $(1, 1)$ in the Cartesian plane have distance $\sqrt{(0-1)^2 + (0-1)^2} = \sqrt{2}$, thus $\sqrt{2}$ is a real number. Any rational number is clearly a real number – slope is a physical quantity, and the line from $(0, 0)$ to (b, a) has slope a/b . In ancient Greece, Pythagoras – who has sometimes been described as the first pure Mathematician, believed that every real quantity was in fact rational, a belief that we now know to be false. The numbers π and $\sqrt{2}$ mentioned above are not rational numbers. For the moment it is useful to recall a practical method for distinguishing between rational numbers and real quantities that are not rational – consider their decimal expansions. If the reader is unfamiliar with the result to which we are alluding, we urge you to experiment. Use a calculator or (even better) a computer algebra package to find the decimal expansions of various quantities. Try π , $\sqrt{2}$, $1/7$, $2/5$, $16/17$, $1/2$ and a few other quantities of your own choice. Given that we have already said that the first two of these are not rational, try to determine the pattern. What is it about the decimal expansions that distinguishes rational quantities from reals that aren't rational?

Given that we can't give a precise definition of a real number at this point it is perhaps surprising that we *can* define the set \mathbb{C} of complex numbers with precision (modulo the fact that we define them in terms of \mathbb{R}).

$$\mathbb{C} = \{a + bi \mid a \in \mathbb{R} \text{ and } b \in \mathbb{R} \text{ and } i^2 = -1\}$$

Translating this bit of mathematics into English we get:

\mathbb{C}	=	{
The complex numbers	are defined to be	the set of all
$a + bi$		
expressions of the form a plus b times i		such that

$a \in \mathbb{R}$	and	$b \in \mathbb{R}$
a is an element of the reals	and	b is an element of the reals

and	$i^2 = -1$	}
and	i has the property that its square is negative one.	

We sometimes denote a complex number using a single variable (by convention, either late alphabet Roman letters or Greek letters). Suppose that we've defined $z = a + bi$. The single letter z denotes the entire complex number. We can extract the individual components of this complex number by talking about the real and imaginary parts of z . Specifically, $Re(z) = a$ is called the real part of z , and $Im(z) = b$ is called the imaginary part of z .

Complex numbers are added and multiplied as if they were binomials (polynomials with just two terms) where i is treated as if it were the variable – except that we use the algebraic property that i 's square is -1. For example, to add the complex numbers $1 + 2i$ and $3 - 6i$ we just think of the binomials $1 + 2x$ and $3 - 6x$. Of course we normally write a binomial with the term involving the variable coming first, but this is just a convention. The sum of those binomials would be $4 - 4x$ and so the sum of the given complex numbers is $4 - 4i$. This sort of operation is fairly typical and is called *component-wise* addition. To multiply complex numbers we have to recall how it is that we multiply binomials. This is the well-known FOIL rule (first, outer, inner, last). For example the product of $3 - 2x$ and $4 + 3x$ is $(3 \cdot 4) + (3 \cdot 3x) + (-2x \cdot 4) + (-2x \cdot 3x)$ this expression simplifies to $12 + x - 6x^2$. The analogous calculation with complex numbers looks just the same, until we get to the very last stage where, in simplifying, we use the fact that $i^2 = -1$.

$$\begin{aligned} & (3 - 2i) \cdot (4 + 3i) \\ &= (3 \cdot 4) + (3 \cdot 3i) + (-2i \cdot 4) + (-2i \cdot 3i) \\ &= 12 + 9i - 8i - 6i^2 \\ &= 12 + i + 6 \\ &= 18 + i \end{aligned}$$

The real numbers have a natural ordering, and hence, so do the other sets that are contained in \mathbb{R} . The complex numbers can't really be put into a well-defined order — which should be bigger, 1 or i ? But we do have a way to, at least partially, accomplish this task. The *modulus* of a complex number is a real number that gives the distance from the origin ($0 + 0i$) of the complex plane, to a given complex number. We indicate the modulus using absolute value bars, and you should note that if a complex number happens to be purely real, the modulus and the usual notion of absolute value coincide. If $z = a + bi$ is a complex number, then its modulus, $\|a + bi\|$, is given by the formula $\sqrt{a^2 + b^2}$.

Several of the sets of numbers we've been discussing can be split up based on the so-called *trichotomy property*: every real number is either positive, negative or zero. In particular, \mathbb{Z} , \mathbb{Q} and \mathbb{R} can have modifiers stuck on so that we can discuss (for example) the negative real numbers, or the positive rational numbers or the integers that aren't negative. To do this, we put superscripts on the set symbols, either a $+$ or a $-$ or the word “noneg.”

So

$$\mathbb{Z}^+ = \{x \in \mathbb{Z} \mid x > 0\}$$

and

$$\mathbb{Z}^- = \{x \in \mathbb{Z} \mid x < 0\}$$

and

$$\mathbb{Z}^{\text{nonneg}} = \{x \in \mathbb{Z} \mid x \geq 0\}.$$

Presumably, we could also use “nonpos” as a superscript to indicate non-positive integers, but this never seems to come up in practice. Also, you should note that \mathbb{Z}^+ is really the same thing as \mathbb{N} , but that $\mathbb{Z}^{\text{nonneg}}$ is different because it contains 0.

We would be remiss in closing this section without discussing the way the sets of numbers we’ve discussed fit together. Simply put, each is contained in the next. \mathbb{N} is contained in \mathbb{Z} , \mathbb{Z} is contained in \mathbb{Q} , \mathbb{Q} is contained in \mathbb{R} , and \mathbb{R} is contained in \mathbb{C} . Geometrically the complex numbers are essentially a two-dimensional plane. The real numbers sit inside this plane just as the x -axis sits inside the usual Cartesian plane – in this context you may hear people talk about “the real line within the complex plane.” It is probably clear how \mathbb{N} lies within \mathbb{Z} , and every integer is certainly a real number. The intermediate set \mathbb{Q} (which contains the integers, and is contained by the reals) has probably the most interesting relationship with the set that contains it. Think of the real line as being solid, like a dark pencil stroke. The rationals are like sand that has been sprinkled very evenly over that line. Every point on the line has bits of sand nearby, but not (necessarily) on top of it.

Exercises — 1.1

1. Each of the quantities indexing the rows of the following table is in one or more of the sets which index the columns. Place a check mark in a table entry if the quantity is in the set.

	N	Z	Q	R	C
17					
π					
$22/7$					
-6					
e^0					
$1 + i$					
$\sqrt{3}$					
i^2					

2. Write the set \mathbb{Z} of integers using a singly infinite listing.

3. Identify each as rational or irrational.

(a) 5021.2121212121...

(b) 0.2340000000...

(c) 12.31331133311133331111...

(d) π

(e) 2.987654321987654321987654321...

4. The “see and say” sequence is produced by first writing a 1, then iterating the following procedure: look at the previous entry and say how many entries there are of each integer and write down what you just said. The first several terms of the “see and say” sequence are 1, 11, 21, 1112, 3112, 211213, 312213, 212223, Comment on the rationality (or irrationality) of the number whose decimal digits are obtained by concatenating the “see and say” sequence.

0.1112111123112211213...

5. Give a description of the set of rational numbers whose decimal expansions terminate. (Alternatively, you may think of their decimal expansions ending in an infinitely-long string of zeros.)
6. Find the first 20 decimal places of π , $3/7$, $\sqrt{2}$, $2/5$, $16/17$, $\sqrt{3}$, $1/2$ and $42/100$. Classify each of these quantity’s decimal expansion as: terminating, having a repeating pattern, or showing no discernible pattern.
7. Consider the process of long division. Does this algorithm give any insight as to why rational numbers have terminating or repeating decimal expansions? Explain.
8. Give an argument as to why the product of two rational numbers is again a rational.
9. Perform the following computations with complex numbers

(a) $(4 + 3i) - (3 + 2i)$

(b) $(1 + i) + (1 - i)$

(c) $(1 + i) \cdot (1 - i)$

(d) $(2 - 3i) \cdot (3 - 2i)$

10. The *conjugate* of a complex number is denoted with a superscript star, and is formed by negating the imaginary part. Thus if $z = 3 + 4i$ then the conjugate of z is $z^* = 3 - 4i$. Give an argument as to why the product of a complex number and its conjugate is a real quantity. (I.e. the imaginary part of $z \cdot z^*$ is necessarily 0, no matter what complex number is used for z .)

1.2 Definitions: Prime numbers

You may have noticed that in Section 1.1 an awful lot of emphasis was placed on whether we had good, precise definitions for things. Indeed, more than once apologies were made for giving imprecise or intuitive definitions. This is because, in Mathematics, definitions are our lifeblood. More than in any other human endeavor, Mathematicians strive for precision. This precision comes with a cost – Mathematics can deal with only the very simplest of phenomena⁴. To laypeople who think of math as being a horribly difficult subject, that last sentence will certainly sound odd, but most professional Mathematicians will be nodding their heads at this point. Hard questions are more properly dealt with by Philosophers than by Mathematicians. Does a cat have a soul? Impossible to say, because neither of the nouns in that question can be defined with any precision. Is the square root of 2 a rational number? Absolutely not! The reason for the certainty we feel in answering this second question is that we know *precisely* what is meant by the phrases “square root of 2” and “rational number.”

We often need to first approach a topic by thinking visually or intuitively, but when it comes to proving our assertions, nothing beats the power of having the “right” definitions around. It may be surprising to learn that the “right” definition often evolves over the years. This happens for the simple reason that some definitions lend themselves more easily to proving assertions. In fact, it is often the case that definitions are inspired by attempts to prove something that fail. In the midst of such a failure, it isn’t uncommon for a Mathematician to bemoan “If only the definition of (fill in the blank) were . . .”, then to realize that it is possible to use that definition or a modification of it. But! When there are several definitions for the same idea they had better agree with one another!

⁴For an intriguing discussion of this point, read Gian Carlo Rota’s book *Indiscrete Thoughts* [14].

Consider the definition of a prime number.

Definition. A prime number is a positive integer, greater than 1, whose only factors are 1 and itself.

You probably first heard this definition in Middle School, if not earlier. It is a perfectly valid definition of what it means for an integer to be prime. In more advanced mathematics, it was found that it was necessary to define a notion of primality for objects other than integers. It turns out that the following statement is essentially equivalent to the definition of “prime” we’ve just given (when dealing with integers), but that it can be applied in more general settings.

Definition. A prime is a quantity p such that whenever p is a factor of some product ab , then either p is a factor of a or p is a factor of b .

Exercise. The number 1 is not considered to be a prime. Does 1 satisfy the above definition?

If you go on to study Number Theory or Abstract Algebra you’ll see how the alternate definition we’ve given needs to be tweaked so that (for example) 1 *wouldn’t* get counted as a prime. The fix isn’t hugely complicated (but it is a *little* complicated) and is a bit beyond our scope right now. . .

Often, it is the case that we can formulate *many* equivalent definitions for some concept. When this happens you may run across the abbreviation TFAE, which stands for “The following are equivalent.” A TFAE proof consists of showing that a host of different statements actually define the same concept.

Since we have been discussing primes in this section (mainly as an example of a concept with more than one equivalent definition), this seems like a reasonable time to make some explorations relative to prime numbers. We’ll begin in the third century B.C.. Eratosthenes of Cyrene was a Greek Mathematician and Astronomer who is remembered to this day for his many

accomplishments. He was a librarian at the great library of Alexandria. He made measurements of the Earth's circumference and the distances of the Sun and Moon that were remarkably accurate, but probably his most remembered achievement is the "sieve" method for finding primes. Indeed, the sieve of Eratosthenes is still of importance in mathematical research. Basically, the sieve method consists of creating a very long list of natural numbers and then crossing off all the numbers that aren't primes (a positive integer that isn't 1, and isn't a prime is called *composite*). This process is carried out in stages. First we circle 2 and then cross off every number that has 2 as a factor – thus we've identified 2 as the first prime number and eliminated a whole bunch of numbers that aren't prime. The first number that hasn't been eliminated at this stage is 3, we circle it (indicating that 3 is the second prime number) and then cross off every number that has 3 as a factor. Note that some numbers (for example, 6 and 12) will have been crossed off more than once! In the third stage of the sieve process, we circle 5, which is the smallest number that hasn't yet been crossed off, and then cross off all multiples of 5. The first three stages in the sieve method are shown in Figure 1.1.

It is interesting to note that the sieve gives us a means of finding all the primes up to p^2 by using the primes up to (but not including) p . For example, to find all the primes less than $13^2 = 169$, we need only use 2, 3, 5, 7 and 11 in the sieve.

Despite the fact that one can find primes using this simple mechanical method, the way that prime numbers are distributed amongst the integers is very erratic. Nearly any statement that purports to show some regularity in the distribution of the primes will turn out to be false. Here are two such false conjectures regarding prime numbers.

Conjecture 1. *Whenever p is a prime number, $2^p - 1$ is also a prime.*

Conjecture 2. *The polynomial $x^2 - 31x + 257$ evaluates to a prime number*

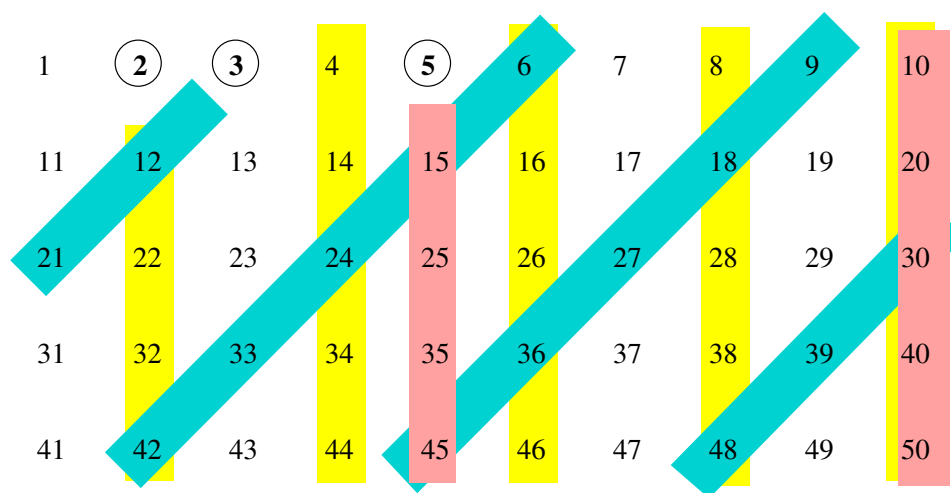


Figure 1.1: The first three stages in the sieve of Eratosthenes. What is the smallest composite number that hasn't been crossed off?

whenever x is a natural number.

In the exercises for this section, you will be asked to explore these statements further.

Prime numbers act as multiplicative building blocks for the rest of the integers. When we disassemble an integer into its building blocks we are finding the *prime factorization* of that number. Prime factorizations are unique. That is, a number is either prime or it has prime factors (possibly raised to various powers) that are uniquely determined – except that they may be re-ordered.

On the next page is a table that contains all the primes that are less than 5000. Study this table and discover the secret of its compactness!

T	0	1	2	3	4	5	6	7	8	9
H										
0	2 3 5 7	1 3 7 9	3 9	1 7	1 3 7	3 9	1 7	1 3 9	3 9	7
1	1 3 7 9	3	7	1 7 9	9	1 7	3 7	3 9	1	1 3 7 9
2		1	3 7 9	3 9	1	1 7	3 9	1 7	1 3	3
3	7	1 3 7		1 7	7 9	3 9	7	3 9	3 9	7
4	1 9	9	1	1 3 9	3 9	7	1 3 7	9	7	1 9
5	3 9		1 3		1 7	7	3 9	1 7	7	3 9
6	1 7	3 7 9		1	1 3 7	3 9	1	3 7	3	1
7	1 9	9	7	3 9	3	1 7	1 9	3	7	7
8	9	1	1 3 7 9	9		3 7 9	3	7	1 3 7	
9	7	1 9	9	7	1 7	3	7	1 7	3	1 7
10	9	3 9	1	1 3 9	9	1	1 3 9		7	1 3 7
11	3 9	7	3 9			1 3	3	1	1 7	3
12	1	3 7	3 9	1 7	9		9	7 9	3 9	1 7
13	1 3 7	9	1 7				1 7	3	1	9
14	9		3 7 9	3 9	7	1 3 9		1	1 3 7 9	3 9
15		1	3	1	3 9	3 9	7	1 9	3	7
16	1 7 9	3 9	1 7	7		3 7	3 7 9			3 7 9
17	9		1 3	3	1 7	3 9		7	3 7 9	
18	1	1	3	1	7		1 7	1 3 7 9	9	
19	1 7	3		1 3	9	1		3 9	7	3 7 9
20	3	1 7	7 9	9		3	3 9		1 3 7 9	9
21		1 3	9	1 7	1 3	3	1	9		
22	3 7	3	1	7 9	3	1	7 9	3	1 7	3 7
23	9	1		3 9	1 7	1 7		1 7	1 3 9	3 9
24		1 7	3	7	1 7	9	7	3 7		
25	3		1	1 9	3 9	1 7		9		1 3
26	9	7	1	3	7	7 9	3	1 7	3 7 9	3 9
27	7	1 3 9	9	1	1 9	3	7	7	9	1 7
28	1 3	9		3 7	3	1 7	1	9	7	7
29	3 9	7	7	9		3 7	3 9	1		9
30	1	1 9	3	7	1 9		1 7	9	3 9	
31	9	9	1	7			3 7 9		1 7	1
32	3 9	7	1 9			1 3 7 9		1		9
33	1 7	3 9	3 9	1	3 7	9	1	1 3	9	1
34	7	3		3	9	7	1 3 7 9			1 9
35		1 7	3 7 9	3 9	1 7	7 9		1	1 3	3
36	7	3 7	3	1 7	3	9		1 3 7		1 7
37	1 9	9	7	3 9			1 7 9	9		3 7
38	3		1 3	3	7	1 3	3	7	1 9	
39	7	1 7 9	3 9	1	3 7		7		9	
40	1 3 7	3 9	1 7		9	1 7		3 9		1 3 9
41		1	7 9	3 9		3 7 9		7		
42	1	1 7 9	9	1	1 3	3 9	1	1 3	3 9	7
43			7	7 9	9	7	3	3		1 7
44	9		1 3		1 7	1 7	3		1 3	3
45	7	3 7 9	3		7 9		1 7		3	1 7
46	3		1	7 9	3 9	1 7	3	3 9		1
47	3		1 3 9	3		1 9			3 7 9	3 9
48	1	3 7		1			1	1 7	9	
49	3 9	9		1 3 7	3	1 7	7 9	3	7	3 9

Exercises — 1.2

1. Find the prime factorizations of the following integers.

(a) 105

(b) 414

(c) 168

(d) 1612

(e) 9177

2. Use the sieve of Eratosthenes to find all prime numbers up to 100.

1	2	3	4	5	6	7	8	9	10
11	12	13	14	15	16	17	18	19	20
21	22	23	24	25	26	27	28	29	30
31	32	33	34	35	36	37	38	39	40
41	42	43	44	45	46	47	48	49	50
51	52	53	54	55	56	57	58	59	60
61	62	63	64	65	66	67	68	69	70
71	72	73	74	75	76	77	78	79	80
81	82	83	84	85	86	87	88	89	90
91	92	93	94	95	96	97	98	99	100

3. What would be the largest prime one would sieve with in order to find all primes up to 400?

4. Characterize the prime factorizations of numbers that are perfect squares.

5. Complete the following table which is related to the conjecture that whenever p is a prime number, $2^p - 1$ is also a prime.

p	$2^p - 1$	prime?	factors
2	3	yes	1 and 3
3	7	yes	1 and 7
5	31	yes	
7	127		
11			

6. Find a counterexample for the conjecture that $x^2 - 31x + 257$ evaluates to a prime number whenever x is a natural number.
7. Use the second definition of “prime” to see that 6 is not a prime. In other words, find two numbers (the a and b that appear in the definition) such that 6 is not a factor of either, but *is* a factor of their product.
8. Use the second definition of “prime” to show that 35 is not a prime.
9. A famous conjecture that is thought to be true (but for which no proof is known) is the Twin Prime conjecture. A pair of primes is said to be twin if they differ by 2. For example, 11 and 13 are twin primes, as are 431 and 433. The Twin Prime conjecture states that there are an infinite number of such twins. Try to come up with an argument as to why 3, 5 and 7 are the only prime triplets.
10. Another famous conjecture, also thought to be true – but as yet unproved, is Goldbach’s conjecture. Goldbach’s conjecture states that every even number greater than 4 is the sum of two odd primes. There is a function $g(n)$, known as the Goldbach function, defined on the positive integers, that gives the number of different ways to write a given number as the sum of two odd primes. For example $g(10) = 2$ since $10 = 5 + 5 = 7 + 3$. Thus another version of Goldbach’s conjecture is that $g(n)$ is positive whenever n is an even number greater than 4.
Graph $g(n)$ for $6 \leq n \leq 20$.

1.3 More scary notation

It is often the case that we want to prove statements that assert something is true for *every* element of a set. For example, “Every number has an additive inverse.” You should note that the truth of that statement is relative, it depends on what is meant by “number.” If we are talking about natural numbers it is clearly false: 3’s additive inverse isn’t in the set under consideration. If we are talking about integers or any of the other sets we’ve considered, the statement is true. A statement that begins with the English words “every” or “all” is called *universally quantified*. It is asserted that the statement holds for *everything* within some universe. It is probably clear that when we are making statements asserting that a thing has an additive inverse, we are not discussing human beings or animals or articles of clothing – we are talking about objects that it is reasonable to add together: numbers of one sort or another. When being careful – and we should always strive to be careful! – it is important to make explicit what universe (known as the *universe of discourse*) the objects we are discussing come from. Furthermore, we need to distinguish between statements that assert that everything in the universe of discourse has some property, and statements that say something about a few (or even just one) of the elements of our universe. Statements of the latter sort are called *existentially quantified*.

Adding to the glossary or translation lexicon we started earlier, there are symbols which describe both these types of quantification. The symbol \forall , an upside-down A, is used for universal quantification, and is usually translated as “for all.” The symbol \exists , a backwards E, is used for existential quantification, it’s translated as “there is” or “there exists.” Lets have a look at a mathematically precise sentence that captures the meaning of the one with which we started this section.

$$\forall x \in \mathbb{Z}, \exists y \in \mathbb{Z}, x + y = 0.$$

Parsing this as we have done before with an English translation in parallel, we get:

$\forall x$	$\in \mathbb{Z}$	$\exists y$
For every number x	in the set of integers	there is a number y
$\in \mathbb{Z}$	$x + y = 0$	
in the integers	having the property that their sum is 0.	

Exercise. Which type of quantification do the following statements have?

1. Every dog has his day.
2. Some days it's just not worth getting out of bed.
3. There's a party in somebody's dorm this Saturday.
4. There's someone for everyone.

A couple of the examples in the exercise above actually have two quantifiers in them. When there are two or more (different) quantifiers in a sentence you have to be careful about keeping their order straight. The following two sentences contain all the same elements except that the words that indicate quantification have been switched. Do they have the same meaning?

For every student in James Woods High School, there is some item of cafeteria food that they like to eat.

There is some item of cafeteria food that every student in James Woods High School likes to eat.

Exercises — 1.3

1. How many quantifiers (and what sorts) are in the following sentence?
“Everybody has *some* friend that thinks they know everything about a sport.”
2. The sentence “Every metallic element is a solid at room temperature.” is false. Why?
3. The sentence “For every pair of (distinct) real numbers there is another real number between them.” is true. Why?
4. Write your own sentences containing four quantifiers. One sentence in which the quantifiers appear $(\forall\exists\forall\exists)$ and another in which they appear $(\exists\forall\exists\forall)$.

1.4 Definitions of elementary number theory

1.4.1 Even and odd

If you divide a number by 2 and it comes out even (i.e. with no remainder) the number is said to be *even*. So the *word* even is related to division. It turns out that the *concept* even is better understood through thinking about multiplication.

Definition. *An integer n is even exactly when there is an integer m such that $n = 2m$.*

You should note that there is a “two-way street” sort of quality to this definition – indeed with most, if not all, definitions. If a number is even, then we are guaranteed the existence of another integer half as big. On the other hand, if we can show that another integer half as big exists, then we know the original number is even. This two-wayness means that the definition is what is known as a *biconditional*, a concept which we’ll revisit in Section 2.2.

A lot of people don’t believe that 0 should be counted as an even number. Now that we are armed with a precise definition, we can answer this question easily. Is there an integer x such that $0 = 2x$? Certainly! let x also be 0. (Notice that in the definition, nothing was said about m and n being distinct from one another.)

An integer is *odd* if it isn’t even. That is, amongst integers, there are only two possibilities: even or odd. We can also define oddness without reference to “even.”

Definition. *An integer n is odd exactly when there is an integer m such that $n = 2m + 1$.*

1.4.2 Decimal and base- n notation

You can also identify even numbers by considering their decimal representation. Recall that each digit in the decimal representation of a number has

a value that depends on its position. For example, the number 3482 really means $3 \cdot 10^3 + 4 \cdot 10^2 + 8 \cdot 10^1 + 2 \cdot 10^0$. This is also known as place notation. The fact that we use the powers of 10 in our place notation is probably due to the fact that most humans have 10 fingers. It is possible to use *any* number in place of 10. In Computer Science there are 3 other bases in common use: 2, 8 and 16 – these are known (respectively) as binary, octal and hexadecimal notation. When denoting a number using some base other than 10, it is customary to append a subscript indicating the base. So, for example, 1011_2 is binary notation meaning $1 \cdot 2^3 + 0 \cdot 2^2 + 1 \cdot 2^1 + 1 \cdot 2^0$ or $8 + 2 + 1 = 11$. No matter what base we are using, the rightmost digit of the number multiplies the base raised to the 0-th power. Any number raised to the 0-th power is 1, and the rightmost digit is consequently known as the units digit. We are now prepared to give some statements that are equivalent to our definition of even. These statements truly don't deserve the designation "theorem," they are immediate consequences of the definition.

Theorem 1.4.1. *An integer is even if the units digit in its decimal representation is one of 0, 2, 4, 6 or 8.*

Theorem 1.4.2. *An integer is even if the units digit in its binary representation is 0.*

For certain problems it is natural to use some particular notational system. For example, the last theorem would tend to indicate that binary numbers are useful in problems dealing with even and odd. Given that there are many different notations that are available to us, it is obviously desirable to have means at our disposal for converting between them. It is possible to develop general rules for converting a base- a number to a base- b number (where a and b are arbitrary) but it is actually more convenient to pick a "standard" base (and since we're human we'll use base-10) and

develop methods for converting between an arbitrary base and the “standard” one. Imagine that in the not-too-distant future we need to convert some numbers from the base-7 system used by the Seven-lobed Amoebazoids from Epsilon Eridani III to the base-12 scheme favored by the Dodecatons of Alpha-Centauri IV. We will need a procedure for converting base-7 to base-10 and another procedure for converting from base-10 to base-12. In the School House Rock episode “Little Twelve Toes” they describe base-12 numeration in a way that is understandable for elementary school children – the digits they use are $\{1, 2, 3, 4, 5, 6, 7, 8, 9, \delta, \epsilon\}$, the last two digits (which are pronounced “dec” and “el”) are necessary since we need single symbols for the things we ordinarily denote using 10 and 11.

Converting from some other base to decimal is easy. You just use the definition of place notation. For example, to find what 451663_7 represents in decimal, just write

$$4 \cdot 7^5 + 5 \cdot 7^4 + 1 \cdot 7^3 + 6 \cdot 7^2 + 6 \cdot 7 + 3 = 4 \cdot 16807 + 5 \cdot 2401 + 1 \cdot 343 + 6 \cdot 49 + 6 \cdot 7 + 3 = 79915.$$

(Everything in the line above can be interpreted as a base-10 number, and no subscripts are necessary for base-10.)

Converting from decimal to some other base is harder. There is an algorithm called “repeated division” that we’ll explore a bit in the exercises for this section. For the moment, just verify that $3\delta 2\epsilon 7_{12}$ is also a representation of the number more conventionally written as 79915.

1.4.3 Divisibility

The notion of being even has an obvious generalization. Suppose we asked whether 3 divided evenly into a given number. Presumably we could make a definition of what it meant to be *threeven*, but rather than doing so (or engaging in any further punnery) we shall instead move to a general definition. We need a notation for the situation when one number divides evenly

into another. There are many ways to describe this situation in English, but essentially just one in “math,” we use a vertical bar – *not* a fraction bar. Indeed the difference between this vertical bar and the fraction symbol ($|$ versus $/$) needs to be strongly stressed. The vertical bar when placed between two numbers is a symbol which asks the question “Does the first number divide evenly (i.e. with no remainder) into the second?” On the other hand the fraction bar asks you to actually carry out some division. The value of $2|5$ is *false*, whereas the value of $2/5$ is $.4$

As was the case in defining even, it turns out that it is best to think of multiplication, not division, when making a formal definition of this concept. Given any two integers n and d we define the symbol $d|n$ by

Definition. $d|n$ exactly when $\exists k \in \mathbb{Z}$ such that $n = kd$.

In spoken language the symbol $d|n$ is translated in a variety of ways:

- d is a divisor of n .
- d divides n evenly.
- d is a factor of n .
- n is an integer multiple of d .

1.4.4 Floor and ceiling

Suppose there is an elevator with a capacity of 1300 pounds. A large group of men who all weigh about 200 pounds want to ascend in it. How many should ride at a time? This is just a division problem, $1300/200$ gives 6.5 men should ride together. Well, obviously putting half a person on an elevator is a bad idea – should we just round-up and let 7 ride together? Not if the 1300 pound capacity rating doesn’t have a safety margin! This is an example of the kind of problem in which the floor function is used. The floor function takes a real number as input and returns the next lower integer.

Suppose after a party we have 43 unopened bottles of beer. We'd like to store them in containers that hold 12 bottles each. How many containers will we need? Again, this is simply a division problem – $43/12 = 3.58\overline{333}$. So we need 3 boxes and another 7 twelfths of a box. Obviously we really need 4 boxes – at least one will have some unused space in it. In this sort of situation we're dealing with the ceiling function. Given a real number, the ceiling function rounds it up to the next integer.

Both of these functions are denoted using symbols that look very much like absolute value bars. The difference lies in some small horizontal strokes.

If x is a real number, its floor is denoted $\lfloor x \rfloor$, and its ceiling is denoted $\lceil x \rceil$. Here are the formal definitions:

Definition. $y = \lfloor x \rfloor$ exactly when $y \in \mathbb{Z}$ and $y \leq x < y + 1$.

Definition. $y = \lceil x \rceil$ exactly when $y \in \mathbb{Z}$ and $y - 1 < x \leq y$.

Basically, the definition of floor says that y is an integer that is less than or equal to x , but $y + 1$ definitely exceeds x . The definition of ceiling can be paraphrased similarly.

1.4.5 Div and mod

In the next section we'll discuss the so-called division algorithm – this may be over-kill since you certainly already know how to do division! Indeed, in the U.S., long division is usually first studied in the latter half of elementary school, and division problems that don't involve a remainder may be found as early as the first grade. Nevertheless, we're going to discuss this process in sordid detail because it gives us a good setting in which to prove relatively easy statements. Suppose you are setting-up a long division problem in which the integer n is being divided by a positive divisor d . (If you want to divide by a negative number, just divide by the corresponding positive number and then throw an extra minus sign on at the end.)

$$\begin{array}{r} q \\ d \overline{) n} \\ \quad \vdots \\ \quad \quad r \end{array}$$

Recall that the answer consists of two parts, a *quotient* q , and a *remainder* r . Of course, r may be zero, but also, the largest r can be is $d - 1$. The assertion that this answer uniquely exists is known as the *quotient-remainder theorem*:

Theorem 1.4.3. *Given integers n and $d > 0$, there are unique integers q and r such that $n = qd + r$ and $0 \leq r < d$.*

The words “div” and “mod” that appear in the title of this subsection provide mathematical shorthand for q and r . Namely, “ $n \bmod d$ ” is a way of expressing the remainder r , and “ $n \operatorname{div} d$ ” is a way of expressing the quotient q .

If two integers, m and n , leave the same remainder when you divide them by d , we say that they are *congruent modulo d* . One could express this by writing $n \bmod d = m \bmod d$, but usually we adopt a shorthand notation

$$n \equiv m \pmod{d}.$$

If one is in a context in which it is completely clear what d is, it’s acceptable to just write $n \equiv m$.

The “mod” operation is used quite a lot in mathematics. When we do computations modulo some number d , (this is known as “modular arithmetic” or, sometimes, “clock arithmetic”) some very nice properties of “mod” come in handy:

$$x + y \bmod d = (x \bmod d + y \bmod d) \bmod d$$

and

$$x \cdot y \bmod d = (x \bmod d \cdot y \bmod d) \bmod d.$$

These rules mean that we can either do the operations first, then reduce the answer mod d or we can do the reduction mod d first and then do the operations (although we may have to do one more round of reduction mod d).

For example, if we are working mod 10, and want to compute $87 \cdot 96 \bmod 10$, we can instead just compute $7 \cdot 6 \bmod 10$, which is 2.

1.4.6 Binomial coefficients

A “binomial” is a polynomial with 2 terms, for example $x + 1$ or $a + b$. The numbers that appear as the coefficients when one raises a binomial to some power are – rather surprisingly – known as binomial coefficients.

Let’s have a look at the first several powers of $a + b$.

$$\begin{aligned}(a + b)^0 &= 1 \\(a + b)^1 &= a + b \\(a + b)^2 &= a^2 + 2ab + b^2\end{aligned}$$

To go much further than the second power requires a bit of work, but try the following

Exercise. Multiply $(a + b)$ and $(a^2 + 2ab + b^2)$ in order to determine $(a + b)^3$. If you feel up to it, multiply $(a^2 + 2ab + b^2)$ times itself in order to find $(a + b)^4$.

Since we’re interested in the coefficients of these polynomials, it’s important to point out that if no coefficient appears in front of a term that means the coefficient is 1.

These binomial coefficients can be placed in an arrangement known as Pascal's triangle ⁵, which provides a convenient way to calculate small binomial coefficients

$$\begin{array}{cccccc}
 & & & & & 1 \\
 & & & & & 1 & 1 \\
 & & & & 1 & 2 & 1 \\
 & & 1 & 3 & 3 & 1 \\
 1 & 4 & 6 & 4 & 1
 \end{array}$$

Figure 1.2: The first 5 rows of Pascal's triangle (which are numbered 0 through 4 ...).

Notice that in the triangle there is a border on both sides containing 1's and that the numbers on the inside of the triangle are the sum of the two numbers above them. You can use these facts to extend the triangle.

Exercise. Add the next two rows to the Pascal triangle in Figure 1.2.

Binomial coefficients are denoted using a somewhat strange looking symbol. The number in the k -th position in row number n of the triangle is denoted $\binom{n}{k}$. This looks a little like a fraction, but the fraction bar is missing. Don't put one in! It's *supposed* to be missing. In spoken English you say " n choose k " when you encounter the symbol $\binom{n}{k}$.

There is a formula for the binomial coefficients – which is nice. Otherwise we'd need to complete a pretty huge Pascal triangle in order to compute something like $\binom{52}{5}$. The formula involves factorial notation. Just to be sure we are all on the same page, we'll define factorials before proceeding.

⁵This triangle was actually known well before Blaise Pascal began to study it, but it carries his name today.

The symbol for factorials is an exclamation point following a number. This is just a short-hand for expressing the product of all the numbers up to a given one. For example $7!$ means $1 \cdot 2 \cdot 3 \cdot 4 \cdot 5 \cdot 6 \cdot 7$. Of course, there's really no need to write the initial 1 — also, for some reason people usually write the product in decreasing order ($7! = 7 \cdot 6 \cdot 5 \cdot 4 \cdot 3 \cdot 2$).

The formula for a binomial coefficient is

$$\binom{n}{k} = \frac{n!}{k! \cdot (n-k)!}.$$

For example

$$\binom{5}{3} = \frac{5!}{3! \cdot (5-3)!} = \frac{1 \cdot 2 \cdot 3 \cdot 4 \cdot 5}{(1 \cdot 2 \cdot 3) \cdot (1 \cdot 2)} = 10.$$

A slightly more complicated example (and one that gamblers are fond of) is

$$\begin{aligned} \binom{52}{5} &= \frac{52!}{5! \cdot (52-5)!} = \frac{1 \cdot 2 \cdot 3 \cdots 52}{(1 \cdot 2 \cdot 3 \cdot 4 \cdot 5) \cdot (1 \cdot 2 \cdot 3 \cdots 47)} \\ &= \frac{48 \cdot 49 \cdot 50 \cdot 51 \cdot 52}{1 \cdot 2 \cdot 3 \cdot 4 \cdot 5} = 2598960. \end{aligned}$$

The reason that a gambler might be interested in the number we just calculated is that binomial coefficients do more than just give us the coefficients in the expansion of a binomial. They also can be used to compute how many ways one can choose a subset of a given size from a set. Thus $\binom{52}{5}$ is the number of ways that one can get a 5 card hand out of a deck of 52 cards.

Exercise. *There are seven days in a week. In how many ways can one choose a set of three days (per week)?*

Exercises — 1.4

1. An integer n is *doubly-even* if it is even, and the integer m guaranteed to exist because n is even is itself even. Is 0 doubly-even? What are the first 3 positive, doubly-even integers?
2. Dividing an integer by two has an interesting interpretation when using binary notation: simply shift the digits to the right. Thus, $22 = 10110_2$ when divided by two gives 1011_2 which is $8 + 2 + 1 = 11$. How can you recognize a doubly-even integer from its binary representation?
3. The *octal* representation of an integer uses powers of 8 in place notation. The digits of an octal number run from 0 to 7, one never sees 8's or 9's. How would you represent 8 and 9 as octal numbers? What octal number comes immediately after 777_8 ? What (decimal) number is 777_8 ?
4. One method of converting from decimal to some other base is called *repeated division*. One divides the number by the base and records the remainder – one then divides the quotient obtained by the base and records the remainder. Continue dividing the successive quotients by the base until the quotient is smaller than the base. Convert 3267 to base-7 using repeated division. Check your answer by using the meaning of base-7 place notation. (For example 54321_7 means $5 \cdot 7^4 + 4 \cdot 7^3 + 3 \cdot 7^2 + 2 \cdot 7^1 + 1 \cdot 7^0$.)
5. State a theorem about the octal representation of even numbers.
6. In hexadecimal (base-16) notation one needs 16 “digits,” the ordinary digits are used for 0 through 9, and the letters A through F are used to give single symbols for 10 through 15. The first 32 natural number in hexadecimal are: 1,2,3,4,5,6,7,8,9,A,B,C,D,E,F,10,11,12,13,14,15,16,17,18,19,1A, 1B,1C,1D,1E,1F,20.

Write the next 10 hexadecimal numbers after AB .

Write the next 10 hexadecimal numbers after FA .

7. For conversion between the three bases used most often in Computer Science we can take binary as the “standard” base and convert using a table look-up. Each octal digit will correspond to a binary triple, and each hexadecimal digit will correspond to a 4-tuple of binary numbers. Complete the following tables. (As a check, the 4-tuple next to A in the table for hexadecimal should be 1010 – which is nice since A is really 10 so if you read that as “ten-ten” it is a good aid to memory.)

octal	binary
0	000
1	001
2	
3	
4	
5	
6	
7	

hexadecimal	binary
0	0000
1	0001
2	0010
3	
4	
5	
6	
7	
8	
9	
A	
B	
C	
D	
E	
F	

8. Use the tables from the previous problem to make the following conversions.

- (a) Convert 757_8 to binary.
 - (b) Convert 1007_8 to hexadecimal.
 - (c) Convert 100101010110_2 to octal.
 - (d) Convert 111101000110101 to hexadecimal.
 - (e) Convert $FEED_{16}$ to binary.
 - (f) Convert $FFFFFF_{16}$ to octal.
9. It is a well known fact that if a number is divisible by 3, then 3 divides the sum of the (decimal) digits of that number. Is this result true in base 7? Do you think this result is true in *any* base?
10. Suppose that 340 pounds of sand must be placed into bags having a 50 pound capacity. Write an expression using either floor or ceiling notation for the number of bags required.
11. True or false?

$$\lfloor \frac{n}{d} \rfloor < \lceil \frac{n}{d} \rceil$$

for all integers n and $d > 0$. Support your claim.

12. What is the value of $\lceil \pi \rceil^2 - \lceil \pi^2 \rceil$?
13. Assuming the symbols n, d, q and r have meanings as in the quotient-remainder theorem (see page 29 of GIAM). Write expressions for q and r , in terms of n and d using floor and/or ceiling notation.
14. Calculate the following quantities:
- (a) $3 \pmod{5}$
 - (b) $37 \pmod{7}$
 - (c) $1000001 \pmod{100000}$

- (d) $6 \text{ div } 6$
- (e) $7 \text{ div } 6$
- (f) $1000001 \text{ div } 2$

15. Calculate the following binomial coefficients:

- (a) $\binom{3}{0}$
- (b) $\binom{7}{7}$
- (c) $\binom{13}{5}$
- (d) $\binom{13}{8}$
- (e) $\binom{52}{7}$

16. An ice cream shop sells the following flavors: chocolate, vanilla, strawberry, coffee, butter pecan, mint chocolate chip and raspberry. How many different bowls of ice cream – with three scoops – can they make?

1.5 Some algorithms of elementary number theory

An *algorithm* is simply a set of clear instructions for achieving some task. The Persian mathematician and astronomer Al-Khwarizmi⁶ was a scholar at the House of Wisdom in Baghdad who lived in the 8th and 9th centuries A.D. He is remembered for his algebra treatise *Hisab al-jabr w'al-muqabala* from which we derive the very word “algebra,” and a text on the Hindu-Arabic numeration scheme.

Al-Khwarizmi also wrote a treatise on Hindu-Arabic numerals. The Arabic text is lost but a Latin translation, *Algoritmi de numero Indorum* (in English *Al-Khwarizmi on the Hindu Art of Reckoning*) gave rise to the word algorithm deriving from his name in the title. [12]

While the study of algorithms is more properly a subject within Computer Science, a student of Mathematics can derive considerable benefit from it.

There is a big difference between an algorithm description intended for human consumption and one meant for a computer⁷. The two favored human-readable forms for describing algorithms are pseudocode and flowcharts. The former is text-based and the latter is visual. There are many different modules from which one can build algorithmic structures: for-next loops, do-while loops, if-then statements, goto statements, switch-case structures, etc. We'll use a minimal subset of the choices available.

- Assignment statements

⁶Abu Ja'far Muhammad ibn Musa al-Khwarizmi

⁷The whole history of Computer Science could be described as the slow advance whereby computers have become able to utilize more and more abstracted descriptions of algorithms. Perhaps in the not-too-distant future machines will be capable of understanding instruction sets that currently require human interpreters.

- If-then control statements
- Goto statements
- Return

We take the view that an algorithm is something like a function, it takes for its input a list of parameters that describe a particular case of some general problem, and produces as its output a solution to that problem. (It should be noted that there are other possibilities – some programs require that the variable in which the output is to be placed be handed them as an input parameter, others have no specific output, their purpose is achieved as a side-effect.) The intermediary between input and output is the algorithm instructions themselves and a set of so-called local variables which are used much the way scrap paper is used in a hand calculation – intermediate calculations are written on them, but they are tossed aside once the final answer has been calculated.

Assignment statements allow us to do all kinds of arithmetic operations (or rather to think of these types of operations as being atomic.) In actuality even a simple procedure like adding two numbers requires an algorithm of sorts, we'll avoid such a fine level of detail. Assignments consist of evaluating some (possibly quite complicated) formula in the inputs and local variables and assigning that value to some local variable. The two uses of the phrase “local variable” in the previous sentence do not need to be distinct, thus $x = x + 1$ is a perfectly legal assignment.

If-then control statements are decision makers. They first calculate a Boolean expression (this is just a fancy way of saying something that is either **true** or **false**), and send program flow to different locations depending on that result. A small example will serve as an illustration. Suppose that in the body of an algorithm we wish to check if 2 variables, x and y are equal, and if they are, increment x by 1. This is illustrated in Figure 1.3 both in pseudocode and as a flowchart.

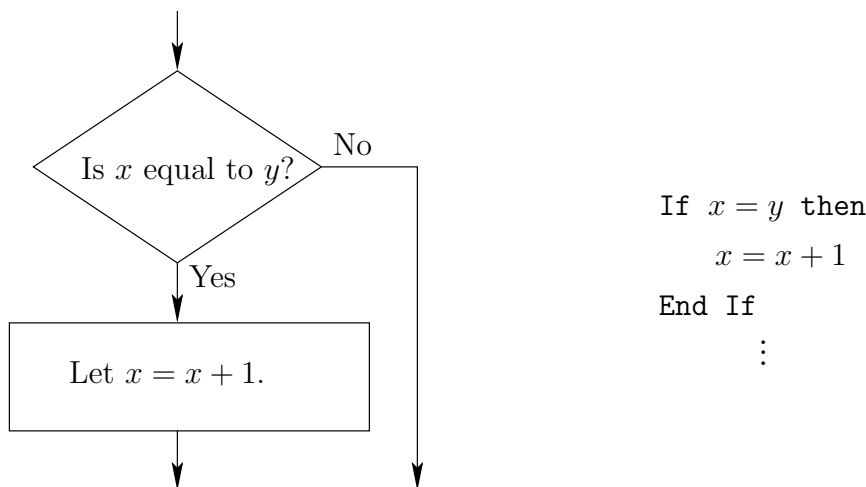


Figure 1.3: A small example in pseudocode and as a flowchart

Notice the use of indentation in the pseudocode example to indicate the statements that are executed if the Boolean expression is true. These examples also highlight the difference between the two senses that the word “equals” (and the symbol $=$) has. In the Boolean expression the sense is that of *testing* equality, in the assignment statements (as the name implies) an *assignment* is being made. In many programming languages this distinction is made explicit, for instance in the C language equality testing is done via the symbol “ $==$ ” whereas assignment is done using a single equals sign ($=$). In Mathematics the equals sign usually indicates equality testing, when the assignment sense is desired the word “let” will generally precede the equality.

While this brief introduction to the means of notating algorithms is by no means complete, it is hopefully sufficient for our purpose which is solely to introduce two algorithms that are important in elementary number theory. The division algorithm, as presented here, is simply an explicit version of the process one follows to calculate a quotient and remainder using long division. The procedure we give is unusually inefficient – with very little thought one

could devise an algorithm that would produce the desired answer using many fewer operations – however the main point here is purely to show that division can be accomplished by essentially mechanical means. The Euclidean algorithm is far more interesting both from a theoretical and a practical perspective. The Euclidean algorithm computes the greatest common divisor (gcd) of two integers. The gcd of two numbers a and b is denoted $\text{gcd}(a, b)$ and is the largest integer that divides both a and b evenly.

A pseudocode outline of the division algorithm is as follows:

```
Algorithm: Division
Inputs: integers  $n$  and  $d$ .
Local variables:  $q$  and  $r$ .

Let  $q = 0$ .
Let  $r = n$ .
Label 1.
If  $r < d$  then
    Return  $q$  and  $r$ .
End If
Let  $q = q + 1$ .
Let  $r = r - d$ .
Goto 1.
```

This same algorithm is given in flowchart form in Figure 1.4.

Note that in a flowchart the action of a “Goto” statement is clear because an arrow points to the location where program flow is being redirected. In pseudocode a “Label” statement is required which indicates a spot where flow can be redirected via subsequent “Goto” statements. Because of the potential for confusion in complicated algorithms that involve multitudes of Goto statements and their corresponding Labels, this sort of redirection is now deprecated in virtually all popular programming environments.

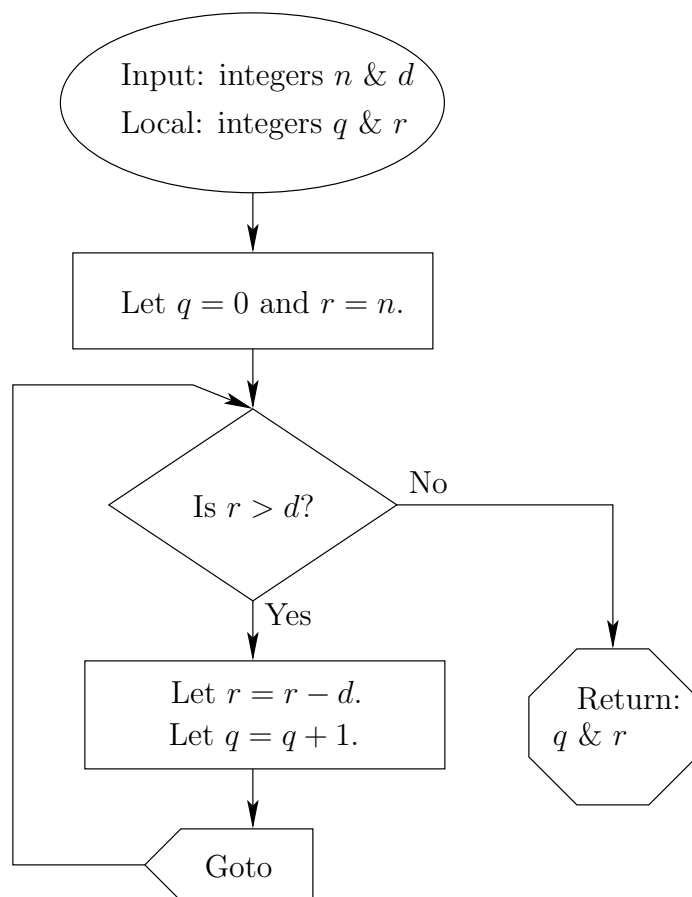


Figure 1.4: The division algorithm in flowchart form.

Before we move on to describe the Euclidean algorithm it might be useful to describe more explicitly what exactly it's *for*. Given a pair of integers, a and b , there are two quantities that it is important to be able to compute, the *least common multiple* or lcm, and the *greatest common divisor* or gcd. The lcm also goes by the name *lowest common denominator* because it is the smallest denominator that could be used as a common denominator in the process of adding two fractions that had a and b in their denominators. The

gcd and the lcm are related by the formula

$$\text{lcm}(a, b) = \frac{ab}{\text{gcd}(a, b)},$$

so they are essentially equivalent as far as representing a computational challenge.

The Euclidean algorithm depends on a rather extraordinary property of the gcd. Suppose that we are trying to compute $\text{gcd}(a, b)$ and that a is the larger of the two numbers. We first feed a and b into the division algorithm to find q and r such that $a = qb + r$. It turns out that b and r have the *same* gcd as did a and b . In other words, $\text{gcd}(a, b) = \text{gcd}(b, r)$, furthermore these numbers are smaller than the ones we started with! This is nice because it means we're now dealing with an easier version of the same problem. In designing an algorithm it is important to formulate a clear *ending criterion*, a condition that tells you you're done. In the case of the Euclidean algorithm, we know we're done when the remainder r comes out 0.

So, here, without further ado is the Euclidean algorithm in pseudocode. A flowchart version is given in Figure 1.5.

```
Algorithm: Euclidean
Inputs: integers  $a$  and  $b$ .
Local variables:  $q$  and  $r$ .

Label 1.
Let  $(q, r) = \text{Division}(a, b)$ .
If  $r = 0$  then
    Return  $b$ .
End If
Let  $a = b$ .
Let  $b = r$ .
Goto 1.
```

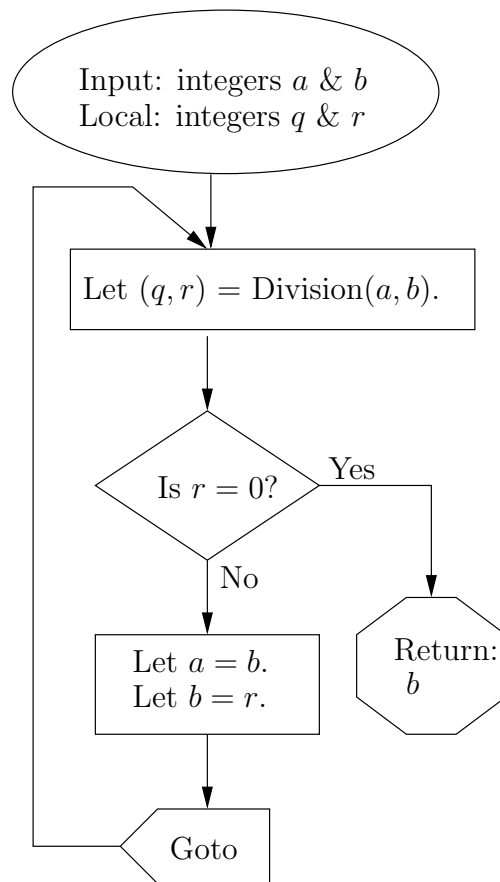


Figure 1.5: The Euclidean algorithm in flowchart form.

It should be noted that for small numbers one can find the gcd and lcm quite easily by considering their factorizations into primes. For the moment consider numbers that factor into primes but not into prime powers (that is, their factorizations don't involve exponents). The gcd is the product of the primes that are in common between these factorizations (if there are no primes in common it is 1). The lcm is the product of all the distinct primes that appear in the factorizations. As an example, consider 30 and 42. The factorizations are $30 = 2 \cdot 3 \cdot 5$ and $42 = 2 \cdot 3 \cdot 7$. The primes that are common to both factorizations are 2 and 3, thus $\text{gcd}(30, 42) = 2 \cdot 3 = 6$. The set of all the primes that appear in either factorization is $\{2, 3, 5, 7\}$ so $\text{lcm}(30, 42) = 2 \cdot 3 \cdot 5 \cdot 7 = 210$.

The technique just described is of little value for numbers having more than about 50 decimal digits because it rests *a priori* on the ability to find the prime factorizations of the numbers involved. Factoring numbers is easy enough if they're reasonably small, especially if some of their prime factors are small, but in general the problem is considered so difficult that many cryptographic schemes are based on it.

Exercises — 1.5

1. Trace through the division algorithm with inputs $n = 27$ and $d = 5$, each time an assignment statement is encountered write it out. How many assignments are involved in this particular computation?
2. Find the gcd's and lcm's of the following pairs of numbers.

a	b	$\gcd(a, b)$	$\text{lcm}(a, b)$
110	273		
105	42		
168	189		

3. Formulate a description of the gcd of two numbers in terms of their prime factorizations in the general case (when the factorizations may include powers of the primes involved).
4. Trace through the Euclidean algorithm with inputs $a = 3731$ and $b = 2730$, each time the assignment statement that calls the division algorithm is encountered write out the expression $a = qb + r$. (With the actual values involved !)

1.6 Rational and irrational numbers

When we first discussed the rational numbers in Section 1.1 we gave the following definition, which isn't quite right.

$$\mathbb{Q} = \left\{ \frac{a}{b} \mid a \in \mathbb{Z} \text{ and } b \in \mathbb{Z} \text{ and } b \neq 0 \right\}$$

We are now in a position to fix the problem.

So what was the problem after all? Essentially this: there are many expressions formed with one integer written above another (with an intervening fraction bar) that represent the exact same rational number. For example $\frac{3}{6}$ and $\frac{14}{28}$ are distinct things that appear in the set defined above, but we all know that they both represent the rational number $\frac{1}{2}$. To eliminate this problem with our definition of the rationals we need to add an additional condition that ensures that such duplicates don't arise. It turns out that what we want is for the numerators and denominators of our fractions to have *no* factors in common. Another way to say this is that the a and b from the definition above should be chosen so that $\gcd(a, b) = 1$. A pair of numbers whose gcd is 1 are called *relatively prime*.

We're ready, at last, to give a good, precise definition of the set of rational numbers. (Although it should be noted that we're not quite done fiddling around; an even better definition will be given in Section 6.3.)

$$\mathbb{Q} = \left\{ \frac{a}{b} \mid a, b \in \mathbb{Z} \text{ and } b \neq 0 \text{ and } \gcd(a, b) = 1 \right\}.$$

As we have in the past, let's parse this with an English translation in parallel.

\mathbb{Q}	=	{
The rational numbers	are defined to be	the set of all

$\frac{a}{b}$		$a, b \in \mathbb{Z}$
fractions of the form a over b	such that	a and b are integers
and $b \neq 0$	and	$\gcd(a, b) = 1$
and b is non-zero	and	a and b are relatively prime.

Finally, we are ready to face a fundamental problem that was glossed-over in Section 1.1. We defined two sets back then, \mathbb{Q} and \mathbb{R} , the hidden assumption that one makes in asserting that there are two of something is that the two things are distinct. Is this really the case? The reals have been defined (unrigorously) as numbers that measure the magnitudes of physical quantities, so another way to state the question is this: Are there physical quantities (for example lengths) that are *not* rational numbers?

The answer is that *yes* there are numbers that measure lengths which are not rational numbers. With our new and improved definition of what is meant by a rational number we are ready to *prove* that there is at least one length that can't be expressed as a fraction. Using the Pythagorean theorem it's easy to see that the length of the diagonal of a unit square is $\sqrt{2}$. The proof that $\sqrt{2}$ is not rational is usually attributed to the followers of Pythagoras (but probably not to Pythagoras himself). In any case it is a result of great antiquity. The proof is of a type known as *reductio ad absurdum*⁸. We show that a given assumption leads logically to an absurdity, a statement that *can't* be true, then we know that the original assumption must itself be false. This method of proof is a bit slippery; one has to first assume the *exact opposite* of what one hopes to prove and then argue (on purpose) towards a ridiculous conclusion.

Theorem 1.6.1. *The number $\sqrt{2}$ is not in the set \mathbb{Q} of rational numbers.*

Before we can actually give the proof we should prove an intermediary

⁸Reduction to an absurdity – better known these days as proof by contradiction.

result – but we won't, we'll save this proof for the student to do later (heh, heh, heh. . .). These sorts of intermediate results, things that don't deserve to be called theorems themselves, but that aren't entirely self-evident are known as lemmas. It is often the case that in an attempt at proving a statement we find ourselves in need of some small fact. Perhaps it even seems to be true but it's not clear. In such circumstances, good form dictates that we first state and prove the lemma then proceed on to our theorem and its proof. So, here, without its proof is the lemma we'll need.

Lemma 1.6.2. *If the square of an integer is even, then the original integer is even.*

Given that thoroughness demands that we fill in this gap by actually proving the lemma at a later date, we can now proceed with the proof of our theorem.

Proof: Suppose to the contrary that $\sqrt{2}$ is a rational number. Then by the definition of the set of rational numbers, we know that there are integers a and b having the following properties: $\sqrt{2} = \frac{a}{b}$ and $\gcd(a, b) = 1$.

Consider the expression $\sqrt{2} = \frac{a}{b}$. By squaring both sides of this we obtain

$$2 = \frac{a^2}{b^2}.$$

This last expression can be rearranged to give

$$a^2 = 2b^2$$

An immediate consequence of this last equation is that a^2 is an even number. Using the lemma above we now know that a is an

even integer and hence that there is an integer m such that $a = 2m$. Substituting this last expression into the previous equation gives

$$(2m)^2 = 2b^2,$$

thus,

$$4m^2 = 2b^2,$$

so

$$2m^2 = b^2.$$

This tells us that b^2 is even, and hence (by the lemma), b is even.

Finally, we have arrived at the desired absurdity because if a and b are both even then $\gcd(a, b) \geq 2$, but, on the other hand, one of our initial assumptions is that $\gcd(a, b) = 1$.

Q.E.D.

Exercises — 1.6

1. Rational Approximation is a field of mathematics that has received much study. The main idea is to find rational numbers that are very good approximations to given irrationals. For example, $22/7$ is a well-known rational approximation to π . Find good rational approximations to $\sqrt{2}$, $\sqrt{3}$, $\sqrt{5}$ and e .
2. The theory of base- n notation that we looked at in the sub-section on base- n can be extended to deal with real and rational numbers by introducing a decimal point (which should probably be re-named in accordance with the base) and adding digits to the right of it. For instance 1.1011 is binary notation for $1 \cdot 2^0 + 1 \cdot 2^{-1} + 0 \cdot 2^{-2} + 1 \cdot 2^{-3} + 1 \cdot 2^{-4}$ or $1 + \frac{1}{2} + \frac{1}{8} + \frac{1}{16} = 1\frac{11}{16}$.
Consider the binary number $.1010010001000010000010000001 \dots$, is this number rational or irrational? Why?
3. If a number x is even, it's easy to show that its square x^2 is even. The lemma that went unproved in this section asks us to start with a square (x^2) that is even and deduce that the UN-squared number (x) is even. Perform some numerical experimentation to check whether this assertion is reasonable. Can you give an argument that would prove it?
4. The proof that $\sqrt{2}$ is irrational can be generalized to show that \sqrt{p} is irrational for every prime number p . What statement would be equivalent to the lemma about the parity of x and x^2 in such a generalization?
5. Write a proof that $\sqrt{3}$ is irrational.

1.7 Relations

One of the principle ways in which mathematical writing differs from ordinary writing is in its incredible brevity. For instance, a Ph.D. thesis for someone in the humanities would be very suspicious if its length were less than 300 pages, whereas it would be quite acceptable for a math doctoral student to submit a thesis amounting to less than 100 pages. Indeed, the usual criteria for a doctoral thesis (or indeed any scholarly work in mathematics) is that it be “new, true and interesting.” If one can prove a truly interesting, novel result in a single page – they’ll probably hand over the sheepskin.

How is this great brevity achieved? By inserting single symbols in place of a whole paragraph’s worth of words! One class of symbols in particular has immense power – so-called relational symbols. When you place a relational symbol between two expressions, you create a sentence that says the relation *holds*. The period at the end of the last sentence should probably be pronounced! “The relation holds, period!” In other words when you write down a mathematical sentence involving a relation, you are asserting the relation is True (the capital T is intentional). This is why it’s okay to write “ $2 < 3$ ” but it’s *not* okay to write “ $3 < 2$.” The symbol $<$ is a relation symbol and you are only supposed to put it between two things when they actually bear this relation to one another.

The situation becomes slightly more complicated when we have variables in relational expressions, but before we proceed to consider that complication let’s make a list of the relations we’ve seen to date:

$$=, <, >, \leq, \geq, |, \text{ and } \equiv \pmod{m}.$$

Each of these, when placed between numbers, produces a statement that is either true or false. Ordinarily we wouldn’t write down the false ones, instead we should express that we know the relation *doesn’t* hold by negating the relation symbol (often by drawing a slash through it, but some of the

symbols above are negations of others).

So what about expressions involving variables and these relation symbols? For example what does $x < y$ really mean? Okay, I know that you know what $x < y$ means but, philosophically, a relation symbol involving variables is doing something that you may have only been vaguely aware of in the past – it is introducing a *supposition*. Watch out for relation symbols involving variables! Whenever you encounter them it means the rules of the game are being subtly altered – up until the point where you see $x < y$, x and y are just two random numbers, but after that point we must suppose that x is the smaller of the two.

The relations we've discussed so far are *binary* relations, that is, they go in between *two* numbers. There are also higher order relations. For example, a famous ternary relation (a relationship between three things) is the notion of “betweenness.” If A , B and C are three points which all lie on a single line, we write $A \star B \star C$ if B falls somewhere on the line segment \overline{AC} . So the symbol $A \star B \star C$ is shorthand for the sentence “Point B lies somewhere in between points A and C on the line determined by them.”

There is a slightly silly tendency these days to define functions as being a special class of relations. (This is slightly silly not because it's wrong – indeed, functions are a special type of relation – but because it's the least intuitive approach possible, and it is usually foisted-off on middle or high school students.) When this approach is taken, we first define a relation to be *any* set of ordered pairs and then state a restriction on the ordered pairs that may be in a relation if it is to be a function. Clearly what these Algebra textbook authors are talking about are *binary* relations, a ternary relation would actually be a set of ordered triples, and higher order relations might involve ordered 4-tuples or 5-tuples, etc. A couple of small examples should help to clear up this connection between a relation symbol and some set of tuples.

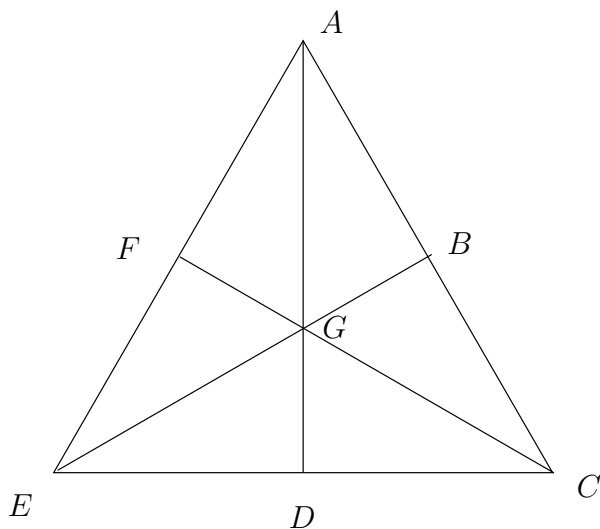
Consider the numbers from 1 to 5 and the less-than relation, $<$. As a set

of ordered pairs, this relation is the set

$$\{(1, 2), (1, 3), (1, 4), (1, 5), (2, 3), (2, 4), (2, 5), (3, 4), (3, 5), (4, 5)\}.$$

The pairs that are *in* the relation are those such that the first is smaller than the second.

An example involving the ternary relation “betweenness” can be had from the following diagram.



The betweenness relation on the points in this diagram consists of the following triples.

$$\{(A, B, C), (A, G, D), (A, F, E), (B, G, E), (C, B, A), (C, G, F), (C, D, E), \\ (D, G, A), (E, D, C), (E, G, B), (E, F, A), (F, G, C)\}.$$

Exercise. When thinking of a function as a special type of relation, the pairs are of the form $(x, f(x))$. That is, they consist of an input and the corresponding output. What is the restriction that must be placed on the pairs in a relation if it is to be a function? (Hint: think about the so-called vertical line test.)

Exercises — 1.7

1. Consider the numbers from 1 to 10. Give the set of pairs of these numbers that corresponds to the divisibility relation.
2. The *domain* of a function (or binary relation) is the set of numbers appearing in the first coordinate. The *range* of a function (or binary relation) is the set of numbers appearing in the second coordinate.

Consider the set $\{0, 1, 2, 3, 4, 5, 6\}$ and the function $f(x) = x^2 \pmod{7}$. Express this function as a relation by explicitly writing out the set of ordered pairs it contains. What is the range of this function?

3. What relation on the numbers from 1 to 10 does the following set of ordered pairs represent?

$$\begin{aligned} &\{(1, 1), (1, 2), (1, 3), (1, 4), (1, 5), (1, 6), (1, 7), (1, 8), (1, 9), (1, 10), \\ &\quad (2, 2), (2, 3), (2, 4), (2, 5), (2, 6), (2, 7), (2, 8), (2, 9), (2, 10), \\ &\quad (3, 3), (3, 4), (3, 5), (3, 6), (3, 7), (3, 8), (3, 9), (3, 10), \\ &\quad (4, 4), (4, 5), (4, 6), (4, 7), (4, 8), (4, 9), (4, 10), \\ &\quad (5, 5), (5, 6), (5, 7), (5, 8), (5, 9), (5, 10), \\ &\quad (6, 6), (6, 7), (6, 8), (6, 9), (6, 10), \\ &\quad (7, 7), (7, 8), (7, 9), (7, 10), \\ &\quad (8, 8), (8, 9), (8, 10), \\ &\quad (9, 9), (9, 10), \\ &\quad (10, 10)\} \end{aligned}$$

4. Draw a five-pointed star, label all 10 points. There are 40 triples of these labels that satisfy the betweenness relation. List them.

5. Sketch a graph of the relation

$$\{(x, y) \mid x, y \in \mathbb{R} \text{ and } y > x^2\}.$$

6. A function $f(x)$ is said to be *invertible* if there is another function $g(x)$ such that $g(f(x)) = x$ for all values of x . (Usually, the inverse function, $g(x)$ would be denoted $f^{-1}(x)$.) Suppose a function is presented to you as a relation – that is, you are just given a set of pairs. How can you distinguish whether the function represented by this list of input/output pairs is invertible? How can you produce the inverse (as a set of ordered pairs)?
7. There is a relation known as “has color” which goes from the set

$$F = \{\textit{orange, cherry, pumpkin, banana}\}$$

to the set

$$C = \{\textit{orange, red, green, yellow}\}.$$

What pairs are in “has color”?

Chapter 2

Logic and quantifiers

If at first you don't succeed, try again. Then quit. There's no use being a damn fool about it. –W. C. Fields

2.1 Predicates and Logical Connectives

In every branch of Mathematics there are special, atomic, notions that defy precise definition. In Geometry, for example, the atomic notions are points, lines and their incidence. Euclid defines a point as “that which has no part” – people can argue (and have argued) incessantly over what exactly is meant by this. Is it essentially saying that anything without volume, area or length of some sort is a point? In modern times it has been recognized that any formal system of argumentation has to have such elemental, undefined, concepts – and that Euclid’s apparent lapse in precision comes from an attempt to hide this basic fact. The notion of “point” can’t really *be* defined. All we can do is point (no joke intended) at a variety of points and hope that our audience will absorb the same concept of point that we hold via the process of induction¹.

The atomic concepts in Set Theory are “set”, “element” and “member-

¹inference of a generalized conclusion from particular instances – compare DEDUCTION

ship”. The atomic concepts in Logic are “true”, “false”, “sentence” and “statement”.

Regarding *true* and *false*, we hope there is no uncertainty as to their meanings. *Sentence* also has a well-understood meaning that most will agree on – a syntactically correct ordered collection of words such as “Johnny was a football player.” or “Red is a color.” or “This is a sentence which does not refer to itself.” A *statement* is a sentence which is either true or false. In other words, a statement is a sentence whose truth value is *definite*, in more other words, it is always possible to decide – one way or the other – whether a statement is true or false.² The first example of a sentence given above (“Johnny was a football player”) is not a statement – the problem is that it is ambiguous unless we know who Johnny is. If it had said “Johnny Unitas was a football player.” then it would have been a statement. If it had said “Johnny Appleseed was a football player.” it would also have been a statement, just not a true one.

Ambiguity is only one reason that a sentence may not be a statement. As we consider more complex sentences, it may be the case that the truth value of a given sentence simply cannot be decided. One of the most celebrated mathematical results of the 20th century is Kurt Gödel’s “Incompleteness Theorem.” An important aspect of this theory is the proof that in any axiomatic system of mathematical thought there must be undecidable sentences – statements which can neither be proved nor disproved from the axioms³. Simple sentences (e.g. those of the form subject-verb-object) have little chance of being undecidable for this reason, so we will next look at ways of building more complex sentences from simple components.

²Although, as a practical matter it may be almost impossibly difficult to do so! For instance it is certainly either true or false that I ate eggs for breakfast on my 21st birthday – but I don’t remember, and short of building a time machine, I don’t know how you could find out.

³There are trivial systems that are complete, but if a system is sufficiently complicated that it contains “interesting” statements it can’t be complete.

Let's start with an example. Suppose I come up to you in some windowless room and make the statement: "The sun is shining but it's raining!" You decide to investigate my claim and determine its veracity. Upon reaching a room that has a view of the exterior there are four possible combinations of sunniness and/or precipitation that you may find. That is, the atomic predicates "The sun is shining" and "It is raining" can each be true or false independently of one another. In the following table we introduce a convention used throughout the remainder of this book – that true is indicated with a capital letter T and false is indicated with the Greek letter ϕ (which is basically a Greek F, and is a lot harder to mistake for a T than an F is.)

The sun is shining	It is raining
T	T
T	ϕ
ϕ	T
ϕ	ϕ

Each row of the above table represents a possible state of the outside world. Suppose you observe the conditions given in the last row, namely that it is neither sunny, nor is it raining – you would certainly conclude that I am not to be trusted. I.e. my statement, the compounding of "The sun is shining" and "It is raining" (with the word "but" in between as a connector) is false. If you think about it a bit, you'll agree that this so-called *compound sentence* is true only in the case that both of its component pieces are true. This underscores an amusing linguistic point: "but" and "and" have exactly the same meaning! More precisely, they *denote* the same thing, they have subtly different connotations however – "but" indicates that both of the statements it connects are true and that the speaker is surprised by this state of affairs.

In Mathematics we distinguish two main connectives for hooking-up simple sentences into compound ones. The *conjunction* of two sentences is the

compound sentence made by sticking the word “and” between them. The *disjunction* of two sentences is formed by placing an “or” between them. Conjunctions are true only when both components are true. Disjunctions are false only when both components are false.

As usual, mathematicians have developed an incredibly terse, compact notation for these ideas.⁴ First, we represent an entire sentence by a single letter – traditionally, a capital letter. This is called a *predicate variable*. For example, following the example above, we could denote the sentence “The sun is shining” by the letter S . Similarly, we could make the assignment $R =$ “It is raining.” The conjunction and disjunction of these sentences can then be represented using the symbols $S \wedge R$ and $S \vee R$, respectively. As a mnemonic, note that the connective in $S \wedge R$ looks very much like the capital letter A (as in And).

To display, very succinctly, the effect of these two connectives we can use so-called truth tables. In a truth table we list all possible truth values of the predicate variables and then enumerate the truth values of some compound sentence. For the conjunction and disjunction connectors we have (respectively):

A	B	$A \wedge B$		A	B	$A \vee B$	
T	T	T		T	T	T	
T	ϕ	ϕ	and	T	ϕ	T	.
ϕ	T	ϕ		ϕ	T	T	
ϕ	ϕ	ϕ		ϕ	ϕ	ϕ	

In addition to these connectors we need a modifier (called *negation*) that acts on individual sentences. The negation of a sentence A is denoted by $\neg A$, and its truth value is exactly the opposite of A 's truth value. The negation of a sentence is also known as the *denial* of a sentence. A truth table for the negation operator is somewhat trivial but we include it here for completeness.

⁴One begins to suspect that mathematicians form an unusually lazy sub-species of humanity.

A	$\neg A$
T	ϕ
ϕ	T

These three simple tools (and, or & not) are sufficient to create extraordinarily complex sentences out of basic components. The way these pieces interrelate is a bit reminiscent of algebra, in fact the study of these logical operators (or any operators that act like them) is called *Boolean Algebra*⁵. There are distinct differences between Boolean and ordinary algebra however. In regular algebra we have the binary connectors $+$ (plus) and \cdot (times), and the unary negation operator $-$, these are certainly analogous to \wedge , \vee & \neg , but there are certain consequences of the fact that multiplication is effectively repeated addition that simply don't hold for the Boolean operators. For example, there is a well-defined precedence between \cdot and $+$. In parsing the expression $4 \cdot 5 + 3$ we all know that the multiplication is to be done first. There is no such rule governing order of operations between \wedge and \vee , so an expression like $A \wedge B \vee C$ is simply ambiguous – it *must* have parentheses inserted in order to show the order, either $(A \wedge B) \vee C$ or $A \wedge (B \vee C)$. Another distinction between ordinary and Boolean algebra is exponentiation. If there *were* exponents in Boolean algebra, we'd need two different kinds – one for repeated conjunction and another for repeated disjunction.

Exercise. *Why is it that there is no such thing as exponentiation in the algebra of Logic?*

While there are many differences between Boolean algebra and the usual, garden-variety algebra, there are also many similarities. For instance, the associative, commutative and distributive laws of Algebra all have versions that work in the Boolean case.

⁵In honor of George Boole, whose 1854 book *An investigation into the Laws of Thought* inaugurated the subject.

A very handy way of visualizing Boolean expressions is given by digital logic circuit diagrams. To discuss these diagrams we must make a brief digression into Electronics. One of the most basic components inside an electronic device is a transistor, this is a component that acts like a switch for electricity, but the switch itself is controlled by electricity. In Figure 2.1 we see the usual schematic representation of a transistor. If voltage is applied to the wire labeled z , the transistor becomes conductive, and current may flow from x to y .

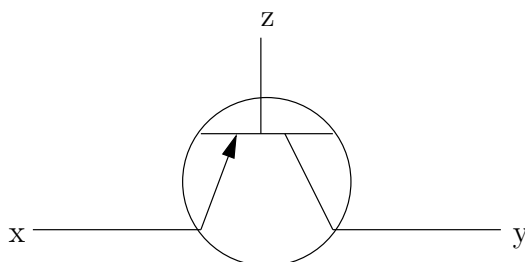


Figure 2.1: A schematic representation of a transistor.

Suppose that two transistors are connected as in Figure 2.2 (this is called a *series* connection). In order for current to flow from x to y we must have voltage applied to *both* the wires labeled z and w . In other words, this circuit effectively creates the **and** operation – assuming voltage is always applied to x , if z **and** w are energized then the output at y will be energized.

When two transistors are connected in parallel (this is illustrated in Figure 2.3) current can flow from x to y when either (or *both*) of the wires at z and w have voltage applied. This brings up a point which is confusing for some: in common speech the use of the word “or” often has the sense known as *exclusive or* (a.k.a. xor), when we say “X or Y” we mean “Either X or Y, but not both.” In Electronics and Mathematics, *or* always has the non-exclusive (better known as inclusive) sense.

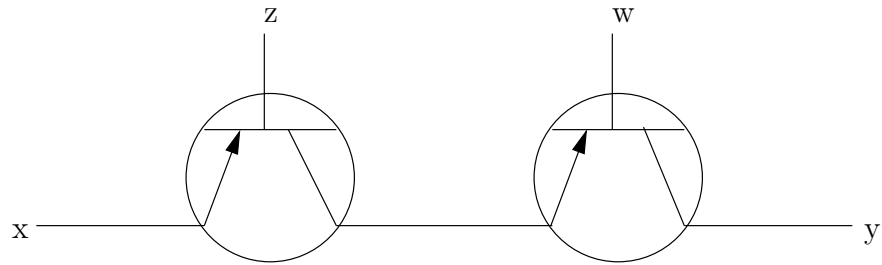


Figure 2.2: The connection of two transistors in series provides an implementation of the *and* operator.

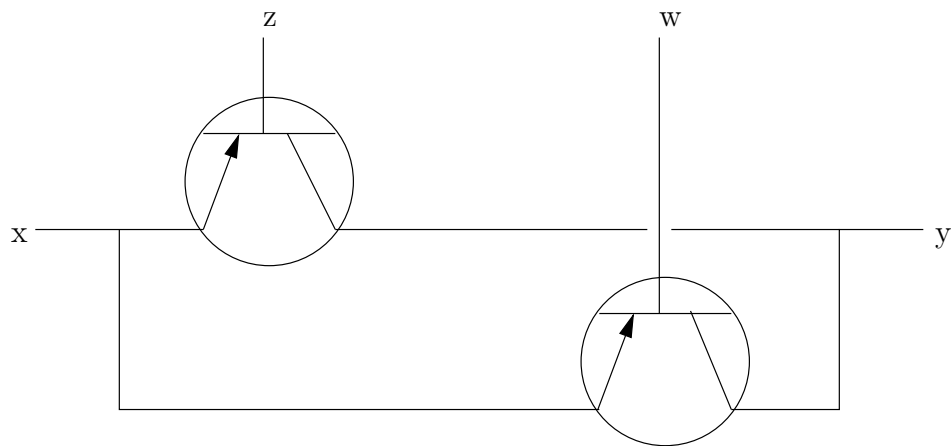
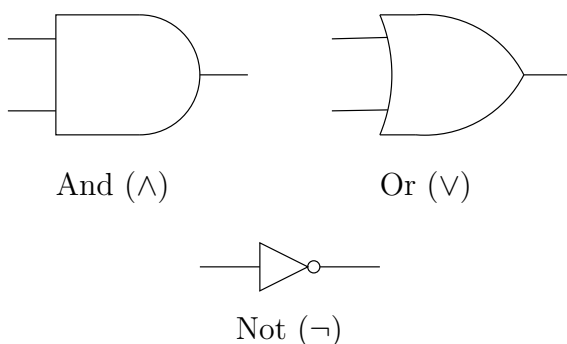


Figure 2.3: The connection of two transistors in parallel provides an implementation of the *or* operator.

As a sort of graphical shorthand, electronics engineers use the symbols below to indicate and-gates, or-gates & not-gates (better known as negators).



An and-gate has two transistors inside it that are wired in series – if both the inputs are energized the output will be too. An or-gate has two transistors in parallel inside it. Not-gates involve magic – when their input is not on, their output *is* and vice versa.

Using this graphical “language” one can make schematic representations of logical expressions. Some find that tracing such diagrams makes understanding the structure of a Boolean expression easier. For example, in Figure 2.4 we illustrate 2 of the possible ways that the conjunction of four predicate variables can be parenthesized. In fact, when a multitude of predicates are joined by the same connective, the way in which the expression is parenthesized is unimportant, thus one often sees a further shorthand — gates with more than 2 inputs.

A common task for an electronics designer is to come up with a digital logic circuit having a prescribed input/output table. Note that an input/output table for a logic circuit is entirely analogous with a truth table for a compound sentence in Logic — except that we use 0’s and 1’s rather than T’s and ϕ ’s.

Suppose that we wanted to design a circuit that would have the following input/output table.

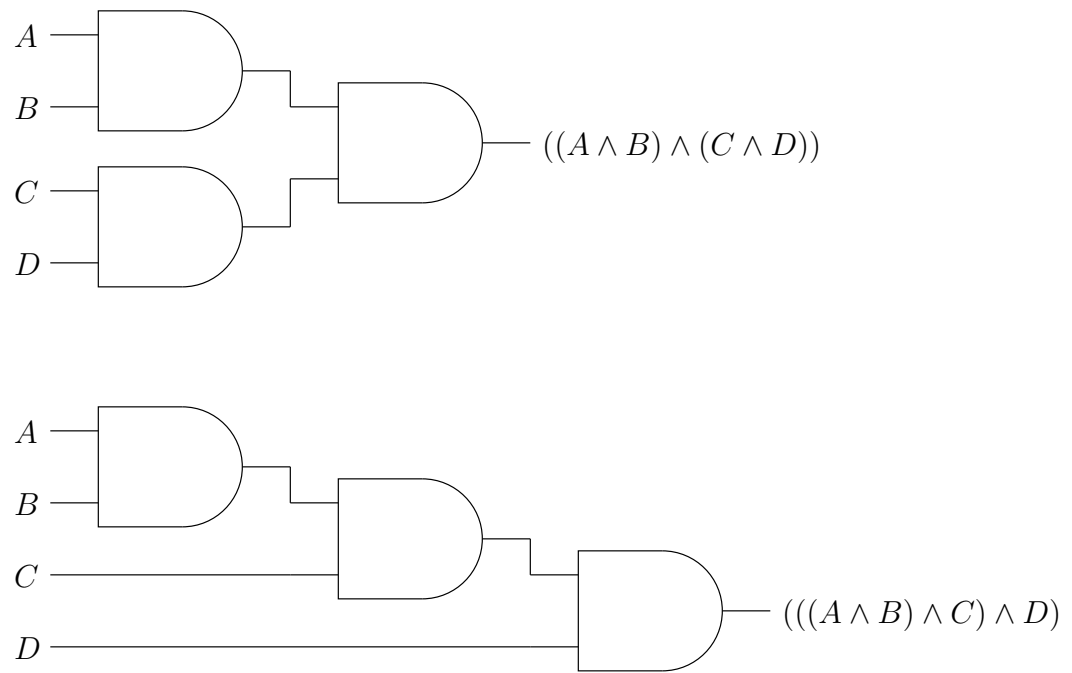
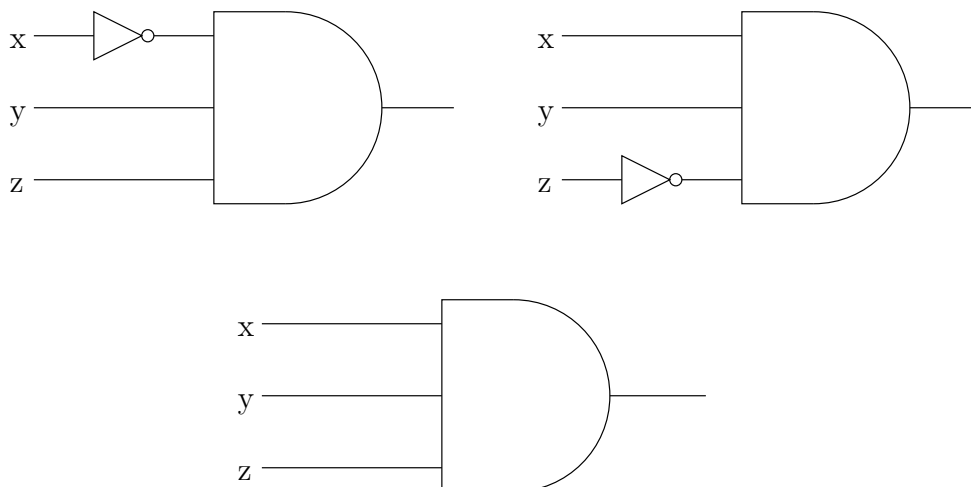


Figure 2.4: Two of the possible ways to parenthesize the conjunction of four statement variables – expressed as digital logic circuits.

x	y	z	out
0	0	0	0
0	0	1	0
0	1	0	0
0	1	1	1
1	0	0	0
1	0	1	0
1	1	0	1
1	1	1	1

A systematic method for accomplishing such a design task involves a notion called *disjunctive normal form*. A Boolean expression is in disjunctive normal form if it consists of the disjunction of one or more statements, each of which consists entirely of conjunctions of predicate variables and/or their negations. In other words, the *or* of a bunch of *ands*. In terms of digital logic circuits, the *ands* we're talking about are called *recognizers*. For example, the following 3-input and-gates recognize the input states in the 4th, 7th and 8th rows of the i/o table above. (These are the rows where the output is supposed to be 1.)



In Figure 2.5 we illustrate how to create a circuit whose i/o table is as above using these recognizers.

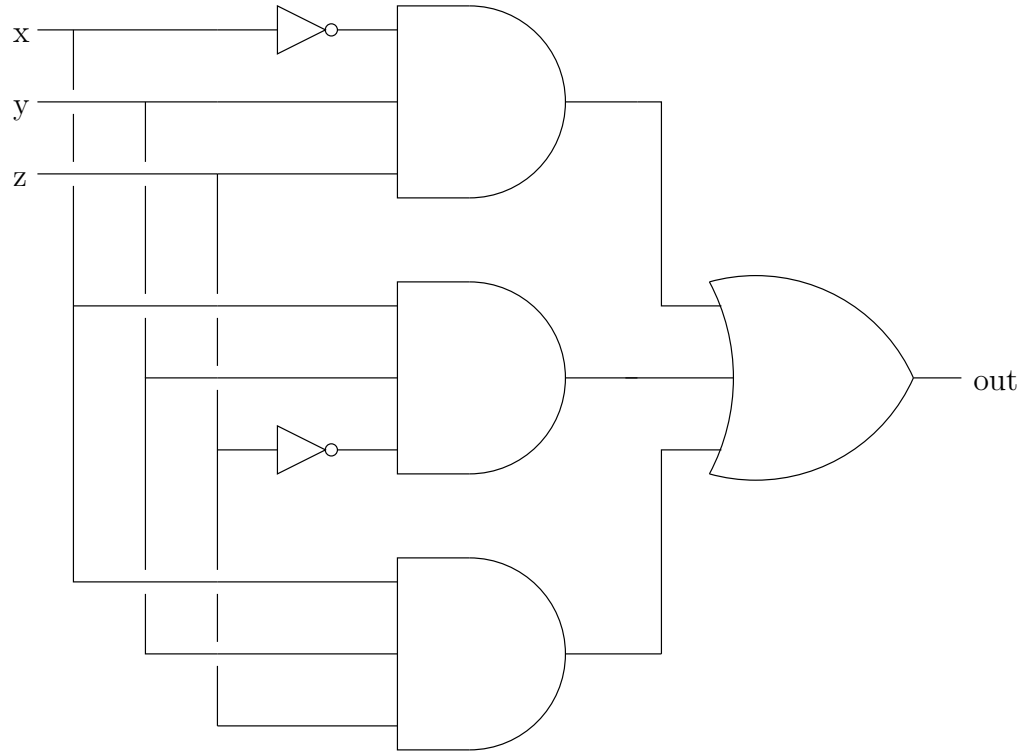


Figure 2.5: A digital logic circuit built using disjunctive normal form. The output of this circuit is $(\neg x \wedge y \wedge z) \vee (x \wedge y \wedge \neg z) \vee (x \wedge y \wedge z)$.

Exercises — 2.1

1. Design a digital logic circuit (using and, or & not gates) that implements an exclusive or.
2. Consider the sentence “This is a sentence which does not refer to itself.” which was given in the beginning of this chapter as an example. Is this sentence a statement? If so, what is its truth value?
3. Consider the sentence “This sentence is false.” Is this sentence a statement?
4. Complete truth tables for each of the sentences $(A \wedge B) \vee C$ and $A \wedge (B \vee C)$. Does it seem that these sentences have the same logical content?
5. There are two other logical connectives that are used somewhat less commonly than \vee and \wedge . These are the Scheffer stroke and the Peirce arrow – written $|$ and \downarrow , respectively — they are also known as NAND and NOR.

The truth tables for these connectives are:

A	B	$A B$		A	B	$A \downarrow B$
T	T	ϕ		T	T	ϕ
T	ϕ	T	and	T	ϕ	ϕ
ϕ	T	T		ϕ	T	ϕ
ϕ	ϕ	T		ϕ	ϕ	T

Find an expression for $(A \wedge \neg B) \vee C$ using only these new connectives (as well as negation and the variable symbols themselves).

6. The famous logician Raymond Smullyan devised a family of logical puzzles around a fictitious place he called “the Island of Knights and Knaves.” The inhabitants of the island are either knaves, who always make false statements, or knights, who always make truthful statements.

In the most famous knight/knave puzzle, you are in a room which has only two exits. One leads to certain death and the other to freedom. There are two individuals in the room, and you know that one of them is a knight and the other is a knave, but you don't know which. Your challenge is to determine the door which leads to freedom by asking a single question.

2.2 Implication

Suppose a mother makes the following statement to her child: “If you finish your peas, you’ll get dessert.”

This is a compound sentence made up of the two simpler sentences $P =$ “You finish your peas” and $D =$ “You’ll get dessert.” It is an example of a type of compound sentence called a *conditional*. Conditionals are if-then type statements. In ordinary language the word “then” is often elided (as is the case with our example above). Another way of phrasing the “If P then D.” relationship is to use the word “implies” — although it would be a rather uncommon mother who would say “Finishing your peas implies that you will receive dessert.”

As was the case in the previous section, there are four possible situations and we must consider each to decide the truth/falsity of this conditional statement. The peas may or may not be finished, and independently, the dessert may or may not be proffered.

Suppose the child finishes the peas and the mother comes across with the dessert. Clearly, in this situation the mother’s statement was true. On the other hand, if the child finishes the hated peas and yet does not receive a treat, it is just as obvious that the mother has lied! What do we say about the mother’s veracity in the case that the peas go unfinished? Here, Mom gets a break. She can either hold firm and deliver no dessert, or she can be a softy and give out unearned sweets – in either case, we can’t accuse her of telling a falsehood. The statement she made had to do *only* with the eventualities following total pea consumption, she said nothing about what happens if the peas go uneaten.

A conditional statement’s components are called the *antecedent* (this is the “if” part, as in “finish your peas”) and the *consequent* (this is the “then” part, as in “get dessert”). The discussion in the last paragraph was intended to make the point that when the antecedent is false, we should consider the

conditional to be true. Conditionals that are true because their antecedents are false are said to be *vacuously true*. The conditional involving an antecedent A and a consequent B is expressed symbolically using an arrow: $A \implies B$. Here is a truth table for this connective.

A	B	$A \implies B$
T	T	T
T	ϕ	ϕ
ϕ	T	T
ϕ	ϕ	T

Exercise. Note that this truth table is similar to the truth table for $A \vee B$ in that there is only a single row having a ϕ in the last column. For $A \vee B$ the ϕ occurs in the 4th row and for $A \implies B$ it occurs in the 2nd row. This suggests that by suitably modifying things (replacing A or B by their negations) we could come up with an “or” statement that had the same meaning as the conditional. Try it!

It is fairly common that conditionals are used to express threats, as in the peas/dessert example. Another common way to express a threat is to use a disjunction – “Finish your peas, or you won’t get dessert.” If you’ve been paying attention (and did the last exercise), you will notice that this is *not* the disjunction that should have the same meaning as the original conditional. There is probably no mother on Earth who would say “Don’t finish your peas, or you get dessert!” to her child (certainly not if she expects to be understood). So what’s going on here?

The problem is that “Finish your peas, or you won’t get dessert.” has the same logical content as “If you get dessert then you finished your peas.” (Notice that the roles of the antecedent and consequent have been switched.) And, while this last sentence sounds awkward, it is probably a more accurate reflection of what the mother intended. The problem *really* is that people are incredibly sloppy with their conditional statements! A lot of people secretly

want the 3rd row of the truth table for \implies to have a ϕ in it, and it simply doesn't! The operator that results if we do make this modification is called the biconditional, and is expressed in English using the phrase “if and only if” (which leads mathematicians to the abbreviation “iff” much to the consternation of spell-checking programs everywhere). The biconditional is denoted using an arrow that points both ways. Its truth table follows.

A	B	$A \iff B$
T	T	T
T	ϕ	ϕ
ϕ	T	ϕ
ϕ	ϕ	T

Please note, that while we like to strive for precision, we do not necessarily recommend the use of phrases such as “You will receive dessert if, and only if, you finish your peas.” with young children.

Since conditional sentences are often confused with the sentence that has the roles of antecedent and consequent reversed, this switched-around sentence has been given a name: it is the *converse* of the original statement. Another conditional that is distinct from (but related to) a given conditional is its *inverse*. This sort of sentence probably had to be named because of a very common misconception, many people think that the way to negate an if-then proposition is to negate its parts. Algebraically, this looks reasonable – sort of a distributive law for logical negation over implications – $\neg(A \implies B) = \neg A \implies \neg B$. Sadly, this reasonable looking assertion can't possibly be true; since implications have just one ϕ in a truth table, the negation of an implication must have three – but the statement with the \neg 's on the *parts* of the implication is going to only have a single ϕ in *its* truth table.

To recap, the converse of an implication has the pieces (antecedent and consequent) switched about. The inverse of an implication has the pieces negated. Neither of these is the same as the original implication. Oddly, this

is one of those times when two wrongs *do* make a right. If you start with an implication, form its converse, then take the inverse of that, you get a statement having exactly the same logical meaning as the original. This new statement is called the *contrapositive*.

This information is displayed in Table 2.1

	converses		
	↔		
	$A \implies B$	$B \implies A$	
inverses	$\neg A \implies \neg B$	$\neg B \implies \neg A$	
	↕		

Table 2.1: The relationship between a conditional statement, its converse, its inverse and its contrapositive.

One final piece of advice about conditionals: don't confuse logical if-then relationships with causality. Many of the if-then sentences we run into in ordinary life describe cause and effect: "If you cut the green wire the bomb will explode." (Okay, that one is an example from the ordinary life of a bomb squad technician, but ...) It is usually best to think of the if-then relationships we find in Logic as divorced from the flow of time, the fact that $A \implies B$ is logically the same as $\neg A \vee B$ lends credence to this point of view.

Exercises — 2.2

1. The transitive property of equality says that if $a = b$ and $b = c$ then $a = c$. Does the implication arrow satisfy a transitive property? If so, state it.
2. Complete truth tables for the compound sentences $A \implies B$ and $\neg A \vee B$.
3. Complete a truth table for the compound sentence $A \implies (B \implies C)$ and for the sentence $(A \implies B) \implies C$. What can you conclude about conditionals and the associative property?
4. Determine a sentence using the *and* connector (\wedge) that gives the negation of $A \implies B$.
5. Rewrite the sentence “Fix the toilet or I won’t pay the rent!” as a conditional.
6. Why is it that the sentence “If pigs can fly, I am the king of Mesopotamia.” true?
7. Express the statement $A \implies B$ using the Peirce arrow and/or the Scheffer stroke. (See Exercise 5 in the previous section.)
8. Find the contrapositives of the following sentences.
 - (a) If you can’t do the time, don’t do the crime.
 - (b) If you do well in school, you’ll get a good job.
 - (c) If you wish others to treat you in a certain way, you must treat others in that fashion.
 - (d) If it’s raining, there must be clouds.
 - (e) If $a_n \leq b_n$, for all n and $\sum_{n=0}^{\infty} b_n$ is a convergent series, then $\sum_{n=0}^{\infty} a_n$ is a convergent series.

9. What are the converse and inverse of “If you watch my back, I’ll watch your back.”?
10. The integral test in Calculus is used to determine whether an infinite series converges or diverges: Suppose that $f(x)$ is a positive, decreasing, real-valued function with $\lim_{x \rightarrow \infty} f(x) = 0$, if the improper integral $\int_0^{\infty} f(x)$ has a finite value, then the infinite series $\sum_{n=1}^{\infty} f(n)$ converges. The integral test should be envisioned by letting the series correspond to a right-hand Riemann sum for the integral, since the function is decreasing, a right-hand Riemann sum is an underestimate for the value of the integral, thus

$$\sum_{n=1}^{\infty} f(n) < \int_0^{\infty} f(x).$$

Discuss the meanings of and (where possible) provide justifications for the inverse, converse and contrapositive of the conditional statement in the integral test.

11. On the Island of Knights and Knaves (see page 67) you encounter two individuals named Locke and Demosthenes.

Locke says, “Demosthenes is a knave.”

Demosthenes says “Locke and I are knights.”

Who is a knight and who a knave?

2.3 Logical equivalences

Some logical statements are “the same.” For example, in the last section, we discussed the fact that a conditional and its contrapositive have the same logical content. Wouldn’t we be justified in writing something like the following?

$$A \implies B = \neg B \implies \neg A$$

Well, one pretty serious objection to doing that is that the equals sign ($=$) has already got a job; it is used to indicate that two numerical quantities are the same. What we’re doing here is really sort of a different thing! Nevertheless, there is a concept of “sameness” between certain compound statements, and we need a symbolic way of expressing it. There are two notations in common use. The notation that seems to be preferred by logicians is the biconditional (\iff). The notation we’ll use in the rest of this book is an equals sign with a bit of extra decoration on it (\cong).

Thus we can either write

$$(A \implies B) \iff (\neg B \implies \neg A)$$

or

$$A \implies B \cong \neg B \implies \neg A.$$

I like the latter, but use whichever form you like – no one will have any problem understanding either.

The formal definition of *logical equivalence*, which is what we’ve been describing, is this: two compound sentences are logically equivalent if in a truth table (that contains all possible combinations of the truth values of the predicate variables in its rows) the truth values of the two sentences are equal in every row.

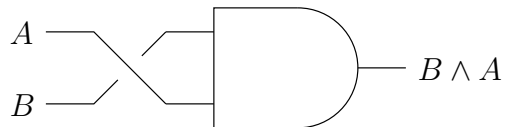
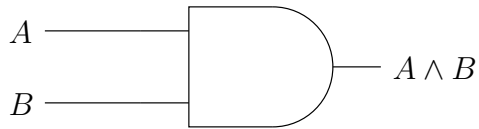
Exercise. Consider the two compound sentences $A \vee B$ and $A \vee (\neg A \wedge B)$. There are a total of 2 predicate variables between them, so a truth table with 4 rows will suffice. Fill out the missing entries in the truth table and determine whether the statements are equivalent.

A	B	$A \vee B$	$A \vee (\neg A \wedge B)$
T	T		
T	ϕ		
ϕ	T		
ϕ	ϕ		

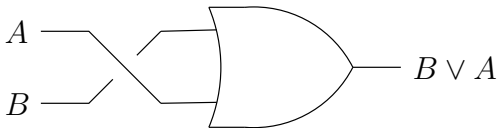
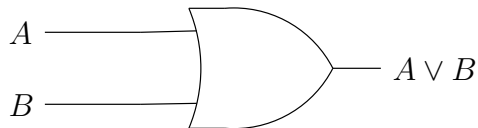
One could, in principle, verify all logical equivalences by filling out truth tables. Indeed, in the exercises for this section we will ask you to develop a certain facility at this task. While this activity can be somewhat fun, and many of my students want the filling-out of truth tables to be a significant portion of their midterm exam, you will probably eventually come to find it somewhat tedious. A slightly more mature approach to logical equivalences is this: use a set of basic equivalences – which themselves may be verified via truth tables – as the basic *rules* or *laws* of logical equivalence, and develop a strategy for converting one sentence into another using these rules. This process will feel very familiar, it is like “doing” algebra, but the rules one is allowed to use are subtly different.

First we have the *commutative laws*, one each for conjunction and disjunction. It’s worth noting that there *isn’t* a commutative law for implication.

The commutative property of conjunction says that $A \wedge B \cong B \wedge A$. This is quite an apparent statement from the perspective of linguistics. Surely it’s the same thing to say “the weather is cold and snowy” as it is to say “the weather is snowy and cold.” This commutative property is also clear from the perspective of digital logic circuits.

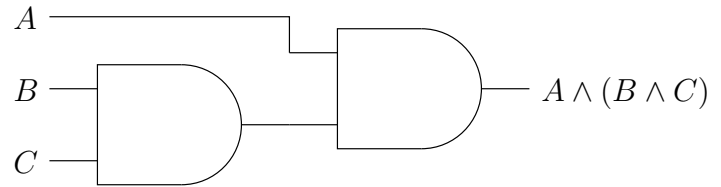
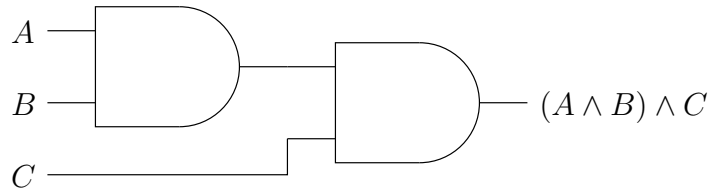


The commutative property of disjunctions is equally transparent from the perspective of a circuit diagram.

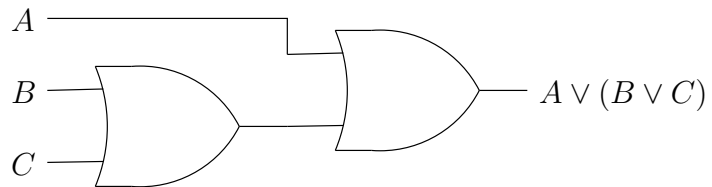
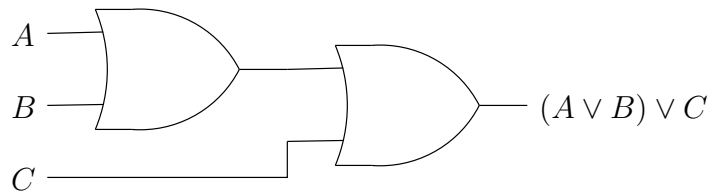


The *associative laws* also have something to do with what order operations are done. One could think of the difference in the following terms: Commutative properties involve spatial or physical order and the associative properties involve temporal order. The associative law of addition could be used to say we'll get the same result if we add 2 and 3 first, then add 4, or if we add 2 to the sum of 3 and 4 (i.e. that $(2+3)+4$ is the same as $2+(3+4)$.) Note that physically, the numbers are in the same order (2 then 3 then 4) in both expressions but that the parentheses indicate a precedence in *when* the plus signs are evaluated.

The associative law of conjunction states that $A \wedge (B \wedge C) \cong (A \wedge B) \wedge C$. In visual terms, this means the following two circuit diagrams are equivalent.



The associative law of disjunction states that $A \vee (B \vee C) \cong (A \vee B) \vee C$. Visually, this looks like:



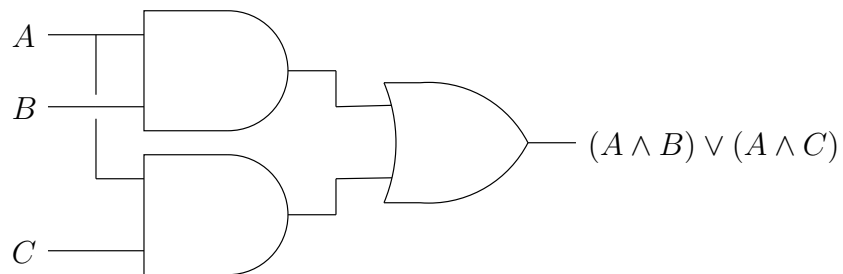
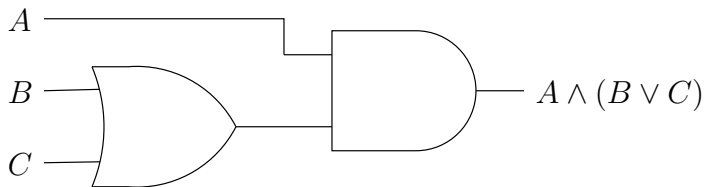
Exercise. In a situation where both associativity and commutativity pertain the symbols involved can appear in any order and with any reasonable paren-

thesization. In how many different ways can the sum $2 + 3 + 4$ be expressed? Only consider expression that are fully parenthesized.

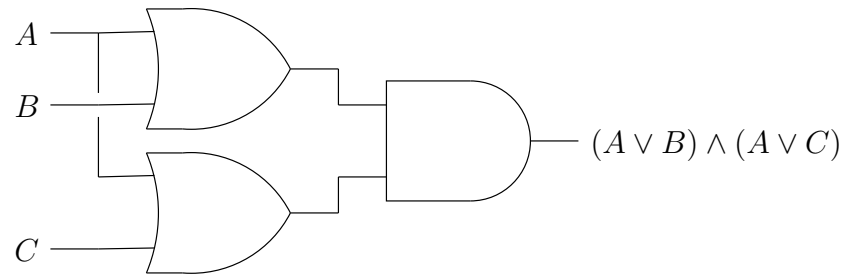
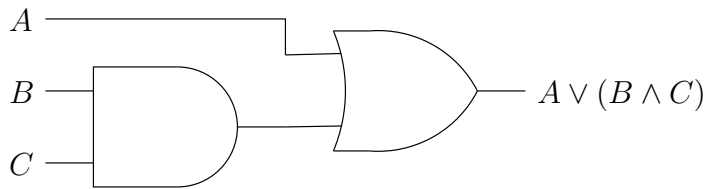
The next type of basic logical equivalences we'll consider are the so-called *distributive laws*. Distributive laws involve the interaction of two operations, when we distribute multiplication over a sum, we effectively replace one instance of an operand *and the associated operator*, with two instances, as is illustrated below.

$$2 * (3 + 4) = (2 * 3) + (2 * 4)$$

The logical operators \wedge and \vee each distribute over the other. Thus we have the distributive law of conjunction over disjunction, which is expressed in the equivalence $A \wedge (B \vee C) \cong (A \wedge B) \vee (A \wedge C)$ and in the following digital logic circuit diagram.



We also have the distributive law of disjunction over conjunction which is given by the equivalence $A \vee (B \wedge C) \cong (A \vee B) \wedge (A \vee C)$ and in the circuit diagram:



Traditionally, the laws we've just stated would be called *left*-distributive laws and we would also need to state that there are *right*-distributive laws that apply. Since, in the current setting, we have already said that the commutative law is valid, this isn't really necessary.

Exercise. State the right-hand versions of the distributive laws.

The next set of laws we'll consider come from trying to figure out what the distribution of a minus sign over a sum $-(x + y) = -x + -y$ should correspond to in Boolean algebra. At first blush one might assume the analogous thing in Boolean algebra would be something like $\neg(A \wedge B) \cong \neg A \wedge \neg B$, but we can easily dismiss this by looking at a truth table.

A	B	$\neg(A \wedge B)$	$\neg A \wedge \neg B$
T	T	ϕ	ϕ
T	ϕ	T	ϕ
ϕ	T	T	ϕ
ϕ	ϕ	T	T

What actually works is a set of rules known as DeMorgan's laws, which basically say that you distribute the negative sign but you also must change the operator. As logical equivalences, DeMorgan's laws are

$$\neg(A \wedge B) \cong \neg A \vee \neg B$$

and

$$\neg(A \vee B) \cong \neg A \wedge \neg B.$$

In ordinary arithmetic there are two notions of "inverse." The *negative* of a number is known as its additive inverse and the *reciprocal* of a number is its multiplicative inverse. These notions lead to a couple of equations,

$$x + -x = 0$$

and

$$x \cdot \frac{1}{x} = 1.$$

Boolean algebra only has one "inverse" concept, the denial of a predicate (i.e. logical negation), but the equations above have analogues, as do the symbols 0 and 1 that appear in them. First, consider the Boolean expression $A \vee \neg A$. This is the logical *or* of a statement and its exact opposite; when one is true the other is false and vice versa. But, the disjunction $A \vee \neg A$, is always true! We use the symbol t (which stands for *tautology*) to represent a compound sentence whose truth value is always true. A tautology (t) is to

Boolean algebra something like a zero (0) is to arithmetic. Similar thinking about the Boolean expression $A \wedge \neg A$ leads to the definition of the symbol c (which stands for *contradiction*) to represent a sentence that is always false. The rules we have been discussing are known as *complementarity laws*:

$$A \vee \neg A \cong t \quad \text{and} \quad A \wedge \neg A \cong c$$

Now that we have the special logical sentences represented by t and c we can present the so-called *identity laws*, $A \wedge t \cong A$ and $A \vee c \cong A$. If you “and” a statement with something that is always true, this new compound has the exact same truth values as the original. If you “or” a statement with something that is always false, the new compound statement is also unchanged from the original. Thus performing a conjunction with a tautology has no effect – sort of like multiplying by 1. Performing a disjunction with a contradiction also has no effect – this is somewhat akin to adding 0.

The number 0 has a special property: $0 \cdot x = 0$ is an equation that holds no matter what x is. This is known as a domination property. Note that there isn’t a dominance rule that involves 1. On the Boolean side, *both* the symbols t and c have related domination rules.

$$A \vee t \cong t \quad \text{and} \quad A \wedge c \cong c$$

In mathematics the word *idempotent* is used to describe situations where a power of a thing may be equal to that thing. For example, because $(-1)^3 = -1$, we say that -1 is an idempotent. Both of the Boolean operations have idempotence relations that just always work (regardless of the operand). In ordinary algebra idempotents are very rare (0, 1 and -1 are the only ones that come to mind), but in Boolean algebra *every* statement is equivalent to its square – where the square of A can be interpreted either as $A \wedge A$ or as $A \vee A$.

$$A \vee A \cong A \quad \text{and} \quad A \wedge A \cong A$$

There are a couple of properties of the logical negation operator that should be stated, though probably they seem self-evident. If you form the denial of a denial, you come back to the same thing as the original; also the symbols c and t are negations of one another.

$$\neg(\neg A) \cong A \quad \text{and} \quad \neg t \cong c$$

Finally, we should mention a really strange property, called *absorption*, which states that the expressions $A \wedge (A \vee B)$ and $A \vee (A \wedge B)$ don't actually have anything to do with B at all! Both of the preceding statements are equivalent to A .

$$A \wedge (A \vee B) \cong A \quad \text{and} \quad A \vee (A \wedge B) \cong A$$

In Table 2.2, we have collected all of these basic logical equivalences in one place.

	Conjunctive version	Disjunctive version	Algebraic analog
Commutative laws	$A \wedge B \cong B \wedge A$	$A \vee B \cong B \vee A$	$2 + 3 = 3 + 2$
Associative laws	$A \wedge (B \wedge C)$ $\cong (A \wedge B) \wedge C$	$A \vee (B \vee C)$ $\cong (A \vee B) \vee C$	$2 + (3 + 4)$ $= (2 + 3) + 4$
Distributive laws	$A \wedge (B \vee C) \cong$ $(A \wedge B) \vee (A \wedge C)$	$A \vee (B \wedge C) \cong$ $(A \vee B) \wedge (A \vee C)$	$2 \cdot (3 + 4)$ $= (2 \cdot 3 + 2 \cdot 4)$
DeMorgan's laws	$\neg(A \wedge B)$ $\cong \neg A \vee \neg B$	$\neg(A \vee B)$ $\cong \neg A \wedge \neg B$	none
Double negation	$\neg(\neg A) \cong A$	same	$-(-2) = 2$
Complementarity	$A \wedge \neg A \cong c$	$A \vee \neg A \cong t$	$2 + (-2) = 0$
Identity laws	$A \wedge t \cong A$	$A \vee c \cong A$	$7 + 0 = 7$
Domination	$A \wedge c \cong c$	$A \vee t \cong t$	$7 \cdot 0 = 0$
Idempotence	$A \wedge A \cong A$	$A \vee A \cong A$	$1 \cdot 1 = 1$
Absorption	$A \wedge (A \vee B) \cong A$	$A \vee (A \wedge B) \cong A$	none

Table 2.2: Basic logical equivalences.

Exercises — 2.3

1. There are 3 operations used in basic algebra (addition, multiplication and exponentiation) and thus there are potentially 6 different distributive laws. State all 6 “laws” and determine which 2 are actually valid. (As an example, the distributive law of addition over multiplication would look like $x + (y \cdot z) = (x + y) \cdot (x + z)$, this isn’t one of the true ones.)
2. Use truth tables to verify or disprove the following logical equivalences.
 - (a) $(A \wedge B) \vee B \cong (A \vee B) \wedge B$
 - (b) $A \wedge (B \vee \neg A) \cong A \wedge B$
 - (c) $(A \wedge \neg B) \vee (\neg A \wedge \neg B) \cong (A \vee \neg B) \wedge (\neg A \vee \neg B)$
 - (d) The absorption laws.
3. Draw pairs of related digital logic circuits that illustrate DeMorgan’s laws.
4. Find the negation of each of the following and simplify as much as possible.
 - (a) $(A \vee B) \iff C$
 - (b) $(A \vee B) \implies (A \wedge B)$
5. Because a conditional sentence is equivalent to a certain disjunction, and because DeMorgan’s law tells us that the negation of a disjunction is a conjunction, it follows that the negation of a conditional is a conjunction. Find denials (the negation of a sentence is often called its “denial”) for each of the following conditionals.
 - (a) “If you smoke, you’ll get lung cancer.”

- (b) “If a substance glitters, it is not necessarily gold.”
 - (c) “If there is smoke, there must also be fire.”
 - (d) “If a number is squared, the result is positive.”
 - (e) “If a matrix is square, it is invertible.”
6. The so-called “ethic of reciprocity” is an idea that has come up in many of the world’s religions and philosophies. Below are statements of the ethic from several sources. Discuss their logical meanings and determine which (if any) are logically equivalent.
- (a) “One should not behave towards others in a way which is disagreeable to oneself.” Mencius VII.A.4 (Hinduism)
 - (b) “None of you [truly] believes until he wishes for his brother what he wishes for himself.” Number 13 of Imam “Al-Nawawi’s Forty Hadiths.” (Islam)
 - (c) “And as ye would that men should do to you, do ye also to them likewise.” Luke 6:31, King James Version. (Christianity)
 - (d) “What is hateful to you, do not to your fellow man. This is the law: all the rest is commentary.” Talmud, Shabbat 31a. (Judaism)
 - (e) “An it harm no one, do what thou wilt” (Wicca)
 - (f) “What you would avoid suffering yourself, seek not to impose on others.” (the Greek philosopher Epictetus – first century A.D.)
 - (g) “Do not do unto others as you expect they should do unto you. Their tastes may not be the same.” (the Irish playwright George Bernard Shaw – 20th century A.D.)
7. You encounter two natives of the land of knights and knaves. Fill in an explanation for each line of the proofs of their identities.

- (a) Natasha says, “Boris is a knave.”
 Boris says, “Natasha and I are knights.”

Claim: Natasha is a knight, and Boris is a knave.

Proof: If Natasha is a knave, then Boris is a knight.
 If Boris is a knight, then Natasha is a knight.
 Therefore, if Natasha is a knave, then Natasha is a knight.
 Hence Natasha is a knight.
 Therefore, Boris is a knave.

Q.E.D.

- (b) Bonaparte says “I am a knight and Wellington is a knave.”
 Wellington says “I would tell you that B is a knight.”

Claim: Bonaparte is a knight and Wellington is a knave.

Proof: Either Wellington is a knave or Wellington is a knight.
 If Wellington is a knight it follows that Bonaparte is a knight.
 If Bonaparte is a knight then Wellington is a knave.
 So, if Wellington is a knight then Wellington is a knave (which is impossible!)
 Thus, Wellington is a knave.
 Since Wellington is a knave, his statement “I would tell you that Bonaparte is a knight” is false.
 So Wellington would in fact tell us that Bonaparte is a knave.
 Since Wellington is a knave we conclude that Bonaparte is a knight.
 Thus Bonaparte is a knight and Wellington is a knave (as

claimed).

Q.E.D.

2.4 Two-column proofs

If you've ever spent much time trying to check someone else's work in solving an algebraic problem, you'd probably agree that it would be a help to know what they were *trying* to do in each step. Most people have this fairly vague notion that they're allowed to "do the same thing on both sides" and they're allowed to simplify the sides of the equation separately – but more often than not, several different things get done on a given line, mistakes get made, and it can be nearly impossible to figure out what went wrong and where.

Now, after all, the beauty of math is supposed to lie in its crystal clarity, so this sort of situation is really unacceptable. It may be an impossible goal to get "the average Joe" to perform algebraic manipulations with clarity, but those of us who aspire to become mathematicians must certainly hold ourselves to a higher standard. Two-column proofs are usually what is meant by a "higher standard" when we are talking about relatively mechanical manipulations – like doing algebra, or more to the point, proving logical equivalences. Now don't despair! You will not, in a mathematical career, be expected to provide two-column proofs very often. In fact, in more advanced work one tends to not give *any* sort of proof for a statement that lends itself to a two-column approach. But, if you find yourself writing "As the reader can easily verify, Equation 17 holds. . ." in a paper, or making some similar remark to your students, you are *morally obligated* to being able to produce a two-column proof.

So what, exactly, is a two-column proof? In the left column you show your work, being careful to go one step at a time. In the right column you provide a justification for each step.

We're going to go through a couple of examples of two-column proofs in the context of proving logical equivalences. One thing to watch out for: if you're trying to prove a given equivalence, and the first thing you write down is that very equivalence, *it's wrong!* This would constitute the logical

error known as “begging the question” also known as “circular reasoning.” It’s clearly not okay to try to demonstrate some fact by first *asserting the very same fact*. Nevertheless, there is (for some unknown reason) a powerful temptation to do this very thing. To avoid making this error, we will not put any equivalences on a single line. Instead we will start with one side or the other of the statement to be proved, and modify it using known rules of equivalence, until we arrive at the other side.

Without further ado, let’s provide a proof of the equivalence $A \wedge (B \vee \neg A) \cong A \wedge B$.⁶

$$\begin{aligned}
 & A \wedge (B \vee \neg A) && \text{distributive law} \\
 \cong & (A \wedge B) \vee (A \wedge \neg A) && \text{complementarity} \\
 \cong & (A \wedge B) \vee c && \text{identity law} \\
 \cong & (A \wedge B)
 \end{aligned}$$

We have assembled a nice, step-by-step sequence of equivalences – each justified by a known law – that begins with the left-hand side of the statement to be proved and ends with the right-hand side. That’s an irrefutable proof!

In the next example we’ll highlight a slightly sloppy habit of thought that tends to be problematic. People usually (at first) associate a direction with the basic logical equivalences. This is reasonable for several of them because one side is markedly simpler than the other. For example, the domination rule would normally be used to replace a part of a statement that looked like “ $A \wedge c$ ” with the simpler expression “ c ”. There is a certain amount of strategization necessary in doing these proofs, and I usually advise people to start with the more complicated side of the equivalence to be proved. It

⁶This equivalence should have been verified using truth tables in the exercises from the previous section.

just feels right to work in the direction of making things simpler, but there are times when one has to take one step back before proceeding two steps forward. . .

Let's have a look at another equivalence: $A \wedge (B \vee C) \cong (A \wedge (B \vee C)) \vee (A \wedge C)$. There are many different ways in which valid steps can be concatenated to convert one side of this equivalence into the other, so a subsidiary goal is to find a proof that uses the least number of steps. Following my own advice, I'll start with the right-hand side of this one.

$$\begin{aligned}
 & (A \wedge (B \vee C)) \vee (A \wedge C) && \\
 & && \text{distributive law} \\
 \cong & ((A \wedge B) \vee (A \wedge C)) \vee (A \wedge C) && \\
 & && \text{associative law} \\
 \cong & (A \wedge B) \vee ((A \wedge C) \vee (A \wedge C)) && \\
 & && \text{idempotence} \\
 \cong & (A \wedge B) \vee (A \wedge C) && \\
 & && \text{distributive law} \\
 \cong & A \wedge (B \vee C) &&
 \end{aligned}$$

Note that in the example we've just done, the two applications of the distributive law go in opposite directions as far as their influence on the complexity of the expressions are concerned.

Exercises — 2.4

Write two-column proofs that verify each of the following logical equivalences.

1. $A \vee (A \wedge B) \cong A \wedge (A \vee B)$

2. $(A \wedge \neg B) \vee A \cong A$

3. $A \vee B \cong A \vee (\neg A \wedge B)$

4. $\neg(A \vee \neg B) \vee (\neg A \wedge \neg B) \cong \neg A$

5. $A \cong A \wedge ((A \vee \neg B) \vee (A \vee B))$

6. $(A \wedge \neg B) \wedge (\neg A \vee B) \cong c$

7. $A \cong A \wedge (A \vee (A \wedge (B \vee C)))$

8. $\neg(A \wedge B) \wedge \neg(A \wedge C) \cong \neg A \vee (\neg B \wedge \neg C)$

2.5 Quantified statements

All of the statements discussed in the previous sections were of the “completely unambiguous” sort; that is, they didn’t have any *unknowns* in them. As a reader of this text, it’s a sure bet that you’ve mastered Algebra and are firmly convinced of the utility of x and y . Admittedly, we’ve used variables to refer to sentences (or sentence fragments) themselves, but we’ve said that sentences that had variables *in them* were ambiguous and didn’t even deserve to be called logical statements. The notion of *quantification* allows us to use the power of variables within a sentence without introducing ambiguity.

Consider the sentence “There are exactly 7 odd primes less than 20.” This sentence has some kind of ambiguity in it (because it doesn’t mention the primes explicitly) and yet it certainly seems to have a definite truth value! The reason its truth value is known (by the way, it is T) is that the sentence is quantified. “ X is an odd prime less than 20.” is an ambiguous sentence, but “There are exactly 7 distinct X ’s that are odd primes less than 20.” is not. This example represents a fairly unusual form of quantification. Usually, we take away the ambiguity of a sentence having a variable in it by asserting one of two levels of quantification: “this is true at least once” or “this is always true”. We’ve actually seen the symbols (\exists and \forall) for these concepts already (in Section 1.3).

An *open sentence* is one that has variables in it. We represent open sentences using a sort of functional notation to show what variables are in them.

Examples:

- i) $P(x) = “2^{2^x} + 1 \text{ is a prime.}”$
- ii) $Q(x, y) = “x \text{ is prime or } y \text{ is a divisor of } x.”$
- iii) $L(f, c, l) = “\text{The function } f \text{ has limit } l \text{ at } c, \text{ if and only if, for every positive number } \epsilon, \text{ there is a positive number } \delta \text{ such that whenever}”$

$|x - c| < \delta$ it follows that $|f(x) - l| < \epsilon$.”

That last example certainly is a doozy! At first glance it would appear to have more than three variables in it, and indeed it does! In order of appearance, we have f , l , c , ϵ , δ and x – the last three variables that appear (ϵ , δ and x) are said to be *bound*. A variable in an open sentence is bound if it is in the scope of a quantifier. Bound variables don’t need to be mentioned in the argument list of the sentence. Unfortunately, when sentences are given in natural languages the quantification status of a variable may not be clear. For example in the third sentence above, the variable δ is easily seen to be in the scope of the quantifier \exists because of the words “there is a positive number” that precede it. Similarly, ϵ is universally quantified (\forall) because the phrase “for every positive number” appears before it. What is the status of x ? Is it really bound? The answers to such questions may not be clear at first, but after some thought you should be able to decide that x is universally quantified.

Exercise. *What word in example iii) indicates that x is in the scope of a \forall quantifier?*

It is not uncommon, in advanced Mathematics, to encounter compound sentences involving dozens of variables and 4 or 5 levels of quantification. Such sentences seem hopelessly complicated at first sight – the key to understanding them is to determine each variable’s quantification status explicitly and to break things down into simpler sub-parts.

For instance, in understanding example iii) above, it might be useful to define some new open sentences:

$$D(x, c, \delta) = “|x - c| < \delta”$$

$$E(f, x, l, \epsilon) = “|f(x) - l| < \epsilon”$$

Furthermore, it’s often handy to replace an awkward phrase (such as “the limit of f at c is l ”) with symbols when possible.

Example iii) now looks like

$$\lim_{x \rightarrow c} f(x) = l \iff \forall \epsilon > 0 \exists \delta > 0 \forall x D(x, c, \delta) \implies E(f, x, l, \epsilon).$$

The sentence $D(x, c, \delta)$ is usually interpreted as saying that “ x is close to c ” (where δ tells you *how* close.) The sentence $E(f, x, l, \epsilon)$ could be expressed informally as “ $f(x)$ is close to l ” (again, ϵ serves to make the word “close” more exact).

It’s instructive to write this sentence one last time, *completely* in symbols and without the abbreviations we created for saying that x is near c and $f(x)$ is near l :

$$\lim_{x \rightarrow c} f(x) = l \iff \forall \epsilon > 0 \exists \delta > 0 \forall x (|x - c| < \delta) \implies (|f(x) - l| < \epsilon).$$

It would not be unfair to say that developing the facility to read, and understand, this hieroglyph (and others like it) constitutes the first several weeks of a course in Real Analysis.

Let us turn back to another of the examples (of an open sentence) from the beginning of this section. $P(x) = “2^{2^x} + 1$ is a prime.”

In the 17th century, Pierre de Fermat made the conjecture⁷ that $\forall x \in \mathbb{N}, P(x)$. No doubt, this seemed reasonable to Fermat because the numbers given by this formula (they are called Fermat numbers in his honor) are all primes – at first! Fermat numbers are conventionally denoted with a subscripted letter F, $F_n = 2^{2^n} + 1$, the first five Fermat numbers are prime.

$$\begin{aligned} F_0 &= 2^{2^0} + 1 = 3 \\ F_1 &= 2^{2^1} + 1 = 5 \\ F_2 &= 2^{2^2} + 1 = 17 \\ F_3 &= 2^{2^3} + 1 = 257 \\ F_4 &= 2^{2^4} + 1 = 65537 \end{aligned}$$

Fermat probably computed that $F_5 = 4294967297$, and we can well imagine that he checked that this number was not divisible by any small primes.

⁷Fermat’s more famous conjecture, that $x^n + y^n = z^n$ has no non-trivial integer solutions if n is an integer with $n > 2$ was discovered after his death.

Of course, this was well before the development of effective computing machinery, so we shouldn't blame Fermat for not noticing that $4294967297 = 641 \cdot 6700417$. This remarkable feat of factoring can be replicated in seconds on a modern computer, however it was done first by Leonhard Euler in 1732! There is quite a lot of literature concerning the primeness and/or compositeness of Fermat numbers. So far, all the Fermat numbers between F_5 and F_{32} (inclusive) have been shown to be composite. One might be tempted to conjecture that only the first five Fermat numbers are prime, however this temptation should be resisted . . .

Let us set aside, for the moment, further questions about Fermat numbers. Suppose we define the set U (for 'Universe') by $U = \{0, 1, 2, 3, 4\}$. Then the assertion, " $\forall x \in U, P(x)$." is certainly true. You should note that the only variable in this sentence is x , and that the variable is bound – it is universally quantified. Open sentences that have all variables bound are *statements*. It is possible (in principle, and in finite universes, in practice) to check the truth value of such sentences. Indeed, the sentence " $\forall x \in U, P(x)$ " has the same logical content as " $P(0) \wedge P(1) \wedge P(2) \wedge P(3) \wedge P(4)$ ". Both happen to be true, but the real point here is to note that a universally quantified sentence can be thought of instead as a conjunction.

Exercise. Define a new set U by $U = \{0, 1, 2, 3, 4, 5\}$. Write a sentence using disjunctions that is equivalent to " $\exists x \in U, \neg P(x)$."

Even when we are dealing with infinite universes, it is possible to think of universally quantified sentences in terms of conjunctions, and existentially quantified sentences in terms of disjunctions. For example, a quick look at the graphs should be sufficient to convince you that " $x > \ln x$ " is a sentence that is true for all x values in \mathbb{R}^+ . There is a notation, reminiscent of so-called sigma notation for sums, that can be used to express this universally quantified sentence as a conjunction.

$$\forall x \in \mathbb{R}^+, x > \ln x \cong \bigwedge_{x \in \mathbb{R}^+} x > \ln x$$

A similar notation exists for disjunctions. Purely as an example, consider the following problem from recreational math: Find a four digit number that is an integer multiple of its reversal. (By reversal, we mean the four digit number with the digits in the opposite order – for example, the reversal of 1234 is 4321.) The sentence⁸ that states that this question has a solution is

$$\exists abcd \in \mathbb{Z}, \exists k \in \mathbb{Z}, abcd = k \cdot dcba$$

This could be expressed instead as the disjunction of 9000 statements, or more compactly as

$$\bigvee_{1000 \leq abcd \leq 9999} \exists k \in \mathbb{Z}, abcd = k \cdot dcba.$$

Exercise. *The existential statement above is true because $8712 = 4 \cdot 2178$. There is one other solution – find it!*

An important, or at least useful, talent for a Mathematics student to develop is the ability to negate quantified sentences. The major reason for this is the fact that the contrapositive of a conditional sentence is logically equivalent to it. This leads to a method known as “proof by contraposition” which can be expressed succinctly by the advice:

“If you get stuck, try writing down the contrapositive.”

Since writing down the contrapositive will often involve finding the negation of a quantified sentence, let’s try a few.

⁸This sentence uses what is commonly referred to as an “abuse of notation” in order to avoid an unnecessarily complex problem statement. One should not necessarily avoid such abuses if one’s readers can be expected to easily understand what is meant, any more than one should completely eschew the splitting of infinitives.

Our universe of discourse⁹ will be $P = \{\text{Manny, Moe, Jack}\}$. Consider the sentence “ $\forall x \in P, x$ starts with M.” The equivalent sentence expressed conjunctively is

$$\begin{aligned} &(\text{Manny starts with M}) \wedge \\ &(\text{Moe starts with M}) \wedge \\ &(\text{Jack starts with M}). \end{aligned}$$

The negation of this sentence (by DeMorgan’s law) is a disjunction:

$$\begin{aligned} &(\text{Manny doesn't start with M}) \vee \\ &(\text{Moe doesn't start with M}) \vee \\ &(\text{Jack doesn't start with M}) \end{aligned}$$

Finally, this disjunction of three sentences can be converted into a single sentence, existentially quantified over P :

$$\text{“}\exists x \in P, \neg(x \text{ starts with M).”}$$

The discussion in the previous paragraphs justifies some laws of Logic which should be thought of as generalizations of DeMorgan’s laws:

$$\neg(\forall x \in U, P(x)) \cong \exists x \in U, \neg P(x)$$

and

$$\neg(\exists x \in U, P(x)) \cong \forall x \in U, \neg P(x).$$

It’s equally valid to think of these rules in a way that’s divorced from DeMorgan’s laws. To show that a universal sentence is *false*, it suffices to show that an existential sentence involving a negation of the original is true.¹⁰

⁹The Pep Boys – Manny, Moe and Jack – are hopefully known to some readers as the mascots of a chain of automotive supply stores.

¹⁰To show that it is not the case that every Pep boy’s name starts with ‘M’, one only

needs to demonstrate that there is a Pep boy (Jack) whose name doesn't start with 'M'.

Exercises — 2.5

1. There is a common variant of the existential quantifier, $\exists!$, if you write $\exists! x, P(x)$ you are asserting that there is a *unique* element in the universe that makes $P(x)$ true. Determine how to negate the sentence $\exists! x, P(x)$.

2. The order in which quantifiers appear is important. Let $L(x, y)$ be the open sentence “ x is in love with y .” Discuss the meanings of the following quantified statements and find their negations.

(a) $\forall x \exists y L(x, y)$.

(b) $\exists x \forall y L(x, y)$.

(c) $\forall x \forall y L(x, y)$.

(d) $\exists x \exists y L(x, y)$.

3. Determine a useful denial of:

$$\forall \epsilon > 0 \exists \delta > 0 \forall x (|x - c| < \delta) \implies (|f(x) - l| < \epsilon).$$

The denial above gives a criterion for saying $\lim_{x \rightarrow c} f(x) \neq l$.

4. A *Sophie Germain prime* is a prime number p such that the corresponding odd number $2p + 1$ is also a prime. For example 11 is a Sophie Germain prime since $23 = 2 \cdot 11 + 1$ is also prime. Almost all Sophie Germain primes are congruent to 5 (mod 6), nevertheless, there are exceptions – so the statement “There are Sophie Germain primes that are not 5 mod 6.” is true. Verify this.

5. Alvin, Betty, and Charlie enter a cafeteria which offers three different entrees, turkey sandwich, veggie burger, and pizza; four different beverages, soda, water, coffee, and milk; and two types of desserts, pie and

pudding. Alvin takes a turkey sandwich, a soda, and a pie. Betty takes a veggie burger, a soda, and a pie. Charlie takes a pizza and a soda. Based on this information, determine whether the following statements are true or false.

- (a) \forall people p , \exists dessert d such that p took d .
- (b) \exists person p such that \forall desserts d , p did not take d .
- (c) \forall entrees e , \exists person p such that p took e .
- (d) \exists entree e such that \forall people p , p took e .
- (e) \forall people p , p took a dessert \iff p did not take a pizza.
- (f) Change one word of statement 5d so that it becomes true.
- (g) Write down the negation of 5a and compare it to statement 5b. Hopefully you will see that they are the same! Does this make you want to modify one or both of your answers to 5a and 5b?

2.6 Deductive reasoning and Argument forms

Deduction is the process by which we determine new truths from old. It is sometimes claimed that nothing truly new can come from deduction, the truth of a statement that is arrived at by deductive processes was lying (perhaps hidden somewhat) within the hypotheses. This claim is something of a canard, as any Sherlock Holmes aficionado can tell you, the statements that can sometimes be deduced from others can be remarkably surprising. A better argument against deduction is that it is a relatively ineffective way for most human beings to discover new truths – for that purpose inductive processes are superior for the majority of us. Nevertheless, if a chain of deductive reasoning leading from known hypotheses to a particular conclusion can be exhibited, the truth of the conclusion is *unassailable*. For this reason, mathematicians have latched on to deductive reasoning as *the* tool for, if not discovering our theorems, communicating them to others.

The word “argument” has a negative connotation for many people because it seems to have to do with *disagreement*. Arguments within mathematics (as well as many other scholarly areas), while they may be impassioned, should not involve discord. A mathematical argument is a sequence of logically connected statements designed to produce *agreement* as to the validity of a proposition. This “design” generally follows one of two possibilities, inductive reasoning or deductive reasoning. In an inductive argument a long list of premises is presented whose truths are considered to be apparent to all, each of which provides evidence that the desired conclusion is true. So an inductive argument represents a kind of statistical thing, you have all these statements that are true each of which indicates that the conclusion is most likely true. . . A strong inductive argument amounts to what attorneys call a “preponderance of the evidence.” Occasionally a person who has been convicted of a crime based on a preponderance of the evidence is later found to be innocent. This usually happens when new evidence is discovered that

incontrovertibly proves (i.e. shows through deductive means) that he or she cannot be guilty. In a nutshell: inductive arguments can be wrong.

In contrast a deductive argument can only turn out to be wrong under certain well-understood circumstances.

Like an inductive argument, a deductive argument is essentially just a long sequence of statements; but there is some additional structure. The last statement in the list is the *conclusion* – the statement to be proved – those occurring before it are known as *premises*. Premises may be further subdivided into (at least) five sorts: axioms, definitions, previously proved theorems, hypotheses and deductions. Axioms and definitions are often glossed over, indeed, they often go completely unmentioned (but rarely *unused*) in a proof. In the interest of brevity this is quite appropriate, but conceptually, you should think of an argument as being based off of the axioms for the particular area you are working in, and its standard definitions. A rote knowledge of all the other theorems proved up to the one you are working with would generally be considered excessive, but completely memorizing the axioms and standard definitions of a field is essential. Hypotheses are a funny class of premises – they are things which can be assumed true for the sake of the current argument. For example, if the statement you are trying to prove is a conditional, then the antecedent may be assumed true (if the antecedent is false, then the conditional is automatically true!). You should always be careful to list all hypotheses explicitly, and at the end of your proof make sure that each and every hypothesis got used somewhere along the way. If a hypothesis really isn't necessary then you have proved a more general statement (that's a good thing).

Finally, deductions – I should note that the conclusion is also a deduction – obey a very strict rule: every deduction follows from the premises that have already been written down (this includes axioms and definitions that probably won't actually have been written, hypotheses and all the deductions made up to this point) by one of the so-called rules of inference.

Each of the rules of inference actually amounts to a logical tautology that has been re-expressed as a sort of re-writing rule. Each rule of inference will be expressed as a list of logical sentences that are assumed to be among the premises of the argument, a horizontal bar, followed by the symbol \therefore (which is usually voiced as the word “therefore”) and then a new statement that can be placed among the deductions.

For example, one (very obvious) rule of inference is

$$\frac{A \wedge B}{\therefore B}$$

This rule is known as *conjunctive simplification*, and is equivalent to the tautology $(A \wedge B) \implies B$.

The *modus ponens* rule¹¹ is one of the most useful.

$$\frac{A \quad A \implies B}{\therefore B}$$

Modus ponens is related to the tautology $(A \wedge (A \implies B)) \implies B$.

Modus tollens is the rule of inference we get if we put modus ponens through the “contrapositive” wringer.

$$\frac{\neg B \quad A \implies B}{\therefore \neg A}$$

Modus tollens is related to the tautology $(\neg B \wedge (A \implies B)) \implies \neg A$.

Modus ponens and modus tollens are also known as *syllogisms*. A syllogism is an argument form wherein a deduction follows from two premises.

¹¹Latin for “method of affirming”, the related *modus tollens* rule means “method of denying.”

There are two other common syllogisms, *hypothetical syllogism* and *disjunctive syllogism*.

Hypothetical syllogism basically asserts a transitivity property for implications.

$$\frac{A \implies B \quad B \implies C}{\therefore A \implies C}$$

Disjunctive syllogism can be thought of as a statement about alternatives, but be careful to remember that in Logic, the disjunction always has the inclusive sense.

$$\frac{A \vee B \quad \neg B}{\therefore A}$$

Exercise. Convert the $A \vee B$ that appears in the premises of the disjunctive syllogism rule into an equivalent conditional. How is the new argument form related to *modus ponens* and/or *modus tollens*?

The word “dilemma” usually refers to a situation in which an individual is faced with an impossible choice. A cute example known as the Crocodile’s dilemma is as follows:

A crocodile captures a little boy who has strayed too near the river. The child’s father appears and the crocodile tells him “Don’t worry, I shall either release your son or I shall eat him. If you can say, in advance, which I will do, then I shall release him.” The father responds, “You will eat my son.” What should the crocodile do?

In logical arguments the word dilemma is used in another sense having to do with certain rules of inference. *Constructive dilemma* is a rule of inference having to do with the conclusion that one of two possibilities must hold.

$$\begin{array}{c}
 A \implies B \\
 C \implies D \\
 A \vee C \\
 \hline
 \therefore B \vee D
 \end{array}$$

Destructive dilemma is often not listed among the rules of inference because it can easily be obtained by using the constructive dilemma and replacing the implications with their contrapositives.

$$\begin{array}{c}
 A \implies B \\
 C \implies D \\
 \neg B \vee \neg D \\
 \hline
 \therefore \neg A \vee \neg C
 \end{array}$$

In Table 2.3, the ten most common rules of inference are listed. Note that all of these are equivalent to tautologies that involve conditionals (as opposed to biconditionals), every one of the basic logical equivalences that we established in Section 2.3 is really a tautology involving a biconditional, collectively these are known as the “rules of replacement.” In an argument, any statement allows us to infer a logically equivalent statement. Or, put differently, we could replace any premise with a different, but logically equivalent, premise. You might enjoy trying to determine a minimal set of rules of inference, that together with the rules of replacement would allow one to form all of the same arguments as the ten rules in Table 2.3.

Name	Form
Modus ponens	$\begin{array}{l} A \\ A \implies B \\ \hline \therefore B \end{array}$
Modus tollens	$\begin{array}{l} \neg B \\ A \implies B \\ \hline \therefore \neg A \end{array}$
Hypothetical syllogism	$\begin{array}{l} A \implies B \\ B \implies C \\ \hline \therefore A \implies C \end{array}$
Disjunctive syllogism	$\begin{array}{l} A \vee B \\ \neg B \\ \hline \therefore A \end{array}$
Constructive dilemma	$\begin{array}{l} A \implies B \\ C \implies D \\ A \vee C \\ \hline \therefore B \vee D \end{array}$

Table 2.3: The rules of inference.

Name	Form
Destructive dilemma	$ \begin{array}{l} A \implies B \\ C \implies D \\ \neg B \vee \neg D \\ \hline \therefore \neg A \vee \neg C \end{array} $
Conjunctive simplification	$ \begin{array}{l} A \wedge B \\ \hline \therefore A \end{array} $
Conjunctive addition	$ \begin{array}{l} A \\ B \\ \hline \therefore A \wedge B \end{array} $
Disjunctive addition	$ \begin{array}{l} A \\ \hline \therefore A \vee B \end{array} $
Absorption	$ \begin{array}{l} A \implies B \\ \hline \therefore A \implies (A \wedge B) \end{array} $

Table 2.3: The rules of inference. (continued)

Exercises — 2.6

1. In the movie “Monty Python and the Holy Grail” we encounter a medieval villager who (with a bit of prompting) makes the following argument.

If she weighs the same as a duck, then she’s made of wood.

If she’s made of wood then she’s a witch.

Therefore, if she weighs the same as a duck, she’s a witch.

Which rule of inference is he using?

2. In constructive dilemma, the antecedent of the conditional sentences are usually chosen to represent opposite alternatives. This allows us to introduce their disjunction as a tautology. Consider the following proof that there is never any reason to worry (found on the walls of an Irish pub).

Either you are sick or you are well.

If you are well there’s nothing to worry about.

If you are sick there are just two possibilities:

Either you will get better or you will die.

If you are going to get better there’s nothing to worry about.

If you are going to die there are just two possibilities:

Either you will go to Heaven or to Hell.

If you go to Heaven there is nothing to worry about. If you go to Hell, you’ll be so busy shaking hands with all your friends there won’t be time to worry . . .

Identify the three tautologies that are introduced in this “proof.”

3. For each of the following arguments, write it in symbolic form and determine which rules of inference are used.

- (a) You are either with us, or you're against us. And you don't appear to be with us. So, that means you're against us!
- (b) All those who had cars escaped the flooding. Sandra had a car – therefore, Sandra escaped the flooding.
- (c) When Johnny goes to the casino, he always gambles 'til he goes broke. Today, Johnny has money, so Johnny hasn't been to the casino recently.
- (d) (A non-constructive proof that there are irrational numbers a and b such that a^b is rational.) Either $\sqrt{2}^{\sqrt{2}}$ is rational or it is irrational. If $\sqrt{2}^{\sqrt{2}}$ is rational, we let $a = b = \sqrt{2}$. Otherwise, we let $a = \sqrt{2}^{\sqrt{2}}$ and $b = \sqrt{2}$. (Since $\sqrt{2}^{\sqrt{2}^{\sqrt{2}}} = 2$, which is rational.) It follows that in either case, there are irrational numbers a and b such that a^b is rational.

2.7 Validity of arguments and common errors

An argument is said to be *valid* or to have a *valid form* if each deduction in it can be justified with one of the rules of inference listed in the previous section. The *form* of an argument might be valid, but still the conclusion may be false if some of the premises are false. So to show that an argument is good we have to be able to do two things: show that the argument is *valid* (i.e. that every step can be justified) and that the argument is *sound* which means that all the premises are true. If you start off with a false premise, you can prove *anything*!

Consider, for example the following “proof” that $2 = 1$.

Suppose that a and b are two real numbers such that $a = b$.

	by hypothesis, a and b are equal, so
$a^2 = ab$	
	subtracting b^2 from both sides
$a^2 - b^2 = ab - b^2$	
	factoring both sides
$(a + b)(a - b) = b(a - b)$	
	canceling $(a - b)$ from both sides
$a + b = b$	

Now let a and b both have a particular value, $a = b = 1$, and we see that $1 + 1 = 1$, i.e. $2 = 1$.

This argument is not sound (thank goodness!) because one of the premises – actually the bad premise appears as one of the justifications of a step – is false. You can argue with perfect logic to achieve complete nonsense if you include false premises.

Exercise. *It is not true that you can always cancel the same thing from both sides of an equation. Under what circumstances is such cancellation disallowed?*

So, how can you tell if an argument has a valid form? Use a truth table. As an example, we'll verify that the rule of inference known as "destructive dilemma" is valid using a truth table. This argument form contains 4 predicate variables so the truth table will have 16 rows. There is a column for each of the variables, the premises of the argument and its conclusion.

A	B	C	D	$A \implies B$	$C \implies D$	$\neg B \vee \neg D$	$\neg A \vee \neg C$
T	T	T	T	T	T	ϕ	ϕ
T	T	T	ϕ	T	ϕ	T	ϕ
T	T	ϕ	T	T	T	ϕ	T
T	T	ϕ	ϕ	T	T	T	T
T	ϕ	T	T	ϕ	T	T	ϕ
T	ϕ	T	ϕ	ϕ	ϕ	T	ϕ
T	ϕ	ϕ	T	ϕ	T	T	T
T	ϕ	ϕ	ϕ	ϕ	T	T	T
ϕ	T	T	T	T	T	ϕ	T
ϕ	T	T	ϕ	T	ϕ	T	T
ϕ	T	ϕ	T	T	T	ϕ	T
ϕ	T	ϕ	ϕ	T	T	T	T
ϕ	ϕ	T	T	T	T	T	T
ϕ	ϕ	T	ϕ	T	ϕ	T	T
ϕ	ϕ	ϕ	T	T	T	T	T
ϕ	ϕ	ϕ	ϕ	T	T	T	T

Now, mark the lines in which all of the premises of this argument form are true. You should note that *in every single situation in which all the premises are true* the conclusion is also true. That's what makes "destructive dilemma" – and all of its friends – a rule of inference. Whenever all the

premises are true so is the conclusion. You should also notice that there are several rows in which the conclusion is true but some one of the premises isn't. That's okay too, isn't it reasonable that the conclusion of an argument can be true, but at the same time the particulars of the argument are unconvincing?

As we've noted earlier, an argument by deductive reasoning can go wrong in only certain well-understood ways. Basically, either the form of the argument is invalid, or at least one of the premises is false. Avoiding false premises in your arguments can be trickier than it sounds – many statements that sound appealing or intuitively clear are actually counter-factual. The other side of the coin, being sure that the *form* of your argument is valid, seems easy enough – just be sure to only use the rules of inference as found in Table 2.3. Unfortunately most arguments that you either read or write will be in prose, rather than appearing as a formal list of deductions. When dealing with that setting – using natural rather than formalized language – making errors in form is quite common.

Two invalid forms are usually singled out for criticism, the *converse error* and the *inverse error*. In some sense these two apparently different ways to screw up are really the same thing. Just as a conditional statement and its contrapositive are known to be equivalent, so too are the other related statements – the converse and the inverse – equivalent. The converse error consists of mistaking the implication in a modus ponens form for its converse.

The converse error:

$$\frac{B}{A \implies B} \\ \therefore A$$

Consider, for a moment the following argument.

If a rhinoceros sees something on fire, it will stomp on it.

A rhinoceros stomped on my duck.

Therefore, the rhino must have thought that my duck was on fire.

It *is* true that rhinoceroses have an instinctive desire to extinguish fires. Also, we can well imagine that if someone made this ridiculous argument that their duck must actually have been crushed by a rhino. But, is the conclusion that the duck was on fire justified? Not really, what the first part of the argument asserts is that “(on fire) implies (rhino stomping)” but couldn’t a rhino stomp on something for other reasons? Perhaps the rhino was just ill-tempered. Perhaps the duck was just horrifically unlucky.

The closer the conditional is to being a biconditional, the more reasonable sounding is an argument exhibiting the converse error. Indeed, if the argument actually contains a biconditional, the “converse error” is not an error at all.

The following is a perfectly valid argument, that (sadly) has a false premise.

You will get an A in your Foundations class if and only if you read Dr. Fields’ book.

You read Dr. Fields’ book.

Therefore, you will get an A in Foundations.

Suppose that we try changing the major premise of that last argument to something more believable.

If you read Dr. Fields’ book, you will pass your Foundations class.

You did not read Dr. Fields’ book.

Therefore, you will not pass Foundations.

This last argument exhibits the so-called *inverse error*. It is by no means meant as a guarantee, but nevertheless, it seems reasonable that if someone reads this book they will pass a course on this material. The second premise

is also easy to envision as true, although the “you” that it refers to obviously isn’t *you*, because *you* are reading this book! But even if we accept the premises as true, the conclusion doesn’t follow. A person might have read some other book that addressed the requisite material in an exemplary way.

Notice that the names for these two errors are derived from the change that would have to be made to convert them to modus ponens. For example, the inverse error is depicted formally by:

$$\frac{\neg A \quad A \implies B}{\therefore \neg B}$$

If we replaced the conditional in this argument form by its *inverse* ($\neg A \implies \neg B$) then the revised argument would be modus ponens. Similarly, if we replace the conditional in an argument that suffers from the converse error by its converse, we’ll have modus ponens.

Exercises — 2.7

1. Determine the logical form of the following arguments. Use symbols to express that form and determine whether the form is valid or invalid. If the form is invalid, determine the type of error made. Comment on the soundness of the argument as well, in particular, determine whether any of the premises are questionable.
 - (a) All who are guilty are in prison.
George is not in prison.
Therefore, George is not guilty.
 - (b) If one eats oranges one will have high levels of vitamin C.
You do have high levels of vitamin C.
Therefore, you must eat oranges.
 - (c) All fish live in water.
The mackerel is a fish.
Therefore, the mackerel lives in water.
 - (d) If you're lazy, don't take math courses.
Everyone is lazy.
Therefore, no one should take math courses.
 - (e) All fish live in water.
The octopus lives in water.
Therefore, the octopus is a fish.
 - (f) If a person goes into politics, they are a scoundrel.
Harold has gone into politics.
Therefore, Harold is a scoundrel.
2. Below is a rule of inference that we call extended elimination.

$$\frac{\begin{array}{l} (A \vee B) \vee C \\ \neg A \\ \neg B \end{array}}{\therefore C}$$

Use a truth table to verify that this rule is valid.

3. If we allow quantifiers and open sentences in an argument form we get a couple of new argument forms. Arguments involving existentially quantified premises are rare – the new forms we are speaking of are called “universal modus ponens” and “universal modus tollens.” The minor premises may also be quantified or they may involve particular elements of the universe of discourse – this leads us to distinguish argument subtypes that are termed “universal” and “particular.”

$$\forall x, A(x) \implies B(x)$$

For example $\frac{A(p)}{\therefore B(p)}$ is the particular form of universal

modus ponens (here, p is not a variable – it stands for some particular

$$\forall x, A(x) \implies B(x)$$

element of the universe of discourse) and $\frac{\forall x, \neg B(x)}{\therefore \forall x, \neg A(x)}$ is the

$$\therefore \forall x, \neg A(x)$$

universal form of (universal) modus tollens.

Reexamine the arguments from problem (1), determine their forms (including quantifiers) and whether they are universal or particular.

4. Identify the rule of inference being used.

(a) The Buley Library is very tall.

Therefore, either the Buley Library is very tall or it has many levels underground.

(b) The grass is green.

The sky is blue.

Therefore, the grass is green and the sky is blue.

(c) g has order 3 or it has order 4.

If g has order 3, then g has an inverse.

If g has order 4, then g has an inverse.

Therefore, g has an inverse.

(d) x is greater than 5 and x is less than 53.

Therefore, x is less than 53.

(e) If $a|b$, then a is a perfect square.

If $a|b$, then b is a perfect square.

Therefore, if $a|b$, then a is a perfect square and b is a perfect square.

5. Read the following proof that the sum of two odd numbers is even. Discuss the rules of inference used.

Proof: Let x and y be odd numbers. Then $x = 2k + 1$ and $y = 2j + 1$ for some integers j and k . By algebra,

$$x + y = 2k + 1 + 2j + 1 = 2(k + j + 1).$$

Note that $k + j + 1$ is an integer because k and j are integers.

Hence $x + y$ is even.

Q.E.D.

6. Sometimes in constructing a proof we find it necessary to “weaken” an inequality. For example, we might have already deduced that $x < y$ but what we need in our argument is that $x \leq y$. It is okay to deduce $x \leq y$

from $x < y$ because the former is just shorthand for $x < y \vee x = y$. What rule of inference are we using in order to deduce that $x \leq y$ is true in this situation?

Chapter 3

Proof techniques I — Standard methods

Love is a snowmobile racing across the tundra and then suddenly it flips over, pinning you underneath. At night, the ice weasels come. —Matt Groening

3.1 Direct proofs of universal statements

If you form the product of 4 consecutive numbers, the result will be one less than a perfect square. Try it!

$$1 \cdot 2 \cdot 3 \cdot 4 = 24 = 5^2 - 1$$

$$2 \cdot 3 \cdot 4 \cdot 5 = 120 = 11^2 - 1$$

$$3 \cdot 4 \cdot 5 \cdot 6 = 360 = 19^2 - 1$$

It always works!

The three calculations that we've carried out above constitute an inductive argument in favor of the result. If you like we can try a bunch of further

examples,

$$13 \cdot 14 \cdot 15 \cdot 16 = 43680 = 209^2 - 1$$

$$14 \cdot 15 \cdot 16 \cdot 17 = 571200 = 239^2 - 1$$

but really, no matter how many examples we produce, we haven't *proved* the statement — we've just given evidence.

Generally, the first thing to do in proving a universal statement like this is to rephrase it as a conditional. The resulting statement is a *Universal Conditional Statement* or a UCS. The reason for taking this step is that the *hypotheses* will then be clear — they form the antecedent of the UCS. So, while you won't have really made any progress in the proof by taking this advice, you will at least know what tools you have at hand. Taking the example we started with, and rephrasing it as a UCS we get

$$\forall a, b, c, d \in \mathbb{Z}, (a, b, c, d \text{ consecutive}) \implies \exists k \in \mathbb{Z}, a \cdot b \cdot c \cdot d = k^2 - 1$$

The antecedent of the UCS is that a, b, c and d must be *consecutive*. By concentrating our attention on what it means to be consecutive, we should quickly realize that the original way we thought of the problem involved a red herring. We don't need to have variables for all four numbers; because they are consecutive, a uniquely determines the other three. Finally we have a version of the statement that we'd like to prove that should lend itself to our proof efforts.

Theorem 3.1.1.

$$\forall a \in \mathbb{Z}, \exists k \in \mathbb{Z}, a(a + 1)(a + 2)(a + 3) = k^2 - 1.$$

In this simplistic example, the only thing we need to do is come up with a value for k given that we know what a is. In other words, a “proof” of this statement involves doing some algebra.

Without further ado...

Proof: Suppose that a is a particular but arbitrarily chosen integer. Consider the product of the 4 consecutive integers, a , $a + 1$, $a + 2$ and $a + 3$. We would like to show that this product is one less than the square of an integer k . Let k be $a^2 + 3a + 1$.

First, note that

$$a(a + 1)(a + 2)(a + 3) = a^4 + 6a^3 + 11a^2 + 6a.$$

Then, note that

$$\begin{aligned} k^2 - 1 &= (a^2 + 3a + 1)^2 - 1 \\ &= (a^4 + 6a^3 + 11a^2 + 6a + 1) - 1 \\ &= a^4 + 6a^3 + 11a^2 + 6a. \end{aligned}$$

Q.E.D.

Now, if you followed the algebra above, (none of which was particularly difficult) the proof stands as a completely valid argument showing the truth of our proposition, but this is *very* unsatisfying! All the real work was concealed in one stark little sentence: “Let k be $a^2 + 3a + 1$.” Where on Earth did that particular value of k come from? The answer to that question should hopefully convince you that there is a huge difference between *devising* a proof and *writing* one. A good proof can sometimes be somewhat akin to a good demonstration of magic, a magician doesn’t reveal the inner workings of his trick, neither should a mathematician feel guilty about leaving out some of the details behind the work! Heck, there are plenty of times when

you just have to *guess* at something, but if your guess works out, you can write a perfectly correct proof.

In devising the proof above, we multiplied out the consecutive numbers and then realized that we'd be done if we could find a polynomial in a whose square was $a^4 + 6a^3 + 11a^2 + 6a + 1$. Now, obviously, we're going to need a quadratic polynomial, and because the leading term is a^4 and the constant term is 1, it should be of the form $a^2 + ma + 1$. Squaring this gives $a^4 + 2ma^3 + (m^2 + 2)a^2 + 2ma + 1$ and comparing that result with what we want, we pretty quickly realize that m had better be 3. So it wasn't magic after all!

This seems like a good time to make a comment on polynomial arithmetic. Many people give up (or go searching for a computer algebra system) when dealing with products of anything bigger than binomials. This is a shame because there is an easy method using a table for performing such multiplications. As an example, in devising the previous proof we needed to form the product $a(a+1)(a+2)(a+3)$, now we can use the distributive law or the infamous F.O.I.L rule to multiply pairs of these, but we still need to multiply $(a^2 + a)$ with $(a^2 + 5a + 6)$. Create a table that has the terms of these two polynomials as its row and column headings.

	a^2	$5a$	6
a^2			
a			

Now, fill in the entries of the table by multiplying the corresponding row and column headers.

	a^2	$5a$	6
a^2	a^4	$5a^3$	$6a^2$
a	a^3	$5a^2$	$6a$

Finally add up all the entries of the table, combining any like terms.

You should note that the F.O.I.L rule is just a mnemonic for the case when the table has 2 rows and 2 columns.

Okay, let's get back to doing proofs. We are going to do a lot of proofs involving the concepts of elementary number theory so, as a convenience, all of the definitions that were made in Chapter 1 are gathered together in Table 3.1.

Even

$$\forall n \in \mathbb{Z},$$

$$n \text{ is even} \iff \exists k \in \mathbb{Z}, n = 2k$$

Odd

$$\forall n \in \mathbb{Z},$$

$$n \text{ is odd} \iff \exists k \in \mathbb{Z}, n = 2k + 1$$

Divisibility

$$\forall n \in \mathbb{Z}, \forall d > 0 \in \mathbb{Z},$$

$$d | n \iff \exists k \in \mathbb{Z}, n = kd$$

Floor

$$\forall x \in \mathbb{R},$$

$$y = \lfloor x \rfloor \iff y \in \mathbb{Z} \wedge y \leq x < y + 1$$

Ceiling

$$\forall x \in \mathbb{R},$$

$$y = \lceil x \rceil \iff y \in \mathbb{Z} \wedge y - 1 < x \leq y$$

Quotient-remainder theorem, Div and Mod

$$\forall n, d > 0 \in \mathbb{Z},$$

$$\exists! q, r \in \mathbb{Z}, n = qd + r \wedge 0 \leq r < d$$

$$n \operatorname{div} d = q$$

$$n \operatorname{mod} d = r$$

Prime

$$\forall p \in \mathbb{Z}$$

$$p \text{ is prime} \iff$$

$$(p > 1) \wedge (\forall x, y \in \mathbb{Z}^+, p = xy \implies x = 1 \vee y = 1)$$

Table 3.1: The definitions of elementary number theory restated.

In this section we are concerned with direct proofs of universal statements. Such statements come in two flavors – those that appear to involve conditionals, and those that don't:

Every prime greater than two is odd.

versus

For all integers n , if n is a prime greater than two, then n is odd.

These two forms can readily be transformed one into the other, so we will always concentrate on the latter. A direct proof of a UCS always follows a form known as “generalizing from the generic particular.” We are trying to prove that $\forall x \in U, P(x) \implies Q(x)$. The argument (in skeletal outline) will look like:

Proof: Suppose that a is a particular but arbitrary element of U such that $P(a)$ holds.

⋮

Therefore $Q(a)$ is true.

Thus we have shown that for all x in U , $P(x) \implies Q(x)$.

Q.E.D.

Okay, so this outline is pretty crappy. It tells you how to start and end a direct proof, but those obnoxious dot-dot-dots in the middle are where all the real work has to go. If I could tell you (even in outline) how to fill in those dots, that would mean mathematical proof isn't really a very interesting activity to engage in. Filling in those dots will sometimes (rarely) be obvious, more often it will be extremely challenging; it will require great creativity, loads of concentration, you'll call on all your previous mathematical experiences, and you will most likely experience a certain degree of anguish. Just

remember that your sense of accomplishment is proportional to the difficulty of the puzzles you attempt. So let's attempt another...

In Table 3.1 one of the very handy notions defined is that of the *floor* of a real number.

$$y = \lfloor x \rfloor \iff (y \in \mathbb{Z} \wedge y \leq x < y + 1).$$

There is a sad tendency for people to apply old rules in new situations just because of a chance similarity in the notation. The brackets used in notating the floor function look very similar to ordinary parentheses, so the following “rule” is often proposed

$$\lfloor x + y \rfloor = \lfloor x \rfloor + \lfloor y \rfloor$$

Exercise. Find a counterexample to the previous “rule.”

What is (perhaps) surprising is that if one of the numbers involved is an integer then the “rule” really works.

Theorem 3.1.2.

$$\forall x \in \mathbb{R}, \forall n \in \mathbb{Z}, \lfloor x + n \rfloor = \lfloor x \rfloor + \lfloor n \rfloor$$

Since the floor of an integer *is* that integer, we could restate this as $\lfloor x + n \rfloor = \lfloor x \rfloor + n$.

Now, let's try rephrasing this theorem as a UCS: If x is a real number and n is an integer, then $\lfloor x + n \rfloor = \lfloor x \rfloor + n$. This is bad ... it appears that the only hypotheses that we can use involve what kinds of numbers x and n are — our hypotheses aren't particularly potent. The next most useful ally in constructing proofs are the definitions of the concepts involved. The quantity $\lfloor x \rfloor$ appears in the theorem, let's make use of the definition:

$$a = \lfloor x \rfloor \iff a \in \mathbb{Z} \wedge a \leq x < a + 1.$$

The only other floor function that appears in the statement of the theorem (perhaps even more prominently) is $\lfloor x + n \rfloor$, here, the definition gives us

$$b = \lfloor x + n \rfloor \iff b \in \mathbb{Z} \wedge b \leq x + n < b + 1.$$

These definitions are our only available tools so we'll certainly *have* to make use of them, and it's important to notice that that is a good thing; the definitions allow us to work with something well-understood (the inequalities that appear within them) rather than with something new and relatively suspicious (the floor notation). Putting the proof of this statement together is an exercise in staring at the two definitions above and noting how one can be converted into the other. It is also a testament to the power of *naming* things.

Proof: Suppose that x is a particular but arbitrary real number and that n is a particular but arbitrary integer. Let $a = \lfloor x \rfloor$. By the definition of the floor function it follows that a is an integer and $a \leq x < a + 1$. By adding n to each of the parts of this inequality we deduce a new (and equally valid) inequality, $a + n \leq x + n < a + n + 1$. Note that $a + n$ is an integer and the inequality above together with this fact constitute precisely the definition of $a + n = \lfloor x + n \rfloor$. Finally, recalling that $a = \lfloor x \rfloor$ (by assumption), and rewriting, we obtain the desired result

$$\lfloor x + n \rfloor = \lfloor x \rfloor + n.$$

Q.E.D.

As we've seen in the examples presented in this section, coming up with a proof can sometimes involve a bit of ingenuity. But, sometimes, there is a "follow your nose" sort of approach that will allow you to devise a valid

argument without necessarily displaying any great leaps of genius! Here are a few pieces of advice about proof-writing:

- Before anything else, determine precisely what hypotheses you can use.
- Jot down the definitions of *anything* in the statement of the theorem.
- There are 26 letters at your disposal (and even more if you know Greek) (and you can always throw on subscripts!) don't be stingy with letters. The nastiest mistake you can make is to use the same variable for two different things.
- Please write a rough draft first. Write two drafts! Even if you can write beautiful, lucid prose on the first go around, it won't fly when it comes to organizing a proof.
- The statements in a proof are supposed to be logical statements. That means they should be Boolean (statements that are either true or false). An algebraic expression all by itself doesn't count, an inequality or an equality does.
- Don't say "if" when you mean "since." Really! If you start a proof about rational numbers like so:

Proof: Suppose that x is a particular but arbitrary rational number. If x is a rational number, it follows that ...

people are going to look at you funny. What's the point of *supposing* that x is rational, then acting as if you're in doubt of that fact by writing "if"? You mean "since."

- Mark off the beginning and the end of your proofs as a hint to your readers. In this book we start off a proof by writing *Proof:* in italics and we end every proof with the abbreviation Q.E.D.¹

¹*Quod erat demonstrandum* or "(that) which was to be demonstrated." some authors

We'll close this section with a word about axioms. The axioms in any given area of math are your most fundamental tools. Axioms don't need to be proved – we are supposed to just accept them! A very common problem for beginning proofwriters is telling the difference between statements that are axiomatic and statements that require some proof. For instance, in the exercises for this section there is a problem that asks us to prove that the sum of two rational numbers is rational. Doesn't this seem like it might be one of the axioms of rational numbers? Is it really something that *can* be proved? Well, we know how the process of adding rational numbers works: we put the fractions over a common denominator and then just add numerators. Do you see how adding fractions really rests on our ability to add the numerators (which are integers). So, in doing that exercise you can use the fact (indeed, you'll need to use the fact) that the sum of two integers is an integer. So how about *that* statement? Is it necessary to prove that adding integers produces an integer? As a matter of fact it *is* necessary since the structure of the integers rests on a foundation known as the Peano axioms for the naturals – and the Peano axioms *don't* include one that guarantees that the sum of two naturals is also a natural. If you are tempted to trace this whole thing back, to “find out how deep the rabbit hole goes,” I commend you. But, if you just want to be able to get on with doing your homework problems, I sympathize with that sentiment too. Let's agree that integers behave the way we've come to expect – if you add or multiply integers the result will be an integer.

prefer placing a small rectangle at the end of their proofs, but Q.E.D. seems more pompous.

Exercises — 3.1

1. Every prime number greater than 3 is of one of the two forms $6k + 1$ or $6k + 5$. What statement(s) could be used as hypotheses in proving this theorem?
2. Prove that 129 is odd.
3. Prove that the sum of two rational numbers is a rational number.
4. Prove that the sum of an odd number and an even number is odd.
5. Prove that if the sum of two integers is even, then so is their difference.
6. Prove that for every real number x , $\frac{2}{3} < x < \frac{3}{4} \implies \lfloor 12x \rfloor = 8$.
7. Prove that if x is an odd integer, then x^2 is of the form $4k + 1$ for some integer k .
8. Prove that for all integers a and b , if a is odd and $6 \mid (a + b)$, then b is odd.
9. Prove that $\forall x \in \mathbb{R} \ x \notin \mathbb{Z} \implies \lfloor x \rfloor + \lfloor -x \rfloor = -1$.
10. Define the *evenness* of an integer n by:

$$\text{evenness}(n) = k \iff 2^k \mid n \wedge 2^{k+1} \nmid n$$

State and prove a theorem concerning the evenness of products.

11. Suppose that a , b and c are integers such that $a \mid b$ and $b \mid c$. Prove that $a \mid c$.
12. Suppose that a , b , c and d are integers with $a \neq c$. Further, suppose that x is a real number satisfying the equation

$$\frac{ax + b}{cx + d} = 1.$$

Show that x is rational. Where is the hypothesis $a \neq c$ used?

13. Show that if two positive integers a and b satisfy $a \mid b$ and $b \mid a$ then they are equal.

3.2 More direct proofs

In creating a direct proof we need to look at our hypotheses, consider the desired conclusion, and develop a strategy for transforming A into B. Quite often you'll find it easy to make several deductions from the hypotheses, but none of them seems to be headed in the direction of the desired conclusion. The usual advice at this stage is "Try working backwards from the conclusion."²

There is a lovely result known as the "arithmetic-geometric mean inequality" whose proof epitomizes this approach. Basically this inequality compares two different ways of getting an "average" between two real numbers. The *arithmetic mean* of two real numbers a and b is the one you're probably used to, $(a+b)/2$. Many people just call this the "mean" of a and b without using the modifier "arithmetic" but as we'll see, our notion of what intermediate value to use in between two numbers is dependent on context. Consider the following two sequences of numbers (both of which have a missing entry)

$$2 \ 9 \ 16 \ 23 \ _ \ 37 \ 44$$

and

$$3 \ 6 \ 12 \ 24 \ _ \ 96 \ 192.$$

How should we fill in the blanks?

The first sequence is an *arithmetic sequence*. Arithmetic sequences are characterized by the property that the difference between successive terms is a constant. The second sequence is a *geometric sequence*. Geometric sequences have the property that the ratio of successive terms is a constant. The blank in the first sequence should be filled with the arithmetic mean

²Some people refer to this as the forwards-backwards method, since you work backwards from the conclusion, but also forwards from the premises, in the hopes of meeting somewhere in the middle.

of the surrounding entries $(23 + 37)/2 = 30$. The blank in the second sequence should be filled using the geometric mean of *its* surrounding entries: $\sqrt{24 \cdot 96} = 48$.

Given that we accept the utility of having two inequivalent concepts of *mean* that can be used in different contexts, it is interesting to see how these two means compare to one another. The arithmetic-geometric mean inequality states that the arithmetic mean is always bigger.

$$\forall a, b \in \mathbb{R}, \quad a, b \geq 0 \implies \frac{a+b}{2} \geq \sqrt{ab}$$

In proving this statement we have little choice but to work backwards from the conclusion because the only hypothesis we have to work with is that a and b are non-negative real numbers – which isn't a particularly potent tool. But what should we do? There isn't a good response to that question, we'll just have to try a bunch of different things and hope that something will work out. When we finally get around to writing up our proof though, we'll have to rearrange the statements in the opposite order from the way they were discovered. This means that we would be ill-advised to make any uni-directional inferences, we should strive to make biconditional connections between our statements (or else try to intentionally make converse errors).

The first thing that appeals to your humble author is to eliminate both the fractions and the radicals. . .

$$\frac{a+b}{2} \geq \sqrt{ab}$$

$$\iff a+b \geq 2\sqrt{ab}$$

$$\iff (a+b)^2 \geq 4ab$$

$$\iff a^2 + 2ab + b^2 \geq 4ab$$

One of the steps above involves squaring both sides of an inequality. We need to ask ourselves if this step is really reversible. In other words, is the following conditional true?

$$\forall x, y \in \mathbb{R}^{\text{nonneg}}, x \geq y \implies \sqrt{x} \geq \sqrt{y}$$

Exercise. *Provide a justification for the previous implication.*

What should we try next? There's really no good justification for this but experience working with quadratic polynomials either in equalities or inequalities leads most people to try "moving everything to one side," that is, manipulating things so that one side of the equation or inequality is zero.

$$a^2 + 2ab + b^2 \geq 4ab$$

$$\iff a^2 - 2ab + b^2 \geq 0$$

Whoa! We're done! Do you see why? If not, I'll give you one hint: the square of any real number is greater than or equal to zero.

Exercise. *Re-assemble all of the steps taken in the previous few paragraphs into a proof of the arithmetic-geometric mean inequality.*

Exercises — 3.2

1. Suppose you have a savings account which bears interest compounded monthly. The July statement shows a balance of \$ 2104.87 and the September statement shows a balance \$ 2125.97. What would be the balance on the (missing) August statement?
2. Recall that a quadratic equation $ax^2 + bx + c = 0$ has two real solutions if and only if the discriminant $b^2 - 4ac$ is positive. Prove that if a and c have different signs then the quadratic equation has two real solutions.
3. Prove that if $x^3 - x^2$ is negative then $3x + 4 < 7$.
4. Prove that for all integers a, b , and c , if $a|b$ and $a|(b + c)$, then $a|c$.
5. Show that if x is a positive real number, then $x + \frac{1}{x} \geq 2$.
6. Prove that for all real numbers a, b , and c , if $ac < 0$, then the quadratic equation $ax^2 + bx + c = 0$ has two real solutions.
Hint: The quadratic equation $ax^2 + bx + c = 0$ has two real solutions if and only if $b^2 - 4ac > 0$ and $a \neq 0$.
7. Show that $\binom{n}{k} \cdot \binom{k}{r} = \binom{n}{r} \cdot \binom{n-r}{k-r}$ (for all integers r, k and n with $r \leq k \leq n$).
8. In proving the *product rule* in Calculus using the definition of the derivative, we might start our proof with:

$$\begin{aligned} & \frac{d}{dx} (f(x) \cdot g(x)) \\ &= \lim_{h \rightarrow 0} \frac{f(x+h) \cdot g(x+h) - f(x) \cdot g(x)}{h} \end{aligned}$$

The last two lines of our proof should be:

$$\begin{aligned} &= \lim_{h \rightarrow 0} \frac{f(x+h) - f(x)}{h} \cdot g(x) + f(x) \cdot \lim_{h \rightarrow 0} \frac{g(x+h) - g(x)}{h} \\ &= \frac{d}{dx} (f(x)) \cdot g(x) + f(x) \cdot \frac{d}{dx} (g(x)) \end{aligned}$$

Fill in the rest of the proof.

3.3 Indirect proofs: contradiction and contraposition

Suppose we are trying to prove that all thrackles are polycyclic³. A *direct* proof of this would involve looking up the definition of what it means to be a thrackle, and of what it means to be polycyclic, and somehow discerning a way to convert whatever thrackle's logical equivalent is into the logical equivalent of polycyclic. As happens fairly often, there may be no obvious way to accomplish this task. Indirect proof takes a completely different tack. Suppose you had a thrackle that wasn't polycyclic, and furthermore, show that this supposition leads to something truly impossible. Well, if it's impossible for a thrackle to *not* be polycyclic, then it must be the case that all of them *are*. Such an argument is known as *proof by contradiction*.

Quite possibly the sweetest indirect proof known is Euclid's proof that there are an infinite number of primes.

Theorem 3.3.1. (*Euclid*) *The set of all prime numbers is infinite.*

Proof: Suppose on the contrary that there are only a finite number of primes. This finite set of prime numbers could, in principle, be listed in ascending order.

$$\{p_1, p_2, p_3, \dots, p_n\}$$

Consider the number N formed by adding 1 to the product of all of these primes.

$$N = 1 + \prod_{k=1}^n p_k$$

³Both of these strange sounding words represent real mathematical concepts, however, they don't have anything to do with one another.

Clearly, N is much larger than the largest prime p_n , so N cannot be a prime number itself. Thus N must be a product of some of the primes in the list. Suppose that p_j is one of the primes that divides N . Now notice that, by construction, N would leave remainder 1 upon division by p_j . This is a contradiction since we cannot have both $p_j \mid N$ and $p_j \nmid N$.

Since the supposition that there are only finitely many primes leads to a contradiction, there must indeed be an infinite number of primes.

Q.E.D.

If you are working on proving a UCS and the direct approach seems to be failing you may find that another indirect approach, proof by contraposition, will do the trick. In one sense this proof technique isn't really all that indirect; what one does is determine the contrapositive of the original conditional and then prove *that* directly. In another sense this method *is* indirect because a proof by contraposition can usually be recast as a proof by contradiction fairly easily.

The easiest proof I know of using the method of contraposition (and possibly the nicest example of this technique) is the proof of the lemma we stated in Section 1.6 in the course of proving that $\sqrt{2}$ wasn't rational. In case you've forgotten we needed the fact that whenever x^2 is an even number, so is x .

Let's first phrase this as a UCS.

$$\forall x \in \mathbb{Z}, x^2 \text{ even} \implies x \text{ even}$$

Perhaps you tried to prove this result earlier. If so you probably came across the conceptual problem that all you have to work with is the evenness

of x^2 which doesn't give you much ammunition in trying to show that x is even. The contrapositive of this statement is:

$$\forall x \in \mathbb{Z}, x \text{ not even} \implies x^2 \text{ not even}$$

Now, since x and x^2 are integers, there is only one alternative to being even – so we can re-express the contrapositive as

$$\forall x \in \mathbb{Z}, x \text{ odd} \implies x^2 \text{ odd.}$$

Without further ado, here is the proof:

Theorem 3.3.2.

$$\forall x \in \mathbb{Z}, x^2 \text{ even} \implies x \text{ even}$$

Proof: This statement is logically equivalent to

$$\forall x \in \mathbb{Z}, x \text{ odd} \implies x^2 \text{ odd}$$

so we prove that instead.

Suppose that x is a particular but arbitrarily chosen integer such that x is odd. Since x is odd, there is an integer k such that $x = 2k + 1$. It follows that $x^2 = (2k + 1)^2 = 4k^2 + 4k + 1 = 2(2k^2 + 2k) + 1$. Finally, we see that x^2 must be odd because it is of the form $2m + 1$, where $m = 2k^2 + 2k$ is clearly an integer.

Q.E.D.

Let's have a look at a proof of the same statement done by contradiction.

Proof: We wish to show that

$$\forall x \in \mathbb{Z}, x^2 \text{ even} \implies x \text{ even.}$$

Suppose to the contrary that there is an integer x such that x^2 is even but x is odd.⁴ Since x is odd, there is an integer m such that $x = 2m + 1$. Therefore, by simple arithmetic, we obtain $x^2 = 4m^2 + 4m + 1$ which is clearly odd. This is a contradiction because (by assumption) x^2 is even.

Q.E.D.

The main problem in applying the method of proof by contradiction is that it usually involves “cleverness.” You have to come up with some reason why the presumption that the theorem is false leads to a contradiction – and this may or may not be obvious. More than any other proof technique, proof by contradiction demands that we use drafts and rewriting. After monkeying around enough that we find a way to reach a contradiction, we need to go back to the beginning of the proof and highlight the feature that we will eventually contradict! After all, we want it to look like our proofs are completely clear, concise and reasonable even if their formulation caused us some sort of Gordian-level mental anguish.

We’ll end this section with an example from Geometry.

Theorem 3.3.3. *Among all triangles inscribed in a fixed circle, the one with maximum area is equilateral.*

Proof: We’ll proceed by contradiction. Suppose to the contrary that there is a triangle, $\triangle ABC$, inscribed in a circle having maximum area that is not equilateral. Since $\triangle ABC$ is not equilateral, there are two sides of it that are not equal. Without loss of generality, suppose that sides \overline{AB} and \overline{BC} have different lengths. Consider the remaining side (\overline{AC}) to be the base of this triangle. We can construct another triangle $\triangle AB'C$, also inscribed

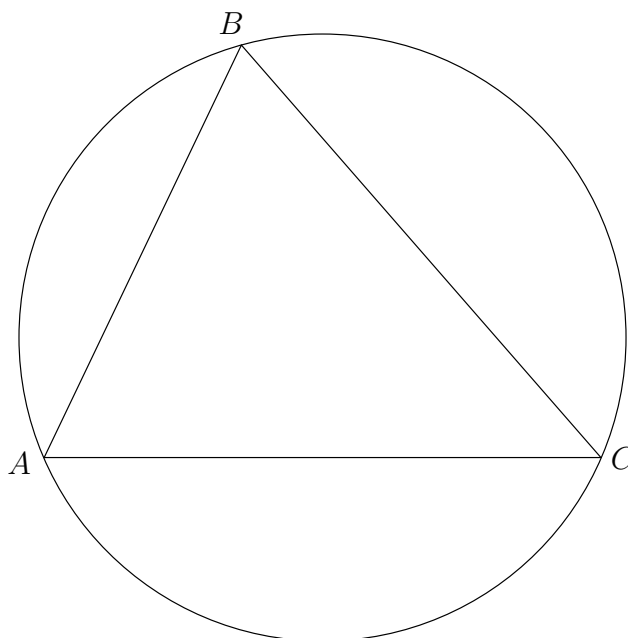
⁴Recall that the negation of a UCS is an existentially quantified conjunction.

in our circle, and also having \overline{AC} as its base, having a greater altitude than $\triangle ABC$ — since the area of a triangle is given by the formula $bh/2$ (where b is the base, and h is the altitude), this triangle's area is evidently greater than that of $\triangle ABC$. This is a contradiction since $\triangle ABC$ was presumed to have maximal area.

We leave the actual construction $\triangle AB'C$ to the following exercise.

Q.E.D.

Exercise. *Where should we place the point B' in order to create a triangle $\triangle AB'C$ having greater area than any triangle such as $\triangle ABC$ which is not isosceles?*



Exercises — 3.3

1. Prove that if the cube of an integer is odd, then that integer is odd.
2. Prove that whenever a prime p does not divide the square of an integer, it also doesn't divide the original integer. ($p \nmid x^2 \implies p \nmid x$)
3. Prove (by contradiction) that there is no largest integer.
4. Prove (by contradiction) that there is no smallest positive real number.
5. Prove (by contradiction) that the sum of a rational and an irrational number is irrational.
6. Prove (by contraposition) that for all integers x and y , if $x + y$ is odd, then $x \neq y$.
7. Prove (by contraposition) that for all real numbers a and b , if ab is irrational, then a is irrational or b is irrational.
8. A *Pythagorean triple* is a set of three natural numbers, a , b and c , such that $a^2 + b^2 = c^2$. Prove that, in a Pythagorean triple, at least one of a and b is even. Use either a proof by contradiction or a proof by contraposition.
9. Suppose you have 2 pairs of positive real numbers whose products are 1. That is, you have (a, b) and (c, d) in \mathbb{R}^2 satisfying $ab = cd = 1$. Prove that $a < c$ implies that $b > d$.

3.4 Disproofs

The idea of a “disproof” is really just semantics – in order to disprove a statement we need to *prove* its negation.

So far we’ve been discussing proofs quite a bit, but have paid very little attention to a really huge issue. If the statements we are attempting to prove are false, no proof is ever going to be possible. Really, a prerequisite to developing a facility with proofs is developing a good “lie detector.” We need to be able to guess, or quickly ascertain, whether a statement is true or false. If we are given a universally quantified statement the first thing to do is try it out for some random elements of the universe we’re working in. If we happen across a value that satisfies the statement’s hypotheses but doesn’t satisfy the conclusion, we’ve found what is known as a *counterexample*.

Consider the following statement about integers and divisibility:

Conjecture 3.

$$\forall a, b, c \in \mathbb{Z}, a|bc \implies a|b \vee a|c.$$

This is phrased as a UCS, so the hypothesis is clear, we’re looking for three integers so that the first divides the product of the other two. In the following table we have collected several values for a , b and c such that $a|bc$.

a	b	c	$a b \vee a c ?$
2	7	6	yes
2	4	5	yes
3	12	11	yes
3	5	15	yes
5	4	15	yes
5	10	3	yes
7	2	14	yes

Exercise. As noted in Section 1.2 the statement above is related to whether or not a is prime. Note that in the table, only prime values of a appear. This is a rather broad hint. Find a counterexample to Conjecture 3.

There can be times when the search for a counterexample starts to feel really futile. Would you think it likely that a statement about natural numbers could be true for (more than) the first 50 numbers and yet still be false?

Conjecture 4.

$$\forall n \in \mathbb{Z}^+ \quad n^2 - 79n + 1601 \text{ is prime.}$$

Exercise. Find a counterexample to Conjecture 4

Hidden within Euclid's proof of the infinitude of the primes is a sequence. Recall that in the proof we deduced a contradiction by considering the number N defined by

$$N = 1 + \prod_{k=1}^n p_k.$$

Define a sequence by

$$N_n = 1 + \prod_{k=1}^n p_k,$$

where $\{p_1, p_2, \dots, p_n\}$ are the actual first n primes. The first several values of this sequence are:

n	N_n
1	$1 + (2) = 3$
2	$1 + (2 \cdot 3) = 7$
3	$1 + (2 \cdot 3 \cdot 5) = 31$
4	$1 + (2 \cdot 3 \cdot 5 \cdot 7) = 211$
5	$1 + (2 \cdot 3 \cdot 5 \cdot 7 \cdot 11) = 2311$
\vdots	\vdots

Now, in the proof, we deduced a contradiction by noting that N_n is much larger than p_n , so if p_n is the largest prime it follows that N_n can't be prime – but what really appears to be the case (just look at that table!) is that N_n actually *is* prime for all n .

Exercise. Find a counterexample to the conjecture that $1 + \prod_{k=1}^n p_k$ is itself always a prime.

Exercises — 3.4

1. Find a polynomial that assumes only prime values for a reasonably large range of inputs.
2. Find a counterexample to the conjecture that $\forall a, b, c \in \mathbb{Z}, a | bc \implies a | b \vee a | c$ using only powers of 2.
3. The alternating sum of factorials provides an interesting example of a sequence of integers.

$$1! = 1$$

$$2! - 1! = 1$$

$$3! - 2! + 1! = 5$$

$$4! - 3! + 2! - 1! = 19$$

et cetera

Are they all prime? (After the first two 1's.)

4. It has been conjectured that whenever p is prime, $2^p - 1$ is also prime. Find a minimal counterexample.
5. True or false: The sum of any two irrational numbers is irrational. Prove your answer.
6. True or false: There are two irrational numbers whose sum is rational. Prove your answer.
7. True or false: The product of any two irrational numbers is irrational. Prove your answer.

8. True or false: There are two irrational numbers whose product is rational. Prove your answer.
9. True or false: Whenever an integer n is a divisor of the square of an integer, m^2 , it follows that n is a divisor of m as well. (In symbols, $\forall n \in \mathbb{Z}, \forall m \in \mathbb{Z}, n \mid m^2 \implies n \mid m$.) Prove your answer.
10. In an exercise in Section 3.2 we proved that the quadratic equation $ax^2 + bx + c = 0$ has two solutions if $ac < 0$. Find a counterexample which shows that this implication cannot be replaced with a biconditional.

3.5 Even more direct proofs: By cases and By exhaustion

Proof by exhaustion is the least attractive proof method from an aesthetic perspective. An exhaustive proof consists of literally (and exhaustively) checking every element of the universe to see if the given statement is true for it. Usually, of course, this is impossible because the universe of discourse is infinite; but when the universe of discourse is finite, one certainly can't argue the validity of an exhaustive proof.

In the last few decades the introduction of powerful computational assistance for mathematicians has led to a funny situation. There is a growing list of important results that have been “proved” by exhaustion using a computer. Important examples of this phenomenon are the non-existence of a projective plane of order 10[10] and the only known value of a Ramsey number for hypergraphs[13].

Proof by cases is subtly different from exhaustive proof – for one thing a valid proof by cases can be used in an infinite universe. In a proof by cases one has to divide the universe of discourse into a finite number of sets⁵ and then provide a separate proof for each of the cases. A great many statements about the integers can be proved using the division of integers into even and odd. Another set of cases that is used frequently is the finite number of possible remainders obtained when dividing by an integer d . (Note that even and odd correspond to the remainders 0 and 1 obtained after division by 2.)

A very famous instance of proof by cases is the computer-assisted proof of the four color theorem. The four color theorem is a result known to map makers for quite some time that says that 4 colors are always sufficient to color the nations on a map in such a way that countries sharing a boundary are always colored differently. Figure 3.1 shows one instance of an arrange-

⁵It is necessary to provide an argument that this list of cases is complete! I.e. that every element of the universe falls into one of the cases.

ment of nations that requires at least four different colors, the theorem says that four colors are *always* enough. It should be noted that real cartographers usually reserve a fifth color for oceans (and other water) and that it is possible to conceive of a map requiring five colors if one allows the nations to be non-contiguous. In 1977, Kenneth Appel and Wolfgang Haken proved the four color theorem by reducing the infinitude of possibilities to 1,936 separate cases and analyzing each of these with a computer. The inelegance of a proof by cases is probably proportional to some power of the number of cases, but in any case, this proof is generally considered somewhat inelegant. Ever since the proof was announced there has been an ongoing effort to reduce the number of cases (currently the record is 633 cases – still far too many to be checked through without a computer) or to find a proof that does not rely on cases. For a good introductory article on the four color theorem see[6].

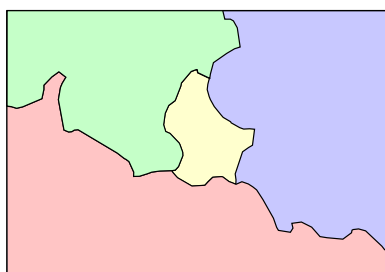


Figure 3.1: The nations surrounding Luxembourg show that sometimes 4 colors are required in cartography.

Most exhaustive proofs of statements that aren't trivial tend to either be (literally) too exhausting or to seem rather contrived. One example of a situation in which an exhaustive proof of some statement exists is when the statement is thought to be universally true but no general proof is known – yet the statement has been checked for a large number of cases. Goldbach's conjecture is one such statement. Christian Goldbach [4] was a mathemati-

cian born in Königsberg Prussia, who, curiously, did *not* make the conjecture⁶ which bears his name. In a letter to Leonard Euler, Goldbach conjectured that every odd number greater than 5 could be expressed as the sum of three primes (nowadays this is known as the weak Goldbach conjecture). Euler apparently liked the problem and replied to Goldbach stating what is now known as Goldbach's conjecture: Every even number greater than 2 can be expressed as the sum of two primes. This statement has been lying around since 1742, and a great many of the world's best mathematicians have made their attempts at proving it – to no avail! (Well, actually a lot of progress has been made but the result still hasn't been proved.) It's easy to verify the Goldbach conjecture for relatively small even numbers, so what *has* been done is/are proofs by exhaustion of Goldbach's conjecture restricted to finite universes. As of this writing, the conjecture has been verified to be true of all even numbers less than 2×10^{17} .

Whenever an exhaustive proof, or a proof by cases exists for some statement it is generally felt that a direct proof would be more esthetically pleasing. If you are in a situation that doesn't admit such a direct proof, you should at least seek a proof by cases using the minimum possible number of cases. For example, consider the following theorem and proof.

Theorem 3.5.1. $\forall n \in \mathbb{Z}$ n^2 is of the form $4k$ or $4k + 1$ for some $k \in \mathbb{Z}$.

Proof: We will consider the four cases determined by the four possible residues mod 4.

case i) If $n \equiv 0 \pmod{4}$ then there is an integer m such that $n = 4m$. It follows that $n^2 = (4m)^2 = 16m^2$ is of the form $4k$ where k is $4m^2$.

case ii) If $n \equiv 1 \pmod{4}$ then there is an integer m such that $n = 4m + 1$. It follows that $n^2 = (4m + 1)^2 = 16m^2 + 8m + 1$ is of the form $4k + 1$ where k is $4m^2 + 2m$.

⁶This conjecture was discussed previously in the exercises of Section 1.2

case iii) If $n \equiv 2 \pmod{4}$ then there is an integer m such that $n = 4m + 2$. It follows that $n^2 = (4m + 2)^2 = 16m^2 + 16m + 4$ is of the form $4k$ where k is $4m^2 + 4m + 1$.

case iv) If $n \equiv 3 \pmod{4}$ then there is an integer m such that $n = 4m + 3$. It follows that $n^2 = (4m + 3)^2 = 16m^2 + 24m + 9$ is of the form $4k + 1$ where k is $4m^2 + 6m + 2$.

Since these four cases exhaust the possibilities and since the desired result holds in each case, our proof is complete.

Q.E.D.

While the proof just stated is certainly valid, the argument is inelegant since a smaller number of cases would suffice.

Exercise. *The previous theorem can be proved using just two cases. Do so.*

We'll close this section by asking you to determine an exhaustive proof where the complexity of the argument is challenging but not *too* impossible.

Graph pebbling is an interesting concept originated by the famous combinatorialist Fan Chung. A “graph” (as the term is used here) is a collection of places or locations which are known as “nodes,” some of which are joined by paths or connections which are known as “edges.” Graphs have been studied by mathematicians for about 400 years, and many interesting problems can be put in this setting. Graph pebbling is a crude version of a broader problem in resource management – often a resource actually gets used in the process of transporting it. Think of the big tanker trucks that are used to transport gasoline. What do they run on? Well, actually they probably burn diesel — but the point is that in order to move the fuel around we have to consume some of it. Graph pebbling takes this to an extreme: in order to move one pebble we must consume one pebble.

Imagine that a bunch of pebbles are randomly distributed on the nodes of a graph, and that we are allowed to do *graph pebbling moves* – we remove two pebbles from some node and place a single pebble on a node that is connected to it. See Figure 3.3.

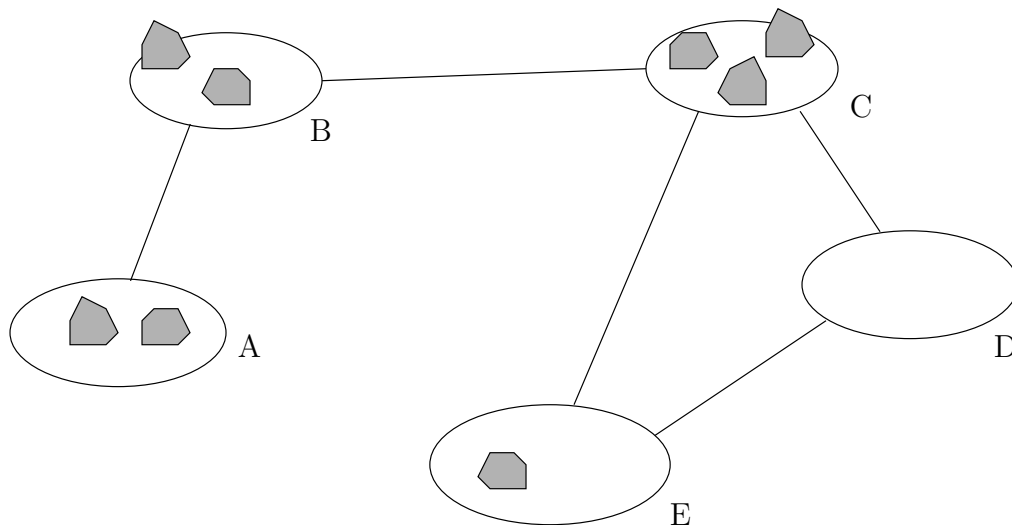


Figure 3.2: In graph pebbling problems a collection of pebbles are distributed on the nodes of a graph. There is no significance to the particular graph that is shown here, or to the arrangement of pebbles – we are just giving an example.

For any particular graph, we can ask for its *pebbling number*, ρ . This is the smallest number so that if ρ pebbles are distributed *in any way whatsoever* on the nodes of the graph, it will be possible to use pebbling moves so as to get a pebble to any node.

For example, consider the triangle graph – three nodes which are all mutually connected. The pebbling number of this graph is 3. If we start with one pebble on each node we are already done; if there is a node that has two pebbles on it, we can use a pebbling move to reach either of the other two nodes.

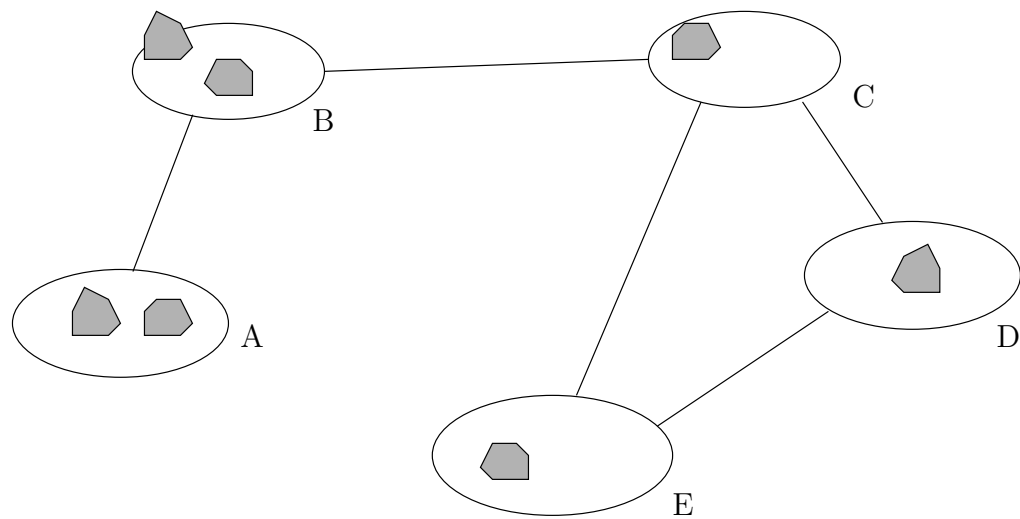


Figure 3.3: A graph pebbling move takes two pebbles off of a node and puts one of them on an adjacent node (the other is discarded). Notice how node C, which formerly held 3 pebbles, now has only 1 and that a pebble is now present on node D where previously there was none.

Exercise. *There is a graph C_5 which consists of 5 nodes connected in a circular fashion. Determine its pebbling number. Prove your answer exhaustively.*

Hint: the pebbling number must be greater than 4 because if one pebble is placed on each of 4 nodes the configuration is unmovable (we need to have two pebbles on a node in order to be able to make a pebbling move at all) and so the 5th node can never be reached.

Exercises — 3.5

1. Prove that if n is an odd number then $n^4 \pmod{16} = 1$.
2. Prove that every prime number other than 2 and 3 has the form $6q + 1$ or $6q + 5$ for some integer q . (Hint: this problem involves thinking about cases as well as contrapositives.)
3. Show that the sum of any three consecutive integers is divisible by 3.
4. There is a graph known as K_4 that has 4 nodes and there is an edge between every pair of nodes. The pebbling number of K_4 has to be at least 4 since it would be possible to put one pebble on each of 3 nodes and not be able to reach the remaining node using pebbling moves. Show that the pebbling number of K_4 is actually 4.
5. Find the pebbling number of a graph whose nodes are the corners and whose edges are the, uhmm, edges of a cube.
6. A *vampire number* is a $2n$ digit number v that factors as $v = xy$ where x and y are n digit numbers and the digits of v are the union of the digits in x and y in some order. The numbers x and y are known as the “fangs” of v . To eliminate trivial cases, pairs of trailing zeros are disallowed.

Show that there are no 2-digit vampire numbers.

Show that there are seven 4-digit vampire numbers.
7. Lagrange’s theorem on representation of integers as sums of squares says that every positive integer can be expressed as the sum of at most 4 squares. For example, $79 = 7^2 + 5^2 + 2^2 + 1^2$. Show (exhaustively) that 15 can not be represented using fewer than 4 squares.
8. Show that there are exactly 15 numbers x in the range $1 \leq x \leq 100$ that can’t be represented using fewer than 4 squares.

9. The *trichotomy property* of the real numbers simply states that every real number is either positive or negative or zero. Trichotomy can be used to prove many statements by looking at the three cases that it guarantees. Develop a proof (by cases) that the square of any real number is non-negative.

10. Consider the game called “binary determinant tic-tac-toe”⁷ which is played by two players who alternately fill in the entries of a 3×3 array. Player One goes first, placing 1’s in the array and player Zero goes second, placing 0’s. Player One’s goal is that the final array have determinant 1, and player Zero’s goal is that the determinant be 0. The determinant calculations are carried out mod 2.

Show that player Zero can always win a game of binary determinant tic-tac-toe by the method of exhaustion.

⁷ This question was problem A4 in the 63rd annual William Lowell Putnam Mathematics Competition (2002). There are three collections of questions and answers from previous Putnam exams available from the MAA [1, 7, 9]

3.6 Proofs and disproofs of existential statements

From a certain point of view, there is no need for the current section. If we are proving an existential statement we are *disproving* some universal statement. (Which has already been discussed.) Similarly, if we are trying to disprove an existential statement, then we are actually *proving* a related universal statement. Nevertheless, sometimes the way a theorem is stated emphasizes the existence question over the corresponding universal – and so people talk about proving and disproving existential statements as a separate issue from universal statements.

Proofs of existential questions come in two basic varieties: constructive and non-constructive. Constructive proofs are conceptually the easier of the two – you actually name an example that shows the existential question is true. For example:

Theorem 3.6.1. *There is an even prime.*

Proof: The number 2 is both even and prime.

Q.E.D.

Exercise. *The Fibonacci numbers are defined by the initial values $F(0) = 1$ and $F(1) = 1$ and the recursive formula $F(n + 1) = F(n) + F(n - 1)$ (to get the next number in the series you add the last and the penultimate).*

n	$F(n)$
0	1
1	1
2	2
3	3
4	5
5	8
⋮	⋮

Prove that there is a Fibonacci number that is a perfect square.

A non-constructive existence proof is trickier. One approach is to argue by contradiction – if the thing we’re seeking doesn’t exist that will lead to an absurdity. Another approach is to outline a search algorithm for the desired item and provide an argument as to why it cannot fail!

A particularly neat approach is to argue using dilemma. This is my favorite non-constructive existential theorem/proof.

Theorem 3.6.2. *There are irrational numbers α and β such that α^β is rational.*

Proof: If $\sqrt{2}^{\sqrt{2}}$ is rational then we are done. (Let $\alpha = \beta = \sqrt{2}$.) Otherwise, let $\alpha = \sqrt{2}^{\sqrt{2}}$ and $\beta = \sqrt{2}$. The result follows because $\left(\sqrt{2}^{\sqrt{2}}\right)^{\sqrt{2}} = \sqrt{2}^{(\sqrt{2}\sqrt{2})} = \sqrt{2}^2 = 2$, which is clearly rational.

Q.E.D.

Many existential proofs involve a property of the natural numbers known as the well-ordering principle. The well-ordering principle is sometimes abbreviated WOP. If a set has WOP it doesn’t mean that the set is ordered in a particularly good way, but rather that its subsets are like wells – the kind one hoists water out of with a bucket on a rope. You needn’t be concerned with WOP in general at this point, but notice that the subsets of the natural numbers have a particularly nice property – any non-empty set of natural numbers must have a least element (much like every water well has a bottom).

Because the natural numbers have the well-ordering principle we can prove that there is a least natural number with property X by simply finding *any* natural number with property X – by doing that we’ve shown that the set of natural numbers with property X is non-empty and that’s the only hypothesis the WOP needs.

For example, in the exercises in Section 3.5 we introduced vampire numbers. A *vampire number* is a $2n$ digit number v that factors as $v = xy$ where x and y are n digit numbers and the digits of v are the union of the digits in x and y in some order. The numbers x and y are known as the “fangs” of v . To eliminate trivial cases, pairs of trailing zeros are disallowed.

Theorem 3.6.3. *There is a smallest 6-digit vampire number.*

Proof: The number 125460 is a vampire number (in fact this is the smallest example of a vampire number with two sets of fangs: $125460 = 204 \cdot 615 = 246 \cdot 510$). Since the set of 6-digit vampire numbers is non-empty, the well-ordering principle of the natural numbers allows us to deduce that there is a smallest 6-digit vampire number.

Q.E.D.

This is quite an interesting situation in that we know there is a smallest 6-digit vampire number without having any idea what it is!

Exercise. *Show that 102510 is the smallest 6-digit vampire number.*

There are quite a few occasions when we need to prove statements involving the unique existence quantifier ($\exists!$). In such instances we need to do just a little bit more work. We need to show existence – either constructively or non-constructively – and we also need to show uniqueness. To give an example of a unique existence proof we’ll return to a concept first discussed in Section 1.5 and finish-up some business that was glossed-over there.

Recall the Euclidean algorithm that was used to calculate the greatest common divisor of two integers a and b (which we denote $\gcd(a, b)$). There is a rather important question concerning algorithms known as the “halting problem.” Does the program eventually halt, or does it get stuck in an infinite loop? We know that the Euclidean algorithm halts (and outputs the correct result) because we know the following unique existence result.

$$\forall a, b \in \mathbb{Z}^+, \exists! d \in \mathbb{Z}^+ \text{ such that } d = \gcd(a, b)$$

Now, before we can prove this result, we'll need a precise definition for $\gcd(a, b)$. Firstly, a gcd must be a *common divisor* which means it needs to divide both a and b . Secondly, among all the common divisors, it must be the *largest*. This second point is usually addressed by requiring that every other common divisor divides the gcd. Finally we should note that a gcd is always positive, for whenever a number divides another number so does its negative, and whichever of those two is positive will clearly be the greater! This allows us to extend the definition of gcd to all integers, but things are conceptually easier if we keep our attention restricted to the positive integers.

Definition. *The greatest common divisor, or gcd, of two positive integers a and b is a positive integer d such that $d \mid a$ and $d \mid b$ and if c is any other positive integer such that $c \mid a$ and $c \mid b$ then $c \mid d$.*

$$\forall a, b, c, d \in \mathbb{Z}^+ \quad d = \gcd(a, b) \iff d \mid a \wedge d \mid b \wedge (c \mid a \wedge c \mid b \implies c \mid d)$$

Armed with this definition, let's return our attention to proving the unique existence of the gcd. The uniqueness part is easier so we'll do that first. We argue by contradiction. Suppose that there were two different numbers d and d' satisfying the definition of $\gcd(a, b)$. Put d' in the place of c in the definition to see that $d' \mid d$. Similarly, we can deduce that $d \mid d'$ and if two numbers each divide into the other, they must be equal. This is a contradiction since we assumed d and d' were different.

For the existence part we'll need to define a set – known as the \mathbb{Z} -module generated by a and b – that consists of all numbers of the form $xa + yb$ where x and y range over the integers.

This set has a very nice geometric character that often doesn't receive the attention it deserves. Every element of a \mathbb{Z} -module generated by two

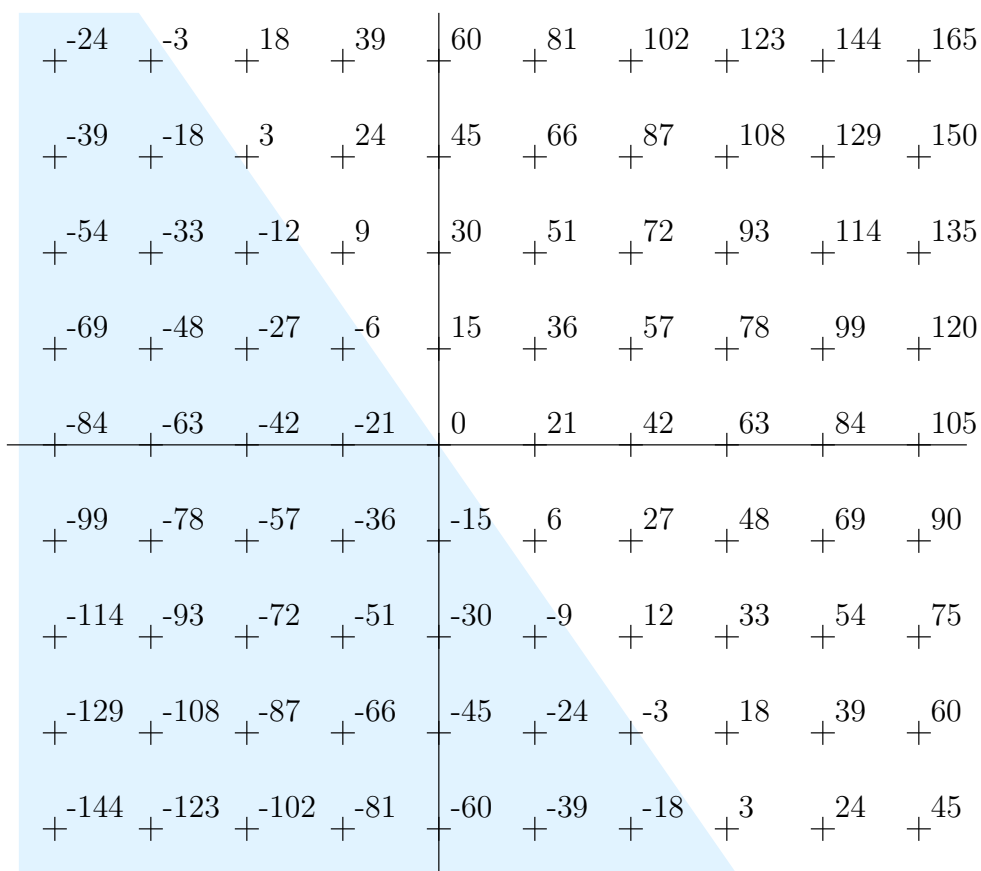


Figure 3.4: The \mathbb{Z} -module generated by 21 and 15. The number $21x + 15y$ is printed by the point (x, y) .

numbers (15 and 21 in the example) corresponds to a point in the Euclidean plane. As indicated in Figure 3.4 there is a dividing line between the positive and negative elements in a \mathbb{Z} -module. It is also easy to see that there are many repetitions of the same value at different points in the plane.

Exercise. *The value 0 clearly occurs in a \mathbb{Z} -module when both x and y are themselves zero. Find another pair of (x, y) values such that $21x + 15y$ is zero. What is the slope of the line which separates the positive values from the negative in our \mathbb{Z} -module?*

In thinking about this \mathbb{Z} -module, and perusing Figure 3.4, you may have noticed that the smallest positive number in the \mathbb{Z} -module is 3. If you hadn't noticed that, look back and verify that fact now.

Exercise. *How do we know that some smaller positive value (a 1 or a 2) doesn't occur somewhere in the Euclidean plane?*

What we've just observed is a particular instance of a general result.

Theorem 3.6.4. *The smallest positive number in the \mathbb{Z} -module generated by a and b is $d = \gcd(a, b)$.*

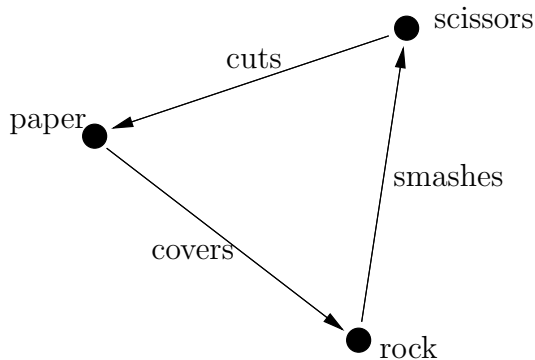
Proof: Suppose that d is the smallest positive number in the \mathbb{Z} -module $\{xa + yb \mid x, y \in \mathbb{Z}\}$. There are particular values of x and y (which we will distinguish with over-lines) such that $d = \bar{x}a + \bar{y}b$. Now, it is easy to see that if c is any common divisor of a and b then $c \mid d$, so what remains to be proved is that d itself is a divisor of both a and b . Consider dividing d into a . By the division algorithm there are uniquely determined numbers q and r such that $a = qd + r$ with $0 \leq r < d$. We will show that $r = 0$. Suppose, to the contrary, that r is positive. Note that we can write r as $r = a - qd = a - q(\bar{x}a + \bar{y}b) = (1 - q\bar{x})a - q\bar{y}b$. The last equality shows that r is in the \mathbb{Z} -module under consideration, and so, since

d is the smallest positive integer in this \mathbb{Z} -module it follows that $r \geq d$ which contradicts the previously noted fact that $r < d$. Thus, $r = 0$ and so it follows that $d \mid a$. An entirely analogous argument can be used to show that $d \mid b$ which completes the proof that $d = \gcd(a, b)$.

Q.E.D.

Exercises — 3.6

1. Show that there is a perfect square that is the sum of two perfect squares.
2. Show that there is a perfect cube that is the sum of three perfect cubes.
3. Show that the WOP doesn't hold in the integers. (This is an existence proof, you show that there is a subset of \mathbb{Z} that doesn't have a smallest element.)
4. Show that the WOP doesn't hold in \mathbb{Q}^+ .
5. In the proof of Theorem 3.6.4 we weaseled out of showing that $d \mid b$. Fill in that part of the proof.
6. Give a proof of the unique existence of q and r in the division algorithm.
7. A *digraph* is a drawing containing a collection of points that are connected by arrows. The game known as *scissors-paper-rock* can be represented by a digraph that is *balanced* (each point has the same number of arrows going out as going in). Show that there is a balanced digraph having 5 points.



Chapter 4

Sets

No more turkey, but I'd like some more of the bread it ate. –Hank Ketcham

4.1 Basic notions of set theory

In modern mathematics there is an area called Category theory¹ which studies the relationships between different areas of mathematics. More precisely, the founders of category theory noticed that essentially the same theorems and proofs could be found in many different mathematical fields – with only the names of the structures involved changed. In this sort of situation one can make what is known as a *categorical* argument in which one proves the desired result in the abstract, without reference to the details of any particular field. In effect this allows one to prove many theorems at once – all you need to convert an abstract categorical proof into a concrete one relevant to a particular area is a sort of key or lexicon to provide the correct names for things. Now, category theory probably shouldn't really be studied until you have a background that includes enough different fields that you can make sense of their categorical correspondences. Also, there are a good many

¹The classic text by Saunders Mac Lane [11] is still considered one of the best introductions to Category theory.

mathematicians who deride category theory as “abstract nonsense.” But, as someone interested in developing a facility with proofs, you should be on the lookout for categorical correspondences. If you ever hear yourself utter something like “well, the proof of *that* goes just like the proof of the (insert weird technical-sounding name here) theorem” you are probably noticing a categorical correspondence.

Okay, so category theory won’t be of much use to you until much later in your mathematical career (if at all), and one could argue that it doesn’t really save that much effort. Why not just do two or three different proofs instead of learning a whole new field so we can combine them into one? Nevertheless, category theory is being mentioned here at the beginning of the chapter on sets. Why?

We are about to see our first example of a categorical correspondence. Logic and Set theory are different aspects of the same thing. To describe a set people often quote Kurt Gödel – “A set is a Many that allows itself to be thought of as a One.” (Note how the attempt at defining what is really an elemental, undefinable concept ends up sounding rather mystical.) A more practical approach is to think of a set as the collection of things that make some open sentence *true*.²

Recall that in Logic the atomic concepts were “true”, “false”, “sentence” and “statement.” In Set theory, they are “set”, “element” and “membership.” These concepts (more or less) correspond to one another. In most books, a set is denoted either using the letter M (which stands for the German word “menge”) or early alphabet capital roman letters – A , B , C , *et cetera*. Here, we will often emphasize the connection between sets and open sentences in Logic by using a subscript notation. The set that corresponds to the open sentence $P(x)$ will be denoted S_P , we call S_P the *truth set* of $P(x)$.

²This may sound less metaphysical, but this statement is also faulty because it defines “set” in terms of “collection” – which will of course be defined elsewhere as “the sort of things of which sets are one example.”

$$S_P = \{x \mid P(x)\}$$

On the other hand, when we have a set given in the absence of any open sentence, we'll be happy to use the early alphabet, capital roman letters convention – or frankly, any other letters we feel like! Whenever we have a set A given, it is easy to state a logical open sentence that would correspond to it. The membership question: $M_A(x) =$ “Is x in the set A ?” Or, more succinctly, $M_A(x) = “x \in A”$. Thus the atomic concept “true” from Logic corresponds to the answer “yes” to the membership question in Set theory (and of course “false” corresponds to “no”).

There are many interesting foundational issues which we are going to sidestep in our current development of Set theory. For instance, recall that in Logic we always worked inside some “universe of discourse.” As a consequence of the approach we are taking now, all of our set theoretic work will be done within some unknown “universal” set. Attempts at specifying (*a priori*) a universal set for doing mathematics within are doomed to failure. In the early days of the twentieth century they attempted to at least get Set theory itself on a firm footing by defining the universal set to be “the set of all sets” – an innocuous sounding idea that had funny consequences (we'll investigate this in Section 4.5).

In Logic we had “sentences” and “statements,” the latter were distinguished as having definite truth values. The corresponding thing in Set theory is that sets have the property that we can always tell whether a given object is or is not in them. If it ever becomes necessary to talk about “sets” where we're not really sure what's in them we'll use the term *collection*.

You should think of a set as being an *unordered* collection of things, thus $\{\text{popover}, 1, \text{froggy}\}$ and $\{1, \text{froggy}, \text{popover}\}$ are two ways to represent the same set. Also, a set either contains, or doesn't contain, a given element. It doesn't make sense to have an element in a set multiple times. By convention, if an element is listed more than once when a set is listed we ignore

the repetitions. So, the sets $\{1, 1\}$ and $\{1\}$ are really the same thing. If the notion of a set containing multiple instances of its elements is needed there is a concept known as a *multiset* that is studied in Combinatorics. In a multiset, each element is preceded by a so-called *repetition number* which may be the special symbol ∞ (indicating an unlimited number of repetitions). The multiset concept is useful when studying puzzles like “How many ways can the letters of MISSISSIPPI be rearranged?” because the letters in MISSISSIPPI can be expressed as the multiset $\{1 \cdot M, 4 \cdot I, 2 \cdot P, 4 \cdot S\}$. With the exception of the following exercise, in the remainder of this chapter we will only be concerned with sets, never multisets.

Exercise. (*Not for the timid!*) *How many ways can the letters of MISSISSIPPI be arranged?*

If a computer scientist were seeking a data structure to implement the notion of “set,” he’d want a sorted list where repetitions of an entry were somehow disallowed. We’ve already noted that a set should be thought of as an unordered collection, and yet it’s been asserted that a *sorted* list would be the right vehicle for representing a set on a computer. Why? One reason is that we’d like to be able to tell (quickly) whether two sets are the same or not. If the elements have been presorted it’s easier.

Consider the difficulty in deciding whether the following two sets are equal.

$$S_1 = \{\spadesuit, 1, e, \pi, \diamond, A, \Omega, h, \oplus, \epsilon\}$$

$$S_2 = \{A, 1, \epsilon, \pi, e, s, \oplus, \spadesuit, \Omega, \diamond\}$$

If instead we compare them after they've been sorted, the job is much easier.

$$S_1 = \{1, A, \diamond, e, \epsilon, h, \Omega, \oplus, \pi, \spadesuit\}$$

$$S_2 = \{1, A, \diamond, e, \epsilon, \Omega, \oplus, \pi, s, \spadesuit\}$$

This business about ordered versus unordered comes up fairly often so it's worth investing a few moments to figure out how it works. If a collection of things that is inherently unordered is handed to us we generally *put* them in an order that is pleasing to us. Consider receiving five cards from the dealer in a card game, or extracting seven letters from the bag in a game of Scrabble. If, on the other hand, we receive a collection where order is important we certainly *may not* rearrange them. Imagine someone receiving the telephone number of an attractive other but writing it down with the digits sorted in increasing order!

Exercise. Consider a universe consisting of just the first 5 natural numbers $U = \{1, 2, 3, 4, 5\}$. How many different sets having 4 elements are there in this universe? How many different ordered collections of 4 elements are there?

The last exercise suggests an interesting question. If you have a universal set of some fixed (finite) size, how many different sets are there? Obviously you can't have any more elements in a set than are in your universe. What's the smallest possible size for a set? Many people would answer 1 – which isn't unreasonable! – after all a set is supposed to be a collection of things, and is it really possible to have a *collection* with nothing in it? The standard answer is 0 however, mostly because it makes a certain counting formula work out nicely. A set with one element is known as a *singleton set* (note the use of the indefinite article). A set with no elements is known as the *empty set* (note the definite article). There are as many singletons as there

are elements in your universe. They aren't the same though, for example $1 \neq \{1\}$. There is only one empty set and it is denoted \emptyset – irrespective of the universe we are working in.

Let's have a look at a small example. Suppose we have a universal set with 3 elements, without loss of generality, $\{1, 2, 3\}$. It's possible to construct a set, whose elements are all the possible sets in this universe. This set is known as the *power set* of the universal set. Indeed, we can construct the power set of *any* set A and we denote it with the symbol $\mathcal{P}(A)$. Returning to our example we have

$$\begin{aligned} \mathcal{P}(\{1, 2, 3\}) = \{ & \emptyset, \\ & \{1\}, \{2\}, \{3\}, \\ & \{1, 2\}, \{1, 3\}, \{2, 3\}, \\ & \{1, 2, 3\} \}. \end{aligned}$$

Exercise.

Find the power sets $\mathcal{P}(\{1, 2\})$ and $\mathcal{P}(\{1, 2, 3, 4\})$.

Conjecture a formula for the number of elements (these are, of course, sets) in $\mathcal{P}(\{1, 2, \dots, n\})$.

Hint: If your conjectured formula is correct you should see why these sets are named as they are.

One last thing before we end this section. The size (a.k.a. cardinality) of a set is just the number of elements in it. We use the very same symbol for cardinality as we do for the absolute value of a numerical entity. There should really never be any confusion. If A is a set then $|A|$ means that we should count how many things are in A . If A isn't a set then we are talking about the ordinary absolute value

Exercises — 4.1

1. What is the power set of \emptyset ? Hint: if you got the last exercise in the chapter you'd know that this power set has $2^0 = 1$ element.
2. Try iterating the power set operator. What is $\mathcal{P}(\mathcal{P}(\emptyset))$? What is $\mathcal{P}(\mathcal{P}(\mathcal{P}(\emptyset)))$?
3. Determine the following cardinalities.
 - (a) $A = \{1, 2, \{3, 4, 5\}\}$ $|A| = \underline{\hspace{2cm}}$
 - (b) $B = \{\{1, 2, 3, 4, 5\}\}$ $|B| = \underline{\hspace{2cm}}$
4. What, in Logic, corresponds the notion \emptyset in Set theory?
5. What, in Set theory, corresponds to the notion t (a tautology) in Logic?
6. What is the truth set of the proposition $P(x) =$ “3 divides x and 2 divides x ”?
7. Find a logical open sentence such that $\{0, 1, 4, 9, \dots\}$ is its truth set.
8. How many singleton sets are there in the power set of $\{a, b, c, d, e\}$? “Doubleton” sets?
9. How many 8 element subsets are there in

$$\mathcal{P}(\{a, b, c, d, e, f, g, h, i, j, k, l, m, n, o, p\})?$$

10. How many singleton sets are there in the power set of $\{1, 2, 3, \dots, n\}$?

4.2 Containment

There are two notions of being “inside” a set. A thing may be an *element* of a set, or may be contained as a subset. Distinguishing these two notions of inclusion is essential. One difficulty that sometimes complicates things is that a set may contain other sets *as elements*. For instance, as we saw in the previous section, the elements of a power set are themselves sets.

A set A is a *subset* of another set B if all of A 's elements are also in B . The terminology *superset* is used to refer to B in this situation, as in “The set of all real-valued functions in one real variable is a superset of the polynomial functions.” The subset/superset relationship is indicated with a symbol that should be thought of as a stylized version of the less-than-or-equal sign, when A is a subset of B we write $A \subseteq B$.

We say that A is a *proper subset* of B if B has some elements that aren't in A , and in this situation we write $A \subset B$ or if we really want to emphasize the fact that the sets are not equal we can write $A \subsetneq B$. By the way, if you want to emphasize the superset relationship, all of these symbols can be turned around. So for example $A \supseteq B$ means that A is a superset of B although they could potentially be equal.

As we've seen earlier, the symbol \in is used between an element of a set and the set that it's in. The following exercise is intended to clarify the distinction between \in and \subseteq .

Exercise. Let $A = \{1, 2, \{1\}, \{a, b\}\}$. Which of the following are true?

- | | |
|-----------------------------|------------------------------|
| i) $\{a, b\} \subseteq A$. | vi) $\{1\} \subseteq A$. |
| ii) $\{a, b\} \in A$. | vii) $\{1\} \in A$. |
| iii) $a \in A$. | viii) $\{2\} \in A$. |
| iv) $1 \in A$. | ix) $\{2\} \subseteq A$. |
| v) $1 \subseteq A$. | x) $\{\{1\}\} \subseteq A$. |

Another perspective that may help clear up the distinction between \in and \subseteq is to consider what they correspond to in Logic. The “element of” symbol \in is used to construct open sentences that embody the membership question – thus it corresponds to single sentences in Logic. The “set containment” symbol \subseteq goes between two *sets* and so whatever it corresponds to in Logic should be something that can appropriately be inserted between two sentences. Let’s run through a short example to figure out what that might be. To keep things simple we’ll work inside the universal set $U = \{1, 2, 3, \dots, 50\}$. Let T be the subset of U consisting of those numbers that are divisible by 10, and let F be those that are divisible by 5.

$$T = \{10, 20, 30, 40, 50\}$$

$$F = \{5, 10, 15, 20, 25, 30, 35, 40, 45, 50\}$$

Hopefully it is clear that \subseteq can be inserted between these two sets like so: $T \subseteq F$. On the other hand we can re-express the sets T and F using set-builder notation in order to see clearly what their membership questions are.

$$T = \{x \in U \mid 10 \mid x\}$$

$$F = \{x \in U \mid 5 \mid x\}$$

What logical operator fits nicely between $10 \mid x$ and $5 \mid x$? Well, of course, it’s the implication arrow. It’s easy to verify that $10 \mid x \implies 5 \mid x$, and it’s equally easy to note that the other direction doesn’t work, $5 \mid x \not\implies 10 \mid x$ — for instance, 5 goes evenly into 15, but 10 doesn’t.

The general statement is: if A and B are sets, and $M_A(x)$ and $M_B(x)$ are their respective membership questions, then $A \subseteq B$ corresponds precisely to $\forall x \in U, M_A(x) \implies M_B(x)$.

Now to many people (me included!) this looks funny at first, \subseteq in Set theory corresponds to \implies in Logic. It seems like both of these symbols

are arrows of a sort – but they point in opposite directions! Personally, I resolve the apparent discrepancy by thinking about the “strength” of logical predicates. One predicate is stronger than another if it puts more conditions on the elements that would make it true. For example, “ x is doubly-even” is stronger than “ x is (merely) even.” Now, the stronger statement implies the weaker (assuming of course that they are stronger and weaker versions of the same idea). If a number is doubly-even (i.e. divisible by 4) then it is certainly even – but the converse is certainly not true, 6 is even but *not* doubly-even. Think of all this in terms of sets now. Which set contains the other, the set of doubly-even numbers or the set of even numbers? Clearly the set that corresponds to more stringent membership criteria is smaller than the set that corresponds to less restrictive criteria, thus the set defined by a weak membership criterion contains the one having a stronger criterion.

If we are asked to prove that one set is contained in another as a subset, $A \subseteq B$, there are two ways to proceed. We may either argue by thinking about elements, or (although this amounts to the same thing) we can show that A 's membership criterion implies B 's membership criterion.

Exercise. Consider S , the set of perfect squares and F , the set of perfect fourth powers. Which is contained in the other? Can you prove it?

We'll end this section with a fairly elementary proof – mainly just to illustrate how one should proceed in proving that one set is contained in another.

Let D represent the set of all integers that are divisible by 9,

$$D = \{x \in \mathbb{Z} \mid \exists k \in \mathbb{Z}, x = 9k\}.$$

Let C represent the set of all integers that are divisible by 3,

$$C = \{x \in \mathbb{Z} \mid \exists k \in \mathbb{Z}, x = 3k\}.$$

The set D is contained in C . Let's prove it!

Proof: Suppose that x is an arbitrary element of D . From the definition of D it follows that there is an integer k such that $x = 9k$. We want to show that $x \in C$, but since $x = 9k$ it is easy to see that $x = 3(3k)$ which shows (since $3k$ is clearly an integer) that x is in C .

Q.E.D.

Exercises — 4.2

1. Insert either \in or \subseteq in the blanks in the following sentences (in order to produce true sentences).
 - i) 1 _____ $\{3, 2, 1, \{a, b\}\}$
 - ii) $\{a\}$ _____ $\{a, \{a, b\}\}$
 - iii) $\{a, b\}$ _____ $\{3, 2, 1, \{a, b\}\}$
 - iv) $\{\{a, b\}\}$ _____ $\{a, \{a, b\}\}$
2. Suppose that p is a prime, for each n in \mathbb{Z}^+ , define the set $P_n = \{x \in \mathbb{Z}^+ \mid p^n \mid x\}$. Conjecture and prove a statement about the containments between these sets.
3. Provide a counterexample to dispel the notion that a subset must have fewer elements than its superset.
4. We have seen that $A \subseteq B$ corresponds to $M_A \implies M_B$. What corresponds to the contrapositive statement?
5. Determine two sets A and B such that both of the sentences $A \in B$ and $A \subseteq B$ are true.
6. Prove that the set of perfect fourth powers is contained in the set of perfect squares.

4.3 Set operations

In this section we'll continue to develop the correspondence between Logic and Set theory.

The logical connectors \wedge and \vee correspond to the set-theoretic notions of union (\cup) and intersection (\cap). The symbols are designed to provide a mnemonic for the correspondence; the Set theory symbols are just rounded versions of those from Logic.

Explicitly, if $P(x)$ and $Q(x)$ are open sentences, then the *union* of the corresponding truth sets S_P and S_Q is defined by

$$S_P \cup S_Q = \{x \in U \mid P(x) \vee Q(x)\}.$$

Exercise. *Suppose two sets A and B are given. Re-express the previous definition of “union” using their membership criteria, $M_A(x) = “x \in A”$ and $M_B(x) = “x \in B.”$*

The union of more than two sets can be expressed using a big union symbol. For example, consider the family of real intervals defined by $I_n = (n, n + 1]$.³ There's an interval for every integer n . Also, every real number is in one of these intervals. The previous sentence can be expressed as

$$\mathbb{R} = \bigcup_{n \in \mathbb{Z}} I_n.$$

The intersection of two sets is conceptualized as “what they have in common” but the precise definition is found by considering conjunctions,

$$A \cap B = \{x \in U \mid x \in A \wedge x \in B\}.$$

Exercise. *With reference to two open sentences $P(x)$ and $Q(x)$, define the intersection of their truth sets, $S_P \cap S_Q$.*

³The elements of I_n can also be distinguished as the solution sets of the inequalities $n < x \leq n + 1$.

There is also a “big” version of the intersection symbol. Using the same family of intervals as before,

$$\emptyset = \bigcap_{n \in \mathbb{Z}} I_n.$$

Of course the intersection of any distinct pair of these intervals is empty so the statement above isn’t particularly strong.

Negation in Logic corresponds to complementation in Set theory. The *complement* of a set A is usually denoted by \overline{A} (although some prefer a superscript c – as in A^c), this is the set of all things that *aren’t* in A . In thinking about complementation one quickly sees why the importance of working within a well-defined universal set is stressed. Consider the set of all math textbooks. Obviously the complement of this set would contain texts in English, Engineering and Evolution – but that statement is implicitly assuming that the universe of discourse is “textbooks.” It’s equally valid to say that a very long sequence of zeros and ones, a luscious red strawberry, and the number $\sqrt{\pi}$ are not math textbooks and so these things are all elements of the complement of the set of all math textbooks. What is really a concern for us is the issue of whether or not the complement of a set is well-defined, that is, can we tell for sure whether a given item is or is not in the complement of a set. This question is decidable exactly when the membership question for the original set is decidable. Many people think that the main reason for working within a fixed universal set is that we then have well-defined complements. The real reason that we accept this restriction is to ensure that both membership criteria, $M_A(x)$ and $M_{\overline{A}}(x)$, are decidable open sentences. As an example of the sort of strangeness that can crop up, consider that during the time that I, as the author of this book, was writing the last paragraph, this text was nothing more than a very long sequence of zeros and ones in the memory of my computer...

Every rule that we learned in Chapter 2 (see Table 2.2) has a set-theoretic

equivalent. These set-theoretic versions are expressed using equalities (i.e. the symbol $=$ in between two sets) which is actually a little bit funny if you think about it. We normally use $=$ to mean that two numbers or variables have the same numerical magnitude, as in $12^2 = 144$, we are doing something altogether different when we use that symbol between two sets, as in $\{1, 2, 3\} = \{\sqrt{1}, \sqrt{4}, \sqrt{9}\}$, but people seem to be used to this so there's no sense in quibbling.

Exercise. *Develop a useful definition for set equality. In other words, come up with a (quantified) logical statement that means the same thing as “ $A = B$ ” for two arbitrary sets A and B .*

Exercise. *What symbol in Logic should go between the membership criteria $M_A(x)$ and $M_B(x)$ if A and B are equal sets?*

In Table 4.1 the rules governing the interactions between the set theoretic operations are collected.

We are now in a position somewhat similar to when we jumped from proving logical assertions with truth tables to doing two-column proofs. We have two different approaches for showing that two sets are equal. We can do a so-called “element chasing” proof (to show $A = B$, assume $x \in A$ and prove $x \in B$ and then vice versa). Or, we can construct a proof using the basic set equalities given in Table 4.1. Often the latter can take the form of a two-column proof.

	Intersection version	Union version
Commutative laws	$A \cap B = B \cap A$	$A \cup B = B \cup A$
Associative laws	$A \cap (B \cap C)$ $= (A \cap B) \cap C$	$A \cup (B \cup C)$ $= (A \cup B) \cup C$
Distributive laws	$A \cap (B \cup C) =$ $(A \cap B) \cup (A \cap C)$	$A \cup (B \cap C) =$ $(A \cup B) \cap (A \cup C)$
DeMorgan's laws	$\overline{A \cap B}$ $= \overline{A} \cup \overline{B}$	$\overline{A \cup B}$ $= \overline{A} \cap \overline{B}$
Double complement	$\overline{\overline{A}} = A$	same
Complementarity	$A \cap \overline{A} = \emptyset$	$A \cup \overline{A} = U$
Identity laws	$A \cap U = A$	$A \cup \emptyset = A$
Domination	$A \cap \emptyset = \emptyset$	$A \cup U = U$
Idempotence	$A \cap A = A$	$A \cup A = A$
Absorption	$A \cap (A \cup B) = A$	$A \cup (A \cap B) = A$

Table 4.1: Basic set theoretic equalities.

Before we proceed much further in our study of set theory it would be a good idea to give you an example. We're going to prove the same assertion in two different ways — once via element chasing and once using the basic set theoretic equalities from Table 4.1.

The statement we'll prove is $A \cup B = A \cup (\bar{A} \cap B)$.

First, by chasing elements:

Proof: Suppose x is an element of $A \cup B$. By the definition of union we know that

$$x \in A \vee x \in B.$$

The conjunctive identity law and the fact that $x \in A \vee x \notin A$ is a tautology gives us an equivalent logical statement:

$$(x \in A \vee x \notin A) \wedge (x \in A \vee x \in B).$$

Finally, this last statement is equivalent to

$$x \in A \vee (x \notin A \wedge x \in B)$$

which is the definition of $x \in A \cup (\bar{A} \cap B)$.

On the other hand, if we assume that $x \in A \cup (\bar{A} \cap B)$, it follows that

$$x \in A \vee (x \notin A \wedge x \in B).$$

Applying the distributive law, disjunctive complementarity and the identity law, in sequence we obtain

$$\begin{aligned}
& x \in A \vee (x \notin A \wedge x \in B) \\
\cong & (x \in A \vee x \notin A) \wedge (x \in A \vee x \in B) \\
\cong & t \wedge (x \in A \vee x \in B) \\
\cong & x \in A \vee x \in B
\end{aligned}$$

The last statement in this chain of logical equivalences provides the definition of $x \in A \cup B$.

Q.E.D.

A two-column proof of the same statement looks like this:

Proof:

$A \cup B$	Given
$= U \cap (A \cup B)$	Identity law
$= (A \cup \bar{A}) \cap (A \cup B)$	Complementarity
$= (A \cup (\bar{A} \cap B))$	Distributive law

Q.E.D.

There are some notions within Set theory that don't have any clear parallels in Logic. One of these is essentially a generalization of the concept of "complements." If you think of the set \bar{A} as being the difference between the universal set U and the set A you are on the right track. The *difference* between two sets is written $A \setminus B$ (sadly, sometimes this is denoted using the ordinary subtraction symbol $A - B$) and is defined by

$$A \setminus B = A \cap \bar{B}.$$

The difference, $A \setminus B$, consists of those elements of A that aren't in B . In some developments of Set theory, the difference of sets is defined first and then complementation is defined by $\overline{A} = U \setminus A$.

The difference of sets (like the difference of real numbers) is not a commutative operation. In other words $A \setminus B \neq B \setminus A$ (in general). It is possible to define an operation that acts somewhat like the difference, but that *is* commutative. The *symmetric difference* of two sets is denoted using a triangle (really a capital Greek delta)

$$A \Delta B = (A \setminus B) \cup (B \setminus A).$$

Exercise. Show that $A \Delta B = (A \cup B) \setminus (A \cap B)$.

Come on! You read right past that exercise without even pausing!

What? You say you *did* try it and it was too hard?

Okay, just for you (and this time only) I've prepared an aid to help you through...

On the next page is a two-column proof of the result you need to prove, but the lines of the proof are all scrambled. Make a copy and cut out all the pieces and then glue them together into a valid proof.

So, no more excuses, just do it!

$= (A \cap \bar{B}) \cup (B \cap \bar{A})$	identity law
$= (A \cup B) \cap \overline{(A \cap B)}$	def. of relative difference
$(A \cup B) \setminus (A \cap B)$	Given
$= ((A \cap \bar{A}) \cup (A \cap \bar{B})) \cup ((B \cap \bar{A}) \cup (B \cap \bar{B}))$	distributive law
$= (A \setminus B) \cup (B \setminus A)$	def. of relative difference
$= (A \cap \overline{(A \cap B)}) \cup (B \cap \overline{(A \cap B)})$	distributive law
$= A \Delta B$	def. of symmetric difference
$= (A \cap (\bar{A} \cup \bar{B})) \cup (B \cap (\bar{A} \cup \bar{B}))$	DeMorgan's law
$= (\emptyset \cup (A \cap \bar{B})) \cup ((B \cap \bar{A}) \cup \emptyset)$	complementarity

Exercises — 4.3

1. Let $A = \{1, 2, \{1, 2\}, b\}$ and let $B = \{a, b, \{1, 2\}\}$. Find the following:

(a) $A \cap B$

(b) $A \cup B$

(c) $A \setminus B$

(d) $B \setminus A$

(e) $A \Delta B$

2. In a standard deck of playing cards one can distinguish sets based on face-value and/or suit. Let $A, 2, \dots, 9, 10, J, Q$ and K represent the sets of cards having the various face-values. Also, let $\heartsuit, \spadesuit, \clubsuit$ and \diamondsuit be the sets of cards having the possible suits. Find the following

(a) $A \cap \heartsuit$

(b) $A \cup \heartsuit$

(c) $J \cap (\spadesuit \cup \heartsuit)$

(d) $K \cap \heartsuit$

(e) $A \cap K$

(f) $A \cup K$

3. Do element-chasing proofs (show that an element is in the left-hand side if and only if it is in the right-hand side) to prove each of the following set equalities.

(a) $\overline{A \cap B} = \overline{A} \cup \overline{B}$

(b) $A \cup B = A \cup (\overline{A} \cap B)$

(c) $A \Delta B = (A \cup B) \setminus (A \cap B)$

$$(d) (A \cup B) \setminus C = (A \setminus C) \cup (B \setminus C)$$

4. For each positive integer n , we'll define an interval I_n by

$$I_n = [-n, 1/n).$$

Find the union and intersection of all the intervals in this infinite family.

$$\bigcup_{n \in \mathbb{N}} I_n =$$

$$\bigcap_{n \in \mathbb{N}} I_n =$$

5. There is a set X such that, for all sets A , we have $X \triangle A = A$. What is X ?
6. There is a set Y such that, for all sets A , we have $Y \triangle A = \overline{A}$. What is Y ?
7. In proving a set-theoretic identity, we are basically showing that two sets are equal. One reasonable way to proceed is to show that each is contained in the other. Prove that $A \cap (B \cup C) = (A \cap B) \cup (A \cap C)$ by showing that $A \cap (B \cup C) \subseteq (A \cap B) \cup (A \cap C)$ and $(A \cap B) \cup (A \cap C) \subseteq A \cap (B \cup C)$.
8. Prove that $A \cup (B \cap C) = (A \cup B) \cap (A \cup C)$ by showing that $A \cup (B \cap C) \subseteq (A \cup B) \cap (A \cup C)$ and $(A \cup B) \cap (A \cup C) \subseteq A \cup (B \cap C)$.
9. Prove the set-theoretic versions of DeMorgan's laws using the technique discussed in the previous problems.
10. The previous technique (showing that $A = B$ by arguing that $A \subseteq B \wedge B \subseteq A$) will have an outline something like

Proof: First we will show that $A \subseteq B$.

Towards that end, suppose $x \in A$.

\vdots

Thus $x \in B$.

Now, we will show that $B \subseteq A$.

Suppose that $x \in B$.

\vdots

Thus $x \in A$.

Therefore $A \subseteq B \wedge B \subseteq A$ so we conclude that $A = B$.

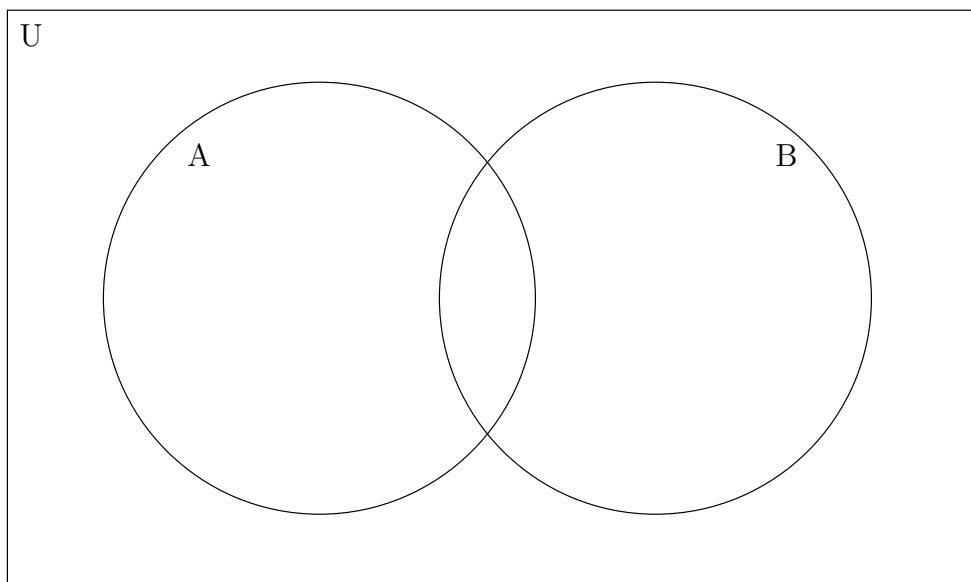
Q.E.D.

Formulate a proof that $A \Delta B = (A \cup B) \setminus (A \cap B)$.

4.4 Venn diagrams

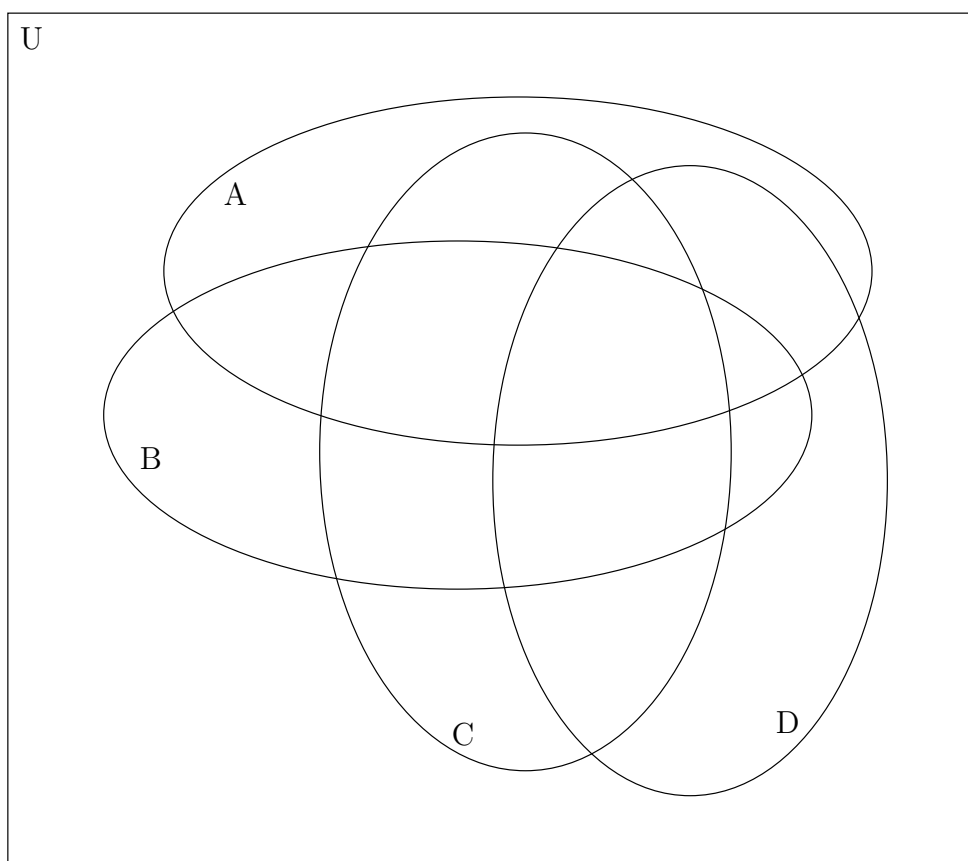
Hopefully, you've seen Venn diagrams before, but possibly you haven't thought deeply about them. Venn diagrams take advantage of an obvious but important property of closed curves drawn in the plane. They divide the points in the plane into two sets, those that are inside the curve and those that are outside! (Forget for a moment about the points that are on the curve.) This seemingly obvious statement is known as the *Jordan curve theorem*, and actually requires some details. A *Jordan curve* is the sort of curve you might draw if you are required to end where you began and you are required not to cross-over any portion of the curve that has already been drawn. In technical terms such a curve is called *continuous*, *simple* and *closed*. The Jordan curve theorem is one of those statements that hardly seems like it needs a proof, but nevertheless, the proof of this statement is probably the best-remembered work of the famous French mathematician Camille Jordan.

The prototypical Venn diagram is the picture that looks something like the view through a set of binoculars.



In a Venn diagram the universe of discourse is normally drawn as a rect-

angular region inside of which all the action occurs. Each set in a Venn diagram is depicted by drawing a simple closed curve – typically a circle, but not necessarily! For instance, if you want to draw a Venn diagram that shows all the possible intersections among four sets, you’ll find it’s impossible with (only) circles.

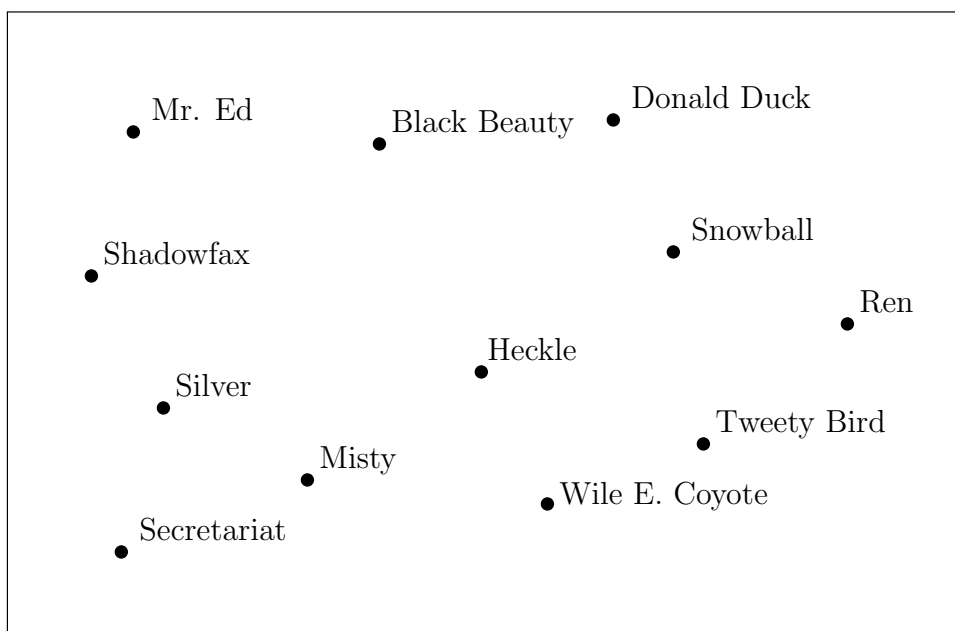


Exercise. *Verify that the diagram above has regions representing all 16 possible intersections of 4 sets.*

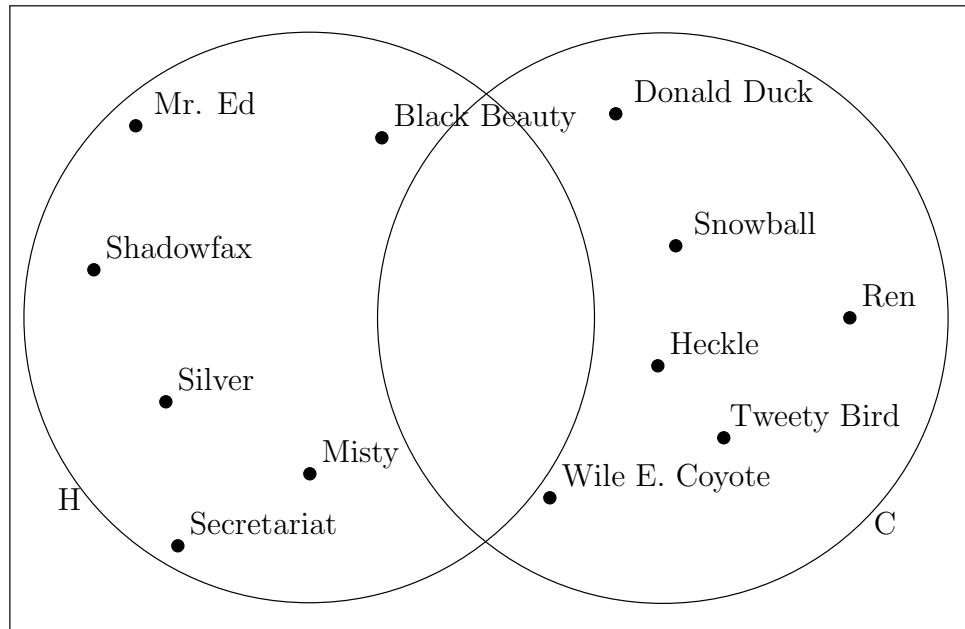
There is a certain “zen” to Venn diagrams that must be internalized, but once you have done so they can be used to think very effectively about the relationships between sets. The main deal is that the points inside of one of the simple closed curves are not necessarily in the set – only *some* of the

points inside a simple closed curve are in the set, and we don't know precisely where they are! The various simple closed curves in a Venn diagram divide the universe up into a bunch of regions. It might be best to think of these regions as fenced-in areas in which the elements of a set mill about, much like domesticated animals in their pens. One of our main tools in working with Venn diagrams is to deduce that certain of these regions don't contain any elements – we then mark that region with the emptyset symbol (\emptyset).

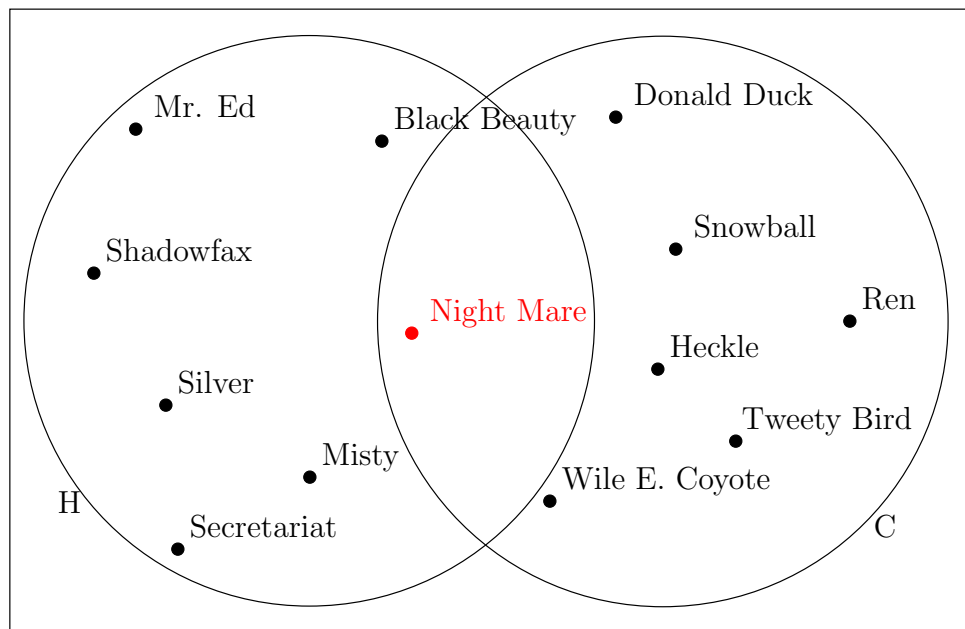
Here is a small example of a finite universe.



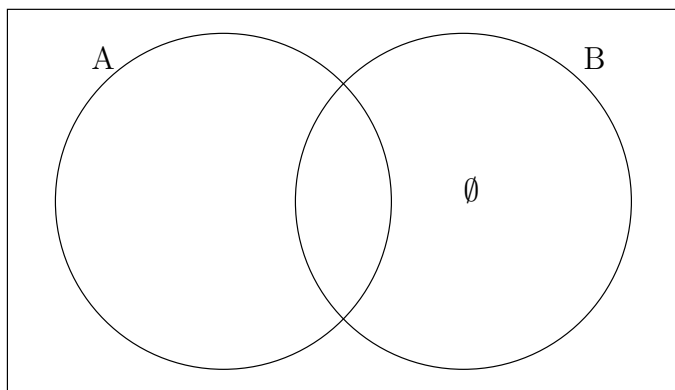
And here is the same universe with some Jordan curves used to encircle two subsets.



This picture might lead us to think that the set of cartoon characters and the set of horses are disjoint, so we thought it would be nice to add one more element to our universe in order to dispel that notion.

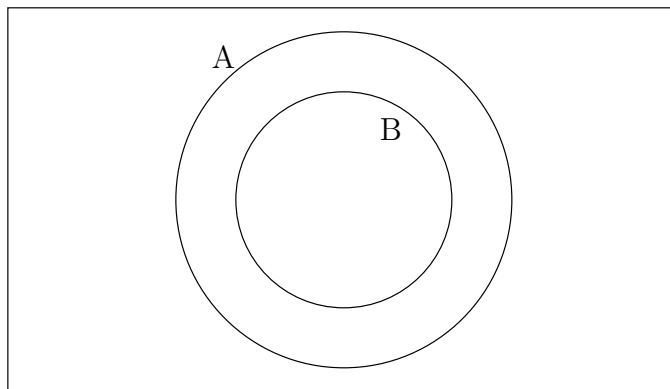
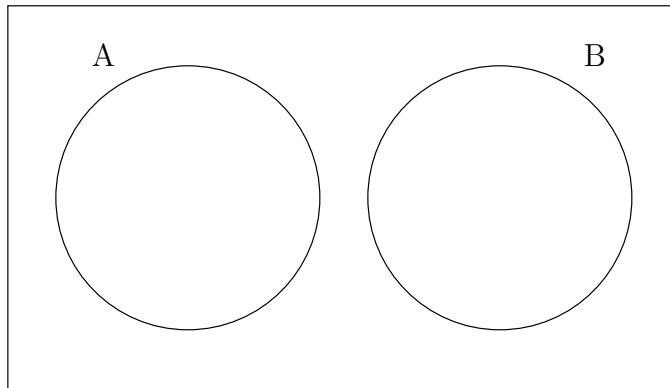


Suppose we have two sets A and B and we're interested in proving that $B \subseteq A$. The job is done if we can show that all of B 's elements are actually in the eye-shaped region that represents the intersection $A \cap B$. It's equivalent if we can show that the region marked with \emptyset in the following diagram is actually empty.



Let's put all this together. The inclusion $B \subseteq A$ corresponds to the logical sentence $M_B \implies M_A$. We know that implications are equivalent to OR statements, so $M_B \implies M_A \cong \neg M_B \vee M_A$. The notion that the region we've indicated above is empty is written as $\overline{A} \cap B = \emptyset$, in logical terms this is $\neg M_A \wedge M_B \cong c$. Finally, we apply DeMorgan's law and a commutation to get $\neg M_B \vee M_A \cong t$. You should take note of the convention that when you see a logical sentence just written on the page (as is the case with $M_B \implies M_A$ in the first sentence of this paragraph) what's being asserted is that the sentence is *universally true*. Thus, writing $M_B \implies M_A$ is the same thing as writing $M_B \implies M_A \cong t$.

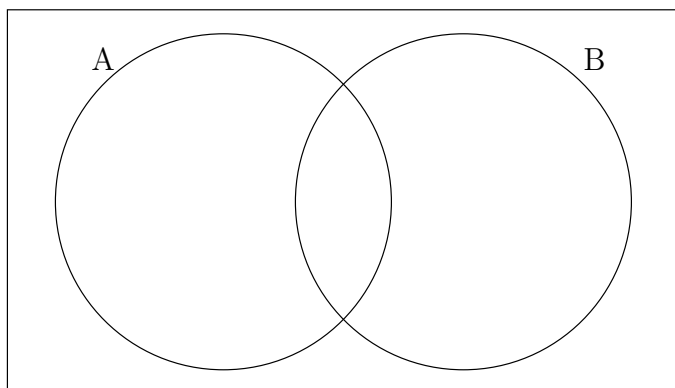
One can use information that is known *a priori* when drawing a Venn diagram. For instance if two sets are known to be disjoint, or if one is known to be contained in the other, we can draw Venn diagrams like the following.



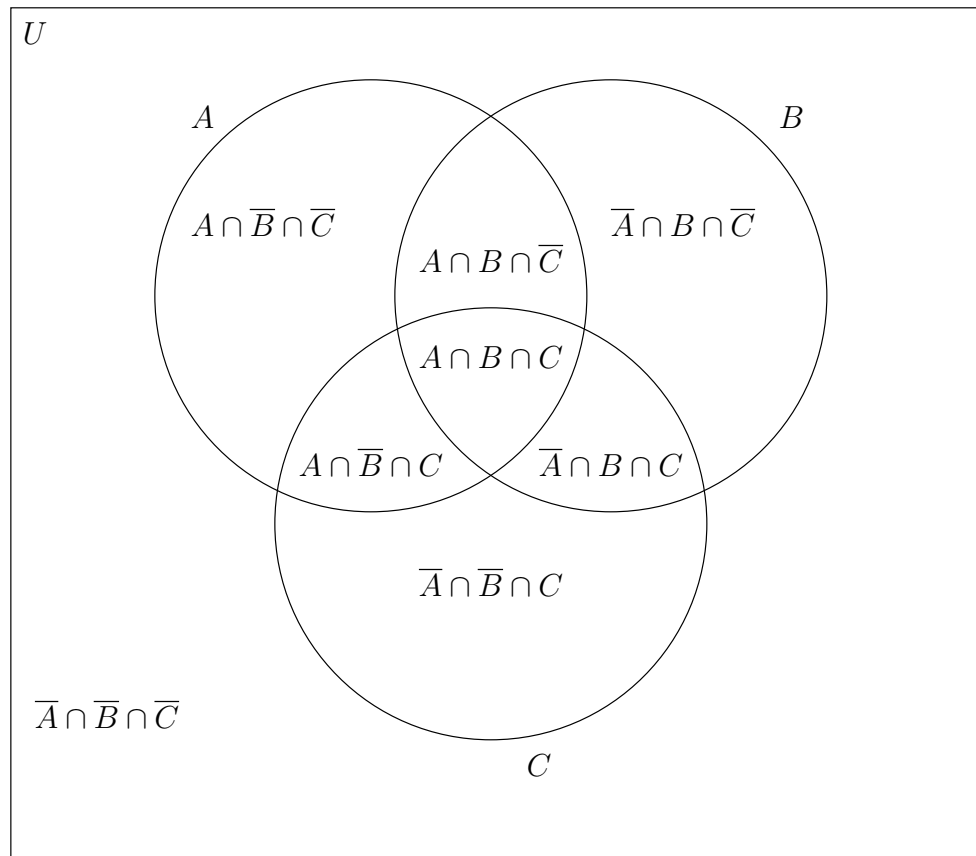
However, both of these situations can also be dealt with by working with Venn diagrams in which the sets are in *general position* – which in this situation means that every possible intersection is shown – and then marking any empty regions with \emptyset .

Exercise. *On a Venn diagram for two sets in general position, indicate the empty regions when*

- a) *The sets are disjoint.*
- b) *A is contained in B.*



There is a connection, perhaps obvious, between the regions we see in a Venn diagram with sets in general position and the recognizers we studied in the section on digital logic circuits. In fact both of these topics have to do with *disjunctive normal form*. In a Venn diagram with k sets, we are seeing the universe of discourse broken up into the union of 2^k regions each of which corresponds to an intersection of either one of the sets or its complement. An arbitrary expression involving set-theoretic symbols and these k sets is true in certain of these 2^k regions and false in the others. We have put the arbitrary expression in disjunctive normal form when we express it as a union of the intersections that describe those regions.



Exercises — 4.4

1. Venn diagrams are usually made using simple closed curves with no further restrictions. Try creating Venn diagrams for 3, 4 and 5 sets (in general position) using rectangular simple closed curves.
2. We call a curve *rectilinear* if it is made of line segments that meet at right angles. Use rectilinear simple closed curves to create a Venn diagram for 5 sets.
3. Argue as to why rectilinear curves will suffice to build any Venn diagram.
4. Find the disjunctive normal form of $A \cap (B \cup C)$.
5. Find the disjunctive normal form of $(A \Delta B) \Delta C$
6. The prototypes for the *modus ponens* and *modus tollens* argument forms are the following:

All men are mortal.	All men are mortal.
Socrates is a man.	Zeus is not mortal.
Therefore Socrates is	Therefore Zeus is not a
mortal.	man.

 Illustrate these arguments using Venn diagrams.
7. Use Venn diagrams to convince yourself of the validity of the following containment statement

$$(A \cap B) \cup (C \cap D) \subseteq (A \cup C) \cap (B \cup D).$$

Now prove it!

8. Use Venn diagrams to show that the following set equivalence is false.

$$(A \cup B) \cap (C \cup D) = (A \cup C) \cap (B \cup D)$$

4.5 Russell's Paradox

There is no Nobel prize category for mathematics.⁴ Alfred Nobel's will called for the awarding of annual prizes in physics, chemistry, physiology or medicine, literature, and peace. Later, the "Bank of Sweden Prize in Economic Sciences in Memory of Alfred Nobel" was created and certainly several mathematicians have won what is improperly known as the Nobel prize in Economics. But, there is no Nobel prize in Mathematics *per se*. There is an interesting urban myth that purports to explain this lapse: Alfred Nobel's wife either left him for, or had an affair with a mathematician — so Nobel, the inventor of dynamite and an immensely wealthy and powerful man, when he decided to endow a set of annual prizes for "those who, during the preceding year, shall have conferred the greatest benefit on mankind" pointedly left out mathematicians.

One major flaw in this theory is that Nobel was never married.

In all likelihood, Nobel simply didn't view mathematics as a field which provides benefits for mankind — at least not directly. The broadest division within mathematics is between the "pure" and "applied" branches. Just precisely where the dividing line between these spheres lies is a matter of opinion, but it can be argued that it is so far to one side that one may as well call an applied mathematician a physicist (or chemist, or biologist, or economist, or . . .). One thing is clear, Nobel believed to a certain extent in the utilitarian ethos. The value of a thing (or a person) is determined by how useful it is (or they are), which makes it interesting that one of the few mathematicians to win a Nobel prize was Bertrand Russell (the 1950 prize in Literature "in recognition of his varied and significant writings in which he champions humanitarian ideals and freedom of thought").

⁴There are prizes considered equivalent to the Nobel in stature — the Fields Medal, awarded every four years by the International Mathematical Union to up to four mathematical researchers under the age of forty, and the Abel Prize, awarded annually by the King of Norway.

Bertrand Russell was one of the twentieth century's most colorful intellectuals. He helped revolutionize the foundations of mathematics, but was perhaps better known as a philosopher. It's hard to conceive of *anyone* who would characterize Russell as an applied mathematician!

Russell was an ardent anti-war and anti-nuclear activist. He achieved a status (shared with Albert Einstein, but very few others) as an eminent scientist who was also a powerful moral authority. Russell's mathematical work was of a very abstruse foundational sort; he was concerned with the idea of reducing all mathematical thought to Logic and Set theory.

In the beginning of our investigations into Set theory we mentioned that the notion of a "set of all sets" leads to something paradoxical. Now we're ready to look more closely into that remark and hopefully gain an understanding of Russell's paradox.

By this point you should be okay with the notion of a set that contains other sets, but would it be okay for a set to contain *itself*? That is, would it make sense to have a set defined by

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

The set A has three elements, 1, 2 and itself. So we could write

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

and then

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

and then

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

et cetera.

This obviously seems like a problem. Indeed, often paradoxes seem to be caused by self-reference of this sort. Consider

The sentence in this box is false.

So a reasonable alternative is to “do” math among the sets that don’t exhibit this particular pathology.

Thus, inside the set of all sets we are singling out a particular subset that consists of sets which don’t contain themselves.

$$\mathcal{S} = \{A \mid A \text{ is a set} \wedge A \notin A\}$$

Now within the universal set we’re working in (the set of all sets) there are only two possibilities: a given set is either in \mathcal{S} or it is in its complement $\overline{\mathcal{S}}$. Russell’s paradox comes about when we try to decide which of these alternatives pertains to \mathcal{S} itself, the problem is that each alternative leads us to the other!

If we assume that $\mathcal{S} \in \mathcal{S}$, then it must be the case that \mathcal{S} satisfies the membership criterion for \mathcal{S} . Thus, $\mathcal{S} \notin \mathcal{S}$.

On the other hand, if we assume that $\mathcal{S} \notin \mathcal{S}$, then we see that \mathcal{S} does indeed satisfy the membership criterion for \mathcal{S} . Thus $\mathcal{S} \in \mathcal{S}$.

Russell himself developed a workaround for the paradox which bears his name. Together with Alfred North Whitehead he published a 3 volume work entitled *Principia Mathematica*⁵ [17]. In the Principia, Whitehead and Russell develop a system known as *type theory* which sets forth principles for avoiding problems like Russell’s paradox. Basically, a set and its elements are of different “types” and so the notion of a set being contained in itself (as an element) is disallowed.

⁵Isaac Newton also published a 3 volume work which is often cited by this same title, *Philosophiae Naturalis Principia Mathematica*.

Exercises — 4.5

1. Verify that $(A \implies \neg A) \wedge (\neg A \implies A)$ is a logical contradiction in two ways: by filling out a truth table and using the laws of logical equivalence.
2. One way out of Russell's paradox is to declare that the collection of sets that don't contain themselves as elements is not a set itself. Explain how this circumvents the paradox.

Chapter 5

Proof techniques II — Induction

Who was the guy who first looked at a cow and said, "I think I'll drink whatever comes out of these things when I squeeze 'em!"? –Bill Watterson

5.1 The principle of mathematical induction

The Principle of Mathematical Induction (PMI) may be the least intuitive proof method available to us. Indeed, at first, PMI may feel somewhat like grabbing yourself by the seat of your pants and lifting yourself into the air. Despite the indisputable fact that proofs by PMI often feel like magic, we need to convince you of the validity of this proof technique. It is one of the most important tools in your mathematical kit!

The simplest argument in favor of the validity of PMI is simply that it is axiomatic. This may seem somewhat unsatisfying, but the axioms for the natural number system, known as the Peano axioms, include one that justifies PMI. The Peano axioms will not be treated thoroughly in this book, but here they are:

- i) There is a least element of \mathbb{N} that we denote by 0.

- ii) Every natural number a has a successor denoted by $s(a)$. (Intuitively, think of $s(a) = a + 1$.)
- iii) There is no natural number whose successor is 0. (In other words, -1 isn't in \mathbb{N} .)
- iv) Distinct natural numbers have distinct successors. ($a \neq b \implies s(a) \neq s(b)$)
- v) If a subset of the natural numbers contains 0 and also has the property that whenever $a \in S$ it follows that $s(a) \in S$, then the subset S is actually equal to \mathbb{N} .

The last axiom is the one that justifies PMI. Basically, if 0 is in a subset, and the subset has this property about successors¹, then 1 must be in it. But if 1 is in it, then 1's successor (2) must be in it. And so on ...


The subset ends up having every natural number in it.

Exercise. *Verify that the following symbolic formulation has the same content as the version of the 5th Peano axiom given above.*

$$\forall S \subseteq \mathbb{N} (0 \in S) \wedge (\forall a \in \mathbb{N}, a \in S \implies s(a) \in S) \implies S = \mathbb{N}$$

On August 16th 2003, Ma Lihua of Beijing, China earned her place in the record books by single-handedly setting up an arrangement of dominoes standing on end (actually, the setup took 7 weeks and was almost ruined by some cockroaches in the Singapore Expo Hall) and toppling them. After the first domino was tipped over it took about six minutes before 303,621 out of the 303,628 dominoes had fallen. (One has to wonder what kept those other 7 dominoes upright ...)

¹Whenever a number is in it, the number's successor must be in it.



This is the model one should keep in mind when thinking about PMI: domino toppling. In setting up a line of dominoes, what do we need to do in order to ensure that they will all fall when the toppling begins? Every domino must be placed so that it will hit and topple its successor. This is exactly analogous to $(a \in S \implies s(a) \in S)$. (Think of S having the membership criterion, $x \in S =$ “ x will have fallen when the toppling is over.”) The other thing that has to happen (barring the action of cockroaches) is for someone to knock over the first domino. This is analogous to $0 \in S$.

Rather than continuing to talk about subsets of the naturals, it will be convenient to recast our discussion in terms of infinite families of logical statements. If we have a sequence of statements, (one for each natural number) $P_0, P_1, P_2, P_3, \dots$ we can prove them *all* to be true using PMI. We have to do two things. First – and this is usually the easy part – we must show that P_0 is true (i.e. the first domino *will* get knocked over). Second, we must show, for every possible value of k , $P_k \implies P_{k+1}$ (i.e. each domino will

knock down its successor). These two parts of an inductive proof are known, respectively, as the *basis* and the *inductive step*.

An outline for a proof using PMI:

Theorem $\forall n \in \mathbb{N}, P_n$

Proof: (By induction)

Basis:

\vdots (Here we must show
that P_0 is true.)

Inductive step:

\vdots (Here we must show
that $\forall k, P_k \implies P_{k+1}$
is true.)

Q.E.D.

Soon we'll do an actual example of an inductive proof, but first we have to say something *REALLY IMPORTANT* about such proofs. Pay attention! This is *REALLY IMPORTANT*! When doing the second part of an inductive proof (the inductive step), you are proving a UCS, and if you recall how that's done, you start by assuming the antecedent is true. But the particular UCS we'll be dealing with is $\forall k, P_k \implies P_{k+1}$. That means that in the course of proving $\forall n, P_n$ we have to *assume* that P_k is true. Now this sounds very much like the error known as "circular reasoning," especially as many authors don't even use different letters (n versus k in our outline) to distinguish the two statements. (And, quite honestly, we only introduced the variable k to assuage a certain lingering guilt regarding circular reasoning.) The sentence $\forall n, P_n$ is what we're trying to prove, but the sentence we need to prove in order to do that is $\forall k, P_k \implies P_{k+1}$. This is subtly different – in proving

that $\forall k, P_k \implies P_{k+1}$ (which is a UCS!) we assume that P_k is true *for some particular value of k* .

The sentence P_k is known as the *inductive hypothesis*. Think about it this way: If we were doing an entirely separate proof of $\forall n, P_n \implies P_{n+1}$, it would certainly be fair to use the inductive hypothesis, and *once that proof was done*, it would be okay to quote that result in an inductive proof of $\forall n, P_n$. Thus we can compartmentalize our way out of the difficulty!

Okay, so on to an example. In Section 4.1 we discovered a formula relating the sizes of a set A and its power set $\mathcal{P}(A)$. If $|A| = n$ then $|\mathcal{P}(A)| = 2^n$. What we've got here is an infinite family of logical sentences, one for each value of n in the natural numbers,

$$|A| = 0 \implies |\mathcal{P}(A)| = 2^0,$$

$$|A| = 1 \implies |\mathcal{P}(A)| = 2^1,$$

$$|A| = 2 \implies |\mathcal{P}(A)| = 2^2,$$

$$|A| = 3 \implies |\mathcal{P}(A)| = 2^3,$$

et cetera.

This is exactly the sort of situation in which we use induction.

Theorem 5.1.1. *For all finite sets A , $|A| = n \implies |\mathcal{P}(A)| = 2^n$.*

Proof: Let $n = |A|$ and proceed by induction on n .

Basis: Suppose A is a finite set and $|A| = 0$, it follows that $A = \emptyset$. The power set of \emptyset is $\{\emptyset\}$ which is a set having 1 element. Note that $2^0 = 1$.

Inductive step: Suppose that A is a finite set with $|A| = k + 1$. Choose some particular element of A , say a , and note that we can divide the subsets of A (i.e. elements of $\mathcal{P}(A)$) into two categories, those that contain a and those that don't.

Let $S_1 = \{X \in \mathcal{P}(A) \mid a \in X\}$ and let $S_2 = \{X \in \mathcal{P}(A) \mid a \notin X\}$.

We have created two sets that contain all the elements of $\mathcal{P}(A)$, and which are disjoint from one another. In symbolic form, $S_1 \cup S_2 = \mathcal{P}(A)$ and $S_1 \cap S_2 = \emptyset$. It follows that $|\mathcal{P}(A)| = |S_1| + |S_2|$.

Notice that S_2 is actually the power set of the k -element set $A \setminus \{a\}$. By the inductive hypothesis, $|S_2| = 2^k$. Also, notice that each set in S_1 corresponds uniquely to a set in S_2 if we just remove the element a from it. This shows that $|S_1| = |S_2|$. Putting this all together we get that $|\mathcal{P}(A)| = 2^k + 2^k = 2(2^k) = 2^{k+1}$.

Q.E.D.

Here are a few pieces of advice about proofs by induction:

- Statements that can be proved inductively don't always start out with P_0 . Sometimes P_1 is the first statement in an infinite family. Sometimes it's P_5 . Don't get hung up about something that could be handled by renumbering things.
- In your final write-up you only need to prove the initial case (whatever it may be) for the basis, but it is a good idea to try the first several cases while you are in the "draft" stage. This can provide insights into how to prove the inductive step, and it may also help you avoid a classic error in which the inductive approach fails essentially just because there is a gap between two of the earlier dominoes.²
- It is a good idea to write down somewhere just what it is that needs to be proved in the inductive step — just don't make it look like you're assuming what needs to be shown. For instance in the proof above it might have been nice to start the inductive step with a comment

²See exercise 2, the classic fallacious proof that all horses are the same color.

along the following lines, “What we need to show is that under the assumption that any set of size k has a power set of size 2^k , it follows that a set of size $k + 1$ will have a power set of size 2^{k+1} .”

We’ll close this section with a short discussion about nothing.

When we first introduced the natural numbers (\mathbb{N}) in Chapter 1 we decided to follow the convention that the smallest natural number is 1. You may have noticed that the Peano axioms mentioned in the beginning of this section treat 0 as the smallest natural number. Many people follow Dr. Peano’s convention, but we’re going to stick with our original interpretation:

$$\mathbb{N} = \{1, 2, 3, \dots\}$$

Despite our stubbornness, we are forced to admit that many inductive proofs are made easier by treating the “first” case as being in truth the one numbered 0. We’ll use the symbol $\mathbb{Z}^{\text{nonneg}}$ to indicate the set $\mathbb{N} \cup \{0\}$.

Exercises — 5.1

1. Consider the sequence of number that are 1 greater than a multiple of 4. (Such numbers are of the form $4j + 1$.)

$$1, 5, 9, 13, 17, 21, 25, 29, \dots$$

The sum of the first several numbers in this sequence can be expressed as a polynomial.

$$\sum_{j=0}^n 4j + 1 = 2n^2 + 3n + 1$$

Complete the following table in order to provide evidence that the formula above is correct.

n	$\sum_{j=0}^n 4j + 1$	$2n^2 + 3n + 1$
0	1	1
1	$1 + 5 = 6$	$2 \cdot 1^2 + 3 \cdot 1 + 1 = 6$
2	$1 + 5 + 9 =$	
3		
4		

2. What is wrong with the following inductive proof of “all horses are the same color.”?

Theorem Let H be a set of n horses, all horses in H are the same color.

Proof: We proceed by induction on n .



Basis: Suppose H is a set containing 1 horse. Clearly this horse is the same color as itself.

Inductive step: Given a set of $k + 1$ horses H we can construct two sets of k horses. Suppose $H = \{h_1, h_2, h_3, \dots, h_{k+1}\}$. Define $H_a = \{h_1, h_2, h_3, \dots, h_k\}$ (i.e. H_a contains just the first k horses) and $H_b = \{h_2, h_3, h_4, \dots, h_{k+1}\}$ (i.e. H_b contains the last k horses). By the inductive hypothesis both these sets contain horses that are “all the same color.” Also, all the horses from h_2 to h_k are in both sets so both H_a and H_b contain only horses of this (same) color. Finally, we conclude that all the horses in H are the same color.

Q.E.D.

3. For each of the following theorems, write the statement that must be proved for the basis – then prove it, if you can!

- (a) The sum of the first n positive integers is $(n^2 + n)/2$.
- (b) The sum of the first n (positive) odd numbers is n^2 .
- (c) If n coins are flipped, the probability that all of them are “heads” is $1/2^n$.
- (d) Every $2^n \times 2^n$ chessboard – with one square removed – can be tiled perfectly³ by L-shaped trominoes. (A trominoe is like a domino but made up of 3 little squares. There are two kinds, straight

 and L-shaped . This problem is only concerned with the L-shaped trominoes.)

³Here, “perfectly tiled” means that every trominoe covers 3 squares of the chessboard (nothing hangs over the edge) and that every square of the chessboard is covered by some trominoe.

4. Suppose that the rules of the game for PMI were changed so that one did the following:
- Basis. Prove that $P(0)$ is true.
 - Inductive step. Prove that for all k , P_k implies P_{k+2}

Explain why this would not constitute a valid proof that P_n holds for all natural numbers n . How could we change the basis in this outline to obtain a valid proof?

5. If we wanted to prove statements that were indexed by the integers,

$$\forall z \in \mathbb{Z}, P_z,$$

what changes should be made to PMI?

5.2 Formulas for sums and products

Gauss, when only a child, found a formula for summing the first 100 natural numbers (or so the story goes...). This formula, and his clever method for justifying it, can be easily generalized to the sum of the first n naturals. While learning calculus, notably during the study of Riemann sums, one encounters other summation formulas. For example, in approximating the integral of the function $f(x) = x^2$ from 0 to 100 one needs the sum of the first 100 *squares*. For this reason, somewhere in almost every calculus book one will find the following formulas collected:

$$\begin{aligned}\sum_{j=1}^n j &= \frac{n(n+1)}{2} \\ \sum_{j=1}^n j^2 &= \frac{n(n+1)(2n+1)}{6} \\ \sum_{j=1}^n j^3 &= \frac{n^2(n+1)^2}{4}.\end{aligned}$$

A really industrious author might also include the sum of the fourth powers. Jacob Bernoulli (a truly industrious individual) got excited enough to find formulas for the sums of the first ten powers of the naturals. Actually, Bernoulli went much further. His work on sums of powers lead to the definition of what are now known as Bernoulli numbers and let him calculate $\sum_{j=1}^{1000} j^{10}$ in about seven minutes – long before the advent of calculators! In [16, p. 320], Bernoulli is quoted:

With the help of this table it took me less than half of a quarter of an hour to find that the tenth powers of the first 1000 numbers being added together will yield the sum

91, 409, 924, 241, 424, 243, 424, 241, 924, 242, 500.

To the beginning calculus student, the beauty of the above relationships may be somewhat dimmed by the memorization challenge that they represent. It is fortunate then, that the right-hand side of the third formula is just the square of the right-hand side of the first formula. And of course, the right-hand side of the first formula is something that can be deduced by a six year old child (provided that he is a super-genius!) This happy coincidence leaves us to apply most of our rote memorization energy to formula number two, because the first and third formulas are related by the following rather bizarre-looking equation,

$$\sum_{j=1}^n j^3 = \left(\sum_{j=1}^n j \right)^2.$$

The sum of the cubes of the first n numbers is the square of their sum.

For completeness we should include the following formula which should be thought of as the sum of the zeroth powers of the first n naturals.

$$\sum_{j=1}^n 1 = n$$

Exercise. Use the above formulas to approximate the integral

$$\int_{x=0}^{10} x^3 - 2x + 3 dx$$

Our challenge today is not to merely memorize these formulas but to prove their validity. We'll use PMI.

Before we start in on a proof, it's important to figure out where we're trying to go. In proving the formula that Gauss discovered by induction we need to show that the $k + 1$ -th version of the formula holds, assuming

that the k -th version does. Before proceeding on to read the proof do the following

Exercise. Write down the $k+1$ -th version of the formula for the sum of the first n naturals. (You have to replace every n with a $k+1$.)

Theorem 5.2.1.

$$\forall n \in \mathbb{Z}^{nonneg}, \sum_{j=1}^n j = \frac{n(n+1)}{2}$$

Proof: We proceed by induction on n .

Basis: Notice that when $n = 0$ the sum on the left-hand side has no terms in it! This is known as an empty sum, and by definition, an empty sum's value is 0. Also, when $n = 0$ the formula on the right-hand side becomes $(0 \cdot 1)/2$ and this is 0 as well.⁴

Inductive step: Consider the sum on the left-hand side of the $k+1$ -th version of our formula.

$$\sum_{j=1}^{k+1} j$$

We can separate out the last term of this sum.

$$= (k+1) + \sum_{j=1}^k j$$

Next, we can use the inductive hypothesis to replace the sum (the part that goes from 1 to k) with a formula.

$$= (k+1) + \frac{k(k+1)}{2}$$

⁴If you'd prefer to avoid the "empty sum" argument, you can choose to use $n = 1$ as the basis case. The theorem should be restated so the universe of discourse is *positive* naturals.

From here on out it's just algebra . . .

$$\begin{aligned} &= \frac{2(k+1)}{2} + \frac{k(k+1)}{2} \\ &= \frac{2(k+1) + k(k+1)}{2} \\ &= \frac{(k+1) \cdot (k+2)}{2}. \end{aligned}$$

Q.E.D.

Notice how the inductive step in this proof works. We start by writing down the left-hand side of P_{k+1} , we pull out the last term so we've got the left-hand side of P_k (plus something else), then we apply the inductive hypothesis and do some algebra until we arrive at the right-hand side of P_{k+1} . Overall, we've just transformed the left-hand side of the statement we wish to prove into its right-hand side.

There is another way to organize the inductive steps in proofs like these that works by manipulating entire equalities (rather than just one side or the other of them).

Inductive step (alternate): By the inductive hypothesis, we can write

$$\sum_{j=1}^k j = \frac{k(k+1)}{2}.$$

Adding $(k+1)$ to both side of this yields

$$\sum_{j=1}^{k+1} j = (k+1) + \frac{k(k+1)}{2}.$$

Next, we can simplify the right-hand side of this to obtain

$$\sum_{j=1}^{k+1} j = \frac{(k+1)(k+2)}{2}.$$

Q.E.D.

Oftentimes one can save considerable effort in an inductive proof by creatively using the factored form during intermediate steps. On the other hand, sometimes it is easier to just simplify everything completely, and also, completely simplify the expression on the right-hand side of $P(k+1)$ and then verify that the two things are equal. This is basically just another take on the technique of “working backwards from the conclusion.” Just remember that in writing-up your proof you need to make it look as if you reasoned directly from the premises to the conclusion. We’ll illustrate what we’ve been discussing in this paragraph while proving the formula for the sum of the squares of the first n naturals.

Theorem 5.2.2.

$$\forall n \in \mathbb{Z}^{nonneg}, \sum_{j=1}^n j^2 = \frac{n(n+1)(2n+1)}{6}$$

Proof: We proceed by induction on n .

Basis: When $n = 1$ the sum has only one term, $1^2 = 1$. On the other hand, the formula is $\frac{1(1+1)(2 \cdot 1 + 1)}{6} = 1$. Since these are equal, the basis is proved.

Inductive step:

Before proceeding with the inductive step, in this box, we will figure out what the right-hand side of our theorem looks like when n is replaced with $k + 1$:

$$\begin{aligned} & \frac{(k+1)((k+1)+1)(2(k+1)+1)}{6} \\ &= \frac{(k+1)(k+2)(2k+3)}{6} \\ &= \frac{(k^2+3k+2)(2k+3)}{6} \\ &= \frac{2k^3+9k^2+13k+6}{6}. \end{aligned}$$

By the inductive hypothesis,

$$\sum_{j=1}^k j^2 = \frac{k(k+1)(2k+1)}{6}.$$

Adding $(k+1)^2$ to both sides of this equation gives

$$(k+1)^2 + \sum_{j=1}^k j^2 = \frac{k(k+1)(2k+1)}{6} + (k+1)^2.$$

Thus,

$$\sum_{j=1}^{k+1} j^2 = \frac{k(k+1)(2k+1)}{6} + \frac{6(k+1)^2}{6}.$$

Therefore,

$$\begin{aligned}
\sum_{j=1}^{k+1} j^2 &= \frac{(k^2 + k)(2k + 1)}{6} + \frac{6(k^2 + 2k + 1)}{6} \\
&= \frac{(2k^3 + 3k^2 + k) + (6k^2 + 12k + 6)}{6} \\
&= \frac{2k^3 + 9k^2 + 13k + 6}{6} \\
&= \frac{(k^2 + 3k + 2)(2k + 3)}{6} \\
&= \frac{(k + 1)(k + 2)(2k + 3)}{6} \\
&= \frac{(k + 1)((k + 1) + 1)(2(k + 1) + 1)}{6}.
\end{aligned}$$

This proves the inductive step, so the result is true.

Q.E.D.

Notice how the last four lines of the proof are the same as those in the box above containing our scratch work? (Except in the reverse order.)

We'll end this section by demonstrating one more use of this technique. This time we'll look at a formula for a product rather than a sum.

Theorem 5.2.3.

$$\forall n \geq 2 \in \mathbb{Z}, \prod_{j=2}^n \left(1 - \frac{1}{j^2}\right) = \frac{n+1}{2n}.$$

Before preceding with the proof let's look at an example (although this has nothing to do with proving anything, it's really not a bad idea – it can keep you from wasting a lot of time trying to prove something that isn't actually true!) When $n = 4$ the product is

$$\left(1 - \frac{1}{2^2}\right) \cdot \left(1 - \frac{1}{3^2}\right) \cdot \left(1 - \frac{1}{4^2}\right).$$

This simplifies to

$$\left(1 - \frac{1}{4}\right) \cdot \left(1 - \frac{1}{9}\right) \cdot \left(1 - \frac{1}{16}\right) = \left(\frac{3}{4}\right) \cdot \left(\frac{8}{9}\right) \cdot \left(\frac{15}{16}\right) = \frac{360}{576}.$$

The formula on the right-hand side is

$$\frac{4+1}{2 \cdot 4} = \frac{5}{8}.$$

Well! These two expressions are *clearly* not equal to one another... What? You say they are? Just give me a second with my calculator...

Alright then. I guess we can't dodge doing the proof...

Proof: (Using mathematical induction on n .)

Basis: When $n = 2$ the product has only one term, $1 - 1/2^2 = 3/4$. On the other hand, the formula is $\frac{2+1}{2 \cdot 2} = 3/4$. Since these are equal, the basis is proved.

Inductive step:

Let k be a particular but arbitrarily chosen integer such that

$$\prod_{j=2}^k \left(1 - \frac{1}{j^2}\right) = \frac{k+1}{2k}.$$

Multiplying⁵ both sides by the $k+1$ -th term of the product gives

$$\left(1 - \frac{1}{(k+1)^2}\right) \cdot \prod_{j=2}^k \left(1 - \frac{1}{j^2}\right) = \frac{k+1}{2k} \cdot \left(1 - \frac{1}{(k+1)^2}\right).$$

⁵Really, the only reason I'm doing this silly proof is to point out to you that when you're doing the inductive step in a proof of a formula for a **product**, you don't add to both sides anymore, you **multiply**. You see that, right? Well, consider yourself to have been pointed out to or ... oh, whatever.


Thus

$$\begin{aligned}\prod_{j=2}^{k+1} \left(1 - \frac{1}{j^2}\right) &= \frac{k+1}{2k} \cdot \left(1 - \frac{1}{(k+1)^2}\right) \\ &= \frac{k+1}{2k} - \frac{(k+1)}{2k(k+1)^2} \\ &= \frac{k+1}{2k} - \frac{(1)}{2k(k+1)} \\ &= \frac{(k+1)^2 - 1}{2k(k+1)} \\ &= \frac{k^2 + 2k}{2k(k+1)} \\ &= \frac{k(k+2)}{2k(k+1)} \\ &= \frac{k+2}{2(k+1)}.\end{aligned}$$

Q.E.D.

Exercises — 5.2

1. Write an inductive proof of the formula for the sum of the first n cubes.
2. Find a formula for the sum of the first n fourth powers.
3. The sum of the first n natural numbers is sometimes called the n -th triangular number T_n . Triangular numbers are so-named because one can represent them with triangular shaped arrangements of dots.



The first several triangular numbers are 1, 3, 6, 10, 15, et cetera.

Determine a formula for the sum of the first n triangular numbers $\left(\sum_{i=1}^n T_i\right)$ and prove it using PMI.

4. Consider the alternating sum of squares:

$$\begin{aligned}
 &1 \\
 &1 - 4 = -3 \\
 &1 - 4 + 9 = 6 \\
 &1 - 4 + 9 - 16 = -10 \\
 &\text{et cetera}
 \end{aligned}$$

Guess a general formula for $\sum_{i=1}^n (-1)^{i-1} i^2$, and prove it using PMI.

5. Prove the following formula for a product.

$$\prod_{i=2}^n \left(1 - \frac{1}{i}\right) = \frac{1}{n}$$

6. Prove $\sum_{j=0}^n (4j + 1) = 2n^2 + 3n + 1$ for all integers $n \geq 0$.
7. Prove $\sum_{i=1}^n \frac{1}{(2i-1)(2i+1)} = \frac{n}{2n+1}$ for all natural numbers n .
8. The *Fibonacci numbers* are a sequence of integers defined by the rule that a number in the sequence is the sum of the two that precede it.

$$F_{n+2} = F_n + F_{n+1}$$

The first two Fibonacci numbers (actually the zeroth and the first) are both 1.

Thus, the first several Fibonacci numbers are

$$F_0 = 1, F_1 = 1, F_2 = 2, F_3 = 3, F_4 = 5, F_5 = 8, F_6 = 13, F_7 = 21, \text{ et cetera}$$

Use mathematical induction to prove the following formula involving Fibonacci numbers.

$$\sum_{i=0}^n (F_i)^2 = F_n \cdot F_{n+1}$$

5.3 Divisibility statements and other proofs using PMI

There is a very famous result known as Fermat's Little Theorem. This would probably be abbreviated FLT except for two things. In science fiction FLT means "faster than light travel" and there is *another* theorem due to Fermat that goes by the initials FLT: Fermat's Last Theorem. Fermat's last theorem states that equations of the form $a^n + b^n = c^n$, where n is a positive natural number, only have integer solutions that are trivial (like $0^3 + 1^3 = 1^3$) when n is greater than 2. When n is 1, there are lots of integer solutions. When n is 2, there are still plenty of integer solutions – these are the so-called Pythagorean triples, for example 3,4 & 5 or 5,12 & 13. It is somewhat unfair that this statement is known as Fermat's last *theorem* since he didn't prove it (or at least we can't be sure that he proved it). Five years after his death, Fermat's son published a translated⁶ version of Diophantus's *Arithmetica* containing his father's notations. One of those notations – near the place where Diophantus was discussing the equation $x^2 + y^2 = z^2$ and its solution in whole numbers – was the statement of what is now known as Fermat's last theorem as well as the following claim:

Cuius rei demonstrationem mirabilem sane detexi hanc marginis
exiguitas non caperet.

In English:

I have discovered a truly remarkable proof of this that the margin
of this page is too small to contain.

Between 1670 and 1994 a lot of famous mathematicians worked on FLT but never found the "demonstrationem mirabilem." Finally in 1994, Andrew

⁶The translation from Greek into Latin was done by Claude Bachet.

Wiles of Princeton announced a proof of FLT, but in Wiles's own words, his is "a twentieth century proof" it can't be the proof Fermat had in mind.

These days most people believe that Fermat was mistaken. Probably he thought a proof technique that works for small values of n could be generalized. It remains a tantalizing question, can a proof of FLT using only methods available in the 17th century be accomplished?

Part of the reason that so many people spent so much effort on FLT over the centuries is that Fermat had an excellent record as regards being correct about his theorems and proofs. The result known as Fermat's little theorem is an example of a theorem and proof that Fermat got right. It is probably known as his "little" theorem because its statement is very short, but it is actually a fairly deep result.

Theorem 5.3.1 (Fermat's Little Theorem). *For every prime number p , and for all integers x , the p -th power of x and x itself are congruent mod p . Symbolically:*

$$x^p \equiv x \pmod{p}$$

A slight restatement of Fermat's little theorem is that p is always a divisor of $x^p - x$ (assuming p is a prime and x is an integer). Math professors enjoy using their knowledge of Fermat's little theorem to cook up divisibility results that can be proved using mathematical induction. For example, consider the following:

$$\forall n \in \mathbb{Z}^{\text{nonneg}}, 3 \mid (n^3 + 2n + 6).$$

This is really just the $p = 3$ case of Fermat's little theorem with a little camouflage added: $n^3 + 2n + 6 = (n^3 - n) + 3(n + 2)$. But let's have a look at proving this statement using PMI.

Theorem 5.3.2. $\forall n \in \mathbb{Z}^{\text{nonneg}}, 3 \mid (n^3 + 2n + 6)$

Proof: (By mathematical induction)

Basis: Clearly $3 \mid 6$.

Inductive step:

(We need to show that $3 \mid (k^3 + 2k + 6) \implies 3 \mid ((k + 1)^3 + 2(k + 1) + 6)$.)

Consider the quantity $(k + 1)^3 + 2(k + 1) + 6$.

$$\begin{aligned} & (k + 1)^3 + 2(k + 1) + 6 \\ &= (k^3 + 3k^2 + 3k + 1) + (2k + 2) + 6 \\ &= (k^3 + 2k + 6) + 3k^2 + 3k + 3 \\ &= (k^3 + 2k + 6) + 3(k^2 + k + 1). \end{aligned}$$

By the inductive hypothesis, 3 is a divisor of $k^3 + 2k + 6$ so there is an integer m such that $k^3 + 2k + 6 = 3m$. Thus,

$$\begin{aligned} & (k + 1)^3 + 2(k + 1) + 6 \\ &= 3m + 3(k^2 + k + 1) \\ &= 3(m + k^2 + k + 1). \end{aligned}$$

This equation shows that 3 is a divisor of $(k + 1)^3 + 2(k + 1) + 6$, which is the desired conclusion.

Q.E.D.

Exercise. Devise an inductive proof of the statement, $\forall n \in \mathbb{Z}^{\text{nonneg}}, 5 \mid x^5 + 4x - 10$.

There is one other subtle trick for devising statements to be proved by PMI that you should know about. An example should suffice to make it

clear. Notice that 7 is equivalent to 1 (mod 6), it follows that any power of 7 is also 1 (mod 6). So, if we subtract 1 from some power of 7 we will have a number that is divisible by 6.

The proof (by PMI) of a statement like this requires another subtle little trick. Somewhere along the way in the proof you'll need the identity $7 = 6 + 1$.

Theorem 5.3.3.

$$\forall n \in \mathbb{Z}^{nonneg}, 6 \mid 7^n - 1$$

Proof: (By PMI)

Basis: Note that $7^0 - 1$ is 0 and also that $6 \mid 0$.

Inductive step:

(We need to show that if $6 \mid 7^k - 1$ then $6 \mid 7^{k+1} - 1$.)

Consider the quantity $7^{k+1} - 1$.

$$\begin{aligned} 7^{k+1} - 1 &= 7 \cdot 7^k - 1 \\ &= (6 + 1) \cdot 7^k - 1 \\ &= 6 \cdot 7^k + 1 \cdot 7^k - 1 \\ &= 6(7^k) + (7^k - 1) \end{aligned}$$

By the inductive hypothesis, $6 \mid 7^k - 1$ so there is an integer m such that $7^k - 1 = 6m$. It follows that

$$7^{k+1} - 1 = 6(7^k) + 6m.$$

So, clearly, 6 is a divisor of $7^{k+1} - 1$.

Q.E.D.

Mathematical induction can often be used to prove inequalities. There are quite a few examples of families of statements where there is an inequality for every natural number. Often such statements seem to be *obviously* true and yet devising a proof can be illusive. If such is the case, try using PMI. One hint: it is fairly typical that the inductive step in a PMI proof of an inequality will involve reasoning that isn't particularly sharp. Just remember that if you have an inequality and you make the big side even bigger, the resulting statement is certainly still true!

Consider the sequences 2^n and $n!$.

n	0	1	2	3
2^n	1	2	4	8
$n!$	1	1	2	6

As the table illustrates, for small values of n , $2^n > n!$. But from $n = 4$ onward the inequality is reversed.

Theorem 5.3.4.

$$\forall n \geq 4 \in \mathbb{Z}^{nonneg}, 2^n < n!$$

Proof: (By mathematical induction)

Basis: When $n = 4$ we have $2^4 < 4!$, which is certainly true ($16 < 24$).

Inductive step: Suppose that k is a natural number with $k > 4$, and that $2^k < k!$. Multiply the left hand side of this inequality by 2 and the right hand side by $k + 1$ ⁷ to get

$$2 \cdot 2^k < (k + 1) \cdot k!$$

So

⁷It might be smoother to justify this step by first proving the lemma that $\forall a, b, c, d \in \mathbb{R}^+, a < b \wedge c < d \implies ac < bd$.

$$2^{k+1} < (k + 1)!.$$

Q.E.D.

The observant Calculus student will certainly be aware of the fact that, asymptotically, exponential functions grow faster than polynomial functions. That is, if you have a base b which is greater than 1, the function b^x is eventually larger than any polynomial $p(x)$. This may seem a bit hard to believe if $b = 1.001$ and $p(x) = 500x^{10}$. The graph of $y = 1.001^x$ is practically indistinguishable from the line $y = 1$ (at first), whereas the graph of $y = 500x^{10}$ has already reached the astronomical value of five trillion (5,000,000,000,000) when x is just 10. Nevertheless, the exponential will eventually outstrip the polynomial. We can use the methods of this section to get started on proving the fact mentioned above. Consider the two sequences n^2 and 2^n .

n	0	1	2	3	4	5	6
n^2	0	1	4	9	16	25	36
2^n	1	2	4	8	16	32	64

If we think of a “race” between the sequences n^2 and 2^n , notice that 2^n starts out with the lead. The two sequences are tied when $n = 2$. Briefly, n^2 goes into the lead but they are tied again when $n = 4$. After that it would appear that 2^n recaptures the lead for good. Of course we’re making a rather broad presumption – is it really true that n^2 never catches up with 2^n again? Well, if we’re right then the following theorem should be provable:

Theorem 5.3.5. *For all natural numbers n , if $n \geq 4$ then $n^2 \leq 2^n$.*

Proof:

Basis: When $n = 4$ we have $4^2 \leq 2^4$, which is true since both numbers are 16.

Inductive step: (In the inductive step we assume that $k^2 \leq 2^k$ and then show that $(k + 1)^2 \leq 2^{k+1}$.)

The inductive hypothesis tells us that

$$k^2 \leq 2^k.$$

If we add $2k + 1$ to the left-hand side of this inequality and 2^k to the right-hand side we will produce the desired inequality. Thus our proof will follow provided that we know that $2k + 1 \leq 2^k$. Indeed, it is sufficient to show that $2k + 1 \leq k^2$ since we already know (by the inductive hypothesis) that $k^2 \leq 2^k$.

So the result remains in doubt unless you can complete the exercise that follows. . .

Q.E.D.???

Exercise. *Prove the lemma: For all $n \in \mathbb{Z}^{nonneg}$, if $n \geq 4$ then $2n + 1 \leq n^2$.*

Exercises — 5.3

Give inductive proofs of the following

1. $\forall x \in \mathbb{Z}^{\text{nonneg}}, 3 \mid x^3 - x$
2. $\forall x \in \mathbb{Z}^{\text{nonneg}}, 3 \mid x^3 + 5x$
3. $\forall x \in \mathbb{Z}^{\text{nonneg}}, 11 \mid x^{11} + 10x$
4. $\forall n \in \mathbb{Z}^{\text{nonneg}}, 3 \mid 4^n - 1$
5. $\forall n \in \mathbb{Z}^{\text{nonneg}}, 6 \mid (3n^2 + 3n - 12)$
6. $\forall n \in \mathbb{Z}^{\text{nonneg}}, 5 \mid (n^5 - 5n^3 + 14n)$
7. $\forall n \in \mathbb{Z}^{\text{nonneg}}, 4 \mid (13^n + 4n - 1)$
8. $\forall n \in \mathbb{Z}^{\text{nonneg}}, 7 \mid 8^n + 6$
9. $\forall n \in \mathbb{Z}^{\text{nonneg}}, 6 \mid 2n^3 - 2n - 14$
10. $\forall n \geq 3 \in \mathbb{Z}^{\text{nonneg}}, 3n^2 + 3n + 1 < 2n^3$
11. $\forall n > 3 \in \mathbb{Z}^{\text{nonneg}}, n^3 < 3^n$
12. $\forall n \geq 3 \in \mathbb{Z}^{\text{nonneg}}, n^3 + 3 > n^2 + 3n + 1$
13. $\forall x \geq 4 \in \mathbb{Z}^{\text{nonneg}}, x^2 2^x \leq 4^x$

5.4 The strong form of mathematical induction

The strong form of mathematical induction (a.k.a. the principle of complete induction, PCI; also a.k.a. course-of-values induction) is so-called because the hypotheses one uses are stronger. Instead of showing that $P_k \implies P_{k+1}$ in the inductive step, we get to assume that all the statements numbered smaller than P_{k+1} are true. To make life slightly easier we'll renumber things a little. The statement that needs to be proved is

$$\forall k(P_0 \wedge P_1 \wedge \dots \wedge P_{k-1}) \implies P_k.$$

An outline of a strong inductive proof is:

Theorem $\forall n \in \mathbb{Z}^{\text{nonneg}}, P_n$

Proof: (By complete induction)

Basis:

(Technically, a PCI
proof doesn't require a
basis. We recommend
that you show that P_0
is true anyway.)

Inductive step:

(Here we must show
that $\forall k, \left(\bigwedge_{i=0}^{k-1} P_i\right) \implies$
 P_k is true.)

Q.E.D.

It's fairly common that we won't truly need all of the statements from P_0 to P_{k-1} to be true, but just one of them (and we don't know *a priori* which

one). The following is a classic result; the proof that all numbers greater than 1 have prime factors.

Theorem 5.4.1. *For all natural numbers n , $n > 1$ implies n has a prime factor.*

Proof: (By strong induction) Consider an arbitrary natural number $n > 1$. If n is prime then n clearly has a prime factor (itself), so suppose that n is not prime. By definition, a composite natural number can be factored, so $n = a \cdot b$ for some pair of natural numbers a and b which are both greater than 1. Since a and b are factors of n both greater than 1, it follows that $a < n$ (it is also true that $b < n$ but we don't need that ...). The inductive hypothesis can now be applied to deduce that a has a prime factor p . Since $p | a$ and $a | n$, by transitivity $p | n$. Thus n has a prime factor.

Q.E.D.

Exercises — 5.4

Give inductive proofs of the following

1. A “postage stamp problem” is a problem that (typically) asks us to determine what total postage values can be produced using two sorts of stamps. Suppose that you have 3¢ stamps and 7¢ stamps, show (using strong induction) that any postage value 12¢ or higher can be achieved. That is,

$$\forall n \in \mathbb{Z}^{\text{nonneg}}, n \geq 12 \implies \exists x, y \in \mathbb{Z}^{\text{nonneg}}, n = 3x + 7y.$$

2. Show that any integer postage of 12¢ or more can be made using only 4¢ and 5¢ stamps.
3. The polynomial equation $x^2 = x + 1$ has two solutions, $\alpha = \frac{1+\sqrt{5}}{2}$ and $\beta = \frac{1-\sqrt{5}}{2}$. Show that the Fibonacci number F_n is less than or equal to α^n for all $n \geq 0$.

Chapter 6

Relations and functions

If evolution really works, how come mothers only have two hands? –Milton Berle

6.1 Relations

A *relation* in mathematics is a symbol that can be placed between two numbers (or variables) to create a logical statement (or open sentence). The main point here is that the insertion of a relation symbol between two numbers creates a statement whose value is either true or false. For example, we have previously seen the divisibility symbol (\mid) and noted the common error of mistaking it for the division symbol ($/$); one of these tells us to perform an arithmetic operation, the other asks us whether *if* such an operation were performed there would be a remainder. There are many other symbols that we have seen which have this characteristic, the most important is probably $=$, but there are lots: \neq , $<$, \leq , $>$, \geq all work this way – if we place them between two numbers we get a Boolean thing, it's either true or false. If, instead of numbers, we think of placing sets on either side of a relation symbol, then $=$, \subseteq and \supseteq are valid relation symbols. If we think of placing logical expressions on either side of a relation then, honestly, *any* of the logical sym-

bols is a relation, but we normally think of \wedge and \vee as operators and give things like \equiv , \implies and \iff the status of relations.

In the examples we've looked at the things on either side of a relation are of the same type. This is usually, but not always, the case. The prevalence of relations with the same kind of things being compared has even lead to the aphorism "Don't compare apples and oranges." Think about the symbol \in for a moment. As we've seen previously, it isn't usually appropriate to put *sets* on either side of this, we might have numbers or other objects on the left and sets on the right. Let's look at a small example. Let $A = \{1, 2, 3, a, b\}$ and let $B = \{\{1, 2, a\}, \{1, 3, 5, 7, \dots\}, \{1\}\}$. The "element of" relation, \in , is a *relation from A to B*.

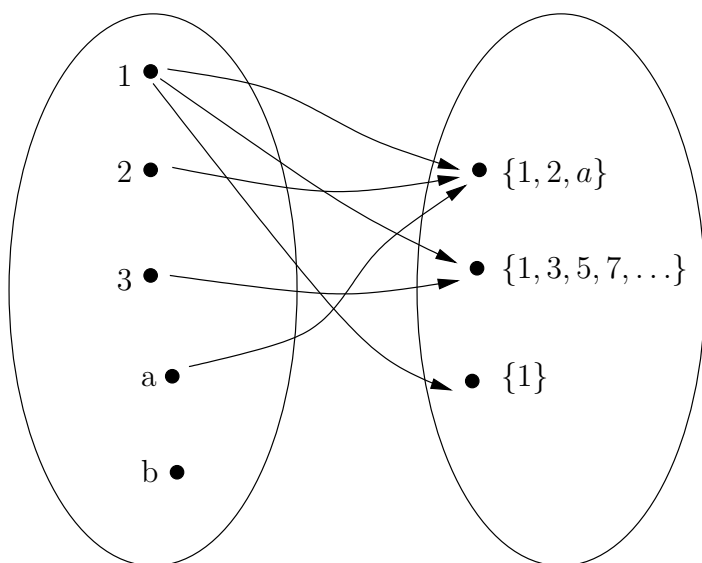


Figure 6.1: The "element of" relation is an example of a relation that goes *from* one set *to* a different set.

A diagram such as we have given in Figure 6.1 seems like a very natural thing. Such pictures certainly give us an easy visual tool for thinking about relations. But we should point out certain hidden assumptions. First, they'll only work if we are dealing with finite sets, or sets like the odd numbers in

our example (sets that are infinite but could in principle be listed). Second, by drawing the two sets separately, it seems that we are assuming they are not only different, but *disjoint*. The sets not only need not be disjoint, but often (most of the time!) we have relations that go from a set to itself so the sets in a picture like this may be identical. In Figure 6.2 we illustrate the divisibility relation on the set of all divisors of 6 — this is an example in which the sets on either side of the relation are the same. Notice the linguistic distinction, we can talk about either “a relation from A to B ” (when there are really two different sets) or “a relation on A ” (when there is only one).

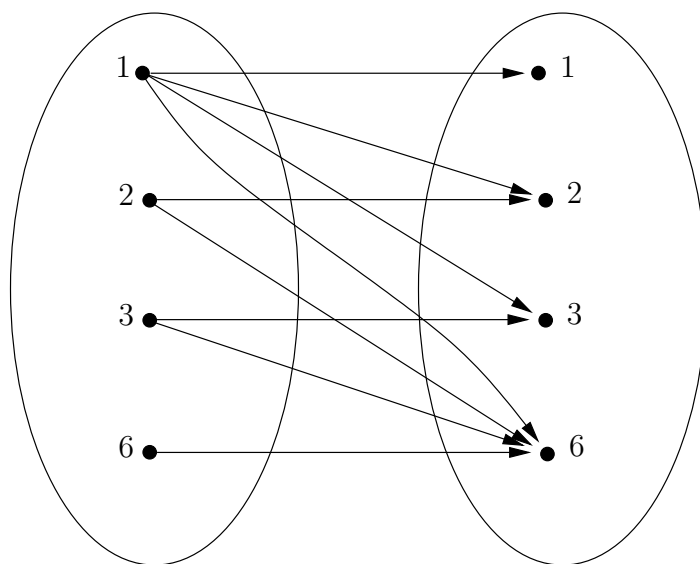
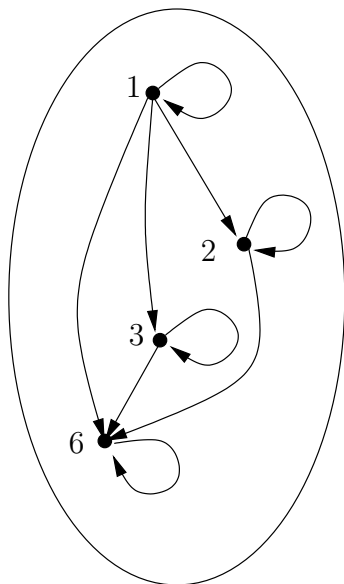


Figure 6.2: The “divides” relation is an example of a relation that goes from a set to itself. In this example we say that we have a relation *on* the set of divisors of 6.

Purists will note that it is really inappropriate to represent the same set in two different places in a Venn diagram. The diagram in Figure 6.2 should really look like this:



Indeed, this representation is definitely preferable, although it may be more crowded. A picture such as this is known as the *directed graph* (a.k.a. *digraph*) of the relation.

Recall that when we were discussing sets we said the best way to describe a set is simply to list all of its elements. Well, what is the best way to describe a relation? In the same spirit, it would seem we should explicitly list all the things that make the relation true. But it takes a *pair* of things, one to go on the left side and one to go on the right, to make a relation true (or for that matter false!). Also it should be evident that order is important in this context, for example $2 < 3$ is true but $3 < 2$ isn't. The identity of a relation is so intimately tied up with the set of ordered pairs that make it true, that when dealing with abstract relations we *define them* as sets of ordered pairs.

Given two sets, A and B , the *Cartesian product of A and B* is the set of all ordered pairs (a, b) where a is in A and b is in B . We denote the Cartesian product using the symbol \times .

$$A \times B = \{(a, b) \mid a \in A \wedge b \in B\}$$

From here on out in your mathematical career you'll need to take note of the context that the symbol \times appears in. If it appears between numbers go ahead and multiply, but if it appears between sets you're doing something different – forming the Cartesian product.

The familiar x - y plane, is often called the Cartesian plane. This is done for two reasons. Rene Descartes, the famous mathematician and philosopher, was the first to consider coordinatizing the plane and thus is responsible for our current understanding of the relationship between geometry and algebra. Rene Descartes' name is also memorialized in the definition of the Cartesian product of sets, and the plane is nothing more than the product $\mathbb{R} \times \mathbb{R}$. Indeed, the plane provided the very first example of the concept that was later generalized to the Cartesian product of sets.

Exercise. Suppose $A = \{1, 2, 3\}$ and $B = \{a, b, c\}$. Is $(a, 1)$ in the Cartesian product $A \times B$? List all elements of $A \times B$.

In the abstract, we can define a relation as *any* subset of an appropriate Cartesian product. So an abstract relation R from a set A to a set B is just some subset of $A \times B$. Similarly, a relation R on a set S is defined by a subset of $S \times S$. This definition looks a little bit strange when we apply it to an actual (concrete) relation that we already know about. Consider the relation “less than.” To describe “less than” as a subset of a Cartesian product we must write

$$< = \{(x, y) \in \mathbb{R} \times \mathbb{R} \mid y - x \in \mathbb{R}^+\}.$$

This looks funny.

Also, if we have defined some relation $R \subseteq A \times B$, then in order to say that a particular pair, (a, b) , of things make the relation true we have to write

$$aRb.$$

This looks funny too.

Despite the strange appearances, these examples do express the correct way to deal with relations.

Let's do a completely made-up example. Suppose A is the set $\{a, e, i, o, u\}$ and B is the set $\{r, s, t, l, n\}$ and we define a relation from A to B by

$$R = \{(a, s), (a, t), (a, n), (e, t), (e, l), (e, n), (i, s), (i, t), (o, r), (o, n), (u, s)\}.$$

Then, for example, because $(e, t) \in R$ we can write eRt . We indicate the negation of the concept that two elements are related by drawing a slash through the name of the relation, for example the notation \neq is certainly familiar to you, as is $\not\leq$ (although in this latter case we would normally write \geq instead). We can denote the fact that (a, l) is not a pair that makes the relation true by writing $a\not Rl$.

We should mention another way of visualizing relations. When we are dealing with a relation on \mathbb{R} , the relation is actually a subset of $\mathbb{R} \times \mathbb{R}$, that means we can view the relation as a subset of the x - y plane. In other words, we can graph it. The graph of the “ $<$ ” relation is given in Figure 6.3.

A relation on any set that is a subset of \mathbb{R} can likewise be graphed. The graph of the “ $|$ ” relation is given in Figure 6.4.

Eventually, we will get around to defining functions as relations that have a certain nice property. For the moment, we'll just note that some of the operations that you are used to using with functions also apply with relations. When one function “undoes” what another function “does” we say the functions are inverses. For example, the function $f(x) = 2x$ (i.e. doubling) and the function $g(x) = x/2$ (halving) are inverse functions because, no matter what number we start with, if we double it and then halve that result, we end up with the original number. The inverse of a relation R is written R^{-1} and it consists of the reversals of the pairs in R ,

$$R^{-1} = \{(b, a) \mid (a, b) \in R\}.$$

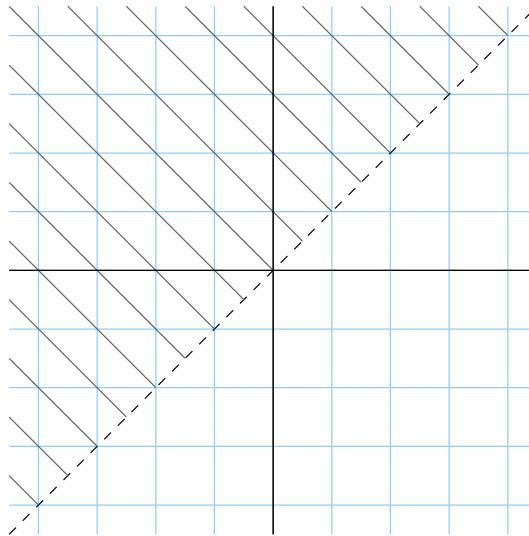


Figure 6.3: The “less than” relation can be viewed as a subset of $\mathbb{R} \times \mathbb{R}$, i.e. it can be graphed.

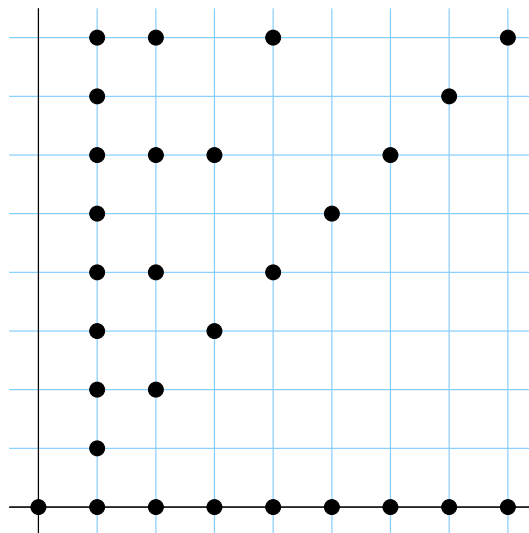


Figure 6.4: The divisibility relation can be graphed. Only those points (as indicated) with integer coordinates are in the graph.

This can also be expressed by writing

$$bR^{-1}a \iff aRb.$$

The process of “doing one function and then doing another” is known as functional composition. For instance, if $f(x) = 2x + 1$ and $g(x) = \sqrt{x}$, then we can compose them (in two different orders) to obtain either $f(g(x)) = 2\sqrt{x} + 1$ or $g(f(x)) = \sqrt{2x + 1}$. When composing functions there is an “intermediate result” that you get by applying the first function to your input, and then you calculate the second function’s value at the intermediate result. (For example, in calculating $g(f(4))$ we get the intermediate result $f(4) = 9$ and then we go on to calculate $g(9) = 3$.)

The definition of the *composite* of two relations focuses very much on this idea of the intermediate result. Suppose R is a relation from A to B and S is a relation from B to C then the composite $S \circ R$ is given by

$$S \circ R = \{(a, c) \mid \exists b \in B, (a, b) \in R \wedge (b, c) \in S\}.$$

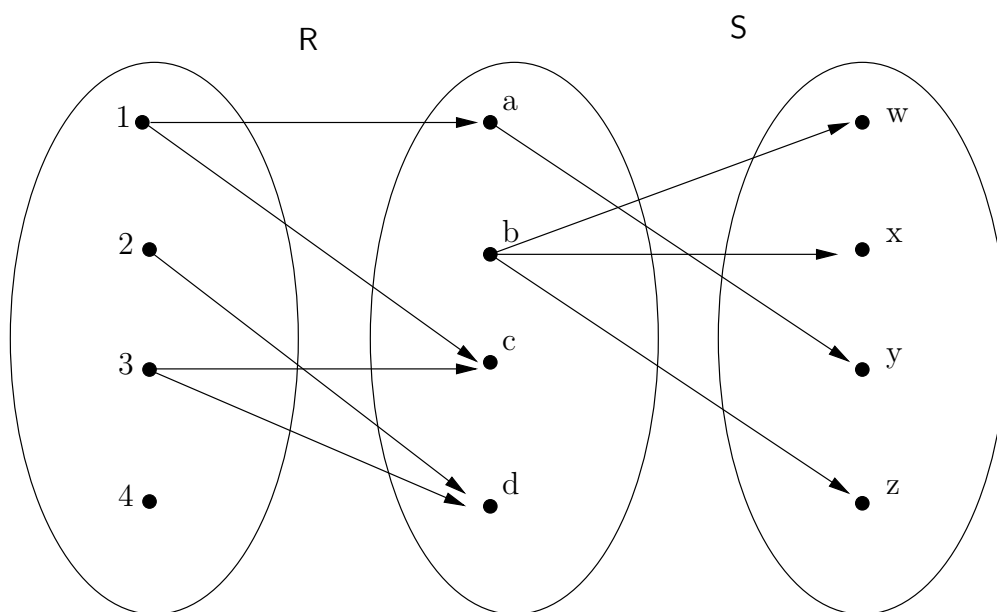
In this definition, b is the “intermediate result,” if there is no such b that serves to connect a to c then (a, c) won’t be in the composite. Also, notice that this is the composition R first, then S , but it is written as $S \circ R$ – watch out for this! The compositions of relations should be read from right to left. This convention makes sense when you consider functional composition, $f(g(x))$ means g first, then f so if we use the “little circle” notation for the composition of relations we have $f \circ g(x) = f(g(x))$ which is nice because the symbols f and g appear in the same order. But beware! there are atavists out there who write their compositions the other way around.

You should probably have a diagram like the following in mind while thinking about the composition of relations. Here, we have the set $A = \{1, 2, 3, 4\}$, the set B is $\{a, b, c, d\}$ and $C = \{w, x, y, z\}$. The relation R goes from A to B and consists of the following set of pairs,

$$R = \{(1, a), (1, c), (2, d), (3, c), (3, d)\}.$$

And

$$S = \{(a, y), (b, w), (b, x), (b, z)\}.$$



Exercise. Notice that the composition $R \circ S$ is impossible (or, more properly, it is empty). Why?

What is the (only) pair in the composition $S \circ R$?

Exercises — 6.1

1. The *lexicographic order*, $<_{\text{lex}}$, is a relation on the set of all words, where $x <_{\text{lex}} y$ means that x would come before y in the dictionary. Consider just the three letter words like “iff”, “fig”, “the”, et cetera. Come up with a usable definition for $x_1x_2x_3 <_{\text{lex}} y_1y_2y_3$.
2. What is the graph of “=” in $\mathbb{R} \times \mathbb{R}$?
3. The *inverse* of a relation R is denoted R^{-1} . It contains exactly the same ordered pairs as R but with the order switched. (So technically, they aren’t *exactly* the same ordered pairs . . .)

$$R^{-1} = \{(b, a) \mid (a, b) \in R\}$$

Define a relation S on $\mathbb{R} \times \mathbb{R}$ by $S = \{(x, y) \mid y = \sin x\}$. What is S^{-1} ? Draw a single graph containing S and S^{-1} .

4. The “socks and shoes” rule is a very silly little mnemonic for remembering how to invert a composition. If we think of undoing the process of putting on our socks and shoes (that’s socks first, then shoes) we have to first remove our shoes, *then* take off our socks.

The socks and shoes rule is valid for relations as well.

Prove that $(S \circ R)^{-1} = R^{-1} \circ S^{-1}$.

6.2 Properties of relations

There are two special classes of relations that we will study in the next two sections, equivalence relations and ordering relations. The prototype for an equivalence relation is the ordinary notion of numerical equality, $=$. The prototypical ordering relation is \leq . Each of these has certain salient properties that are the root causes of their importance. In this section we will study a compendium of properties that a relation may or may not have.

A relation that has three of the properties we'll discuss:

1. reflexivity
2. symmetry
3. transitivity

is said to be an equivalence relation; it will in some ways resemble $=$.

A relation that has another set of three properties:

1. reflexivity
2. anti-symmetry
3. transitivity

is called an ordering relation; it will resemble \leq .

Additionally, there is a property known as irreflexivity that many relations have.

There are a total of 5 properties that we have named, and we will discuss them all more thoroughly. But first, we'll state the formal definitions. Take note that these properties are all stated for a relation that goes from a set to itself, indeed, most of them wouldn't even make sense if we tried to define them for a relation from a set to a different set.

<p>A relation R on a set S is reflexive iff</p> $\forall a \in S, \quad aRa$ <p>“Everything is related to itself.”</p>
<p>A relation R on a set S is irreflexive iff</p> $\forall a \in S, \quad a \not R a$ <p>“Nothing is related to itself.”</p>
<p>A relation R on a set S is symmetric iff</p> $\forall a, b \in S, \quad aRb \implies bRa$ <p>“No one-way streets.”</p>
<p>A relation R on a set S is anti-symmetric iff</p> $\forall a, b \in S, \quad aRb \wedge bRa \implies a = b$ <p>“Only one-way streets.”</p>
<p>A relation R on a set S is transitive iff</p> $\forall a, b, c \in S, \quad aRb \wedge bRc \implies aRc$ <p>“Whenever there’s a roundabout route, there’s a direct route.”</p>

Table 6.1: Properties that relations may (or may not) have.

The digraph of a relation that is reflexive will have little loops at every vertex. The digraph of a relation that is irreflexive will contain no loops at all. Hopefully it is clear that these concepts represent extreme opposite possibilities – they are *not* however negations of one another.

Exercise. Find the logical denial of the property that says a relation is reflexive

$$\neg(\forall a \in S, \quad aRa).$$

How does this differ from the defining property for “irreflexive”?

If a relation R is defined on some subset S of the reals, then it can be graphed in the Euclidean plane. Reflexivity for R can be interpreted in terms

of the line L defined by the equation $y = x$. Every point in $(S \times S) \cap L$ must be in R . A similar statement can be made concerning the irreflexive property. If a relation R is irreflexive its graph completely avoids the line $y = x$.

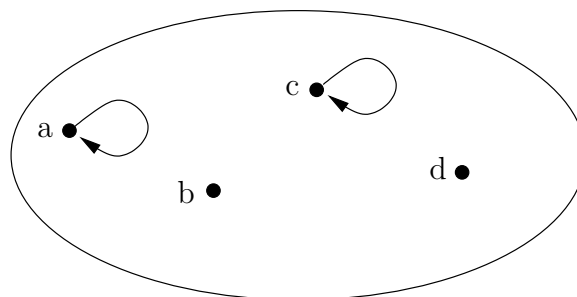
Note that the reflexive and irreflexive properties are defined with a single quantified variable. Symmetry and anti-symmetry require two universally quantified variables for their definitions.

A relation R on a set S is **symmetric** iff

$$\forall a, b \in S, \quad aRb \implies bRa.$$

This can be interpreted in terms of digraphs as follows: If a connection from a to b exists in the digraph of R , then there must also be a connection from b to a . In Table 6.1 this is interpreted as “no one-way streets” and while that’s not quite what it says, that *is* the effect of this definition. Since *if* a connection exists in one direction, there must also be a connection in the other direction, it follows that we will never see a one-way connection.

Because most of the properties we are studying are defined using conditional statements it is often the case that a relation has a property for vacuous reasons. When the “if” part doesn’t happen, there’s no need for its corresponding “then” part to happen either – the conditional is still true. In the context of our discussion on the symmetry property of a relation this means that the following digraph *is* the digraph of a symmetric relation (although it is neither reflexive nor irreflexive).



Anti-symmetry is described as meaning “only one-way streets” but the definition is given as:

A relation R on a set S is **anti-symmetric** iff

$$\forall a, b \in S, \quad aRb \wedge bRa \implies a = b.$$

It may be hard at first to understand why the definition we use for anti-symmetry is the one above. If one wanted to insure that there were never two-way connections between elements of the set it might seem easier to define anti-symmetry as follows:

(Alternate definition) A relation R on a set S is **anti-symmetric** iff

$$\forall a, b \in S, \quad aRb \implies b \not R a.$$

This definition may seem more straight-forward, but it turns out the original definition is easier to use in proofs. We need to convince ourselves that the (first) definition really accomplishes what we want. Namely, if a relation R satisfies the property that $\forall a, b \in S, \quad aRb \wedge bRa \implies a = b$, then there will not actually be any pair of elements that are related in both orders. One way to think about it is this: suppose that a and b are distinct elements of S and that both aRb and bRa are true. The property now guarantees that $a = b$ which contradicts the notion that a and b are distinct. This is a miniature proof by contradiction; if you assume there *are* a pair of distinct elements that are related in both orders you get a contradiction, so there *aren't*!

A funny thing about the anti-symmetry property is this: When it is true of a relation it is *always* vacuously true! The property is engineered in such a way that when it is true, it forces that the statement in its antecedent never really happens.

Transitivity is an extremely useful property as witnessed by the fact that both equivalence relations and ordering relations must have this property.

When speaking of the transitive property of equality we say “Two things that are equal to a third, are equal to each other.” When dealing with ordering we may encounter statements like the following. “Since ‘Aardvark’ precedes ‘Bulwark’ in the dictionary, and since ‘Bulwark’ precedes ‘Catastrophe’, it is plainly true that ‘Aardvark’ comes before ‘Catastrophe’ in the dictionary.”

Again, the definition of transitivity involves a conditional. Also, transitivity may be viewed as the most complicated of the properties we’ve been studying; it takes three universally quantified variables to state the property.

A relation R on a set S is **transitive** iff

$$\forall a, b, c \in S, \quad aRb \wedge bRc \quad \implies \quad aRc$$

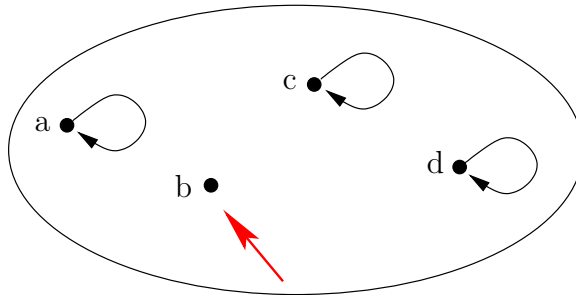
We paraphrased transitivity as “Whenever there’s a roundabout route, there’s a direct route.” In particular, what the definition says is that *if* there’s a connection from a to b and from b to c (the roundabout route from a to c) then there must be a connection from a to c (the direct route).

You’ll really need to watch out for relations that are transitive for vacuous reasons. So long as one never has three elements a , b and c with aRb and bRc the statement that defines transitivity is automatically true.

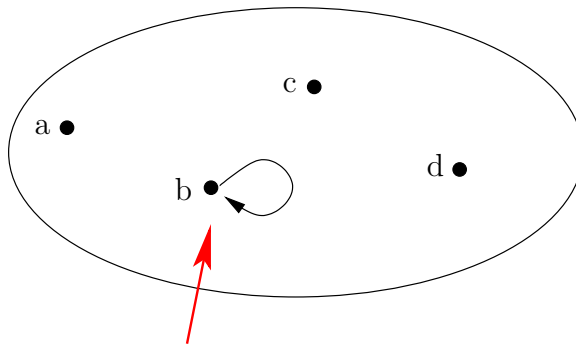
A very useful way of thinking about these various properties that relations may have is in terms of what *doesn’t* happen when a relation has them. Before we proceed, it is important that you do the following

Exercise. *Find logical negations for the formal properties defining each of the five properties.*

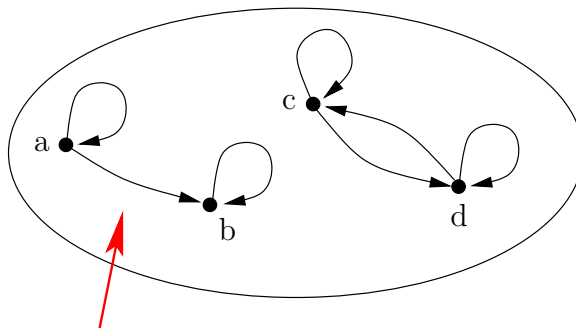
If a relation R is reflexive we will never see a node that doesn't have a loop.



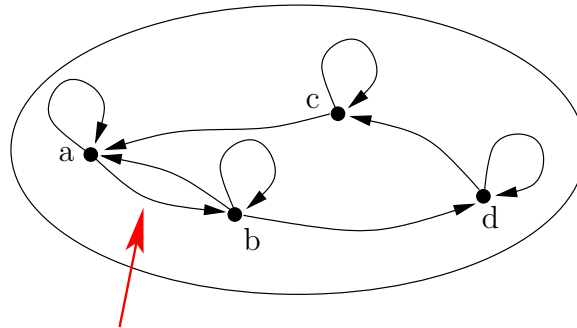
If a relation R is irreflexive we will never see a node that *does* have a loop!



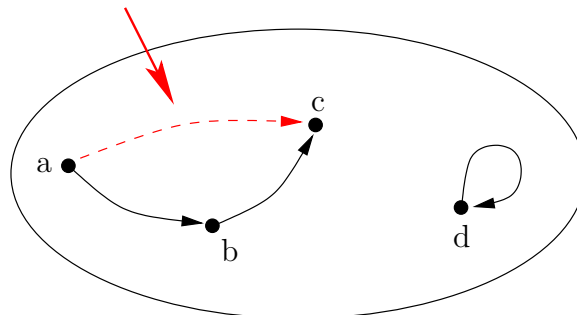
If a relation R is symmetric we will never see a pair of nodes that are connected in one direction only.



If a relation R is anti-symmetric we will never see a pair of nodes that are connected in both directions.



If a relation R is transitive the thing we will never see is a bit harder to describe. There will never be a pair of arrows meeting head to tail *without* there also being an arrow going from the tail of the first to the head of the second.



Exercises — 6.2

1. Consider the relation S defined by $S = \{(x, y) \mid x \text{ is smarter than } y\}$.
Is S symmetric or anti-symmetric? Explain.
2. Consider the relation A defined by $A = \{(x, y) \mid x \text{ has the same astrological sign as } y\}$.
Is A symmetric or anti-symmetric? Explain.
3. Explain why both of the relations just described (in problems 1 and 2) have the transitive property.
4. For each of the five properties, name a relation that has it and a relation that doesn't.

6.3 Equivalence relations

The main idea of an equivalence relation is that it is something like equality, but not quite. Usually there is some property that we can name, so that equivalent things share that property. For example Albert Einstein and Adolf Eichmann were two entirely different human beings, if you consider all the different criteria that one can use to distinguish human beings there is little they have in common. But, if the only thing one was interested in was a person's initials, one would have to say that Einstein and Eichmann were equivalent. Future examples of equivalence relations will be less frivolous. . . But first, the formal definition:

Definition. *A relation R on a set S is an equivalence relation iff R is reflexive, symmetric and transitive.*

Probably the most important equivalence relation we've seen to date is "equivalence mod m " which we will denote using the symbol \equiv_m . This relation may even be more interesting than actual equality! The reason for this seemingly odd statement is that "equivalence mod m " gives us non-trivial equivalence classes. Equivalence classes are one of the most potent ideas in modern mathematics and it's essential that you understand them, so we'll start with an example. Consider equivalence mod 5. What other numbers is (say) 11 equivalent to? There are many! Any number that leaves the same remainder as 11 when we divide it by 5. This collection is called the equivalence class of 11 and is usually denoted using an overline — $\overline{11}$, another notation that is often seen for the set of things equivalent to 11 is $11/\equiv_5$.

$$\overline{11} = \{\dots, -9, -4, 1, 6, 11, 16, \dots\}$$

It's easy to see that we will get the exact same set if we choose any other element of the equivalence class (in place of 11), which leads us to an infinite list of set equalities,

$$\bar{1} = \bar{6} = \bar{11} = \dots$$

And similarly,

$$\bar{2} = \bar{7} = \bar{12} = \dots$$

In fact, there are really just 5 different sets that form the equivalence classes mod 5: $\bar{0}$, $\bar{1}$, $\bar{2}$, $\bar{3}$, and $\bar{4}$. (Note: we have followed the usual convention of using the smallest possible non-negative integers as the representatives for our equivalence classes.)

What we've been discussing here is one of the first examples of a *quotient structure*. We start with the integers and “mod out” by an equivalence relation. In doing so, we “move to the quotient” which means (in this instance) that we go from \mathbb{Z} to a much simpler set having only five elements: $\{\bar{0}, \bar{1}, \bar{2}, \bar{3}, \bar{4}\}$. In moving to the quotient we will generally lose a lot of information, but greatly highlight some particular feature – in this example, properties related to divisibility by 5.

Given some equivalence relation R defined on a set S the set of equivalence classes of S under R is denoted S/R (which is read “ S mod R ”). This use of the slash – normally reserved for division – shouldn't cause any confusion since those aren't numbers on either side of the slash but rather a set and a relation. This notation may also clarify why some people denote the equivalence classes above by $0/ \equiv_5$, $1/ \equiv_5$, $2/ \equiv_5$, $3/ \equiv_5$ and $4/ \equiv_5$.

The set of equivalence classes forms a *partition* of the set S .

Definition. A partition P of a set S is a set of sets such that

$$S = \bigcup_{X \in P} X \quad \text{and} \quad \forall X, Y \in P, X \neq Y \implies X \cap Y = \emptyset.$$

In words, if you take the union of all the pieces of the partition you'll get the set S , and any pair of sets from the partition that aren't identical

are disjoint. Partitions are an inherently useful way of looking at things, although in the real world there are often problems (sets we thought were disjoint turn out to have elements in common, or we discover something that doesn't fit into any of the pieces of our partition), in mathematics we usually find that partitions do just what we would want them to do. Partitions divide some set up into a number of convenient pieces in such a way that we're guaranteed that every element of the set is in one of the pieces and also so that none of the pieces overlap. Partitions are a useful way of dissecting sets, and equivalence relations (via their equivalence classes) give us an easy way of creating partitions – usually with some additional structure to boot! The properties that make a relation an equivalence relation (reflexivity, symmetry and transitivity) are designed to ensure that equivalence classes exist and do provide us with the desired partition. For the beginning proof writer this all may seem very complicated, but take heart! Most of the work has already been done for you by those who created the general theory of equivalence relations and quotient structures. All you have to do (usually) is prove that a given relation is an equivalence relation by verifying that it is indeed reflexive, symmetric and transitive. Let's have a look at another example.

In Number Theory, the square-free part of an integer is what remains after we divide-out the largest perfect square that divides it. (This is also known as the *radical* of an integer.) The following table gives the square-free part, $sf(n)$, for the first several values of n .

n	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
$sf(n)$	1	2	3	1	5	6	7	2	1	10	11	3	13	14	15	1	17	2	19	5

It's easy to compute the square-free part of an integer if you know its prime factorization – just reduce all the exponents mod 2. For example¹

¹This is the size of largest sporadic finite simple group, known as “the Monster.”

$$\begin{aligned}
& 80801742479451287588645990496171075700575436800000000 \\
& = 2^{46} \cdot 3^{20} \cdot 5^9 \cdot 7^6 \cdot 11^2 \cdot 13^3 \cdot 17 \cdot 19 \cdot 23 \cdot 29 \cdot 31 \cdot 41 \cdot 47 \cdot 59 \cdot 71
\end{aligned}$$

the square-free part of this number is

$$\begin{aligned}
& 5 \cdot 13 \cdot 17 \cdot 19 \cdot 23 \cdot 29 \cdot 31 \cdot 41 \cdot 47 \cdot 59 \cdot 71 \\
& = 3504253225343845
\end{aligned}$$

which, while it is still quite a large number, is certainly a good bit smaller than the original!

We will define an equivalence relation S on the set of natural numbers by using the square-free part:

$$\forall x, y \in \mathbb{Z}^{\text{nonneg}}, xSy \iff sf(x) = sf(y)$$

In other words, two natural numbers will be S -related if they have the same square-free parts.

Exercise. *What is $1/S$?*

Before we proceed to the proof that S is an equivalence relation we'd like you to be cognizant of a bigger picture as you read. Each of the three parts of the proof will have a similar structure. We will show that S has one of the three properties by using the fact that $=$ has that property. In more advanced work this entire proof could be omitted or replaced by the phrase “ S inherits reflexivity, symmetry and transitivity from equality, and is therefore an equivalence relation.” (Nice trick isn't it? But before you're allowed to use it you have to show that you can do it the hard way ...)

Theorem 6.3.1. *The relation S defined by*

$$\forall x, y \in \mathbb{Z}^{\text{nonneg}}, xSy \iff sf(x) = sf(y)$$

is an equivalence relation on $\mathbb{Z}^{\text{nonneg}}$.

Proof: We must show that \mathbf{S} is reflexive, symmetric and transitive.

reflexive — (Here we must show that $\forall x \in \mathbb{Z}^{\text{nonneg}}, x\mathbf{S}x$.) Let x be an arbitrary natural number. Since $sf(x) = sf(x)$ (this is the reflexive property of $=$) it follows from the definition of \mathbf{S} that $x\mathbf{S}x$.

symmetric — (Here we must show that $\forall x, y \in \mathbb{Z}^{\text{nonneg}}, x\mathbf{S}y \implies y\mathbf{S}x$.) Let x and y be arbitrary natural numbers, and further suppose that $x\mathbf{S}y$. Since $x\mathbf{S}y$, it follows from the definition of \mathbf{S} that $sf(x) = sf(y)$, obviously then $sf(y) = sf(x)$ (this is the symmetric property of $=$) and so $y\mathbf{S}x$.

transitive — (Here we must show that $\forall x, y, z \in \mathbb{Z}^{\text{nonneg}}, x\mathbf{S}y \wedge y\mathbf{S}z \implies x\mathbf{S}z$.) Let x, y and z be arbitrary natural numbers, and further suppose that both $x\mathbf{S}y$ and $y\mathbf{S}z$. From the definition of \mathbf{S} we deduce that $sf(x) = sf(y)$ and $sf(y) = sf(z)$. Clearly, $sf(x) = sf(z)$ (this deduction comes from the transitive property of $=$), so $x\mathbf{S}z$.

Q.E.D.

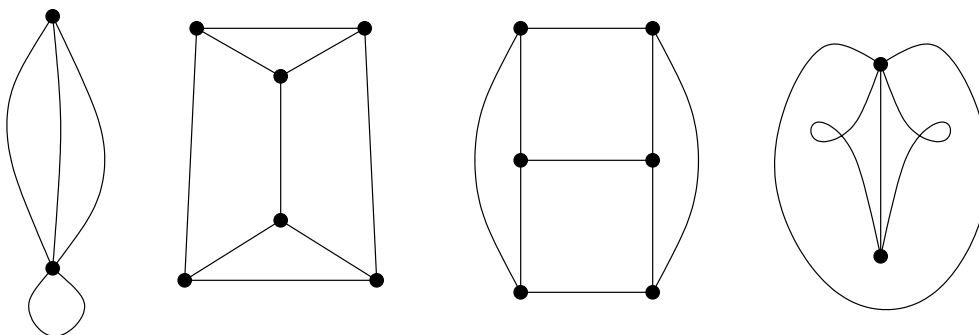
We'll end this section with an example of an equivalence relation that doesn't "inherit" the three properties from equality.

A *graph* is a mathematical structure consisting of two sets, a set V of points (a.k.a. vertices) and a set² E of edges. The elements of E may be either ordered or unordered pairs from V . If E consists of ordered pairs we have a *directed graph* or *digraph* – the diagrams we have been using to visualize relations! If E consists of unordered pairs then we are dealing with

²Technically, E is a so-called multiset in many instances – there may be several edges that connect the same pair of vertices.

an *undirected graph*. Since the undirected case is actually the more usual, if the word “graph” appears without a modifier it is assumed that we are talking about an undirected graph.

The previous paragraph gives a relatively precise definition of a graph in terms of sets, however the real way to think of graphs is in terms of diagrams where a set of dots are connected by paths. (The paths will, of course, need to have arrows on them in digraphs.) Below are a few examples of the diagrams that are used to represent graphs.



Two graphs are said to be *isomorphic* if they represent the same connections. There must first of all be a one-to-one correspondence between the vertices of the two graphs, and further, a pair of vertices in one graph are connected by some number of edges if and only if the corresponding vertices in the other graph are connected by the same number of edges.

Exercise. *The four examples of graphs above actually are two pairs of isomorphic graphs. Which pairs are isomorphic?*

This word “isomorphic” has a nice etymology. It means “same shape.” Two graphs are isomorphic if they have the same shape. We don’t have the tools right now to do a formal proof (in fact we need to look at some further prerequisites before we can really precisely define isomorphism), but isomorphism of graphs is an equivalence relation. Let’s at least verify this informally.

Reflexivity Is a graph isomorphic to itself? That is, does a graph have the “same shape” as itself? Clearly!

Symmetry If graph A is isomorphic to graph B , is it also the case that graph B is isomorphic to graph A ? I.e. if A has the “same shape” as B , doesn't B have the same shape as A ? Of course!

Transitivity Well . . .the answer here is going to be “Naturally!” but let's wait to delve into this issue when we have a usable formal definition for graph isomorphism. The question at this stage should be clear though: If A is isomorphic to B and B is isomorphic to C , then isn't A isomorphic to C ?

Exercises — 6.3

1. Consider the relation A defined by

$$A = \{(x, y) \mid x \text{ has the same astrological sign as } y\}.$$

Show that A is an equivalence relation. What equivalence class under A do you belong to?

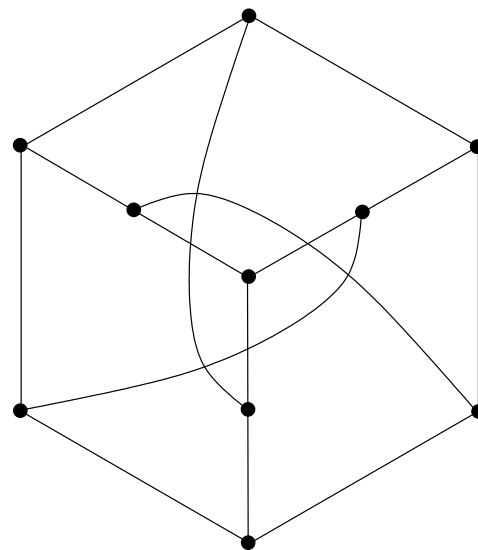
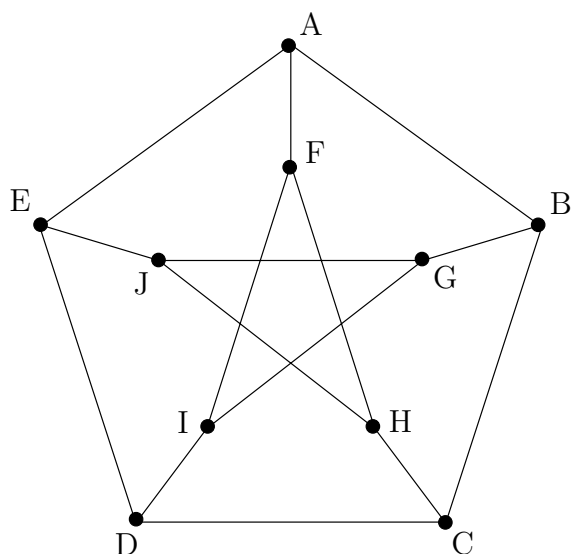
2. Define a relation \square on the integers by $x \square y \iff x^2 = y^2$. Show that \square is an equivalence relation. List the equivalence classes x/\square for $0 \leq x \leq 5$.

3. Define a relation A on the set of all words by

$$w_1 A w_2 \iff w_1 \text{ is an anagram of } w_2.$$

Show that A is an equivalence relation. (Words are anagrams if the letters of one can be re-arranged to form the other. For example, 'ART' and 'RAT' are anagrams.)

4. The two diagrams below both show a famous graph known as the Petersen graph. The picture on the left is the usual representation which emphasizes its five-fold symmetry. The picture on the right highlights the fact that the Petersen graph also has a three-fold symmetry. Label the right-hand diagram using the same letters (A through J) in order to show that these two representations are truly isomorphic.



5. We will use the symbol \mathbb{Z}^* to refer to the set of all integers *except* 0. Define a relation Q on the set of all pairs in $\mathbb{Z} \times \mathbb{Z}^*$ (pairs of integers where the second coordinate is non-zero) by $(a, b)Q(c, d) \iff ad = bc$. Show that Q is an equivalence relation.
6. The relation Q defined in the previous problem partitions the set of all pairs of integers into an interesting set of equivalence classes. Explain why

$$\mathbb{Q} = (\mathbb{Z} \times \mathbb{Z}^*)/Q.$$

Ultimately, this is the “right” definition of the set of rational numbers!

7. Reflect back on the proof in problem 5. Note that we were fairly careful in assuring that the second coordinate in the ordered pairs is non-zero. (This was the whole reason for introducing the \mathbb{Z}^* notation.) At what point in the argument did you use this hypothesis?

6.4 Ordering relations

The prototype for ordering relations is \leq . Although a case could be made for using $<$ as the prototypical ordering relation. These two relations differ in one important sense: \leq is reflexive and $<$ is irreflexive. Various authors, having made different choices as to which of these is the more prototypical, have defined ordering relations in slightly different ways. The majority view seems to be that an ordering relation is reflexive (which means that ordering relations are modeled after \leq). We would really like to take the contrary position – we always root for the underdog – but one of our favorite ordering relation (divisibility) is reflexive and it would be eliminated if we made the other choice³. So...

Definition. *A relation R on a set S is an ordering relation iff R is reflexive, anti-symmetric and transitive.*

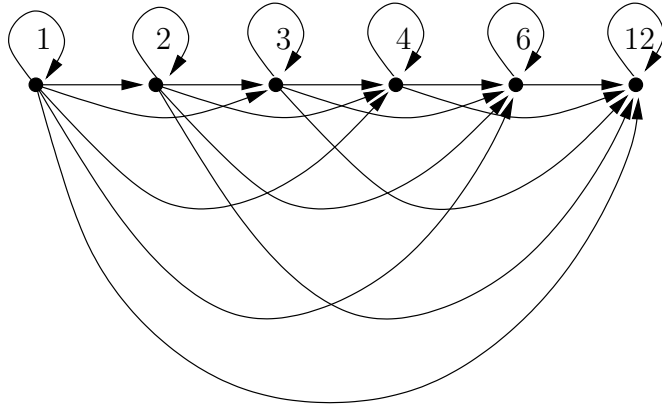
Now, we've used \leq to decide what properties an ordering relation should have, but we should point out that most ordering relations don't do nearly as good a job as \leq does. The \leq relation imposes what is known as a *total order* on the sets that it acts on (you should note that it can't be used to compare complex numbers, but it can be placed between reals or any of the sets of numbers that are contained in \mathbb{R} .) Most ordering relations only create what is known as a *partial order* on the sets they act on. In a total ordering (a.k.a. a linear ordering) every pair of elements can be compared and we can use the ordering relation to decide which order they go in. In a partial ordering there may be elements that are incomparable.

Definition. *If x and y are elements of a set S and R is an ordering relation on S then we say x and y are comparable if $xRy \vee yRx$.*

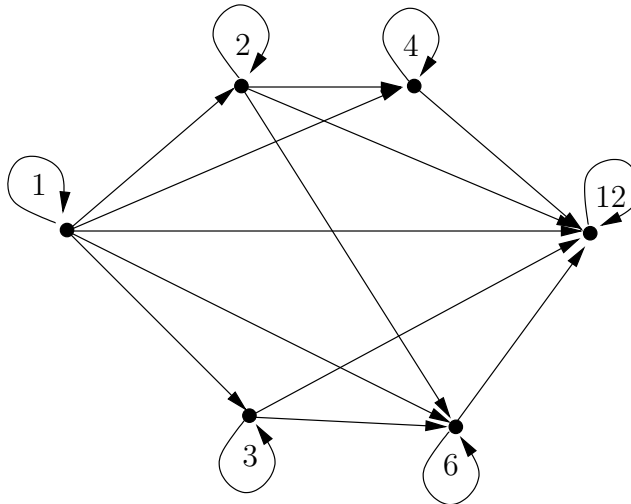
³If you insist on making the other choice, you will have a “strict ordering relation” a.k.a. an “irreflexive ordering relation”

Definition. If x and y are elements of a set S and R is an ordering relation on S then we say x and y are incomparable if neither xRy nor yRx is true.

Consider the set $S = \{1, 2, 3, 4, 6, 12\}$. If we look at the relation \leq on this set we get the following digraph.



On the other hand, perhaps you noticed these numbers are the divisors of 12. The divisibility relation will give us our first example of a partial order.



Exercise. Which elements in the above partial order are incomparable?

A set together with an ordering relation creates a mathematical structure known as a *partially ordered set*. Since that is a bit of a mouthful, the

abbreviated form *poset* is actually heard more commonly. If one wishes to refer to a poset it is necessary to identify both the set and the ordering relation. Thus, if S is a set and R is an ordering relation, we write (S, R) to denote the corresponding poset.

The digraphs given above for two posets having the same underlying set provide an existence proof – the same set may have different orders imposed upon it. They also highlight another issue – these digraphs for ordering relations get pretty crowded! Hasse diagrams for posets (named after the famous German mathematician Helmut Hasse) are a way of displaying all the information in a poset’s digraph, but much more succinctly. There are features of a Hasse diagram that correspond to each of the properties that an ordering relation must have.

Since ordering relations are always reflexive, there will always be loops at every vertex in the digraph. In a Hasse diagram we leave out the loops.

Since ordering relations are anti-symmetric, every edge in the digraph will go in one direction or the other. In a Hasse diagram we arrange the vertices so that that direction is *upward* – that way we can leave out all the arrowheads without losing any information.

The final simplification that we make in creating a Hasse diagram for a poset has to do with the transitivity property – we leave out any connections that could be deduced because of transitivity.

Hasse diagrams for the two orderings that we’ve been discussing are shown in Figure 6.5

Often there is some feature of the elements of the set being ordered that allows us to arrange a Hasse diagram in “ranks.” For example, consider $\mathcal{P}(\{1, 2, 3\})$, the set of all subsets of a three element set – this set can be partially ordered using the \subseteq relation. (Technically, we should verify that this relation is reflexive, anti-symmetric and transitive before proceeding, but by now you know why subset containment is denoted using a rounded version of \leq .) Subsets of the same size can’t possibly be included one in the other

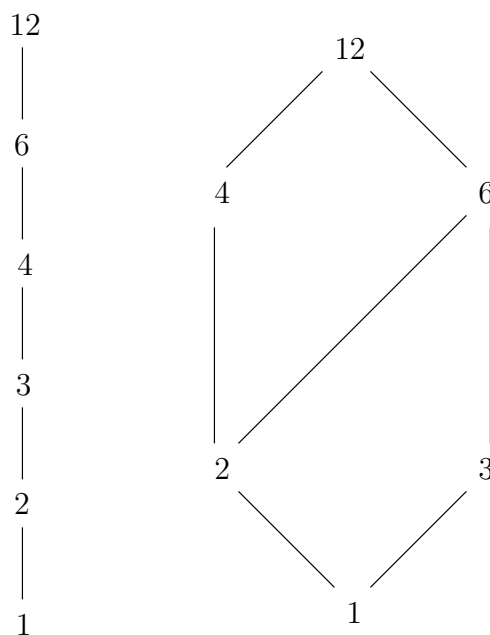


Figure 6.5: Hasse diagrams of the set $\{1, 2, 3, 4, 6, 12\}$ totally ordered by \leq and partially ordered by $|$.

unless they happen to be equal! This allows us to draw the Hasse diagram for this set with the nodes arranged in four rows. (See Figure 6.6.)

Exercise. Try drawing a Hasse diagram for the partially ordered set

$$(\mathcal{P}(\{1, 2, 3, 4\}), \subseteq).$$

Posets like $(\mathcal{P}(\{1, 2, 3\}), \subseteq)$ that can be laid out in ranks are known as *graded posets*. Things in a graded poset that have the same rank are always incomparable.

Definition. A graded poset is a triple (S, R, ρ) , where S is a set, R is an ordering relation, and ρ is a function from S to \mathbb{Z} .

In the example we've been considering (the graded poset of subsets of a set partially ordered by set inclusion), the grading function ρ takes a subset

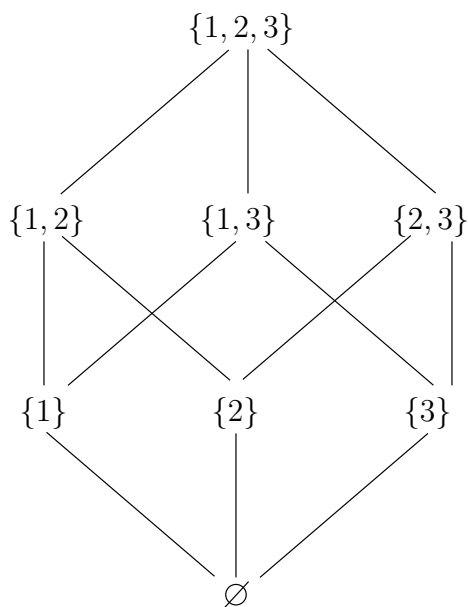


Figure 6.6: Hasse diagram for the power set of $\{1, 2, 3\}$ partially ordered by set containment.

to its size. That is, $\rho(A) = |A|$. Another nice example of a graded poset is the set of divisors of some number partially ordered by the divisibility relation (\mid). In this case the grading function takes a number to its total degree – the sum of all the exponents appearing in its prime factorization. In Figure 6.7 we show the poset of divisors of 72 and indicate the grading.

We will end this section by giving a small collection of terminology relevant to partially ordered sets.

A *chain* in a poset is a subset of the elements, all of which are comparable. If you restrict your attention to a chain within a poset, you will be looking at a total order. An *antichain* in a poset is a subset of the elements, none of which are comparable. Thus, for example, a subset of elements having the same rank (in a graded poset) is an antichain. Chains and antichains are said to be *maximal* if it is not possible to add further elements to them (whilst maintaining the properties that make them chains and/or antichains). An

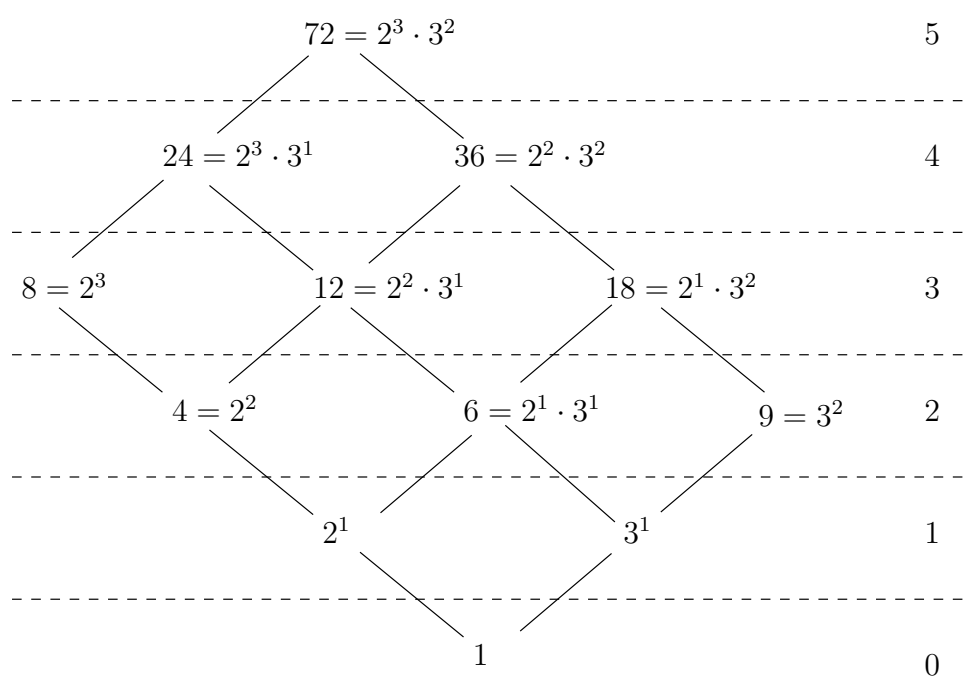


Figure 6.7: Hasse diagram for the divisors of 72, partially ordered by divisibility. This is a graded poset.

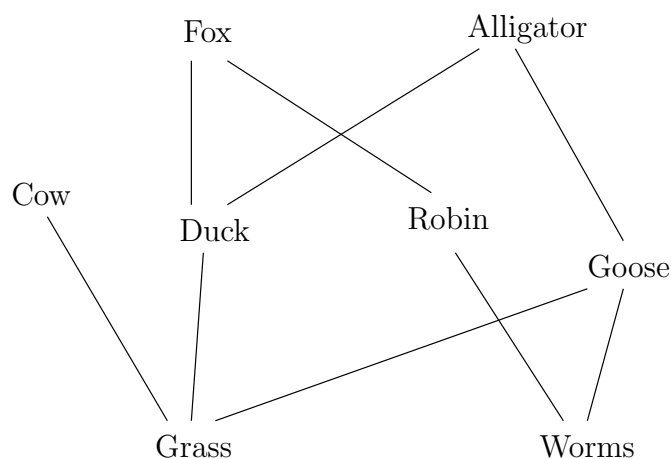
element x , that appears above another element y – and connected to it – in a Hasse diagram is said to *cover* it. In this situation you may also say that x is an *immediate successor* of y . A *maximal element* is an element that is not covered by any other element. Similarly, a *minimal element* is an element that is not a cover of any other element. If a chain is maximal, it follows that it must contain both a maximal and a minimal element (with respect to the surrounding poset). The collection of all maximal elements forms an antichain, as does (separately) the collection of all minimal elements. Finally, we have the notions of *greatest element* (a.k.a. top) and *least element* (a.k.a. bottom) – the greatest element is greater than every other element in the poset, the least element is smaller than every other element. Please be careful to distinguish these concepts from maximal and minimal elements – a greatest element is automatically maximal, and a least element is always minimal, but it is possible to have a poset with no greatest element that nevertheless has one or more maximal elements, and it is possible to have a poset with no least element that has one or more minimal elements.

In the poset of divisors of 72, the subset $\{2, 6, 12, 24\}$ is a chain. Since it would be possible to add both 1 and 72 to this chain and still have a chain, this chain is not maximal. (But, of course, $\{1, 2, 6, 12, 24, 72\}$ is.) On the other hand, $\{8, 12, 18\}$ is an antichain (indeed, this is a maximal antichain). This poset has both a top and a bottom – 1 is the least element and 72 is the greatest element. Notice that the elements which cover 1 (the least element) are the prime divisors of 72.

Exercises — 6.4

1. In population ecology there is a partial order “predates” which basically means that one organism feeds upon another. Strictly speaking this relation is not transitive; however, if we take the point of view that when a wolf eats a sheep, it is also eating some of the grass that the sheep has fed upon, we see that in a certain sense it is transitive. A chain in this partial order is called a “food chain” and so-called apex predators are said to “sit atop the food chain”. Thus “apex predator” is a term for a maximal element in this poset. When poisons such as mercury and PCBs are introduced into an ecosystem, they tend to collect disproportionately in the apex predators – which is why pregnant women and young children should not eat shark or tuna but sardines are fine.

Below is a small example of an ecology partially ordered by “predates”

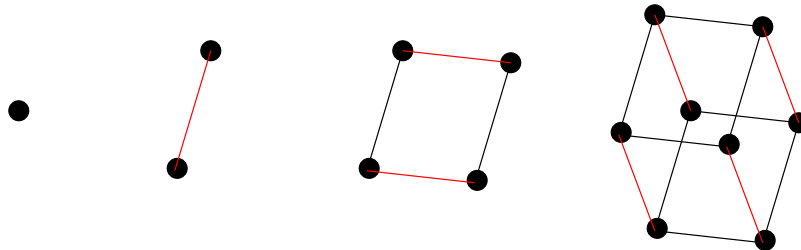


Find the largest antichain in this poset.

2. Referring to the poset given in exercise 1, match the following.

- | | |
|-------------------------------|--------------------------|
| 1. An (non-maximal) antichain | a. Grass |
| 2. A maximal antichain | b. Goose |
| 3. A maximal element | c. Fox |
| 4. A (non-maximal) chain | d. {Grass, Duck} |
| 5. A maximal chain | e. There isn't one! |
| 6. A cover for "Worms" | f. {Fox, Alligator, Cow} |
| 7. A least element | g. {Cow, Duck, Goose} |
| 8. A minimal element | h. {Worms, Robin, Fox} |

3. The graph of the edges of a cube is one in an infinite sequence of graphs. These graphs are defined recursively by "Make two copies of the previous graph then join corresponding nodes in the two copies with edges." The 0-dimensional 'cube' is just a single point. The 1-dimensional cube is a single edge with a node at either end. The 2-dimensional cube is actually a square and the 3-dimensional cube is what we usually mean when we say "cube."



Make a careful drawing of a *hypercube* – which is the name of the graph that follows the ordinary cube in this sequence.

4. Label the nodes of a hypercube with the divisors of 210 in order to produce a Hasse diagram of the poset determined by the divisibility relation.
5. Label the nodes of a hypercube with the subsets of $\{a, b, c, d\}$ in order to produce a Hasse diagram of the poset determined by the subset containment relation.
6. Complete a Hasse diagram for the poset of divisors of 11025 (partially ordered by divisibility).
7. Find a collection of sets so that, when they are partially ordered by \subseteq , we obtain the same Hasse diagram as in the previous problem.

6.5 Functions

The concept of a function is one of the most useful abstractions in mathematics. In fact it is an abstraction that can be further abstracted! For instance an *operator* is an entity which takes functions as inputs and produces functions as outputs, thus an operator is to functions as functions themselves are to numbers. There are many operators that you have certainly encountered already – just not by that name. One of the most famous operators is “differentiation,” when you take the derivative of some function, the answer you obtain is another function. If two different people are given the same differentiation problem and they come up with different answers, we *know* that at least one of them has made a mistake! Similarly, if two calculations of the value of a function are made for the same input, they *must* match.

The property we are discussing used to be captured by saying that a function needs to be “well-defined.” The old school definition of a function was:

Definition. A function f is a well-defined rule, that, given any input value x produces a unique output⁴ value $f(x)$.

A more modern definition of a function is the following.

Definition. A function is a binary relation which does not contain distinct pairs having the same initial element.

When we think of a function as a special type of binary relation, the pairs that are “in” the function have the form $(x, f(x))$, that is, they consist of an input and the corresponding output.

We have gotten relatively used to relations “on” a set, but recall that the more general situation is that a binary relation is a subset of $A \times B$. In this setting, if the relation is actually a function f , we say that f is a function

⁴The use of the notation $f(x)$ to indicate the output of function f associated with input x was instituted by Leonard Euler, and so it is known as Euler notation.

from A to B . Now, quite often there are input values that simply don't work for a given function (for instance the well-known "you can't take the square root of a negative" rule). Also, it is often the case that certain outputs just can't happen. So, when dealing with a function as a relation contained in $A \times B$ there are actually four sets that are of interest – the sets A and B (of course) but also some sets that we'll denote by A' and B' . The set A' consists of those elements of A that actually appear as the first coordinate of a pair in the relation f . The set B' consists of those elements of B that actually appear as the second coordinate of a pair in the relation f . A generic example of how these four sets might look is given in Figure 6.8.

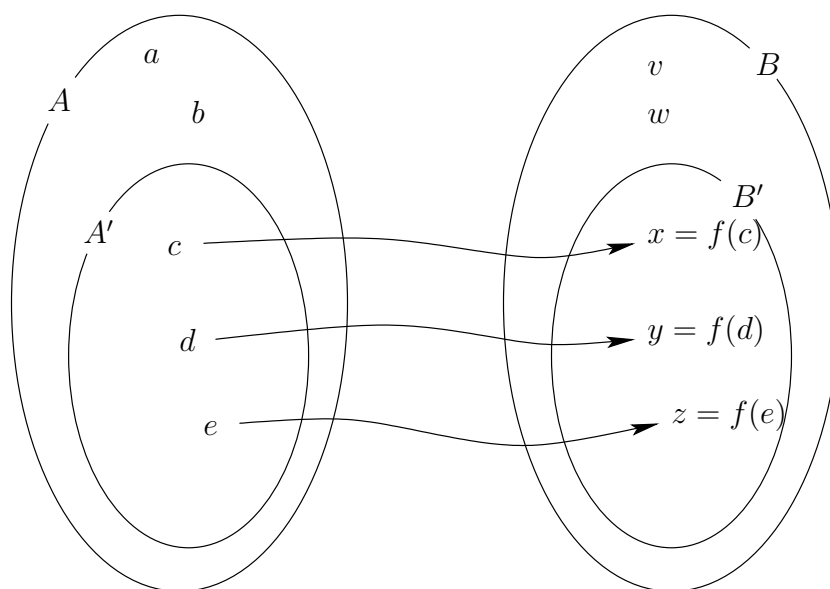


Figure 6.8: The sets related to an arbitrary function.

Sadly, only three of the sets we have just discussed are known to the mathematical world. The set we have denoted A' is called the *domain* of the function f . The set we have denoted B' is known as the *range* of the function f . The set we have denoted B is called the *codomain* of the function f . The set we have been calling A does not have a name. In fact, the formal definition of the term "function" has been rigged so that there is no difference

between the sets A and A' . This seems a shame, if you think of range and domain as being primary, doesn't it seem odd that we have a way to refer to a superset of the range (i.e. the codomain) but no way of referring to a superset of the domain?

Nevertheless, this is just the way it is ... There is only one set on the input side – the domain of our function.

The domain of any relation is expressed by writing $\text{Dom}(R)$. Which is defined as follows.

Definition. *If R is a relation from A to B then $\text{Dom}(R)$ is a subset of A defined by*

$$\text{Dom}(R) = \{a \in A \mid \exists b \in B, (a, b) \in R\}$$

We should point out that the notation just given for the domain of a relation R , ($\text{Dom}(R)$) has analogs for the other sets that are involved with a relation. We write $\text{Cod}(R)$ to refer to the codomain of the relation, and $\text{Rng}(R)$ to refer to the range.

Since we are now thinking of functions as special classes of relations, it follows that a function is just a set of ordered pairs. This means that the identity of a function is tied up, not just with a formula that gives the output for a given input, but also with what values can be used for those inputs. Thus the function $f(x) = 2x$ defined on \mathbb{R} is a completely different animal from the function $f(x) = 2x$ defined on $\mathbb{Z}^{\text{nonneg}}$. If you really want to specify a function precisely you must give its domain as well as a formula for it. Usually, one does this by writing a formula, then a semicolon, then the domain. (E.g. $f(x) = x^2; \quad x \geq 0$.)

Okay, so, finally, we are prepared to give the real definition of a function.

Definition. *If A and B are sets, then f is a function from A to B (which is expressed symbolically by $f : A \rightarrow B$), if and only if f is a subset of $A \times B$, $\text{Dom}(f) = A$ and $((a, b) \in f \wedge (a, c) \in f \implies b = c$.*

Recapping, a function *must* have its domain equal to the set A where its inputs come from. This is sometimes expressed by saying that a function is *defined* on its domain. A function's range and codomain may be different however. In the event that the range and codomain *are* the same ($\text{Cod}(R) = \text{Rng}(R)$) we have a rather special situation and the function is graced by the appellation "surjection." The term "onto" is also commonly used to describe a surjective function.

Exercise. *There is an expression in mathematics, "Every function is onto its range." that really doesn't say very much. Why not?*

If one has elements x and y , of the domain and codomain, (respectively) and $y = f(x)$ ⁵ then one may say that " y is the image of x " or that " x is a preimage of y ." Take careful note of the articles used in these phrases – we say " y is **the** image of x " but " x is **a** preimage of y ." This is because y is uniquely determined by x , but not vice versa. For example, since the squares of 2 and -2 are both 4, if we consider the function $f(x) = x^2$, the image of (say) 2 is 4, but a preimage for 4 could be either 2 or -2 .

It would be pleasant if there were a nice way to refer to the preimage of some element, y , of the range. One notation that you have probably seen before is " $f^{-1}(y)$." There is a major difficulty with writing down such a thing. By writing " f^{-1} " you are making a rather vast presumption – that there actually is a function that serves as an inverse for f . Usually, there is not.

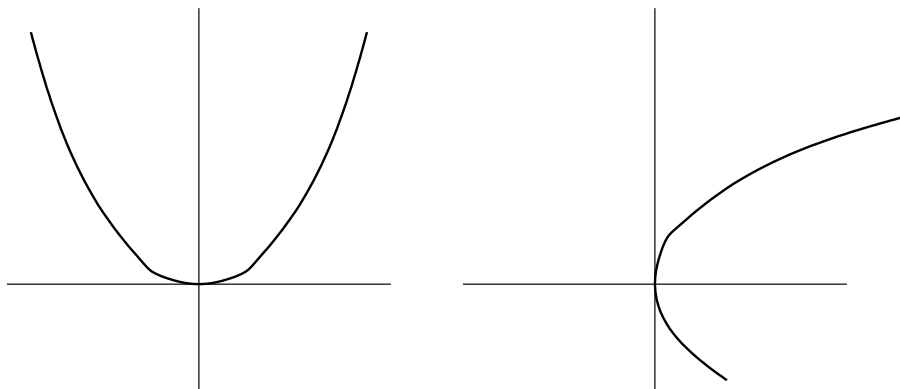
One can define an inverse for any relation, the inverse is formed by simply exchanging the elements in the ordered pairs that make up R .

Definition. *The inverse relation of a relation R is denoted R^{-1} and*

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

⁵Or, equivalently, $(x, y) \in f$.

In terms of graphs, the inverse and the original relation are related by being reflections in the line $y = x$. It is possible for one, both, or neither of these to be functions. The canonical example to keep in mind is probably $f(x) = x^2$ and its inverse.



The graph that we obtain by reflecting $y = f(x) = x^2$ in the line $y = x$ doesn't pass the vertical line test and so it is the graph of (merely) a relation – not of a function. The function $g(x) = \sqrt{x}$ that we all know and love is not truly the inverse of $f(x)$. In fact this function is defined to make a specific (and natural) choice – it returns the positive square root of a number. But this leads to a subtle problem; if we start with a negative number (say -3) and square it we get a positive number (9) and if we then come along and take the square root we get another positive number (3). This is problematic since we didn't end up where we started which is what ought to happen if we apply a function followed by its inverse.

We'll try to handle the general situation in a bit, but for the moment let's consider the nice case: when the inverse of a function is also a function. When exactly does this happen? Well, we have just seen that the inverse of a function doesn't necessarily pass the vertical line test, and it turns out that that is the predominant issue. So, under what circumstances does the inverse pass the vertical line test? When the original function passes the so-called horizontal line test (every horizontal line intersects the graph at most

once). Thinking again about $f(x) = x^2$, there are some horizontal lines that miss the graph entirely, but all horizontal lines of the form $y = c$ where c is positive will intersect the graph twice. There are many functions that *do* pass the horizontal line test, for instance, consider $f(x) = x^3$. Such functions are known as *injections*, this is the same thing as saying a function is “one-to-one.” Injective functions can be inverted – the domain of the inverse function of f will only be the range, $\text{Rng}(f)$, which as we have seen may fall short of the being the entire codomain, since $\text{Rng}(f) \subseteq \text{Cod}(f)$.

Let’s first define injections in a way that is divorced from thinking about their graphs.

Definition. A function $f(x)$ is an injection iff for all pairs of inputs x_1 and x_2 , if $f(x_1) = f(x_2)$ then $x_1 = x_2$.

This is another of those defining properties that is designed so that when it is true it is vacuously true. An injective function never takes two distinct inputs to the same output. Perhaps the cleanest way to think about injective functions is in terms of preimages – when a function is injective, preimages are unique. Actually, this is a good time to mention something about surjective functions and preimages – if a function is surjective, every element of the codomain *has* a preimage. So, if a function has both of these properties it means that every element of the codomain has one (and only one) preimage.

A function that is both injective and surjective (one-to-one and onto) is known as a *bijection*. Bijections are tremendously important in mathematics since they provide a way of perfectly matching up the elements of two sets. You will probably spend a good bit of time in the future devising maps between sets and then proving that they are bijections, so we will start practicing that skill now. . .

Ordinarily, we will show that a function is a bijection by proving separately that it is both a surjection and an injection.

To show that a function is surjective we need to show that it is possible

to find a preimage for every element of the codomain. If we happen to know what the inverse function is, then it is easy to find a preimage for an arbitrary element. In terms of the taxonomy for proofs that was introduced in Chapter 3, we are talking about a constructive proof of an existential statement. A function f is surjective iff $\forall y \in \text{Cod}(f), \exists x \in \text{Dom}(f), y = f(x)$, so to prove surjectivity is to find the x that “works” for an arbitrary y . If this is done by literally naming x , we have proved the statement constructively.

To show that a function is an injection, we traditionally prove that the property used in the definition of an injective function is true. Namely, we suppose that x_1 and x_2 are distinct elements of $\text{Dom}(f)$ and that $f(x_1) = f(x_2)$ and then we show that actually $x_1 = x_2$. This is in the spirit of a proof by contradiction – if there were actually distinct elements that get mapped to the same value then f would *not* be injective, but by deducing that $x_1 = x_2$ we are contradicting that presumption and so, are showing that f is indeed an injection.

Let’s start by looking at a very simple example, $f(x) = 2x - 1$; $x \in \mathbb{Z}^{\text{nonneg}}$. Clearly this function is not a surjection if we are thinking that $\text{Cod}(f) = \mathbb{Z}^{\text{nonneg}}$ since the outputs are always odd. Let $\mathcal{O} = \{1, 3, 5, 7, \dots\}$ be the set of odd naturals.

Theorem 6.5.1. *The function $f : \mathbb{Z}^{\text{nonneg}} \longrightarrow \mathcal{O}$ defined by $f(x) = 2x - 1$ is a bijection from $\mathbb{Z}^{\text{nonneg}}$ to \mathcal{O} .*

Proof: First we will show that f is surjective. Consider an arbitrary element y of the set \mathcal{O} . Since $y \in \mathcal{O}$ it follows that y is both positive and odd. Thus there is an integer k , such that $y = 2k + 1$, but also $y > 0$. From this it follows that $2k + 1 > 0$ and so $k > -1/2$. Since k is also an integer, this last inequality implies that $k \in \mathbb{Z}^{\text{nonneg}}$. (Recall that $\mathbb{Z}^{\text{nonneg}} = \{0, 1, 2, 3, \dots\}$.) We can easily verify that a preimage for y is $k + 1$, since $f(k + 1) = 2(k + 1) - 1 = 2k + 2 - 1 = 2k + 1 = y$.

Next we show that f is injective. Suppose that there are two input values, x_1 and x_2 such that $f(x_1) = f(x_2)$. Then $2x_1 - 1 = 2x_2 - 1$ and simple algebra leads to $x_1 = x_2$.

Q.E.D.

For a slightly more complicated example consider the function from $\mathbb{Z}^{\text{nonneg}}$ to \mathbb{Z} defined by

$$f(x) = \begin{cases} x/2 & \text{if } x \text{ is even} \\ -(x-1)/2 & \text{if } x \text{ is odd} \end{cases}$$

This function does quite a handy little job, it matches up the natural numbers and the integers in pairs. Every even natural gets matched with a positive integer and every odd natural (except 1) gets matched with a negative integer (1 gets paired with 0). This function is really doing something remarkable – common sense would seem to indicate that the integers must be a larger set than the naturals (after all $\mathbb{Z}^{\text{nonneg}}$ is completely contained inside of \mathbb{Z}), but the function f defined above serves to show that these two sets are *exactly the same size!*

Theorem 6.5.2. *The function f defined above is bijective.*

Proof: First we will show that f is surjective.

It suffices to find a preimage for an arbitrary element of \mathbb{Z} . Suppose that y is a particular but arbitrarily chosen integer. There are two cases to consider: $y \leq 0$ and $y > 0$.

If $y > 0$ then $x = 2y$ is a preimage for y . This follows easily since $x = 2y$ is obviously even and so x 's image will be defined by the first case in the definition of f . Thus $f(x) = f(2y) = (2y)/2 = y$.

If $y \leq 0$ then $x = 1 - 2y$ is a preimage for y . Clearly, $1 - 2y$ is odd whenever y is an integer, thus this value for x will fall into

the second case in the definition of f . So, $f(x) = f(1 - 2y) = -((1 - 2y) - 1)/2 = -(-2y)/2 = y$.

Since the cases $y > 0$ and $y \leq 0$ are exhaustive (that is, every y in \mathbb{Z} falls into one or the other of these cases), and we have found a preimage for y in both cases, it follows that f is surjective.

Next, we will show that f is injective.

Suppose that x_1 and x_2 are elements of $\mathbb{Z}^{\text{nonneg}}$ and that $f(x_1) = f(x_2)$. Consider the following three cases: x_1 and x_2 are both even, both odd, or have opposite parity.

If x_1 and x_2 are both even, then by the definition of f we have $f(x_1) = x_1/2$ and $f(x_2) = x_2/2$ and since these functional values are equal, we have $x_1/2 = x_2/2$. Doubling both sides of this leads to $x_1 = x_2$.

If x_1 and x_2 are both odd, then by the definition of f we have $f(x_1) = -(x_1 - 1)/2$ and $f(x_2) = -(x_2 - 1)/2$ and since these functional values are equal, we have $-(x_1 - 1)/2 = -(x_2 - 1)/2$. A bit more algebra (doubling, negating and adding one to both sides) leads to $x_1 = x_2$.

If x_1 and x_2 have opposite parity, we will assume w.l.o.g. that x_1 is even and x_2 is odd. The equality $f(x_1) = f(x_2)$ becomes $x_1/2 = -(x_2 - 1)/2$. Note that $x_1 \geq 2$ so $f(x_1) = x_1/2 \geq 1$. Also, note that $x_2 \geq 1$ so

$$\begin{aligned} x_2 - 1 &\geq 0 \\ (x_2 - 1)/2 &\geq 0 \\ -(x_2 - 1)/2 &\leq 0 \\ f(x_2) &\leq 0 \end{aligned}$$

therefore we have a contradiction since it is impossible for the two values $f(x_1)$ and $f(x_2)$ to be equal while $f(x_1) \geq 1$ and $f(x_2) \leq 0$. Since the last case under consideration leads to a contradiction, it follows that x_1 and x_2 never have opposite parities, and so the first two cases are exhaustive – in both of those cases we reached the desired conclusion that $x_1 = x_2$ so it follows that f is injective.

Q.E.D.

We'll conclude this section by mentioning that the ideas of “image” and “preimage” can be extended to sets. If S is a subset of $\text{Dom}(f)$ then the *image of S under f* is denoted $f(S)$ and

$$f(S) = \{y \mid \exists x \in \text{Dom}(f), x \in S \wedge y = f(x)\}.$$

Similarly, if T is a subset of $\text{Rng}(f)$ we can define something akin to the preimage. The *inverse image of the set T under the function f* is denoted $f^{-1}(T)$ and

$$f^{-1}(T) = \{x \mid \exists y \in \text{Cod}(f), y \in T \wedge y = f(x)\}.$$

Essentially, we have extended the function f so that it goes between the power sets of its codomain and range! This new notion gives us some elegant ways of restating what it means to be surjective and injective.

A function f is surjective iff $f(\text{Dom}(f)) = \text{Cod}(f)$.

A function f is injective iff the inverse images of singletons are always singletons. That is,

$$\forall y \in \text{Rng}(f), \exists x \in \text{Dom}(f), f^{-1}(\{y\}) = \{x\}.$$

Exercises — 6.5

1. For each of the following functions, give its domain, range and a possible codomain.

(a) $f(x) = \sin(x)$

(b) $g(x) = e^x$

(c) $h(x) = x^2$

(d) $m(x) = \frac{x^2+1}{x^2-1}$

(e) $n(x) = \lfloor x \rfloor$

(f) $p(x) = \langle \cos(x), \sin(x) \rangle$

2. Find a bijection from the set of odd squares, $\{1, 9, 25, 49, \dots\}$, to the non-negative integers, $\mathbb{Z}^{\text{nonneg}} = \{0, 1, 2, 3, \dots\}$. Prove that the function you just determined is both injective and surjective. Find the inverse function of the bijection above.
3. The natural logarithm function $\ln(x)$ is defined by a definite integral with the variable x in the upper limit.

$$\ln(x) = \int_{t=1}^x \frac{1}{t} dt.$$

From this definition we can deduce that $\ln(x)$ is strictly increasing on its entire domain, $(0, \infty)$. Why is this true?

We can use the above definition with $x = 2$ to find the value of $\ln(2) \approx .693$. We will also take as given the following rule (which is valid for all logarithmic functions).

$$\ln(a^b) = b \ln(a)$$

Use the above information to show that there is neither an upper bound nor a lower bound for the values of the natural logarithm. These facts together with the information that \ln is strictly increasing show that $\text{Rng}(\ln) = \mathbb{R}$.

4. Georg Cantor developed a systematic way of listing the rational numbers. By “listing” a set one is actually developing a bijection from $\mathbb{Z}^{\text{nonneg}}$ to that set. The method known as “Cantor’s Snake” creates a bijection from the naturals to the non-negative rationals. First we create an infinite table whose rows are indexed by positive integers and whose columns are indexed by non-negative integers – the entries in this table are rational numbers of the form “column index” / “row index.” We then follow a snake-like path that zig-zags across this table – whenever we encounter a rational number that we haven’t seen before (in lower terms) we write it down. This is indicated in the diagram below by circling the entries.

	0	1	2	3	4	5	6	7	8
1	0/1	1/1	2/1	3/1	4/1	5/1	6/1	7/1	8/1
2	0/2	1/2	2/2	3/2	4/2	5/2	6/2	7/2	8/2
3	0/3	1/3	2/3	3/3	4/3	5/3	6/3	7/3	8/3
4	0/4	1/4	2/4	3/4	4/4	5/4	6/4	7/4	8/4
5	0/5	1/5	2/5	3/5	4/5	5/5	6/5	7/5	8/5
6	0/6	1/6	2/6	3/6	4/6	5/6	6/6	7/6	8/6
7	0/7	1/7	2/7	3/7	4/7	5/7	6/7	7/7	8/7
8	0/8	1/8	2/8	3/8	4/8	5/8	6/8	7/8	8/8

Effectively this gives us a function f which produces the rational number that would be found in a given position in this list. For example $f(1) = 0/1$, $f(2) = 1/1$ and $f(5) = 1/3$.

What is $f(26)$? What is $f(30)$? What is $f^{-1}(3/4)$? What is $f^{-1}(6/7)$?

6.6 Special functions

There are a great many functions that fail the horizontal line test which we nevertheless seem to have inverse functions for. For example, x^2 fails HLT but \sqrt{x} is a pretty reasonable inverse for it – one just needs to be careful about the “plus or minus” issue. Also, $\sin x$ fails HLT pretty badly; any horizontal line $y = c$ with $-1 \leq c \leq 1$ will hit $\sin x$ infinitely many times. But look! Right here on my calculator is a button labeled “ \sin^{-1} .”⁶ This apparent contradiction can be resolved using the notion of restriction.

Definition. *Given a function f and a subset D of its domain, the restriction of f to D is denoted $f|_D$ and*

$$f|_D = \{(x, y) \mid x \in D \wedge (x, y) \in f\}.$$

The way we typically use restriction is to eliminate any regions in $\text{Dom}(f)$ that cause f to fail to be one-to-one. That is, we choose a subset $D \subseteq \text{Dom}(f)$ so that $f|_D$ is an injection. This allows us to invert the restricted version of f . There can be problems in doing this, but if we are careful about how we choose D , these problems are usually resolvable.

Exercise. *Suppose f is a function that is not one-to-one, and D is a subset of $\text{Dom}(f)$ such that $f|_D$ is one-to-one. The restricted function $f|_D$ has an inverse which we will denote by g . Note that g is a function from $\text{Rng}(f|_D)$ to D . Which of the following is always true:*

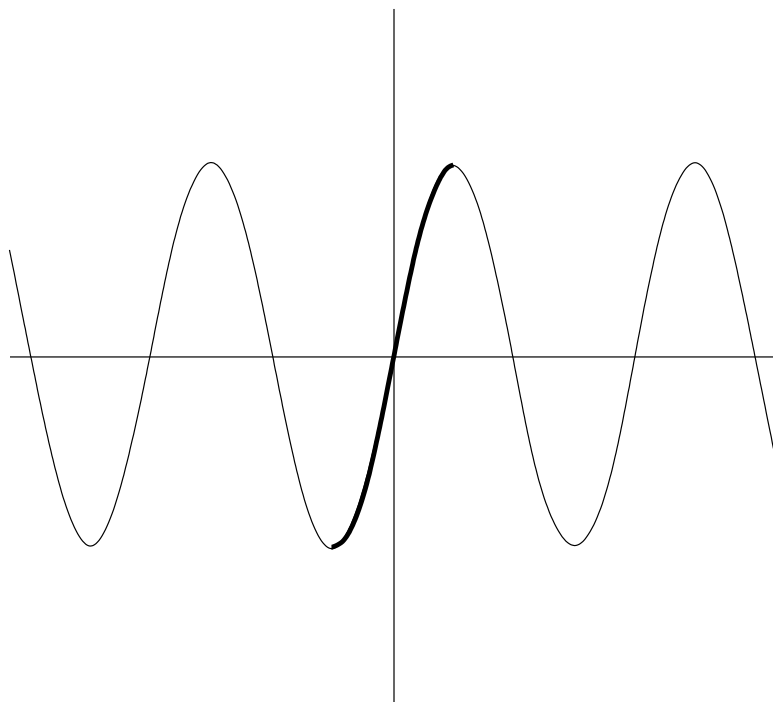
$$f(g(x)) = x \quad \text{or} \quad g(f(x)) = x?$$

Technically, when we do the process outlined above (choose a domain D so that the restriction $f|_D$ is invertible, and find that inverse) the function we get is a *right inverse* for f .

⁶It might be labeled “asin” instead. The old-style way to refer to the inverse of a trig function was arc-whatever. So the inverse of sine was arcsine, the inverse of tangent was arctangent.

Let's take a closer look at the inverse sine function. This should help us to really understand the "right inverse" concept.

A glance at the graph of $y = \sin x$ will certainly convince us that this function is not injective, but the portion of the graph shown in bold below passes the horizontal line test.



If we restrict the domain of the sine function to the closed interval $[-\pi/2, \pi/2]$, we have an invertible function. The inverse of this restricted function is the function we know as $\sin^{-1}(x)$ or $\arcsin(x)$. The domain and range of $\sin^{-1}(x)$ are (respectively) the intervals $[-1, 1]$ and $[-\pi/2, \pi/2]$.

Notice that if we choose a number x in the range $-1 \leq x \leq 1$ and apply the inverse sine function to it, we will get a number between $-\pi/2$ and $\pi/2$ – i.e. a number we can interpret as an *angle* in radian measure. If we then proceed to calculate the sine of this angle, we will get back our original number x .

On the other hand, if we choose an angle first, then take the sine of it to get a number in $[-1, 1]$ and then take the inverse sine of *that*, we will only end up with the same angle we started with **if** we chose the original angle so that it lay in the interval $[-\pi/2, \pi/2]$.

Exercise. We get a right inverse for the cosine function by restricting it to the interval $[0, \pi]$. What are the domain and range of \cos^{-1} ?

The *winding map* is a function that goes from \mathbb{R} to the unit circle in the x - y plane, defined by

$$W(t) = (\cos t, \sin t).$$

One can think of this map as literally winding the infinitely long real line around and around the circle. Obviously, this is not an injection – there are an infinite number of values of t that get mapped to (for instance) the point $(1, 0)$, t can be any integer multiple of 2π .

Exercise. What is the set $W^{-1}(\{(0, 1)\})$?

If we restrict W to the half-open interval $[0, 2\pi)$ the restricted function $W|_{[0, 2\pi)}$ is an injection. The inverse function is not easy to write down, but it is possible to express (in terms of the inverse functions of sine and cosine) if we consider the four cases determined by what quadrant a point on the unit circle may lie in.

Exercise. Suppose (x, y) represents a point on the unit circle. If (x, y) happens to lie on one of the coordinate axes we have

$$\begin{aligned}W^{-1}((1, 0)) &= 0 \\W^{-1}((0, 1)) &= \pi/2 \\W^{-1}((-1, 0)) &= \pi \\W^{-1}((0, -1)) &= 3\pi/2.\end{aligned}$$

If neither x nor y is zero, there are four cases to consider. Write an expression for $W^{-1}((x, y))$ using the cases (i) $x > 0 \wedge y > 0$, (ii) $x < 0 \wedge y > 0$, (iii) $x < 0 \wedge y < 0$ and (iv) $x > 0 \wedge y < 0$.

This last example that we have done (the winding map) was unusual in that the outputs were ordered pairs. In thinking of this map as a relation (that is, as a set of ordered pairs) we have an ordered pair in which the second element is an ordered pair! Just for fun, here is another way of expressing the winding map:

$$W = \{(t, (\cos t, \sin t)) \mid t \in \mathbb{R}\}$$

When dealing with very complicated expressions involving ordered pairs, or more generally, ordered n -tuples, it is useful to have a way to refer succinctly to the pieces of a tuple.

Let's start by considering the set $P = \mathbb{R} \times \mathbb{R}$ — i.e. P is the x - y plane. There are two functions, whose domain is P that “pick out” the x , and/or y coordinate. These functions are called π_1 and π_2 , π_1 is the projection onto the first coordinate and π_2 is the projection onto the second coordinate.⁷

Definition. *The function $\pi_1 : \mathbb{R} \times \mathbb{R} \longrightarrow \mathbb{R}$ known as projection onto the first coordinate is defined by*

⁷Don't think of the usual $\pi \approx 3.14159$ when looking at π_1 and π_2 . These functions are named as they are because π is the Greek letter corresponding to ‘p’ which stands for “projection.”

$$\pi_1((x, y)) = x.$$

The definition of π_2 is entirely analogous.

You should note that these projection functions are *very* bad as far as being one-to-one is concerned. For instance, the preimage of 1 under the map π_1 consists of all the points on the vertical line $x = 1$. That's a lot of preimages! These guys are so far from being one-to-one that it seems impossible to think of an appropriate restriction that would become invertible. Nevertheless, there is a function that provides a right inverse for both π_1 and π_2 . Now, these projection maps go from $\mathbb{R} \times \mathbb{R}$ to \mathbb{R} so an inverse needs to be a map from \mathbb{R} to $\mathbb{R} \times \mathbb{R}$. What is a reasonable way to produce a *pair* of real numbers if we have a single real number in hand? There are actually many ways one could proceed, but one reasonable choice is to create a pair where the input number appears in both coordinates. This is the so-called *diagonal map*, $d : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$, defined by $d(a) = (a, a)$.

Exercise. Which of the following is always true,

$$d(\pi_1((x, y))) = (x, y) \quad \text{or} \quad \pi_1(d(x)) = x?$$

There are a few other functions that it will be convenient to introduce at this stage. All of them are aspects of the characteristic function of a subset, so we'll start with that.

Whenever we have a subset/superset relationship, $S \subseteq D$, it is possible to define a function whose codomain is $\{0, 1\}$ which performs a very useful task – if an input x is in the set S the function will indicate this by returning 1, otherwise it will return 0. The function which has this behavior is known as 1_S , and is called the *characteristic function of the subset S* (There are those who use the term *indicator function of S* for 1_S .) By definition, D is the domain of this function.

$$1_S : D \longrightarrow \{0, 1\}$$

$$1_S(x) = \begin{cases} 1 & \text{if } x \in S \\ 0 & \text{otherwise} \end{cases}$$

Exercise. If you have the characteristic function of a subset S , how can you create the characteristic function of its complement, \overline{S} .

A characteristic function may be thought of as an embodiment of a membership criterion. The logical open sentence “ $x \in S$ ” being true is the same thing as the equation “ $1_S(x) = 1$.” There is a notation, growing in popularity, that does the same thing for an arbitrary open sentence. The *Iverson bracket* notation uses the shorthand $[P(x)]$ to represent a function that sends any x that makes $P(x)$ true to 1, and any inputs that make $P(x)$ false will get sent to 0.

$$[P(x)] = \begin{cases} 1 & \text{if } P(x) \\ 0 & \text{otherwise} \end{cases}$$

The Iverson brackets can be particularly useful in expressing and simplifying sums. For example, we can write $\sum_{i=1}^{24} [2|i]$ to find the number of even natural numbers less than 25. Similarly, we can write $\sum_{i=1}^{24} [3|i]$ to find the number of natural numbers less than 25 that are divisible by 3.

Exercise. What does the following formula count?

$$\sum_{i=1}^{24} [2|i] + [3|i] - [6|i]$$

There is a much more venerable notation known as the *Kronecker delta* that can be thought of as a special case of the idea inherent in Iverson brackets. We write δ_{ij} as a shorthand for a function that takes two inputs, $\delta(i, j)$ is 1 if and only if i and j are equal.

$$\delta_{ij} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{otherwise} \end{cases}$$

The corresponding Iverson bracket would simply be $[i = j]$.

We'll end this section with a function that will be especially important in Chapter 8. If we have an arbitrary subset of the natural numbers, we can associate it with an infinite string of 0's and 1's. By sticking a decimal point in front of such a thing, we get binary notation for a real number in the interval $[0, 1]$. There is a subtle problem that we'll deal with when we study this function in more detail in Chapter 8 — some real numbers can be expressed in two different ways in base 2. For example, $1/2$ can either be written as $.1$ or as $.0\bar{1}$ (where, as usual, the overline indicates a pattern that repeats forever). For the moment, we are talking about defining a function ϕ whose domain is $\mathcal{P}(\mathbb{Z}^{\text{nonneg}})$ and whose codomain is the set of all infinite binary strings. Let us denote these binary expansions by $.b_1b_2b_3b_4\dots$. Suppose A is a subset of $\mathbb{Z}^{\text{nonneg}}$, then the binary expansion associated with A will be determined by $b_i = 1_A(i)$. (Alternatively, we can use the Iverson bracket notation: $b_i = [i \in A]$.)

The function ϕ defined in the last paragraph turns out to be a bijection — given a subset we get a unique binary expansion, and given binary expansion we get (using ϕ^{-1}) a unique subset of $\mathbb{Z}^{\text{nonneg}}$.

A few examples will probably help to clarify this function's workings. Consider the set $\{1, 2, 3\} \subseteq \mathbb{Z}^{\text{nonneg}}$, the binary expansion that this corresponds to will have 1's in the first three positions after the decimal — $\phi(\{1, 2, 3\}) = .111$ this is the number written $.875$ in decimal. The infinite repeating binary number $.0\bar{1}$ is the base-2 representation of $1/3$, it is easy to see that $.0\bar{1}$ is the image of the set of odd naturals, $\{1, 3, 5, \dots\}$.

Exercise. Find the binary representation for the real number which is the image of the set of even numbers under ϕ .

Exercise. Find the binary representation for the real number which is the image of the set of triangular numbers under ϕ . (Recall that the triangular numbers are $T = \{1, 3, 6, 10, 15, \dots\}$.)

Exercises — 6.6

1. The n -th triangular number, denoted $T(n)$, is given by the formula $T(n) = (n^2 + n)/2$. If we regard this formula as a function from \mathbb{R} to \mathbb{R} , it fails the horizontal line test and so it is not invertible. Find a suitable restriction so that T is invertible.
2. The usual algebraic procedure for inverting $T(x) = (x^2 + x)/2$ fails. Use your knowledge of the geometry of functions and their inverses to find a formula for the inverse. (Hint: it may be instructive to first invert the simpler formula $S(x) = x^2/2$ — this will get you the right vertical scaling factor.)
3. What is $\pi_2(W(t))$?
4. Find a right inverse for $f(x) = |x|$.
5. In three-dimensional space we have projection functions that go onto the three coordinate axes (π_1 , π_2 and π_3) and we also have projections onto coordinate planes. For example, $\pi_{12} : \mathbb{R} \times \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R} \times \mathbb{R}$, defined by

$$\pi_{12}((x, y, z)) = (x, y)$$

is the projection onto the x - y coordinate plane.

The triple of functions $(\cos t, \sin t, t)$ is a parametric expression for a helix. Let $H = \{(\cos t, \sin t, t) \mid t \in \mathbb{R}\}$ be the set of all points on the helix. What is the set $\pi_{12}(H)$? What are the sets $\pi_{13}(H)$ and $\pi_{23}(H)$?

6. Consider the set $\{1, 2, 3, \dots, 10\}$. Express the characteristic function of the subset $S = \{1, 2, 3\}$ as a set of ordered pairs.
7. If S and T are subsets of a set D , what is the product of their characteristic functions $1_S \cdot 1_T$?

8. Evaluate the sum

$$\sum_{i=1}^{10} \frac{1}{i} \cdot [i \text{ is prime}].$$

Chapter 7

Proof techniques III — Combinatorics

Tragedy is when I cut my finger. Comedy is when you fall into an open sewer and die. —Mel Brooks

7.1 Counting

Many results in mathematics are answers to “How many . . .” questions.

“How many subsets does a finite set have?”

“How many handshakes will transpire when n people first meet?”

“How many functions are there from a set of size n to a set of size m ?”

The title of this section, “Counting,” is not intended to evoke the usual process of counting sheep, or counting change. What we want is to be able to count some collection *in principle* so that we will be able to discover a formula for its size.

There are two principles that will be indispensable in counting things. These principles are simple, yet powerful, and they have been named in the most unimaginative way possible. The “multiplication rule” which tells us when we should multiply, and the “addition rule” which tells us when we

should add.

Before we describe these principles in detail, we'll have a look at a simpler problem which is most easily described by an example: How many integers are there in the list $(7, 8, 9, \dots, 44)$? We could certainly write down all the integers from 7 to 44 (inclusive) and then count them – although this wouldn't be the best plan if the numbers 7 and 44 were replaced with (say) 7, 045, 356 and 22, 355, 201. A method that does lead to a generalized ability to count the elements of a finite sequence arises if we think carefully about what exactly a finite sequence *is*.

Definition. A sequence from a set S is a function from $\mathbb{Z}^{\text{nonneg}}$ to S .

Definition. A finite sequence from a set S is a function from $\{0, 1, 2, \dots, n\}$ to S , where n is some particular (finite) integer.

Now it is easy to see that there are $n+1$ elements in the set $\{0, 1, 2, \dots, n\}$ so counting the elements of a finite sequence will be easy if we can determine the function involved and figure out what n is by inverting it (n is an inverse image for the last element in a listing of the sequence).

In the example that we started with, the function is $f(x) = x + 7$. We can sum up the process that allows us to count the sequence by saying “there is a one-to-one correspondence between the lists

$$(7, 8, 9, \dots, 44)$$

and

$$(0, 1, 2, \dots, 37)$$

and the later has 38 entries.”

More generally, if there is a list of consecutive numbers beginning with k and ending with n , there will be $n - k + 1$ entries in the list. Lists of consecutive integers represent a relatively simple type of finite sequence. Usually we would have some slightly more interesting function that we'd need to invert.

The following exercise involves inverting the function $(x + 5)^2$.

Exercise. *How many integers are in the list $(25, 36, 49, \dots, 10000)$?*

We will have a lot more practice with counting the elements of sequences in the exercises at the end of this section, let's continue on our tour of counting by having a look at the addition rule.

The addition rule says that it is appropriate to add if we can partition a collection into *disjoint* pieces. In other words, if a set S is the union of two or more subsets and these subsets are mutually disjoint, we can find the size of S by adding the sizes of the subsets.

In the game Yahtzee, one rolls 5 dice and (optionally) performs a second roll of some or all of the dice. The object is to achieve several final configurations that are modeled after the hands in Poker. In particular, one configuration, known as a “full house,” is achieved by having two of one number and three of another. (Colloquially, we say “three-of-a-kind plus a pair is a full house.”)

Now, we could use Yahtzee “hands” to provide us with a whole collection of counting problems once we have our basic counting principles, but for the moment we just want to make a simple (and obvious) point about “full houses” – the pair is either smaller or larger than the three-of-a-kind. This means we can partition the set of all possible full houses into two disjoint sets – the full houses consisting of a small pair and a larger three-of-a-kind and those where the pair is larger than the three-of-a-kind. If we can find some way of counting these two cases separately, then the total number of full houses will be the sum of these numbers.

The multiplication rule gives us a way of counting things by thinking about how we might construct them. The numbers that are multiplied are the number of choices we have in the construction process. Surprisingly often, the number of choices we can make in a given stage of constructing some configuration is independent of the choices that have gone before – if

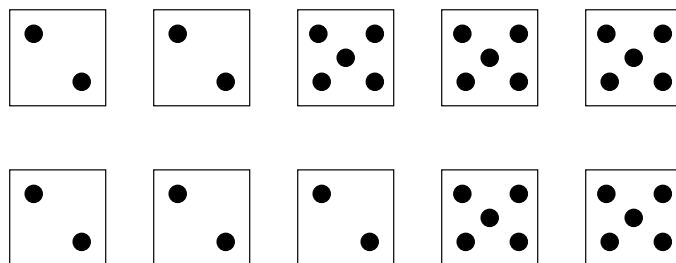


Figure 7.1: In Yahtzee, a full house may consist of a pair and a larger three-of-a-kind, or vice versa.

this is not the case the multiplication rule may not apply.

If some object can be constructed in k stages, and if in the first stage we have n_1 choices as to how to proceed, in the second stage we have n_2 choices, *et cetera*. Then the total number of such objects is the product $n_1 n_2 \cdots n_k$.

A *permutation of an n -set* (w.l.o.g. $\{1, 2, \dots, n\}$) is an ordered n -tuple where each entry is a distinct element of the n -set. Generally, a permutation may be regarded as a bijection from an n -set to itself. Our first use of the multiplication rule will be to count the total number of permutations of $\{1, 2, 3, \dots, n\}$.

Let's start by counting the permutations of $\{1, 2, 3\}$. A permutation will be a 3-tuple containing the numbers 1, 2 and 3 in some order. We will think about building such a thing in three stages. First, we must select a number to go in the first position – there are 3 choices. Having made that choice, there will only be two possibilities for the number in the second position. Finally there is just one number remaining to put in the third position¹. Thus there are $3 \cdot 2 \cdot 1 = 6$ permutations of a 3 element set.

The general rule is that there are $n!$ permutations of $\{1, 2, \dots, n\}$.

There are times when configurations that are like permutations (in that

¹People may say you have “no choice” in this last situation, but what they mean is that you have only one choice.

they are ordered and have no duplicates) but don't consist of all n numbers are useful.

Definition. A k -permutation from an n -set is an ordered selection of k distinct elements from a set of size n .

There are certain natural limitations on the value of k , for instance k can't be negative – although (arguably) k can be 0, it makes more sense to think of k being at least 1. Also, if k exceeds n we won't be able to find *any* k -permutations, since it will be impossible to meet the “distinct” requirement. If k and n are equal, there is no difference between a k -permutation and an ordinary permutation. Therefore, we ordinarily restrict k to lie in the range $0 < k < n$.

The notation $P(n, k)$ is used for the total number of k -permutations of a set of size n . For example, $P(4, 2)$ is 12, since there are twelve different ordered pairs having distinct entries where the entries come from $\{1, 2, 3, 4\}$.

Exercise. Write down all twelve 2-permutations of the 4-set $\{1, 2, 3, 4\}$.

Counting k -permutations using the multiplication rule is easy. We build a k -permutation in k stages. In stage 1, we pick the first element in the permutation – there are n possible choices. In stage 2, we pick the second element – there are now only $n - 1$ choices since we may not repeat the first entry. We keep going like this until we've picked k entries. The number $P(n, k)$ is the product of k numbers beginning with n and descending down to $n - k + 1$. To verify that $n - k + 1$ is really the right lower limit, check that there are indeed k entries in the sequence

$$(n, n - 1, n - 2, \dots, n - k + 1).$$

This verification may be easier if we rewrite the sequence as

$$(n - 0, n - 1, n - 2, \dots, n - (k - 1)).$$

Let's have a look at another small example – $P(8, 4)$. There will be 8 choices for the first entry in a 4-tuple, 7 choices for the second entry, 6 choices for the third entry and 5 choices for the last entry. (Note that $5 = 8 - 4 + 1$.) Thus $P(8, 4) = 8 \cdot 7 \cdot 6 \cdot 5 = 1680$.

Finally, we should take note that it is relatively easy to express $P(n, k)$ using factorials. If we divide a number factorialized by some smaller number factorialized, we will get a descending product just like those above.

Exercise. *What factorial would we divide $8!$ by in order to get $P(8, 4)$?*

The general rule is that $P(n, k) = \frac{n!}{(n-k)!}$.

If we were playing a card game in which we were dealt 5 cards from a deck of 52, we would receive our cards in the form of $P(52, 5) = 52 \cdot 51 \cdot 50 \cdot 49 \cdot 48 = 311875200$ ordered 5-tuples. Normally, we don't really care about what order the cards came to us in. In a card game one ordinarily begins sorting the cards so as to see what hand one has – this is a sure sign that the order the cards were dealt is actually immaterial. How many different orders can five cards be put in? The answer to this question is $5! = 120$ since what we are discussing is nothing more than a permutation of a set of size 5. Thus, if we say that there are 311,875,200 different possible hands in 5-card poker, we are over-counting things by quite a bit! Any given hand will appear 120 times in that tabulation, which means the right value is $311875200/120 = 2598960$.

Okay, so there around 2.6 million different hands in 5-card poker. Unless you plan to become a gambler this isn't really that useful of a piece of information – but if you generalize what we've done in the paragraph above, you'll have found a way to count unordered collections of a given size taken from a set.

A *k-combination from an n-set* is an unordered selection, without repetitions, of k things out of n . This is the exact same thing as a subset of size k of a set of size n , and the number of such things is denoted by several

different notations – $C(n, k)$, nCk and $\binom{n}{k}$ among them². We can come up with a formula for $C(n, k)$ by a slightly roundabout argument. Suppose we think of counting the k -permutations of n things using the multiplication rule in a different way than we have previously. We'll build a k -permutation in two stages. First we'll choose k symbols to put into our permutation – which can be done in $C(n, k)$ ways. And second, we'll put those k symbols into a particular order – which can be accomplished in $k!$ ways. Thus $P(n, k) = C(n, k) \cdot k!$. Since we already know that $P(n, k) = \frac{n!}{(n-k)!}$, we can substitute and solve to obtain

$$C(n, k) = \frac{n!}{k! \cdot (n - k)!}.$$

It is possible to partition many counting problems into 4 “types” based on the answers to two questions:

Is order important in the configurations being counted?

Are we allowed to have repeated elements in a configuration?

Suppose that we are in the general situation of selecting k things out of a set of size n . It should be possible to write formulas involving n and k in the four cells of the following table.

		Does order matter?	
		Yes	No
Are repeats okay?	No		
	Yes		

²Watch out for the $\binom{n}{k}$ notation, it is easy to confuse it with the fraction $(\frac{n}{k})$. They are not the same — the fraction bar is *supposed* to be missing in $\binom{n}{k}$.

Ordered with repetition

Selecting a PIN number³ for your bank account is a good example of the kind of problem that is dealt with in the lower left part of the table. Obviously, the order in which you key-in the digits of your PIN is important. If one's number is 1356, it won't do to put in 6531! Also there is no reason that we couldn't have repeated digits in a PIN. (Although someone who chooses a PIN like 3333 is taking a bit of a security risk! A bad guy looking over your shoulder may easily discern what your PIN is.) A PIN is an ordered selection of 4 things out of 10, where repetition is allowed. There are 10^4 possible PINs. We can determine this by thinking of the multiplication principle – there are 10 choices for the first digit of our PIN, since repetition is okay there are still 10 choices for our second digit, then (still) 10 choices for the third digit as well as the fourth digit.

In general, when selecting k things out of n possibilities, where order counts and repetition is allowed, there are n^k possible selections.

Ordered without repetition

Suppose that one wishes to come up with a password for a computer account. There are 52 letters (both upper and lower case) 10 numerals and 32 symbols and punctuation marks – for a total of 94 different characters that may be used. Some system administrators can be very paranoid about passwords that might be guessable – for instance no password that appears in a dictionary should ever be used on a system where security is a concern. Suppose that your system administrator will reject any password that has repeated symbols, and that passwords must have 8 characters. How many passwords are possible?

This is an instance of a counting problem where we are selecting 8 things out of a set of size 94 – clearly order is important and the system adminis-

³The phrase “PIN number” is redundant. The ‘N’ in PIN stands for “number.” Anyway, a PIN is a four digit (secret) number used to help ensure that automated banking (such as withdrawing your life's savings) is only done by an authorized individual.

trator's restriction means that we may not have repeats. The multiplication rule tells us that there are $94 \cdot 93 \cdot 92 \cdot 91 \cdot 90 \cdot 89 \cdot 88 \cdot 87 = 4488223369069440$ different passwords. And in the general case (selecting k things out of a set of size n , without repetition, and with order counting) there will be $n!/(n-k)!$ possibilities. This is the number we have denoted previously by $P(n, k)$.

Unordered without repetition

This is also a case that we've considered previously. If we are choosing k things out of n and order is unimportant and there can be no repetitions, then what we are describing is a k -subset of the n -set. There are $C(n, k) = \frac{n!}{k!(n-k)!}$ distinct subsets. Here, we'll give an example that doesn't sound like we're talking about counting subsets of a particular size. (Although we really are!)

How many different sequences of 6 strictly increasing numbers can we choose from $\{1, 2, 3, \dots, 20\}$?

Obviously, listing all such sequences would be an arduous task. We might start with $(1, 2, 3, 4, 5, 6)$ and try to proceed in some orderly fashion to $(15, 16, 17, 18, 19, 20)$, but unfortunately there are 38,760 such sequences so unless we enlist the aid of a computer we are unlikely to finish this job in a reasonable time. The number we've just given (38,760) is $C(20, 6)$ and so it would seem that we're claiming that this problem is really unordered selection without repetition of 6 things out of 20. Well, actually, some parts of this are clearly right – we are selecting 6 things from a set of size 20, and because our sequences are supposed to be *strictly* increasing there will be no repetitions – but, a strictly increasing sequence is clearly ordered and the formula we are using is for unordered collections.

By specifying a particular ordering (strictly increasing) on the sequences we are counting above, we are actually removing the importance of order. Put another way: if order really mattered, the symbols 1 through 6 could be put into 720 different orders – but we only want to count one of those possibilities. Put another other way: there is a one-to-one correspondence between a 6-subset of $\{1, 2, 3, \dots, 20\}$ and a strictly increasing sequence. Just

make sure the subset is written in increasing order!

Okay, at this point we have filled-in three out of the four cells in our table.

		Does order matter?	
		Yes	No
Are repeats okay?	No	$P(n, k) = \frac{n!}{(n-k)!}$	$C(n, k) = \frac{n!}{k!(n-k)!}$
	Yes	n^k	

What kinds of things are we counting in the lower right part of the table? Unordered selections of k things out of n possibilities where there may (or may not!) be repetitions. The game Yahtzee provides a nice example of this type of configuration. When we roll 5 dice, we do not do so one-at-a-time, rather, we roll them as a group – the dice are indistinguishable so there is no way to order our set of 5 outcomes. In fact, it would be quite reasonable to, after one's roll, arrange the die in (say) increasing order. We'll repeat a bit of advice that was given previously: if one is free to rearrange a configuration to suit one's needs, that is a clue that order is *not* important in the configurations under consideration. Finally, are repetitions allowed? The outcomes in Yahtzee are 5 numbers from the set $\{1, 2, 3, 4, 5, 6\}$, and while it is possible to have no repetitions, that is a pretty special outcome!

In general, the same number can appear on two, or several, or even *all 5* of the die⁴.

So, how many different outcomes are there when one rolls five dice? To answer this question it will be helpful to think about how we might express such an outcome. Since order is unimportant, we can choose to put the numbers that appear on the individual die in whatever order we like. We may as well place them in increasing order. There will be 5 numbers and each number is between 1 and 6. We can list the outcomes systematically by starting with an all-ones Yahtzee:

(1,1,1,1,1)	(1,1,1,1,2)	(1,1,1,1,3)	(1,1,1,1,4)	(1,1,1,1,5)	(1,1,1,1,6)
(1,1,1,2,2)	(1,1,1,2,3)	(1,1,1,2,4)	(1,1,1,2,5)	(1,1,1,2,6)	(1,1,1,3,3)
(1,1,1,3,4)	(1,1,1,3,5)	(1,1,1,3,6)	(1,1,1,4,4)	(1,1,1,4,5)	(1,1,1,4,6)
(1,1,1,5,5)	(1,1,1,5,6)	(1,1,1,6,6)	(1,1,2,2,2)	(1,1,2,2,3)	(1,1,2,2,4)
(1,1,2,2,5)	(1,1,2,2,6)	(1,1,2,3,3)	(1,1,2,3,4)	(1,1,2,3,5)	(1,1,2,3,6)
(1,1,2,4,4)	(1,1,2,4,5)	(1,1,2,4,6)	(1,1,2,5,5)	(1,1,2,5,6)	(1,1,2,6,6)
(1,1,3,3,3)	(1,1,3,3,4)	(1,1,3,3,5)	(1,1,3,3,6)	(1,1,3,4,4)	(1,1,3,4,5)
(1,1,3,4,6)	(1,1,3,5,5)	(1,1,3,5,6)	(1,1,3,6,6)	(1,1,4,4,4)	(1,1,4,4,5)
(1,1,4,4,6)	(1,1,4,5,5)	(1,1,4,5,6)	(1,1,4,6,6)	(1,1,5,5,5)	(1,1,5,5,6)
(1,1,5,6,6)	(1,1,6,6,6)	(1,2,2,2,2)	(1,2,2,2,3)	(1,2,2,2,4)	(1,2,2,2,5)
(1,2,2,2,6)	(1,2,2,3,3)	(1,2,2,3,4)	(1,2,2,3,5)	(1,2,2,3,6)	(1,2,2,4,4)
(1,2,2,4,5)	(1,2,2,4,6)	(1,2,2,5,5)	(1,2,2,5,6)	(1,2,2,6,6)	(1,2,3,3,3)
(1,2,3,3,4)	(1,2,3,3,5)	(1,2,3,3,6)	(1,2,3,4,4)	(1,2,3,4,5)	(1,2,3,4,6)
(1,2,3,5,5)	(1,2,3,5,6)	(1,2,3,6,6)	(1,2,4,4,4)	(1,2,4,4,5)	(1,2,4,4,6)
(1,2,4,5,5)	(1,2,4,5,6)	(1,2,4,6,6)	(1,2,5,5,5)	(1,2,5,5,6)	(1,2,5,6,6)
(1,2,6,6,6)	(1,3,3,3,3)	(1,3,3,3,4)	(1,3,3,3,5)	(1,3,3,3,6)	(1,3,3,4,4)
(1,3,3,4,5)	(1,3,3,4,6)	(1,3,3,5,5)	(1,3,3,5,6)	(1,3,3,6,6)	(1,3,4,4,4)
(1,3,4,4,5)	(1,3,4,4,6)	(1,3,4,5,5)	(1,3,4,5,6)	(1,3,4,6,6)	(1,3,5,5,5)
(1,3,5,5,6)	(1,3,5,6,6)	(1,3,6,6,6)	(1,4,4,4,4)	(1,4,4,4,5)	(1,4,4,4,6)

⁴When this happens you are supposed to jump in the air and yell “Yahtzee!”

(1,4,4,5,5)	(1,4,4,5,6)	(1,4,4,6,6)	(1,4,5,5,5)	(1,4,5,5,6)	(1,4,5,6,6)
(1,4,6,6,6)	(1,5,5,5,5)	(1,5,5,5,6)	(1,5,5,6,6)	(1,5,6,6,6)	(1,6,6,6,6)
(2,2,2,2,2)	(2,2,2,2,3)	(2,2,2,2,4)	(2,2,2,2,5)	(2,2,2,2,6)	(2,2,2,3,3)
(2,2,2,3,4)	(2,2,2,3,5)	(2,2,2,3,6)	(2,2,2,4,4)	(2,2,2,4,5)	(2,2,2,4,6)
(2,2,2,5,5)	(2,2,2,5,6)	(2,2,2,6,6)	(2,2,3,3,3)	(2,2,3,3,4)	(2,2,3,3,5)
(2,2,3,3,6)	(2,2,3,4,4)	(2,2,3,4,5)	(2,2,3,4,6)	(2,2,3,5,5)	(2,2,3,5,6)
(2,2,3,6,6)	(2,2,4,4,4)	(2,2,4,4,5)	(2,2,4,4,6)	(2,2,4,5,5)	(2,2,4,5,6)
(2,2,4,6,6)	(2,2,5,5,5)	(2,2,5,5,6)	(2,2,5,6,6)	(2,2,6,6,6)	(2,3,3,3,3)
(2,3,3,3,4)	(2,3,3,3,5)	(2,3,3,3,6)	(2,3,3,4,4)	(2,3,3,4,5)	(2,3,3,4,6)
(2,3,3,5,5)	(2,3,3,5,6)	(2,3,3,6,6)	(2,3,4,4,4)	(2,3,4,4,5)	(2,3,4,4,6)
(2,3,4,5,5)	(2,3,4,5,6)	(2,3,4,6,6)	(2,3,5,5,5)	(2,3,5,5,6)	(2,3,5,6,6)
(2,3,6,6,6)	(2,4,4,4,4)	(2,4,4,4,5)	(2,4,4,4,6)	(2,4,4,5,5)	(2,4,4,5,6)
(2,4,4,6,6)	(2,4,5,5,5)	(2,4,5,5,6)	(2,4,5,6,6)	(2,4,6,6,6)	(2,5,5,5,5)
(2,5,5,5,6)	(2,5,5,6,6)	(2,5,6,6,6)	(2,6,6,6,6)	(3,3,3,3,3)	(3,3,3,3,4)
(3,3,3,3,5)	(3,3,3,3,6)	(3,3,3,4,4)	(3,3,3,4,5)	(3,3,3,4,6)	(3,3,3,5,5)
(3,3,3,5,6)	(3,3,3,6,6)	(3,3,4,4,4)	(3,3,4,4,5)	(3,3,4,4,6)	(3,3,4,5,5)
(3,3,4,5,6)	(3,3,4,6,6)	(3,3,5,5,5)	(3,3,5,5,6)	(3,3,5,6,6)	(3,3,6,6,6)
(3,4,4,4,4)	(3,4,4,4,5)	(3,4,4,4,6)	(3,4,4,5,5)	(3,4,4,5,6)	(3,4,4,6,6)
(3,4,5,5,5)	(3,4,5,5,6)	(3,4,5,6,6)	(3,4,6,6,6)	(3,5,5,5,5)	(3,5,5,5,6)
(3,5,5,6,6)	(3,5,6,6,6)	(3,6,6,6,6)	(4,4,4,4,4)	(4,4,4,4,5)	(4,4,4,4,6)
(4,4,4,5,5)	(4,4,4,5,6)	(4,4,4,6,6)	(4,4,5,5,5)	(4,4,5,5,6)	(4,4,5,6,6)
(4,4,6,6,6)	(4,5,5,5,5)	(4,5,5,5,6)	(4,5,5,6,6)	(4,5,6,6,6)	(4,6,6,6,6)
(5,5,5,5,5)	(5,5,5,5,6)	(5,5,5,6,6)	(5,5,6,6,6)	(5,6,6,6,6)	(6,6,6,6,6)

Whew . . . err, I mean, Yahtzee!

You can describe a generic element of the above listing by saying “It starts with some number of 1’s (which may be zero), then there are some 2’s (again, it might be that there are zero 2’s), then some (possibly none) 3’s, then some 4’s (or maybe not), then some 5’s (I think you probably get the idea) and finally some 6’s (sorry for all the parenthetical remarks).”

We could, of course, actually count the outcomes as listed above (there are 252) but that would be pretty dull – and it wouldn't get us any closer to solving such problems in general. To count things like Yahtzee rolls it will turn out that we can count something related but much simpler – blank-comma arrangements. For the Yahtzee problem we count arrangements of 5 blanks and 5 commas. That is, things like $_ _ , _ , _ , _ , _$, and $_ _ _ _ _ , , , ,$, and $, , , _ _ _ _ _ , ,$. These arrangements of blanks and commas correspond uniquely to Yahtzee rolls – the commas serve to separate different numerical values and the blanks are where we would write-in the 5 outcomes on the die.

Convince yourself that there really is a one-to-one correspondence between Yahtzee outcomes and arrangements of 5 blanks and 5 commas by doing the following

Exercise. *What Yahtzee rolls correspond to the following blank-comma arrangements?*

$_ , _ , _ , _ , _ , _$ $_ _ , _ _ _ , , , ,$ $, , , , , _ _ _ _ _$

What blank-comma arrangements correspond to the following Yahtzee outcomes?

$\{2, 3, 4, 5, 6\}$ $\{3, 3, 3, 3, 4\}$ $\{5, 5, 6, 6, 6\}$

It may seem at first that this blank-comma thing is okay, but that we're still no closer to answering the question we started with. It may seem that way until you realize how easy it is to count these blank-comma arrangements! You see, there are 10 symbols in one of these blank-comma arrangements and if we choose positions for (say) the commas, the blanks will have to go into the other positions – thus every 5-subset of $\{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}$ gives us a blank-comma arrangement and every one of *them* gives us a Yahtzee outcome. That is why there are $C(10, 5) = 252$ outcomes listed in the giant tabulation above.

In general, when we are selecting k things from a set of size n (with repetition and without order) we will need to consider blank-comma arrangements having k blanks and $n - 1$ commas. As an aid to memory, consider that when you actually write-out the elements of a set it takes one fewer commas than there are elements – for example $\{1, 2, 3, 4\}$ has 4 elements but we only need 3 commas to separate them. The general answer to our problem is either $C(k + n - 1, k)$ or $C(k + n - 1, n - 1)$, depending on whether you want to think about selecting positions for the k blanks or for the $n - 1$ commas. It turns out that these binomial coefficients are equal so there’s no problem with the apparent ambiguity.

So, finally, our table of counting formulas is complete. We’ll produce it here one more time and, while we’re at it, ditch the $C(n, k)$ notation in favor of the more usual “binomial coefficient” notation $\binom{n}{k}$.

		Does order matter?	
		Yes	No
Are repeats okay?	No	$P(n, k) = \frac{n!}{(n-k)!}$	$\binom{n}{k} = \frac{n!}{k!(n-k)!}$
	Yes	n^k	$\binom{n+k-1}{k}$

Exercises — 7.1

1. Determine the number of entries in the following sequences.
 - (a) (999, 1000, 1001, ... 2006)
 - (b) (13, 15, 17, ... 199)
 - (c) (13, 19, 25, ... 601)
 - (d) (5, 10, 17, 26, 37, ... 122)
 - (e) (27, 64, 125, 216, ... 8000)
 - (f) (7, 11, 19, 35, 67, ... 131075)
2. How many “full houses” are there in Yahtzee? (A full house is a pair together with a three-of-a-kind.)
3. In how many ways can you get “two pairs” in Yahtzee?
4. Prove that the binomial coefficients $\binom{n+k-1}{k}$ and $\binom{n+k-1}{n-1}$ are equal.
5. The “Cryptographer’s alphabet” is used to supply small examples in coding and cryptography. It consists of the first 6 letters, $\{a, b, c, d, e, f\}$. How many “words” of length up to 6 can be made with this alphabet? (A word need not actually be a word in English, for example both “fed” and “dfe” would be words in the sense we are using the term.)
6. How many “words” are there of length 4, with distinct letters from the Cryptographer’s alphabet, in which the letters appear in increasing order alphabetically? (“Acef” would be one such word, but “cafe” would not.)
7. How many “words” are there of length 4 from the Cryptographer’s alphabet, with repeated letters allowed, in which the letters appear in non-decreasing order alphabetically?

8. How many subsets does a finite set have?
9. How many handshakes will transpire when n people first meet?
10. How many functions are there from a set of size n to a set of size m ?
11. How many relations are there from a set of size n to a set of size m ?

7.2 Parity and Counting arguments

This section is concerned with two very powerful elements of the proof-making arsenal: “Parity” is a way of referring to the result of an even/odd calculation; Counting arguments most often take the form of counting some collection in two different ways – and then comparing those results. These techniques have little to do with one another, but when they are applicable they tend to produce really elegant little arguments.

In (very) early computers and business machines, paper cards were used to store information. A so-called “punch card” or “Hollerith card” was used to store binary information by means of holes punched into it. Paper tape was also used in a similar fashion. A typical paper tape format would involve 8 positions in rows across the tape that might or might not be punched, often a column of smaller holes would appear as well which did not store information but were used to drive the tape through the reading mechanism on a sprocket. Tapes and cards could be “read” either by small sets of electrical contacts which would touch through a punched hole or be kept separate if the position wasn’t punched, or by using a photo-detector to sense whether light could pass through the hole or not. The mechanisms for reading and writing on these paper media were amazingly accurate, and allowed early data processing machines to use just a couple of large file cabinets to store what now fits in a jump drive one can wear on a necklace. (About 10 or 12 cabinets could hold a gigabyte of data).

Paper media was ideally suited to storing binary information, but of course most of the real data people needed to store and process would be alphanumeric⁵. There were several encoding schemes that served to translate between the character sets that people commonly used and the binary numerals that could be stored on paper. One of these schemes still survives

⁵“Alphanumeric” is a somewhat antiquated term that refers to information containing both alphabetic characters and numeric characters – as well as punctuation marks, etc.

today – ASCII. The American Standard Code for Information Interchange uses 7-bit binary numerals to represent characters, so it contains 128 different symbols. This is more than enough to represent both upper- and lower-case letters, the 10 numerals, and the punctuation marks – many of the remaining spots in the ASCII code were used to contain so-called “control characters” that were associated with functionality that appeared on old-fashioned tele-type equipment – things like “ring the bell,” “move the carriage backwards one space,” “move the carriage to the next line,” etc. These control characters are why modern keyboards still have a modifier key labeled “Ctrl” on them. The following listing gives the decimal and binary numerals from 0 to 127 and the ASCII characters associated with them – the non-printing and control characters have a 2 or 3 letter mnemonic designation.

0	0000 0000	NUL	64	0100 0000	@
1	0000 0001	SOH	65	0100 0001	A
2	0000 0010	STX	66	0100 0010	B
3	0000 0011	ETX	67	0100 0011	C
4	0000 0100	EOT	68	0100 0100	D
5	0000 0101	ENQ	69	0100 0101	E
6	0000 0110	ACK	70	0100 0110	F
7	0000 0111	BEL	71	0100 0111	G
8	0000 1000	BS	72	0100 1000	H
9	0000 1001	TAB	73	0100 1001	I
10	0000 1010	LF	74	0100 1010	J
11	0000 1011	VT	75	0100 1011	K
12	0000 1100	FF	76	0100 1100	L
13	0000 1101	CR	77	0100 1101	M
14	0000 1110	SO	78	0100 1110	N
15	0000 1111	SI	79	0100 1111	O
16	0001 0000	DLE	80	0101 0000	P
17	0001 0001	DC1	81	0101 0001	Q
18	0001 0010	DC2	82	0101 0010	R
19	0001 0011	DC3	83	0101 0011	S
20	0001 0100	DC4	84	0101 0100	T
21	0001 0101	NAK	85	0101 0101	U
22	0001 0110	SYN	86	0101 0110	V
23	0001 0111	ETB	87	0101 0111	W
24	0001 1000	CAN	88	0101 1000	X
25	0001 1001	EM	89	0101 1001	Y

26	0001 1010	SUB	90	0101 1010	Z
27	0001 1011	ESC	91	0101 1011	[
28	0001 1100	FS	92	0101 1100	\
29	0001 1101	GS	93	0101 1101]
30	0001 1110	RS	94	0101 1110	^
31	0001 1111	US	95	0101 1111	_
32	0010 0000		96	0110 0000	`
33	0010 0001	!	97	0110 0001	a
34	0010 0010	"	98	0110 0010	b
35	0010 0011	#	99	0110 0011	c
36	0010 0100	\$	100	0110 0100	d
37	0010 0101	%	101	0110 0101	e
38	0010 0110	&	102	0110 0110	f
39	0010 0111	'	103	0110 0111	g
40	0010 1000	(104	0110 1000	h
41	0010 1001)	105	0110 1001	i
42	0010 1010	*	106	0110 1010	j
43	0010 1011	+	107	0110 1011	k
44	0010 1100	,	108	0110 1100	l
45	0010 1101	-	109	0110 1101	m
46	0010 1110	.	110	0110 1110	n
47	0010 1111	/	111	0110 1111	o
48	0011 0000	0	112	0111 0000	p
49	0011 0001	1	113	0111 0001	q
50	0011 0010	2	114	0111 0010	r
51	0011 0011	3	115	0111 0011	s
52	0011 0100	4	116	0111 0100	t
53	0011 0101	5	117	0111 0101	u
54	0011 0110	6	118	0111 0110	v
55	0011 0111	7	119	0111 0111	w
56	0011 1000	8	120	0111 1000	x
57	0011 1001	9	121	0111 1001	y
58	0011 1010	:	122	0111 1010	z
59	0011 1011	;	123	0111 1011	{
60	0011 1100	<	124	0111 1100	
61	0011 1101	=	125	0111 1101	}
62	0011 1110	>	126	0111 1110	~
63	0011 1111	?	127	0111 1111	DEL

Now it only takes 7 bits to encode the 128 possible values in the ASCII system, which can easily be verified by noticing that the left-most bits in all

of the binary representations above are 0. Most computers use 8 bit words or “bytes” as their basic units of information, and the fact that the ASCII code only requires 7 bits lead someone to think up a use for that additional bit. It became a “parity check bit.” If the seven bits of an ASCII encoding have an odd number of 1’s, the parity check bit is set to 1 — otherwise, it is set to 0. The result of this is that, subsequently, all of the 8 bit words that encode ASCII data will have an even number of 1’s. This is an example of a so-called error detecting code known as the “even code” or the “parity check code.” If data is sent over a noisy telecommunications channel, or is stored in fallible computer memory, there is some small but calculable probability that there will be a “bit error.” For instance, one computer might send 10000111 (which is the ASCII code that says “ring the bell”) but another machine across the network might receive 10100111 (the 3rd bit from the left has been received in error) now if we are only looking at the rightmost seven bits we will think that the ASCII code for a single quote has been received, but if we note that this piece of received data has an odd number of ones we’ll realize that something is amiss. There are other more advanced coding schemes that will let us not only *detect* an error, but (within limits) *correct* it as well! This rather amazing feat is what makes wireless telephony (not to mention communications with deep space probes — whoops! I mentioned it) work.

The concept of parity can be used in many settings to prove some fairly remarkable results.

In Section 6.3 we introduced the idea of a graph. This notion was first used by Leonhard Euler to solve a recreational math problem posed by the citizens of Königsberg, Prussia (this is the city now known as Kaliningrad, Russia.) Königsberg was situated at a place where two branches of the Pregel river⁶ come together – there is also a large island situated near this confluence. By Euler’s time, the city of Königsberg covered this island as well

⁶Today, this river is known as the Pregolya.

as the north and south banks of the river and also the promontory where the branches came together. A network of seven bridges had been constructed to connect all these land masses. The townsfolk are alleged to have become enthralled by the question of whether it was possible to leave one's home and take a walk through town which crossed each of the bridges exactly once and, finally, return to one's home.

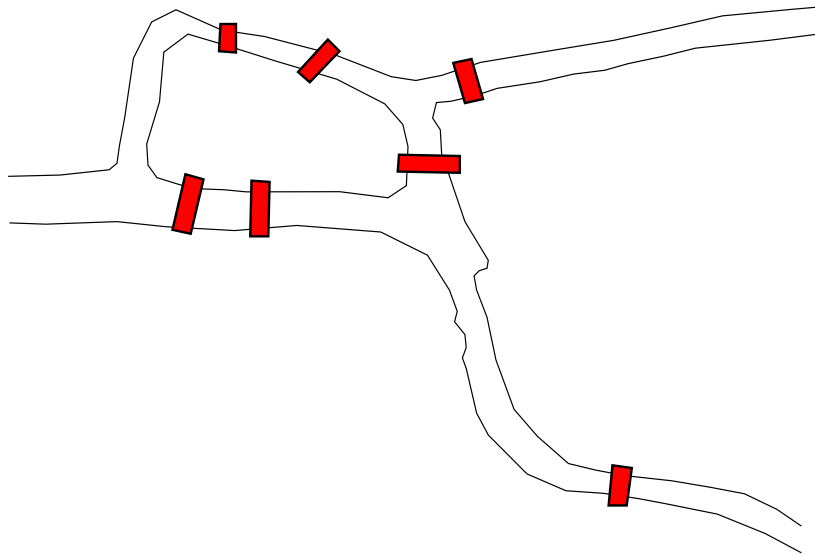


Figure 7.2: A simplified map of Königsberg, Prussia circa 1736.

Euler settled the question (it can't be done) by converting the map of Königsberg into a graph and then making some simple observations about the parities of the nodes in this graph. The degree of a node in a graph is the number of edges that are incident with it (if a so-called "loop edge" is present it adds two to the node's degree). The "parity of a node" is shorthand for the "parity of the *degree* of the node."

The graph of Königsberg has 4 nodes: one of degree 5 and three of degree 3. All the nodes are odd. Would it be possible to either modify Königsberg or come up with an entirely new graph having some even nodes? Well, the answer to that is easy – just tear down one of the bridges, and two of

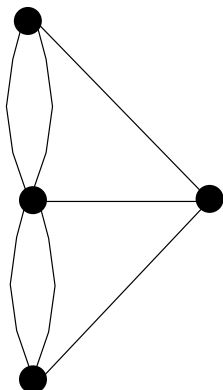


Figure 7.3: Euler’s solution of the “seven bridges of Königsberg problem” involved representing the town as an undirected graph.

the nodes will have their degree changed by one; they’ll both become even. Notice that, by removing one edge, we effected the parity of two nodes. Is it possible to create a graph with four nodes in which just one of them is even? More generally, given any short list of natural numbers, is it possible to draw a graph whose degrees are the listed numbers?

Exercise. *Try drawing graphs having the following lists of vertex degrees. (In some cases it will be impossible...)*

- $\{1, 1, 2, 3, 3\}$
- $\{1, 2, 3, 5\}$
- $\{1, 2, 3, 4\}$
- $\{4, 4, 4, 4, 5\}$
- $\{3, 3, 3, 3\}$
- $\{3, 3, 3, 3, 3\}$

When it is possible to create a graph with a specified list of vertex degrees, it is usually easy to do. Of course, when it's impossible you struggle a bit. . . To help get things rolling (just in case you haven't really done the exercise) I'll give a hint – for the first list it is possible to draw a graph, for the second it is not. Can you distinguish the pattern? What makes one list of vertex degrees reasonable and another not?

Exercise. (*If you didn't do the last exercise, stop being such a lame-o and try it now. BTW, if you did do it, good for you! You can either join with me now in sneering at all those people who are scurrying back to do the last one, or try the following:*)

Figure out a way to distinguish a sequence of numbers that can be the degree sequence of some graph from the sequences that cannot be.

Okay, now if you're reading this sentence you should know that every other list of vertex degrees above is impossible, you should have graphs drawn in the margin here for the 1st, 3rd and 5th degree sequences, and you may have discovered some version of the following

Theorem 7.2.1. *In an undirected graph, the number of vertices having an odd degree is even.*

A slightly pithier statement is: All graphs have an even number of odd nodes.

We'll leave the proof of this theorem to the exercises but most of the work is done in proving the following equivalent result.

Theorem 7.2.2. *In an undirected graph the sum of the degrees of the vertices is even.*

Proof: The sum of the degrees of all the vertices in a graph G ,

$$\sum_{v \in V(G)} \deg(v),$$

counts every edge of G exactly twice.

Thus,

$$\sum_{v \in V(G)} \deg(v) = 2 \cdot |E(G)|.$$

In particular we see that this sum is even.

Q.E.D.

The question of whether a graph having a given list of vertex degrees can exist comes down to an elegant little argument using both of the techniques in the title of this section. We count the edge set of the graph in two ways – once in the usual fashion and once by summing the vertex degrees; we also note that since this latter count is actually a double count we can bring in the concept of parity.

Another perfectly lovely argument involving parity arises in questions concerning whether or not it is possible to tile a pruned chessboard with dominoes. We’ve seen dominoes before in Section 5.1 and we’re just hoping you’ve run across chessboards before. Usually a chessboard is 8×8 , but we would like to adopt a more liberal interpretation that a chessboard can be any rectangular grid of squares we might choose.⁷ Suppose that we have a supply of dominoes that are of just the right size that if they are laid on a chessboard they perfectly cover two adjacent squares. Our first question is quite simple. Is it possible to perfectly tile an $m \times n$ chessboard with such dominoes?

First let’s specify the question a bit more. By “perfectly tiling” a chessboard we mean that every domino lies (fully) on the board, covering precisely

⁷The game known as “draughts” in the UK and “checkers” in the US is played on an 8×8 board, but (for example) international draughts is played on a 10×10 board and Canadian checkers is played on a 12×12 board.

two squares, and that every square of the board is covered by a domino.

The answer is straightforward. If at least one of m or n is even it can be done. A necessary condition is that the number of squares be even (since every domino covers two squares) and so, if both m and n are odd we will be out of luck.



A “pruned board” is obtained by either literally removing some of the squares or perhaps by marking them as being off limits in some way. When we ask questions about perfect tilings of pruned chessboards things get more interesting and the notion of parity can be used in several ways.

Here are two tiling problems regarding square chessboards:

1. An even-sided square board (e.g. an ordinary 8×8 board) with diagonally opposite corners pruned.
2. An odd-sided board with one square pruned.

Both of these situations satisfy the basic necessary condition that the number of squares on the board must be even. You may be able to determine another “parity” approach to these tiling problems by attempting the following

Exercise. *Below are two five-by-five chessboards each having a single square pruned. One can be tiled by dominoes and the other cannot. Which is which?*



The pattern of black and white squares on a chessboard is an example of a sort of artificial parity, if we number the squares of the board appropriately then the odd squares will be white and the even squares will be black. We are used to chessboards having this alternating black/white pattern on them, but nothing about these tiling problems required that structure⁸ If we were used to monochromatic chessboards, we might never solve the previous two problems – unless of course we invented the coloring scheme in order to solve them. An odd-by-odd chessboard has more squares of one color than of the other. An odd-by-odd chessboard needs to have a square pruned in order for it to be possible for it to be tiled by dominoes – but if the wrong colored square is pruned it will *still* be impossible. Each domino covers two squares – one of each color! (So the pruned board must have the same number of white squares as black.)

We'll close this section with another example of the technique of counting in two ways.

A magic square of order n is a square $n \times n$ array containing the numbers $1, 2, 3, \dots, n^2$. The numbers must be arranged in such a way that every row and every column sum to the same number – this value is known as the magic sum.

For example, the following is an order 3 magic square.

1	6	8
5	7	3
9	2	4

The definition of a magic square requires that the rows and columns sum to the same number but says nothing about what that number must be. It is conceivable that we could produce magic squares (of the same order)

⁸Nothing about chess requires this structure either, but it does let us do some error checking. For instance, bishops always end up on the same color they left from and knights always switch colors as they move.

having different magic sums. This is *conceivable*, but in fact the magic sum is determined completely by n .

Theorem 7.2.3. *A magic square of order n has a magic sum equal to $\frac{n^3 + n}{2}$.*

Proof: We count the total of the entries in the magic square in two ways. The sum of all the entries in the magic square is

$$S = 1 + 2 + 3 + \dots + n^2.$$

Using the formula for the sum of the first k naturals ($\sum_{i=1}^k i = \frac{k^2+k}{2}$) and evaluating at n^2 gives

$$S = \frac{n^4 + n^2}{2}.$$

On the other hand, if the magic sum is M , then each of the n rows has numbers in it which sum to M so

$$S = nM.$$

By equating these different expressions for S and solving for M , we prove the desired result:

$$nM = \frac{n^4 + n^2}{2},$$

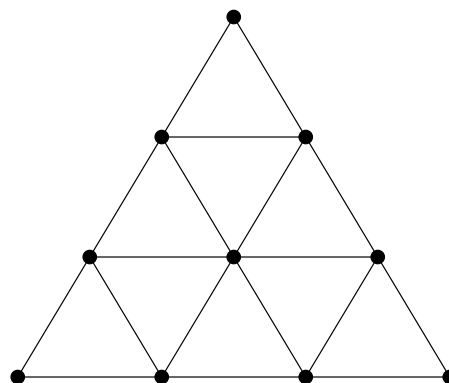
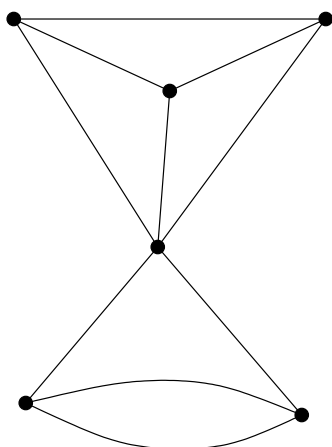
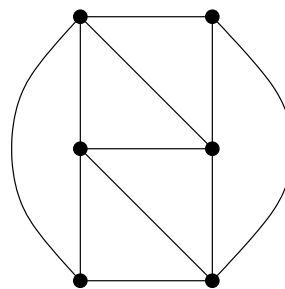
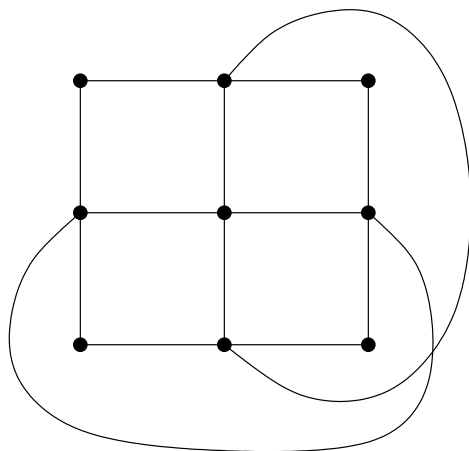
therefore

$$M = \frac{n^3 + n}{2}.$$

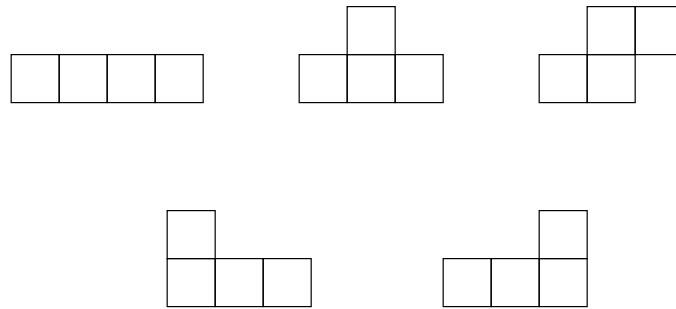
Q.E.D.

Exercises — 7.2

1. A walking tour of Königsberg such as is described in this section, or more generally, a circuit through an arbitrary graph that crosses each edge precisely once and begins and ends at the same node is known as an *Eulerian circuit*. An *Eulerian path* also crosses every edge of a graph exactly once but it begins and ends at distinct nodes. For each of the following graphs determine whether an Eulerian circuit or path is possible, and if so, draw it.



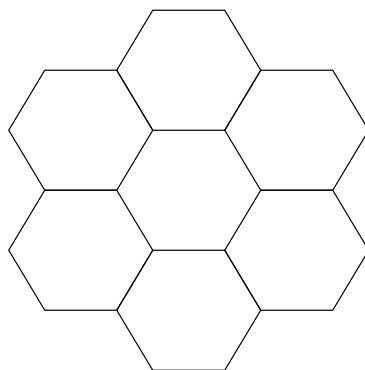
2. Complete the proof of the fact that “Every graph has an even number of odd nodes.”
3. Provide an argument as to why an 8×8 chessboard with two squares pruned from diagonally opposite corners cannot be tiled with dominoes.
4. Prove that, if n is odd, any $n \times n$ chessboard with a square the same color as one of its corners pruned can be tiled by dominoes.
5. The five tetrominoes (familiar to players of the video game Tetris) are relatives of dominoes made up of four small squares.



All together these five tetrominoes contain 20 squares so it is conceivable that they could be used to tile a 4×5 chessboard. Prove that this is actually impossible.

6. State necessary and sufficient conditions for the existence of an Eulerian circuit in a graph.
7. State necessary and sufficient conditions for the existence of an Eulerian path in a graph.

8. Construct magic squares of order 4 and 5.
9. A magic hexagon of order 2 would consist of filling-in the numbers from 1 to 7 in the hexagonal array below. The magic condition means that each of the 9 “lines” of adjacent hexagons would have the same sum. Is this possible?



10. Is there a magic hexagon of order 3?

7.3 The pigeonhole principle

The word “pigeonhole” can refer to a hole in which a pigeon roosts (i.e. pretty much what it sounds like) or a series of roughly square recesses in a desk in which one could sort correspondence (see Figure 7.4).

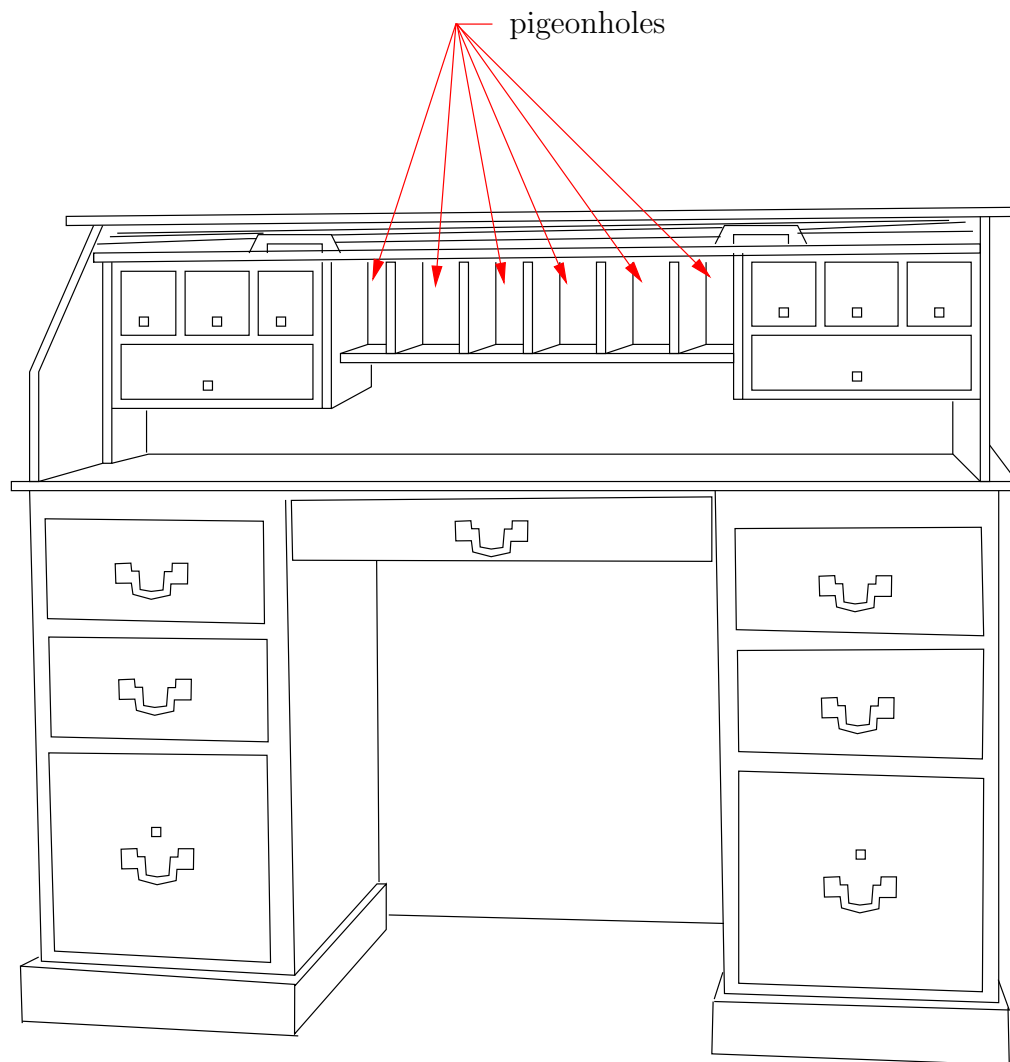


Figure 7.4: Pigeonholes in an old-fashioned roll top desk could be used to sort letters.

Whether you prefer to think of roosting birds or letters being sorted, the first and easiest version of the pigeonhole principle is that if you have more “things” than you have “containers” there must be a container holding at least two things.

If we have 6 pigeons who are trying to roost in a coop with 5 pigeonholes, two birds will have to share.

If we have 7 letters to sort and there are 6 pigeonholes in our desk, we will have to put two letters in the same compartment.

The “things” and the “containers” don’t necessarily have to be interpreted in the strict sense that the “things” go *into* the “containers.” For instance, a nice application of the pigeonhole principle is that if there are at least 13 people present in a room, some pair of people will have been born in the same month. In this example the things are the people and the containers are the months of the year.

The abstract way to phrase the pigeonhole principle is:

Theorem 7.3.1. *If f is a function such that $|\text{Dom}(f)| > |\text{Rng}(f)|$ then f is not injective.*

The proof of this statement is an easy example of proof by contradiction so we’ll include it here.

Proof: Suppose to the contrary that f is a function with $|\text{Dom}(f)| > |\text{Rng}(f)|$ and that f is injective. Of course f is onto its range, so since we are presuming that f is injective it follows that f is a bijection between $\text{Dom}(f)$ and $\text{Rng}(f)$. Therefore (since f provides a one-to-one correspondence) $|\text{Dom}(f)| = |\text{Rng}(f)|$. This clearly contradicts the statement that $|\text{Dom}(f)| > |\text{Rng}(f)|$.

Q.E.D.

For a statement with an almost trivial proof the pigeonhole principle is very powerful. We can use it to prove a host of existential results – some are fairly silly, others very deep. Here are a few examples:

There are two people (who are not bald) in New York City having exactly the same number of hairs on their heads.

There are two books in (insert your favorite library) that have the same number of pages.

Given n married couples (so $2n$ people) if we choose $n + 1$ people we will be forced to choose both members of some couple.

Suppose we select $n + 1$ numbers from the set $\{1, 2, 3, \dots, 2n\}$, we will be forced to have chosen two numbers such that one is divisible by the other.

We can come up with stronger forms of the pigeonhole principle by considering pigeonholes with capacities. Suppose we have 6 pigeonholes in a desk, each of which can hold 10 letters. What number of letters will guarantee that one of the pigeonholes is full? The largest number of letters we could have without having 10 in some pigeonhole is $9 \cdot 6 = 54$, so if there are 55 letters we must have 10 letters in some pigeonhole.

More generally, if we have n containers, each capable of holding m objects, than if there are $n \cdot (m - 1) + 1$ objects placed in the containers, we will be assured that one of the containers is at capacity.

The ordinary pigeonhole principle is the special case $m = 2$ of this stronger version.

There is an even stronger version, which ordinarily is known as the “strong form of the pigeonhole principle.” In the strong form, we have pigeonholes with an assortment of capacities.

Theorem 7.3.2. *If there are n containers having capacities $m_1, m_2, m_3, \dots, m_n$, and there are $1 + \sum_{i=1}^n (m_i - 1)$ objects placed in them, then for some i , container i has (at least) m_i objects in it.*

Proof: If no container holds its full capacity, then the largest

the total of the objects could be is $\sum_{i=1}^n (m_i - 1)$.

Q.E.D.

Exercises — 7.3

1. The statement that there are two non-bald New Yorkers with the same number of hairs on their heads requires some careful estimates to justify it. Please justify it.
2. A mathematician, who always rises earlier than her spouse, has developed a scheme – using the pigeonhole principle – to ensure that she always has a matching pair of socks. She keeps only blue socks, green socks and black socks in her sock drawer – 10 of each. So as not to wake her husband she must select some number of socks from her drawer in the early morning dark and take them with her to the adjacent bathroom where she dresses. What number of socks does she choose?
3. If we select 1001 numbers from the set $\{1, 2, 3, \dots, 2000\}$ it is certain that there will be two numbers selected such that one divides the other. We can prove this fact by noting that every number in the given set can be expressed in the form $2^k \cdot m$ where m is an odd number and using the pigeonhole principle. Write-up this proof.
4. Given any set of 53 integers, show that there are two of them having the property that either their sum or their difference is evenly divisible by 103.
5. Prove that if 10 points are placed inside a square of side length 3, there will be 2 points within $\sqrt{2}$ of one another.
6. Prove that if 10 points are placed inside an equilateral triangle of side length 3, there will be 2 points within 1 of one another.
7. Prove that in a simple graph (an undirected graph with no loops or parallel edges) having n nodes, there must be two nodes having the same degree.

$$\begin{aligned}(x + y)^6 &= (x + y)^3 \cdot (x + y)^3 \\ &= (x^3 + 3x^2y + 3xy^2 + y^3) \cdot (x^3 + 3x^2y + 3xy^2 + y^3)\end{aligned}$$

For this product I wouldn't even *think* about the distributive law, I'd jump to the table method right away:

	x^3	$+3x^2y$	$+3xy^2$	$+y^3$
x^3				
$+3x^2y$				
$+3xy^2$				
$+y^3$				

In the end you should obtain

$$x^6 + 6x^5y + 15x^4y^2 + 20x^3y^3 + 15x^2y^4 + 6xy^5 + y^6.$$

Now all of this is a lot of work and it's really much easier to notice the form of the answer: The exponent on x starts at 6 and descends with each successive term down to 0. The exponent on y starts at 0 and ascends to 6. The coefficients in the answer are the numbers in the sixth row of Pascal's triangle.

Finally, the form of Pascal's triangle makes it really easy to extend. A number in the interior of the triangle is always the sum of the two above it (on either side). Numbers that aren't in the interior of the triangle are always 1.

We showed rows 0 through 6 above. Rows 7 and 8 are

$$\begin{array}{cccccccc} 1 & 7 & 21 & 35 & 35 & 21 & 7 & 1 \\ 1 & 8 & 28 & 56 & 70 & 56 & 28 & 8 & 1. \end{array}$$

With this information in hand, it becomes nothing more than a matter of copying down the answer to compute

$$(x+y)^8 = x^8 + 8x^7y + 28x^6y^2 + 56x^5y^3 + 70x^4y^4 + 56x^3y^5 + 28x^2y^6 + 8xy^7 + y^8.$$

Exercise. Given the method using Pascal's triangle for computing $(x+y)^n$ we can use substitution to determine more general binomial powers.

Find $(x^4 + x^2)^5$.

All of the above hinges on the fact that one can compute a binomial coefficient by summing the two that appear to either side and above it in Pascal's triangle. This fact is the fundamental relationship between binomial coefficients – it is usually called Pascal's formula.

Theorem 7.4.1. For all natural numbers n and k with $0 < k \leq n$,

$$\binom{n}{k} = \binom{n-1}{k} + \binom{n-1}{k-1}.$$

We are going to prove it twice.

Proof: (The first proof is a combinatorial argument.)

There are $\binom{n}{k}$ subsets of size k of the set $N = \{1, 2, 3, \dots, n\}$. We will partition these k -subsets into two disjoint cases: those that contain the final number, n , and those that do not.

Let

$$A = \{S \subseteq N \mid |S| = k \wedge n \notin S\}$$

and, let

$$B = \{S \subseteq N \mid |S| = k \wedge n \in S\}.$$

Since the number n is either in a k -subset or it isn't, these sets are disjoint and exhaustive. So the addition rule tells us that

$$\binom{n}{k} = |A| + |B|.$$

The set A is really just the set of all k -subsets of the $(n-1)$ -set $\{1, 2, 3, \dots, n-1\}$, so $|A| = \binom{n-1}{k}$.

Any of the sets in B can be obtained by adjoining the element n to a $k-1$ subset of the $(n-1)$ -set $\{1, 2, 3, \dots, n-1\}$, so $|B| = \binom{n-1}{k-1}$.

Substituting gives us the desired result.

Q.E.D.

Proof: (The second proof is algebraic in nature.)

Consider the sum

$$\binom{n-1}{k} + \binom{n-1}{k-1}.$$

Applying the formula we deduced in Section 7.1 we get

$$\begin{aligned} & \binom{n-1}{k} + \binom{n-1}{k-1} \\ &= \frac{(n-1)!}{k!(n-1-k)!} + \frac{(n-1)!}{(k-1)!((n-1)-(k-1))!} \\ &= \frac{(n-1)!}{k!(n-k-1)!} + \frac{(n-1)!}{(k-1)!(n-k)!} \end{aligned}$$

A common denominator for these fractions is $k!(n-k)!$. (We will have to multiply the top and bottom of the first fraction by $(n-k)$ and the top and bottom of the second fraction by k .)

$$\begin{aligned}
 &= \frac{(n-k)(n-1)!}{k!(n-k)(n-k-1)!} + \frac{k(n-1)!}{k(k-1)!(n-k)!} \\
 &= \frac{(n-k)(n-1)!}{k!(n-k)!} + \frac{k(n-1)!}{k!(n-k)!} \\
 &= \frac{(n-k)(n-1)! + k(n-1)!}{k!(n-k)!} \\
 &= \frac{(n-k+k)(n-1)!}{k!(n-k)!} \\
 &= \frac{(n)(n-1)!}{k!(n-k)!} \\
 &= \frac{n!}{k!(n-k)!}.
 \end{aligned}$$

We recognize the final expression as the definition of $\binom{n}{k}$, so we have proved that

$$\binom{n-1}{k} + \binom{n-1}{k-1} = \binom{n}{k}.$$

Q.E.D.

There are quite a few other identities concerning binomial coefficients that can also be proved in (at least) two ways. We will provide one or two

other examples and leave the rest to you in the exercises for this section.

Theorem 7.4.2. For all natural numbers n and k with $0 < k \leq n$,

$$k \cdot \binom{n}{k} = n \cdot \binom{n-1}{k-1}.$$

Let's try a purely algebraic approach first.

Proof:

Using the formula for the value of a binomial coefficient we get

$$k \cdot \binom{n}{k} = k \cdot \frac{n!}{k!(n-k)!}.$$

We can do some cancellation to obtain

$$k \cdot \binom{n}{k} = \frac{n!}{(k-1)!(n-k)!}.$$

Finally we factor-out an n to obtain

$$k \cdot \binom{n}{k} = n \cdot \frac{(n-1)!}{(k-1)!(n-k)!},$$

since $(n-k)$ is the same thing as $((n-1) - (k-1))$ we have

$$k \cdot \binom{n}{k} = n \cdot \frac{(n-1)!}{(k-1)!((n-1) - (k-1))!} = n \cdot \binom{n-1}{k-1}$$

Q.E.D.

A combinatorial argument usually involves counting *something* in two ways. What could that something be? Well, if you see a product in some formula you should try to imagine what the multiplication rule would say in that particular circumstance.

Proof: Consider the collection of all subsets of size k taken from $N = \{1, 2, 3, \dots, n\}$ in which one of the elements has been marked to distinguish it from the others in some way.⁹

We can count this collection in two ways using the multiplication rule.

Firstly, we could select a k -subset in $\binom{n}{k}$ ways and then from among the k elements of the subset we could select one to be marked. By this analysis there are $\binom{n}{k} \cdot k$ elements in our collection.

Secondly, we could select an element from the n -set which will be the “marked” element of our subset, and then choose the additional $k - 1$ elements from the remaining $n - 1$ elements of the n -set. By this analysis there are $n \cdot \binom{n-1}{k-1}$ elements in the collection we have been discussing.

Thus,

$$k \cdot \binom{n}{k} = n \cdot \binom{n-1}{k-1}$$

Q.E.D.

The final result that we’ll talk about actually has (at least) three proofs. One of which suffers from the fault that it is “like swatting a fly with a sledge hammer.”

The result concerns the sum of all the numbers in some row of Pascal’s triangle.

Theorem 7.4.3. *For all natural numbers n and k with $0 < k \leq n$,*

⁹ For example, a committee of k individuals one of whom has been chosen as chairperson, is an example of the kind of entity we are discussing.

$$\sum_{k=0}^n \binom{n}{k} = 2^n$$

Our sledge hammer is a powerful result known as the binomial theorem which is a formalized statement of the material we began this section with.

Theorem 7.4.4 (The Binomial Theorem). *For all natural numbers n , and real numbers x and y ,*

$$(x + y)^n = \sum_{k=0}^n \binom{n}{k} x^{n-k} y^k.$$

We won't be proving this result just now. But, the following proof is a proof of the previous theorem using this more powerful result.

Proof: Substitute $x = y = 1$ in the binomial theorem.

Q.E.D.

Our second proof will be combinatorial. Let us re-iterate that a combinatorial proof usually consists of counting some collection in two different ways. The formula we have in this example contains a sum, so we should search for a collection of things that can be counted using the addition rule.

Proof: The set of all subsets of $N = \{1, 2, 3, \dots, n\}$, which we denote by $\mathcal{P}(N)$, can be partitioned into $n + 1$ sets based on the sizes of the subsets,

$$\mathcal{P}(N) = S_0 \cup S_1 \cup S_2 \cup \dots \cup S_n,$$

where $S_k = \{S \mid S \subseteq N \wedge |S| = k\}$ for $0 \leq k \leq n$. Since no subset of N can appear in two different parts of the partition (a subset's size is unique) and every subset of N appears in one of the parts of the partition (the sizes of subsets are all in the range from 0 to n). The addition principle tells us that

$$|\mathcal{P}(N)| = |S_0| + |S_1| + \cup |S_2| + \dots + |S_n|.$$

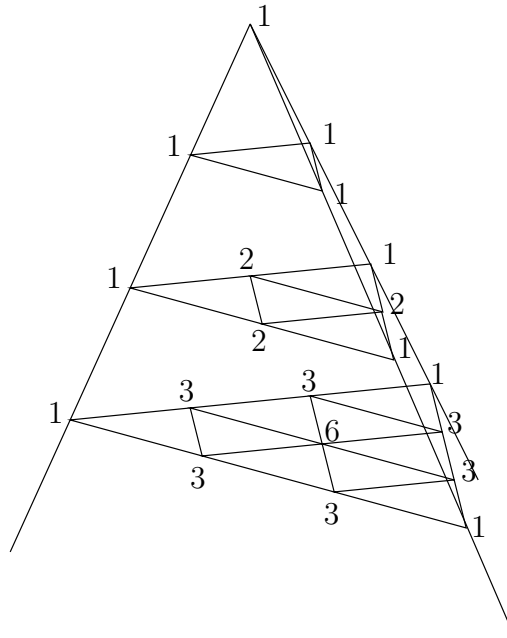
We have previously proved that $|\mathcal{P}(N)| = 2^n$ and we know that $|S_k| = \binom{n}{k}$ so it follows that

$$2^n = \binom{n}{0} + \binom{n}{1} + \binom{n}{2} + \dots + \binom{n}{n}.$$

Q.E.D.

Exercises — 7.4

1. Use the binomial theorem (with $x = 1000$ and $y = 1$) to calculate 1001^6 .
2. Find $(2x + 3)^5$.
3. Find $(x^2 + y^2)^6$.
4. The following diagram contains a 3-dimensional analog of Pascal's triangle that we might call "Pascal's tetrahedron." What would the next layer look like?



5. The student government at Lagrange High consists of 24 members chosen from amongst the general student body of 210. Additionally, there is a steering committee of 5 members chosen from amongst those in student government. Use the multiplication rule to determine two different formulas for the total number of possible governance structures.

6. Prove the identity

$$\binom{n}{k} \cdot \binom{k}{r} = \binom{n}{r} \cdot \binom{n-r}{k-r}$$

combinatorially.

7. Prove the binomial theorem.

$$\forall n \in \mathbb{Z}^{\text{nonneg}}, \forall x, y \in \mathbb{R}, (x + y)^n = \sum_{k=0}^n \binom{n}{k} x^{n-k} y^k$$

Chapter 8

Cardinality

The very existence of flame-throwers proves that some time, somewhere, someone said to themselves, “You know, I want to set those people over there on fire, but I’m just not close enough to get the job done.” –George Carlin

8.1 Equivalent sets

We have seen several interesting examples of equivalence relations already, and in this section we will explore one more: we’ll say two sets are equivalent if they have the same number of elements. Usually, an equivalence relation has the effect that it highlights one characteristic of the objects being studied, while ignoring all the others. Equivalence of sets brings the issue of size (a.k.a. cardinality) into sharp focus while, at the same time, it forgets all about the many other features of sets. Sets that are equivalent (under the relation we are discussing) are sometimes said to be *equinumerous*¹.

A couple of examples may be in order.

- If $A = \{1, 2, 3\}$ and $B = \{a, b, c\}$ then A and B are equivalent.

¹Perversely, there are also those who use the term *equipollent* to indicate that sets are the same size. This term actually applies to logical statements that are deducible from one another.

- Since the empty set is unique – \emptyset is the only set having 0 elements – it follows that there are no other sets equivalent to it.
- Every singleton set² is equivalent to every other singleton set.

Hopefully these examples are relatively self-evident. Unfortunately, that very self-evidence may tend to make you think that this notion of equivalence isn't all that interesting — nothing could be further from the truth! The notion of equivalence of sets becomes really interesting when we study infinite sets. Once we have the right definition in hand we will be able to prove some truly amazing results. For instance, the sets $\mathbb{Z}^{\text{nonneg}}$ and \mathbb{Q} turn out to be equivalent. Since the naturals are wholly contained in the rationals this is (to say the least) counter-intuitive! Coming up with the “right” definition for this concept is crucial.

We could make the following:

Definition. (*Well . . . not quite.*) For all sets A and B , we say A and B are equivalent, and write $A \equiv B$ iff $|A| = |B|$.

The problem with this definition is that it is circular. We're trying to come up with an equivalence relation so that the equivalence classes will represent the various cardinalities of sets (i.e. their sizes) and we define the relation in terms of cardinalities. We won't get anything new from this.

Georg Cantor was the first person to develop the modern notion of the equivalence of sets. His early work used the notion implicitly, but when he finally developed the concept of one-to-one correspondences in an explicit way he was able to prove some amazing facts. The phrase “one-to-one correspondence” has a fairly impressive ring to it, but one can discover what it means by just thinking carefully about what it means to count something.

Consider the solmization syllables used for the notes of the major scale in music; they form the set {do, re, mi, fa, so, la, ti}. What are we doing when

²Recall that a singleton set is a set having just one element.

we count this set (and presumably come up with a total of 7 notes)? We first point at ‘do’ while saying ‘one,’ then point at ‘re’ while saying ‘two,’ et cetera. In a technical sense we are creating a one-to-one correspondence between the set containing the seven syllables and the special set $\{1, 2, 3, 4, 5, 6, 7\}$. You should notice that this one-to-one correspondence is by no means unique. For instance we could have counted the syllables in reverse — a descending scale, or in some funny order — a little melody using each note once. The fact that there are seven syllables in the solmization of the major scale is equivalent to saying that there exists a one-to-one correspondence between the syllables and the special set $\{1, 2, 3, 4, 5, 6, 7\}$. Saying “there exists” in this situation may seem a bit weak since in fact there are $7! = 5040$ correspondences, but “there exists” is what we really want here. What exactly is a one-to-one correspondence? Well, we’ve actually seen such things before — a one-to-one correspondence is really just a bijective function between two sets. We’re finally ready to write a definition that Georg Cantor would approve of.

Definition. *For all sets A and B , we say A and B are equivalent, and write $A \equiv B$ iff there exists a one-to-one (and onto) function f , with $\text{Dom}(f) = A$ and $\text{Rng}(f) = B$.*

Somewhat more succinctly, one can just say the sets are equivalent iff there is a bijection between them.

We are going to ask you to prove that the above definition defines an equivalence relation in the exercises for this section. In order to give you a bit of a jump start on that proof we’ll outline what the proof that the relation is symmetric should look like.

To show that the relation is symmetric we must assume that A and B are sets with $A \equiv B$ and show that this implies that $B \equiv A$. According to the definition above this means that we’ll need to locate a function (that is one-to-one) from B to A . On the other hand, since it is given that $A \equiv B$, the definition tells

us that there actually is an injective function, f , from A to B . The inverse function f^{-1} would do exactly what we'd like (namely form a map from B to A) assuming that we can show that f^{-1} has the right properties. We need to know that f^{-1} is a function (remember that in general the inverse of a function is only a relation) and that it is one-to-one. That f^{-1} is a function is a consequence of the fact that f is one-to-one. That f^{-1} is one-to-one is a consequence of the fact that f is a function.

The above is just a sketch of a proof. In the exercise you'll need to fill in the rest of the details as well as provide similar arguments for reflexivity and transitivity.

For each possible finite cardinality k , there are many, many sets having that cardinality, but there is one set that stands out as the most basic – the set of numbers from 1 to k . For each cardinality $k > 0$, we use the symbol $\mathbb{Z}_k^{\text{nonneg}}$ to indicate this set:

$$\mathbb{Z}_k^{\text{nonneg}} = \{1, 2, 3, \dots, k\}.$$

The finite cardinalities are the equivalence classes (under the relation of set equivalence) containing the empty set and the sets $\mathbb{Z}_k^{\text{nonneg}}$. Of course there are also infinite sets! The prototype for an infinite set would have to be the entire set $\mathbb{Z}^{\text{nonneg}}$. The long-standing tradition is to use the symbol \aleph_0 ³ for the cardinality of sets having the same size as $\mathbb{Z}^{\text{nonneg}}$, alternatively, such sets are known as “countable.” One could make a pretty good argument that it is the finite sets that are actually countable! After all it would literally take forever to count the natural numbers! We have to presume that the people who instituted this terminology meant for “countable” to mean “countable, in principle” or “countable if you're willing to let me keep counting forever”

³The Hebrew letter (capital) aleph with a subscript zero – usually pronounced “aleph naught.”

or maybe “countable if you can keep counting faster and faster and are capable of ignoring the speed of light limitations on how fast your lips can move.” Worse yet, the term “countable” has come to be used for sets whose cardinalities are either finite *or* the size of the naturals. If we want to refer specifically to the infinite sort of countable set most mathematicians use the term *denumerable* (although this is not universal) or *countably infinite*. Finally, there are sets whose cardinalities are bigger than the naturals. In other words, there are sets such that no one-to-one correspondence with $\mathbb{Z}^{\text{nonneg}}$ is possible. We don’t mean that people have looked for one-to-one correspondences between such sets and $\mathbb{Z}^{\text{nonneg}}$ and haven’t been able to find them – we literally mean that it can’t be done; and it is has been proved that it can’t be done! Sets having cardinalities that are this ridiculously huge are known as *uncountable*.

Exercises — 8.1

1. Name four sets in the equivalence class of $\{1, 2, 3\}$.
2. Prove that set equivalence is an equivalence relation.
3. Construct a Venn diagram showing the relationships between the sets of sets which are finite, infinite, countable, denumerable and uncountable.
4. Place the sets $\mathbb{Z}^{\text{nonneg}}$, \mathbb{R} , \mathbb{Q} , \mathbb{Z} , $\mathbb{Z} \times \mathbb{Z}$, \mathbb{C} , $\mathbb{Z}_{2007}^{\text{nonneg}}$ and \emptyset ; somewhere on the Venn diagram above. (Note to students (and graders): there are no wrong answers to this question, the point is to see what your intuition about these sets says at this point.)

8.2 Examples of set equivalence

There is an ancient conundrum about what happens when an irresistible force meets an immovable object. In a similar spirit there are sometimes heated debates among young children concerning which super-hero will win a fight. Can Wolverine take Batman? What about the Incredible Hulk versus the Thing? Certainly Superman is at the top of the heap in this ordering. Or is he? Would the man of steel even engage in a fight with a female super-hero, say Wonder Woman? (Remember the 1950's sensibilities of Clark Kent's alter ego.)

To many people the current topic will seem about as sensible as the school-yard discussions just alluded to. We are concerned with knowing whether one infinite set is bigger than another, or are they the same size. There are generally three reasons that people disdain to consider such questions. The first is that, like super-heros, infinite sets are just products of the imagination. The second is that there can be no difference because "infinite is infinite" – once you get to the size we call infinity, you can't add something to that to get to a bigger infinity. The third is that the answers to questions like this are not going to earn me big piles of money so "who cares?"

Point one is actually pretty valid. Physicists have determined that we appear to inhabit a universe of finite scope, containing a finite number of subatomic particles, so in reality there can be no infinite sets. Nevertheless, the axioms we use to study many fields in mathematics guarantee that the objects of consideration are indeed infinite in number. Infinity appears as a concept even when we know it can't appear in actuality. Point two, the "there's only one size of infinity" argument is wrong. We'll see an informal argument showing that there are at least two sizes of infinity, and a more formal theorem that shows there is actually an infinite hierarchy of infinities in Section [8.3](#)

Point three, "who cares?" is in some sense the toughest of all to deal with.

Hopefully you'll enjoy the clever arguments to come for their own intrinsic beauty. But, if you can figure a way to make big piles of money using this stuff that would be nice too.

Let's get started.

Which set is bigger – the natural numbers, $\mathbb{Z}^{\text{nonneg}}$ or the set, $\mathbb{E}^{\text{nonneg}}$, of non-negative even numbers? Both are clearly infinite, so the “infinity is infinity” camp might lead to the correct conclusion through invalid reasoning. On the other hand, the even numbers are contained in the natural numbers so there's a pretty compelling case for saying the evens are somehow smaller than the naturals. The mathematically rigorous way to show that these sets have the same cardinality is by displaying a one-to-one correspondence. Given an even number how can we produce a natural to pair it with? And, given a natural how can we produce an even number to pair with it? The map $f : N \rightarrow \mathbb{E}^{\text{nonneg}}$ defined by $f(x) = 2x$ is clearly a function, and just about as clearly, injective⁴. Is the map f also a surjection? In other words, is every non-negative even number the image of some natural under f ? Given some non-negative even number e we need to be able to come up with an x such that $f(x) = e$. Well, since e is an even number, by the definition of “even” we know that there is an integer k such that $e = 2k$ and since e is either zero or positive it follows that k must also be either 0 or positive. It turns out that k is actually the x we are searching for. Put more succinctly, every non-negative even number $2k$ has a preimage, k , under the map f . So f maps $\mathbb{Z}^{\text{nonneg}}$ surjectively onto $\mathbb{E}^{\text{nonneg}}$. Now the sets we've just considered,

$$\mathbb{Z}^{\text{nonneg}} = \{0, 1, 2, 3, 4, 5, 6, \dots\}$$

and

$$\mathbb{E}^{\text{nonneg}} = \{0, 2, 4, 6, 8, 10, 12, \dots\}$$

⁴If x and y are different numbers that map to the same value, then $f(x) = f(y)$ so $2x = 2y$. But we can cancel the 2's and derive that $x = y$, which is a contradiction.

both have the feature that they can be listed – at least in principle. There is a first element, followed by a second element, followed by a third element, et cetera, in each set. The next set we'll look at, \mathbb{Z} , can't be listed so easily. To list the integers we need to let the dot-dot-dots go both forward (towards positive infinity) and backwards (towards negative infinity),

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

To show that the integers are actually equinumerous with the natural numbers (which is what we're about to do – and by the way, isn't that pretty remarkable?) we need, essentially, to figure out a way to list the integers in a singly infinite list. Using the symbol \pm we can arrange for a singly infinite listing, and if you think about what the symbol \pm means you'll probably come up with

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

This singly infinite listing of the integers does the job we're after in a sense – it displays a one-to-one correspondence with $\mathbb{Z}^{\text{nonneg}}$. In fact any singly infinite listing can be thought of as displaying a one-to-one correspondence with $\mathbb{Z}^{\text{nonneg}}$ – the first entry (or should we say zeroth entry?) in the list is corresponded with 0, the second entry is corresponded with 1, and so on.

$$\begin{array}{cccccccc} 0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 & \dots \\ \updownarrow & \updownarrow & \updownarrow & \updownarrow & \updownarrow & \updownarrow & \updownarrow & & \\ 0 & 1 & -1 & 2 & -2 & 3 & -3 & 4 & \dots \end{array}$$

To make all of this precise we need to be able to explicitly give the one-to-one correspondence. It isn't enough to have a picture of it – we need a formula. Notice that the negative integers are all paired with even naturals and the positive integers are all paired with odd naturals. This observation leads us to a piecewise definition for a function that gives the bijection we seek

$$f(x) = \begin{cases} -x/2 & \text{if } x \text{ is even} \\ (x+1)/2 & \text{if } x \text{ is odd} \end{cases}.$$

By the way, notice that since 0 is even it falls into the first case, and fortunately that formula gives the “right” value.

Exercise. *The inverse function, f^{-1} , must also be defined piecewise, but based on whether the input is positive or negative. Define the inverse function.*

The examples we’ve done so far have shown that the integers, the natural numbers and the even naturals all have the same cardinality. This is the first infinite cardinal number, known as \aleph_0 . In a certain sense we could view both of the equivalences we’ve shown as demonstrating that $2 \cdot \infty = \infty$. Our next example will lend credence to the rule: $\infty \cdot \infty = \infty$. The Cartesian product of two finite sets (the set of all ordered pairs with entries from the sets in question) has cardinality equal to the product of the cardinalities of the sets. What do you suppose will happen if we let the sets be infinite? For instance, what is the cardinality of $\mathbb{Z}^{\text{nonneg}} \times \mathbb{Z}^{\text{nonneg}}$? Consider this: the subset of ordered pairs that start with a 0 can be thought of as a copy of $\mathbb{Z}^{\text{nonneg}}$ sitting inside this Cartesian product. In fact the subset of ordered pairs starting with any particular number gives another copy of $\mathbb{Z}^{\text{nonneg}}$ inside $\mathbb{Z}^{\text{nonneg}} \times \mathbb{Z}^{\text{nonneg}}$. There are infinitely many copies of $\mathbb{Z}^{\text{nonneg}}$ sitting inside of $\mathbb{Z}^{\text{nonneg}} \times \mathbb{Z}^{\text{nonneg}}$! This just really ought to get us to a larger cardinality. The surprising result that it *doesn’t* involves an idea sometimes known as “Cantor’s Snake” – a trick that allows us to list the elements of $\mathbb{Z}^{\text{nonneg}} \times \mathbb{Z}^{\text{nonneg}}$ in a singly infinite list⁵. You can visualize the set $\mathbb{Z}^{\text{nonneg}} \times \mathbb{Z}^{\text{nonneg}}$ as the points having integer coordinates in the first quadrant (together with the origin and the positive x and y axes).

⁵Cantor’s snake was originally created to show that $\mathbb{Q}^{\text{nonneg}}$ and $\mathbb{Z}^{\text{nonneg}}$ are equinumerous. This function was introduced in the exercises for Section 6.5. The version we are presenting here avoids certain complications.

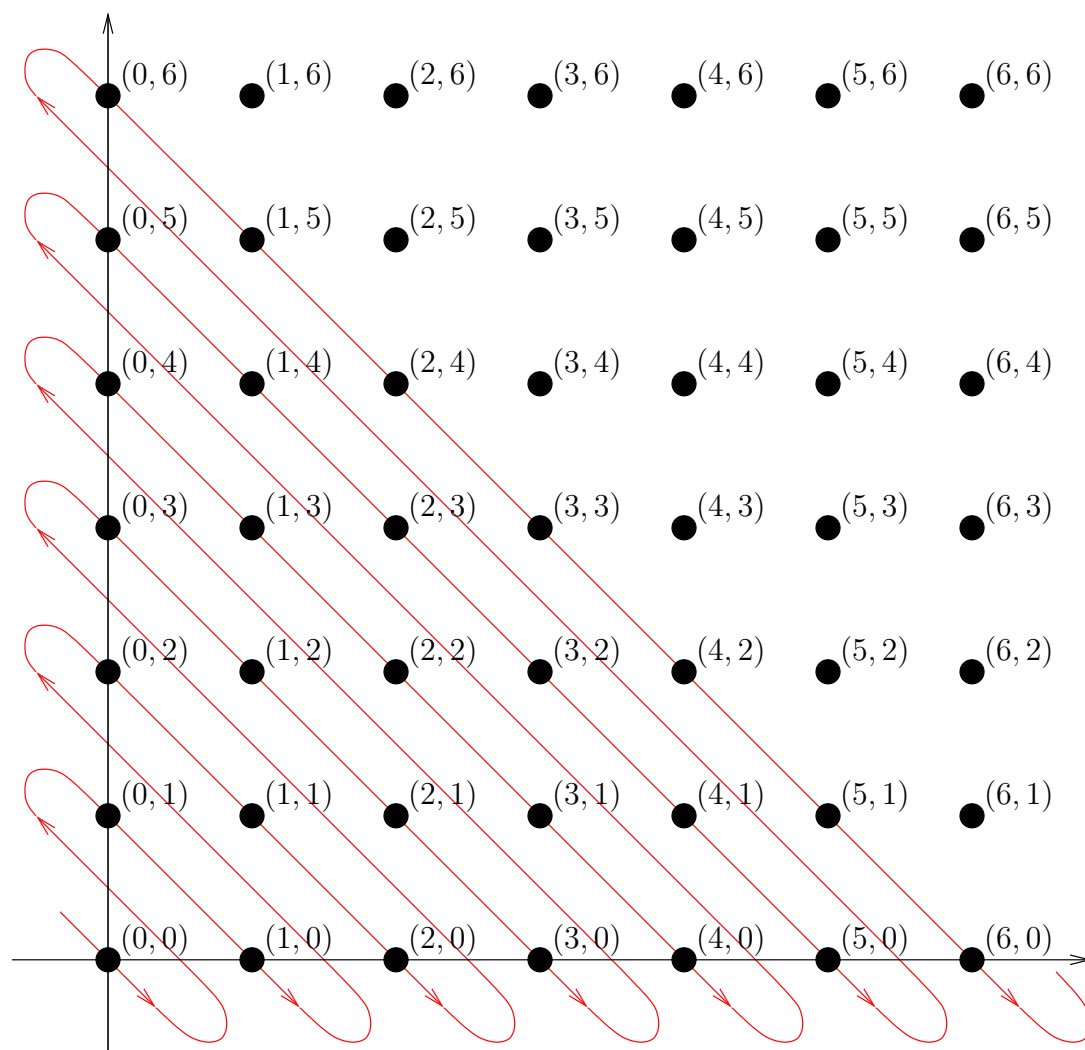


Figure 8.1: Cantor's snake winds through the set $\mathbb{Z}^{\text{noneg}} \times \mathbb{Z}^{\text{noneg}}$ encountering its elements one after the other.

This set of points and the path through them known as Cantor's snake is shown in Figure 8.1.

The diagram in Figure 8.1 gives a visual form of the one-to-one correspondence we seek. In tabular form we would have something like the following.

$$\begin{array}{cccccccccc}
 0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & \dots \\
 \updownarrow & \updownarrow & \updownarrow & \updownarrow & \updownarrow & \updownarrow & \updownarrow & \updownarrow & \updownarrow & \\
 (0, 0) & (0, 1) & (1, 0) & (0, 2) & (1, 1) & (2, 0) & (0, 3) & (1, 2) & (2, 1) & \dots
 \end{array}$$

We need to produce a formula. In truth, we should really produce two formulas. One that takes an ordered pair (x, y) and produces a number n . Another that takes a number n and produces an ordered pair (x, y) . The number n tells us where the pair (x, y) lies in our infinite listing. There is a problem though: the second formula (that gives the map from $\mathbb{Z}^{\text{nonneg}}$ to $\mathbb{Z}^{\text{nonneg}} \times \mathbb{Z}^{\text{nonneg}}$) is really hard to write down – it's easier to describe the map algorithmically. A simple observation will help us to deduce the various formulas. The ordered pairs along the y -axis (those of the form $(0, \text{something})$) correspond to triangular numbers. In fact the pair $(0, n)$ will correspond to the n -th triangular number, $T(n) = (n^2 + n)/2$. The ordered pairs along the descending slanted line starting from $(0, n)$ all have the feature that the sum of their coordinates is n (because as the x -coordinate is increasing, the y -coordinate is decreasing). So, given an ordered pair (x, y) , the number corresponding to the position at the upper end of the slanted line it is on (which will have coordinates $(0, x + y)$) will be $T(x + y)$, and the pair (x, y) occurs in the listing exactly x positions after $(0, x + y)$. Thus, the function $f : \mathbb{Z}^{\text{nonneg}} \times \mathbb{Z}^{\text{nonneg}} \rightarrow N$ is given by

$$f(x, y) = x + T(x + y) = x + \frac{(x + y)^2 + (x + y)}{2}.$$

To go the other direction – that is, to take a position in the listing and derive an ordered pair – we need to figure out where a given number lies relative to the triangular numbers. For instance, try to figure out what (x, y) pair position number 13 will correspond with. Well, the next smaller triangular number is 10 which is $T(4)$, so 13 will be the number of an ordered pair along the descending line whose y -intercept is 4. In fact, 13 will be paired with an ordered pair having a 3 in the x -coordinate (since 13 is 3 larger than 10) so it follows that $f^{-1}(13) = (3, 1)$.

Of course we need to generalize this procedure. One of the hardest parts of finding that generalization is finding the number 4 in the above example (when we just happen to notice that $T(4) = 10$). What we're really doing there is inverting the function $T(n)$. Finding an inverse for $T(n) = (n^2 + n)/2$ was the essence of one of the exercises in Section 6.6. The parabola $y = (x^2 + x)/2$ has roots at 0 and -1 and is scaled by a factor of $1/2$ relative to the "standard" parabola $y = x^2$. Its vertex is at $(-1/2, -1/8)$. The graph of the inverse relation is, of course, obtained by reflecting through the line $y = x$ and by considering scaling and horizontal/vertical translations we can deduce a formula for a function that gives a right inverse for T ,

$$T^{-1}(x) = \sqrt{2x + 1/4} - 1/2.$$

So, given n , a position in the listing, we calculate $A = \lfloor \sqrt{2n + 1/4} - 1/2 \rfloor$. The x -coordinate of our ordered pair is $n - T(A)$ and the y -coordinate is $A - x$. It is not pretty, but the above discussion can be translated into a formula for f^{-1} .

$$f^{-1}(n) = \left(n - \frac{\lfloor \sqrt{2n + 1/4} - 1/2 \rfloor^2 + \lfloor \sqrt{2n + 1/4} - 1/2 \rfloor}{2}, \right. \\ \left. \lfloor \sqrt{2n + 1/4} - 1/2 \rfloor - n + \frac{\lfloor \sqrt{2n + 1/4} - 1/2 \rfloor^2 + \lfloor \sqrt{2n + 1/4} - 1/2 \rfloor}{2} \right).$$

When restricted to the appropriate sets (f 's domain is restricted to $\mathbb{Z}^{\text{noneg}} \times \mathbb{Z}^{\text{noneg}}$ and f^{-1} 's domain is restricted to $\mathbb{Z}^{\text{noneg}}$), these functions are two-sided inverses for one another. That fact is sufficient to prove that f is bijective. So far we have shown that the sets $\mathbb{E}^{\text{noneg}}$, $\mathbb{Z}^{\text{noneg}}$, \mathbb{Z} and $\mathbb{Z}^{\text{noneg}} \times \mathbb{Z}^{\text{noneg}}$ all have the same cardinality — \aleph_0 . We plan to provide an argument that there actually are other infinite cardinals in the next section. Before leaving the present topic (examples of set equivalence) we'd like to present another nice technique for deriving the bijective correspondences we use to show that sets

are equivalent – geometric constructions. Consider the set of points on the line segment $[0, 1]$. Now consider the set of points on the line segment $[0, 2]$. This second line segment, being twice as long as the first, must have a lot more points on it. Right?

Well, perhaps you're getting used to this sort of thing. . . The interval $[0, 1]$ is a subset of the interval $[0, 2]$, but since both represent infinite sets of points it's possible they actually have the same cardinality. We can prove that this is so using a geometric technique. We position the line segments appropriately and then use projection from a carefully chosen point to develop a bijection. Imagine both intervals as lying on the x -axis in the x - y plane. Shift the smaller interval up one unit so that it lies on the line $y = 1$. Now, use projection from the point $(0, 2)$, to visualize the correspondence see Figure 8.2

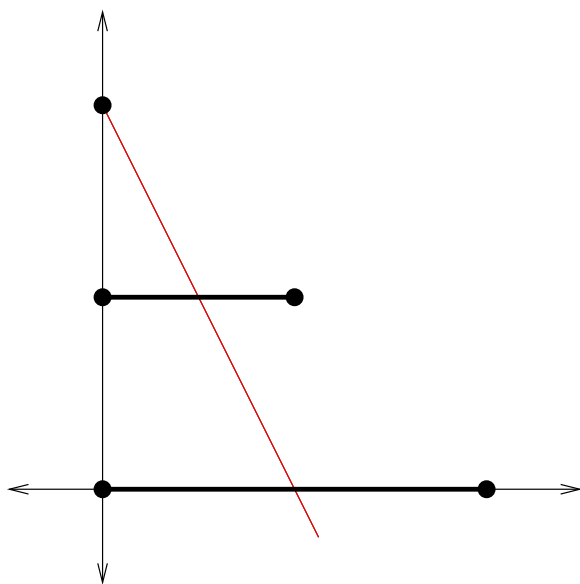


Figure 8.2: Projection from a point can be used to show that intervals of different lengths contain the same number of points.

By considering appropriate projections we can prove that any two arbitrary intervals (say $[a, b]$ and $[c, d]$) have the same cardinalities! It also isn't all that hard to derive a formula for a bijective function between two

intervals.

$$f(x) = c + \frac{(x - a)(d - c)}{(b - a)}$$

There are other geometric constructions which we can use to show that there are the same number of points in a variety of entities. For example, consider the upper half of the unit circle (Remember the unit circle from Trig? All points (x, y) satisfying $x^2 + y^2 = 1$.) This is a semi-circle having a radius of 1, so the arclength of said semi-circle is π . It isn't hard to imagine that this semi-circular arc contains the same number of points as an interval of length π , and we've already argued that all intervals contain the same number of points. . . But, a nice example of geometric projection — vertical projection (a.k.a. π_1) — can be used to show that (for example) the interval $(-1, 1)$ and the portion of the unit circle lying in the upper half-plane are equinumerous.

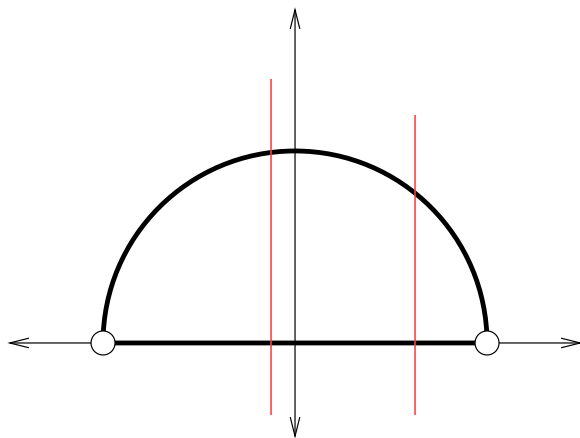


Figure 8.3: Vertical projection provides a bijective correspondence between an interval and a semi-circle.

Once the bijection is understood geometrically it is fairly simple to provide formulas. To go from the semi-circle to the interval, we just forget about the y -coordinate:

$$f(x, y) = x.$$

To go in the other direction we need to recompute the missing y-value:

$$f^{-1}(x) = (x, \sqrt{1 - x^2}).$$

Now we're ready to put some of these ideas together in order to prove something really quite remarkable. It may be okay to say that line segments of different lengths are equinumerous – although one's intuition still balks at the idea that a line a mile long only has the same number of points on it as a line an inch long (or, if you prefer, make that a centimeter versus a kilometer). Would you believe that the entire line – that is the infinitely extended line – has no more points on it than a tiny little segment? You should be ready to prove this one yourself.

Exercise. *Find a point such that projection from that point determines a one-to-one correspondence between the portion of the unit circle in the upper half plane and the line $y = 1$.*

In the exercises from Section 8.1 you were supposed to show that set equivalence is an equivalence relation. Part of that proof should have been showing that the relation is transitive, and that really just comes down to showing that the composition of two bijections is itself a bijection. If you didn't make it through that exercise give it another try now, but whether or not you can finish that proof it should be evident what that transitivity means to us in the current situation. Any pair of line segments are the same size – a line segment (i.e. an interval) and a semi-circle are the same size – the semi-circle and an infinite line are the same size – transitivity tells us that an infinitely extended line has the same number of points as (for example) the interval $(0, 1)$.

Exercises — 8.2

1. Prove that positive numbers of the form $3k + 1$ are equinumerous with positive numbers of the form $4k + 2$.
2. Prove that $f(x) = c + \frac{(x - a)(d - c)}{(b - a)}$ provides a bijection from the interval $[a, b]$ to the interval $[c, d]$.
3. Prove that any two circles are equinumerous (as sets of points).
4. Determine a formula for the bijection from $(-1, 1)$ to the line $y = 1$ determined by vertical projection onto the upper half of the unit circle, followed by projection from the point $(0, 0)$.
5. It is possible to generalize the argument that shows a line segment is equivalent to a line to higher dimensions. In two dimensions we would show that the unit disk (the interior of the unit circle) is equinumerous with the entire plane $\mathbb{R} \times \mathbb{R}$. In three dimensions we would show that the unit ball (the interior of the unit sphere) is equinumerous with the entire space $\mathbb{R}^3 = \mathbb{R} \times \mathbb{R} \times \mathbb{R}$. Here we would like you to prove the two-dimensional case.

Gnomonic projection is a style of map rendering in which a portion of a sphere is projected onto a plane that is tangent to the sphere. The sphere's center is used as the point to project from. Combine vertical projection from the unit disk in the x - y plane to the upper half of the unit sphere $x^2 + y^2 + z^2 = 1$, with gnomonic projection from the unit sphere to the plane $z = 1$, to deduce a bijection between the unit disk and the (infinite) plane.

8.3 Cantor's theorem

Many people believe that the result known as Cantor's theorem says that the real numbers, \mathbb{R} , have a greater cardinality than the natural numbers, $\mathbb{Z}^{\text{nonneg}}$. That isn't quite right. In fact Cantor's theorem is a much broader statement, one of whose consequences is that $|\mathbb{R}| > |\mathbb{Z}^{\text{nonneg}}|$. Before we go on to discuss Cantor's theorem in full generality, we'll first explore it, essentially, in this simplified form. Once we know that $|\mathbb{R}| \neq |\mathbb{Z}^{\text{nonneg}}|$, we'll be in a position to explore a lot of interesting issues relative to the infinite. In particular, this result means that there are at least two cardinal numbers that are infinite – thus the “infinity is infinity” idea will be discredited. Once we have the full power of Cantor's theorem, we'll see just how completely wrong that concept is.

To show that some pair of sets are not equivalent it is necessary to show that there cannot be a one-to-one correspondence between them. Ordinarily, one would try to argue by contradiction in such a situation. That is what we'll need to do to show that the reals and the naturals are not equinumerous. We'll presume that they are in fact the same size and try to reach a contradiction.

What exactly does the assumption that \mathbb{R} and $\mathbb{Z}^{\text{nonneg}}$ are equivalent mean? It means there is a one-to-one correspondence, that is, a bijective function from \mathbb{R} to $\mathbb{Z}^{\text{nonneg}}$. In a nutshell, it means that it is possible to list all the real numbers in a singly-infinite list. Now, it is certainly possible to make an infinite list of real numbers (since $\mathbb{Z}^{\text{nonneg}} \subseteq \mathbb{R}$, by listing the naturals themselves we are making an infinite list of reals!). The problem is that we would need to be sure that every real number is on the list somewhere. In fact, since we've used a geometric argument to show that the interval $(0, 1)$ and the set \mathbb{R} are equinumerous, it will be sufficient to presume that there is an infinite list containing all the numbers in the interval $(0, 1)$.

Exercise. Notice that, for example, $\pi - 3$ is a real number in $(0, 1)$. Make

a list of 10 real numbers in the interval $(0, 1)$. Make sure that at least 5 of them are not rational.

In the previous exercise, you've started the job, but we need to presume that it is truly possible to complete this job. That is, we must presume that there really is an infinite list containing every real number in the interval $(0, 1)$.

Once we have an infinite list containing every real number in the interval $(0, 1)$ we have to face up to a second issue. What does it really mean to list a particular real number? For instance if $e - 2$ is in the seventh position on our list, is it OK to write " $e - 2$ " there or should we write "0.7182818284590452354..."? Clearly it would be simpler to write " $e - 2$ " but it isn't necessarily possible to do something of that kind for every real number – on the other hand, writing down the decimal expansion is a problem too; in a certain sense, "most" real numbers in $(0, 1)$ have infinitely long decimal expansions. There is also another problem with decimal expansions; they aren't unique. For example, there is really no difference between the finite expansion 0.5 and the infinitely long expansion $0.4\bar{9}$.

Rather than writing something like " $e - 2$ " or "0.7182818284590452354...", we are going to in fact write ".1011011111100001010100010110001010001010...". In other words, we are going to write the base-2 expansions of the real numbers in our list. Now, the issue of non-uniqueness is still there in binary, and in fact if we were to stay in base-10 it would be possible to plug a certain gap in our argument – but the binary version of this argument has some especially nice features. Every binary (or for that matter decimal) expansion corresponds to a unique real number, but it doesn't work out so well the other way around — there are sometimes two different binary expansions that correspond to the same real number. There is a lovely fact that we are not going to prove (you may get to see this result proved in a course in Real Analysis) that points up the problem. Whenever two different binary expansions represent the same real number, one of them is a terminating ex-

pansion (it ends in infinitely many 0's) and the other is an infinite expansion (it ends in infinitely many 1's). We won't prove this fact, but the gist of the argument is a proof by contradiction — you may be able to get the point by studying Figure 8.4. (Try to see how it would be possible to find a number in between two binary expansions that didn't end in all-zeros and all-ones.)

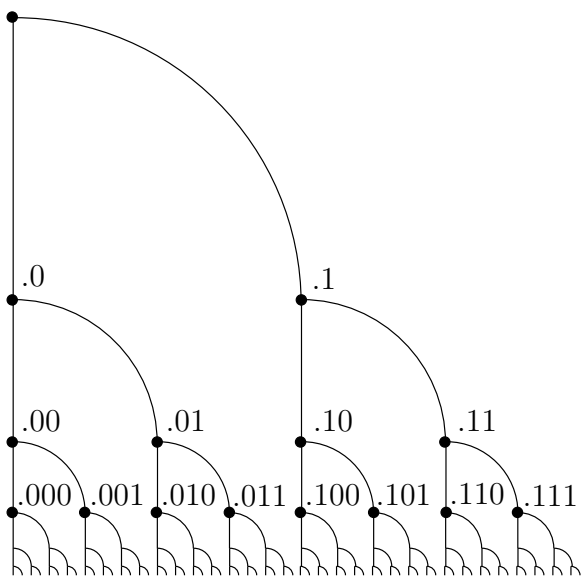


Figure 8.4: The base-2 expansions of reals in the interval $[0, 1]$ are the leaves of an infinite tree.

So, instead of showing that the set of reals in $(0, 1)$ can't be put in one-to-one correspondence with $\mathbb{Z}^{\text{nonneg}}$, what we're really going to do is show that their binary expansions can't be put in one-to-one correspondence with $\mathbb{Z}^{\text{nonneg}}$. Since there are an infinite number of reals that have two different binary expansions this doesn't really do the job as advertised at the beginning of this section. (Perhaps you are getting used to our wily ways by now — yes, this does mean that we're going to ask you to do the real proof in the exercises.) The set of binary numerals, $\{0, 1\}$, is an instance of a mathematical structure known as a field; basically, that means that it's possible to add, subtract, multiply and divide (but not divide by 0) with them. We are only mentioning

this fact so that you'll understand why the set $\{0, 1\}$ is often referred to as \mathbb{F}_2 . We're only mentioning that fact so that you'll understand why we call the set of all possible binary expansion \mathbb{F}_2^∞ . Finally, we're only mentioning *that* fact so that we'll have a succinct way of expressing this set. Now we can write " \mathbb{F}_2^∞ " rather than "the set of all possible infinitely-long binary sequences."

Suppose we had a listing of all the elements of \mathbb{F}_2^∞ . We would have an infinite list of things, each of which is itself an infinite list of 0's and 1's.

So what? We need to proceed from here to find a contradiction.

This argument that we've been edging towards is known as Cantor's diagonalization argument. The reason for this name is that our listing of binary representations looks like an enormous table of binary digits and the contradiction is deduced by looking at the diagonal of this infinite-by-infinite table. The diagonal is itself an infinitely long binary string — in other words, the diagonal can be thought of as a binary expansion itself. If we take the complement of the diagonal, (switch every 0 to a 1 and vice versa) we will also have a thing that can be regarded as a binary expansion and this binary expansion can't be one of the ones on the list! This bit-flipped version of the diagonal is different from the first binary expansion in the first position, it is different from the second binary expansion in the second position, it is different from the third binary expansion in the third position, and so on. The very presumption that we could list all of the elements of \mathbb{F}_2^∞ allows us to construct an element of \mathbb{F}_2^∞ that could not be on the list!

This argument has been generalized many times, so this is the first in a class of things known as diagonal arguments. Diagonal arguments have been used to settle several important mathematical questions. There is a valid diagonal argument that even does what we'd originally set out to do: prove that $\mathbb{Z}^{\text{nonneg}}$ and \mathbb{R} are not equinumerous. Strangely, the argument can't be made to work in binary, and since you're going to be asked to write it up in the exercises, we want to point out one of the potential pitfalls. If we were to use a diagonal argument to show that $(0, 1)$ isn't countable, we would start

by assuming that every element of $(0, 1)$ was written down in a list. For most real numbers in $(0, 1)$ we could write out their binary representation uniquely, but for some we would have to make a choice: should we write down the representation that terminates, or the one that ends in infinitely-many 1's? Suppose we choose to use the terminating representations, then none of the infinite binary strings that end with all 1's will be on the list. It's possible that the thing we get when we complement the diagonal is one of these (unlisted) binary strings so we don't *necessarily* have a contradiction. If we make the other choice – use the infinite binary representation when we have a choice – there is a similar problem. You may think that our use of binary representations for real numbers was foolish in light of the failure of the argument to “go through” in binary. Especially since, as we've alluded to, it can be made to work in decimal. The reason for our apparent stubbornness is that these infinite binary strings do something else that's very nice. An infinitely long binary sequence can be thought of as the indicator function of a subset of \mathbb{N} . For example, $.001101010001$ is the indicator of $\{2, 3, 5, 7, 11\}$.

Exercise. Complete the table.

<i>binary expansion</i>	<i>subset of $\mathbb{Z}^{\text{nonneg}}$</i>
.1	$\{0\}$
.0111	
	$\{2, 4, 6\}$
$\overline{.01}$	
	$\{3k + 1 \mid k \in \mathbb{Z}^{\text{nonneg}}\}$

The set, \mathbb{F}_2^∞ , we've been working with is in one-to-one correspondence with the power set of the natural numbers, $\mathcal{P}(\mathbb{Z}^{\text{nonneg}})$. When viewed in this light, the proof we did above showed that the power set of $\mathbb{Z}^{\text{nonneg}}$ has an infinite cardinality strictly greater than that of $\mathbb{Z}^{\text{nonneg}}$ itself. In other words, $\mathcal{P}(\mathbb{Z}^{\text{nonneg}})$ is uncountable.

What Cantor's theorem says is that this always works. If A is any set, and $\mathcal{P}(A)$ is its power set then $|A| < |\mathcal{P}(A)|$. In a way, this more general theorem is easier to prove than the specific case we just handled.

Theorem 8.3.1 (Cantor). *For all sets A , A is not equivalent to $\mathcal{P}(A)$.*

Proof: Suppose that there is a set A that can be placed in one-to-one correspondence with its power set. Then there is a bijective function $f : A \rightarrow \mathcal{P}(A)$. We will deduce a contradiction by constructing a subset of A (i.e. a member of $\mathcal{P}(A)$) that cannot be in the range of f .

Let $S = \{x \in A \mid x \notin f(x)\}$.

If S is in the range of f , there is a preimage y such that $S = f(y)$. But, if such a y exists then the membership question, $y \in S$, must either be true or false. If $y \in S$, then because $S = f(y)$, and S consists of those elements that are not in their images, it follows that $y \notin S$. On the other hand, if $y \notin S$ then $y \notin f(y)$ so (by the definition of S) it follows that $y \in S$. Either possibility leads to the other, which is a contradiction.

Q.E.D.

Cantor's theorem guarantees that there is an infinite hierarchy of infinite cardinal numbers. Let's put it another way. People have sought a construction that, given an infinite set, could be used to create a strictly larger set. For instance, the Cartesian product works like this if our sets are finite — $A \times A$ is strictly bigger than A when A is a finite set. But, as we've already seen, this is not necessarily so if A is infinite (remember the "snake" argument that $\mathbb{Z}^{\text{noneg}}$ and $\mathbb{Z}^{\text{noneg}} \times \mathbb{Z}^{\text{noneg}}$ are equivalent). The real import of Cantor's theorem is that taking the power set of a set *does* create a set of larger cardinality. So we get an infinite tower of infinite cardinalities, starting with $\aleph_0 = |\mathbb{Z}^{\text{noneg}}|$, by successively taking power sets.

$$\aleph_0 = |\mathbb{Z}^{\text{nonneg}}| < |\mathcal{P}(\mathbb{Z}^{\text{nonneg}})| < |\mathcal{P}(\mathcal{P}(\mathbb{Z}^{\text{nonneg}}))| < |\mathcal{P}(\mathcal{P}(\mathcal{P}(\mathbb{Z}^{\text{nonneg}})))| < \dots$$

Exercises — 8.3

1. Determine a substitution rule – a consistent way of replacing one digit with another along the diagonal so that a diagonalization proof showing that the interval $(0, 1)$ is uncountable will work in decimal. Write up the proof.
2. Can a diagonalization proof showing that the interval $(0, 1)$ is uncountable be made workable in base-3 (ternary) notation?
3. In the proof of Cantor's theorem we construct a set S that cannot be in the image of a presumed bijection from A to $\mathcal{P}(A)$. Suppose $A = \{1, 2, 3\}$ and f determines the following correspondences: $1 \longleftrightarrow \emptyset$, $2 \longleftrightarrow \{1, 3\}$ and $3 \longleftrightarrow \{1, 2, 3\}$. What is S ?
4. An argument very similar to the one embodied in the proof of Cantor's theorem is found in the Barber's paradox. This paradox was originally introduced in the popular press in order to give laypeople an understanding of Cantor's theorem and Russell's paradox. It sounds somewhat sexist to modern ears. (For example, it is presumed without comment that the Barber is male.)

In a small town there is a Barber who shaves those men (and only those men) who do not shave themselves. Who shaves the Barber?

Explain the similarity to the proof of Cantor's theorem.

5. Cantor's theorem, applied to the set of all sets leads to an interesting paradox. The power set of the set of all sets is a collection of sets, so it must be contained in the set of all sets. Discuss the paradox and determine a way of resolving it.

6. Verify that the final deduction in the proof of Cantor's theorem, " $(y \in S \implies y \notin S) \wedge (y \notin S \implies y \in S)$," is truly a contradiction.

8.4 Dominance

We've said a lot about the equivalence relation determined by Cantor's definition of set equivalence. We've also, occasionally, written things like $|A| < |B|$, without being particularly clear about what that means. It's now time to come clean. There is actually a (perhaps) more fundamental notion used for comparing set sizes than equivalence — dominance. Dominance is an ordering relation on the class of all sets. One should probably really define dominance first and then define set equivalence in terms of it. We haven't followed that plan for (at least) two reasons. First, many people may want to skip this section — the results of this section depend on the difficult Cantor-Bernstein-Schröder theorem⁶. Second, we will later take the view that dominance should really be considered to be an ordering relation on the set of all cardinal numbers — i.e. the equivalence classes of the set equivalence relation — not on the collection of all sets. From that perspective, set equivalence really needs to be defined *before* dominance.

One set is said to dominate another if there is a function from the latter *into* the former. More formally, we have the following

Definition. *If A and B are sets, we say “ A dominates B ” and write $|A| > |B|$ iff there is an injective function f with domain B and codomain A .*

It is easy to see that this relation is reflexive and transitive. The Cantor-Bernstein-Schröder theorem proves that it is also anti-symmetric — which means dominance is an ordering relation. Be advised that there is an abuse of terminology here that one must be careful about — what are the domain and range of the “dominance” relation? The definition would lead us to think that sets are the things that go on either side of the “dominance”

⁶This theorem has been known for many years as the Schröder-Bernstein theorem, but, lately, has had Cantor's name added as well. Since Cantor proved the result before the other gentlemen this is fitting. It is also known as the Cantor-Bernstein theorem (leaving out Schröder) which doesn't seem very nice.

relation, but the notation is a bit more honest, “ $|A| > |B|$ ” indicates that the things really being compared are the cardinal numbers of sets (not the sets themselves). Thus anti-symmetry for this relation is

$$(|A| > |B|) \wedge (|B| > |A|) \implies (|A| = |B|).$$

In other words, if A dominates B and vice versa, then A and B are equivalent sets — a strict interpretation of anti-symmetry for this relation might lead to the conclusion that A and B are actually the same set, which is clearly an absurdity.

Naturally, we want to prove the Cantor-Bernstein-Schröder theorem (which we’re going to start calling the C-B-S theorem for brevity), but first it’ll be instructive to look at some of its consequences. Once we have the C-B-S theorem we get a very useful shortcut for proving set equivalences. Given sets A and B , if we can find injective functions going between them in both directions, we’ll know that they’re equivalent. So, for example, we can use C-B-S to prove that the set of all infinite binary strings and the set of reals in $(0, 1)$ really are equinumerous. (In case you had some remaining doubt. . .)

It is easy to dream up an injective function from $(0, 1)$ to \mathbb{F}_2^∞ : just send a real number to its binary expansion, and if there are two, make a consistent choice — let’s say we’ll take the non-terminating expansion.

There is a cute thought-experiment called Hilbert’s Hotel that will lead us to a technique for developing an injective function in the other direction. Hilbert’s Hotel has \aleph_0 rooms. If any countable collection of guests show up there will be enough rooms for everyone. Suppose you arrive at Hilbert’s hotel one dark and stormy evening and the “No Vacancy” light is on — there are already a denumerable number of guests there — every room is full. The clerk sees you dejectedly considering your options, trying to think of another hotel that might still have rooms when, clearly, a *very* large convention is in town. He rushes out and says “My friend, have no fear! Even though we have no vacancies, there is always room for one more at our establishment.”

He goes into the office and makes the following announcement on the PA system. “Ladies and Gentlemen, in order to accommodate an incoming guest, please vacate your room and move to the room numbered one higher. Thank you.” There is an infinite amount of grumbling, but shortly you find yourself occupying room number 1.

To develop an injection from \mathbb{F}_2^∞ to $(0, 1)$ we’ll use “room number 1” to separate the binary expansions that represent the same real number. Move all the digits of a binary expansion down by one, and make the first digit 0 for (say) the terminating expansions and 1 for the non-terminating ones. Now consider these expansions as real numbers — all the expansions that previously coincided are now separated into the intervals $(0, 1/2)$ and $(1/2, 1)$. Notice how funny this map is, there are now many, many, (infinitely-many) real numbers with no preimages. For instance, only a subset of the rational numbers in $(0, 1/2)$ have preimages. Nevertheless, the map is injective, so C-B-S tells us that \mathbb{F}_2^∞ and $(0, 1)$ are equivalent. There are quite a few different proofs of the C-B-S theorem. The one Cantor himself wrote relies on the axiom of choice. The axiom of choice was somewhat controversial when it was introduced, but these days most mathematicians will use it without qualms. What it says (essentially) is that it is possible to make an infinite number of choices. More precisely, it says that if we have an infinite set consisting of non-empty sets, it is possible to select an element out of each set. If there is a definable rule for picking such an element (as is the case, for example, when we selected the nonterminating decimal expansion whenever there was a choice in defining the injection from $(0, 1)$ to \mathbb{F}_2^∞) the axiom of choice isn’t needed. The usual axioms for set theory were developed by Zermelo and Frankel, so you may hear people speak of the ZF axioms. If, in addition, we want to specifically allow the axiom of choice, we are in the ZFC axiom system. If it’s possible to construct a proof for a given theorem without using the axiom of choice, almost everyone would agree that that is preferable. On the other hand, a proof of the C-B-S theorem, which necessarily must be able

to deal with uncountably infinite sets, will have to depend on some sort of notion that will allow us to deal with huge infinities.

The proof we will present here⁷ is attributed to Julius König. König was a contemporary of Cantor's who was (initially) very much respected by him. Cantor came to dislike König after the latter presented a well-publicized (and ultimately wrong) lecture claiming the continuum hypothesis was false. Apparently the continuum hypothesis was one of Cantor's favorite ideas, because he seems to have construed König's lecture as a personal attack. Anyway...

König's proof of C-B-S doesn't use the axiom of choice, but it does have its own strangeness: a function that is not necessarily computable — that is, a function for which (for certain inputs) it may not be possible to compute an output in a finite amount of time! Except for this oddity, König's proof is probably the easiest to understand of all the proofs of C-B-S. Before we get too far into the proof it is essential that we understand the basic setup. The Cantor-Bernstein-Schröder theorem states that whenever A and B are sets and there are injective functions $f : A \rightarrow B$ and $g : B \rightarrow A$, then it follows that A and B are equivalent. Saying A and B are equivalent means that we can find a bijective function between them. So, to prove C-B-S, we hypothesize the two injections and somehow we must construct the bijection.

Figure 8.5 has a presumption in it — that A and B are countable — which need not be the case. Nevertheless, it gives us a good picture to work from. The basic hypotheses, that A and B are sets and we have two functions, one from A into B and another from B into A , are shown. We will have to build our bijective function in a piecewise manner. If there is a non-empty intersection between A and B , we can use the identity function for that part of the domain of our bijection. So, without loss of generality, we can presume that A and B are disjoint. We can use the functions f and g to create infinite sequences, which alternate back and forth between A and B , containing any

⁷We first encountered this proof in a Wikipedia article[3].

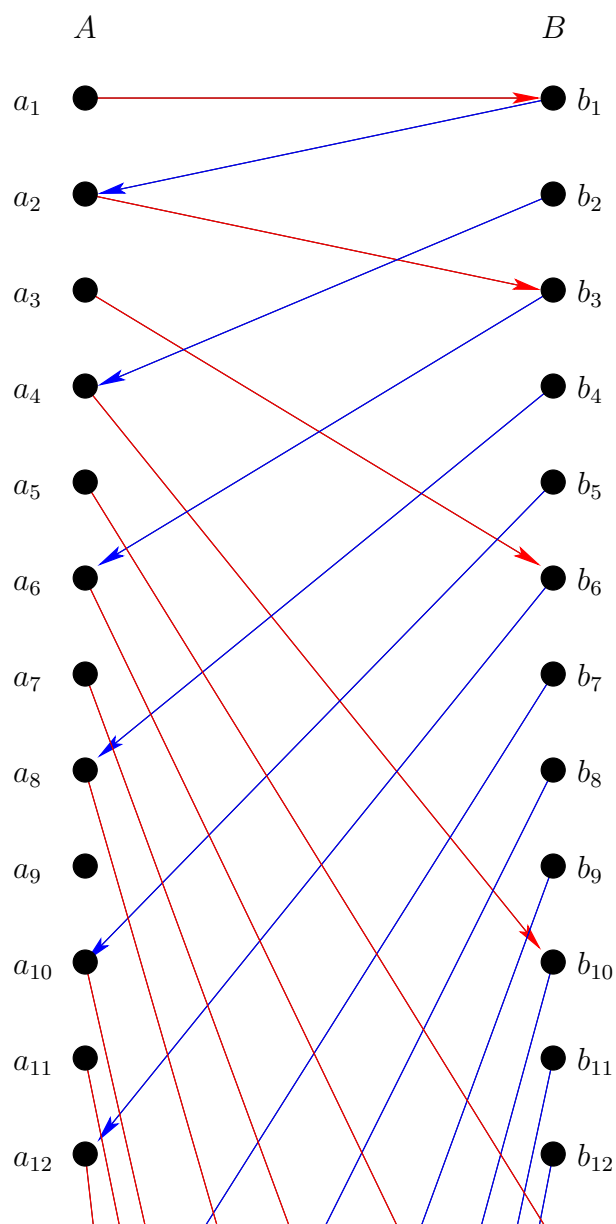


Figure 8.5: Hypotheses for proving the Cantor-Bernstein-Schröder theorem: two sets with injective functions going both ways.

particular element. Suppose $a \in A$ is an arbitrary element. Since f is defined on all of A , we can compute $f(a)$. Now since $f(a)$ is an element of B , and g is defined on all of B , we can compute $g(f(a))$, and so on. Thus, we get the infinite sequence

$$a, \quad f(a), \quad g(f(a)), \quad f(g(f(a))), \quad \dots$$

If the element a also happens to be the image of something under g (this may or may not be so — since g isn't necessarily onto) then we can also extend this sequence to the left. Indeed, it may be possible to extend the sequence infinitely far to the left, or, this process may stop when one of f^{-1} or g^{-1} fails to be defined.

$$\dots g^{-1}(f^{-1}(g^{-1}(a))), \quad f^{-1}(g^{-1}(a)), \quad g^{-1}(a), \quad a, \quad f(a), \quad g(f(a)), \quad f(g(f(a))), \quad \dots$$

Now, every element of the disjoint union of A and B is in one of these sequences. Also, it is easy to see that these sequences are either disjoint or identical. Taking these two facts together it follows that these sequences form a partition of $A \cup B$. We'll define a bijection $\phi : A \cup B \rightarrow A \cup B$ by deciding what it must do on these sequences. There are four possibilities for how the sequences we've just defined can play out. In extending them to the left, we may run into a place where one of the inverse functions needed isn't defined — or not. We say a sequence is an A -stopper, if, in extending to the left, we end up on an element of A that has no preimage under g (see Figure 8.6). Similarly, we can define a B -stopper. If the inverse functions are always defined within a given sequence there are also two possibilities; the sequence may be finite (and so it must be cyclic in nature) or the sequence may be truly infinite.

Finally, here is a definition for ϕ .

$$\phi(x) = \begin{cases} g^{-1}(x) & \text{if } x \text{ is in a } B\text{-stopper} \\ f(x) & \text{otherwise} \end{cases}$$

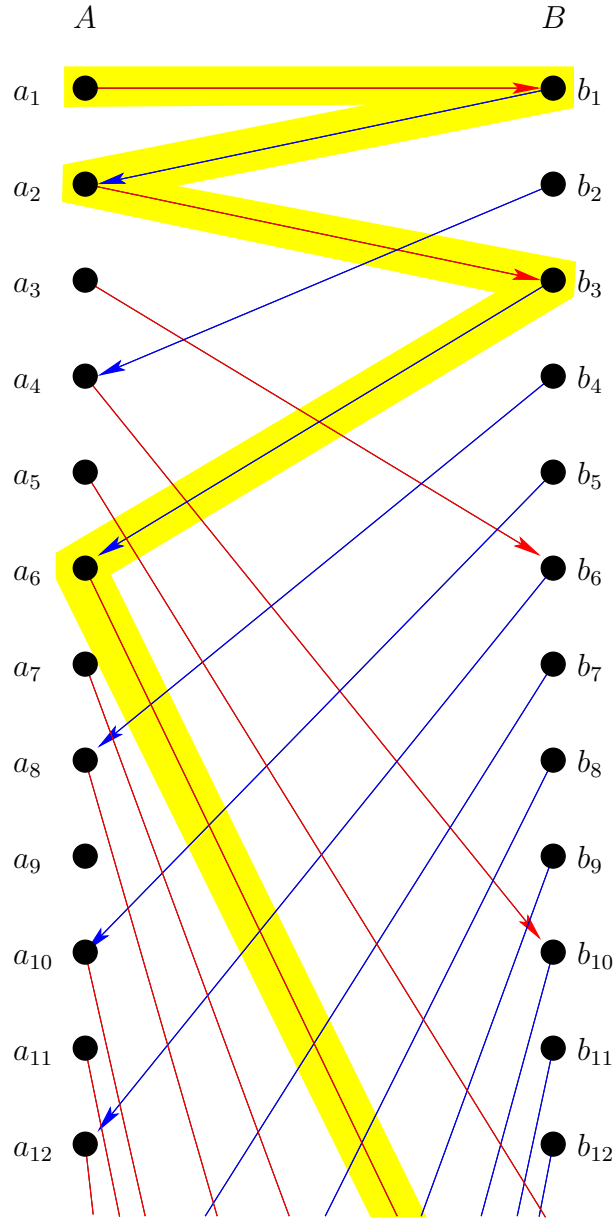


Figure 8.6: An A -stopper is an infinite sequence that terminates to the left in A .

Notice that if a sequence is either cyclic or infinite it doesn't matter whether we use f or g^{-1} since both will be defined for all elements of such sequences. Also, certainly f will work if we are in an A -stopper. The function we've just created is perfectly well-defined, but it may take arbitrarily long to determine whether we have an element of a B -stopper, as opposed to an element of an infinite sequence. We cannot determine whether we're in an infinite versus a finite sequence in a prescribed finite number of steps.

Exercises — 8.4

1. How could the clerk at the Hilbert Hotel accommodate a countable number of new guests?
2. Let F be the collection of all real-valued functions defined on the real line. Find an injection from \mathbb{R} to F . Do you think it is possible to find an injection going the other way? In other words, do you think that F and \mathbb{R} are equivalent? Explain.
3. Fill in the details of the proof that dominance is an ordering relation. (You may simply cite the C-B-S theorem in proving anti-symmetry.)
4. We can inject \mathbb{Q} into \mathbb{Z} by sending $\pm\frac{a}{b}$ to $\pm 2^a 3^b$. Use this and another obvious injection to (in light of the C-B-S theorem) reaffirm the equivalence of these sets.

8.5 The continuum hypothesis and the generalized continuum hypothesis

The word “continuum” in the title of this section is used to indicate sets of points that have a certain continuity property. For example, in a real interval it is possible to move from one point to another, in a smooth fashion, without ever leaving the interval. In a range of rational numbers this is not possible, because there are irrational values in between every pair of rationals. There are many sets that behave as a continuum – the intervals (a, b) or $[a, b]$, the entire real line \mathbb{R} , the x-y plane $\mathbb{R} \times \mathbb{R}$, a volume in 3-dimensional space (or for that matter the entire space \mathbb{R}^3). It turns out that all of these sets have the same size.

The cardinality of the continuum, denoted \mathfrak{c} , is the cardinality of all of the sets above.

In the previous section we mentioned the continuum hypothesis and how angry Cantor became when someone (König) tried to prove it was false. In this section we’ll delve a little deeper into what the continuum hypothesis says and even take a look at CH’s big brother, GCH. Before doing so, it seems like a good idea to look into the equivalences we’ve asserted about all those sets above which (if you trust us) have the cardinality \mathfrak{c} .

We’ve already seen that an interval is equivalent to the entire real line but the notion that the entire infinite Cartesian plane has no more points in it than an interval one inch long defies our intuition. Our conception of dimensionality leads us to think that things of higher dimension must be larger than those of lower dimension. This preconception is false as we can see by demonstrating that a 1×1 square can be put in one-to-one correspondence with the unit interval. Let $S = \{(x, y) \mid 0 < x < 1 \wedge 0 < y < 1\}$ and let I be the open unit interval $(0, 1)$. We can use the Cantor-Bernstein-Schroeder theorem to show that S and I are equinumerous – we just need to find injections from I to S and vice versa. Given an element r in I we can map

it injectively to the point (r, r) in S . To go in the other direction, consider a point (a, b) in S and write out the decimal expansions of a and b :

$$a = 0.a_1a_2a_3a_4a_5 \dots$$

$$b = 0.b_1b_2b_3b_4b_5 \dots$$

as usual, if there are two decimal expansions for a and/or b we will make a consistent choice – say the infinite one.

From these decimal expansions, we can create the decimal expansion of a number in I by interleaving the digits of a and b . Let

$$s = 0.a_1b_1a_2b_2a_3b_3 \dots$$

be the image of (a, b) . If two different points get mapped to the same value s then both points have x and y coordinates that agree in every position of their decimal expansion (so they must really be equal). It is a little bit harder to create a bijective function from S to I (and thus to show the equivalence directly, without appealing to C-B-S). The problem is that, once again, we need to deal with the non-uniqueness of decimal representations of real numbers. If we make the choice that, whenever there is a choice to be made, we will use the non-terminating decimal expansions for our real numbers there will be elements of I not in the image of the map determined by interleaving digits (for example 0.15401050902060503 is the interleaving of the digits after the decimal point in $\pi = 3.141592653 \dots$ and $1/2 = 0.5$, this is clearly an element of I but it can't be in the image of our map since $1/2$ should be represented by $0.4\bar{9}$ according to our convention. If we try other conventions for dealing with the non-uniqueness it is possible to find other examples that show simple interleaving will not be surjective. A slightly more subtle approach is required.

Presume that all decimal expansions are non-terminating (as we can, WLOG) and use the following approach: Write out the decimal expansion of the coordinates of a point (a, b) in S . Form the digits into blocks with as

many 0's as possible followed by a non-zero digit. Finally, interleave these blocks.

For example if

$$a = 0.124520047019902\dots$$

and

$$b = 0.004015648000031\dots$$

we would separate the digits into blocks as follows:

$$a = 0.1 \quad 2 \quad 4 \quad 5 \quad 2 \quad 004 \quad 7 \quad 01 \quad 9 \quad 9 \quad 02\dots$$

and

$$b = 0.004 \quad 01 \quad 5 \quad 6 \quad 4 \quad 8 \quad 00003 \quad 1\dots$$

and the number formed by interleaving them would be

$$s = 0.10042014556240048\dots$$

We've shown that the unit square, S , and the unit interval, I , have the same cardinality. These arguments can be extended to show that all of $R \times R$ also has this cardinality (**c**).

So now let's turn to the continuum hypothesis.

We mentioned earlier in this chapter that the cardinality of $\mathbb{Z}^{\text{nonneg}}$ is denoted \aleph_0 . The fact that that capital letter aleph is wearing a subscript ought to make you wonder what other aleph-sub-something-or-others there are out there. What is \aleph_1 ? What about \aleph_2 ? Cantor presumed that there was a sequence of cardinal numbers (which is itself, of course, infinite) that give all of the possible infinities. The smallest infinite set that anyone seems to be able to imagine is $\mathbb{Z}^{\text{nonneg}}$, so Cantor called that cardinality \aleph_0 . What ever

the “next” infinite cardinal is, is called \aleph_1 . It’s conceivable that there actually isn’t a “next” infinite cardinal after \aleph_0 — it might be the case that the collection of infinite cardinal numbers isn’t well-ordered! In any case, if there *is* a “next” infinite cardinal, what is it? Cantor’s theorem shows that there is a way to build *some* infinite cardinal bigger than \aleph_0 — just apply the power set construction. The continuum hypothesis just says that this bigger cardinality that we get by applying the power set construction *is* that “next” cardinality we’ve been talking about.

To re-iterate, we’ve shown that the power set of $\mathbb{Z}^{\text{noneg}}$ is equivalent to the interval $(0, 1)$ which is one of the sets whose cardinality is \mathbf{c} . So the continuum hypothesis, the thing that got Georg Cantor so very heated up, comes down to asserting that

$$\aleph_1 = \mathbf{c}.$$

There really should be a big question mark over that. A *really* big question mark. It turns out that the continuum hypothesis lives in a really weird world. . . To this day, no one has the least notion of whether it is true or false. But wait! That’s not all! The real weirdness is that it would appear to be *impossible* to decide. Well, that’s not *so* bad — after all, we talked about undecidable sentences way back in the beginning of Chapter 2. Okay, so here’s the ultimate weirdness. It has been *proved* that one can’t prove the continuum hypothesis. It has also been *proved* that one can’t disprove the continuum hypothesis.

Having reached this stage in a book about proving things I hope that the last two sentences in the previous paragraph caused some thought along the lines of “well, ok, with respect to what axioms?” to run through your head. So, if you did think something along those lines pat yourself on the back. And if you *didn’t* then recognize that you need to start thinking that way — things are proved or disproved only in a relative way, it depends what axioms you allow yourself to work with. The usual axioms for mathematics

are called ZFC; the Zermelo-Frankel set theory axioms together with the axiom of choice. The “ultimate weirdness” we’ve been describing about the continuum hypothesis is a result due to a gentleman named Paul Cohen that says “CH is independent of ZFC.” More pedantically – it is impossible to either prove or disprove the continuum hypothesis within the framework of the ZFC axiom system.

It would be really nice to end this chapter by mentioning Paul Cohen, but there is one last thing we’d like to accomplish — explain what GCH means. So here goes.

The generalized continuum hypothesis says that the power set construction is basically the only way to get from one infinite cardinality to the next. In other words GCH says that not only does $\mathcal{P}(\mathbb{Z}^{\text{nonneg}})$ have the cardinality known as \aleph_1 , but every other aleph number can be realized by applying the power set construction a bunch of times. Some people would express this symbolically by writing

$$\forall n \in \mathbb{Z}^{\text{nonneg}}, \quad \aleph_{n+1} = 2^{\aleph_n}.$$

I’d really rather not bring this chapter to a close with that monstrosity so instead I think I’ll just say

Paul Cohen.

Hah! I did it! I ended the chapter by sayi... Hunh? Oh.

Paul Cohen.

Chapter 9

Proof techniques IV — Magic

If you can keep your head when all about you are losing theirs, it's just possible you haven't grasped the situation. –Jean Kerr

The famous mathematician Paul Erdős is said to have believed that God has a Book in which all the really elegant proofs are written. The greatest praise that a collaborator¹ could receive from Erdős was that they had discovered a “Book proof.” It is not easy or straightforward for a mere mortal to come up with a Book proof but notice that, since the Book is inaccessible to the living, all the Book proofs of which we are aware were constructed by ordinary human beings. In other words, it's not impossible!

The title of this final chapter is intended to be whimsical – there is no real magic involved in any of the arguments that we'll look at. Nevertheless, if you reflect a bit on the mental processes that must have gone into the development of these elegant proofs, perhaps you'll agree that there is something magical there.

At a minimum we hope that you'll agree that they are beautiful – they

¹The collaborators of Paul Erdős were legion. His collaborators, and their collaborators, and *their* collaborators, etc. are organized into a tree structure according to their so-called Erdős number, see [5].

are proofs from the Book².

Acknowledgment: Several of the topics in this section were unknown to the author until he visited the excellent mathematics website maintained by Alexander Bogomolny at

<http://www.cut-the-knot.org/>

²There is a lovely book entitled “Proofs from the Book” [2] that has a nice collection of Book proofs.

9.1 Morley's miracle

Probably you have heard of the impossibility of trisecting an angle. (Hold on for a quick rant about the importance of understanding your hypotheses. . .) What's *actually* true is that you can't trisect a generic angle if you accept the restriction of using the old-fashioned tools of Euclidean geometry: the compass and straight-edge. There are a lot of constructions that can't be done using just a straight-edge and compass – angle trisection, duplication of a cube³, squaring a circle, constructing a regular heptagon, *et cetera*.

If you allow yourself to use a *ruler* – i.e. a straight-edge with marks on it (indeed you really only need two marks a unit distance apart) then angle trisection *can* be done via what is known as a neusis construction. Nevertheless, because of the central place of Euclid's *Elements* in mathematical training throughout the centuries, and thereby, a very strong predilection towards that which *is* possible via compass and straight-edge alone, it is perhaps not surprising that a perfectly beautiful result that involved trisecting angles went undiscovered until 1899, when Frank Morley stated his Trisector Theorem. There is much more to this result than we will state here – so much more that the name “Morley's Miracle” that has been given to the Trisector theorem is truly justified – but even the simple, initial part of this beautiful theory is arguably miraculous! To learn more about Morley's theorem and its extension see [8].

So, let's state the theorem!

³Duplicating the cube is also known as the Delian problem – the problem comes from a pronouncement by the oracle of Apollo at Delos that a plague afflicting the Athenians would be lifted if they built an altar to Apollo that was twice as big as the existing altar. The existing altar was a cube, one meter on a side, so they carefully built a two meter cube – but the plague raged on. Apparently what Apollo wanted was a cube that had double the *volume* of the present altar – its side length would have to be $\sqrt[3]{2} \approx 1.25992$ and since this was Greece and it was around 430 B.C. and there were no electronic calculators, they were basically just screwed.

Start with an arbitrary triangle $\triangle ABC$. Trisect each of its angles to obtain a diagram something like that in Figure 9.1.

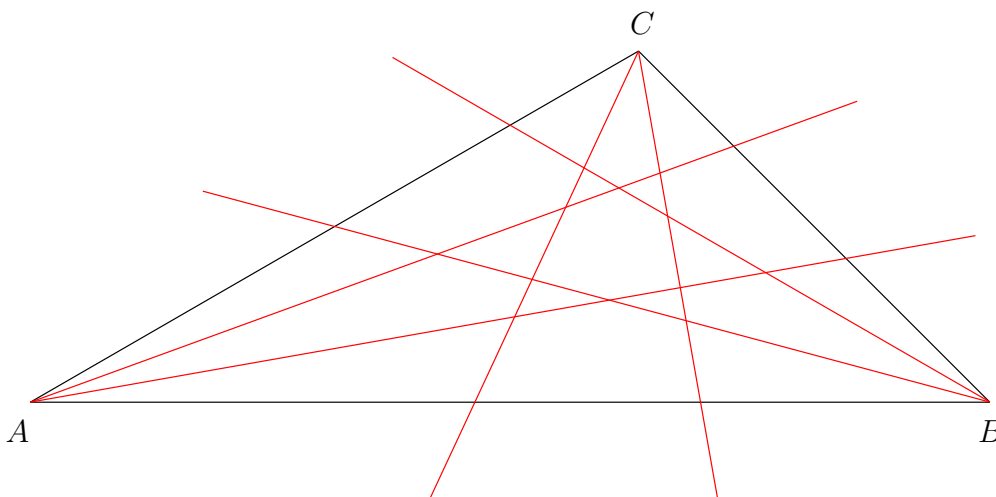


Figure 9.1: The setup for Morley’s Miracle – start with an arbitrary triangle and trisect each of its angles.

The six angle trisectors that we’ve just drawn intersect one another in quite a few points.

Exercise. *You could literally count the number of intersection points between the angle trisectors on the diagram, but you should also be able to count them (perhaps we should say “double-count them”) combinatorially. Give it a try!*

Among the points of intersection of the angle trisectors there are three that we will single out – the intersections of adjacent trisectors. In Figure 9.2 the intersection of adjacent trisectors are indicated, additionally, we have connected them together to form a small triangle in the center of our original triangle.

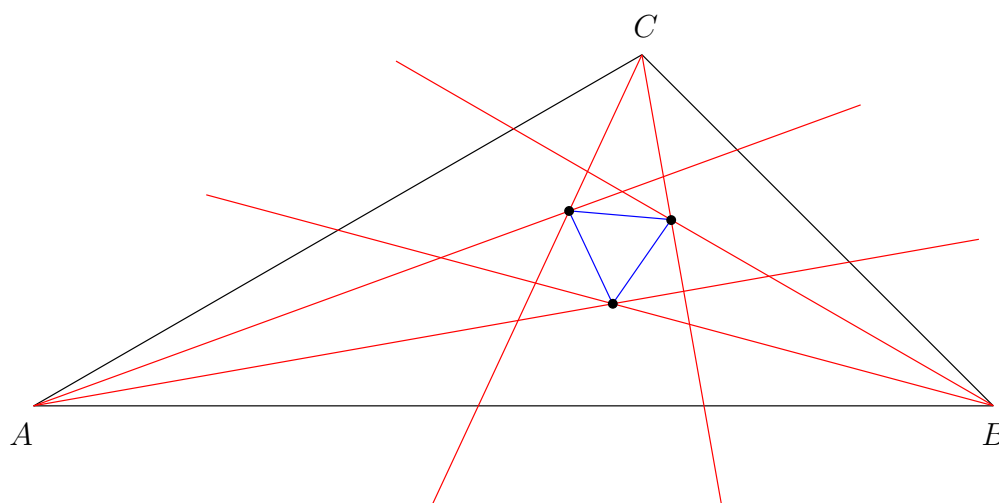


Figure 9.2: A triangle is formed whose vertices are the intersections of the adjacent trisectors of the angles of $\triangle ABC$.

Are you ready for the miraculous part? Okay, here goes!

Theorem 9.1.1. *The points of intersection of the adjacent trisectors in an arbitrary triangle $\triangle ABC$ form the vertices of an equilateral triangle.*

In other words, that little blue triangle in Figure 9.2 that kind of *looks* like it might be equilateral actually does have all three sides equal to one another. Furthermore, it doesn't matter what triangle we start with, if we do the construction above we'll get a perfect $60^\circ - 60^\circ - 60^\circ$ triangle in the middle!

Sources differ, but it is not clear whether Morley ever proved his theorem. The first valid proof (according to R. K. Guy in [8] was published in 1909 by M. Satyanarayana [15]. There are now *many* other proofs known, for instance the cut-the-knot website (<http://www.cut-the-knot.org/>) exposts no less than nine different proofs. The proof by Satyanarayana used trigonometry. The proof we'll look at here is arguably the shortest ever produced and it is due to John Conway. It is definitely a "Book proof"!

Let us suppose that an arbitrary triangle $\triangle ABC$ is given. We want to show that the triangle whose vertices are the intersections of the adjacent trisectors is equilateral – this triangle will be referred to as the *Morley triangle*. Let's also denote by A , B and C the measures of the angles of $\triangle ABC$. (This is what is generally known as an “abuse of notation” – we are intentionally confounding the vertices (A , B and C) of the triangle with the measure of the angles at those vertices.) It turns out that it is fairly hard to reason from our knowledge of what the angles A , B and C are to deduce that the Morley triangle is equilateral. How does the following plan sound: suppose we construct a triangle, that definitely *does* have an equilateral Morley triangle, whose angles also happen to be A , B and C . Such a triangle would be similar⁴ to the original triangle $\triangle ABC$ – if we follow the similarity transform from the constructed triangle back to $\triangle ABC$ we will see that their Morley triangles must coincide; thus if one is equilateral so is the other!

One of the features of Conway's proof that leads to its great succinctness and beauty is his introduction of some very nice notation. Since we are dealing with angle trisectors, let a , b and c be angles such that $3a = A$, $3b = B$ and $3c = C$. Furthermore, let a superscript star denote the angle that is $\pi/3$ (or 60° if you prefer) greater than a given angle. So, for example,

$$a^* = a + \pi/3$$

and

$$a^{**} = a + 2\pi/3.$$

Now, notice that the sum $a + b + c$ must be $\pi/3$. This is an immediate consequence of $A + B + C = \pi$ which is true for any triangle in the plane. It

⁴In Geometry, two objects are said to be *similar* if one can be made to exactly coincide with the other after a series of rigid translations, rotations and scalings. In other words, they have the same shape if you allow for differences in scale and are allowed to slide them around and spin them about as needed.

follows that by distributing two stars amongst the three numbers a , b and c we will come up with three quantities which sum to π . In other words, there are Euclidean triangles having the following triples as their vertex angles:

$$\begin{array}{ll} (a, b, c^{**}) & (a, b^*, c^*) \\ (a, b^{**}, c) & (a^*, b^*, c) \\ (a^{**}, b, c) & (a^*, b, c^*) \end{array}$$

Exercise. What would a triangle whose vertex angles are $(0^*, 0^*, 0^*)$ be?

In a nutshell, Conway's proof consists of starting with an equilateral triangle of unit side length, adding appropriately scaled versions of the six triangles above and ending up with a figure (having an equilateral Morley triangle) similar to $\triangle ABC$. The generic picture is given in Figure 9.3. Before we can really count this argument as a proof, we need to say a bit more about what the phrase "appropriately scaled" means. In order to appropriately scale the triangles (the small acute ones) that appear green in Figure 9.3 we have a relatively easy job – just scale them so that the side opposite the trisected angle has length one; that way they will join perfectly with the central equilateral triangle.

The triangles (these are the larger obtuse ones) that appear purple in 9.3 are a bit more puzzling. Ostensibly, we have two different jobs to accomplish – we must scale them so that both of the edges that they will share with green triangles have the correct lengths. How do we know that this won't require two different scaling factors? Conway also developed an elegant argument that handles this question as well. Consider the purple triangle at the bottom of the diagram in Figure 9.3 – it has vertex angles (a, b, c^{**}) . It is possible to construct triangles similar (via reflections) to the adjacent green triangles (a, b^*, c^*) and (a^*, b, c^*) *inside* of triangle (a, b, c^{**}) . To do this just construct two lines that go through the top vertex (where the angle c^{**} is) that cut the opposite edge at the angle c^* in the two possible senses – these two lines will

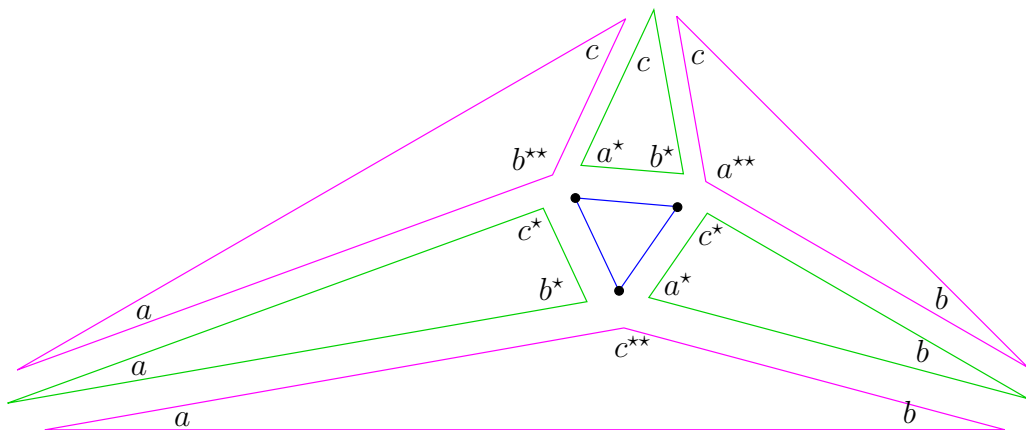


Figure 9.3: Conway's proof involves putting these pieces together to obtain a triangle (with an equilateral Morley triangle) that is similar to $\triangle ABC$.

coincide if it should happen that c^* is precisely $\pi/2$ but generally there will be two and it is evident that the two line segments formed have the same length. We scale the purple triangle so that this common length will be 1. See Figure 9.4.

Exercise. *If it should happen that $c^* = \pi/2$, what can we say about C ?*


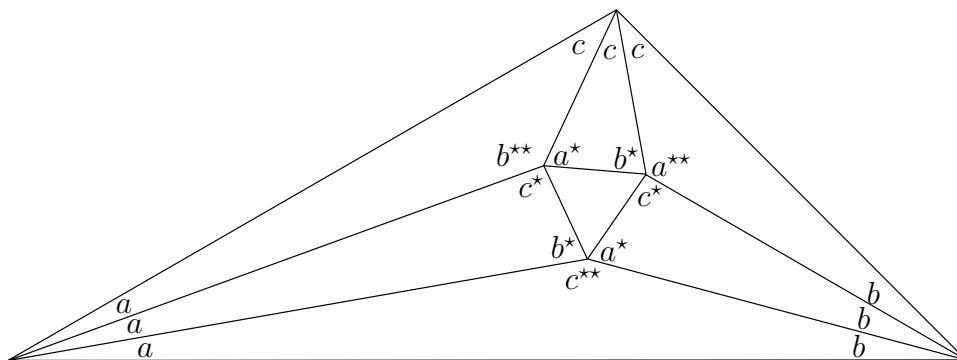


Figure 9.4: The scaling factor for the obtuse triangles in Conway's puzzle proof is determined so that the segments constructed in their midpoints have unit length.

Of course the other two obtuse triangles can be handled in a similar way.

Exercises — 9.1

1. What value should we get if we sum all of the angles that appear around one of the interior vertices in the finished diagram? Verify that all three have the correct sum.



2. In this section we talked about similarity. Two figures in the plane are similar if it is possible to turn one into the other by a sequence of mappings: a translation, a rotation and a scaling.

Geometric similarity is an equivalence relation. To fix our notation, let $T(x, y)$ represent a generic translation, $R(x, y)$ a rotation and $S(x, y)$ a scaling – thus a generic similarity is a function from \mathbb{R}^2 to \mathbb{R}^2 that can be written in the form $S(R(T(x, y)))$.

Discuss the three properties of an equivalence relation (reflexivity, symmetry and transitivity) in terms of geometric similarity.

9.2 Five steps into the void

In this section we'll talk about another Book proof also due to John Conway. This proof serves as an introduction to a really powerful general technique – the idea of an invariant. An invariant is some sort of quantity that one can calculate that itself doesn't change as other things are changed. Of course different situations have different invariant quantities.

The setup here is simple and relatively intuitive. We have a bunch of checkers on a checkerboard – in fact we have an infinite number of checkers, but not filling up the whole board, they completely fill an infinite half-plane which we could take to be the set

$$S = \{(x, y) \mid x \in \mathbb{Z} \wedge y \in \mathbb{Z} \wedge y \leq 0\}.$$

See Figure 9.5.

Think of these checkers as an army and the upper half-plane is “enemy territory.” Our goal is to move one of our soldiers into enemy territory as far as possible. The problem is that our “soldiers” move the way checkers do, by jumping over another man (who is then removed from the board). It's clear that we can get someone into enemy territory – just take someone in the second row and jump a guy in the first row. It is also easy enough to see that it is possible to get a man two steps into enemy territory – we could bring two adjacent men a single step into enemy territory, have one of them jump the other and then a man from the front rank can jump over him.

Exercise. *The strategy just stated uses 4 men (in the sense that they are removed from the board – 5 if you count the one who ends up two steps into enemy territory as well). Find a strategy for moving someone two steps into enemy territory that is more efficient – that is, involves fewer jumps.*

Exercise. *Determine the most efficient way to get a man three steps into enemy territory. An actual checkers board and pieces (or some coins, or rocks) might come in handy.*

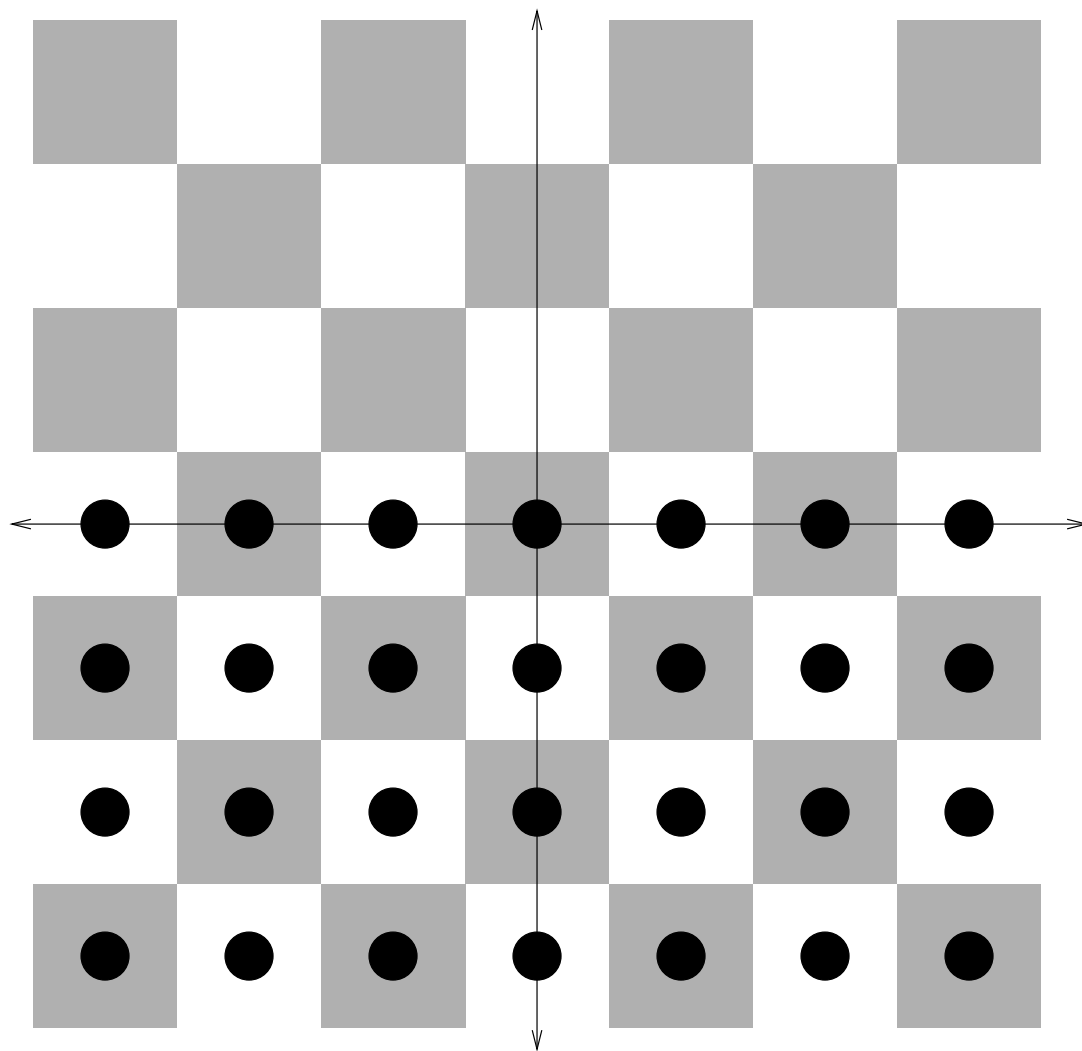


Figure 9.5: An infinite number of checkers occupying the integer lattice points such that $y \leq 0$.

We'll count the man who ends up some number of steps above the x -axis, as well as all the pieces who get jumped and removed from the board as a measure of the efficiency of a strategy. If you did the last exercise correctly you should have found that eight men are the minimum required to get 3 steps into enemy territory. So far, the number of men required to get a given distance into enemy territory seems to always be a power of 2.

# of steps	# of men
1	2
2	4
3	8

As a picture is sometimes literally worth one thousand words, we include here 3 figures illustrating the moves necessary to put a scout 1, 2 and 3 steps into the void.

In order to show that 8 men are sufficient to get a scout 3 steps into enemy territory, we show that it is possible to reproduce the configuration that can place a man two steps in – shifted up by one unit.

You may be surprised to learn that the pattern of 8 men which are needed to get someone three steps into the void can be re-created – shifted up by one unit – using just 12 men. This means that we can get a man 4 steps into enemy territory using $12 + 8 = 20$ men. You were expecting 16 weren't you? (I know *I* was!)

Exercise. *Determine how to get a marker 4 steps into the void.*

The *real* surprise is that it is simply impossible to get a man five steps into enemy territory. So the sequence we've been looking at actually goes

$$2, 4, 8, 20, \infty.$$

The proof of this surprising result works by using a fairly simple, but clever, strategy. We assign a numerical value to a set of men that is dependent on their positions – then we show that this value never increases when we

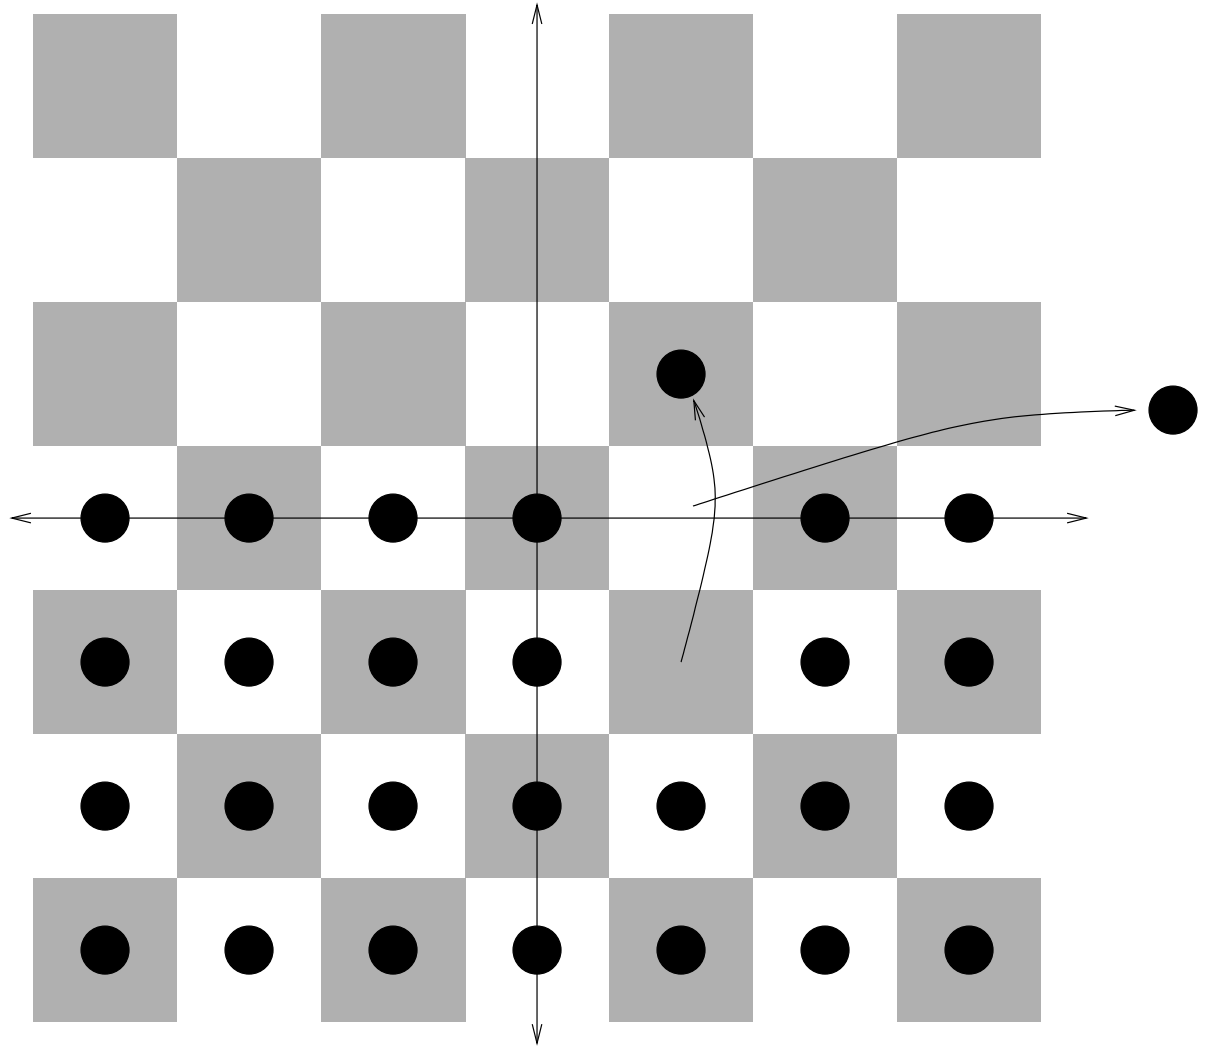


Figure 9.6: One man is sacrificed in order to move a scout one step into enemy territory.

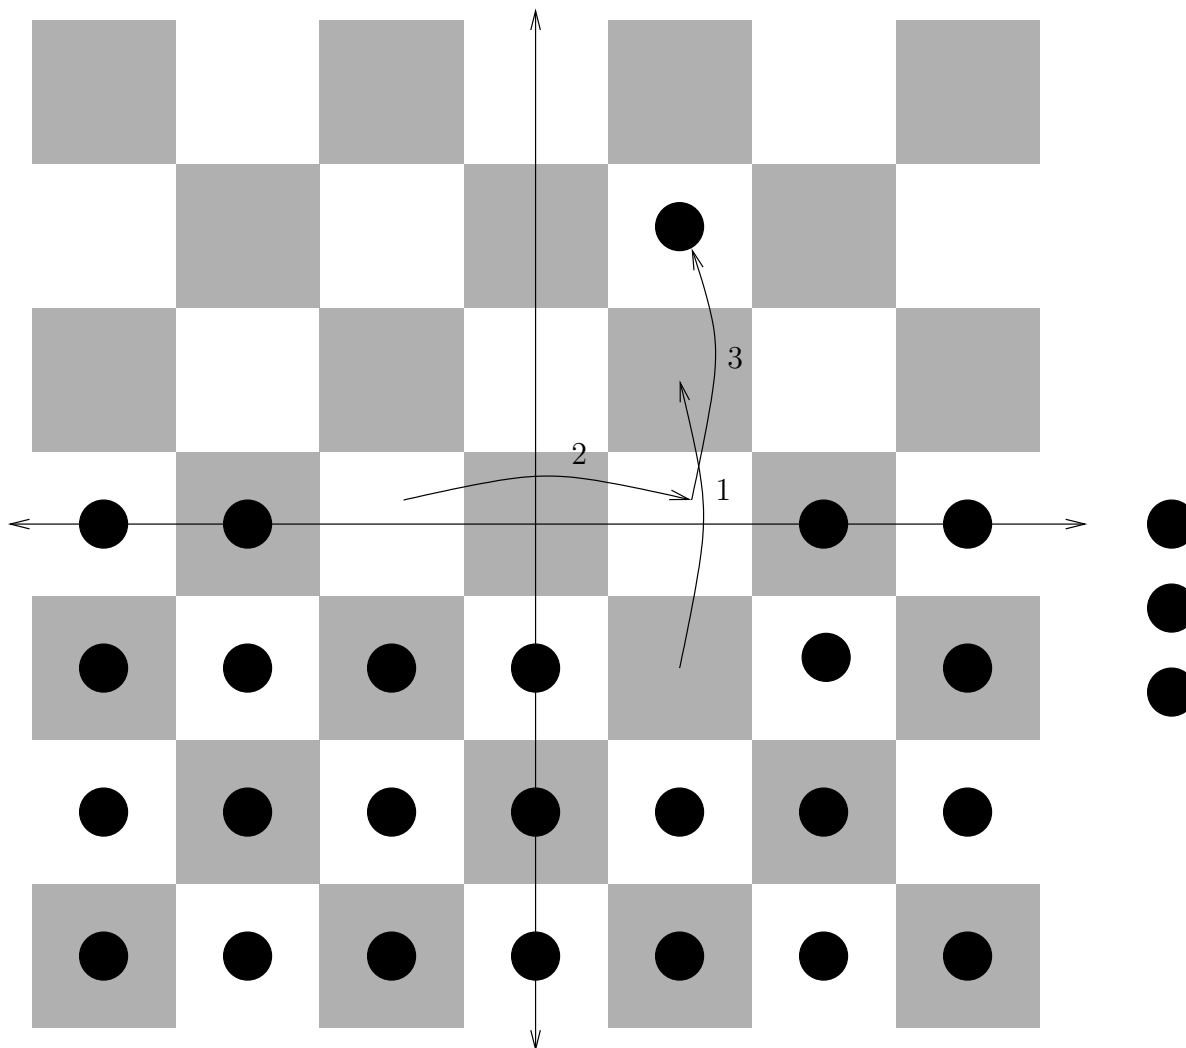


Figure 9.7: Three men are sacrificed in order to move a scout two steps into enemy territory.

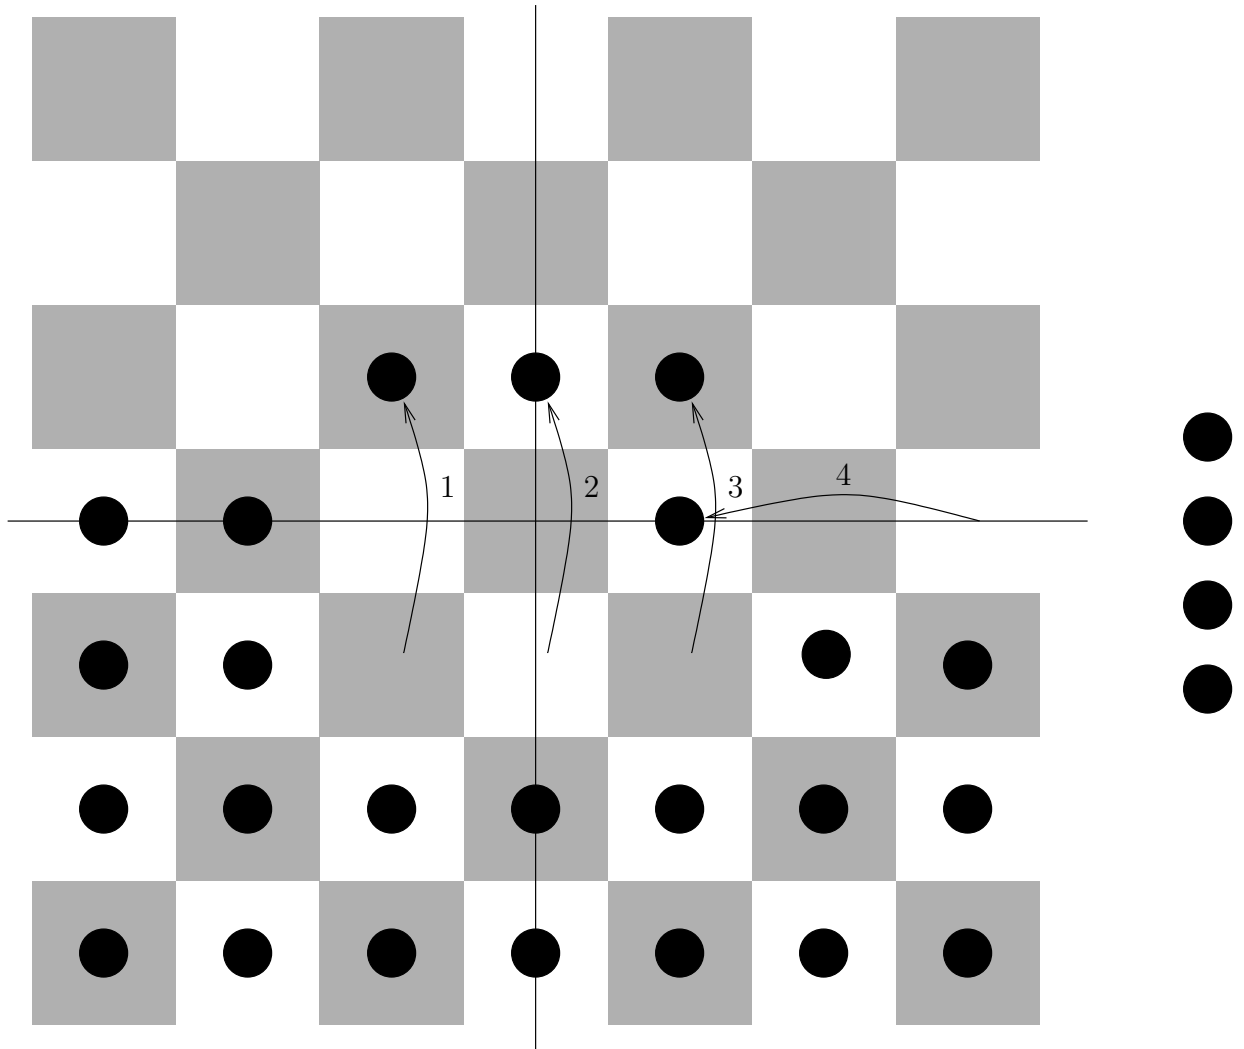


Figure 9.8: Eight men are needed to get a scout 3 steps into the void.

make “checker jumping” moves – finally we note that the value assigned to a man in position $(0, 5)$ is equal to the value of the entire original set of men (that is, with *all* the positions in the lower half-plane occupied). This is a pretty nice strategy, but how exactly are we going to assign these numerical values?


A man’s value is related to his distance from the point $(0, 5)$ in what is often called “the taxicab metric.” We don’t use the straight-line distance, but rather determine the number of blocks we will have to drive in the north-south direction and in the east-west direction and add them together. The value of a set of men is the sum of their individual values. Since we need to deal with the value of the set of men that completely fills the lower half-plane, we are going to have to have most of these values be pretty tiny! To put it in a more mature and dignified manner: the infinite sum of the values of the men in our army must be convergent.

We’ve previously seen geometric series which have convergent sums. Recall the formula for such a sum is

$$\sum_{k=0}^{\infty} ar^k = \frac{a}{1-r},$$

where a is the initial term of the sum and r is the common ratio between terms.

Conway’s big insight was to associate the powers of some number r with the positions on the board – r^k goes on the squares that are distance k from the target location. If we have a man who is actually *at* the target location, he will be worth r^0 or 1. We need to arrange for two things to happen: the sum of all the powers of r in the initial setup of the board must be less than or equal to 1, and checker-jumping moves should result in the total value of a set of men going down or (at worst) staying the same. These goals push us in different directions: In order for the initial sum to be less than 1, we would like to choose r to be fairly small. In order to have checker-jumping

Figure 9.9: The taxicab distance to $(0, 5)$.

moves we need to choose r to be (relatively) larger. Is there a value of r that does the trick? Can we find a balance between these competing desires?

Think about the change in the value of our invariant as a checker jumping move gets made. See Figure 9.10.

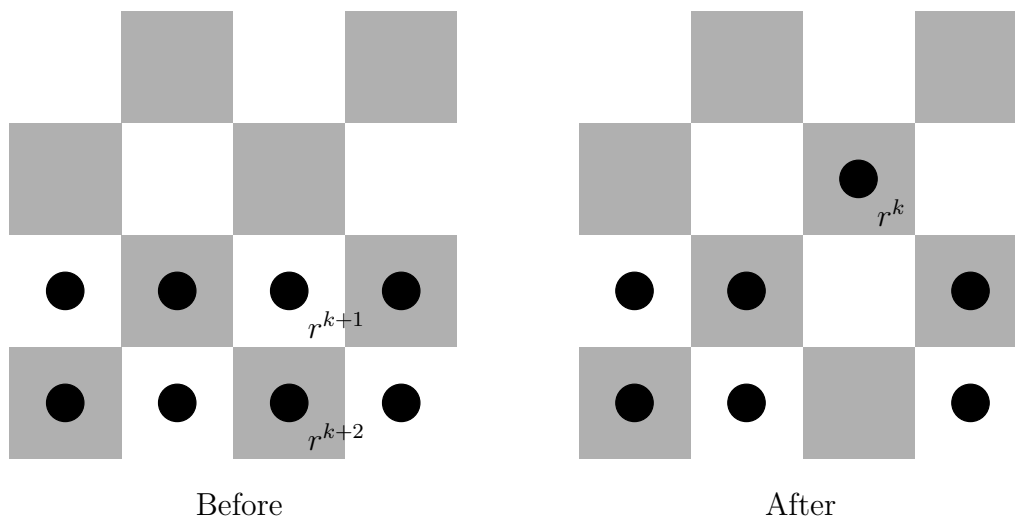


Figure 9.10: In making a checker-jumping move, two men valued r^{k+1} and r^{k+2} are replaced by a single man valued r^k .

If we choose r so that $r^{k+2} + r^{k+1} \leq r^k$ then the checker-jumping move will at worst leave the total sum fixed. Note that so long as $r < 1$ a checker-jumping move that takes us away from the target position will certainly *decrease* the total sum.

As is often the case, we'll analyze the inequality by looking instead at the corresponding equality. What value of r makes $r^{k+2} + r^{k+1} = r^k$? The answer is that r must be a root of the quadratic equation $x^2 + x - 1$.

Exercise. *Do the algebra to verify the previous assertion.*

Exercise. *Find the value of r that solves the above equation.*

Hopefully you used the quadratic formula to solve the previous exercise. You should of course have found two solutions, $-1.618033989\dots$ and

.618033989... , these decimal approximations are actually $-\phi$ and $1/\phi$, where $\phi = \frac{1 + \sqrt{5}}{2}$ is the famous “golden ratio”. If we are hoping for the sum over all the occupied positions of r^k to be convergent, we need $|r| < 1$, so the negative solution is extraneous and so the inequality $r^{k+2} + r^{k+1} \leq r^k$ is true in the interval $[1/\phi, 1)$.

Next we want to look at the value of this invariant when “men” occupy all of the positions with $y \leq 0$. By looking at Figure 9.9 you can see that there is a single square with value r^5 , there are 3 squares with value r^6 , there are 5 squares with value r^7 , *et cetera*. The sum, S , of the values of all the initially occupied positions is

$$S = r^5 \cdot \sum_{k=0}^{\infty} (2k+1)r^k.$$

We have previously seen how to solve for the value of an infinite sum involving powers of r . In the expression above we have powers of r but also multiplied by odd numbers. Can we solve something like this?

Let’s try the same trick that works for a geometric sum. Let

$$T = \sum_{k=0}^{\infty} (2k+1)r^k = 1 + 3r + 5r^2 + 7r^3 + \dots$$

Note that

$$rT = \sum_{k=0}^{\infty} (2k+1)r^{k+1} = r + 3r^2 + 5r^3 + 7r^4 + \dots$$

and it follows that

$$T - rT = 1 + 2 \sum_{k=1}^{\infty} r^k = 1 + 2r + 2r^2 + 2r^3 + 2r^4 + \dots$$

A bit more algebra (and the formula for the sum of a geometric series) leads us to

$$T = \frac{1}{1-r} \left(1 + \frac{2r}{1-r} \right),$$

which simplifies to

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

Finally, recall that we are really interested in $S = r^5 \cdot T$, or

$$S = \frac{r^5 + r^6}{(1-r)^2}.$$

It is interesting to proceed from this expression for S , using the fact that r satisfies $x^2 = 1 - x$, to obtain the somewhat amazing fact that $S = 1$.

The fact that $S = 1$ has an extraordinary consequence. In order to get a single checker to the position $(0, 5)$ we would need to use *everybody!*

For a set consisting of just a single checker positioned at $(0, 5)$ the value of our invariant is 1. On the other hand, the set consisting of the entire army lined up on and below the x -axis also yields a 1. Every checker move either does not change the value of the invariant or reduces it. The best we could possibly hope for is that there would be no need for moves of the sort that reduce the invariant – nevertheless we still could not get a man to $(0, 5)$ in a finite number of moves.

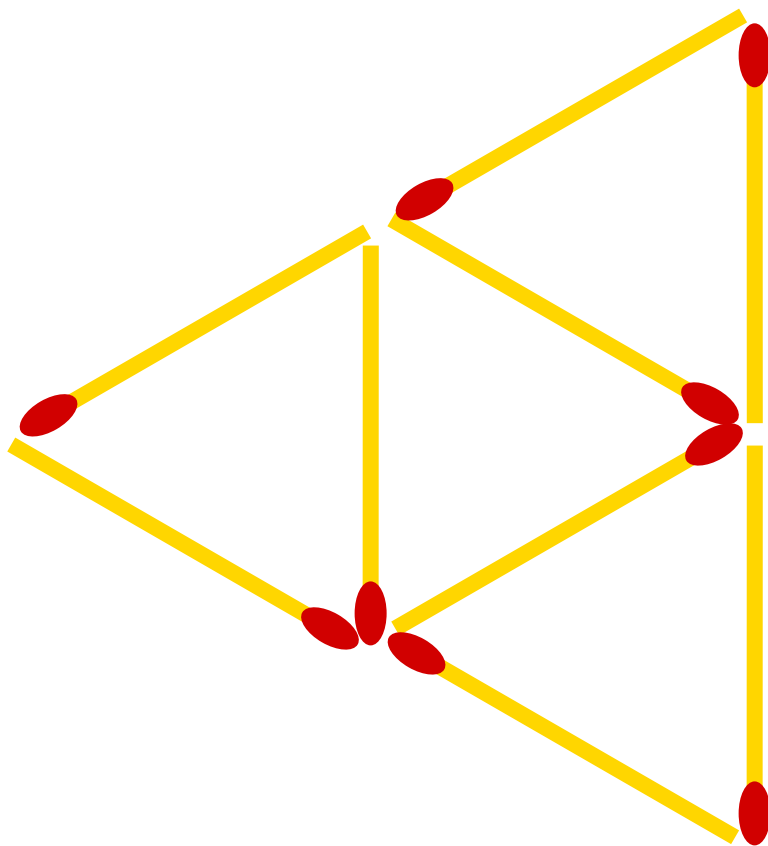
Exercises — 9.2

1. Do the algebra (and show all your work!) to prove that invariant defined in this section actually has the value 1 for the set of all the men occupying the x -axis and the lower half-plane.
2. “Escape of the clones” is a nice puzzle, originally proposed by Maxim Kontsevich. The game is played on an infinite checkerboard restricted to the first quadrant – that is the squares may be identified with points having integer coordinates (x, y) with $x > 0$ and $y > 0$. The “clones” are markers (checkers, coins, small rocks, whatever...) that can move in only one fashion – if the squares immediately above and to the right of a clone are empty, then it can make a “clone move.” The clone moves one space up and a copy is placed in the cell one to the right. We begin with three clones occupying cells $(1, 1)$, $(2, 1)$ and $(1, 2)$ – we’ll refer to those three checkerboard squares as “the prison.” The question is this: can these three clones escape the prison?

You must either demonstrate a sequence of moves that frees all three clones or provide an argument that the task is impossible.

9.3 Monge’s circle theorem

There’s a nice sequence of matchstick puzzles that starts with “Use nine non-overlapping matchsticks to form 4 triangles (all of the same size.” It’s not that hard, and after a while most people come up with



The kicker comes when you next ask them to “use six matches to form 4 (equal sized) triangles.” There’s a picture of the solution to this new puzzle at the back of this section. The answer involves thinking three-dimensionally, so – with that hint – give it a try for a while before looking in the back.

Monge’s circle theorem has nothing to do with matchsticks, but it is a *sweet* example of a proof that works by moving to a higher dimension. People often talk about “thinking outside of the box” when discussing critical

thinking, but the mathematical idea of moving to a higher dimension is even more powerful. When we have a “box” in 2-dimensional space which we then regard as sitting in a 3-dimensional space we find that the box doesn't even *have* an inside or an outside anymore! We get “outside the box” by literally erasing the notion that there *is* an inside of the box!

The setup for Monge's circle theorem consists of three random circles drawn in the plane. Well, to be honest they can't be entirely random – we can't allow a circle that is entirely inside another circle. Because, if a circle was entirely inside another, there would be no external tangents and Monge's circle theorem is about external tangents.

I could probably write a few hundred words to explain the concept of external tangents to a pair of circles, or you could just have a look at Figure 9.11. So, uhmm, just have a look. . .

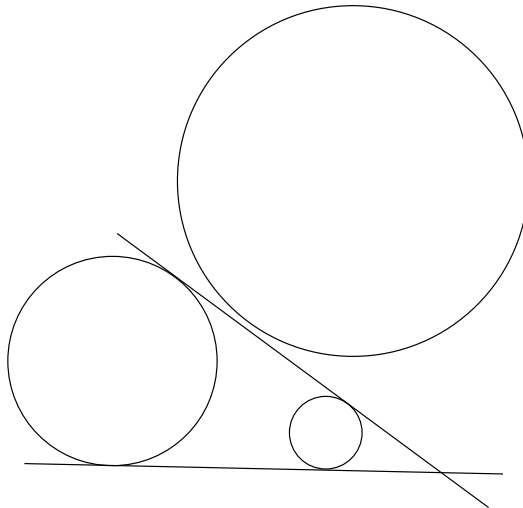


Figure 9.11: The setup for Monge's circle theorem: three randomly placed circles – we are also showing the external tangents to one pair of circles.

Notice how the external tangents⁵ to two of the circles meet in a point? Unless the circles just happen to have exactly the same size (And what are the odds of that?) this is going to be the case. Each pair of external tangents are going to meet in a point. There are three such pairs of external tangents and so they determine three points. I suppose, since these three points are determined in a fairly complicated way from three randomly chosen circles, that we would expect the three points to be pretty much random. Monge's circle theorem says that that isn't so.

Theorem 9.3.1 (Monge's Circle Theorem). *If three circles of different radii in the Euclidean plane are chosen so that no circle lies in the interior of another, the three pairs of external tangents to these circles meet in points which are collinear.*

In Figure 9.12 we see a complete example of Monge's Circle theorem in action. There are three random circles. There are three pairs of external tangents. The three points determined by the intersection of the pairs of external tangents lie on a line (shown dashed in the figure).

We won't even try to write-up a formal proof of the circle theorem. Not that it can't be done – it's just that you can probably get the point better via an informal discussion.

The main idea is simply to move to 3-dimensional space. Imagine the original flat plane containing our three random circles as being the plane $z = 0$ in Euclidean 3-space. Replace the three circles by three spheres of the same radius and having the same centers – clearly the intersections of these spheres with the plane $z = 0$ will be our original circles. While pairs of circles are encompassed by two lines (the external tangents that we've been discussing so much), when we have a pair of spheres in 3-space, they are encompassed by a cone which lies tangent to both spheres⁶. Notice that the

⁵The reason I keep saying “external tangents” is that there are also *internal* tangents.

⁶As before, when the spheres happen to have identical radii we get a degenerate case – the cone becomes a cylinder.

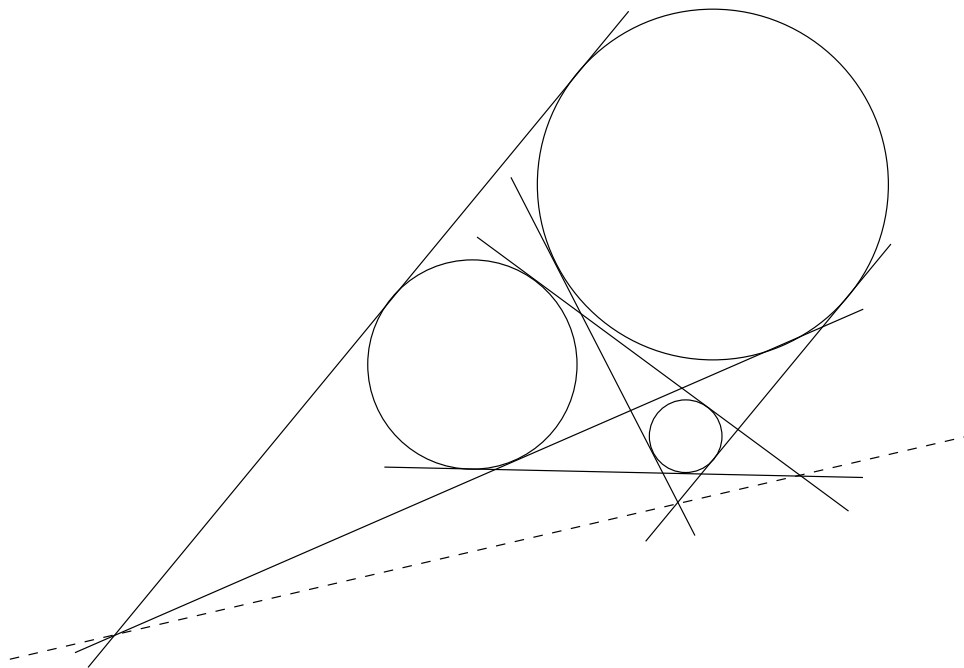


Figure 9.12: An example of Monge's circle theorem. The three pairs of external tangents to the circles intersect in points which are collinear.

cones that lie tangent to a pair of spheres intersect the plane precisely in those infamous external tangents.

Well, okay, we've moved to 3-d. We've replaced our circles with spheres and our external tangents with tangent cones. The points of intersection of the external tangents are now the tips of the cones. But, what good has this all done? Is there any reason to believe that the tips of those cones lie in a line?

Actually, yes! There is a plane that touches all three spheres tangentially. Actually, there are two such planes, one that touches them all on their upper surfaces and one that touches them all on their lower surfaces. Oh damn! There are actually *lots* of planes that are tangent to all three spheres but only one that lies above the three of them. That plane intersects the plane $z = 0$ in a line – nothing fancy there; any pair of non-parallel planes will intersect in a line (and the only way the planes we are discussing would be parallel is if all three spheres just happened to be the same size). But that plane also lies tangent to the cones that envelope our spheres and so that plane (as well as the plane $z = 0$) contains the tips of the cones!

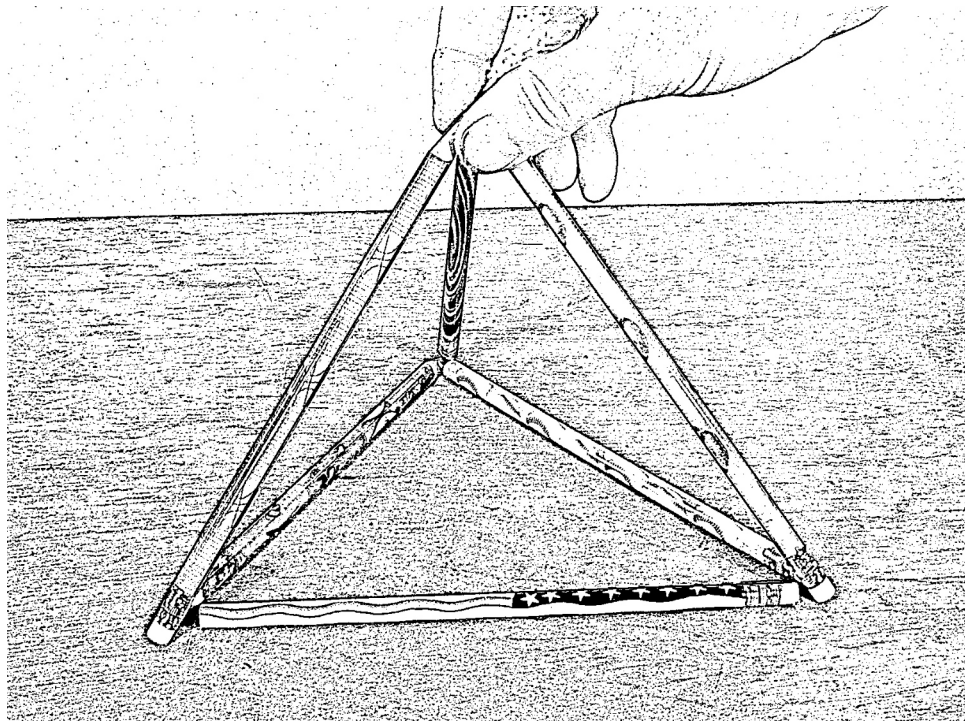


Figure 9.13: Six matchstick (actually, pencils are a lot easier to hold) can be arranged three-dimensionally to create four triangles.

Exercises — 9.3

1. There is a scenario where the proof we have sketched for Monge's circle theorem doesn't really work. Can you envision it? Hint: consider two relatively large spheres and one that is quite small.

Bibliography

- [1] R. E. Greenwood A. M. Gleason and L. M. Kelly. *The William Lowell Putnam Mathematical Competition Problems & Solutions: 1938-1964*. The Mathematical Association of America, Reissued 2003.
- [2] Martin Aigner and Gunter M. Ziegler. *Proofs from THE BOOK*. Springer-Verlag, 2nd edition, 2001.
- [3] Wikipedia contributors. Cantor-bernstein-schroeder theorem. Wikipedia, the free encyclopedia. http://en.wikipedia.org/wiki/Cantor-Bernstein-Schroeder_theorem.
- [4] Wikipedia contributors. Christian Goldbach. Wikipedia, the free encyclopedia. <http://en.wikipedia.org/wiki/Goldbach>.
- [5] Wikipedia contributors. Erdos number. Wikipedia, the free encyclopedia. http://en.wikipedia.org/wiki/Erdos_number.
- [6] Wikipedia contributors. The four color theorem. Wikipedia, the free encyclopedia. http://en.wikipedia.org/wiki/Four_color_theorem.
- [7] Leonard F. Klosinski Gerald L. Alexanderson and Loren C. Larson. *The William Lowell Putnam Mathematical Competitions Problems & Solutions: 1965 - 1984*. The Mathematical Association of America, Reissued 2003.

- [8] Richard K. Guy. The lighthouse theorem, Morley & Malfatti – a budget of paradoxes. *American Mathematical Monthly*, 2007.
- [9] Bjorn Poonen Kiran S. Kedlaya and Ravi Vakil. *The William Lowell Putnam Mathematical Competition 1985-2000: Problems Solutions, and Commentary*. The Mathematical Association of America, 2002.
- [10] C. W. H. Lam. The search for a finite projective plane of order 10. <http://www.cec.m.sfu.ca/organics/papers/lam/paper/html/paper.html>.
- [11] Saunders MacLane. *Categories for the Working Mathematician*. Springer-Verlag, 2nd edition, 1998.
- [12] John J. O'Connor and Edmund F. Robertson. <http://www-history.mcs.st-andrews.ac.uk/history/index.html>. The MacTutor History of Mathematics archive.
- [13] Stanislaw Radziszowski. Small ramsey numbers. <http://www.combinatorics.org/Surveys/ds1.pdf>.
- [14] Gian-Carlo Rota. *Indiscrete Thoughts*. Birkhäuser, 1997.
- [15] M. Satyanarayana. none given. *Math. Quest. Educ. Times (New Series)*, 1909.
- [16] D. J. Struik. *A Source Book in Mathematics, 1200-1800*. Princeton University Press, 1986.
- [17] Alfred North Whitehead and Bertrand Russell. *Principia Mathematica*. Cambridge University Press, 1910.

Index

- absorption, 82, 83
- addition rule, 295
- Aleph-naught, 342
- algorithm, 36
- alphanumeric, 309
- and gates, 62
- antecedent, 68
- anti-symmetry, 243
- antichain, 264
- Appel, Kenneth, 149
- arithmetic mean, 132
- arithmetic sequence, 132
- arithmetic-geometric mean inequality, 132
- ASCII, 310
- associative law, 59, 76, 83
- atomic concepts, 55

- begging the question, 89
- biconditional, 70
- bijection, 275
- binary relation, 51
- binomial, 328
- binomial coefficients, 29
- Blaise Pascal, 30

- Boole, George, 59
- bottom, in a poset, 266
- bound variables, 93

- Cantor's Snake, 348
- cardinality, 170
- Cartesian product, 236
- Category theory, 165
- ceiling function, 27
- chain, 264
- characteristic function, 287
- Chung, Fan, 151
- circular reasoning, 89
- Cohen, Paul, 379
- combination, 298
- commutative law, 59, 75, 83
- complement, 178
- complementarity law, 83
- complementarity laws, 81
- complex numbers, 5
- component-wise operations, 6
- composite, 14
- composition, of functions, 240
- composition, of relations, 240
- compound sentence, 57

- conditional statement, 68
- congruence, 28
- conjunction, 57
- conjunctive simplification, 103
- consequent, 68
- constructive dilemma, 104
- contradiction, 81
- contrapositive, 71
- converse, 70
- converse error, 112
- Conway, John, 385
- countably infinite, 343
- counterexample, 143
- cover, in a poset, 266
- Crocodile's dilemma, 104

- deduction, 101
- deductive argument, 102
- degree, 313
- DeMorgan's law, 83
- DeMorgan's laws, 80
- denumerable, 343
- Descartes, Rene, 237
- destructive dilemma, 105, 111
- diagonal map, 287
- difference (of sets), 182
- digraph, 164, 236, 255
- direct proofs, 125
- directed graph, 236
- disjunction, 58
- disjunctive normal form, 64, 194

- disjunctive syllogism, 104
- distributive law, 59, 78, 83
- divisibility, 26
- division algorithm, 38
- domain, 53
- domination law, 83
- doubly-even, 32
- duck, flaming, 113

- empty set, 169
- empty sum, 213
- equinumerous, 339
- equivalence class, 251
- equivalence relation, 251
- Eratosthenes of Cyrene, 13
- Erdos number, 381
- Erdos, Paul, 381
- error detecting code, 312
- Euclidean algorithm, 41
- Euler, Leonhard, 95, 150, 312
- Eulerian circuit, 320
- Eulerian path, 320
- evenness, 130
- exclusive or, 60
- existential quantification, 20

- factorials, 30
- Fermat numbers, 94
- Fermat's last theorem, 222
- Fermat's little theorem, 222
- Fermat, Pierre de, 94

- Fibonacci numbers, 221
- finite sequence, 294
- floor function, 26
- flowchart, 36
- form (of an argument), 112
- forwards-backwards method, 132
- four color theorem, 148

- Gödel, Kurt, 56, 166
- general position, 193
- generalizing from the generic particular, 125
- geometric mean, 133
- geometric sequence, 132
- Goldbach's conjecture, 19, 149
- Goldbach, Christian, 149
- golden ratio, 399
- graded poset, 263
- graph, 255
- graph isomorphism, 256
- graph pebbling, 151
- greatest common divisor, 40
- greatest common divisor, gcd, 159
- greatest element, in a poset, 266

- Haken, Wolfgang, 149
- Hasse diagrams, 262
- Hasse, Helmut, 262
- Hollerith card, 309
- hypercube, 268
- hypotheses, 102

- hypothetical syllogism, 104

- idempotence, 83
- idempotent, 81
- identity law, 83
- identity laws, 81
- iff, 70
- image, of a set, 279
- imaginary part, 6
- inclusive or, 60
- Incompleteness Theorem, 56
- indicator function, 287
- indirect proof, 137
- induction, 55, 201
- inductive argument, 101
- inductive hypothesis, 204
- infinitude of the primes, 137
- injection, 275
- integers, 2
- intersection, 177
- inverse, 70
- inverse error, 112
- inverse image, of a set, 279
- inverse relation, 242, 273
- inverse, of a relation, 238
- invertible function, 54
- Iverson bracket, 288

- Jordan curve, 188
- Jordan curve theorem, 188
- Jordan, Camille, 188

- Königsberg, 150, 312
- Kaliningrad, 312
- Knights and Knaves, 67
- Kronecker delta, 288
- Kronecker, Leopold, 1
- laws of logical equivalence, 75
- least common multiple, 40
- least element, in a poset, 266
- lemmas, 47
- lexicographic order, 242
- Lihua, Ma, 202
- logic gates, 62
- logical equivalence, 74
- Luxembourg, 149
- magic square, 318
- maximal element, in a poset, 266
- minimal element, in a poset, 266
- modulus, of a complex number, 7
- modus ponens, 103
- modus tollens, 103
- Monge's circle theorem, 402
- Morley triangle, 386
- Morley's theorem, 383
- multiplication rule, 295
- multiset, 168
- NAND, 66
- natural numbers, 1
- negation, 58
- neusis construction, 383
- Newton, Isaac, 4
- noneg, 7
- NOR, 66
- not gates, 62
- octal representation, 32
- open sentence, 92
- operator, 270
- or gates, 62
- ordering relation, 260
- parallel connection, 60
- parity check code, 312
- partial order, 260
- partially ordered set, 261
- partition, 252
- Pascal's triangle, 30, 328
- Peano axioms, 201
- Peirce arrow, 66
- permutation, 296
- Petersen graph, 258
- pigeonhole, 323
- pigeonhole principle, 324
- pigeonhole principle, strong form, 325
- PIN, 300
- place notation, 24
- polynomial multiplication, 122
- poset, 262
- power set, 170
- predicate variable, 58
- Pregel, Pregolya, 312

- premise, 102
- prime factorization, 15
- prime numbers, 13
- product rule, 135
- projection, 286
- projective plane of order 10, 148
- proof by cases, 148
- proof by contradiction, 137
- proof by contraposition, 96, 138
- proof by exhaustion, 148
- proper subset, 172
- Properties of relations, 244
- pseudocode, 36
- punch card, 309
- Pythagoras, 5
- Pythagorean triple, 142

- quantification, 92
- quod erat demonstrandum, 128
- quotient structure, 252
- quotient-remainder theorem, 28

- radical, of an integer, 253
- Ramsey number, 148
- range, 53
- Rational approximation, 49
- rational numbers, 2
- real part, 6
- reals, 4
- recognizers, 64
- reductio ad absurdum, 46

- reflexivity, 243
- relations, 50
- relative primality, 45
- repeated division algorithm, 32
- repetition number, 168
- restriction, of a function, 283
- right inverse, 283
- rules of inference, 102, 105, 107
- rules of replacement, 83, 105

- Scheffer stroke, 66
- sentence, 56
- sequence, 294
- series connection, 60
- set theoretic equalities, 180
- set-builder notation, 3
- sieve of Eratosthenes, 14
- similarity transform, 386
- singleton set, 169
- Smullyan, Raymond, 67
- Sophie Germain prime, 99
- soundness (of an argument), 110
- square-free part, of an integer, 253
- statement, 56
- subset, 172
- successor, 266
- superset, 172
- syllogism, 103
- symmetric difference, 183
- symmetry, 243

tautology, 80
ternary relation, 51
tetromino, 321
TFAE, 13
top, in a poset, 266
total order, 260
transistor, 60
transitivity, 243
triangular numbers, 220
trichotomy, 7
trichotomy property, 156
truth set, 166
truth table, 58
Twin Prime conjecture, 19
two-column proof, 88

uncountable, 343
union, 177
unique existence, 99, 159
universal conditional statement, 120
universal quantification, 20
universal set, 167
universe of discourse, 20, 167, 188

vacuous truth, 69
valid argument form, 110
vampire number, 155, 159
Venn diagram, 188

weak Goldbach conjecture, 150
weasels, ice, 119
well-ordering principle, 158, 164

William Lowell Putnam Mathematics
 Competition, 156
winding map, 285
Yahtzee, 295
Z-module, 160