
Logic and Proof

Release 0.1

Jeremy Avigad, Robert Y. Lewis, and Floris van Doorn

Dec 06, 2018

CONTENTS

1	Introduction	1
1.1	Mathematical Proof	1
1.2	Symbolic Logic	2
1.3	Interactive Theorem Proving	4
1.4	The Semantic Point of View	5
1.5	Goals Summarized	6
1.6	About this Textbook	6
2	Propositional Logic	7
2.1	A Puzzle	7
2.2	A Solution	7
2.3	Rules of Inference	8
2.4	The Language of Propositional Logic	15
2.5	Exercises	16
3	Natural Deduction for Propositional Logic	17
3.1	Derivations in Natural Deduction	17
3.2	Examples	19
3.3	Forward and Backward Reasoning	20
3.4	Reasoning by Cases	22
3.5	Some Logical Identities	23
3.6	Exercises	24
4	Propositional Logic in Lean	25
4.1	Expressions for Propositions and Proofs	25
4.2	More commands	28
4.3	Building Natural Deduction Proofs	30
4.4	Forward Reasoning	34
4.5	Definitions and Theorems	36
4.6	Additional Syntax	38
4.7	Exercises	39
5	Classical Reasoning	41
5.1	Proof by Contradiction	41
5.2	Some Classical Principles	43
5.3	Exercises	45
6	Semantics of Propositional Logic	47
6.1	Truth Values and Assignments	47
6.2	Truth Tables	50
6.3	Soundness and Completeness	51

6.4	Exercises	53
7	First Order Logic	55
7.1	Functions, Predicates, and Relations	55
7.2	The Universal Quantifier	57
7.3	The Existential Quantifier	59
7.4	Relativization and Sorts	60
7.5	Equality	61
7.6	Exercises	62
8	Natural Deduction for First Order Logic	65
8.1	Rules of Inference	65
8.2	The Universal Quantifier	65
8.3	The Existential Quantifier	67
8.4	Equality	68
8.5	Counterexamples and Relativized Quantifiers	69
8.6	Exercises	71
9	First Order Logic in Lean	73
9.1	Functions, Predicates, and Relations	73
9.2	Using the Universal Quantifier	77
9.3	Using the Existential Quantifier	79
9.4	Equality and calculational proofs	82
9.5	Exercises	85
10	Semantics of First Order Logic	91
10.1	Interpretations	91
10.2	Truth in a Model	92
10.3	Examples	93
10.4	Validity and Logical Consequence	95
10.5	Soundness and Completeness	95
10.6	Exercises	96
11	Sets	99
11.1	Elementary Set Theory	99
11.2	Calculations with Sets	102
11.3	Indexed Families of Sets	106
11.4	Cartesian Product and Power Set	107
11.5	Exercises	108
12	Sets in Lean	111
12.1	Basics	111
12.2	Some Identities	115
12.3	Indexed Families	117
12.4	Power Sets	120
12.5	Exercises	121
13	Relations	123
13.1	Order Relations	123
13.2	More on Orderings	125
13.3	Equivalence Relations and Equality	126
13.4	Exercises	128
14	Relations in Lean	129
14.1	Order Relations	129

14.2	Orderings on Numbers	133
14.3	Equivalence Relations	133
14.4	Exercises	135
15	Functions	137
15.1	The Function Concept	137
15.2	Injective, Surjective, and Bijective Functions	139
15.3	Functions and Subsets of the Domain	141
15.4	Functions and Relations	143
15.5	Exercises	143
16	Functions in Lean	145
16.1	Functions and Symbolic Logic	145
16.2	Second- and Higher-Order Logic	146
16.3	Functions in Lean	147
16.4	Defining the Inverse Classically	150
16.5	Functions and Sets in Lean	151
16.6	Exercises	154
17	The Natural Numbers and Induction	157
17.1	The Principle of Induction	157
17.2	Variants of Induction	159
17.3	Recursive Definitions	162
17.4	Defining Arithmetic Operations	164
17.5	Arithmetic on the Natural Numbers	166
17.6	The Integers	168
17.7	Exercises	169
18	The Natural Numbers and Induction in Lean	171
18.1	Induction and Recursion in Lean	171
18.2	Defining the Arithmetic Operations in Lean	174
18.3	Exercises	175
19	Elementary Number Theory	177
19.1	The Quotient-Remainder Theorem	177
19.2	Divisibility	178
19.3	Prime Numbers	181
19.4	Modular Arithmetic	182
19.5	Properties of Squares	184
19.6	Exercises	185
20	Combinatorics	187
20.1	Finite Sets and Cardinality	187
20.2	Counting Principles	188
20.3	Ordered Selections	190
20.4	Combinations and Binomial Coefficients	191
20.5	The Inclusion-Exclusion Principle	194
20.6	Exercises	195
21	The Real Numbers	197
21.1	The Number Systems	197
21.2	Quotient Constructions	198
21.3	Constructing the Real Numbers	200
21.4	The Completeness of the Real Numbers	201
21.5	An Alternative Construction	202

21.6 Exercises	203
22 The Infinite	205
22.1 Equinumerosity	205
22.2 Countably Infinite Sets	206
22.3 Cantor's Theorem	209
22.4 An Alternative Definition of Finiteness	210
22.5 The Cantor-Bernstein Theorem	211
22.6 Exercises	211
23 Axiomatic Foundations	213
23.1 Basic Axioms for Sets	213
23.2 The Axiom of Infinity	215
23.3 The Remaining Axioms	217
23.4 Type Theory	219
23.5 Exercises	221

INTRODUCTION

1.1 Mathematical Proof

Although there is written evidence of mathematical activity in Egypt as early as 3000 BC, many scholars locate the birth of mathematics proper in ancient Greece around the sixth century BC, when deductive proof was first introduced. Aristotle credited Thales of Miletus with recognizing the importance of not just what we know but how we know it, and finding grounds for knowledge in the deductive method. Around 300 BC, Euclid codified a deductive approach to geometry in his treatise, the *Elements*. Through the centuries, Euclid's axiomatic style was held as a paradigm of rigorous argumentation, not just in mathematics, but in philosophy and the sciences as well.

Here is an example of an ordinary proof, in contemporary mathematical language. It establishes a fact that was known to the Pythagoreans.

Theorem. $\sqrt{2}$ is irrational, which is to say, it cannot be expressed as a fraction a/b , where a and b are integers.

Proof. Suppose $\sqrt{2} = a/b$ for some pair of integers a and b . By removing any common factors, we can assume a/b is in lowest terms, so that a and b have no factor in common. Then we have $a = \sqrt{2}b$, and squaring both sides, we have $a^2 = 2b^2$.

The last equation implies that a^2 is even, and since the square of an odd number is odd, a itself must be even as well. We therefore have $a = 2c$ for some integer c . Substituting this into the equation $a^2 = 2b^2$, we have $4c^2 = 2b^2$, and hence $2c^2 = b^2$. This means that b^2 is even, and so b is even as well.

The fact that a and b are both even contradicts the fact that a and b have no common factor. So the original assumption that $\sqrt{2} = a/b$ is false.

In the next example, we focus on the natural numbers,

$$\mathbb{N} = \{0, 1, 2, \dots\}.$$

A natural number n greater than or equal to 2 is said to be *composite* if it can be written as a product $n = m \cdot k$ where neither m nor k is equal to 1, and *prime* otherwise. Notice that if $n = m \cdot k$ witnesses the fact that n is composite, then m and k are both smaller than n . Notice also that, by convention, 0 and 1 are considered neither prime nor composite.

Theorem. Every natural number greater than or equal to 2 can be written as a product of primes.

Proof. We proceed by induction on n . Let n be any natural number greater than 2. If n is prime, we are done; we can consider n itself as a product with one term. Otherwise, n is composite, and we can write $n = m \cdot k$ where m and k are smaller than n and greater than 1. By the inductive hypothesis, each of m and k can be written as a product of primes, say $m = p_1 \cdot p_2 \cdot \dots \cdot p_u$ and $k = q_1 \cdot q_2 \cdot \dots \cdot q_v$. But then we have

$$n = m \cdot k = p_1 \cdot p_2 \cdot \dots \cdot p_u \cdot q_1 \cdot q_2 \cdot \dots \cdot q_v,$$

a product of primes, as required.

Later, we will see that more is true: every natural number greater than 2 can be written as a product of primes in a unique way, a fact known as the *fundamental theorem of arithmetic*.

The first goal of this course is to teach you to write clear, readable mathematical proofs. We will do this by considering a number of examples, but also by taking a reflective point of view: we will carefully study the components of mathematical language and the structure of mathematical proofs, in order to gain a better understanding of how they work.

1.2 Symbolic Logic

Toward understanding how proofs work, it will be helpful to study a subject known as “symbolic logic,” which provides an idealized model of mathematical language and proof. In the *Prior Analytics*, the ancient Greek philosopher set out to analyze patterns of reasoning, and developed the theory of the *syllogism*. Here is one instance of a syllogism:

Every man is an animal.

Every animal is mortal.

Therefore every man is mortal.

Aristotle observed that the correctness of this inference has nothing to do with the truth or falsity of the individual statements, but, rather, the general pattern:

Every A is B.

Every B is C.

Therefore every A is C.

We can substitute various properties for A, B, and C; try substituting the properties of being a fish, being a unicorn, being a swimming creature, being a mythical creature, etc. The various statements that result may come out true or false, but all the instantiations will have the following crucial feature: if the two hypotheses come out true, then the conclusion comes out true as well. We express this by saying that the inference is *valid*.

Although the patterns of language addressed by Aristotle’s theory of reasoning are limited, we have him to thank for a crucial insight: we can classify valid patterns of inference by their logical form, while abstracting away specific content. It is this fundamental observation that underlies the entire field of symbolic logic.

In the seventeenth century, Leibniz proposed the design of a *characteristica universalis*, a universal symbolic language in which one would express any assertion in a precise way, and a *calculus ratiocinator*, a “calculus of thought” which would express the precise rules of reasoning. Leibniz himself took some steps to develop such a language and calculus, but much greater strides were made in the nineteenth century, through the work of Boole, Frege, Peirce, Schroeder, and others. Early in the twentieth century, these efforts blossomed into the field of mathematical logic.

If you consider the examples of proofs in the last section, you will notice that some terms and rules of inference are specific to the subject matter at hand, having to do with numbers and the properties of being prime, composite, even, odd, and so on. But there are other terms and rules of inference that are not domain specific, such as those related to the words “every,” “some,” “and,” and “if ... then.” The goal of symbolic logic is to identify these core elements of reasoning and argumentation and explain how they work, as well as to explain how more domain-specific notions are introduced and used.

To that end, we will introduce symbols for key logical notions, including the following:

- $A \rightarrow B$, “if A then B ”
- $A \wedge B$, “ A and B ”
- $A \vee B$, “ A or B ”
- $\neg A$, “not A ”
- $\forall x A$, “for every x , A ”
- $\exists x A$, “for some x , A ”

We will then provide a formal proof system that will let us establish, deductively, that certain entailments between such statements are valid.

The proof system we will use is a version of *natural deduction*, a type of proof system introduced by Gerhard Gentzen in the 1930s to model informal styles of argument. In this system, the fundamental unit of judgment is the assertion that a statement, A , follows from a finite set of hypotheses, Γ . This is written as $\Gamma \vdash A$. If Γ and Δ are two finite sets of hypotheses, we will write Γ, Δ for the *union* of these two sets, that is, the set consisting of all the hypotheses in each. With these conventions, the rule for the conjunction symbol can be expressed as follows:

$$\frac{\Gamma \vdash A \quad \Delta \vdash B}{\Gamma, \Delta \vdash A \wedge B}$$

This should be interpreted as saying: assuming A follows from the hypotheses Γ , and B follows from the hypotheses Δ , $A \wedge B$ follows from the hypotheses in both Γ and Δ .

We will see that one can write such proofs more compactly leaving the hypotheses implicit, so that the rule above is expressed as follows:

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

In this format, a snippet of the first proof in the previous section might be rendered as follows:

$$\frac{\frac{\frac{}{\neg\text{even}(b)}}{\neg\text{even}(b)} \quad \frac{\forall x (\neg\text{even}(x) \rightarrow \neg\text{even}(x^2))}{\neg\text{even}(b) \rightarrow \neg\text{even}(b^2)}}{\neg\text{even}(b^2)} \quad \text{even}(b^2)}{\perp}}{\text{even}(b)}$$

The complexity of such proofs can quickly grow out of hand, and complete proofs of even elementary mathematical facts can become quite long. Such systems are not designed for writing serious mathematics.

Rather, they provide idealized models of mathematical inference, and insofar as they capture something of the structure of an informal proof, they enable us to study the properties of mathematical reasoning.

The second goal of this course is to help you understand natural deduction, as an example of a formal deductive system.

1.3 Interactive Theorem Proving

Early work in mathematical logic aimed to show that ordinary mathematical arguments could be modeled in symbolic calculi, at least in principle. As noted above, complexity issues limit the range of what can be accomplished in practice; even elementary mathematical arguments require long derivations that are hard to write and hard to read, and do little to promote understanding of the underlying mathematics.

Since the end of the twentieth century, however, the advent of computational proof assistants has begun to make complete formalization feasible. Working interactively with theorem proving software, users can construct formal derivations of complex theorems that can be stored and checked by computer. Automated methods can be used to fill in small gaps by hand, verify long calculations axiomatically, or fill in long chains of inferences deterministically. The reach of automation is currently fairly limited, however. The strategy used in interactive theorem proving is to ask users to provide just enough information for the system to be able to construct and check a formal derivation. This typically involves writing proofs in a sort of “programming language” that is designed with that purpose in mind. For example, here is a short proof in the *Lean* theorem prover:

```
section
variables (P Q : Prop)

theorem my_theorem : P ∧ Q → Q ∧ P :=
assume h : P ∧ Q,
have P, from and.left h,
have Q, from and.right h,
show Q ∧ P, from and.intro <Q> <P>

end
```

If you are reading the present text in online form, you will find a button above the formal “proof script” that says “try it!” Pressing the button opens the proof in an editor window and runs a version of Lean inside your browser to process the proof, turn it into an axiomatic derivation, and verify its correctness. You can experiment by varying the text in the editor; any errors will be noted in the window to the right.

Proofs in Lean can access a library of prior mathematical results, all verified down to axiomatic foundations. A goal of the field of interactive theorem proving is to reach the point where any contemporary theorem can be verified in this way. For example, here is a formal proof that the square root of two is irrational, following the model of the informal proof presented above:

```
import data.nat.prime
open nat

theorem sqrt_two_irrational {a b : ℕ} (co : gcd a b = 1) :
  a^2 ≠ 2 * b^2 :=
assume h : a^2 = 2 * b^2,
have 2 | a^2,
  by simp [h],
have 2 | a,
```

(continues on next page)

(continued from previous page)

```

    from prime.dvd_of_dvd_pow prime_two this,
exists.elim this $
assume (c : nat) (aeq : a = 2 * c),
have 2 * (2 * c^2) = 2 * b^2,
  by simp [eq.symm h, aeq];
  simp [pow_succ, mul_comm, mul_assoc, mul_left_comm],
have 2 * c^2 = b^2,
  from eq_of_mul_eq_mul_left dec_trivial this,
have 2 | b^2,
  by simp [eq.symm this],
have 2 | b,
  from prime.dvd_of_dvd_pow prime_two this,
have 2 | gcd a b,
  from dvd_gcd <2 | a> <2 | b>,
have 2 | (1 : N),
  by simp * at *,
show false, from absurd <2 | 1> dec_trivial

```

The third goal of this course is to teach you to write elementary proofs in Lean. The facts that we will ask you to prove in Lean will be more elementary than the informal proofs we will ask you to write, but our intent is that formal proofs will model and clarify the informal proof strategies we will teach you.

1.4 The Semantic Point of View

As we have presented the subject here, the goal of symbolic logic is to specify a language and rules of inference that enable us to get at the truth in a reliable way. The idea is that the symbols we choose denote objects and concepts that have a fixed meaning, and the rules of inference we adopt enable us to draw true conclusions from true hypotheses.

One can adopt another view of logic, however, as a system where some symbols have a fixed meaning, such as the symbols for “and,” “or,” and “not,” and others have a meaning that is taken to vary. For example, the expression $P \wedge (Q \vee R)$, read “ P and either Q or R ,” may be true or false *depending on the basic assertions that P , Q , and R stand for*. More precisely, the truth of the compound expression depends only on whether the component symbols denote expressions that are true or false. For example, if P , Q , and R stand for “seven is prime,” “seven is even,” and “seven is odd,” respectively, then the expression is true. If we replace “seven” by “six,” the statement is false. More generally, the expression comes out true whenever P is true and at least one of Q and R is true, and false otherwise.

From this perspective, logic is not so much a language for asserting truth, but a language for describing possible states of affairs. In other words, logic provides a specification language, with expressions that can be true or false depending on how we interpret the symbols that are allowed to vary. For example, if we fix the meaning of the basic predicates, the statement “there is a red block between two blue blocks” may be true or false of a given “world” of blocks, and we can take the expression to describe the set of worlds in which it is true. Such a view of logic is important in computer science, where we use logical expressions to select entries from a database matching certain criteria, to specify properties of hardware and software systems, or to assert constraints that we would like a constraint solver to satisfy.

There are important connections between the syntactic / deductive point of view on the one hand, and the semantic / model-theoretic point of view on the other. We will explore some of these along the way. For example, we will see that it is possible to view the “valid” assertions as those that are true under all possible interpretations of the non-fixed symbols, and the “valid” inferences as those that maintain truth in all possible states and affairs. From this point of view, a deductive system should only allow us to derive

valid assertions and entailments, a property known as *soundness*. If a deductive system is strong enough to allow us to verify *all* valid assertions and entailments, it is said to be *complete*.

The fourth goal of this course is to convey the semantic view of logic, and to lead you to understand how logical expressions can be used to specify states of affairs.

1.5 Goals Summarized

To summarize, these are the goals of this course:

- You should learn to write clear, “literate,” mathematical proofs.
- You should become comfortable with symbolic logic and the formal modeling of deductive proof.
- You should learn how to use an interactive proof assistant.
- You should understand how to use logic as a precise language for making claims about systems of objects and the relationships between them, and specifying certain states of affairs.

Let us take a moment to consider the relationship between some of these goals. It is important not to confuse the first three. We are dealing with three kinds of mathematical language: ordinary mathematical language, the symbolic representations of mathematical logic, and computational implementations in interactive proof assistants. These are very different things!

Symbolic logic is not meant to replace ordinary mathematical language, and you should not use symbols like \wedge and \vee in ordinary mathematical proofs any more than you would use them in place of the words “and” and “or” in letters home to your parents. Natural languages provide nuances of expression that can convey levels of meaning and understanding that go beyond pattern matching to verify correctness. At the same time, modeling mathematical language with symbolic expressions provides a level of precision that makes it possible to turn mathematical language itself into an object of study. Each has its place, and we hope to get you to appreciate the value of each without confusing the two.

The proof languages used by interactive theorem provers lie somewhere between the two extremes. On the one hand, they have to be specified with enough precision for a computer to process them and act appropriately; on the other hand, they aim to capture some of the higher-level nuances and features of informal language in a way that enables us to write more complex arguments and proofs. Rooted in symbolic logic and designed with ordinary mathematical language in mind, they aim to bridge the gap between the two.

1.6 About this Textbook

Both this online textbook and the *Lean* theorem prover it invokes are new and ongoing projects, and in places they are still rough. Please bear with us! Your feedback will be quite helpful.

PROPOSITIONAL LOGIC

2.1 A Puzzle

The following puzzle, titled “Malice and Alice,” is from George J. Summers’ *Logical Deduction Puzzles*.

Alice, Alice’s husband, their son, their daughter, and Alice’s brother were involved in a murder. One of the five killed one of the other four. The following facts refer to the five people mentioned:

1. A man and a woman were together in a bar at the time of the murder.
2. The victim and the killer were together on a beach at the time of the murder.
3. One of Alice’s two children was alone at the time of the murder.
4. Alice and her husband were not together at the time of the murder.
5. The victim’s twin was not the killer.
6. The killer was younger than the victim.

Which one of the five was the victim?

Take some time to try to work out a solution. (You should assume that the victim’s twin is one of the five people mentioned.) Summers’ book offers the following hint: “First find the locations of two pairs of people at the time of the murder, and then determine who the killer and the victim were so that no condition is contradicted.”

2.2 A Solution

If you have worked on the puzzle, you may have noticed a few things. First, it is helpful to draw a diagram, and to be systematic about searching for an answer. The number of characters, locations, and attributes is finite, so that there are only finitely many possible “states of affairs” that need to be considered. The numbers are also small enough so that systematic search through all the possibilities, though tedious, will eventually get you to the right answer. This is a special feature of logic puzzles like this; you would not expect to show, for example, that every even number greater than two can be written as a sum of primes by running through all the possibilities.

Another thing that you may have noticed is that the question seems to presuppose that there is a unique answer to the question, which is to say, over all the states of affairs that meet the list of conditions, there is only one person who can possibly be the killer. *A priori*, without that assumption, there is a difference between finding *some* person who could have been the victim and showing that that person *had* to be the

victim. In other words, there is a difference between exhibiting some state of affairs that meets the criteria and demonstrating conclusively that no other solution is possible.

The published solution in the book not only produces a state of affairs that meets the criterion, but at the same time proves that this is the only one that does so. It is quoted below, in full.

From (1), (2), and (3), the roles of the five people were as follows: Man and Woman in the bar, Killer and Victim on the beach, and Child alone.

Then, from (4), either Alice's husband was in the bar and Alice was on the beach, or Alice was in the bar and Alice's husband was on the beach.

If Alice's husband was in the bar, the woman he was with was his daughter, the child who was alone was his son, and Alice and her brother were on the beach. Then either Alice or her brother was the victim; so the other was the killer. But, from (5), the victim had a twin, and this twin was innocent. Since Alice and her brother could only be twins to each other, this situation is impossible. Therefore Alice's husband was not in the bar.

So Alice was in the bar. If Alice was in the bar, she was with her brother or her son.

If Alice was with her brother, her husband was on the beach with one of the two children. From (5), the victim could not be her husband, because none of the others could be his twin; so the killer was her husband and the victim was the child he was with. But this situation is impossible, because it contradicts (6). Therefore, Alice was not with her brother in the bar.

So Alice was with her son in the bar. Then the child who was alone was her daughter. Therefore, Alice's husband was with Alice's brother on the beach. From previous reasoning, the victim could not be Alice's husband. But the victim could be Alice's brother because Alice could be his twin.

So *Alice's brother was the victim* and Alice's husband was the killer.

This argument relies on some “extralogical” elements, for example, that a father cannot be younger than his child, and that a parent and his or her child cannot be twins. But the argument also involves a number of common logical terms and associated patterns of inference. In the next section, we will focus on some of the key logical terms occurring in the argument above, words like “and,” “or,” “not,” and “if ... then.”

Our goal is to give an account of the patterns of inference that govern the use of those terms. To that end, using the methods of symbolic logic, we will introduce variables A , B , C , ... to stand for fundamental statements, or *propositions*, and symbols \wedge , \vee , \neg , and \rightarrow to stand for “and,” “or,” “not,” and “if ... then ... ,” respectively. Doing so will let us focus on the way that compound statements are built up from basic ones using the logical terms, while abstracting away from the specific content. We will also adopt a stylized notation for representing inferences as *rules*: the inscription

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

indicates that statement C is a *logical consequence* of A and B .

2.3 Rules of Inference

2.3.1 Implication

The first pattern of inference we will discuss, involving the “if ... then ...” construct, can be hard to discern. Its use is largely implicit in the solution above. The inference in the fourth paragraph, spelled out in greater detail, runs as follows:

If Alice was in the bar, Alice was with her brother or her son.

Alice was in the bar.

Alice was with her brother or son.

This rule is sometimes known as *modus ponens*, or “implication elimination,” since it tells us how to use an implication in an argument. As a rule, it is expressed as follows:

$$\frac{A \rightarrow B \quad A}{B} \rightarrow E$$

Read this as saying that if you have a proof of $A \rightarrow B$, possibly from some hypotheses, and a proof of A , possibly from hypotheses, then combining these yields a proof of B , from the hypotheses in both subproofs.

The rule for deriving an “if ... then” statement is more subtle. Consider the beginning of the third paragraph, which argues that if Alice’s husband was in the bar, then Alice or her brother was the victim. Abstracting away some of the details, the argument has the following form:

Suppose Alice’s husband was in the bar.

Then ...

Then ...

Then Alice or her brother was the victim.

Thus, if Alice’s husband was in the bar, then Alice or her brother was the victim.

This is a form of *hypothetical reasoning*. On the supposition that A holds, we argue that B holds as well. If we are successful, we have shown that A implies B , without supposing A . In other words, the temporary assumption that A holds is “canceled” by making it explicit in the conclusion.

$$\frac{\overline{A}^1 \quad \vdots \quad B}{A \rightarrow B}^1 \rightarrow I$$

The hypothesis is given the label 1; when the introduction rule is applied, the label 1 indicates the relevant hypothesis. The line over the hypothesis indicates that the assumption has been “canceled” by the introduction rule.

2.3.2 Conjunction

As was the case for implication, other logical connectives are generally characterized by their *introduction* and *elimination* rules. An introduction rule shows how to establish a claim involving the connective, while an elimination rule shows how to use such a statement that contains the connective to derive others.

Let us consider, for example, the case of conjunction, that is, the word “and.” Informally, we establish a conjunction by establishing each conjunct. For example, informally we might argue:

Alice’s brother was the victim.

Alice's husband was the killer.

Therefore Alice's brother was the victim and Alice's husband was the killer.

The inference seems almost too obvious to state explicitly, since the word “and” simply combines the two assertions into one. Informal proofs often downplay the distinction. In symbolic logic, the rule reads as follows:

$$\frac{A \quad B}{A \wedge B} \wedge I$$

The two elimination rules allow us to extract the two components:

Alice's husband was in the bar and Alice was on the beach.

So Alice's husband was in the bar.

Or:

Alice's husband was in the bar and Alice was on the beach.

So Alice was on the beach.

In symbols, these patterns are rendered as follows:

$$\frac{A \wedge B}{A} \wedge E_l \quad \frac{A \wedge B}{B} \wedge E_r$$

Here the l and r stand for “left” and “right”.

2.3.3 Negation and Falsity

In logical terms, showing “not A” amounts to showing that A leads to a contradiction. For example:

Suppose Alice's husband was in the bar.

...

This situation is impossible.

Therefore Alice's husband was not in the bar.

This is another form of hypothetical reasoning, similar to that used in establishing an “if ... then” statement: we temporarily assume A, show that leads to a contradiction, and conclude that “not A” holds. In symbols, the rule reads as follows:

$$\frac{\begin{array}{c} \overline{A} \quad 1 \\ \vdots \\ \perp \end{array}}{\neg A} 1 \neg I$$

The elimination rule is dual to these. It expresses that if we have both “A” and “not A,” then we have a contradiction. This pattern is illustrated in the informal argument below, which is implicit in the fourth paragraph of the solution to “Malice and Alice.”

The killer was Alice’s husband and the victim was the child he was with.

So the killer was not younger than his victim.

But according to (6), the killer was younger than his victim.

This situation is impossible.

In symbolic logic, the rule of inference is expressed as follows:

$$\frac{\neg A \quad A}{\perp} \neg E$$

Notice also that in the symbolic framework, we have introduced a new symbol, \perp . It corresponds to natural language phrases like “this is a contradiction” or “this is impossible.”

What are the rules governing \perp ? In the proof system we will introduce in the next chapter, there is no introduction rule; “false” is false, and there should be no way to prove it, other than extract it from contradictory hypotheses. On the other hand, the system provides a rule that allows us to conclude anything from a contradiction:

$$\frac{\perp}{A} \perp E$$

The elimination rule also has the fancy Latin name, *ex falso sequitur quodlibet*, which means “anything you want follows from falsity.”

This elimination rule is harder to motivate from a natural language perspective, but, nonetheless, it is needed to capture common patterns of inference. One way to understand it is this. Consider the following statement:

For every natural number n , if n is prime and greater than 2, then n is odd.

We would like to say that this is a true statement. But if it is true, then it is true of any particular number n . Taking $n = 2$, we have the statement:

If 2 is prime and greater than 2, then 2 is odd.

In this conditional statement, both the antecedent and succedent are false. The fact that we are committed to saying that this statement is true shows that we should be able to prove, one way or another, that the statement 2 is odd follows from the false statement that 2 is prime and greater than 2. The *ex falso* neatly encapsulates this sort of inference.

Notice that if we define $\neg A$ to be $A \rightarrow \perp$, then the rules for negation introduction and elimination are nothing more than implication introduction and elimination, respectively. We can think of $\neg A$ expressed colorfully by saying “if A is true, then pigs have wings,” where “pigs have wings” is stands for \perp .

Having introduced a symbol for “false,” it is only fair to introduce a symbol for “true.” In contrast to “false,” “true” has no elimination rule, only an introduction rule:

$$\overline{\top}$$

Put simply, “true” is true.

2.3.4 Disjunction

The introduction rules for disjunction, otherwise known as “or,” are straightforward. For example, the claim that condition (3) is met in the proposed solution can be justified as follows:

Alice’s daughter was alone at the time of the murder.

Therefore, either Alice’s daughter was alone at the time of the murder, or Alice’s son was alone at the time of the murder.

In symbolic terms, the two introduction rules are as follows:

$$\frac{A}{A \vee B} \vee I_l \quad \frac{B}{A \vee B} \vee I_r$$

Here, again, the l and r stand for “left” and “right”.

The disjunction elimination rule is trickier, but it represents a natural form of case-based hypothetical reasoning. The instances that occur in the solution to “Malice and Alice” are all special cases of this rule, so it will be helpful to make up a new example to illustrate the general phenomenon. Suppose, in the argument above, we had established that either Alice’s brother or her son was in the bar, and we wanted to argue for the conclusion that her husband was on the beach. One option is to argue by cases: first, consider the case that her brother was in the bar, and argue for the conclusion on the basis of that assumption; then consider the case that her son was in the bar, and argue for the same conclusion, this time on the basis of the second assumption. Since the two cases are exhaustive, if we know that the conclusion holds in each case, we know that it holds outright. The pattern looks something like this:

Either Alice’s brother was in the bar, or Alice’s son was in the bar.

Suppose, in the first case, that her brother was in the bar. Then ... Therefore, her husband was on the beach.

On the other hand, suppose her son was in the bar. In that case, ... Therefore, in this case also, her husband was on the beach.

Either way, we have established that her husband was on the beach.

In symbols, this pattern is expressed as follows:

$$\frac{\frac{A}{A \vee B} \quad \frac{B}{A \vee B}}{A \vee B} \quad \frac{\frac{C}{C} \quad \frac{C}{C}}{C} \vee E$$

What makes this pattern confusing is that it requires two instances of nested hypothetical reasoning: in the first block of parentheses, we temporarily assume A , and in the second block, we temporarily assume B . When the dust settles, we have established C outright.

There is another pattern of reasoning that is commonly used with “or,” as in the following example:

Either Alice’s husband was in the bar, or Alice was in the bar.

Alice’s husband was not in the bar.

So Alice was in the bar.

In symbols, we would render this rule as follows:

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

We will see in the next chapter that it is possible to *derive* this rule from the others. As a result, we will *not* take this to be a fundamental rule of inference in our system.

2.3.5 If and only if

In mathematical arguments, it is common to say of two statements, A and B , that “ A holds if and only if B holds.” This assertion is sometimes abbreviated “ A iff B ,” and means simply that A implies B and B implies A . It is not essential that we introduce a new symbol into our logical language to model this connective, since the statement can be expressed, as we just did, in terms of “implies” and “and.” But notice that the length of the expression doubles because A and B are each repeated. The logical abbreviation is therefore convenient, as well as natural.

The conditions of “Malice and Alice” imply that Alice is in the bar if and only if Alice’s husband is on the beach. Such a statement is established by arguing for each implication in turn:

I claim that Alice is in the bar if and only if Alice’s husband is on the beach.

To see this, first suppose that Alice is in the bar.

Then ...

Hence Alice’s husband is on the beach.

Conversely, suppose Alice’s husband is on the beach.

Then ...

Hence Alice is in the bar.

Notice that with this example, we have varied the form of presentation, stating the conclusion first, rather than at the end of the argument. This kind of “signposting” is common in informal arguments, in that it helps guide the reader’s expectations and foreshadow where the argument is going. The fact that formal systems of deduction do not generally model such nuances marks a difference between formal and informal arguments, a topic we will return to below.

The introduction is modeled in natural deduction as follows:

$$\frac{\begin{array}{c} \overline{A}^1 \quad \overline{B}^1 \\ \vdots \quad \vdots \\ B \quad A \end{array}}{A \leftrightarrow B}^1 \leftrightarrow I$$

The elimination rules for iff are unexciting. In informal language, here is the “left” rule:

Alice is in the bar if and only if Alice’s husband is on the beach.

Alice is in the bar.

Hence, Alice’s husband is on the beach.

The “right” rule simply runs in the opposite direction.

Alice is in the bar if and only if Alice’s husband is on the beach.

Alice’s husband is on the beach.

Hence, Alice is in the bar.

Rendered in natural deduction, the rules are as follows:

$$\frac{A \leftrightarrow B \quad A}{B} \leftrightarrow E_l \quad \frac{A \leftrightarrow B \quad B}{A} \leftrightarrow E_r$$

2.3.6 Proof by Contradiction

We saw an example of an informal argument that implicitly uses the introduction rule for negation:

Suppose Alice’s husband was in the bar.

...

This situation is impossible.

Therefore Alice’s husband was not in the bar.

Consider the following argument:

Suppose Alice’s husband was not on the beach.

...

This situation is impossible.

Therefore Alice’s husband was on the beach.

At first glance, you might think this argument follows the same pattern as the one before. But a closer look should reveal a difference: in the first argument, a negation is *introduced* into the conclusion, whereas in the second, it is *eliminated* from the hypothesis. Using negation introduction to close the second argument would yield the conclusion “It is not the case that Alice’s husband was not on the beach.” The rule of inference that replaces the conclusion with the positive statement that Alice’s husband *was* on the beach is called a *proof by contradiction*. (It also has a fancy name, *reductio ad absurdum*, “reduction to an absurdity.”)

It may be hard to see the difference between the two rules, because we commonly take the statement “Alice’s husband was not not on the beach” to be a roundabout and borderline ungrammatical way of saying that Alice’s husband was on the beach. Indeed, the rule is equivalent to adding an axiom that says that for every statement A, “not not A” is equivalent to A.

There is a style of doing mathematics known as “constructive mathematics” that denies the equivalence of “not not A” and A. Constructive arguments tend to have much better computational interpretations; a proof that something is true should provide explicit evidence that the statement is true, rather than evidence that it can’t possibly be false. We will discuss constructive reasoning in a later chapter. Nonetheless, proof by contradiction is used extensively in contemporary mathematics, and so, in the meanwhile, we will use proof by contradiction freely as one of our basic rules.

In natural deduction, proof by contradiction is expressed by the following pattern:

$$\frac{\overline{\neg A}^1}{\vdots} \frac{\perp}{A} \text{RAA,1}$$

The assumption $\neg A$ is canceled at the final inference.

2.4 The Language of Propositional Logic

The language of propositional logic starts with symbols A, B, C, \dots which are intended to range over basic assertions, or propositions, which can be true or false. Compound expressions are built up using parentheses and the logical symbols introduced in the last section. For example,

$$((A \wedge \neg B) \rightarrow \neg(C \vee D))$$

is an example of a propositional formula.

When writing expressions in symbolic logic, we will adopt an order of operations which allows us to drop superfluous parentheses. When parsing an expression:

- Negation binds most tightly.
- Then, conjunctions and disjunctions bind from right to left.
- Finally, implications and bi-implications bind from right to left.

So, for example, the expression $\neg A \vee B \rightarrow C \wedge D$ is understood as $((\neg A) \vee B) \rightarrow (C \wedge D)$.

For example, suppose we assign the following variables:

- A : Alice’s husband was in the bar
- B : Alice was on the beach
- C : Alice was in the bar
- D : Alice’s husband was on the beach

Then the statement “either Alice’s husband was in the bar and Alice was on the beach, or Alice was in the bar and Alice’s husband was on the beach” would be rendered as

$$(A \wedge B) \vee (C \wedge D).$$

Sometimes the appropriate translation is not so straightforward, however. Because natural language is more flexible and nuanced, a degree of abstraction and regimentation is needed to carry out the translation. Sometimes different translations are arguably reasonable. In happy situations, alternative translations will be logically equivalent, in the sense that one can derive each from the other using purely logical rules. In less happy situations, the translations will not be equivalent, in which case the original statement is simply ambiguous, from a logical point of view. In cases like that, choosing a symbolic representation helps clarify the intended meaning.

Consider, for example, a statement like “Alice was with her son on the beach, but her husband was alone.” We might choose variables as follows:

- A : Alice was on the beach
- B : Alice’s son was on the beach
- C : Alice’s husband was alone

In that case, we might represent the statement in symbols as $A \wedge B \wedge C$. Using the word “with” may seem to connote more than the fact that Alice and her son were both on the beach; for example, it seems to connote that they are aware of each others’ presence, interacting, etc. Similarly, although we have translated the word “but” and “and,” the word “but” also conveys information; in this case, it seems to emphasize a contrast, while in other situations, it can be used to assert a fact that is contrary to expectations. In both cases, then, the logical rendering models certain features of the original sentence while abstracting others.

2.5 Exercises

1. Here is another (gruesome) logic puzzle by George J. Summers, called “Murder in the Family.”

Murder occurred one evening in the home of a father and mother and their son and daughter. One member of the family murdered another member, the third member witnessed the crime, and the fourth member was an accessory after the fact.

1. The accessory and the witness were of opposite sex.
2. The oldest member and the witness were of opposite sex.
3. The youngest member and the victim were of opposite sex.
4. The accessory was older than the victim.
5. The father was the oldest member.
6. The murderer was not the youngest member.

Which of the four—father, mother, son, or daughter—was the murderer?

Solve this puzzle, and *write a clear argument* to establish that your answer is correct.

2. Using the mnemonic F (Father), M (Mother), D (Daughter), S (Son), Mu (Murderer), V (Victim), W (Witness), A (Accessory), O (Oldest), Y (Youngest), we can define propositional variables like FM (Father is the Murderer), DV (Daughter is the Victim), FO (Father is Oldest), VY (Victim is Youngest), etc. Notice that only the son or daughter can be the youngest, and only the mother or father can be the oldest.

With these conventions, the first clue can be represented as

$$((FA \vee SA) \rightarrow (MW \vee DW)) \wedge ((MA \vee DA) \rightarrow (FW \vee SW)),$$

in other words, if the father or son was the accessory, then the mother or daughter was the witness, and vice-versa. Represent the other five clues in a similar manner.

Representing the fourth clue is tricky. Try to write down a formula that describes all the possibilities that are not ruled out by the information.

3. Consider the following three hypotheses:
 - Alan likes kangaroos, and either Betty likes frogs or Carl likes hamsters.
 - If Betty likes frogs, then Alan doesn’t like kangaroos.
 - If Carl likes hamsters, then Betty likes frogs.

Write a clear argument to show that these three hypotheses are contradictory.

NATURAL DEDUCTION FOR PROPOSITIONAL LOGIC

Reflecting on the arguments in the previous chapter, we see that, intuitively speaking, some inferences are *valid* and some are not. For example, if, in a chain of reasoning, we had established “ A and B ,” it would seem perfectly reasonable to conclude B . If we had established A , B , and “If A and B then C ,” it would be reasonable to conclude C . On the other hand, if we had established “ A or B ,” we would not be justified in concluding B without further information.

The task of symbolic logic is to develop a precise mathematical theory that explains which inferences are valid and why. There are two general approaches to spelling out the notion of validity. In this chapter, we will consider the *deductive* approach: an inference is valid if it can be justified by fundamental rules of reasoning that reflect the meaning of the logical terms involved. In [Chapter 6](#) we will consider the “semantic” approach: an inference is valid if it is an instance of a pattern that always yields a true conclusion from true hypotheses.

3.1 Derivations in Natural Deduction

We have seen that the language of propositional logic allows us to build up expressions from propositional variables A, B, C, \dots using propositional connectives like \rightarrow , \wedge , \vee , and \neg . We will now consider a formal deductive system that we can use to *prove* propositional formulas. There are a number of such systems on offer; the one we will use is called *natural deduction*, designed by Gerhard Gentzen in the 1930s.

In natural deduction, every proof is a proof from *hypotheses*. In other words, in any proof, there is a finite set of hypotheses $\{B, C, \dots\}$ and a conclusion A , and what the proof shows is that A follows from B, C, \dots

Like formulas, proofs are built by putting together smaller proofs, according to the rules. For instance, the way to read the and-introduction rule

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

is as follows: if you have a proof P_1 of A from some hypotheses, and you have a proof P_2 of B from some hypotheses, then you can put them together using this rule to obtain a proof of $A \wedge B$, which uses all the hypotheses in P_1 together with all the hypotheses in P_2 . For example, this is a proof of $(A \wedge B) \wedge (A \wedge C)$ from three hypotheses, A , B , and C :

$$\frac{\frac{A \quad B}{A \wedge B} \quad \frac{A \quad C}{A \wedge C}}{(A \wedge B) \wedge (A \wedge C)}$$

In some presentations of natural deduction, a proof is written as a sequence of lines in which each line can refer to any previous lines for justification. But here we will adopt a rigid two-dimensional diagrammatic format in which the premises of each inference appear immediately above the conclusion. This makes it easy to look over a proof and check that it is correct: each inference should be the result of instantiating the letters in one of the rules with particular formulas.

One thing that makes natural deduction confusing is that when you put together proofs in this way, hypotheses can be eliminated, or, as we will say, *canceled*. For example, we can apply the implies-introduction rule to the last proof, and obtain the following proof of $B \rightarrow (A \wedge B) \wedge (A \wedge C)$ from only *two* hypotheses, A and C :

$$\frac{\frac{\frac{A}{A \wedge B} \quad \overline{B}^1}{(A \wedge B) \wedge (A \wedge C)}}{B \rightarrow (A \wedge B) \wedge (A \wedge C)}^1$$

Here, we have used the label 1 to indicate the place where the hypothesis B was canceled. Any label will do, though we will tend to use numbers for that purpose.

We can continue to cancel the hypothesis A :

$$\frac{\frac{\frac{\overline{A}^2 \quad \overline{B}^1}{A \wedge B} \quad \frac{\overline{A}^2 \quad C}{A \wedge C}}{(A \wedge B) \wedge (A \wedge C)}^1}{B \rightarrow (A \wedge B) \wedge (A \wedge C)}^1}{A \rightarrow (B \rightarrow (A \wedge B) \wedge (A \wedge C))}^2$$

The result is a proof using only the hypothesis C . We can continue to cancel that hypothesis as well:

$$\frac{\frac{\frac{\frac{\overline{A}^2 \quad \overline{B}^1}{A \wedge B} \quad \frac{\overline{A}^2 \quad \overline{C}^3}{A \wedge C}}{(A \wedge B) \wedge (A \wedge C)}^1}{B \rightarrow (A \wedge B) \wedge (A \wedge C)}^1}{A \rightarrow (B \rightarrow (A \wedge B) \wedge (A \wedge C))}^2}{C \rightarrow (A \rightarrow (B \rightarrow (A \wedge B) \wedge (A \wedge C)))}^3$$

The resulting proof uses no hypothesis at all. In other words, it establishes the conclusion outright.

Notice that in the second step, we canceled two “copies” of the hypothesis A . In natural deduction, we can choose which hypotheses to cancel; we could have canceled either one, and left the other hypothesis *open*. In fact, we can also carry out the implication-introduction rule and cancel *zero* hypotheses. For example, the following is a short proof of $A \rightarrow B$ from the hypothesis B :

$$\frac{B}{A \rightarrow B}$$

In this proof, zero copies of A are canceled.

Also notice that although we are using letters like A , B , and C as propositional variables, in the proofs above we can replace them by any propositional formula. For example, we can replace A by the formula $(D \vee E)$ everywhere, and still have correct proofs. In some presentations of logic, different letters are used for propositional variables and arbitrary propositional formulas, but we will continue to blur the distinction. You can think of A , B , and C as standing for propositional variables or formulas, as you prefer. If you think of them as propositional variables, just keep in mind that in any rule or proof, you can replace every variable by a different formula, and still have a valid rule or proof.

Finally, notice also that in these examples, we have assumed a special rule as the starting point for building proofs. It is called the assumption rule, and it looks like this:

$$A$$

What it means is that at any point we are free to simply assume a formula, A . The single formula A constitutes a one-line proof, and the way to read this proof is as follows: assuming A , we have proved A .

The remaining rules of inference were given in the last chapter, and we summarize them here.

Implication:

$$\frac{\begin{array}{c} \overline{A}^1 \\ \vdots \\ B \end{array}}{A \rightarrow B}^1 \rightarrow I \qquad \frac{A \rightarrow B \quad A}{B} \rightarrow E$$

Conjunction:

$$\frac{A \quad B}{A \wedge B} \wedge I \qquad \frac{A \wedge B}{A} \wedge E_l \qquad \frac{A \wedge B}{B} \wedge E_r$$

Negation:

$$\frac{\begin{array}{c} \overline{A}^1 \\ \vdots \\ \perp \end{array}}{\neg A}^1 \neg I \qquad \frac{\neg A \quad A}{\perp} \neg E$$

Disjunction:

$$\frac{A}{A \vee B} \vee I_l \qquad \frac{B}{A \vee B} \vee I_r \qquad \frac{\begin{array}{c} \overline{A}^1 \quad \overline{B}^1 \\ \vdots \quad \vdots \\ A \vee B \quad C \quad C \end{array}}{C}^1 \vee E$$

Truth and falsity:

$$\frac{\perp}{A} \perp E \qquad \frac{}{\top} \top I$$

Bi-implication:

$$\frac{\begin{array}{c} \overline{A}^1 \quad \overline{B}^1 \\ \vdots \quad \vdots \\ B \quad A \end{array}}{A \leftrightarrow B}^1 \leftrightarrow I \qquad \frac{A \leftrightarrow B \quad A}{B} \leftrightarrow E_l \qquad \frac{A \leftrightarrow B \quad B}{A} \leftrightarrow E_r$$

Reductio ad absurdum (proof by contradiction):

$$\frac{\begin{array}{c} \overline{\neg A}^1 \\ \vdots \\ \perp \end{array}}{A}^1 \text{RAA}$$

3.2 Examples

Let us consider some more examples of natural deduction proofs. In each case, you should think about what the formulas say and which rule of inference is invoked at each step. Also pay close attention to which hypotheses are canceled at each stage. If you look at any node of the tree, what has been established at that point is that the claim follows from all the hypotheses above it that haven't been canceled yet.

The following is a proof of $A \rightarrow C$ from $A \rightarrow B$ and $B \rightarrow C$:

$$\frac{\frac{\overline{A}^1 \quad A \rightarrow B}{B} \quad B \rightarrow C}{A \rightarrow C}^1$$

Intuitively, the formula

$$(A \rightarrow B) \wedge (B \rightarrow C) \rightarrow (A \rightarrow C)$$

“internalizes” the conclusion of the previous proof. The \wedge symbol is used to combine hypotheses, and the \rightarrow symbol is used to express that the right-hand side is a consequence of the left. Here is a proof of that formula:

$$\frac{\frac{\frac{A}{A}^1 \quad \frac{\frac{(A \rightarrow B) \wedge (B \rightarrow C)}{(A \rightarrow B) \wedge (B \rightarrow C)}^2}{B}}{C}^1}{(A \rightarrow B) \wedge (B \rightarrow C) \rightarrow (A \rightarrow C)}^2$$

The next proof shows that if a conclusion, C , follows from A and B , then it follows from their conjunction.

$$\frac{\frac{\frac{A \rightarrow (B \rightarrow C)}{B \rightarrow C}^2 \quad \frac{\frac{A \wedge B}{A}^1}{A \wedge B}^1}{C}^1}{(A \rightarrow (B \rightarrow C)) \rightarrow (A \wedge B \rightarrow C)}^2$$

The conclusion of the next proof can be interpreted as saying that if it is not the case that one of A or B is true, then they are both false. It illustrates the use of the rules for negation.

$$\frac{\frac{\frac{\neg(A \vee B)}{\neg A}^3 \quad \frac{\frac{A}{A \vee B}^1}{\neg(A \vee B)}^3}{\neg A}^1 \quad \frac{\frac{\frac{\neg(A \vee B)}{\neg B}^3 \quad \frac{B}{A \vee B}^2}{\neg B}^2}{\neg A \wedge \neg B}^1}{\neg(A \vee B) \rightarrow \neg A \wedge \neg B}^3$$

Finally, the next two examples illustrate the use of the *ex falso* rule. The first is a derivation of an arbitrary formula B from $\neg A$ and A :

$$\frac{\frac{\neg A \quad A}{\perp}}{B}$$

The second shows that B follows from A and $\neg A \vee B$:

$$\frac{\frac{\frac{\neg A}{\neg A \vee B}^1 \quad A}{\perp}^1}{B}^1$$

In some proof systems, these rules are taken to be part of the system. But we do not need to that with our system: these two examples show that the rules can be *derived* from our other rules.

3.3 Forward and Backward Reasoning

Natural deduction is supposed to represent an idealized model of the patterns of reasoning and argumentation we use, for example, when working with logic puzzles as in the last chapter. There are obvious differences: we describe natural deduction proofs with symbols and two-dimensional diagrams, whereas our informal arguments are written with words and paragraphs. It is worthwhile to reflect on what *is* captured by the model. Natural deduction is supposed to clarify the *form* and *structure* of our logical arguments, describe the appropriate means of justifying a conclusion, and explain the sense in which the rules we use are valid.

Constructing natural deduction proofs can be confusing, but it is helpful to think about *why* it is confusing. We could, for example, decide that natural deduction is not a good model for logical reasoning. Or we might come to the conclusion that the features of natural deduction that make it confusing tell us something interesting about ordinary arguments.

In the “official” description, natural deduction proofs are constructed by putting smaller proofs together to obtain bigger ones. To prove $A \wedge B \rightarrow B \wedge A$, we start with the hypothesis $A \wedge B$. Then we construct, separately, the following two proofs:

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

Then we use these two proofs to construct the following one:

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

Finally, we apply the implies-introduction rule to this proof to cancel the hypothesis and obtain the desired conclusion:

$$\frac{\frac{\frac{A \wedge B}{B}^1 \quad \frac{A \wedge B}{A}^1}{B \wedge A}}{A \wedge B \rightarrow B \wedge A}^1$$

The process is similar to what happens in an informal argument, where we start with some hypotheses, and work forward towards a conclusion.

Suppose Susan is tall and John is happy.

Then, in particular, John is happy.

Also, Susan is tall.

So John is happy and Susan is tall.

Therefore we have shown that if Susan is tall and John is happy, then John is happy and Susan is tall.

However, when we *read* natural deduction proofs, we often read them backward. First, we look at the bottom to see what is being proved. Then we consider the rule that is used to prove it, and see what premises the rule demands. Then we look to see how those claims are proved, and so on. Similarly, when we *construct* a natural deduction proof, we typically work backward as well: we start with the claim we are trying to prove, put that at the bottom, and look for rules to apply.

At times that process breaks down. Suppose we are left with a goal that is a single propositional variable, A . There are no introduction rules that can be applied, so, unless A is a hypothesis, it has to come from an elimination rule. But that underspecifies the problem: perhaps the A comes from applying the and-elimination rule to $A \wedge B$, or from applying the or-elimination rule to C and $C \rightarrow A$. At that point, we look to the hypotheses, and start working forward. If, for example, our hypotheses are C and $C \rightarrow A \wedge B$, we would then work forward to obtain $A \wedge B$ and A .

There is thus a general heuristic for proving theorems in natural deduction:

1. Start by working backward from the conclusion, using the introduction rules. For example, if you are trying to prove a statement of the form $A \rightarrow B$, add A to your list of hypotheses and try to derive B . If you are trying to prove a statement of the form $A \wedge B$, use the and-introduction rule to reduce your task to proving A , and then proving B .
2. When you have run out things to do in the first step, use elimination rules to work forward. If you have hypotheses $A \rightarrow B$ and A , apply modus ponens to derive B . If you have a hypothesis $A \vee B$, use or-elimination to split on cases, considering A in one case and B in the other.

In [Chapter 5](#) we will add one more element to this list: if all else fails, try a proof by contradiction.

The tension between forward and backward reasoning is found in informal arguments as well, in mathematics and elsewhere. When we prove a theorem, we typically reason forward, using assumptions, hypotheses, definitions, and background knowledge. But we also keep the goal in mind, and that helps us make sense of the forward steps.

When we turn to interactive theorem proving, we will see that *Lean* has mechanisms to support both forward and backward reasoning. These form a bridge between informal styles of argumentation and the natural deduction model, and thereby provide a clearer picture of what is going on.

Another confusing feature of natural deduction proofs is that every hypothesis has a *scope*, which is to say, there are only certain points in the proof where an assumption is available for use. Of course, this is also a feature of informal mathematical arguments. Suppose a paragraph begins “Let x be any number less than 100,” argues that x has at most five prime factors, and concludes “thus we have shown that every number less than 100 has at most five factors.” The reference “ x ”, and the assumption that it is less than 100, is only active within the scope of the paragraph. If the next paragraph begins with the phrase “Now suppose x is any number greater than 100,” then, of course, the assumption that x is less than 100 no longer applies.

In natural deduction, a hypothesis is available from the point where it is assumed until the point where it is canceled. We will see that interactive theorem proving languages also have mechanisms to determine the scope of references and hypotheses, and that these, too, shed light on scoping issues in informal mathematics.

3.4 Reasoning by Cases

The rule for eliminating a disjunction is confusing, but we can make sense of it with an example. Consider the following informal argument:

George is either at home or on campus.

If he is at home, he is studying.

If he is on campus, he is with his friends.

Therefore, George is either studying or with his friends.

Let A be the statement that George is at home, let B be the statement that George is on campus, let C be the statement that George is studying, and let D be the statement the George is with his friends. Then the argument above has the following pattern: from $A \vee B$, $A \rightarrow C$, and $B \rightarrow D$, conclude $C \vee D$. In natural deduction, we cannot get away with drawing this conclusion in a single step, but it does not take too much work to flesh it out into a proper proof. Informally, we have to argue as follows.

Georges is either at home or on campus.

Case 1: Suppose he is at home. We know that if he is at home, then he is studying. So, in this case, he is studying. Therefore, in this case, he is either studying or with his friends.

Case 2: Suppose he is on campus. We know that if he is on campus, then he is with his friends. So, in this case, he is with his friends. Therefore, in this case, he is either studying or with his friends.

Either way, George is either studying or with his friends.

The natural deduction proof looks as follows:

$$\frac{A \vee B \quad \frac{A \rightarrow C \quad \overline{A}^1}{C} \quad \frac{B \rightarrow D \quad \overline{B}^1}{D}}{\frac{C \vee D}{C \vee D}^1}^1$$

You should think about how the structure of this proof reflects the informal case-based argument above it.

For another example, here is a proof of $A \wedge (B \vee C) \rightarrow (A \wedge B) \vee (A \wedge C)$:

$$\frac{\frac{A \wedge (B \vee C)}{B \vee C}^2 \quad \frac{\frac{A \wedge (B \vee C)}{A}^2 \quad \overline{B}^1}{A \wedge B} \quad \frac{\frac{A \wedge (B \vee C)}{A}^2 \quad \overline{C}^1}{A \wedge C}}{\frac{(A \wedge B) \vee (A \wedge C)}{(A \wedge B) \vee (A \wedge C)}^1}^2$$

3.5 Some Logical Identities

Two propositional formulas, A and B , are said to be *logically equivalent* if $A \leftrightarrow B$ is provable. Logical equivalences are similar to identities like $x + y = y + x$ that occur in algebra. In particular, one can show that if two formulas are equivalent, then one can substitute one for the other in any formula, and the results will also be equivalent. (Some proof systems take this to be a basic rule, and interactive theorem provers can accommodate it, but we will *not* take it to be a fundamental rule of natural deduction.)

For reference, the following list contains some commonly used propositional equivalences, along with some noteworthy formulas. Think about why, intuitively, these formulas should be true.

1. Commutativity of \wedge : $A \wedge B \leftrightarrow B \wedge A$
2. Commutativity of \vee : $A \vee B \leftrightarrow B \vee A$
3. Associativity of \wedge : $(A \wedge B) \wedge C \leftrightarrow A \wedge (B \wedge C)$
4. Associativity of \vee : $(A \vee B) \vee C \leftrightarrow A \vee (B \vee C)$
5. Distributivity of \wedge over \vee : $A \wedge (B \vee C) \leftrightarrow (A \wedge B) \vee (A \wedge C)$
6. Distributivity of \vee over \wedge : $A \vee (B \wedge C) \leftrightarrow (A \vee B) \wedge (A \vee C)$
7. $(A \rightarrow (B \rightarrow C)) \leftrightarrow (A \wedge B \rightarrow C)$.
8. $(A \rightarrow B) \rightarrow ((B \rightarrow C) \rightarrow (A \rightarrow C))$
9. $((A \vee B) \rightarrow C) \leftrightarrow (A \rightarrow C) \wedge (B \rightarrow C)$
10. $\neg(A \vee B) \leftrightarrow \neg A \wedge \neg B$
11. $\neg(A \wedge B) \leftrightarrow \neg A \vee \neg B$
12. $\neg(A \wedge \neg A)$
13. $\neg(A \rightarrow B) \leftrightarrow A \wedge \neg B$
14. $\neg A \rightarrow (A \rightarrow B)$
15. $(\neg A \vee B) \leftrightarrow (A \rightarrow B)$
16. $A \vee \perp \leftrightarrow A$
17. $A \wedge \perp \leftrightarrow \perp$
18. $A \vee \neg A$
19. $\neg(A \leftrightarrow \neg A)$

20. $(A \rightarrow B) \leftrightarrow (\neg B \rightarrow \neg A)$
21. $(A \rightarrow C \vee D) \rightarrow ((A \rightarrow C) \vee (A \rightarrow D))$
22. $((A \rightarrow B) \rightarrow A) \rightarrow A$

All of these can be derived in natural deduction using the fundamental rules listed in [Section 3.1](#). But some of them require the use of the *reductio ad absurdum* rule, or proof by contradiction, which we have not yet discussed in detail. We will discuss the use of this rule, and other patterns of classical logic, in the [Chapter 5](#).

3.6 Exercises

When constructing proofs in natural deduction, use *only* the list of rules given in [Section 3.1](#).

1. Give a natural deduction proof of $\neg(A \wedge B) \rightarrow (A \rightarrow \neg B)$.
2. Give a natural deduction proof of $(A \rightarrow C) \wedge (B \rightarrow \neg C) \rightarrow \neg(A \wedge B)$.
3. Give a natural deduction proof of $(A \wedge B) \rightarrow ((A \rightarrow C) \rightarrow \neg(B \rightarrow \neg C))$.
4. Take another look at Exercise 3 in the last chapter. Using propositional variables A , B , and C for “Alan likes kangaroos,” “Betty likes frogs” and “Carl likes hamsters,” respectively, express the three hypotheses in the previous problem as symbolic formulas, and then derive a contradiction from them in natural deduction.
5. Give a natural deduction proof of $A \vee B \rightarrow B \vee A$.
6. Give a natural deduction proof of $\neg A \wedge \neg B \rightarrow \neg(A \vee B)$
7. Give a natural deduction proof of $\neg(A \wedge B)$ from $\neg A \vee \neg B$. (You do not need to use proof by contradiction.)
8. Give a natural deduction proof of $\neg(A \leftrightarrow \neg A)$.
9. Give a natural deduction proof of $(\neg A \leftrightarrow \neg B)$ from hypothesis $A \leftrightarrow B$.

PROPOSITIONAL LOGIC IN LEAN

In this chapter, you will learn how to write proofs in Lean. We will start with a purely mechanical translation that will enable you to represent any natural deduction proof in Lean. We will see, however, that such a style of writing proofs is not very intuitive, nor does it yield very readable proofs. It also does not scale well.

We will then consider some mechanisms that Lean offers that support a more forward-directed style of argumentation. Since these proofs look more like informal proofs but can be directly translated to natural deduction, they will help us understand the relationship between the two.

4.1 Expressions for Propositions and Proofs

At its core, Lean is what is known as a *type checker*. This means that we can write expressions and ask the system to check that they are well formed, and also ask the system to tell us what type of object they denote. Try this:

```
variables A B C : Prop
#check A ∧ ¬ B → C
```

In the online version of this text, you can press the “try it!” button to copy the example to an editor window, and then hover over the markers on the text to read the messages.

In the example, we declare three variables ranging over propositions, and ask Lean to check the expression $A \wedge \neg B \rightarrow C$. The output of the `#check` command is `A ∧ ¬ B → C : Prop`, which asserts that $A \wedge \neg B \rightarrow C$ is of type `Prop`. In Lean, every well-formed expression has a type.

The logical connectives are rendered in unicode. The following chart shows you how you can type these symbols in the editor, and also provides ascii equivalents, for the purists among you.

Unicode	Ascii	Emacs
	true	
	false	
¬	not	<code>\not, \neg</code>
∧	∧	<code>\and</code>
∨	∨	<code>\or</code>
→	->	<code>\to, \r, \imp</code>
↔	<->	<code>\iff, \lr</code>
∀	forall	<code>\all</code>
∃	exists	<code>\ex</code>
λ	fun	<code>\lam, \fun</code>
≠	≠	<code>\ne</code>

So far, we have only talked about the first seven items on the list. We will discuss the quantifiers, lambda, and equality later. Try typing some expressions and checking them on your own. You should try changing one of the variables in the example above to D, or inserting a nonsense symbol into the expression, and take a look at the error message that Lean returns.

In addition to declaring variables, if P is any expression of type Prop, we can declare the hypothesis that P is true:

```
variables A B : Prop
variable h : A ∧ ¬ B

#check h
```

Formally, what is going on is that any proposition can be viewed as a type, namely, the type of proofs of that proposition. A hypothesis, or premise, is just a variable of that type. Building proofs is then a matter of writing down expressions of the correct type. For example, if P is any expression of type A ∧ B, then `and.left P` is an expression of type A, and `and.right P` is an expression of type B. In other words, if P is a proof of A ∧ B, and `and.left P` is a name for the proof you get by applying the left elimination rule for and:

$$\frac{\begin{array}{c} \vdots \\ P \end{array}}{A \wedge B} \\ A$$

Similarly, `and.right P` is the proof of B you get by applying the right elimination rule. So, continuing the example above, we can write

```
variables A B : Prop
variable h : A ∧ ¬ B

#check and.left h
#check and.right h
```

The two expressions represent, respectively, these two proofs:

$$\frac{A \wedge \neg B}{A}^h \quad \frac{A \wedge \neg B}{\neg B}^h$$

Notice that in this way of representing natural deduction proofs, there are no “free floating” hypotheses. Every hypothesis has a label. In Lean, we will typically use expressions like `h`, `h1`, `h2`, ... to label hypotheses, but you can use any identifier you want.

If `h1` is a proof of A and `h2` is a proof of B, then `and.intro h1 h2` is a proof of A ∧ B. So we can continue the example above:

```
variables A B : Prop
variable h : A ∧ ¬ B

#check and.intro (and.right h) (and.left h)
```

This corresponds to the following proof:

$$\frac{\frac{A \wedge \neg B}{\neg B}^h \quad \frac{A \wedge \neg B}{A}^h}{\neg B \wedge A}$$

What about implication? The elimination rule is easy: if P_1 is a proof of $A \rightarrow B$ and P_2 is a proof of A then $P_1 P_2$ is a proof of B . Notice that we do not even need to name the rule: you just write P_1 followed by P_2 , as though you are applying the first to the second. If P_1 and P_2 are compound expressions, put parentheses around them to make it clear where each one begins and ends.

```
variables A B C D : Prop

variable h1 : A → (B → C)
variable h2 : D → A
variable h3 : D
variable h4 : B

#check h2 h3
#check h1 (h2 h3)
#check (h1 (h2 h3)) h4
```

Lean adopts the convention that applications associate to the left, so that an expression $h1 h2 h3$ is interpreted as $(h1 h2) h3$. Implications associate to the *right*, so that $A \rightarrow B \rightarrow C$ is interpreted as $A \rightarrow (B \rightarrow C)$. This may seem funny, but it is a convenient way to represent implications that take multiple hypotheses, since an expression $A \rightarrow B \rightarrow C \rightarrow D \rightarrow E$ means that E follows from $A, B, C,$ and D . So the example above could be written as follows:

```
variables A B C D : Prop

variable h1 : A → B → C
variable h2 : D → A
variable h3 : D
variable h4 : B

#check h2 h3
#check h1 (h2 h3)
#check h1 (h2 h3) h4
```

Notice that parentheses are still needed in the expression $h1 (h2 h3)$.

The implication introduction rule is the tricky one, because it can cancel a hypothesis. In terms of Lean expressions, the rule translates as follows. Suppose A and B have type `Prop`, and, assuming h is the premise that A holds, P is proof of B , possibly involving h . Then the expression `assume h : A, P` is a proof of $A \rightarrow B$. For example, we can construct a proof of $A \rightarrow A \wedge A$ as follows:

```
variable A : Prop

#check (assume h : A, and.intro h h)
```

Notice that we no longer have to declare A as a premise. The word `assume` makes the premise local to the expression in parentheses, and after the assumption is made, we can refer to h . Given the assumption $h : A$, `and.intro h h` is a proof of $A \wedge A$, and so the expression `assume h : A, and.intro h h` is a proof of $A \rightarrow A \wedge A$. In this case, we could leave out the parentheses because the expression is unambiguous:

```
variable A : Prop

#check assume h : A, and.intro h h
```

Above, we proved $\neg B \wedge A$ from the premise $A \wedge \neg B$. We can instead obtain a proof of $A \wedge \neg B \rightarrow \neg B \wedge A$ as follows:

```
variables A B : Prop
#check (assume h : A ∧ ¬ B, and.intro (and.right h) (and.left h))
```

All we did was move the premise into a local `assume`.

(By the way, the `assume` command is just alternative syntax for the lambda symbol, so we could also have written this:

```
variables A B : Prop
#check (λ h : A ∧ ¬ B, and.intro (and.right h) (and.left h))
```

You will learn more about the lambda symbol later.)

4.2 More commands

Let us introduce a new Lean command, `example`. This command tells Lean that you are about to prove a theorem, or, more generally, write down an expression of the given type. It should then be followed by the proof or expression itself.

```
variables A B : Prop

example : A ∧ ¬ B → ¬ B ∧ A :=
assume h : A ∧ ¬ B,
and.intro (and.right h) (and.left h)
```

When given this command, Lean checks the expression after the `:=` and makes sure it has the right type. If so, it accepts the expression as a valid proof. If not, it raises an error.

Because the `example` command provides information as to the type of the expression that follows (in this case, the proposition being proved), it sometimes enables us to omit other information. For example, we can leave off the type of the assumption:

```
variables A B : Prop

example : A ∧ ¬ B → ¬ B ∧ A :=
assume h,
and.intro (and.right h) (and.left h)
```

Because Lean knows we are trying to prove an implication with premise $A \wedge \neg B$, it can infer that when we write `assume h`, the identifier `h` labels the assumption $A \wedge \neg B$.

We can also go in the other direction, and provide the system with *more* information, with the word `show`. If A is a proposition and P is a proof, the expression “`show A, from P`” means the same thing as P alone, but it signals the intention that P is a proof of A . When Lean checks this expression, it confirms that P really is a proof of A , before parsing the expression surrounding it. So, in our example, we could also write:

```
variables A B : Prop

example : A ∧ ¬ B → ¬ B ∧ A :=
assume h : A ∧ ¬ B,
show ¬ B ∧ A, from and.intro (and.right h) (and.left h)
```

We could even annotate the smaller expressions `and.right h` and `and.left h`, as follows:

```

variables A B : Prop

example : A ∧ ¬ B → ¬ B ∧ A :=
assume h : A ∧ ¬ B,
show ¬ B ∧ A, from and.intro
  (show ¬ B, from and.right h)
  (show A, from and.left h)

```

This is a good place to mention that Lean generally ignores whitespace, like indentation and returns. We could have written the entire example on a single line. In general, we will adopt conventions for indentation and line breaks that show the structure of proofs and make them easier to read.

Although in the examples above the `show` commands were not necessary, there are a number of good reasons to use this style. First, and perhaps most importantly, it makes the proofs easier for us humans to read. Second, it makes the proofs easier to *write*: if you make a mistake in a proof, it is easier for Lean to figure out where you went wrong and provide a meaningful error message if you make your intentions clear. Finally, proving information in the `show` clause often makes it possible for you to omit information in other places, since Lean can infer that information from your stated intentions.

There are notational variants. Rather than declare variables and premises beforehand, you can also present them as “arguments” to the example, followed by a colon:

```

example (A B : Prop) : A ∧ ¬ B → ¬ B ∧ A :=
assume h : A ∧ ¬ B,
show ¬ B ∧ A, from and.intro (and.right h) (and.left h)

```

There are two more tricks that can help you write proofs in Lean. The first is using `sorry`, which is a magical term in Lean which provides a proof of anything at all. It is also known as “cheating.” But cheating can help you construct legitimate proofs incrementally: if Lean accepts a proof with `sorry`’s, the parts of the proof you have written so far have passed Lean’s checks for correctness. All you need to do is replace each `sorry` with a real proof to complete the task.

```

variables A B : Prop

example : A ∧ ¬ B → ¬ B ∧ A :=
assume h, sorry

example : A ∧ ¬ B → ¬ B ∧ A :=
assume h, and.intro sorry sorry

example : A ∧ ¬ B → ¬ B ∧ A :=
assume h, and.intro (and.right h) sorry

example : A ∧ ¬ B → ¬ B ∧ A :=
assume h, and.intro (and.right h) (and.left h)

```

The second trick is the use of *placeholders*, represented by the underscore symbol. When you write an underscore in an expression, you are asking the system to try to fill in the value for you. This falls short of calling full-blown automation to prove a theorem; rather, you are asking Lean to infer the value from the context. If you use an underscore where a proof should be, Lean typically will *not* fill in the proof, but it will give you an error message that tells you what is missing. This will help you write proof terms incrementally, in a backward-driven fashion. In the example above, try replacing each `sorry` by an underscore, `_`, and take a look at the resulting error messages. In each case, the error tells you what needs to be filled in, and the variables and hypotheses that are available to you at that stage.

One more tip: if you want to delimit the scope of variables or premises introduced with the `variables` command, put them in a block that begins with the word `section` and ends with the word `end`. We will use this mechanism below.

4.3 Building Natural Deduction Proofs

In this section, we describe a mechanical translation from natural deduction proofs, by giving a translation for each natural deduction rule. We have already seen some of the correspondences, but we repeat them all here, for completeness.

4.3.1 Implication

We have already explained that implication introduction is implemented with `assume`, and implication elimination is written as application.

```
variables A B : Prop

example : A → B :=
assume h : A,
show B, from sorry

section
  variable h1 : A → B
  variable h2 : A

  example : B := h1 h2
end
```

Since every example begins by declaring the necessary propositional variables, we will henceforth suppress that declaration in the text.

4.3.2 Conjunction

We have already seen that and-introduction is implemented with `and.intro`, and the elimination rules are `and.left` and `and.right`.

```
section
  variables (h1 : A) (h2 : B)

  example : A ∧ B := and.intro h1 h2
end

section
  variable h : A ∧ B

  example : A := and.left h
  example : B := and.right h
end
```

4.3.3 Disjunction

The or-introduction rules are given by `or.inl` and `or.inr`.

```
section
  variable h : A

  example : A ∨ B := or.inl h
end

section
  variable h : B

  example : A ∨ B := or.inr h
end
```

The elimination rule is the tricky one. To prove C from $A \vee B$, you need three arguments: a proof h of $A \vee B$, a proof of C from A , and a proof of C from B . Using line breaks and indentation to highlight the structure as a proof by cases, we can write it with the following form:

```
section
  variable h : A ∨ B
  variables (ha : A → C) (hb : B → C)
  example : C :=
  or.elim h
    (assume h1 : A,
     show C, from ha h1)
    (assume h1 : B,
     show C, from hb h1)
end
```

Notice that we can reuse the label `h1` in each branch, since, conceptually, the two branches are disjoint.

4.3.4 Negation

Internally, negation $\neg A$ is defined by $A \rightarrow \text{false}$, which you can think of as saying that A implies something impossible. The rules for negation are therefore similar to the rules for implication. To prove $\neg A$, assume A and derive a contradiction.

```
section
  example : ¬ A :=
  assume h : A,
  show false, from sorry
end
```

If you have proved a negation $\neg A$, you can get a contradiction by applying it to a proof of A .

```
section
  variable h1 : ¬ A
  variable h2 : A

  example : false := h1 h2
end
```

4.3.5 Truth and falsity

The *ex falso* rule is called `false.elim`:

```
section
  variable h : false

  example : A := false.elim h
end
```

There isn't much to say about `true` beyond the fact that it is trivially true:

```
example : true := trivial
```

4.3.6 Bi-implication

The introduction rule for “if and only if” is `iff.intro`.

```
example : A ↔ B :=
iff.intro
  (assume h : A,
   show B, from sorry)
  (assume h : B,
   show A, from sorry)
```

As usual, we have chosen indentation to make the structure clear. Notice that the same label, `h`, can be used on both branches, with a different meaning in each, because the scope of an `assume` is limited to the expression in which it appears.

The elimination rules are `iff.elim_left` and `iff.elim_right`:

```
section
  variable h1 : A ↔ B
  variable h2 : A

  example : B := iff.elim_left h1 h2
end

section
  variable h1 : A ↔ B
  variable h2 : B

  example : A := iff.elim_right h1 h2
end
```

Lean recognizes the abbreviation `iff.mp` for `iff.elim_left`, where “mp” stands for “modus ponens”. Similarly, you can use `iff.mpr`, for “modus ponens reverse”, instead of `iff.elim_right`.

4.3.7 Reductio ad absurdum (proof by contradiction)

Finally, there is the rule for proof by contradiction, which we will discuss in greater detail in [Chapter 5](#). It is included for completeness here.

The rule is called `by_contradiction`. It has one argument, which is a proof of `false` from $\neg A$. To use the rule, you have to ask Lean to allow classical reasoning, by writing `open classical`. You can do this at the beginning of the file, or any time before using it. If you say `open classical` in a section, it will remain in scope for that section.

```
section
  open classical

  example : A :=
  by_contradiction
    (assume h : ¬ A,
     show false, from sorry)
end
```

4.3.8 Examples

In the last chapter, we constructed the following proof of $A \rightarrow C$ from $A \rightarrow B$ and $B \rightarrow C$:

$$\frac{\frac{\overline{A}^1 \quad A \rightarrow B}{B} \quad B \rightarrow C}{\frac{C}{A \rightarrow C}^1}$$

We can model this in Lean as follows:

```
variables A B C : Prop

variable h1 : A → B
variable h2 : B → C

example : A → C :=
assume h : A,
show C, from h2 (h1 h)
```

Notice that the hypotheses in the natural deduction proof that are not canceled are declared as variables in the Lean version.

We also constructed the following proof:

$$\frac{\frac{\frac{\overline{A \rightarrow (B \rightarrow C)}^2 \quad \frac{\overline{A \wedge B}^1}{A}}{B \rightarrow C}}{\frac{C}{A \wedge B \rightarrow C}^1} \quad \frac{\overline{A \wedge B}^1}{B}}{(A \rightarrow (B \rightarrow C)) \rightarrow (A \wedge B \rightarrow C)}^2$$

Here is how it is written in Lean:

```
example (A B C : Prop) : (A → (B → C)) → (A ∧ B → C) :=
assume h1 : A → (B → C),
assume h2 : A ∧ B,
show C, from h1 (and.left h2) (and.right h2)
```

This works because `and.left h2` is a proof of `A`, and `and.right h2` is a proof of `B`.

Finally, we constructed the following proof of $A \wedge (B \vee C) \rightarrow (A \wedge B) \vee (A \wedge C)$:

$$\frac{\frac{\frac{A \wedge (B \vee C)}{B \vee C}^2}{\frac{\frac{\frac{A \wedge (B \vee C)}{A}^2}{A \wedge B} \quad \frac{B}{B}^1}{(A \wedge B) \vee (A \wedge C)} \quad \frac{\frac{\frac{A \wedge (B \vee C)}{A}^2}{A \wedge C} \quad \frac{C}{C}^1}{(A \wedge B) \vee (A \wedge C)}^1}{(A \wedge B) \vee (A \wedge C)}^1}{(A \wedge (B \vee C)) \rightarrow ((A \wedge B) \vee (A \wedge C))}^2$$

Here is a version in Lean:

```
example (A B C : Prop) : A ∧ (B ∨ C) → (A ∧ B) ∨ (A ∧ C) :=
assume h1 : A ∧ (B ∨ C),
or.elim (and.right h1)
  (assume h2 : B,
    show (A ∧ B) ∨ (A ∧ C),
    from or.inl (and.intro (and.left h1) h2))
  (assume h2 : C,
    show (A ∧ B) ∨ (A ∧ C),
    from or.inr (and.intro (and.left h1) h2))
```

In fact, bearing in mind that `assume` is alternative syntax for the symbol λ , and that Lean can often infer the type of an assumption, we can make the proof remarkably brief:

```
example (A B C : Prop) : A ∧ (B ∨ C) → (A ∧ B) ∨ (A ∧ C) :=
λ h1, or.elim (and.right h1)
  (λ h2, or.inl (and.intro (and.left h1) h2))
  (λ h2, or.inr (and.intro (and.left h1) h2))
```

The proof is cryptic, though. Using such a style makes proofs hard to write, read, understand, maintain, and debug. In the next section we will describe a remarkably simple device that makes it much easier to understand what is going on.

4.4 Forward Reasoning

Lean supports forward reasoning by allowing you to write proofs using the `have` command.

```
variables A B C : Prop

variable h1 : A → B
variable h2 : B → C

example : A → C :=
assume h : A,
have h3 : B, from h1 h,
show C, from h2 h3
```

Writing a proof with `have h : A, from P, ... h ...` has the same effect as writing `... P ...`. This `have` command checks that `P` is a proof of `A`, and then give you the label `h` to use in place of `P`. Thus the last line of the previous proof can be thought of as abbreviating `show C, from h2 (h1 h)`, since `h3` abbreviates `h1 h`. Such abbreviations can make a big difference, especially when the proof `P` is very long.

There are a number of advantages to using `have`. For one thing, it makes the proof more readable: the example above states `B` explicitly as an auxiliary goal. It can also save repetition: `h3` can be used repeatedly after it is introduced, without duplicating the proof. Finally, it makes it easier to construct and debug the

proof: stating B as an auxiliary goal makes it easier for Lean to deliver an informative error message when the goal is not properly met.

In the last section, we considered the following proof:

```
example (A B C : Prop) : (A → (B → C)) → (A ∧ B → C) :=
assume h1 : A → (B → C),
assume h2 : A ∧ B,
show C, from h1 (and.left h2) (and.right h2)
```

Using `have`, it can be written more perspicuously as follows:

```
example (A B C : Prop) : (A → (B → C)) → (A ∧ B → C) :=
assume h1 : A → (B → C),
assume h2 : A ∧ B,
have h3 : A, from and.left h2,
have h4 : B, from and.right h2,
show C, from h1 h3 h4
```

We can be even more verbose, and add another line:

```
example (A B C : Prop) : (A → (B → C)) → (A ∧ B → C) :=
assume h1 : A → (B → C),
assume h2 : A ∧ B,
have h3 : A, from and.left h2,
have h4 : B, from and.right h2,
have h5 : B → C, from h1 h3,
show C, from h5 h4
```

Adding more information doesn't always make a proof more readable; when the individual expressions are small and easy enough to understand, spelling them out in detail can introduce clutter. As you learn to use Lean, you will have to develop your own style, and use your judgment to decide which steps to make explicit.

Here is how some of the basic inferences look, when expanded with `have`. In the and-introduction rule, it is a matter showing each conjunct first, and then putting them together:

```
example (A B : Prop) : A ∧ B → B ∧ A :=
assume h1 : A ∧ B,
have h2 : A, from and.left h1,
have h3 : B, from and.right h1,
show B ∧ A, from and.intro h3 h2
```

Compare that with this version, which instead states first that we will use the `and.intro` rule, and then makes the two resulting goals explicit:

```
example (A B : Prop) : A ∧ B → B ∧ A :=
assume h1 : A ∧ B,
show B ∧ A, from
  and.intro
    (show B, from and.right h1)
    (show A, from and.left h1)
```

Once again, at issue is only readability. Lean does just fine with the following short version:

```
example (A B : Prop) : A ∧ B → B ∧ A :=
λ h, and.intro (and.right h) (and.left h)
```

When using the or-elimination rule, it is often clearest to state the relevant disjunction explicitly:

```
example (A B C : Prop) : C :=
have h : A ∨ B, from sorry,
show C, from or.elim h
  (assume h1 : A,
   show C, from sorry)
  (assume h2 : B,
   show C, from sorry)
```

Here is a have-structured presentation of an example from the previous section:

```
example (A B C : Prop) : A ∧ (B ∨ C) → (A ∧ B) ∨ (A ∧ C) :=
assume h1 : A ∧ (B ∨ C),
have h2 : A, from and.left h1,
have h3 : B ∨ C, from and.right h1,
show (A ∧ B) ∨ (A ∧ C), from
  or.elim h3
    (assume h4 : B,
     have h5 : A ∧ B, from and.intro h2 h4,
     show (A ∧ B) ∨ (A ∧ C), from or.inl h5)
    (assume h4 : C,
     have h5 : A ∧ C, from and.intro h2 h4,
     show (A ∧ B) ∨ (A ∧ C), from or.inr h5)
```

4.5 Definitions and Theorems

Lean allows us to name definitions and theorems for later use. For example, here is a definition of a new “connective”:

```
def triple_and (A B C : Prop) : Prop :=
A ∧ (B ∧ C)
```

As with the `example` command, it does not matter whether the arguments `A`, `B`, and `C` are declared beforehand with the `variables` command, or with the definition itself. We can then apply the definition to any expressions:

```
variables D E F G : Prop

#check triple_and (D ∨ E) (¬ F → G) (¬ D)
```

Later, we will see more interesting examples of definitions, like the following function from natural numbers to natural numbers, which doubles its input:

```
def double (n : ℕ) : ℕ := n + n
```

What is more interesting right now is that Lean also allows us to name theorems, and use them later, as rules of inference. For example, consider the following theorem:

```
theorem and_commute (A B : Prop) : A ∧ B → B ∧ A :=
assume h, and.intro (and.right h) (and.left h)
```

Once we have defined it, we can use it freely:

```

variables C D E : Prop
variable h1 : C ∧ ¬ D
variable h2 : ¬ D ∧ C → E

example : E := h2 (and_commute C (¬ D) h1)

```

It is annoying in this example that we have to give the arguments C and $\neg D$ explicitly, because they are implicit in $h1$. In fact, Lean allows us to tell this to Lean in the definition of `and_commute`:

```

theorem and_commute {A B : Prop} : A ∧ B → B ∧ A :=
assume h, and.intro (and.right h) (and.left h)

```

Here the squiggly braces indicate that the arguments A and B are *implicit*, which is to say, Lean should infer them from the context when the theorem is used. We can then write the following instead:

```

variables C D E : Prop
variable h1 : C ∧ ¬ D
variable h2 : ¬ D ∧ C → E

example : E := h2 (and_commute h1)

```

Indeed, Lean's library has a theorem, `and_comm`, defined in exactly this way.

By the way, we could avoid the `assume` step in the proof of `and_comm` by making the hypothesis into an argument:

```

theorem and_commute {A B : Prop} (h : A ∧ B) : B ∧ A :=
and.intro (and.right h) (and.left h)

```

The two definitions yield the same result.

Definitions and theorems are important in mathematics; they allow us to build up complex theories from fundamental principles. Lean also accepts the word `lemma` instead of `theorem`.

What is interesting is that in interactive theorem proving, we can even define familiar patterns of inference. For example, all of the following inferences were mentioned in the last chapter:

```

namespace hidden

variables {A B : Prop}

theorem or_resolve_left (h1 : A ∨ B) (h2 : ¬ A) : B :=
or.elim h1
  (assume h3 : A, show B, from false.elim (h2 h3))
  (assume h3 : B, show B, from h3)

theorem or_resolve_right (h1 : A ∨ B) (h2 : ¬ B) : A :=
or.elim h1
  (assume h3 : A, show A, from h3)
  (assume h3 : B, show A, from false.elim (h2 h3))

theorem absurd (h1 : ¬ A) (h2 : A) : B :=
false.elim (h1 h2)

end hidden

```

In fact, Lean’s library defines `or.resolve_left`, `or.resolve_right`, and `absurd`. We used the `namespace` command to avoid naming conflicts, which would have raised an error.

When we ask you to prove basic facts from propositional logic in Lean, as with propositional logic, our goal is to have you learn how to use Lean’s primitives. As a result, for those exercises, you should not use facts from the library. As we move towards real mathematics, however, you can use facts from the library more freely.

4.6 Additional Syntax

In this section, we describe some extra syntactic features of Lean, for power users. The syntactic gadgets are often convenient, and sometimes make proofs look prettier.

For one thing, you can use subscripted numbers with a backslash. For example, you can write `h1` by typing `h\1`. The labels are irrelevant to Lean, so the difference is only cosmetic.

Another feature is that you can omit the label in an `assume` statement, providing an “anonymous” hypothesis. You can then refer back to the last anonymous assumption using the keyword `this`:

```
example : A → A ∨ B :=
assume : A,
show A ∨ B, from or.inl this
```

Alternatively, you can refer back to unlabeled assumptions by putting them in French quotes:

```
example : A → B → A ∧ B :=
assume : A,
assume : B,
show A ∧ B, from and.intro <A> <B>
```

You can also use the word `have` without giving a label, and refer back to them using the same conventions. Here is an example that uses these features:

```
theorem my_theorem {A B C : Prop} :
  A ∧ (B ∨ C) → (A ∧ B) ∨ (A ∧ C) :=
assume h : A ∧ (B ∨ C),
have A, from and.left h,
have B ∨ C, from and.right h,
show (A ∧ B) ∨ (A ∧ C), from
  or.elim <B ∨ C>
    (assume : B,
      have A ∧ B, from and.intro <A> <B>,
      show (A ∧ B) ∨ (A ∧ C), from or.inl this)
    (assume : C,
      have A ∧ C, from and.intro <A> <C>,
      show (A ∧ B) ∨ (A ∧ C), from or.inr this)
```

Another trick is that you can write `h.left` and `h.right` instead of `and.left h` and `and.right h` whenever `h` is a conjunction, and you can write `<h1, h2>` instead of `and.intro h1 h2` whenever Lean can figure out that a conjunction is what you are trying to prove. With these conventions, you can write the following:

```
example (A B : Prop) : A ∧ B → B ∧ A :=
assume h : A ∧ B,
show B ∧ A, from <h.right, h.left>
```

This is nothing more than shorthand for the following:

```
example (A B : Prop) : A ∧ B → B ∧ A :=
assume h : A ∧ B,
show B ∧ A, from and.intro (and.right h) (and.left h)
```

Even more concisely, you can write this:

```
example (A B : Prop) : A ∧ B → B ∧ A :=
assume h, ⟨h.right, h.left⟩
```

You can even take apart a conjunction with an `assume`, so that this works:

```
example (A B : Prop) : A ∧ B → B ∧ A :=
assume ⟨h1, h2⟩, ⟨h2, h1⟩
```

Similarly, if `h` is a biconditional, you can write `h.mp` and `h.mpr` instead of `iff.mp h` and `iff.mpr h`, and you can write `⟨h1, h2⟩` instead of `iff.intro h1 h2`. As a result, Lean understands these proofs:

```
example (A B : Prop) : B ∧ (A ↔ B) → A :=
assume ⟨hB, hAB⟩,
hAB.mpr hB

example (A B : Prop) : A ∧ B ↔ B ∧ A :=
⟨assume ⟨h1, h2⟩, ⟨h2, h1⟩, assume ⟨h1, h2⟩, ⟨h2, h1⟩⟩
```

Finally, you can add comments to your proofs in two ways. First, any text after a double-dash `--` until the end of a line is ignored by the Lean processor. Second, any text between `/-` and `-/` denotes a block comment, and is also ignored. You can nest block comments.

```
/- This is a block comment.
   It can fill multiple lines. -/

example (A : Prop) : A → A :=
assume : A,      -- assume the antecedent
show A, from this -- use it to establish the conclusion
```

4.7 Exercises

Prove the following in Lean:

```
variables A B C D : Prop

example : A ∧ (A → B) → B :=
sorry

example : A → ¬ (¬ A ∧ B) :=
sorry

example : ¬ (A ∧ B) → (A → ¬ B) :=
sorry
```

(continues on next page)

(continued from previous page)

```
example (h1 : A ∨ B) (h2 : A → C) (h3 : B → D) : C ∨ D :=  
sorry
```

```
example (h : ¬ A ∧ ¬ B) : ¬ (A ∨ B) :=  
sorry
```

```
example : ¬ (A ↔ ¬ A) :=  
sorry
```

CLASSICAL REASONING

If we take all the rules of propositional logic we have seen so far and exclude *reductio ad absurdum*, or proof by contradiction, we have what is known as *intuitionistic logic*. In intuitionistic logic, it is possible to view proofs in computational terms: a proof of $A \wedge B$ is a proof of A paired with a proof of B , a proof of $A \rightarrow B$ is a procedure which transforms evidence for A into evidence for B , and a proof of $A \vee B$ is a proof of one or the other, tagged so that we know which is the case. The *ex falso* rule makes sense only because we expect that there is no proof of falsity; it is like the empty data type.

Proof by contradiction does not fit in well with this world view: from a proof of a contradiction from $\neg A$, we are supposed to magically produce a proof of A . We will see that with proof by contradiction, we can prove the following law, known as the *law of the excluded middle*: $\forall A, A \vee \neg A$. From a computational perspective, this says that for every A we can decide whether or not A is true.

Classical reasoning does introduce a number of principles into logic, however, that can be used to simplify reasoning. In this chapter, we will consider these principles, and see how they follow from the basic rules.

5.1 Proof by Contradiction

Remember that in natural deduction, proof by contradiction is expressed by the following pattern:

$$\frac{\overline{\neg A}^1}{\vdots} \frac{\perp}{A}^1$$

The assumption $\neg A$ is canceled at the final inference.

In Lean, the inference is named `by_contradiction`, and since it is a classical rule, we have to use the command `open classical` before it is available. Once we do so, the pattern of inference is expressed as follows:

```
open classical

variable (A : Prop)

example : A :=
by_contradiction
  (assume h : ¬ A,
   show false, from sorry)
```

One of the most important consequences of this rule is a classical principle that we mentioned above, namely, the *law of the excluded middle*, which asserts that the following holds for all A : $A \vee \neg A$. In Lean we denote

this law by **em**. In mathematical arguments, one often splits a proof into two cases, assuming first A and then $\neg A$. Using the elimination rule for disjunction, this is equivalent to using $A \vee \neg A$, which is the excluded middle principle for this particular A .

Here is a proof of **em**, in natural deduction, using proof by contradiction:

$$\frac{\frac{\frac{\perp}{\neg A}^1}{A \vee \neg A} \quad \frac{\overline{A}^1}{A \vee \neg A}}{\neg(A \vee \neg A)}^2 \quad \frac{\perp}{A \vee \neg A}^1}{\neg(A \vee \neg A)}^1$$

Here is the same proof rendered in Lean:

```
open classical

variable (A : Prop)

example : A ∨ ¬ A :=
by_contradiction
  (assume h1 : ¬ (A ∨ ¬ A),
   have h2 : ¬ A, from
     assume h3 : A,
     have h4 : A ∨ ¬ A, from or.inl h3,
     show false, from h1 h4,
   have h5 : A ∨ ¬ A, from or.inr h2,
   show false, from h1 h5)
```

The principle is known as the law of the excluded middle because it says that a proposition A is either true or false; there is no middle ground. As a result, the theorem is named **em** in the Lean library. For any proposition A , **em** A denotes a proof of $A \vee \neg A$, and you are free to use it any time **classical** is open:

```
open classical

example (A : Prop) : A ∨ ¬ A :=
or.elim (em A)
  (assume : A, or.inl this)
  (assume : ¬ A, or.inr this)
```

Or even more simply:

```
open classical

example (A : Prop) : A ∨ ¬ A :=
em A
```

In fact, we can go in the other direction, and use the law of the excluded middle to justify proof by contradiction. You are asked to do this in the exercises.

Proof by contradiction is also equivalent to the principle $\neg\neg A \leftrightarrow A$. The implication from right to left holds intuitionistically; the other implication is classical, and is known as *double-negation elimination*. Here is a proof in natural deduction:

$$\frac{\frac{\frac{}{\neg\neg A} 2}{\perp} 1 \quad \frac{}{\neg A} 1}{A} 1 \quad \frac{\frac{}{\neg A} 1 \quad \frac{}{A} 2}{\perp} 1}{\neg\neg A} 1}{\neg\neg A \leftrightarrow A} 2$$

And here is the corresponding proof in Lean:

```
open classical

example (A : Prop) : ¬ ¬ A ↔ A :=
iff.intro
  (assume h1 : ¬ ¬ A,
   show A, from by_contradiction
    (assume h2 : ¬ A,
     show false, from h1 h2))
  (assume h1 : A,
   show ¬ ¬ A, from assume h2 : ¬ A, h2 h1)
```

In the next section, we will derive a number of classical rules and equivalences. These are tricky to prove. In general, to use classical reasoning in natural deduction, we need to extend the general heuristic presented in Section 3.3 as follows:

1. First, work backward from the conclusion, using the introduction rules.
2. When you have run out things to do in the first step, use elimination rules to work forward.
3. If all else fails, use a proof by contradiction.

Sometimes a proof by contradiction is necessary, but when it isn't, it can be less informative than a direct proof. Suppose, for example, we want to prove $A \wedge B \wedge C \rightarrow D$. In a direct proof, we assume A , B , and C , and work towards D . Along the way, we will derive other consequences of A , B , and C , and these may be useful in other contexts. If we use proof by contradiction, on the other hand, we assume A , B , C , and $\neg D$, and try to prove \perp . In that case, we are working in an inconsistent context; any auxiliary results we may obtain that way are subsumed by the fact that ultimately \perp is a consequence of the hypotheses.

5.2 Some Classical Principles

We have already seen that $A \vee \neg A$ and $\neg\neg A \leftrightarrow A$ are two important theorems of classical propositional logic. In this section we will provide some more theorems, rules, and equivalences. Some will be proved here, but most will be left to you in the exercises. In ordinary mathematics, these are generally used without comment. It is nice to know, however, that they can all be justified using the basic rules of classical natural deduction.

If $A \rightarrow B$ is any implication, the assertion $\neg B \rightarrow \neg A$ is known as the *contrapositive*. Every implication implies its contrapositive, and the other direction is true classically:

$$\frac{\frac{\frac{}{\neg B \rightarrow \neg A}}{\neg A} \quad \frac{}{\neg B} 1}{A} 2}{\frac{\frac{}{\perp} 1}{B} 1}{A \rightarrow B} 2}$$

Here is another example. Intuitively, asserting “if A then B ” is equivalent to saying that it cannot be the case that A is true and B is false. Classical reasoning is needed to get us from the second statement to the first.

$$\frac{\frac{\frac{\overline{\neg(A \wedge \neg B)}}{3} \quad \frac{\overline{A}^2 \quad \overline{\neg B}^1}{A \wedge \neg B}}{\frac{\perp}{B}^1}{A \rightarrow B}^2}{\neg(A \wedge \neg B) \rightarrow (A \rightarrow B)}^3$$

Here are the same proofs, rendered in Lean:

```
open classical

variables (A B : Prop)

example (h : ¬ B → ¬ A) : A → B :=
assume h1 : A,
show B, from
  by_contradiction
    (assume h2 : ¬ B,
     have h3 : ¬ A, from h h2,
     show false, from h3 h1)

example (h : ¬ (A ∧ ¬ B)) : A → B :=
assume : A,
show B, from
  by_contradiction
    (assume : ¬ B,
     have A ∧ ¬ B, from and.intro <A> this,
     show false, from h this)
```

Notice that in the second example, we used an anonymous `assume` and an anonymous `have`. We used the brackets `\f<` and `\f>` to write `<A>`, referring back to the first assumption. The first use of the word `this` refers back to the assumption `¬ B`, while the second one refers back to the `have`.

Knowing that we can prove the law of the excluded middle, it is convenient to use it in classical proofs. Here is an example, with a proof of $(A \rightarrow B) \vee (B \rightarrow A)$:

$$\frac{\frac{\overline{B \vee \neg B}}{\frac{\overline{B}^1}{A \rightarrow B} \quad \frac{\overline{\neg B}^1 \quad \overline{B}^2}{\frac{\perp}{A} \quad B \rightarrow A}^2}}{(A \rightarrow B) \vee (B \rightarrow A)} \quad \frac{\overline{B \vee \neg B}}{(A \rightarrow B) \vee (B \rightarrow A)}^1$$

Here is the corresponding proof in Lean:

```
open classical

variables (A B : Prop)

example : (A → B) ∨ (B → A) :=
or.elim (em B)
  (assume h : B,
   have A → B, from
     assume : A,
     show B, from h,
```

(continues on next page)

(continued from previous page)

```

show (A → B) ∨ (B → A),
  from or.inl this)
(assume h : ¬ B,
  have B → A, from
    assume : B,
    have false, from h this,
    show A, from false.elim this,
  show (A → B) ∨ (B → A),
  from or.inr this)
    
```

Using classical reasoning, implication can be rewritten in terms of disjunction and negation:

$$(A \rightarrow B) \leftrightarrow \neg A \vee B.$$

The forward direction requires classical reasoning.

The following equivalences are known as De Morgan's laws:

- $\neg(A \vee B) \leftrightarrow \neg A \wedge \neg B$
- $\neg(A \wedge B) \leftrightarrow \neg A \vee \neg B$

The forward direction of the second of these requires classical reasoning.

Using these identities, we can always push negations down to propositional variables. For example, we have

$$\begin{aligned}
 \neg(\neg A \wedge B \rightarrow C) &\leftrightarrow \neg(\neg(\neg A \wedge B) \vee C) \\
 &\leftrightarrow \neg\neg(\neg A \wedge B) \wedge \neg C \\
 &\leftrightarrow \neg A \wedge B \wedge \neg C.
 \end{aligned}$$

A formula built up from \wedge , \vee , and \neg in which negations only occur at variables is said to be in *negation normal form*.

In fact, using distributivity laws, one can go on to ensure that all the disjunctions are on the outside, so that the formula is a big or of and's of propositional variables and negated propositional variables. Such a formula is said to be in *disjunctive normal form*. Alternatively, all the and's can be brought to the outside. Such a formula is said to be in *conjunctive normal form*. An exercise below, however, shows that putting formulas in disjunctive or conjunctive normal form can make them much longer.

5.3 Exercises

1. Show how to derive the proof-by-contradiction rule from the law of the excluded middle, using the other rules of natural deduction. In other words, assume you have a proof of \perp from $\neg A$. Using $A \vee \neg A$ as a hypothesis, but *without* using the rule RAA, show how you can go on to derive A .
2. Give a natural deduction proof of $\neg(A \wedge B)$ from $\neg A \vee \neg B$. (You do not need to use proof by contradiction.)
3. Construct a natural deduction proof of $\neg A \vee \neg B$ from $\neg(A \wedge B)$. You can do it as follows:
 1. First, prove $\neg B$, and hence $\neg A \vee \neg B$, from $\neg(A \wedge B)$ and A .
 2. Use this to construct a proof of $\neg A$, and hence $\neg A \vee \neg B$, from $\neg(A \wedge B)$ and $\neg(\neg A \vee \neg B)$.
 3. Use this to construct a proof of a contradiction from $\neg(A \wedge B)$ and $\neg(\neg A \vee \neg B)$.

4. Using proof by contradiction, this gives you a proof of $\neg A \vee \neg B$ from $\neg(A \wedge B)$.
4. Give a natural deduction proof of $\neg A \vee B$ from $A \rightarrow B$. You may use the law of the excluded middle.
5. Put $(A \vee B) \wedge (C \vee D) \wedge (E \vee F)$ in disjunctive normal form, that is, write it as a big “or” of multiple “and” expressions.
6. Prove $\neg (A \wedge B) \rightarrow \neg A \vee \neg B$ by replacing the sorry’s below by proofs.

```

open classical
variables {A B C : Prop}

-- Prove  $\neg (A \wedge B) \rightarrow \neg A \vee \neg B$  by replacing the sorry's below
-- by proofs.

lemma step1 (h1 :  $\neg (A \wedge B)$ ) (h2 : A) :  $\neg A \vee \neg B$  :=
have  $\neg B$ , from sorry,
show  $\neg A \vee \neg B$ , from or.inr this

lemma step2 (h1 :  $\neg (A \wedge B)$ ) (h2 :  $\neg (\neg A \vee \neg B)$ ) : false :=
have  $\neg A$ , from
  assume : A,
  have  $\neg A \vee \neg B$ , from step1 h1 <A>,
  show false, from h2 this,
show false, from sorry

theorem step3 (h :  $\neg (A \wedge B)$ ) :  $\neg A \vee \neg B$  :=
by_contradiction
  (assume h' :  $\neg (\neg A \vee \neg B)$ ,
   show false, from step2 h h')

```

7. Also do these:

```

open classical
variables {A B C : Prop}

example (h :  $\neg B \rightarrow \neg A$ ) :  $A \rightarrow B$  :=
sorry

example (h :  $A \rightarrow B$ ) :  $\neg A \vee B$  :=
sorry

```

SEMANTICS OF PROPOSITIONAL LOGIC

Classically, we think of propositional variables as ranging over statements that can be true or false. And, intuitively, we think of a proof system as telling us what propositional formulas *have to* be true, no matter what the variables stand for. For example, the fact that we can prove C from the hypotheses A , B , and $A \wedge B \rightarrow C$ seems to tell us that whenever the hypotheses are true, then C has to be true as well.

Making sense of this involves stepping outside the system and giving an account of truth—more precisely, the conditions under which a propositional formula is true. This is one of the things that symbolic logic was designed to do, and the task belongs to the realm of *semantics*. Formulas and formal proofs are *syntactic* notions, which is to say, they are represented by symbols and symbolic structures. Truth is a *semantic* notion, in that it ascribes a type of *meaning* to certain formulas.

Syntactically, we were able to ask and answer questions like the following:

- Given a set of hypotheses, Γ , and a formula, A , can we derive A from Γ ?
- What formulas can be derived from Γ ?
- What hypotheses are needed to derive A ?

The questions we consider semantically are different:

- Given an assignment of truth values to the propositional variables occurring in a formula A , is A true or false?
- Is there any truth assignment that makes A true?
- Which are the truth assignments that make A true?

In this chapter, we will not provide a fully rigorous mathematical treatment of syntax and semantics. That subject matter is appropriate to a more advanced and focused course on mathematical logic. But we will discuss semantic issues in enough detail to give you a good sense of what it means to think semantically, as well as a sense of how to make pragmatic use of semantic notions.

6.1 Truth Values and Assignments

The first notion we will need is that of a *truth value*. We have already seen two, namely, “true” and “false.” We will use the symbols **T** and **F** to represent these in informal mathematics. These are the values that \top and \perp are intended to denote in natural deduction, and `true` and `false` are intended to denote in Lean.

In this text, we will adopt a “classical” notion of truth, following our discussion in [Section 5](#). This can be understood in various ways, but, concretely, it comes down to this: we will assume that any proposition is either true or false (but, of course, not both). This conception of truth is what underlies the law of the excluded middle, $A \vee \neg A$. Semantically, we read this sentence as saying “either A is true, or $\neg A$ is true.” Since, in our semantic interpretation, $\neg A$ is true exactly when A is false, the law of the excluded middle says that A is either true or false.

The next notion we will need is that of a *truth assignment*, which is simply a function that assigns a truth value to each element of a set of propositional variables. In this section, we will distinguish between propositional variables and arbitrary formulas by using letters P, Q, R, \dots for the former and A, B, C, \dots for the latter. For example, the function v defined by

- $v(P) := \mathbf{T}$
- $v(Q) := \mathbf{F}$
- $v(R) := \mathbf{F}$
- $v(S) := \mathbf{T}$

is a truth assignment for the set of variables $\{P, Q, R, S\}$.

Intuitively, a truth assignment describes a possible “state of the world.” Going back to the Malice and Alice puzzle, let’s suppose the following letters are shorthand for the statements:

- $P :=$ Alice’s brother was the victim
- $Q :=$ Alice was the killer
- $R :=$ Alice was in the bar

In the world described by the solution to the puzzle, the first and third statements are true, and the second is false. So our truth assignment gives the value \mathbf{T} to P and R , and the value \mathbf{F} to Q .

Once we have a truth assignment v to a set of propositional variables, we can extend it to a *valuation function* \bar{v} , which assigns a value of true or false to every propositional formula that depends only on these variables. The function \bar{v} is defined recursively, which is to say, formulas are evaluated from the bottom up, so that value assigned to a compound formula is determined by the values assigned to its components. Formally, the function is defined as follows:

- $\bar{v}(\top) = \mathbf{T}$.
- $\bar{v}(\perp) = \mathbf{F}$.
- $\bar{v}(\ell) = v(\ell)$, where ℓ is any propositional variable.
- $\bar{v}(\neg A) = \mathbf{T}$ if $\bar{v}(A)$ is \mathbf{F} , and vice versa.
- $\bar{v}(A \wedge B) = \mathbf{T}$ if $\bar{v}(A)$ and $\bar{v}(B)$ are both \mathbf{T} , and \mathbf{F} otherwise.
- $\bar{v}(A \vee B) = \mathbf{T}$ if at least one of $\bar{v}(A)$ and $\bar{v}(B)$ is \mathbf{T} ; otherwise \mathbf{F} .
- $\bar{v}(A \rightarrow B) = \mathbf{T}$ if either $\bar{v}(B)$ is \mathbf{T} or $\bar{v}(A)$ is \mathbf{F} , and \mathbf{F} otherwise. (Equivalently, $\bar{v}(A \rightarrow B) = \mathbf{F}$ if $\bar{v}(A)$ is \mathbf{T} and $\bar{v}(B)$ is \mathbf{F} , and \mathbf{T} otherwise.)

The rules for conjunction and disjunction are easy to understand. “ A and B ” is true exactly when A and B are both true; “ A or B ” is true when at least one of A or B is true.

Understanding the rule for implication is trickier. People are often surprised to hear that any if-then statement with a false hypothesis is supposed to be true. The statement “if I have two heads, then circles are squares” may sound like it ought to be false, but by our reckoning, it comes out true. To make sense of this, think about the difference between the two sentences:

- “If I have two heads, then circles are squares.”
- “If I had two heads, then circles would be squares.”

The second sentence is an example of a *counterfactual* implication. It asserts something about how the world might change, if things were other than they actually are. Philosophers have studied counterfactuals for centuries, but mathematical logic is concerned with the first sentence, a *material* implication. The material implication asserts something about the way the world is right now, rather than the way it might

have been. Since it is false that I have two heads, the statement “if I have two heads, then circles are squares” is true.

Why do we evaluate material implication in this way? Once again, let us consider the true sentence “every natural number that is prime and greater than two is odd.” We can interpret this sentence as saying that all of the (infinitely many) sentences in this list are true:

- If 0 is prime and greater than 2, then 0 is odd.
- If 1 is prime and greater than 2, then 1 is odd.
- If 2 is prime and greater than 2, then 2 is odd.
- If 3 is prime and greater than 2, then 3 is odd.
- ...

The first sentence on this list is a lot like our “two heads” example, since both the hypothesis and the conclusion are false. But since it is an instance of a statement that is true in general, we are committed to assigning it the value **T**. The second sentence is a different: the hypothesis is still false, but here the conclusion is true. Together, these tell us that whenever the hypothesis is false, the conditional statement should be true. The fourth sentence has a true hypothesis and a true conclusion. So from the second and fourth sentences, we see that whenever the conclusion is true, the conditional should be true as well. Finally, it seems clear that the sentence “if 3 is prime and greater than 2, then 3 is even” should *not* be true. This pattern, where the hypothesis is true and the conclusion is false, is the only one for which the conditional will be false.

Let us motivate the semantics for material implication another way, using the deductive rules described in the last chapter. Notice that, if B is true, we can prove $A \rightarrow B$ without any assumptions about A :

$$\frac{B}{A \rightarrow B}$$

This follows from the proper reading of the implication introduction rule: given B , one can always infer $A \rightarrow B$, and then cancel an assumption A , *if there is one*. If A was never used in the proof, the conclusion is simply weaker than it needs to be. This inference is validated in Lean:

```
variables A B : Prop
variable hB : B

example : A → B :=
assume hA : A,
  show B, from hB
```

Similarly, if A is false, we can prove $A \rightarrow B$ without any assumptions about B :

$$\frac{\frac{\neg A \quad \overline{A}^1}{\perp} \quad 1}{A \rightarrow B}^1$$

In Lean:

```
variables A B : Prop
variable hnA : ¬ A

example : A → B :=
assume hA : A,
  show B, from false.elim (hnA hA)
```

Finally, if A is true and B is false, we can prove $\neg(A \rightarrow B)$:

$$\frac{\frac{\neg B \quad \frac{\overline{A \rightarrow B}^1 \quad A}{B}}{\perp}}{\neg(A \rightarrow B)}^1$$

Once again, in Lean:

```
variables A B : Prop
variable hA : A
variable hnB : ¬B

example : ¬ (A → B) :=
assume h : A → B,
have hB : B, from h hA,
show false, from hnB hB
```

Now that we have defined the truth of any formula relative to a truth assignment, we can answer our first semantic question: given an assignment v of truth values to the propositional variables occurring in some formula φ , how do we determine whether or not φ is true? This amounts to evaluating $\bar{v}(\varphi)$, and the recursive definition of φ gives a recipe: we evaluate the expressions occurring in φ from the bottom up, starting with the propositional variables, and using the evaluation of an expression's components to evaluate the expression itself. For example, suppose our truth assignment v makes A and B true and C false. To evaluate $(B \rightarrow C) \vee (A \wedge B)$ under v , note that the expression $B \rightarrow C$ comes out false and the expression $A \wedge B$ comes out true. Since a disjunction “false or true” is true, the entire formula is true.

We can also go in the other direction: given a formula, we can attempt to find a truth assignment that will make it true (or false). In fact, we can use Lean to evaluate formulas for us. In the example that follows, you can assign any set of values to the proposition symbols A , B , C , D , and E . When you run Lean on this input, the output of the `eval` statement is the value of the expression.

```
-- Define your truth assignment here
def A := tt
def B := ff
def C := tt
def D := tt
def E := ff

def test (p : Prop) [decidable p] : string :=
if p then "true" else "false"

#eval test ((A ∧ B) ∨ ¬ C)
#eval test (A → D)
#eval test (C → (D ∨ ¬E))
#eval test (¬(A ∧ B ∧ C ∧ D))
```

Try varying the truth assignments, to see what happens. You can add your own formulas to the end of the input, and evaluate them as well. Try to find truth assignments that make each of the formulas tested above evaluate to true. For an extra challenge, try finding a single truth assignment that makes them all true at the same time.

6.2 Truth Tables

The second and third semantic questions we asked are a little trickier than the first. Given a formula A , is there any truth assignment that makes A true? If so, which truth assignments make A true? Instead

of considering one particular truth assignment, these questions ask us to quantify over *all* possible truth assignments.

Of course, the number of possible truth assignments depends on the number of propositional letters we're considering. Since each letter has two possible values, n letters will produce 2^n possible truth assignments. This number grows very quickly, so we'll mostly look at smaller formulas here.

We'll use something called a *truth table* to figure out when, if ever, a formula is true. On the left hand side of the truth table, we'll put all of the possible truth assignments for the present propositional letters. On the right hand side, we'll put the truth value of the entire formula under the corresponding assignment.

To begin with, truth tables can be used to concisely summarize the semantics of our logical connectives:

A	B	$A \wedge B$	A	B	$A \vee B$	A	B	$A \rightarrow B$
T	T	T	T	T	T	T	T	T
T	F	F	T	F	T	T	F	F
F	T	F	F	T	T	F	T	T
F	F	F	F	F	F	F	F	T

We will leave it to you to write the table for $\neg A$, as an easy exercise.

For compound formulas, the style is much the same. Sometimes it can be helpful to include intermediate columns with the truth values of subformulas:

A	B	C	$A \rightarrow B$	$B \rightarrow C$	$(A \rightarrow B) \vee (B \rightarrow C)$
T	T	T	T	T	T
T	T	F	T	F	T
T	F	T	F	T	T
T	F	F	F	T	T
F	T	T	T	T	T
F	T	F	T	F	T
F	F	T	T	T	T
F	F	F	T	T	T

By writing out the truth table for a formula, we can glance at the rows and see which truth assignments make the formula true. If all the entries in the final column are **T**, as in the above example, the formula is said to be *valid*.

6.3 Soundness and Completeness

Suppose we have a fixed deduction system in mind, such as natural deduction. A propositional formula is said to be *provable* if there is a formal proof of it in that system. A propositional formula is said to be a *tautology*, or *valid*, if it is true under any truth assignment. Provability is a syntactic notion, in that it asserts the existence of a syntactic object, namely, a proof. Validity is a semantic notion, in that it has to do with truth assignments and valuations. But, intuitively, these notions should coincide: both express the idea that a formula A *has* to be true, or is *necessarily* true, and one would expect a good proof system to enable us to derive the valid formulas.

The statement that every provable formula is valid is known as *soundness*. If A is any formula, logicians use the notation $\vdash A$ to express the fact that A is provable and the notation $\models A$ to express that A is valid. (The first symbol is sometimes called a “turnstile” and the second symbol is sometimes called a “double-turnstile.”) With this notation, soundness says that for every propositional formula A , if $\vdash A$, then $\models A$. The converse, which says that every valid formula is provable, is known as *completeness*. In symbolic terms, it says that for every formula A , if $\models A$, then $\vdash A$.

Because of the way we have chosen our inference rules and defined the notion of a valuation, this intuition that the two notions should coincide holds true. In other words, the system of natural deduction we have

presented for propositional logic is sound and complete with respect to truth-table semantics.

These notions of soundness and completeness extend to provability from hypotheses. If Γ is a set of propositional formulas and A is a propositional formula, then A is said to be a *logical consequence* of Γ if, given any truth assignment that makes every formula in Γ true, A is true as well. In this extended setting, soundness says that if A is provable from Γ , then A is a logical consequence of Γ . Completeness runs the other way: if A is a logical consequence of Γ , it is provable from Γ . In symbolic terms, we write $\Gamma \vdash A$ to express that A is provable from the formulas in Γ (or that Γ *proves* A), and we write $\Gamma \models A$ to express that A is a logical consequence of Γ (or that Γ *entails* A). With this notation, soundness says that for every propositional formula A and set of propositional formulas Γ , if $\Gamma \vdash A$ then $\Gamma \models A$, and completeness says that for every A and Γ , if $\Gamma \models A$ then $\Gamma \vdash A$.

Given a set of propositional formulas Γ and a propositional formula A , the previous section gives us a recipe for deciding whether Γ entails A : construct a truth tables for all the formulas in Γ and A , and check whether every A comes out true on every line of the table on which every formula of Γ is true. (It doesn't matter what happens to A on the lines where some formula in Γ is false.)

Notice that with the rules of natural deduction, a formula A is provable from a set of hypotheses $\{B_1, B_2, \dots, B_n\}$ if and only if the formula $B_1 \wedge B_2 \wedge \dots \wedge B_n \rightarrow A$ is provable outright, that is, from no hypotheses. So, at least for finite sets of formulas Γ , the two statements of soundness and completeness are equivalent.

Proving soundness and completeness belongs to the realm of *metatheory*, since it requires us to reason about our methods of reasoning. This is not a central focus of this book: we are more concerned with *using* logic and the notion of truth than with establishing their properties. But the notions of soundness and completeness play an important role in helping us understand the nature of the logical notions, and so we will try to provide some hints here as to why these properties hold for propositional logic.

Proving soundness is easier than proving completeness. We wish to show that whenever A is provable from a set of hypotheses, Γ , then A is a logical consequence of Γ . In a later chapter, we will consider proofs by induction, which allows us to establish a property holds of a general collection of objects by showing that it holds of some “simple” ones and is preserved under the passage to objects that are more complex. In the case of natural deduction, it is enough to show that soundness holds of the most basic proofs—using the assumption rule—and that it is preserved under each rule of inference. The base case is easy: the assumption rule says that A is provable from hypothesis A , and clearly every truth assignment that makes A true makes A true. The inductive steps are not much harder; they involve checking that the rules we have chosen mesh with the semantic notions. For example, suppose the last rule is the and-introduction rule. In that case, we have a proof of A from some hypotheses Γ , and a proof of B from some hypotheses Δ , and we combine these to form a proof of $A \wedge B$ from the hypotheses in $\Gamma \cup \Delta$, that is, the hypotheses in both. Inductively, we can assume that A is a logical consequence of Γ and that B is a logical consequence of Δ . Let v be any truth assignment that makes every formula in $\Gamma \cup \Delta$ true. Then by the inductive hypothesis, we have that it makes A true, and B true as well. By the definition of the valuation function, $\bar{v}(A \wedge B) = \mathbf{T}$, as required.

Proving completeness is harder. It suffices to show that if A is any tautology, then A is provable. One strategy is to show that natural deduction can simulate the method of truth tables. For example, suppose A is build up from propositional variables B and C . Then in natural deduction, we should be able to prove

$$(B \wedge C) \vee (B \wedge \neg C) \vee (\neg B \wedge C) \vee (\neg B \wedge \neg C),$$

with one disjunct for each line of the truth table. Then, we should be able to use each disjunct to “evaluate” each expression occurring in A , proving it true or false in accordance with its valuation, until we have a proof of A itself.

A nicer way to proceed is to express the rules of natural deduction in a way that allows us to work backward from A in search of a proof. In other words, first, we give a procedure for constructing a derivation of A by working backward from A . Then we argue that if the procedure fails, then, at the point where it fails, we can find a truth assignment that makes A false. As a result, if every truth assignment makes A true, the procedure returns a proof of A .

6.4 Exercises

1. Show that $A \rightarrow B$, $\neg A \vee B$, and $\neg(A \wedge \neg B)$ are logically equivalent, by writing out the truth table and showing that they have the same values for all truth assignments.
2. Write out the truth table for $(A \rightarrow B) \wedge (B \wedge C \rightarrow A)$.
3. Show that $A \rightarrow B$ and $\neg B \rightarrow \neg A$ are equivalent, by writing out the truth tables and showing that they have the same values for all truth assignments.
4. Does the following entailment hold?

$$\{A \rightarrow B \vee C, \neg B \rightarrow \neg C\} \models A \rightarrow B$$

Justify your answer by writing out the truth table (sorry, it is long). Indicate clearly the rows where both hypotheses come out true.

FIRST ORDER LOGIC

Propositional logic provides a good start at describing the general principles of logical reasoning, but it does not go far enough. Some of the limitations are apparent even in the “Malice and Alice” example from Chapter 2. Propositional logic does not give us the means to express a general principle that tells us that if Alice is with her son on the beach, then her son is with Alice; the general fact that no child is older than his or her parent; or the general fact that if someone is alone, they are not with someone else. To express principles like these, we need a way to talk about objects and individuals, as well as their properties and the relationships between them. These are exactly what is provided by a more expressive logical framework known as *first-order logic*, which will be the topic of the next few chapters.

7.1 Functions, Predicates, and Relations

Consider some ordinary statements about the natural numbers:

- Every natural number is even or odd, but not both.
- A natural number is even if and only if it is divisible by two.
- If some natural number, x , is even, then so is x^2 .
- A natural number x is even if and only if $x + 1$ is odd.
- Any prime number that is greater than 2 is odd.
- For any three natural numbers x , y , and z , if x divides y and y divides z , then x divides z .

These statements are true, but we generally do not think of them as *logically valid*: they depend on assumptions about the natural numbers, the meaning of the terms “even” and “odd,” and so on. But once we accept the first statement, for example, it seems to be a logical consequence that the number of stairs in the White House is either even or odd, and, in particular, if it is not even, it is odd. To make sense of inferences like these, we need a logical system that can deal with objects, their properties, and relations between them.

Rather than fix a single language once and for all, first-order logic allows us to specify the symbols we wish to use for any given domain of interest. In this section, we will use the following running example:

- The domain of interest is the natural numbers, \mathbb{N} .
- There are objects, 0, 1, 2, 3,
- There are functions, addition and multiplication, as well as the square function, on this domain.
- There are predicates on this domain, “even,” “odd,” and “prime.”
- There are relations between elements of this domain, “equal,” “less than”, and “divides.”

For our logical language, we will choose symbols $1, 2, 3, add, mul, square, even, odd, prime, lt$, and so on, to denote these things. We will also have variables x, y , and z ranging over the natural numbers. Note all of the following.

- Functions can take different numbers of arguments: if x and y are natural numbers, it makes sense to write $mul(x, y)$ and $square(x)$. So mul takes two arguments, and $square$ takes only one.
- Predicates and relations can also be understood in these terms. The predicates $even(x)$ and $prime(x)$ take one argument, while the binary relations $divides(x, y)$ and $lt(x, y)$ take two arguments.
- Functions are different from predicates! A function takes one or more arguments, and returns a *value*. A predicate takes one or more arguments, and is either true or false. We can think of predicates as returning propositions, rather than values.
- In fact, we can think of the constant symbols $1, 2, 3, \dots$ as special sorts of function symbols that take zero arguments. Analogously, we can consider the predicates that take zero arguments to be the constant logical values, \top and \perp .
- In ordinary mathematics, we often use “infix” notation for binary functions and relations. For example, we usually write $x \times y$ or $x \cdot y$ instead of $mul(x, y)$, and we write $x < y$ instead of $lt(x, y)$. We will use these conventions when writing proofs in natural deduction, and they are supported in Lean as well.
- We will treat the equality relation, $x = y$, as a special binary relation that is included in every first-order language.

First-order logic allows us to build complex expressions out of the basic ones. Starting with the variables and constants, we can use the function symbols to build up compound expressions like these:

- $x + y + z$
- $(x + 1) \times y \times y$
- $square(x + y \times z)$

Such expressions are called “terms.” Intuitively, they name objects in the intended domain of discourse.

Now, using the predicates and relation symbols, we can make assertions about these expressions:

- $even(x + y + z)$
- $prime((x + 1) \times y \times y)$
- $square(x + y \times z) = w$
- $x + y < z$

Even more interestingly, we can use propositional connectives to build compound expressions like these:

- $even(x + y + z) \wedge prime((x + 1) \times y \times y)$
- $\neg(square(x + y \times z) = w) \vee x + y < z$
- $x < y \wedge even(x) \wedge even(y) \rightarrow x + 1 < y$

The second one, for example, asserts that either $(x + yz)^2$ is not equal to w , or $x + y$ is less than z . Remember, these are expressions in symbolic logic; in ordinary mathematics, we would express the notions using words like “is even” and “if and only if,” as we did above. We will use notation like this whenever we are in the realm of symbolic logic, for example, when we write proofs in natural deduction. Expressions like these are called *formulas*. In contrast to terms, which name things, formulas *say things*; in other words, they make assertions about objects in the domain of discourse.

7.2 The Universal Quantifier

What makes first-order logic powerful is that it allows us to make general assertions using *quantifiers*. The universal quantifier \forall followed by a variable x is meant to represent the phrase “for every x .” In other words, it asserts that every value of x has the property that follows it. Using the universal quantifier, the examples with which we began the previous section can be expressed as follows:

- $\forall x ((\text{even}(x) \vee \text{odd}(x)) \wedge \neg(\text{even}(x) \wedge \text{odd}(x)))$
- $\forall x (\text{even}(x) \leftrightarrow 2 \mid x)$
- $\forall x (\text{even}(x) \rightarrow \text{even}(x^2))$
- $\forall x (\text{even}(x) \leftrightarrow \text{odd}(x + 1))$
- $\forall x (\text{prime}(x) \wedge x > 2 \rightarrow \text{odd}(x))$
- $\forall x \forall y \forall z (x \mid y \wedge y \mid z \rightarrow x \mid z)$

It is common to combine multiple quantifiers of the same kind, and write, for example, $\forall x, y, z (x \mid y \wedge y \mid z \rightarrow x \mid z)$ in the last expression.

Here are some notes on syntax:

- In symbolic logic, the universal quantifier is usually taken to bind tightly. For example, $\forall x P \vee Q$ is interpreted as $(\forall x P) \vee Q$, and we would write $\forall x (P \vee Q)$ to extend the scope.
- Be careful, however. In other contexts, especially in computer science, people often give quantifiers the *widest* scope possible. This is the case with Lean. For example, $\forall \mathbf{x}, P \vee Q$ is interpreted as $\forall \mathbf{x}, (P \vee Q)$, and we would write $(\forall \mathbf{x}, P) \vee Q$ to limit the scope.
- After the quantifier $\forall x$, the variable x is *bound*. For example, the expression $\forall x (\text{even}(x) \vee \text{odd}(x))$ expresses that every number is even or odd. Notice that the variable x does not appear anywhere in the informal statement. The statement is not about x at all; rather x is a dummy variable, a placeholder that stands for the “thing” referred to within a phrase that begins with the words “every thing.” We think of the expression $\forall x (\text{even}(x) \vee \text{odd}(x))$ as being the same as the expression $\forall y (\text{even}(y) \vee \text{odd}(y))$. Lean also treats these expressions as the same.
- In Lean, the expression $\forall \mathbf{x} \mathbf{y} \mathbf{z}, \mathbf{x} \mid \mathbf{y} \rightarrow \mathbf{y} \mid \mathbf{z} \rightarrow \mathbf{x} \mid \mathbf{z}$ is interpreted as $\forall \mathbf{x} \mathbf{y} \mathbf{z}, \mathbf{x} \mid \mathbf{y} \rightarrow (\mathbf{y} \mid \mathbf{z} \rightarrow \mathbf{x} \mid \mathbf{z})$, with parentheses associated to the *right*. The part of the expression after the universal quantifier can therefore be interpreted as saying “given that \mathbf{x} divides \mathbf{y} and that \mathbf{y} divides \mathbf{z} , \mathbf{x} divides \mathbf{z} .” The expression is logically equivalent to $\forall \mathbf{x} \mathbf{y} \mathbf{z}, \mathbf{x} \mid \mathbf{y} \wedge \mathbf{y} \mid \mathbf{z} \rightarrow \mathbf{x} \mid \mathbf{z}$, but we will see that, in Lean, it is often convenient to express facts like this as an iterated implication.

A variable that is not bound is called *free*. Notice that formulas in first-order logic say things about their free variables. For example, in the interpretation we have in mind, the formula $\forall y (x \leq y)$ says that x is less than or equal to every natural number. The formula $\forall z (x \leq z)$ says exactly the same thing; we can always rename a bound variable, as long as we pick a name that does not clash with another name that is already in use. On the other hand, the formula $\forall y (w \leq y)$ says that w is less than or equal to every natural number. This is an entirely different statement: it says something about w , rather than x . So renaming a *free* variable changes the meaning of a formula.

Notice also that some formulas, like $\forall x, y (x \leq y \vee y \leq x)$, have no free variables at all. Such a formula is called a *sentence*, because it makes an outright assertion, a statement that is either true or false about the intended interpretation. In [Chapter 10](#) we will make the notion of an “intended interpretation” precise, and explain what it means to be “true in an interpretation.” For now, the idea that formulas say things about object in an intended interpretation should motivate the rules for reasoning with such expressions.

In [Chapter 1](#) we proved that the square root of two is irrational. One way to construe the statement is as follows:

For every pair of integers, a and b , if $b \neq 0$, it is not the case that $a^2 = 2b^2$.

The advantage of this formulation is that we can restrict our attention to the integers, without having to consider the larger domain of rationals. In symbolic logic, assuming our intended domain of discourse is the integers, we would express this theorem using the universal quantifier:

$$\forall a, b \ b \neq 0 \rightarrow \neg(a^2 = 2b^2).$$

Notice that we have kept the conventional mathematical notation $b \neq 0$ to say that b is not equal to 0, but we can think of this as an abbreviation for $\neg(b = 0)$. How do we prove such a theorem? Informally, we would use such a pattern:

Let a and b be arbitrary integers, suppose $b \neq 0$, and suppose $a^2 = 2b^2$.

...

Contradiction.

What we are really doing is proving that the universal statement holds, by showing that it holds of “arbitrary” values a and b . In natural deduction, the proof would look something like this:

$$\frac{\frac{\frac{\frac{b \neq 0}{1} \quad \frac{a^2 = 2 \times b^2}{2}}{\vdots} \quad \frac{\perp}{\neg(a^2 = 2 \times b^2)}{2}}{\frac{b \neq 0 \rightarrow \neg(a^2 = 2 \times b^2)}{1}}}{\forall b (b \neq 0 \rightarrow \neg(a^2 = 2 \times b^2))} \quad 1$$

$$\frac{\forall b (b \neq 0 \rightarrow \neg(a^2 = 2 \times b^2))}{\forall a \forall b (b \neq 0 \rightarrow \neg(a^2 = 2 \times b^2))}$$

Notice that after the hypotheses are canceled, we have proved $b \neq 0 \rightarrow \neg(a^2 = 2 \times b^2)$ without making any assumptions about a and b ; at this stage in the proof, they are “arbitrary,” justifying the application of the universal quantifiers in the next two rules.

This example motivates the following rule in natural deduction:

$$\frac{A(x)}{\forall x A(x)}$$

provided x is not free in any uncanceled hypothesis. Here $A(x)$ stands for any formula that (potentially) mentions x . Also remember that if y is any “fresh” variable that does not occur in A , we are thinking of $\forall x A(x)$ as being the same as $\forall y A(y)$.

What about the elimination rule? Suppose we know that every number is even or odd. Then, in an ordinary proof, we are free to assert “ a is even or a is odd,” or “ a^2 is even or a^2 is odd.” In terms of symbolic logic, this amounts to the following inference: from $\forall x (even(x) \vee odd(x))$, we can conclude $even(t) \vee odd(t)$ for any term t . This motivates the elimination rule for the universal quantifier:

$$\frac{\forall x A(x)}{A(t)}$$

where t is an arbitrary term.

In a sense, this feels like the elimination rule for implication; we might read the hypothesis as saying “if x is any thing, then x is even or odd.” The conclusion is obtained by applying it to the fact that n is a thing. Note that, in general, we could replace n by any *term* in the language, like $n(m + 5) + 2$. Similarly, the introduction rule feels like the introduction rule for implication. If we want to show that everything has a certain property, we temporarily let x denote an arbitrary thing, and then show that it has the relevant property.

7.3 The Existential Quantifier

Dual to the universal quantifier is the existential quantifier, \exists , which is used to express assertions such as “some number is even,” or, “between any two even numbers there is an odd number.”

The following statements about the natural numbers assert the existence of some natural number:

- There exists an odd composite number. (Remember that a natural number is *composite* if it is greater than 1 and not prime.)
- Every natural number greater than one has a prime divisor.
- For every n , if n has a prime divisor smaller than n , then n is composite.

These statements can be expressed in first-order logic using the existential quantifier as follows:

- $\exists n (odd(n) \wedge composite(n))$
- $\forall n (n > 1 \rightarrow \exists p (prime(p) \wedge p \mid n))$
- $\forall n ((\exists p (p \mid n \wedge prime(p) \wedge p < n)) \rightarrow composite(n))$

After we write $\exists n$, the variable n is bound in the formula, just as for the universal quantifier. So the formulas $\exists n composite(n)$ and $\exists m composite(m)$ are considered the same.

How do we prove such existential statements? Suppose we want to prove that there exists an odd composite number. To do this, we just present a candidate, and show that the candidate satisfies the required properties. For example, we could choose 15, and then show that 15 is odd and that 15 is composite. Of course, there’s nothing special about 15, and we could have proven it also using a different number, like 9 or 35. The choice of candidate does not matter, as long as it has the required property.

In a natural deduction proof this would look like this:

$$\frac{\begin{array}{c} \vdots \\ odd(15) \wedge composite(15) \end{array}}{\exists n (odd(n) \wedge composite(n))}$$

This illustrates the introduction rule for the existential quantifier:

$$\frac{A(t)}{\exists x A(x)}$$

where t is any term. So to prove an existential formula, we just have to give one particular term for which we can prove that formula. Such term is called a *witness* for the formula.

What about the elimination rule? Suppose that we know that n is some natural number and we know that there exists a prime p such that $p < n$ and $p \mid n$. How can we use this to prove that n is composite? We can reason as follows:

Let p be any prime such that $p < n$ and $p \mid n$.

...

Therefore, n is composite.

First, we assume that there is some p which satisfies the properties p is prime, $p < n$ and $p \mid n$, and then we reason about that p . As with case-based reasoning using “or,” the assumption is only temporary: if we can show that n is composite from that assumption, that we have essentially shown that n is composite assuming the existence of such a p . Notice that in this pattern of reasoning, p should be “arbitrary.” In other words, we should not have assumed anything about p beforehand, we should not make any additional assumptions about p along the way, and the conclusion should not mention p . Only then does it makes sense to say that the conclusion follows from the “mere” existence of a p with the assumed properties.

In natural deduction, the elimination rule is expressed as follows:

$$\frac{\frac{\overline{A(y)}^1}{\vdots} \exists x A(x)}{B}^1$$

Here we require that y is not free in B , and that the only uncanceled hypotheses where y occurs freely are the hypotheses $A(y)$ that are canceled when you apply this rule. Formally, this is what it means to say that y is “arbitrary.” As was the case for or elimination and implication introduction, you can use the hypothesis $A(y)$ multiple times in the proof of B , and cancel all of them at once.

Intuitively, the rule says that you can prove B from the assumption $\exists x A(x)$ by assuming $A(y)$ for a fresh variable y , and concluding, in any number of steps, that B follows. You should compare this rule to the rule for or elimination, which is somewhat analogous. In the following example, we show that if $A(x)$ always implies $\neg B(x)$, then there cannot be an x for which both $A(x)$ and $B(x)$ holds.

7.4 Relativization and Sorts

In first-order logic as we have presented it, there is one intended “universe” of objects of discourse, and the universal and existential quantifiers range over that universe. For example, we could design a language to talk about people living in a certain town, with a relation $loves(x, y)$ to express that x loves y . In such a language, we might express the statement that “everyone loves someone” by writing $\forall x \exists y loves(x, y)$.

You should keep in mind that, at this stage, $loves$ is just a symbol. We have designed the language with a certain interpretation in mind, but one could also interpret the language as making statements about the natural numbers, where $loves(x, y)$ means that x is less than or equal to y . In that interpretation, the sentence

$$\forall x, y, z (loves(x, y) \wedge loves(y, z) \rightarrow loves(x, z))$$

is true, though in the original interpretation it makes an implausible claim about the nature of love triangles. In [Chapter 10](#), we will spell out the notion that the deductive rules of first-order logic enable us to determine the statements that are true in *all* interpretations, just as the rules of propositional logic enable us to determine the statements that are true under all truth assignments.

Returning to the original example, suppose we want to represent the statement that, in our town, all the women are strong and all the men are good looking. We could do that with the following two sentences:

- $\forall x (woman(x) \rightarrow strong(x))$
- $\forall x (man(x) \rightarrow good-looking(x))$

These are instances of *relativization*. The universal quantifier ranges over all the people in the town, but this device gives us a way of using implication to restrict the scope of our statements to men and women, respectively. The trick also comes into play when we render “every prime number greater than two is odd”:

$$\forall x (prime(x) \wedge x > 2 \rightarrow odd(x)).$$

We could also read this more literally as saying “for every number x , if x is prime and x is greater than 2, then x is odd,” but it is natural to read it as a restricted quantifier.

It is also possible to relativize the existential quantifier to say things like “some woman is strong” and “some man is good-looking.” These are expressed as follows:

- $\exists x (woman(x) \wedge strong(x))$

- $\exists x (man(x) \wedge good-looking(x))$

Notice that although we used implication to relativize the universal quantifier, here we need to use conjunction instead of implication. The expression $\exists x (woman(x) \rightarrow strong(x))$ says that there is something with the property that if it is a woman, then it is strong. Classically this is equivalent to saying that there is something which is either not a woman or is strong, which is a funny thing to say.

Now, suppose we are studying geometry, and we want to express the fact that given any two distinct points p and q and any two lines L and M , if L and M both pass through p and q , then they have to be the same. (In other words, there is at most one line between two distinct points.) One option is to design a first-order logic where the intended universe is big enough to include both points and lines, and use relativization:

$$\forall p, q, L, M (point(p) \wedge point(q) \wedge line(L) \wedge line(M) \\ \wedge q \neq p \wedge on(p, L) \wedge on(q, L) \wedge on(p, M) \wedge on(q, M) \rightarrow L = M).$$

But dealing with such predicates is tedious, and there is a mild extension of first-order logic, called *many-sorted first-order logic*, which builds in some of the bookkeeping. In many-sorted logic, one can have different sorts of objects—such as points and lines—and a separate stock of variables and quantifiers ranging over each. Moreover, the specification of function symbols and predicate symbols indicates what sorts of arguments they expect, and, in the case of function symbols, what sort of argument they return. For example, we might choose to have a sort with variables p, q, r, \dots ranging over points, a sort with variables L, M, N, \dots ranging over lines, and a relation $on(p, L)$ relating the two. Then the assertion above is rendered more simply as follows:

$$\forall p, q, L, M (p \neq q \wedge on(p, L) \wedge on(q, L) \wedge on(p, M) \wedge on(q, M) \rightarrow L = M).$$

7.5 Equality

In symbolic logic, we use the expression $s = t$ to express the fact that s and t are “equal” or “identical.” The equality symbol is meant to model what we mean when we say, for example, “Alice’s brother is the victim,” or “ $2 + 2 = 4$.” We are asserting that two different descriptions refer to the same object. Because the notion of identity can be applied to virtually any domain of objects, it is viewed as falling under the province of logic.

Talk of “equality” or “identity” raises messy philosophical questions, however. Am I the same person I was three days ago? Are the two copies of *Huckleberry Finn* sitting on my shelf the same book, or two different books? Using symbolic logic to model identity presupposes that we have in mind a certain way of carving up and interpreting the world. We assume that our terms refer to distinct entities, and writing $s = t$ asserts that the two expressions refer to the same thing. Axiomatically, we assume that equality satisfies the following three properties:

- *reflexivity*: $t = t$, for any term t
- *symmetry*: if $s = t$, then $t = s$
- *transitivity*: if $r = s$ and $s = t$, then $r = t$

These properties are not enough to characterize equality, however. If two expressions denote the same thing, then we should be able to substitute one for any other in any expression. It is convenient to adopt the following convention: if r is any term, we may write $r(x)$ to indicate that the variable x may occur in r . Then, if s is another term, we can thereafter write $r(s)$ to denote the result of replacing s for x in r . The substitution rule for terms thus reads as follows: if $s = t$, then $r(s) = r(t)$.

We already adopted a similar convention for formulas: if we introduce a formula as $A(x)$, then $A(t)$ denotes the result of substituting t for x in A . With this in mind, we can write the rules for equality as follows:

$$\frac{}{t = t} \quad \frac{s = t}{t = s} \quad \frac{r = s \quad s = t}{r = t}$$

$$\frac{s = t}{r(s) = r(t)} \quad \frac{s = t \quad P(s)}{P(t)}$$

Here, the first substitution rule governs terms and the second substitution rule governs formulas. In the next chapter, you will learn how to use them.

Using equality, we can define even more quantifiers.

- We can express “there are at least two elements x such that $A(x)$ holds” as $\exists x \exists y (x \neq y \wedge A(x) \wedge A(y))$.
- We can express “there are at most two elements x such that $A(x)$ holds” as $\forall x \forall y \forall z (A(x) \wedge A(y) \wedge A(z) \rightarrow x = y \vee y = z \vee x = z)$. This states that if we have three elements a for which $A(a)$ holds, then two of them must be equal.
- We can express “there are exactly two elements x such that $A(x)$ holds” as the conjunction of the above two statements.

As an exercise, write out in first order logic the statements that there are at least, at most, and exactly three elements x such that $A(x)$ holds.

In logic, the expression $\exists!x A(x)$ is used to express the fact that there is a *unique* x satisfying $A(x)$, which is to say, there is exactly one such x . As above, this can be expressed as follows:

$$\exists x A(x) \wedge \forall y \forall y' (A(y) \wedge A(y') \rightarrow y = y').$$

The first conjunct says that there is at least one object satisfying A , and the second conjunct says that there is at most one. The same thing can be expressed more concisely as follows:

$$\exists x (A(x) \wedge \forall y (A(y) \rightarrow y = x)).$$

You should think about why this second expression works. In the next chapter we will see that, using the rules of natural deduction, we can prove that these two expressions are equivalent.

7.6 Exercises

1. A *perfect number* is a number that is equal to the sum of its proper divisors, that is, the numbers that divide it, other than itself. For example, 6 is perfect, because $6 = 1 + 2 + 3$.

Using a language with variables ranging over the natural numbers and suitable functions and predicates, write down first-order sentences asserting the following. Use a predicate *perfect* to express that a number is perfect.

- a. 28 is perfect.
- b. There are no perfect numbers between 100 and 200.
- c. There are (at least) two perfect numbers between 200 and 10,000. (Express this by saying that there are perfect numbers x and y between 200 and 10,000, with the property that $x \neq y$.)
- d. Every perfect number is even.
- e. For every number, there is a perfect number that is larger than it. (This is one way to express the statement that there are infinitely many perfect numbers.)

Here, the phrase “between a and b ” is meant to include a and b .

By the way, we do not know whether the last two statements are true. They are open questions.

2. Using a language with variables ranging over people, and predicates $trusts(x, y)$, $politician(x)$, $crazy(x)$, $knows(x, y)$, and $related-to(x, y)$, and $rich(x)$, write down first-order sentences asserting the following:
- Nobody trusts a politician.
 - Anyone who trusts a politician is crazy.
 - Everyone knows someone who is related to a politician.
 - Everyone who is rich is either a politician or knows a politician.

In each case, some interpretation may be involved. Notice that writing down a logical expression is one way of helping to clarify the meaning.

NATURAL DEDUCTION FOR FIRST ORDER LOGIC

8.1 Rules of Inference

In the last chapter, we discussed the language of first-order logic, and the rules that govern their use. We summarize them here:

The universal quantifier:

$$\frac{A(x)}{\forall y A(y)} \forall I \qquad \frac{\forall x A(x)}{A(t)} \forall E$$

In the introduction rule, x should not be free in any uncanceled hypothesis. In the elimination rule, t can be any term that does not clash with any of the bound variables in A .

The existential quantifier:

$$\frac{A(t)}{\exists x A(x)} \exists I \qquad \frac{\begin{array}{c} \overline{A(y)}^1 \\ \vdots \\ B \end{array}}{\exists x A(x)} \exists E$$

In the introduction rule, t can be any term that does not clash with any of the bound variables in A . In the elimination rule, y should not be free in B or any uncanceled hypothesis.

Equality:

$$\frac{}{t = t} \text{ refl} \qquad \frac{s = t}{t = s} \text{ symm} \qquad \frac{r = s \quad s = t}{r = t} \text{ trans}$$

$$\frac{s = t}{r(s) = r(t)} \text{ subst} \qquad \frac{s = t \quad P(s)}{P(t)} \text{ subst}$$

Strictly speaking, only refl and the second substitution rule are necessary. The others can be derived from them.

8.2 The Universal Quantifier

The following example of a proof in natural deduction shows that if, for every x , $A(x)$ holds, and for every x , $B(x)$ holds, then for every x , they both hold:

$$\frac{\frac{\frac{\overline{\forall x A(x)}^1}{A(y)}}{\frac{\overline{\forall x B(x)}^2}{B(y)}}}{A(y) \wedge B(y)} \quad \frac{\overline{\forall y (A(y) \wedge B(y))}}{\forall x B(x) \rightarrow \forall y (A(y) \wedge B(y))}^2}{\forall x A(x) \rightarrow (\forall x B(x) \rightarrow \forall y (A(y) \wedge B(y)))}^1$$

Notice that neither of the assumptions 1 or 2 mention y , so that y is really “arbitrary” at the point where the universal quantifiers are introduced.

Here is another example:

$$\frac{\frac{\frac{\overline{\forall x A(x)}^1}{A(y)}}{A(y) \vee B(y)}}{\forall x (A(x) \vee B(x))} \quad \frac{\overline{\forall x A(x) \rightarrow \forall x (A(x) \vee B(x))}^1}{\forall x A(x) \rightarrow \forall x (A(x) \vee B(x))}^1$$

As an exercise, try proving the following:

$$\forall x (A(x) \rightarrow B(x)) \rightarrow (\forall x A(x) \rightarrow \forall x B(x)).$$

Here is a more challenging exercise. Suppose I tell you that, in a town, there is a (male) barber that shaves all and only the men who do not shave themselves. You can show that this is a contradiction, arguing informally, as follows:

By the assumption, the barber shaves himself if and only if he does not shave himself. Call this statement (*).

Suppose the barber shaves himself. By (*), this implies that he does not shave himself, a contradiction. So, the barber does not shave himself.

But using (*) again, this implies that the barber shaves himself, which contradicts the fact we just showed, namely, that the barber does not shave himself.

Try to turn this into a formal argument in natural deduction.

Let us return to the example of the natural numbers, to see how deductive notions play out there. Suppose we have defined *even* and *odd* in such a way that we can prove:

- $\forall n (\neg \text{even}(n) \rightarrow \text{odd}(n))$
- $\forall n (\text{odd}(n) \rightarrow \neg \text{even}(n))$

Then we can go on to derive $\forall n (\text{even}(n) \vee \text{odd}(n))$ as follows:

$$\frac{\frac{\overline{\text{even}(n) \vee \neg \text{even}(n)}}{\text{even}(n)}^1 \quad \frac{\overline{\text{even}(n) \vee \text{odd}(n)}}{\text{even}(n) \vee \text{odd}(n)}}{\frac{\overline{\text{even}(n) \vee \text{odd}(n)}}{\forall n (\text{even}(n) \vee \text{odd}(n))}^1} \quad \frac{\frac{\overline{\forall n (\neg \text{even}(n) \rightarrow \text{odd}(n))}}{\neg \text{even}(n) \rightarrow \text{odd}(n)} \quad \frac{\overline{\neg \text{even}(n)}}{\neg \text{even}(n)}^1}{\text{odd}(n)} \quad \frac{\overline{\text{even}(n) \vee \text{odd}(n)}}{\text{even}(n) \vee \text{odd}(n)}}^1$$

We can also prove and $\forall n \neg(\text{even}(n) \wedge \text{odd}(n))$:

$$\frac{\frac{\frac{}{odd(n) \rightarrow \neg even(n)}}{\neg even(n)}}{\frac{\frac{\frac{}{even(n) \wedge odd(n)}{odd(n)}}{H}}{even(n)}}{H}}{\frac{\perp}{\neg(even(n) \wedge odd(n))}}{H}}{\forall n \neg(even(n) \wedge odd(n))}$$

As we move from modeling basic rules of inference to modeling actual mathematical proofs, we will tend to shift focus from natural deduction to formal proofs in Lean. Natural deduction has its uses: as a model of logical reasoning, it provides us with a convenient means to study metatheoretic properties such as soundness and completeness. For working *within* the system, however, proof languages like Lean's tend to scale better, and produce more readable proofs.

8.3 The Existential Quantifier

Remember that the intuition behind the elimination rule for the existential quantifier is that if we know $\exists x A(x)$, we can temporarily reason about an arbitrary element y satisfying $A(y)$ in order to prove a conclusion that doesn't depend on y . Here is an example of how it can be used. The next proof says that if we know there is something satisfying both A and B , then we know, in particular, that there is something satisfying A .

$$\frac{\frac{\frac{}{\exists x(A(x) \wedge B(x))}{}^1}{\exists x A(x)}}{\frac{\frac{\frac{}{A(y) \wedge B(y)}{}^2}{A(y)}}{\exists x A(x)}}{}^2}{\exists x(A(x) \wedge B(x)) \rightarrow \exists x A(x)}{}^1$$

The following proof shows that if there is something satisfying either A or B , then either there is something satisfying A , or there is something satisfying B .

$$\frac{\frac{\frac{}{\exists x(A(x) \vee B(x))}{}^1}{\exists x A(x) \vee \exists x B(x)}}{\frac{\frac{\frac{\frac{}{A(y) \vee B(y)}{}^2}{\exists x A(x) \vee \exists x B(x)}}{\exists x A(x)}{}^3}{\exists x A(x) \vee \exists x B(x)}}{}^2}{\exists x(A(x) \vee B(x)) \rightarrow \exists x A(x) \vee \exists x B(x)}{}^1$$

The following example is more involved:

$$\frac{\frac{\frac{\frac{\frac{}{\forall x(A(x) \rightarrow \neg B(x))}{}^1}{A(x) \rightarrow \neg B(x)}}{\neg B(x)}}{\exists x(A(x) \wedge B(x))}{}^2}{\frac{\frac{\frac{\frac{}{A(x) \wedge B(x)}{}^3}{A(x)}}{B(x)}}{\perp}}{}^3}{\frac{\frac{\frac{}{\neg \exists x(A(x) \wedge B(x))}{}^2}{\forall x(A(x) \rightarrow \neg B(x)) \rightarrow \neg \exists x(A(x) \wedge B(x))}{}^1}}{}^1$$

In this proof, the existential elimination rule (the line labeled 3) is used to cancel two hypotheses at the same time. Note that when this rule is applied, the hypothesis $\forall x(A(x) \rightarrow \neg B(x))$ has not yet been canceled.

So we have to make sure that this formula doesn't contain the variable x freely. But this is o.k., since this hypothesis contains x only as a bound variable.

Another example is that if x does not occur in P , then $\exists x P$ is equivalent to P :

$$\frac{\frac{\overline{\exists x P}^1}{P} \quad \frac{\overline{P}^2}{\exists x P}^1}{\exists x P \leftrightarrow P}^1$$

This is short but tricky, so let us go through it carefully. On the left, we assume $\exists x P$ to conclude P . We assume P , and now we can immediately cancel this assumption by existential elimination, since x does not occur in P , so it doesn't occur freely in any assumption or in the conclusion. On the right we use existential introduction to conclude $\exists x P$ from P .

8.4 Equality

Recall the natural deduction rules for equality:

$$\frac{}{t = t} \quad \frac{s = t}{t = s} \quad \frac{r = s \quad s = t}{r = t}$$

$$\frac{s = t}{r(s) = r(t)} \quad \frac{s = t \quad P(s)}{P(t)}$$

Keep in mind that we have implicitly fixed some first-order language, and r , s , and t are any terms in that language. Recall also that we have adopted the practice of using functional notation with terms. For example, if we think of $r(x)$ as the term $(x + y) \times (z + 0)$ in the language of arithmetic, then $r(0)$ is the term $(0 + y) \times (z + 0)$ and $r(u + v)$ is $((u + v) + y) \times (z + 0)$. So one example of the first inference on the second line is this:

$$\frac{u + v = 0}{((u + v) + y) \times (z + 0) = (0 + y) \times (z + 0)}$$

The second axiom on that line is similar, except now $P(x)$ stands for any *formula*, as in the following inference:

$$\frac{u + v = 0 \quad x + (u + v) < y}{x + 0 < y}$$

Notice that we have written the reflexivity axiom, $t = t$, as a rule with no premises. If you use it in a proof, it does not count as a hypothesis; it is built into the logic.

In fact, we can think of the first inference on the second line as a special case of the first. Consider, for example, the formula $((u + v) + y) \times (z + 0) = (x + y) \times (z + 0)$. If we plug $u + v$ in for x , we get an instance of reflexivity. If we plug in 0, we get the conclusion of the first example above. The following is therefore a derivation of the first inference, using only reflexivity and the second substitution rule above:

$$\frac{u + v = 0 \quad \overline{((u + v) + y) \times (z + 0) = ((u + v) + y) \times (z + 0)}}{((u + v) + y) \times (z + 0) = (0 + y) \times (z + 0)}$$

Roughly speaking, we are replacing the second instance of $u + v$ in an instance of reflexivity with 0 to get the conclusion we want.

Equality rules let us carry out calculations in symbolic logic. This typically amounts to using the equality rules we have already discussed, together with a list of general identities. For example, the following identities hold for any real numbers x , y , and z :

- commutativity of addition: $x + y = y + x$
- associativity of addition: $(x + y) + z = x + (y + z)$

- additive identity: $x + 0 = 0 + x = x$
- additive inverse: $-x + x = x + -x = 0$
- multiplicative identity: $x \cdot 1 = 1 \cdot x = x$
- commutativity of multiplication: $x \cdot y = y \cdot x$
- associativity of multiplication: $(x \cdot y) \cdot z = x \cdot (y \cdot z)$
- distributivity: $x \cdot (y + z) = x \cdot y + x \cdot z, \quad (x + y) \cdot z = x \cdot z + y \cdot z$

You should imagine that there are implicit universal quantifiers in front of each statement, asserting that the statement holds for *any* values of $x, y,$ and $z.$ Note that $x, y,$ and z can, in particular, be integers or rational numbers as well. Calculations involving real numbers, rational numbers, or integers generally involve identities like this.

The strategy is to use the elimination rule for the universal quantifier to instantiate general identities, use symmetry, if necessary, to orient an equation in the right direction, and then using the substitution rule for equality to change something in a previous result. For example, here is a natural deduction proof of a simple identity, $\forall x, y, z ((x + y) + z = (x + z) + y),$ using only commutativity and associativity of addition. We have taken the liberty of using a brief name to denote the relevant identities, and combining multiple instances of the universal quantifier introduction and elimination rules into a single step.

$$\frac{\frac{\frac{\text{assoc}}{(x+z)+y = x+(z+y)}}{x+(z+y) = (x+z)+y} \quad \frac{\frac{\text{comm}}{y+z = z+y} \quad \frac{\text{assoc}}{(x+y)+z = x+(y+z)}}{(x+y)+z = x+(z+y)}}{(x+y)+z = (x+z)+y}}{\forall x, y, z ((x+y)+z = (x+z)+y)}$$

There is generally nothing interesting to be learned from carrying out such calculations in natural deduction, but you should try one or two examples to get the hang of it, and then take pleasure in knowing that it is possible.

8.5 Counterexamples and Relativized Quantifiers

Consider the statement:

Every prime number is odd.

In first-order logic, we could formulate this as $\forall p (prime(p) \rightarrow odd(p)).$ This statement is false, because there is a prime number which is even, namely the number 2. This is called a *counterexample* to the statement.

More generally, given a formula $\forall x A(x),$ a counterexample is a value t such that $\neg A(t)$ holds. Such a counterexample shows that the original formula is false, because we have the following equivalence: $\neg \forall x A(x) \leftrightarrow \exists x \neg A(x).$ So if we find a value t such that $\neg A(t)$ holds, then by the existential introduction rule we can conclude that $\exists x \neg A(x),$ and then by the above equivalence we have $\neg \forall x A(x).$ Here is a proof of the equivalence:

$$\frac{\frac{\frac{\frac{\frac{\neg \forall x A(x)}{1}}{\exists x \neg A(x)}{4}}{\perp}{5}}{\forall x A(x)}{5}}{\perp}{4} \quad \frac{\frac{\frac{\frac{\frac{\frac{\exists x \neg A(x)}{1}}{\neg A(y)}{3}}{\perp}{3}}{A(y)}{2}}{\forall x A(x)}{2}}{\exists x \neg A(x)}{1}}{\perp}{3}}{\neg \forall x A(x) \leftrightarrow \exists x \neg A(x)}{1}$$

One remark about the proof: at the step marked by 4 we *cannot* use the existential introduction rule, because at that point our only assumption is $\neg\forall x A(x)$, and from that assumption we cannot prove $\neg A(t)$ for a particular term t . So we use a proof by contradiction there.

As an exercise, prove the “dual” equivalence yourself: $\neg\exists x A(x) \leftrightarrow \forall x \neg A(x)$. This can be done without using proof by contradiction.

In [Chapter 7](#) we saw examples of how to use relativization to restrict the scope of a universal quantifier. Suppose we want to say “every prime number is greater than 1”. In first order logic this can be written as $\forall n(\text{prime}(n) \rightarrow n > 1)$. The reason is that the original statement is equivalent to the statement “for every natural number, if it is prime, then it is greater than 1”. Similarly, suppose we want to say “there exists a prime number greater than 100.” This is equivalent to saying “there exists a natural number which is prime and greater than 100,” which can be expressed as $\exists n (\text{prime}(n) \wedge n > 100)$.

As an exercise you can prove the above results about negations of quantifiers also for relativized quantifiers. Specifically, prove the following statements:

- $\neg\exists x (A(x) \wedge B(x)) \leftrightarrow \forall x (A(x) \rightarrow \neg B(x))$
- $\neg\forall x (A(x) \rightarrow B(x)) \leftrightarrow \exists x(A(x) \wedge \neg B(x))$

For reference, here is a list of valid sentences involving quantifiers:

- $\forall x A \leftrightarrow A$ if x is not free in A
- $\exists x A \leftrightarrow A$ if x is not free in A
- $\forall x (A(x) \wedge B(x)) \leftrightarrow \forall x A(x) \wedge \forall x B(x)$
- $\exists x (A(x) \wedge B) \leftrightarrow \exists x A(x) \wedge B$ if x is not free in B
- $\exists x (A(x) \vee B(x)) \leftrightarrow \exists x A(x) \vee \exists x B(x)$
- $\forall x (A(x) \vee B) \leftrightarrow \forall x A(x) \vee B$ if x is not free in B
- $\forall x (A(x) \rightarrow B) \leftrightarrow (\exists x A(x) \rightarrow B)$ if x is not free in B
- $\exists x (A(x) \rightarrow B) \leftrightarrow (\forall x A(x) \rightarrow B)$ if x is not free in B
- $\forall x (A \rightarrow B(x)) \leftrightarrow (A \rightarrow \forall x B(x))$ if x is not free in A
- $\exists x (A(x) \rightarrow B) \leftrightarrow (A(x) \rightarrow \exists x B)$ if x is not free in B
- $\exists x A(x) \leftrightarrow \neg\forall x \neg A(x)$
- $\forall x A(x) \leftrightarrow \neg\exists x \neg A(x)$
- $\neg\exists x A(x) \leftrightarrow \forall x \neg A(x)$
- $\neg\forall x A(x) \leftrightarrow \exists x \neg A(x)$

All of these can be derived in natural deduction. The last two allow us to push negations inwards, so we can continue to put first-order formulas in negation normal form. Other rules allow us to bring quantifiers to the front of any formula, though, in general, there will be multiple ways of doing this. For example, the formula

$$\forall x A(x) \rightarrow \exists y \forall z B(y, z)$$

is equivalent to both

$$\exists x, y \forall z (A(x) \rightarrow B(y, z))$$

and

$$\exists y \forall z \exists x (A(x) \rightarrow B(y, z)).$$

A formula with all the quantifiers in front is said to be in *prenex* form.

8.6 Exercises

1. Give a natural deduction proof of

$$\forall x (A(x) \rightarrow B(x)) \rightarrow (\forall x A(x) \rightarrow \forall x B(x)).$$

2. Give a natural deduction proof of $\forall x B(x)$ from hypotheses $\forall x (A(x) \vee B(x))$ and $\forall y \neg A(y)$.
3. From hypotheses $\forall x (\text{even}(x) \vee \text{odd}(x))$ and $\forall x (\text{odd}(x) \rightarrow \text{even}(s(x)))$ give a natural deduction proof $\forall x (\text{even}(x) \vee \text{even}(s(x)))$. (It might help to think of $s(x)$ as the function defined by $s(x) = x + 1$.)
4. Give a natural deduction proof of $\exists x A(x) \vee \exists x B(x) \rightarrow \exists x (A(x) \vee B(x))$.
5. Give a natural deduction proof of $\exists x (A(x) \wedge C(x))$ from the assumptions $\exists x (A(x) \wedge B(x))$ and $\forall x (A(x) \wedge B(x) \rightarrow C(x))$.
6. Prove some of the other equivalences in the last section.
7. Consider some of the various ways of expressing “nobody trusts a politician” in first-order logic:

- $\forall x (\text{politician}(x) \rightarrow \forall y (\neg \text{trusts}(y, x)))$
- $\forall x, y (\text{politician}(x) \rightarrow \neg \text{trusts}(y, x))$
- $\neg \exists x, y (\text{politician}(x) \wedge \text{trusts}(y, x))$
- $\forall x, y (\text{trusts}(y, x) \rightarrow \neg \text{politician}(x))$

They are all logically equivalent. Show this for the second and the fourth, by giving natural deduction proofs of each from the other. (As a shortcut, in the \forall introduction and elimination rules, you can introduce / eliminate both variables in one step.)

8. Formalize the following statements, and give a natural deduction proof in which the first three statements appear as (uncancelled) hypotheses, and the last line is the conclusion:
- Every young and healthy person likes baseball.
 - Every active person is healthy.
 - Someone is young and active.
 - Therefore, someone likes baseball.

Use $Y(x)$ for “is young,” $H(x)$ for “is healthy,” $A(x)$ for “is active,” and $B(x)$ for “likes baseball.”

9. Give a natural deduction proof of $\forall x, y, z (x = z \rightarrow (y = z \rightarrow x = y))$ using the equality rules in Section 8.4.
10. Give a natural deduction proof of $\forall x, y (x = y \rightarrow y = x)$ using only these two hypotheses (and none of the new equality rules):
- $\forall x (x = x)$
 - $\forall u, v, w (u = w \rightarrow (v = w \rightarrow u = v))$

(Hint: Choose instantiations of u , v , and w carefully. You can instantiate all the universal quantifiers in one step, as on the last homework assignment.)

11. Give a natural deduction proof of $\neg \exists x (A(x) \wedge B(x)) \leftrightarrow \forall x (A(x) \rightarrow \neg B(x))$
12. Give a natural deduction proof of $\neg \forall x (A(x) \rightarrow B(x)) \leftrightarrow \exists x (A(x) \wedge \neg B(x))$
13. Remember that both the following express $\exists! x A(x)$, that is, the statement that there is a unique x satisfying $A(x)$:
- $\exists x (A(x) \wedge \forall y (A(y) \rightarrow y = x))$

- $\exists x A(x) \wedge \forall y \forall y' (A(y) \wedge A(y') \rightarrow y = y')$

Do the following:

- Give a natural deduction proof of the second, assuming the first as a hypothesis.
- Give a natural deduction proof of the first, assuming the second as a hypothesis.

(Warning: these are long.)

FIRST ORDER LOGIC IN LEAN

9.1 Functions, Predicates, and Relations

In the last chapter, we discussed the language of first-order logic. We will see in the course of this book that Lean’s built-in logic is much more expressive; but it *includes* first-order logic, which is to say, anything that can be expressed (and proved) in first-order logic can be expressed (and proved) in Lean.

Lean is based on a foundational framework called *type theory*, in which every variable is assumed to range elements of some *type*. You can think of a type as being a “universe,” or a “domain of discourse,” in the sense of first-order logic.

For example, suppose we want to work with a first-order language with one constant symbol, one unary function symbol, one binary function symbol, one unary relation symbol, and one binary relation symbol. We can declare a new type U (for “universe”) and the relevant symbols as follows:

```
constant U : Type

constant c : U
constant f : U → U
constant g : U → U → U
constant P : U → Prop
constant R : U → U → Prop
```

We can then use them as follows:

```
variables x y : U

#check c
#check f c
#check g x y
#check g x (f c)

#check P (g x (f c))
#check R x y
```

The `#check` command tells us that the first four expressions have type U , and that the last two have type Prop . Roughly, this means that the first four expressions correspond to terms of first-order logic, and that the last two correspond to formulas.

Note all the following:

- A unary function is represented as an object of type $U \rightarrow U$ and a binary function is represented as an object of type $U \rightarrow U \rightarrow U$, using the same notation as for implication between propositions.

- We write, for example, $f\ x$ to denote the result of applying f to x , and $g\ x\ y$ to denote the result of applying g to x and y , again just as we did when using modus ponens for first-order logic. Parentheses are needed in the expression $g\ x\ (f\ c)$ to ensure that $f\ c$ is parsed as a single argument.
- A unary predicate is presented as an object of type $U \rightarrow \text{Prop}$ and a binary predicate is represented as an object of type $U \rightarrow U \rightarrow \text{Prop}$. You can think of a binary relation R as being a function that assumes two arguments in the universe, U , and returns a proposition.
- We write $P\ x$ to denote the assertion that P holds of x , and $R\ x\ y$ to denote that R holds of x and y .

You may reasonably wonder what difference there is between a constant and a variable in Lean. The following declarations also work:

```
variable U : Type

variable c : U
variable f : U → U
variable g : U → U → U
variable P : U → Prop
variable R : U → U → Prop

variables x y : U

#check c
#check f c
#check g x y
#check g x (f c)

#check P (g x (f c))
#check R x y
```

Although the examples function in much the same way, the `constant` and `variable` commands do very different things. The `constant` command declares a new object, axiomatically, and adds it to the list of objects Lean knows about. In contrast, when it is first executed, the `variable` command does not create anything. Rather, it tells Lean that whenever we enter an expression using the corresponding identifier, it should create a temporary variable of the corresponding type.

Many types are already declared in Lean's standard library. For example, there is a type written `nat` or \mathbb{N} , that denotes the natural numbers:

```
#check nat
#check  $\mathbb{N}$ 
```

You can enter the unicode \mathbb{N} with `\nat` or `\N`. The two expressions mean the same thing.

Using this built-in type, we can model the language of arithmetic, as described in the last chapter, as follows:

```
namespace hidden

constant mul :  $\mathbb{N} \rightarrow \mathbb{N} \rightarrow \mathbb{N}$ 
constant add :  $\mathbb{N} \rightarrow \mathbb{N} \rightarrow \mathbb{N}$ 
constant square :  $\mathbb{N} \rightarrow \mathbb{N}$ 
constant even :  $\mathbb{N} \rightarrow \text{Prop}$ 
constant odd :  $\mathbb{N} \rightarrow \text{Prop}$ 
constant prime :  $\mathbb{N} \rightarrow \text{Prop}$ 
constant divides :  $\mathbb{N} \rightarrow \mathbb{N} \rightarrow \text{Prop}$ 
```

(continues on next page)

(continued from previous page)

```

constant lt : ℕ → ℕ → Prop
constant zero : ℕ
constant one : ℕ

end hidden

```

We have used the `namespace` command to avoid conflicts with identifiers that are already declared in the Lean library. (Outside the namespace, the constant `mul` we just declared is named `hidden.mul`.) We can again use the `#check` command to try them out:

```

namespace hidden

constant mul : ℕ → ℕ → ℕ
constant add : ℕ → ℕ → ℕ
constant square : ℕ → ℕ
constant even : ℕ → Prop
constant odd : ℕ → Prop
constant prime : ℕ → Prop
constant divides : ℕ → ℕ → Prop
constant lt : ℕ → ℕ → Prop
constant zero : ℕ
constant one : ℕ

variables w x y z : ℕ

#check mul x y
#check add x y
#check square x
#check even x

end hidden

```

We can even declare infix notation of binary operations and relations:

```

infix + := add
infix * := mul
infix < := lt

```

(Getting notation for numerals 1, 2, 3, ... is trickier.) With all this in place, the examples above can be rendered as follows:

```

#check even (x + y + z) ∧ prime ((x + one) * y * y)
#check ¬ (square (x + y * z) = w) ∨ x + y < z
#check x < y ∧ even x ∧ even y → x + one < y

```

In fact, all of the functions, predicates, and relations discussed here, except for the “square” function and “prime,” are defined in the core Lean library. They become available to us when we put the commands `import data.nat` and `open nat` at the top of a file in Lean.

```

import data.nat
open nat

constant square : ℕ → ℕ

```

(continues on next page)

(continued from previous page)

```
constant prime : ℕ → Prop
constant even  : ℕ → Prop

variables w x y z : ℕ

#check even (x + y + z) ∧ prime ((x + 1) * y * y)
#check ¬ (square (x + y * z) = w) ∨ x + y < z
#check x < y ∧ even x ∧ even y → x + 1 < y
```

Here, we declare the constants `square` and `prime` axiomatically, but refer to the other operations and predicates in the Lean library. In this book, we will often proceed in this way, telling you explicitly what facts from the library you should use for exercises.

Again, note the following aspects of syntax:

- In contrast to ordinary mathematical notation, in Lean, functions are applied without parentheses or commas. For example, we write `square x` and `add x y` instead of `square(x)` and `add(x, y)`.
- The same holds for predicates and relations: we write `even x` and `lt x y` instead of `even(x)` and `lt(x, y)`, as one might do in symbolic logic.
- The notation `add : ℕ → ℕ → ℕ` indicates that addition assumes two arguments, both natural numbers, and returns a natural number.
- Similarly, the notation `divides : ℕ → ℕ → Prop` indicates that `divides` is a binary relation, which assumes two natural numbers as arguments and forms a proposition. In other words, `divides x y` expresses the assertion that `x` divides `y`.

Lean can help us distinguish between terms and formulas. If we `#check` the expression `x + y + 1` in Lean, we are told it has type `ℕ`, which is to say, it denotes a natural number. If we `#check` the expression `even (x + y + 1)`, we are told that it has type `Prop`, which is to say, it expresses a proposition.

In [Chapter 7](#) we considered many-sorted logic, where one can have multiple universes. For example, we might want to use first-order logic for geometry, with quantifiers ranging over points and lines. In Lean, we can model this as by introducing a new type for each sort:

```
variables Point Line : Type
variable lies_on : Point → Line → Prop
```

We can then express that two distinct points determine a line as follows:

```
#check ∀ (p q : Point) (L M : Line),
  p ≠ q → lies_on p L → lies_on q L → lies_on p M →
  lies_on q M → L = M
```

Notice that we have followed the convention of using iterated implication rather than conjunction in the antecedent. In fact, Lean is smart enough to infer what sorts of objects `p`, `q`, `L`, and `M` are from the fact that they are used with the relation `lies_on`, so we could have written, more simply, this:

```
#check ∀ p q L M, p ≠ q → lies_on p L → lies_on q L →
  lies_on p M → lies_on q M → L = M
```

9.2 Using the Universal Quantifier

In Lean, you can enter the universal quantifier by writing `\all`. The motivating examples from Section 7.1 are rendered as follows:

```
import data.nat
open nat

constant prime : ℕ → Prop
constant even  : ℕ → Prop
constant odd   : ℕ → Prop

#check ∀ x, (even x ∨ odd x) ∧ ¬ (even x ∧ odd x)
#check ∀ x, even x ↔ 2 | x
#check ∀ x, even x → even (x^2)
#check ∀ x, even x ↔ odd (x + 1)
#check ∀ x, prime x ∧ x > 2 → odd x
#check ∀ x y z, x | y → y | z → x | z
```

Remember that Lean expects a comma after the universal quantifier, and gives it the *widest* scope possible. For example, $\forall x, P \vee Q$ is interpreted as $\forall x, (P \vee Q)$, and we would write $(\forall x, P) \vee Q$ to limit the scope. If you prefer, you can use the plain ascii expression `forall` instead of the unicode \forall .

In Lean, then, the pattern for proving a universal statement is rendered as follows:

```
variable U : Type
variable P : U → Prop

example : ∀ x, P x :=
assume x,
show P x, from sorry
```

Read `assume x` as “fix an arbitrary value x of U .” Since we are allowed to rename bound variables at will, we can equivalently write either of the following:

```
variable U : Type
variable P : U → Prop

example : ∀ y, P y :=
assume x,
show P x, from sorry

example : ∀ x, P x :=
assume y,
show P y, from sorry
```

This constitutes the introduction rule for the universal quantifier. It is very similar to the introduction rule for implication: instead of using `assume` to temporarily introduce an assumption, we use `assume` to temporarily introduce a new object, y . (In fact, `assume` and `assume` are both alternate syntax for a single internal construct in Lean, which can also be denoted by λ .)

The elimination rule is, similarly, implemented as follows:

```

variable U : Type
variable P : U → Prop
variable h : ∀ x, P x
variable a : U

example : P a :=
show P a, from h a
    
```

Observe the notation: `P a` is obtained by “applying” the hypothesis `h` to `a`. Once again, note the similarity to the elimination rule for implication.

Here is an example of how it is used:

```

variable U : Type
variables A B : U → Prop

example (h1 : ∀ x, A x → B x) (h2 : ∀ x, A x) : ∀ x, B x :=
assume y,
have h3 : A y, from h2 y,
have h4 : A y → B y, from h1 y,
show B y, from h4 h3
    
```

Here is an even shorter version of the same proof, where we avoid using `have`:

```

example (h1 : ∀ x, A x → B x) (h2 : ∀ x, A x) : ∀ x, B x :=
assume y,
show B y, from h1 y (h2 y)
    
```

You should talk through the steps, here. Applying `h1` to `y` yields a proof of `A y → B y`, which we then apply to `h2 y`, which is a proof of `A y`. The result is the proof of `B y` that we are after.

In the last chapter, we considered the following proof in natural deduction:

$$\frac{\frac{\frac{\overline{\forall x A(x)}^1}{A(y)}}{A(y) \wedge B(y)} \quad \frac{\overline{\forall x B(x)}^2}{B(y)}}{\forall y (A(y) \wedge B(y))} \quad \frac{\forall x B(x) \rightarrow \forall y (A(y) \wedge B(y))}{\forall x A(x) \rightarrow (\forall x B(x) \rightarrow \forall y (A(y) \wedge B(y)))}^1$$

Here is the same proof rendered in Lean:

```

variable U : Type
variables A B : U → Prop

example : (∀ x, A x) → (∀ x, B x) → (∀ x, A x ∧ B x) :=
assume hA : ∀ x, A x,
assume hB : ∀ x, B x,
assume y,
have Ay : A y, from hA y,
have By : B y, from hB y,
show A y ∧ B y, from and.intro Ay By
    
```

Here is an alternative version, using the “anonymous” versions of `have`:

```

variable U : Type
variables A B : U → Prop

example : (∀ x, A x) → (∀ x, B x) → (∀ x, A x ∧ B x) :=
assume hA : ∀ x, A x,
assume hB : ∀ x, B x,
assume y,
have A y, from hA y,
have B y, from hB y,
show A y ∧ B y, from and.intro <A y> <B y>

```

The exercises below ask you to prove the barber paradox, which was discussed in the last chapter. You can do that using only propositional reasoning and the rules for the universal quantifier that we have just discussed.

9.3 Using the Existential Quantifier

In Lean, you can type the existential quantifier, \exists , by writing `\ex`. If you prefer you can use the ascii equivalent, `exists`. The introduction rule is `exists.intro` and requires two arguments: a term, and a proof that that term satisfies the required property.

```

variable U : Type
variable P : U → Prop

example (y : U) (h : P y) : ∃ x, P x :=
exists.intro y h

```

The elimination rule for the existential quantifier is given by `exists.elim`. It follows the form of the natural deduction rule: if we know $\exists x, P x$ and we are trying to prove Q , it suffices to introduce a new variable, y , and prove Q under the assumption that $P y$ holds.

```

variable U : Type
variable P : U → Prop
variable Q : Prop

example (h1 : ∃ x, P x) (h2 : ∀ x, P x → Q) : Q :=
exists.elim h1
  (assume (y : U) (h : P y),
    have h3 : P y → Q, from h2 y,
    show Q, from h3 h)

```

As usual, we can leave off the information as to the data type of y and the hypothesis h after the `assume`, since Lean can figure them out from the context. Deleting the `show` and replacing `h3` by its proof, `h2 y`, yields a short (though virtually unreadable) proof of the conclusion.

```

example (h1 : ∃ x, P x) (h2 : ∀ x, P x → Q) : Q :=
exists.elim h1 (assume y h, h2 y h)

```

The following example uses both the introduction and the elimination rules for the existential quantifier.

```

variable U : Type
variables A B : U → Prop

```

(continues on next page)

Here is a proof of the same implication in Lean:

```
variable U : Type
variables A B : U → Prop

example : (∀ x, A x → ¬ B x) → ¬ ∃ x, A x ∧ B x :=
assume h1 : ∀ x, A x → ¬ B x,
assume h2 : ∃ x, A x ∧ B x,
exists.elim h2 $
assume x (h3 : A x ∧ B x),
have h4 : A x, from and.left h3,
have h5 : B x, from and.right h3,
have h6 : ¬ B x, from h1 x h4,
show false, from h6 h5
```

Here, we use `exists.elim` to introduce a value `x` satisfying `A x ∧ B x`. The name is arbitrary; we could just as well have used `z`:

```
example : (∀ x, A x → ¬ B x) → ¬ ∃ x, A x ∧ B x :=
assume h1 : ∀ x, A x → ¬ B x,
assume h2 : ∃ x, A x ∧ B x,
exists.elim h2 $
assume z (h3 : A z ∧ B z),
have h4 : A z, from and.left h3,
have h5 : B z, from and.right h3,
have h6 : ¬ B z, from h1 z h4,
show false, from h6 h5
```

Here is another example of the exists-elimination rule:

```
variable U : Type
variable u : U
variable P : Prop

example : (∃ x : U, P) ↔ P :=
iff.intro
  (assume h1 : ∃ x, P,
   exists.elim h1 $
     assume x (h2 : P),
     h2)
  (assume h1 : P,
   exists.intro u h1)
```

It is subtle: the proof does not go through if we do not declare a variable `u` of type `U`, even though `u` does not appear in the statement of the theorem. The semantics of first-order logic, discussed in the next chapter, presuppose that the universe is nonempty. In Lean, however, it is possible for a type to be empty, and so the proof above depends on the fact that there is an element `u` in `U`.

These features are all illustrated in the following example:

```
variable U : Type
variables P R : U → Prop
variable Q : Prop
```

(continues on next page)

(continued from previous page)

```
example (h1 :  $\exists x, P x \wedge R x$ ) (h2 :  $\forall x, P x \rightarrow R x \rightarrow Q$ ) : Q :=
let ⟨y, hPy, hRy⟩ := h1 in
show Q, from h2 y hPy hRy
```

9.4 Equality and calculational proofs

In Lean, reflexivity, symmetry, and transitivity are called `eq.refl`, `eq.symm`, and `eq.trans`, and the second substitution rule is called `eq.subst`. Their uses are illustrated below.

```
variable A : Type

variables x y z : A
variable P : A → Prop

example : x = x :=
show x = x, from eq.refl x

example : y = x :=
have h : x = y, from sorry,
show y = x, from eq.symm h

example : x = z :=
have h1 : x = y, from sorry,
have h2 : y = z, from sorry,
show x = z, from eq.trans h1 h2

example : P y :=
have h1 : x = y, from sorry,
have h2 : P x, from sorry,
show P y, from eq.subst h1 h2
```

The rule `eq.refl` above assumes `x` as an argument, because there is no hypothesis to infer it from. All the other rules assume their premises as arguments. Here is an example of equational reasoning:

```
variables (A : Type) (x y z : A)

example : y = x → y = z → x = z :=
assume h1 : y = x,
assume h2 : y = z,
have h3 : x = y, from eq.symm h1,
show x = z, from eq.trans h3 h2
```

This proof can be written more concisely:

```
example : y = x → y = z → x = z :=
assume h1 h2, eq.trans (eq.symm h1) h2
```

Because calculations are so important in mathematics, however, Lean provides more efficient ways of carrying them out. One method is to use the `rewrite` tactic. Typing `begin` and `end` anywhere a proof is expected puts Lean into *tactic mode*, which provides an alternative way of writing a proof: rather than writing it directly, you provide Lean with a list of instructions that show Lean how to construct a proof of the

statement in question. The statement to be proved is called the *goal*, and many instructions make progress by transforming the goal into something that is easier to prove. The `rewrite` command, which carries out a substitution on the goal, is a good example. The previous example can be proved as follows:

```
example : y = x → y = z → x = z :=
assume h1 : y = x,
assume h2 : y = z,
show x = z,
begin
  rewrite ←h1,
  apply h2
end
```

If you put the cursor after the word `begin`, Lean will tell you that the goal at that point is to prove $x = z$. The first command changes the goal $x = z$ to $y = z$; the left-facing arrow before `h1` (which you can enter as `\<-`) tells Lean to use the equation in the reverse direction. If you put the cursor after the comma, Lean shows you the new goal, $y = z$. The `apply` command uses `h2` to complete the proof.

An alternative is to rewrite the goal using `h1` and `h2`, which reduces the goal to $x = x$. When that happens, `rewrite` automatically applies reflexivity.

```
example : y = x → y = z → x = z :=
assume h1 : y = x,
assume h2 : y = z,
show x = z,
begin
  rw ←h1,
  rw h2
end
```

In fact, a sequence of rewrites can be combined, using square brackets:

```
example : y = x → y = z → x = z :=
assume h1 : y = x,
assume h2 : y = z,
show x = z,
begin
  rw [←h1, h2]
end
```

And when you reduce a proof to a single tactic, you can use `by` instead of `begin ... end`.

```
example : y = x → y = z → x = z :=
assume h1 : y = x,
assume h2 : y = z,
show x = z, by rw [←h1, h2]
```

If you put the cursor after the `←h1`, Lean shows you the goal at that point.

We will see in the coming chapters that in ordinary mathematical proofs, one commonly carries out calculations in a format like this:

$$\begin{aligned} t_1 &= t_2 \\ \dots &= t_3 \\ \dots &= t_4 \\ \dots &= t_5. \end{aligned}$$

Lean has a mechanism to model such calculational proofs. Whenever a proof of an equation is expected, you can provide a proof using the identifier `calc`, following by a chain of equalities and justification, in the following form:

```
calc
  e1 = e2    : justification 1
    ... = e3 : justification 2
    ... = e4 : justification 3
    ... = e5 : justification 4
```

The chain can go on as long as needed, and in this example the result is a proof of $e1 = e5$. Each justification is the name of the assumption or theorem that is used. For example, the previous proof could be written as follows:

```
example : y = x → y = z → x = z :=
assume h1 : y = x,
assume h2 : y = z,
calc
  x = y : eq.symm h1
  ... = z : h2
```

As usual, the syntax is finicky; notice that there are no commas in the `calc` expression, and the colons and dots need to be entered exactly in that form. All that varies are the expressions `e1`, `e2`, `e3`, ... and the justifications themselves.

The `calc` environment is most powerful when used in conjunction with `rewrite`, since we can then rewrite expressions with facts from the library. For example, Lean's library has a number of basic identities for the integers, such as these:

```
variables x y z : int

example : x + 0 = x :=
add_zero x

example : 0 + x = x :=
zero_add x

example : (x + y) + z = x + (y + z) :=
add_assoc x y z

example : x + y = y + x :=
add_comm x y

example : (x * y) * z = x * (y * z) :=
mul_assoc x y z

example : x * y = y * x :=
mul_comm x y

example : x * (y + z) = x * y + x * z :=
left_distrib x y z

example : (x + y) * z = x * z + y * z :=
right_distrib x y z
```

You can also write the type of integers as \mathbb{Z} , entered with either `\Z` or `\int`. Notice that, for example, `add_comm` is the theorem $\forall x y, x + y = y + x$. So to instantiate it to $s + t = t + s$, you write `add_comm s t`. Using these axioms, here is the calculation above rendered in Lean, as a theorem about the integers:

```
example (x y z : int) : (x + y) + z = (x + z) + y :=
calc
  (x + y) + z = x + (y + z) : add_assoc x y z
  ... = x + (z + y) : eq.subst (add_comm y z) rfl
  ... = (x + z) + y : eq.symm (add_assoc x z y)
```

Using `rewrite` is more efficient, though at times we have to provide information to specify where the rules are used:

```
example (x y z : int) : (x + y) + z = (x + z) + y :=
calc
  (x + y) + z = x + (y + z) : by rw add_assoc
  ... = x + (z + y) : by rw [add_comm y z]
  ... = (x + z) + y : by rw add_assoc
```

In that case, we can use a single `rewrite`:

```
example (x y z : int) : (x + y) + z = (x + z) + y :=
by rw [add_assoc, add_comm y z, add_assoc]
```

If you `#check` the proof before the sequence of rewrites is sufficient, the error message will display the remaining goal.

Here is another example:

```
variables a b d c : int

example : (a + b) * (c + d) = a * c + b * c + a * d + b * d :=
calc
  (a + b) * (c + d) = (a + b) * c + (a + b) * d : by rw left_distrib
  ... = (a * c + b * c) + (a + b) * d : by rw right_distrib
  ... = (a * c + b * c) + (a * d + b * d) : by rw right_distrib
  ... = a * c + b * c + a * d + b * d : by rw ←add_assoc
```

Once again, we can get by with a shorter proof:

```
example : (a + b) * (c + d) = a * c + b * c + a * d + b * d :=
by rw [left_distrib, right_distrib, right_distrib, ←add_assoc]
```

9.5 Exercises

1. Fill in the sorry.

```
section
  variable A : Type
  variable f : A → A
  variable P : A → Prop
  variable h : ∀ x, P x → P (f x)
```

(continues on next page)

(continued from previous page)

```

-- Show the following:
example :  $\forall y, P y \rightarrow P (f (f y)) :=$ 
  sorry
end

```

2. Fill in the sorry.

```

section
  variable U : Type
  variables A B : U  $\rightarrow$  Prop

  example :  $(\forall x, A x \wedge B x) \rightarrow \forall x, A x :=$ 
    sorry
end

```

3. Fill in the sorry.

```

section
  variable U : Type
  variables A B C : U  $\rightarrow$  Prop

  variable h1 :  $\forall x, A x \vee B x$ 
  variable h2 :  $\forall x, A x \rightarrow C x$ 
  variable h3 :  $\forall x, B x \rightarrow C x$ 

  example :  $\forall x, C x :=$ 
    sorry
end

```

4. Fill in the sorry's below, to prove the barber paradox.

```

open classical -- not needed, but you can use it

-- This is an exercise from Chapter 4. Use it as an axiom here.
axiom not_iff_not_self (P : Prop) :  $\neg (P \leftrightarrow \neg P)$ 

example (Q : Prop) :  $\neg (Q \leftrightarrow \neg Q) :=$ 
  not_iff_not_self Q

section
  variable Person : Type
  variable shaves : Person  $\rightarrow$  Person  $\rightarrow$  Prop
  variable barber : Person
  variable h :  $\forall x, shaves\ barber\ x \leftrightarrow \neg shaves\ x\ x$ 

  -- Show the following:
  example : false :=
    sorry
end

```

5. Fill in the sorry.

```

section
  variable U : Type
  variables A B : U → Prop

  example : (∃ x, A x) → ∃ x, A x ∨ B x :=
  sorry
end

```

6. Fill in the sorry.

```

section
  variable U : Type
  variables A B : U → Prop

  variable h1 : ∀ x, A x → B x
  variable h2 : ∃ x, A x

  example : ∃ x, B x :=
  sorry
end

```

7. Fill in the sorry.

```

variable U : Type
variables A B C : U → Prop

example (h1 : ∃ x, A x ∧ B x) (h2 : ∀ x, B x → C x) :
  ∃ x, A x ∧ C x :=
  sorry

```

8. Complete these proofs.

```

variable U : Type
variables A B C : U → Prop

example : (¬ ∃ x, A x) → ∀ x, ¬ A x :=
  sorry

example : (∀ x, ¬ A x) → ¬ ∃ x, A x :=
  sorry

```

9. Fill in the sorry.

```

variable U : Type
variables R : U → U → Prop

example : (∃ x, ∀ y, R x y) → ∀ y, ∃ x, R x y :=
  sorry

```

10. The following exercise shows that in the presence of reflexivity, the rules for symmetry and transitivity are equivalent to a single rule.

```

theorem foo {A : Type} {a b c : A} : a = b → c = b → a = c :=
  sorry

```

(continues on next page)

(continued from previous page)

```

-- notice that you can now use foo as a rule. The curly braces mean that
-- you do not have to give A, a, b, or c

section
  variable A : Type
  variables a b c : A

  example (h1 : a = b) (h2 : c = b) : a = c :=
    foo h1 h2
end

section
  variable {A : Type}
  variables {a b c : A}

  -- replace the sorry with a proof, using foo and rfl, without using eq.symm.
  theorem my_symm (h : b = a) : a = b :=
    sorry

  -- now use foo and my_symm to prove transitivity
  theorem my_trans (h1 : a = b) (h2 : b = c) : a = c :=
    sorry
end

```

11. Replace each sorry below by the correct axiom from the list.

```

-- these are the axioms for a commutative ring

#check @add_assoc
#check @add_comm
#check @add_zero
#check @zero_add
#check @mul_assoc
#check @mul_comm
#check @mul_one
#check @one_mul
#check @left_distrib
#check @right_distrib
#check @add_left_neg
#check @add_right_neg
#check @sub_eq_add_neg

variables x y z : int

theorem t1 : x - x = 0 :=
  calc
x - x = x + -x : by rw sub_eq_add_neg
  ... = 0      : by rw add_right_neg

theorem t2 (h : x + y = x + z) : y = z :=
  calc

```

(continues on next page)

(continued from previous page)

```

y      = 0 + y      : by rw zero_add
...    = (-x + x) + y : by rw add_left_neg
...    = -x + (x + y) : by rw add_assoc
...    = -x + (x + z) : by rw h
...    = (-x + x) + z : by rw add_assoc
...    = 0 + z      : by rw add_left_neg
...    = z          : by rw zero_add

theorem t3 (h : x + y = z + y) : x = z :=
calc
x      = x + 0      : sorry
...    = x + (y + -y) : sorry
...    = (x + y) + -y : sorry
...    = (z + y) + -y : by rw h
...    = z + (y + -y) : sorry
...    = z + 0      : sorry
...    = z          : sorry

theorem t4 (h : x + y = 0) : x = -y :=
calc
x      = x + 0      : by rw add_zero
...    = x + (y + -y) : by rw add_right_neg
...    = (x + y) + -y : by rw add_assoc
...    = 0 + -y      : by rw h
...    = -y          : by rw zero_add

theorem t5 : x * 0 = 0 :=
have h1 : x * 0 + x * 0 = x * 0 + 0, from
calc
  x * 0 + x * 0 = x * (0 + 0) : sorry
  ... = x * 0          : sorry
  ... = x * 0 + 0      : sorry,
show x * 0 = 0, from t2 _ _ _ h1

theorem t6 : x * (-y) = -(x * y) :=
have h1 : x * (-y) + x * y = 0, from
calc
  x * (-y) + x * y = x * (-y + y) : sorry
  ... = x * 0          : sorry
  ... = 0              : by rw t5 x,
show x * (-y) = -(x * y), from t4 _ _ _ h1

theorem t7 : x + x = 2 * x :=
calc
x + x = 1 * x + 1 * x : by rw one_mul
... = (1 + 1) * x    : sorry
... = 2 * x          : rfl

```

SEMANTICS OF FIRST ORDER LOGIC

In Chapter 6, we emphasized a distinction between the *syntax* and the *semantics* of propositional logic. Syntactic questions have to do with the formal structure of formulas and the conditions under which different types of formulas can be derived. Semantic questions, on the other hand, concern the *truth* of a formula relative to some truth assignment.

As you might expect, we can make a similar distinction in the setting of first order logic. The previous two chapters have focused on syntax, but some semantic ideas have slipped in. Recall the running example with domain of interest \mathbb{N} , constant symbols 0, 1, 2, 3, function symbols *add* and *mul*, and predicate symbols *even*, *prime*, *equals*, *le*, etc. We know that the sentence $\forall y \text{le}(0, y)$ is true in this example, if *le* is interpreted as the less-than-or-equal-to relation on the natural numbers. But if we consider the domain \mathbb{Z} instead of \mathbb{N} , that same formula becomes false. The sentence $\forall y \text{lt}(0, y)$ is also false if we consider the domain \mathbb{N} , but (somewhat perversely) interpret the predicate $\text{lt}(x, y)$ as the relation “*x* is greater than *y*” on the natural numbers.

This indicates that the truth or falsity of a first order sentence can depend on how we interpret the quantifiers and basic relations of the language. But some formulas are true under any interpretation: for instance, $\forall y (\text{le}(0, y) \rightarrow \text{le}(0, y))$ is true under all the interpretations considered in the last paragraph, and, indeed, under any interpretation we choose. A sentence like this is said to be *valid*; this is the analogue of a tautology in propositional logic, which is true under every possible truth assignment.

We can broaden the analogy: a “model” in first order logic is the analogue of a truth assignment in propositional logic. In the propositional case, choosing a truth assignment allowed us to assign truth values to all formulas of the language; now, choosing a model will allow us to assign truth values to all sentences of a first order language. The aim of the next section is to make this notion more precise.

10.1 Interpretations

The symbols of the language in our running example—0, 1, *add*, *prime*, and so on—have very suggestive names. When we interpret sentences of this language over the domain \mathbb{N} , for example, it is clear for which elements of the domain *prime* “should” be true, and for which it “should” be false. But let us consider a first order language that has only two unary predicate symbols *fancy* and *tall*. If we take our domain to be \mathbb{N} , is the sentence $\forall x (\text{fancy}(x) \rightarrow \text{tall}(x))$ true or false?

The answer, of course, is that we don’t have enough information to say. There’s no obvious meaning to the predicates *fancy* or *tall*, at least not when we apply them to natural numbers. To make sense of the sentence, we need to know which numbers are fancy and which ones are tall. Perhaps multiples of 10 are fancy, and even numbers are tall; in this case, the formula is true, since every multiple of 10 is even. Perhaps prime numbers are fancy and odd numbers are tall; then the formula is false, since 2 is fancy but not tall.

We call each of these descriptions an *interpretation* of the predicate symbols *fancy* and *tall* in the domain \mathbb{N} . Formally, an interpretation of a unary predicate *P* in a domain *D* is the set of elements of *D* for which

P is true. For an example, the standard interpretation of *prime* in \mathbb{N} that we used above was just the set of prime natural numbers.

We can interpret constant, function, and relation symbols in a similar way. An interpretation of constant symbol c in domain D is an element of D . An interpretation of a function symbol f with arity n is a function that maps n elements of D to another element of D . An interpretation of a relation symbol R with arity n is the set of n tuples of elements of D for which R is true.

It is important to emphasize the difference between a syntactic predicate symbol (or function symbol, or constant symbol) and the semantic predicate (or function, or object) to which it is interpreted. The former is a symbol, relates to other symbols, and has no meaning on its own until we specify an interpretation. Strictly speaking, it makes no sense to write $prime(3)$, where *prime* is a predicate symbol and 3 is a natural number, since the argument to *prime* is supposed to be a syntactic term. Sometimes we may obscure this distinction, as above when we specified a language with constant symbols 0, 1, and 2. But there is still a fundamental distinction between the objects of the domain and the symbols we use to represent them.

Sometimes, when we interpret a language in a particular domain, it is useful to implicitly introduce new constant symbols into the language to denote elements of this domain. Specifically, for each element a of the domain, we introduce a constant symbol \bar{a} , which is interpreted as a . Then the expression $prime(\bar{3})$ does make sense. Interpreting the predicate symbol *prime* in the natural way, this expression will evaluate to true. We think of $\bar{3}$ as a linguistic “name” that represents the natural number 3, in the same way that the phrase “Aristotle” is a name that represents the ancient Greek philosopher.

10.2 Truth in a Model

Fix a first-order language. Suppose we have chosen a domain D in which to interpret the language, along with an interpretation in D of each of the symbols of that language. We will call this structure—the domain D , paired with the interpretation—a *model* for the language. A model for a first-order language is directly analogous to a truth assignment for propositional logic, because it provides all the information we need to determine the truth value of each sentence in the language.

The procedure for evaluating the truth of a sentence based on a model works the way you think it should, but the formal description is subtle. Recall the difference between *terms* and *assertions* that we made earlier in Chapter 4. Terms, like a , $x + y$, or $f(c)$, are meant to represent objects. A term does not have a truth value, since (for example) it makes no sense to ask whether 3 is true or false. Assertions, like $P(a)$, $R(x, f(y))$, or $a + b > a \wedge prime(c)$, apply predicate or relation symbols to terms to produce statements that could be true or false.

The interpretation of a term in a model is an element of the domain of that model. The model directly specifies how to interpret constant symbols. To interpret a term $f(t)$ created by applying a function symbol to another term, we interpret the term t , and then apply the interpretation of f to this term. (This process makes sense, since the interpretation of f is a function on the domain.) This generalizes to functions of higher arity in the obvious way. We will not yet interpret terms that include free variables like x and y , since these terms do not pick out unique elements of the domain. (The variable x could potentially refer to any object.)

For example, suppose we have a language with two constant symbols, a and b , a unary function symbol f , and a binary function symbol g . Let \mathcal{M} be the model with domain \mathbb{N} , where a and b are interpreted as 3 and 5, respectively, $f(x)$ is interpreted as the function which maps any natural number n to n^2 , and g is the addition function. Then the term $g(f(a), b)$ denotes the natural number $3^2 + 5 = 14$.

Similarly, the interpretation of an assertion is a value \mathbf{T} or \mathbf{F} . For the sake of brevity, we will introduce new notation here: if A is an assertion and \mathcal{M} is a model of the language of A , we write $\mathcal{M} \models A$ to mean that A evaluates to \mathbf{T} in \mathcal{M} , and $\mathcal{M} \not\models A$ to mean that A evaluates to \mathbf{F} . (You can read the symbol \models as “models” or “satisfies” or “validates.”)

To interpret a predicate or relation applied to some terms, we first interpret the terms as objects in the domain, and then see if the interpretation of the relation symbol is true of those objects. To continue with the example, suppose our language also has a relation symbol R , and we extend \mathcal{M} to interpret R as the greater-than-or-equal-to relation. Then we have $\mathcal{M} \not\models R(a, b)$, since 3 is not greater than 5, but $\mathcal{M} \models R(g(f(a), b), b)$, since 14 is greater than 5.

Interpreting expressions using the logical connectives \wedge , \vee , \rightarrow , and \neg works exactly as it did in the propositional setting. $\mathcal{M} \models A \wedge B$ exactly when $\mathcal{M} \models A$ and $\mathcal{M} \models B$, and so on.

We still need to explain how to interpret existential and universal expressions. We saw that $\exists x A$ intuitively meant that there was *some* element of the domain that would make A true, when we “replaced” the variable x with that element. To make this a bit more precise, we say that $\mathcal{M} \models \exists x A$ exactly when there is an element a in the domain of \mathcal{M} such that, when we interpret x as a , then $\mathcal{M} \models A$. To continue the example above, we have $\mathcal{M} \models \exists x (R(x, b))$, since when we interpret x as 6 we have $\mathcal{M} \models R(x, b)$.

More concisely, we can say that $\mathcal{M} \models \exists x A$ when there is an a in the domain of \mathcal{M} such that $\mathcal{M} \models A[\bar{a}/x]$. The notation $A[\bar{a}/x]$ indicates that every occurrence of x in A has been replaced by the symbol \bar{a} .

Finally, remember, $\forall x A$ means that A is true for all possible values of x . We make this precise by saying that $\mathcal{M} \models \forall x A$ exactly when for every element a in the domain of \mathcal{M} , interpreting x as a gives that $\mathcal{M} \models A$. Alternatively, we can say that $\mathcal{M} \models \forall x A$ when for every a in the domain of \mathcal{M} , we have $\mathcal{M} \models A[\bar{a}/x]$. In our example above, $\mathcal{M} \not\models \forall x (R(x, b))$, since when we interpret x as 2 we do not have $\mathcal{M} \models R(x, b)$.

These rules allow us to determine the truth value of any *sentence* in a model. (Remember, a sentence is a formula with no free variables.) There are some subtleties: for instance, we’ve implicitly assumed that our formula doesn’t quantify over the same variable twice, as in $\forall x \exists x A$. But for the most part, the interpretation process tells us to “read” a formula as talking directly about objects in the domain.

10.3 Examples

Take a simple language with no constant symbols, one relation symbol \leq , and one binary function symbol $+$. Our model \mathcal{M} will have domain \mathbb{N} , and the symbols will be interpreted as the standard less-than-or-equal-to relation and addition function.

Think about the following questions before you read the answers below. Remember, our domain is \mathbb{N} , not \mathbb{Z} or any other number system.

1. Is it true that $\mathcal{M} \models \exists x (x \leq x)$? What about $\mathcal{M} \models \forall x (x \leq x)$?
2. Similarly, what about $\mathcal{M} \models \exists x (x + x \leq x)$? $\mathcal{M} \models \forall x (x + x \leq x)$?
3. Do the sentences $\exists x \forall y (x \leq y)$ and $\forall x \exists y (x \leq y)$ mean the same thing? Are they true or false?
4. Can you think of a formula A in this language, with one free variable x , such that $\mathcal{M} \models \forall x A$ but $\mathcal{M} \not\models \exists x A$?

These questions indicate a subtle, and often tricky, interplay between the universal and existential quantifiers. Once you’ve thought about them a bit, read the answers:

1. Both of these statements are true. For the former, we can (for example) interpret x as the natural number 0. Then, $\mathcal{M} \models x \leq x$, so the existential is true. For the latter, pick an arbitrary natural number n ; it is still the case that when we interpret x as n , we have $\mathcal{M} \models x \leq x$.
2. The first statement is true, since we can interpret x as 0. The second statement, though, is false. When we interpret x as 1 (or, in fact, as any natural number besides 0), we see that $\mathcal{M} \not\models x + x \leq x$.
3. These sentences do *not* mean the same thing, although in the specified model, both are true. The first expresses that some natural number is less than or equal to every natural number. This is true: 0 is less than or equal to every natural number. The second sentence says that for every natural number,

there is another natural number at least as big. Again, this is true: every natural number a is less than or equal to a . If we took our domain to be \mathbb{Z} instead of \mathbb{N} , the first sentence would be false, while the second would still be true.

4. The situation described here is impossible in our model. If $\mathcal{M} \models \forall xA$, then $\mathcal{M} \models A[\bar{0}/x]$, which implies that $\mathcal{M} \models \exists xA$. The only time this situation can happen is when the domain of our model is empty.

Now consider a different language with constant symbol 2, predicate symbols *prime* and *odd*, and binary relation $<$, interpreted in the natural way over domain \mathbb{N} . The sentence $\forall x (2 < x \wedge \text{prime}(x) \rightarrow \text{odd}(x))$ expresses the fact that every prime number bigger than 2 is odd. It is an example of *relativization*, discussed in Section 7.4. We can now see semantically how relativization works. This sentence is true in our model if, for every natural number n , interpreting x as n makes the sentence true. If we interpret x as 0, 1, or 2, or as any non-prime number, the hypothesis of the implication is false, and thus $2 < x \wedge \text{prime}(x) \rightarrow \text{odd}(x)$ is true. Otherwise, if we interpret x as a prime number bigger than 2, both the hypothesis and conclusion of the implication are true, and $2 < x \wedge \text{prime}(x) \rightarrow \text{odd}(x)$ is again true. Thus the universal statement holds. It was an example like this that partially motivated our semantics for implication back in Chapter 3; any other choice would make relativization impossible.

For the next example, we will consider models that are given by a rectangular grid of “dots.” Each dot has a color (red, blue, or green) and a size (small or large). We use the letter R to represent a large red dot and r to represent a small red dot, and similarly for G, g, B, b .

The logical language we use to describe our dot world has predicates *red*, *green*, *blue*, *small* and *large*, which are interpreted in the obvious ways. The relation $\text{adj}(x, y)$ is true if the dots referred to by x and y are touching, not on a diagonal. The relations $\text{same-color}(x, y)$, $\text{same-size}(x, y)$, $\text{same-row}(x, y)$, and $\text{same-column}(x, y)$ are also self-explanatory. The relation $\text{left-of}(x, y)$ is true if the dot referred to by x is left of the dot referred to by y , regardless of what rows the dots are in. The interpretations of *right-of*, *above*, and *below* are similar.

Consider the following sentences:

1. $\forall x (\text{green}(x) \vee \text{blue}(x))$
2. $\exists x, y (\text{adj}(x, y) \wedge \text{green}(x) \wedge \text{green}(y))$
3. $\exists x ((\exists z \text{right-of}(z, x)) \wedge (\forall y (\text{left-of}(x, y) \rightarrow \text{blue}(y) \vee \text{small}(y))))$
4. $\forall x (\text{large}(x) \rightarrow \exists y (\text{small}(y) \wedge \text{adj}(x, y)))$
5. $\forall x (\text{green}(x) \rightarrow \exists y (\text{same-row}(x, y) \wedge \text{blue}(y)))$
6. $\forall x, y (\text{same-row}(x, y) \wedge \text{same-column}(x, y) \rightarrow x = y)$
7. $\exists x \forall y (\text{adj}(x, y) \rightarrow \neg \text{same-size}(x, y))$
8. $\forall x \exists y (\text{adj}(x, y) \wedge \text{same-color}(x, y))$
9. $\exists y \forall x (\text{adj}(x, y) \wedge \text{same-color}(x, y))$
10. $\exists x (\text{blue}(x) \wedge \exists y (\text{green}(y) \wedge \text{above}(x, y)))$

We can evaluate them in this particular model:

R	r	g	b
R	b	G	b
B	B	B	b

There they have the following truth values:

1. false

2. true
3. true
4. false
5. true
6. true
7. false
8. true
9. false
10. true

For each sentence, see if you can find a model that makes the sentence true, and another that makes it false. For an extra challenge, try to make all of the sentences true simultaneously. Notice that you can use any number of rows and any number of columns.

10.4 Validity and Logical Consequence

We have seen that whether a formula is true or false often depends on the model we choose. Some formulas, though, are true in every possible model. An example we saw earlier was $\forall y (le(0, y) \rightarrow le(0, y))$. Why is this sentence valid? Suppose \mathcal{M} is an arbitrary model of the language, and suppose a is an arbitrary element of the domain of \mathcal{M} . Either $\mathcal{M} \models le(0, \bar{a})$ or $\mathcal{M} \models \neg le(0, \bar{a})$. In either case, the propositional semantics of implication guarantee that $\mathcal{M} \models le(0, \bar{a}) \rightarrow le(0, \bar{a})$. We often write $\models A$ to mean that A is valid.

In the propositional setting, there is an easy method to figure out if a formula is a tautology or not. Writing the truth table and checking for any rows ending with **F** is algorithmic, and we know from the beginning exactly how large the truth table will be. Unfortunately, we cannot do the same for first-order formulas. Any language has infinitely many models, so a “first-order” truth table would be infinitely long. To make matters worse, even checking whether a formula is true in a single model can be a non-algorithmic task. To decide whether a universal statement like $\forall x P(x)$ is true in a model with an infinite domain, we might have to check whether P is true of infinitely many elements.

This is not to say that we can *never* figure out if a first-order sentence is a tautology. For example, we have argued that $\forall y (lt(0, y) \rightarrow lt(0, y))$ was one. It is just a more difficult question than for propositional logic.

As was the case with propositional logic, we can extend the notion of validity to a notion of logical consequence. Fix a first-order language, L . Suppose Γ is a set of sentences in L , and A is a sentence of L . We will say that A is a *logical consequence* of Γ if every model of Γ is a model of A . This is one way of spelling out that A is a “necessary consequence” of A : under any interpretation, if the hypotheses in Γ come out true, A is true as well.

10.5 Soundness and Completeness

In propositional logic, we saw a close connection between the provable formulas and the tautologies—specifically, a formula is provable if and only if it is a tautology. More generally, we say that a formula A is a logical consequence of a set of hypotheses, Γ , if and only if there is a natural deduction proof of A from Γ . It turns out that the analogous statements hold for first order logic.

The “soundness” direction—the fact that if A is provable from Γ then A is true in any model of Γ —holds for reasons that are similar to the reasons it holds in the propositional case. Specifically, the proof proceeds by showing that each rule of natural deduction preserves the truth in a model.

The completeness theorem for first order logic was first proved by Kurt Gödel in his 1929 dissertation. Another, simpler proof was later provided by Leon Henkin.

Theorem. If a formula A is a logical consequence of a set of sentences Γ , then A is provable from Γ .

Compared to the version for propositional logic, the first order completeness theorem is harder to prove. We will not go into too much detail here, but will indicate some of the main ideas. A set of sentences is said to be *consistent* if you cannot prove a contradiction from those hypotheses. Most of the work in Henkin's proof is done by the following "model existence" theorem:

Theorem. Every consistent set of sentences has a model.

From this theorem, it is easy to deduce the completeness theorem. Suppose there is no proof of A from Γ . Then the set $\Gamma \cup \{\neg A\}$ is consistent. (If we could prove \perp from $\Gamma \cup \{\neg A\}$, then by the *reductio ad absurdum* rule we could prove A from Γ .) By the model existence theorem, that means that there is a model \mathcal{M} of $\Gamma \cup \{\neg A\}$. But this is a model of Γ that is not a model of A , which means that A is not a logical consequence of Γ .

The proof of the model existence theorem is intricate. Somehow, from a consistent set of sentences, one has to build a model. The strategy is to build the model out of syntactic entities, in other words, to use terms in an expanded language as the elements of the domain.

The moral here is much the same as it was for propositional logic. Because we have developed our syntactic rules with a certain semantics in mind, the two exhibit different sides of the same coin: the provable sentences are exactly the ones that are true in all models, and the sentences that are provable from a set of hypotheses are exactly the ones that are true in all models of those hypotheses.

We therefore have another way to answer the question posed in the previous section. To show that a sentence is valid, there is no need to check its truth in every possible model. Rather, it suffices to produce a proof.

10.6 Exercises

1. In a first-order language with a binary relation, $R(x, y)$, consider the following sentences:

- $\exists x \forall y R(x, y)$
- $\exists y \forall x R(x, y)$
- $\forall x, y (R(x, y) \wedge x \neq y \rightarrow \exists z (R(x, z) \wedge R(z, y) \wedge x \neq z \wedge y \neq z))$

For each of the following structures, determine whether of each of those sentences is true or false.

- the structure (\mathbb{N}, \leq) , that is, the interpretation in the natural numbers where R is \leq
- the structure (\mathbb{Z}, \leq)
- the structure (\mathbb{Q}, \leq)
- the structure $(\mathbb{N}, |)$, that is, the interpretation in the natural numbers where R is the "divides" relation
- the structure $(P(\mathbb{N}), \subseteq)$, that is, the interpretation where variables range over sets of natural numbers, where R is interpreted as the subset relation.

2. Create a 4 x 4 “dots” world that makes all of the following sentences true:
- $\forall x (green(x) \vee blue(x))$
 - $\exists x, y (adj(x, y) \wedge green(x) \wedge green(y))$
 - $\exists x (\exists z right-of(z, x) \wedge \forall y (left-of(x, y) \rightarrow blue(y) \vee small(y)))$
 - $\forall x (large(x) \rightarrow \exists y (small(y) \wedge adj(x, y)))$
 - $\forall x (green(x) \rightarrow \exists y (same-row(x, y) \wedge blue(y)))$
 - $\forall x, y (same-row(x, y) \wedge same-column(x, y) \rightarrow x = y)$
 - $\exists x \forall y (adj(x, y) \rightarrow \neg same-size(x, y))$
 - $\forall x \exists y (adj(x, y) \wedge same-color(x, y))$
 - $\exists y \forall x (adj(x, y) \rightarrow same-color(x, y))$
 - $\exists x (blue(x) \wedge \exists y (green(y) \wedge above(x, y)))$
3. Fix a first-order language L , and let A and B be any two sentences in L . Remember that $\models A$ means that A is valid. Unpacking the definitions, show that if $\models A \wedge B$, then $\models A$ and $\models B$.
4. Give a concrete example to show that $\models A \vee B$ does not necessarily imply $\models A$ or $\models B$. In other words, pick a language L and choose particular sentences A and B such that $A \vee B$ is valid but neither A nor B is valid.

We have come to a turning point in this textbook. We will henceforth abandon natural deduction, for the most part, and focus on ordinary mathematical proofs. We will continue to think about how informal mathematics can be represented in symbolic terms, and how the rules of natural deduction play out in the informal setting. But the emphasis will be on writing ordinary mathematical arguments, not designing proof trees. Lean will continue to serve as a bridge between the informal and formal realms.

In this chapter, we consider a notion that has come to play a fundamental role in mathematical reasoning, namely, that of a “set.”

11.1 Elementary Set Theory

In a publication in the journal *Mathematische Annalen* in 1895, the German mathematician Georg Cantor presented the following characterization of the notion of a *set* (or *Menge*, in his terminology):

By a *set* we mean any collection M of determinate, distinct objects (called the *elements* of M) of our intuition or thought into a whole.

Since then, the notion of a set has been used to unify a wide range of abstractions and constructions. Axiomatic set theory, which we will discuss in a later chapter, provides a foundation for mathematics in which everything can be viewed as a set.

On a broad construal, *any* collection can be a set; for example, we can consider the set whose elements are Ringo Star, the number 7, and the set whose only member is the Empire State Building. With such a broad notion of set we have to be careful: Russell’s paradox has us consider the set S of all sets that are not elements of themselves, which leads to a contradiction when we ask whether S is an element of itself. (Try it!) The axioms of set theory tell us which sets exist, and have been carefully designed to avoid paradoxical sets like that of the Russell paradox.

In practice, mathematicians are not so freewheeling in their use of sets. Typically, one fixes a domain such as the natural numbers, and consider subsets of that domain. In other words, we consider sets of numbers, sets of points, sets of lines, and so on, rather than arbitrary “sets.” In this text, we will adopt this convention: when we talk about sets, we are always implicitly talking about sets of elements of some domain.

Given a set A of objects in some domain and an object x , we write $x \in A$ to say that x is an element of A . Cantor’s characterization suggests that whenever we have some property, P , of a domain, we can form the set of elements that have that property. This is denoted using “set-builder notation” as $\{x \mid P(x)\}$. For example, we can consider all the following sets of natural numbers:

- $\{n \mid n \text{ is even}\}$
- $\{n \mid n \text{ is prime}\}$
- $\{n \mid n \text{ is prime and greater than } 2\}$

- $\{n \mid n \text{ can be written as a sum of squares}\}$
- $\{n \mid n \text{ is equal to 1, 2, or 3}\}$

This last set is written more simply $\{1, 2, 3\}$. If the domain is not clear from the context, we can specify it by writing it explicitly, for example, in the expression $\{n \in \mathbb{N} \mid n \text{ is even}\}$.

Using set-builder notation, we can define a number of common sets and operations. The *empty set*, \emptyset , is the set with no elements:

$$\emptyset = \{x \mid \text{false}\}.$$

Dually, we can define the *universal set*, \mathcal{U} , to be the set consisting of every element of the domain:

$$\mathcal{U} = \{x \mid \text{true}\}.$$

Given two sets A and B , we define their *union* to be the set of elements in either one:

$$A \cup B = \{x \mid x \in A \text{ or } x \in B\}.$$

And we define their *intersection* to be the set of elements of both:

$$A \cap B = \{x \mid x \in A \text{ and } x \in B\}.$$

We define the *complement* of a set of A to be the set of elements that are not in A :

$$\overline{A} = \{x \mid x \notin A\}.$$

We define the *set difference* of two sets A and B to be the set of elements in A but not B :

$$A \setminus B = \{x \mid x \in A \text{ and } x \notin B\}.$$

Two sets are said to be equal if they have exactly the same elements. If A and B are sets, A is said to be a *subset* of B , written $A \subseteq B$, if every element of A is an element of B . Notice that A is equal to B if and only if A is a subset of B and B is a subset of A .

Notice also that everything we have said about sets so far is readily representable in symbolic logic. We can render the defining properties of the basic sets and constructors as follows:

- $\forall x (x \in \emptyset \leftrightarrow \perp)$
- $\forall x (x \in \mathcal{U} \leftrightarrow \top)$
- $\forall x (x \in A \cup B \leftrightarrow x \in A \vee x \in B)$
- $\forall x (x \in A \cap B \leftrightarrow x \in A \wedge x \in B)$
- $\forall x (x \in \overline{A} \leftrightarrow x \notin A)$
- $\forall x (x \in A \setminus B \leftrightarrow x \in A \wedge x \notin B)$

The assertion that A is a subset of B can be written $\forall x (x \in A \rightarrow x \in B)$, and the assertion that A is equal to B can be written $\forall x (x \in A \leftrightarrow x \in B)$. These are all *universal* statements, that is, statements with universal quantifiers in front, followed by basic assertions and propositional connectives. What this means is that reasoning about sets formally often amounts to using nothing more than the rules for the universal quantifier together with the rules for propositional logic.

Logicians sometimes describe ordinary mathematical proofs as *informal*, in contrast to the *formal proofs* in natural deduction. When writing informal proofs, the focus is on readability. Here is an example.

Theorem. Let A , B , and C denote sets of elements of some domain. Then $A \cap (B \cup C) = (A \cap B) \cup (A \cap C)$.

Proof. Let x be arbitrary, and suppose x is in $A \cap (B \cup C)$. Then x is in A , and either x is in B or x is in C . In the first case, x is in A and B , and hence in $A \cap B$. In the second case, x is in A and C , and hence $A \cap C$. Either way, we have that x is in $(A \cap B) \cup (A \cap C)$.

Conversely, suppose x is in $(A \cap B) \cup (A \cap C)$. There are now two cases.

First, suppose x is in $A \cap B$. Then x is in both A and B . Since x is in B , it is also in $B \cup C$, and so x is in $A \cap (B \cup C)$.

The second case is similar: suppose x is in $A \cap C$. Then x is in both A and C , and so also in $B \cup C$. Hence, in this case also, x is in $A \cap (B \cup C)$, as required.

Notice that this proof does not look anything like a proof in symbolic logic. For one thing, ordinary proofs tend to favor words over symbols. Of course, mathematics uses symbols all the time, but not in place of words like “and” and “not”; you will rarely, if ever, see the symbols \wedge and \neg in a mathematics textbook, unless it is a textbook specifically about logic.

Similarly, the structure of an informal proof is conveyed with ordinary paragraphs and punctuation. Don’t rely on pictorial diagrams, line breaks, and indentation to convey the structure of a proof. Rather, you should rely on literary devices like signposting and foreshadowing. It is often helpful to present an outline of a proof or the key ideas before delving into the details, and the introductory sentence of a paragraph can help guide a reader’s expectations, just as it does in an expository essay.

Nonetheless, you should be able to see elements of natural deduction implicitly in the proof above. In formal terms, the theorem is equivalent to the assertion

$$\forall x (x \in A \cap (B \cup C) \leftrightarrow x \in (A \cap B) \cup (A \cap C)),$$

and the proof proceeds accordingly. The phrase “let x be arbitrary” is code for the \forall introduction rule, and the form of the rest of the proof is a \leftrightarrow introduction. Saying that x is in $A \cap (B \cup C)$ is implicitly an “and,” and the argument uses \wedge elimination to get $x \in A$ and $x \in B \cup C$. Saying $x \in B \cup C$ is implicitly an “or,” and the proof then splits on cases, depending on whether $x \in B$ or $x \in C$.

Modulo the unfolding of definition of intersection and union in terms of “and” and “or,” the “only if” direction of the previous proof could be represented in natural deduction like this:

$$\frac{\frac{\frac{y \in A \cap (B \cup C)}{y \in B \cup C}^1}{y \in A} \quad \frac{\frac{y \in A \cap (B \cup C)}{y \in B}^1}{y \in C}^2}{\frac{y \in A \cap B}{y \in (A \cap B) \cup (A \cap C)}} \quad \frac{\frac{\frac{y \in A \cap (B \cup C)}{y \in A}^1}{y \in C}^2}{y \in (A \cap B) \cup (A \cap C)}^2}{y \in (A \cap B) \cup (A \cap C)}^2 \quad \vdots_1}{\frac{y \in A \cap (B \cup C) \leftrightarrow y \in (A \cap B) \cup (A \cap C)}{\forall x (x \in A \cap (B \cup C) \leftrightarrow x \in (A \cap B) \cup (A \cap C))}^1}$$

In the next chapter, we will see that this logical structure is made manifest in Lean. But writing long proofs in natural deduction is not the most effective to communicate the mathematical ideas. So our goal here is to teach you to think in terms of natural deduction rules, but express the steps in ordinary English.

Here is another example.

Theorem. $(A \setminus B) \setminus C = A \setminus (B \cup C)$.

Proof. Let x be arbitrary, and suppose x is in $(A \setminus B) \setminus C$. Then x is in $A \setminus B$ but not C , and hence it is in A but not in B or C . This means that x is in A but not $B \cup C$, and so in $A \setminus (B \cup C)$.

Conversely, suppose x is in $A \setminus (B \cup C)$. Then x is in A , but not in $B \cup C$. In particular, x is in neither B nor C , because otherwise it would be in $B \cup C$. So x is in $A \setminus B$, and hence in $(A \setminus B) \setminus C$.

Perhaps the biggest difference between informal proofs and formal proofs is the level of detail. Informal proofs will often skip over details that are taken to be “straightforward” or “obvious,” devoting more effort to spelling out inferences that are novel or unexpected.

Writing a good proof is like writing a good essay. To convince your readers that the conclusion is correct, you have to get them to understand the argument, without overwhelming them with unnecessary details. It helps to have a specific audience in mind. Try speaking the argument aloud to friends, roommates, and family members; if their eyes glaze over, it is unreasonable to expect anonymous readers to do better.

One of the best ways to learn to write good proofs is to *read* good proofs, and pay attention to the style of writing. Pick an example of a textbook that you find especially clear and engaging, and think about what makes it so.

Natural deduction and formal verification can help you understand the components that make a proof *correct*, but you will have to develop an intuitive feel for what makes a proof easy and enjoyable to read.

11.2 Calculations with Sets

Calculation is central to mathematics, and mathematical proofs often involve carrying out a sequence of calculations. Indeed, a calculation can be viewed as a proof in and of itself that two expressions describe the same entity.

In high school algebra, students are often asked to prove identities like the following:

Proposition. $\frac{n(n+1)}{2} + (n+1) = \frac{(n+1)(n+2)}{2}$, for every natural number n .

In some places, students are asked to write proofs like this:

Proof.

$$\begin{aligned} \frac{n(n+1)}{2} + (n+1) &= ? \frac{(n+1)(n+2)}{2} \\ \frac{n^2+n}{2} + \frac{2n+2}{2} &= ? \frac{n^2+3n+2}{2} \\ \frac{n^2+n+2n+2}{2} &= ? \frac{n^2+3n+2}{2} \\ \frac{n^2+3n+2}{2} &= \frac{n^2+3n+2}{2}. \end{aligned}$$

Mathematicians generally cringe when they see this. *Don't do it!* It looks like an instance of forward reasoning, where we start with a complex identity and end up proving $x = x$. Of course, what is really

meant is that each line follows from the next. There is a way of expressing this, with the phrase “it suffices to show.” The following presentation comes closer to mathematical vernacular:

Proof. We want to show

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

To do that, it suffices to show

$$\frac{n^2+n}{2} + \frac{2n+2}{2} = \frac{n^2+3n+2}{2}.$$

For that, it suffices to show

$$\frac{n^2+n+2n+2}{2} = \frac{n^2+3n+2}{2}.$$

But this last equation is clearly true.

The narrative doesn't flow well, however. Sometimes there are good reasons to work backward in a proof, but in this case it is easy to present the proof in a more forward-directed manner. Here is one example:

Proof. Calculating on the left-hand side, we have

$$\begin{aligned} \frac{n(n+1)}{2} + (n+1) &= \frac{n^2+n}{2} + \frac{2n+2}{2} \\ &= \frac{n^2+n+2n+2}{2} \\ &= \frac{n^2+3n+2}{2}. \end{aligned}$$

On the right-hand side, we also have

$$\frac{(n+1)(n+2)}{2} = \frac{n^2+3n+2}{2}.$$

So $\frac{n(n+1)}{2} + (n+1) = \frac{n^2+3n+2}{2}$, as required.

Mathematicians often use the abbreviations “LHS” and “RHS” for “left-hand side” and “right-hand side,” respectively, in situations like this. In fact, here we can easily write the proof as a single forward-directed calculation:

Proof.

$$\begin{aligned} \frac{n(n+1)}{2} + (n+1) &= \frac{n^2+n}{2} + \frac{2n+2}{2} \\ &= \frac{n^2+n+2n+2}{2} \\ &= \frac{n^2+3n+2}{2} \\ &= \frac{(n+1)(n+2)}{2}. \end{aligned}$$

Such a proof is clear, compact, and easy to read. The main challenge to the reader is to figure out what justifies each subsequent step. Mathematicians sometimes annotate such a calculation with additional information, or add a few words of explanation in the text before and/or after. But the ideal situation is to carry out the calculation in small enough steps so that each step is straightforward, and needs no explanation. (And, once again, what counts as “straightforward” will vary depending on who is reading the proof.)

We have said that two sets are equal if they have the same elements. In the previous section, we proved that two sets are equal by reasoning about the elements of each, but we can often be more efficient. Assuming A , B , and C are subsets of some domain \mathcal{U} , the following identities hold:

- $A \cup \bar{A} = \mathcal{U}$
- $A \cap \bar{A} = \emptyset$
- $\overline{\bar{A}} = A$
- $A \cup A = A$
- $A \cap A = A$
- $A \cup \emptyset = A$
- $A \cap \emptyset = \emptyset$
- $A \cup \mathcal{U} = \mathcal{U}$
- $A \cap \mathcal{U} = A$
- $A \cup B = B \cup A$
- $A \cap B = B \cap A$
- $(A \cup B) \cup C = A \cup (B \cup C)$
- $(A \cap B) \cap C = A \cap (B \cap C)$
- $\overline{A \cap B} = \bar{A} \cup \bar{B}$
- $\overline{A \cup B} = \bar{A} \cap \bar{B}$
- $A \cap (B \cup C) = (A \cap B) \cup (A \cap C)$
- $A \cup (B \cap C) = (A \cup B) \cap (A \cup C)$
- $A \cap (A \cup B) = A$
- $A \cup (A \cap B) = A$

This allows us to prove further identities by calculating. Here is an example.

Theorem. Let A and B be subsets of some domain \mathcal{U} . Then $(A \cap \bar{B}) \cup B = A \cup B$.

Proof.

$$\begin{aligned} (A \cap \bar{B}) \cup B &= (A \cup B) \cap (\bar{B} \cup B) \\ &= (A \cup B) \cap \mathcal{U} \\ &= A \cup B. \end{aligned}$$

Here is another example.

Theorem. Let A and B be subsets of some domain \mathcal{U} . Then $(A \setminus B) \cup (B \setminus A) = (A \cup B) \setminus (A \cap B)$.

Proof.

$$\begin{aligned}
 (A \setminus B) \cup (B \setminus A) &= (A \cap \overline{B}) \cup (B \cap \overline{A}) \\
 &= ((A \cap \overline{B}) \cup B) \cap ((A \cap \overline{B}) \cup \overline{A}) \\
 &= ((A \cup B) \cap (\overline{B} \cup B)) \cap ((A \cup \overline{A}) \cap (\overline{B} \cup \overline{A})) \\
 &= ((A \cup B) \cap \mathcal{U}) \cap (\mathcal{U} \cap \overline{B \cap A}) \\
 &= (A \cup B) \cap \overline{(A \cap B)} \\
 &= (A \cup B) \setminus (A \cap B).
 \end{aligned}$$

Classically, you may have noticed that propositions, under logical equivalence, satisfy identities similar to sets. That is no coincidence; both are instances of *boolean algebras*. Here are the identities above translated to the language of a boolean algebra:

- $A \vee \neg A = \top$
- $A \wedge \neg A = \perp$
- $\neg \neg A = A$
- $A \vee A = A$
- $A \wedge A = A$
- $A \vee \perp = A$
- $A \wedge \perp = \perp$
- $A \vee \top = \top$
- $A \wedge \top = A$
- $A \vee B = B \vee A$
- $A \wedge B = B \wedge A$
- $(A \vee B) \vee C = A \vee (B \vee C)$
- $(A \wedge B) \wedge C = A \wedge (B \wedge C)$
- $\neg A \wedge B = \neg A \vee \neg B$
- $\neg A \vee B = \neg A \wedge \neg B$
- $A \wedge (B \vee C) = (A \wedge B) \vee (A \wedge C)$
- $A \vee (B \wedge C) = (A \vee B) \wedge (A \vee C)$
- $A \wedge (A \vee B) = A$
- $A \vee (A \wedge B) = A$

Translated to the language of boolean algebras, the first theorem above is as follows:

Theorem. Let A and B be elements of a boolean algebra. Then $(A \wedge \neg B) \vee B = B$.

Proof.

$$\begin{aligned} (A \wedge \neg B) \vee B &= (A \vee B) \wedge (\neg B \vee B) \\ &= (A \vee B) \wedge \top \\ &= (A \vee B). \end{aligned}$$

11.3 Indexed Families of Sets

If I is a set, we will sometimes wish to consider a *family* $(A_i)_{i \in I}$ of sets indexed by elements of I . For example, we might be interested in a sequence

$$A_0, A_1, A_2, \dots$$

of sets indexed by the natural numbers. The concept is best illustrated by some examples.

- For each natural number n , we can define the set A_n to be the set of people alive today that are of age n . For each age we have the corresponding set. Someone of age 20 is an element of the set A_{20} , while a newborn baby is an element of A_0 . The set A_{200} is empty. This family $(A_n)_{n \in \mathbb{N}}$ is a family of sets indexed by the natural numbers.
- For every real number r we can define B_r to be the set of positive real numbers larger than r , so $B_r = \{x \in \mathbb{R} \mid x > r \text{ and } x > 0\}$. Then $(B_r)_{r \in \mathbb{R}}$ is a family of sets indexed by the real numbers.
- For every natural number n we can define $C_n = \{k \in \mathbb{N} \mid k \text{ is a divisor of } n\}$ as the set of divisors of n .

Given a family $(A_i)_{i \in I}$ of sets indexed by I , we can form its *union*:

$$\bigcup_{i \in I} A_i = \{x \mid x \in A_i \text{ for some } i \in I\}.$$

We can also form the *intersection* of a family of sets:

$$\bigcap_{i \in I} A_i = \{x \mid x \in A_i \text{ for every } i \in I\}.$$

So an element x is in $\bigcup_{i \in I} A_i$ if and only if x is in A_i for *some* i in I , and x is in $\bigcap_{i \in I} A_i$ if and only if x is in A_i for every i in I . These operations are represented in symbolic logic by the existential and the universal quantifiers. We have:

- $\forall x (x \in \bigcup_{i \in I} A_i \leftrightarrow \exists i \in I (x \in A_i))$
- $\forall x (x \in \bigcap_{i \in I} A_i \leftrightarrow \forall i \in I (x \in A_i))$

Returning to the examples above, we can compute the union and intersection of each family. For the first example, $\bigcup_{n \in \mathbb{N}} A_n$ is the set of all living people, and $\bigcap_{n \in \mathbb{N}} A_n = \emptyset$. Also, $\bigcup_{r \in \mathbb{R}} B_r = \mathbb{R}_{>0}$, the set of all positive real numbers, and $\bigcap_{r \in \mathbb{R}} B_r = \emptyset$. For the last example, we have $\bigcup_{n \in \mathbb{N}} C_n = \mathbb{N}$ and $\bigcap_{n \in \mathbb{N}} C_n = \{1\}$, since 1 is a divisor of every natural number.

Suppose that I contains just two elements, say $I = \{c, d\}$. Let $(A_i)_{i \in I}$ be a family of sets indexed by I . Because I has two elements, this family consists of just the two sets A_c and A_d . Then the union and intersection of this family are just the union and intersection of the two sets:

$$\begin{aligned} \bigcup_{i \in I} A_i &= A_c \cup A_d \\ \bigcap_{i \in I} A_i &= A_c \cap A_d. \end{aligned}$$

This means that the union and intersection of two sets are just a special case of the union and intersection of a family of sets.

We also have equalities for unions and intersections of families of sets. Here are a few of them:

- $A \cap \bigcup_{i \in I} B_i = \bigcup_{i \in I} (A \cap B_i)$
- $A \cup \bigcap_{i \in I} B_i = \bigcap_{i \in I} (A \cup B_i)$
- $\overline{\bigcap_{i \in I} A_i} = \bigcup_{i \in I} \overline{A_i}$
- $\overline{\bigcup_{i \in I} A_i} = \bigcap_{i \in I} \overline{A_i}$
- $\bigcup_{i \in I} \bigcup_{j \in J} A_{i,j} = \bigcup_{j \in J} \bigcup_{i \in I} A_{i,j}$
- $\bigcap_{i \in I} \bigcap_{j \in J} A_{i,j} = \bigcap_{j \in J} \bigcap_{i \in I} A_{i,j}$

In the last two lines, $A_{i,j}$ is indexed by two sets I and J . This means that for every $i \in I$ and $j \in J$ we have a set $A_{i,j}$. For the first four equalities, try to figure out what the rule means if the index set I contains two elements.

Let's prove the first identity. Notice how the logical forms of the assertions $x \in A \cap \bigcup_{i \in I} B_i$ and $x \in \bigcup_{i \in I} (A \cap B_i)$ dictate the structure of the proof.

Theorem. Let A be any subset of some domain U , and let $(B_i)_{i \in I}$ be a family of subsets of U indexed by I . Then

$$A \cap \bigcup_{i \in I} B_i = \bigcup_{i \in I} (A \cap B_i).$$

Proof. Suppose x is in $A \cap \bigcup_{i \in I} B_i$. Then x is in A and x is in B_j for some $j \in I$. So x is in $A \cap B_j$, and hence in $\bigcup_{i \in I} (A \cap B_i)$.

Conversely, suppose x is in $\bigcup_{i \in I} (A \cap B_i)$. Then, for some j in I , x is in $A \cap B_j$. Hence x is in A , and since x is in B_j , it is in $\bigcup_{i \in I} B_i$. Hence x is in $A \cap \bigcup_{i \in I} B_i$, as required.

11.4 Cartesian Product and Power Set

The *ordered pair* of two objects a and b is denoted (a, b) . We say that a is the *first component* and b is the *second component* of the pair. Two pairs are only equal if the first component are equal and the second components are equal. In symbols, $(a, b) = (c, d)$ if and only if $a = c$ and $b = d$.

Given two sets A and B , we define the *cartesian product* $A \times B$ of these two sets as the set of all pairs where the first component is an element in A and the second component is an element in B . In set-builder notation this means

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

Note that if A and B are subsets of a particular domain \mathcal{U} , the set $A \times B$ need not be a subset of the same domain. However, it will be a subset of $\mathcal{U} \times \mathcal{U}$.

Some axiomatic foundations take the notion of a pair to be primitive. In axiomatic set theory, it is common to *define* an ordered pair to be a particular set, namely

$$(a, b) = \{\{a\}, \{a, b\}\}.$$

Notice that if $a = b$, this set has only one element:

$$(a, a) = \{\{a\}, \{a, a\}\} = \{\{a\}, \{a\}\} = \{\{a\}\}.$$

The following theorem shows that this definition is reasonable.

Theorem. Using the definition of ordered pairs above, we have $(a, b) = (c, d)$ if and only if $a = c$ and $b = d$.

Proof. If $a = c$ and $b = d$ then clearly $(a, b) = (c, d)$. For the other direction, suppose that $(a, b) = (c, d)$, which means

$$\underbrace{\{\{a\}, \{a, b\}\}}_L = \underbrace{\{\{c\}, \{c, d\}\}}_R.$$

Suppose first that $a = b$. Then $L = \{\{a\}\}$. This means that $\{c\} = \{a\}$ and $\{c, d\} = \{a\}$, from which we conclude that $c = a$ and $d = a = b$.

Now suppose that $a \neq b$. If $\{c\} = \{a, b\}$ then we conclude that a and b are both equal to c , contradicting $a \neq b$. Since $\{c\} \in L$, $\{c\}$ must be equal to $\{a\}$, which means that $a = c$. We know that $\{a, b\} \in R$, and since we know $\{a, b\} \neq \{c\}$, we conclude $\{a, b\} = \{c, d\}$. This means that $b \in \{c, d\}$, since $b \neq a = c$, we conclude that $b = d$.

Hence in both cases we conclude that $a = c$ and $b = d$, proving the theorem.

Using ordered pairs we can define the *ordered triple* (a, b, c) to be $(a, (b, c))$. Then we can prove that $(a, b, c) = (d, e, f)$ if and only if $a = d$, $b = e$ and $c = f$, which you are asked to do in the exercises. We can also define ordered n -tuples, which are sequence of n objects, in a similar way.

Given a set A we can define the *power set* $\mathcal{P}(A)$ to be the set of all subsets of A . In set-builder notation we can write this as

$$\mathcal{P}(A) = \{B \mid B \subseteq A\}.$$

If A is a subset of \mathcal{U} , $\mathcal{P}(A)$ may not be a subset of \mathcal{U} , but it is always a subset of $\mathcal{P}(\mathcal{U})$.

11.5 Exercises

1. Prove the following theorem: Let A , B , and C be sets of elements of some domain. Then $A \cup (B \cap C) = (A \cup B) \cap (A \cup C)$. (Henceforth, if we don't specify natural deduction or Lean, "prove" and "show" mean give an ordinary mathematical proof, using ordinary mathematical language rather than symbolic logic.)
2. Prove the following theorem: Let A and B be sets of elements of some domain. Then $\overline{A \setminus B} = \overline{A} \cup B$.
3. Two sets A and B are said to be *disjoint* if they have no element in common. Show that if A and B are disjoint, $C \subseteq A$, and $D \subseteq B$, then C and D are disjoint.
4. Let A and B be sets. Show $(A \setminus B) \cup (B \setminus A) = (A \cup B) \setminus (A \cap B)$, by showing that both sides have the same elements.
5. Let A , B , and C be subsets of some domain \mathcal{U} . Give a calculational proof of the identity $A \setminus (B \cup C) = (A \setminus B) \setminus C$, using the identities above. Also use the fact that, in general, $C \setminus D = C \cap \overline{D}$.
6. Similarly, give a calculational proof of $(A \setminus B) \cup (A \cap B) = A$.
7. Give calculational proofs of the following:

- $A \setminus B = A \setminus (A \cap B)$
- $A \setminus B = (A \cup B) \setminus B$
- $(A \cap B) \setminus C = (A \setminus C) \cap B$

8. Prove that if $(A_{i,j})_{i \in I, j \in J}$ is a family indexed by two sets I and J , then

$$\bigcup_{i \in I} \bigcap_{j \in J} A_{i,j} \subseteq \bigcap_{j \in J} \bigcup_{i \in I} A_{i,j}.$$

Also, find a family $(A_{i,j})_{i \in I, j \in J}$ where the reverse inclusion does not hold.

9. Prove using calculational reasoning that

$$\left(\bigcup_{i \in I} A_i \right) \cap \left(\bigcup_{j \in J} B_j \right) = \bigcup_{\substack{i \in I \\ j \in J}} (A_i \cap B_j).$$

The notation $\bigcup_{\substack{i \in I \\ j \in J}} (A_i \cap B_j)$ means $\bigcup_{i \in I} \bigcup_{j \in J} (A_i \cap B_j)$.

10. Using the definition $(a, b, c) = (a, (b, c))$, show that $(a, b, c) = (d, e, f)$ if and only if $a = d$, $b = e$ and $c = f$.
11. Prove that $A \times (B \cup C) = (A \times B) \cup (A \times C)$
12. Prove that $(A \cap B) \times (C \cap D) = (A \times C) \cap (B \times D)$. Find an expression for $(A \cup B) \times (C \cup D)$ consisting of unions of cartesian products, and prove that your expression is correct.
13. Prove that that $A \subseteq B$ if and only if $\mathcal{P}(A) \subseteq \mathcal{P}(B)$.

SETS IN LEAN

In the last chapter, we noted that although in axiomatic set theory one considers sets of disparate objects, it is more common in mathematics to consider subsets of some fixed domain, \mathcal{U} . This is the way sets are handled in Lean. For any data type U , Lean gives us a new data type, `set U`, consisting of the sets of elements of U . Thus, for example, we can reason about sets of natural numbers, or sets of integers, or sets of pairs of natural numbers.

12.1 Basics

Given $A : \text{set } U$ and $x : U$, we can write $x \in A$ to state that x is a member of the set A . The character \in can be typed using `\in`.

```
import data.set
open set

variable {U : Type}
variables A B C : set U
variable x : U

#check x ∈ A
#check A ∪ B
#check B \ C
#check C ∩ A
#check ¬C
#check ∅ ⊆ A
#check B ⊆ univ
```

You can type the symbols \subseteq , \emptyset , \cup , \cap , \setminus as `\subeq`, `\empty`, `\un`, `\i`, and `\`, respectively. We have made the type variable U implicit, because it can typically be inferred from context. The universal set is denoted `univ`, and set complementation is denoted with a negation symbol. Basic set-theoretic notions like these are defined in Lean's core library, but additional theorems and notation are available in an auxiliary library that we have loaded with the command `import data.set`, which has to appear at the beginning of a file. The command `open set` lets us refer to a theorem named `set.mem_union` as `mem_union`.

The following pattern can be used to show that A is a subset of B :

```
example : A ⊆ B :=
assume x,
assume h : x ∈ A,
show x ∈ B, from sorry
```

And the following pattern can be used to show that A and B are equal:

```
example : A = B :=
eq_of_subset_of_subset
  (assume x,
    assume h : x ∈ A,
    show x ∈ B, from sorry)
  (assume x,
    assume h : x ∈ B,
    show x ∈ A, from sorry)
```

Alternatively, we can use the following pattern:

```
example : A = B :=
ext (assume x, iff.intro
  (assume h : x ∈ A,
    show x ∈ B, from sorry)
  (assume h : x ∈ B,
    show x ∈ A, from sorry))
```

Here, `ext` is short for “extensionality.” In symbolic terms, it is the following fact:

$$\forall x (x \in A \leftrightarrow x \in B) \rightarrow A = B.$$

This reduces proving $A = B$ to proving $\forall x (x \in A \leftrightarrow x \in B)$, which we can do using \forall and \leftrightarrow introduction.

Moreover, Lean supports the following nifty feature: the defining rules for union, intersection and other operations on sets are considered to hold “definitionally.” This means that the expressions $x \in A \cap B$ and $x \in A \wedge x \in B$ mean the same thing to Lean. This is the same for the other constructions on sets; for example $x \in A \setminus B$ and $x \in A \wedge \neg (x \in B)$ mean the same thing to Lean. You can also write $x \notin B$ for $\neg (x \in B)$, where \notin is written using `\notin`. For the other set constructions, the defining equivalences in the last chapter hold definitionally. The following example illustrates these features.

```
example : ∀ x, x ∈ A → x ∈ B → x ∈ A ∩ B :=
assume x,
assume : x ∈ A,
assume : x ∈ B,
show x ∈ A ∩ B, from and.intro <x ∈ A> <x ∈ B>

example : A ⊆ A ∪ B :=
assume x,
assume : x ∈ A,
show x ∈ A ∪ B, from or.inl this

example : ∅ ⊆ A :=
assume x,
assume : x ∈ ∅,
show x ∈ A, from false.elim <x ∈ (∅ : set U)>
```

Remember from [Section 4.5](#) that we can use `assume` without a label, and refer back to hypotheses using French quotes, entered with `\f<` and `\f>`. We have used this feature in the previous example. Without that feature, we could have written the examples above as follows:

```
example : ∀ x, x ∈ A → x ∈ B → x ∈ A ∩ B :=
assume x,
```

(continues on next page)

(continued from previous page)

```

assume h1 : x ∈ A,
assume h2 : x ∈ B,
show x ∈ A ∩ B, from and.intro h1 h2

example : A ⊆ A ∪ B :=
assume x,
assume h : x ∈ A,
show x ∈ A ∪ B, from or.inl h

example : ∅ ⊆ A :=
assume x,
assume h : x ∈ (∅ : set U),
show x ∈ A, from false.elim h

```

Below, and in the chapters that follow, we will begin to use `assume` and `have` command without labels, but you should feel free to adopt whatever style you prefer.

Notice also that in the last example, we had to annotate the empty set by writing `(∅ : set U)` to tell Lean which empty set we mean. Lean can often infer information like this from the context (for example, from the fact that we are trying to show $x \in A$, where A has type `set U`), but in this case, it needs a bit more help.

Alternatively, we can use versions theorems in the Lean library that are designed specifically for use with sets:

```

example : ∀ x, x ∈ A → x ∈ B → x ∈ A ∩ B :=
assume x,
assume : x ∈ A,
assume : x ∈ B,
show x ∈ A ∩ B, from mem_inter <x ∈ A> <x ∈ B>

example : A ⊆ A ∪ B :=
assume x,
assume h : x ∈ A,
show x ∈ A ∪ B, from mem_union_left B h

example : ∅ ⊆ A :=
assume x,
assume : x ∈ ∅,
show x ∈ A, from absurd this (not_mem_empty x)

```

Remember that `absurd` can be used to prove any fact from two contradictory hypotheses $h1 : P$ and $h2 : \neg P$. Here the `not_mem_empty x` is the fact $x \notin \emptyset$. You can see the statements of the theorems using the `#check` command in Lean:

```

#check @mem_inter
#check @mem_of_mem_inter_left
#check @mem_of_mem_inter_right
#check @mem_union_left
#check @mem_union_right
#check @mem_or_mem_of_mem_union
#check @not_mem_empty

```

Here, the `@` symbol in Lean prevents it from trying to fill in implicit arguments automatically, forcing it to

display the full statement of the theorem.

The fact that Lean can identify sets with their logical definitions makes it easy to prove inclusions between sets:

```
example : A \ B ⊆ A :=
assume x,
assume : x ∈ A \ B,
show x ∈ A, from and.left this

example : A \ B ⊆ -B :=
assume x,
assume : x ∈ A \ B,
have x ∉ B, from and.right this,
show x ∈ -B, from this
```

Once again, we can use versions of the theorems designed specifically for sets:

```
example : A \ B ⊆ A :=
assume x,
assume : x ∈ A \ B,
show x ∈ A, from mem_of_mem_diff this

example : A \ B ⊆ -B :=
assume x,
assume : x ∈ A \ B,
have x ∉ B, from not_mem_of_mem_diff this,
show x ∈ -B, from this
```

The fact that Lean has to unfold definitions means that it can be confused at times. For example, in the proof below, if you replace the last line by `sorry`, Lean has trouble figuring out that you want it to unfold the subset symbol:

```
variable {U : Type}
variables A B : set U

example : A ∩ B ⊆ B ∩ A :=
assume x,
assume h : x ∈ A ∩ B,
have h1 : x ∈ A, from and.left h,
have h2 : x ∈ B, from and.right h,
and.intro h2 h1
```

One workaround is to use the `show` command; in general, providing Lean with such additional information is often helpful. Another is to workaround is to give the theorem a name, which prompts Lean to use a slightly different method of processing the proof, fixing the problem as a lucky side effect.

```
example : A ∩ B ⊆ B ∩ A :=
assume x,
assume h : x ∈ A ∩ B,
have h1 : x ∈ A, from and.left h,
have h2 : x ∈ B, from and.right h,
show x ∈ B ∩ A, from sorry
```

(continues on next page)

(continued from previous page)

```

theorem my_example : A ∩ B ⊆ B ∩ A :=
assume x,
assume h : x ∈ A ∩ B,
have h1 : x ∈ A, from and.left h,
have h2 : x ∈ B, from and.right h,
sorry

```

12.2 Some Identities

Here is the proof of the first identity that we proved informally in the previous chapter:

```

example : A ∩ (B ∪ C) = (A ∩ B) ∪ (A ∩ C) :=
eq_of_subset_subset
  (assume x,
    assume h : x ∈ A ∩ (B ∪ C),
    have x ∈ A, from and.left h,
    have x ∈ B ∪ C, from and.right h,
    or.elim (<x ∈ B ∪ C>)
      (assume : x ∈ B,
        have x ∈ A ∩ B, from and.intro <x ∈ A> <x ∈ B>,
        show x ∈ (A ∩ B) ∪ (A ∩ C), from or.inl this)
      (assume : x ∈ C,
        have x ∈ A ∩ C, from and.intro <x ∈ A> <x ∈ C>,
        show x ∈ (A ∩ B) ∪ (A ∩ C), from or.inr this))
  (assume x,
    assume : x ∈ (A ∩ B) ∪ (A ∩ C),
    or.elim this
      (assume h : x ∈ A ∩ B,
        have x ∈ A, from and.left h,
        have x ∈ B, from and.right h,
        have x ∈ B ∪ C, from or.inl this,
        show x ∈ A ∩ (B ∪ C), from and.intro <x ∈ A> this)
      (assume h : x ∈ A ∩ C,
        have x ∈ A, from and.left h,
        have x ∈ C, from and.right h,
        have x ∈ B ∪ C, from or.inr this,
        show x ∈ A ∩ (B ∪ C), from and.intro <x ∈ A> this))

```

Notice that it is considerably longer than the informal proof in the last chapter, because we have spelled out every last detail. Unfortunately, this does not necessarily make it more readable. Keep in mind that you can always write long proofs incrementally, using `sorry`. You can also break up long proofs into smaller pieces:

```

theorem inter_union_subset : A ∩ (B ∪ C) ⊆ (A ∩ B) ∪ (A ∩ C) :=
assume x,
assume h : x ∈ A ∩ (B ∪ C),
have x ∈ A, from and.left h,
have x ∈ B ∪ C, from and.right h,
or.elim (<x ∈ B ∪ C>)
  (assume : x ∈ B,
    have x ∈ A ∩ B, from and.intro <x ∈ A> <x ∈ B>,

```

(continues on next page)

(continued from previous page)

```

    show x ∈ (A ∩ B) ∪ (A ∩ C), from or.inl this)
  (assume : x ∈ C,
    have x ∈ A ∩ C, from and.intro <x ∈ A> <x ∈ C>,
    show x ∈ (A ∩ B) ∪ (A ∩ C), from or.inr this)

theorem inter_union_inter_subset :
  (A ∩ B) ∪ (A ∩ C) ⊆ A ∩ (B ∪ C) :=
  assume x,
  assume : x ∈ (A ∩ B) ∪ (A ∩ C),
  or.elim this
  (assume h : x ∈ A ∩ B,
    have x ∈ A, from and.left h,
    have x ∈ B, from and.right h,
    have x ∈ B ∪ C, from or.inl this,
    show x ∈ A ∩ (B ∪ C), from and.intro <x ∈ A> this)
  (assume h : x ∈ A ∩ C,
    have x ∈ A, from and.left h,
    have x ∈ C, from and.right h,
    have x ∈ B ∪ C, from or.inr this,
    show x ∈ A ∩ (B ∪ C), from and.intro <x ∈ A> this)

example : A ∩ (B ∪ C) = (A ∩ B) ∪ (A ∩ C) :=
  eq_of_subset_of_subset
  (inter_union_subset A B C)
  (inter_union_inter_subset A B C)

```

Notice that the two theorems depend on the variables A , B , and C , which have to be supplied as arguments when they are applied. They also depend on the underlying type, U , but because the variable U was marked implicit, Lean figures it out from the context.

In the last chapter, we showed $(A \cap \overline{B}) \cup B = B$. Here is the corresponding proof in Lean:

```

example : (A ∩ ¬B) ∪ B = A ∪ B :=
  calc
  (A ∩ ¬B) ∪ B = (A ∪ B) ∩ (¬B ∪ B) : by rw union_distrib_right
  ... = (A ∪ B) ∩ univ      : by rw compl_union_self
  ... = A ∪ B              : by rw inter_univ

```

Translated to propositions, the theorem above states that for every pair of elements A and B in a Boolean algebra, $(A \wedge \neg B) \vee B = B$. Lean allows us to do calculations on propositions as though they are elements of a Boolean algebra, with equality replaced by \leftrightarrow .

```

variables A B : Prop

example : (A ∧ ¬ B) ∨ B ↔ A ∨ B :=
  calc
  (A ∧ ¬ B) ∨ B ↔ (A ∨ B) ∧ (¬ B ∨ B) : by rw and_or_distrib_right
  ... ↔ (A ∨ B) ∧ true      : by rw not_or_self
  ... ↔ (A ∨ B)            : by rw and_true

```

12.3 Indexed Families

Remember that if $(A_i)_{i \in I}$ is a family of sets indexed by I , then $\bigcap_{i \in I} A_i$ denotes the intersection of all the sets A_i , and $\bigcup_{i \in I} A_i$ denotes their union. In Lean, we can specify that A is a family of sets by writing $A : I \rightarrow \text{set } U$ where I is a `Type`. In other words, a family of sets is really a function which for each element i of type I returns a set $A\ i$. We can then define the union and intersection as follows:

```
variables {I U : Type}

def Union (A : I → set U) : set U := { x | ∃ i : I, x ∈ A i }
def Inter (A : I → set U) : set U := { x | ∀ i : I, x ∈ A i }

section
variables (x : U) (A : I → set U)

example (h : x ∈ Union A) : ∃ i, x ∈ A i := h
example (h : x ∈ Inter A) : ∀ i, x ∈ A i := h
end
```

The examples show that Lean can unfold the definitions so that $x \in \text{Inter } A$ can be treated as $\forall i, x \in A\ i$ and $x \in \text{Union } A$ can be treated as $\exists i, x \in A\ i$. To refresh your memory as to how to work with the universal and existential quantifiers in Lean, see [Chapters 9](#). We can then define notation for the indexed union and intersection:

```
notation `∪` binders ``, ` r:(scoped f, Union f) := r
notation `∩` binders ``, ` r:(scoped f, Inter f) := r

variables (A : I → set U) (x : U)

example (h : x ∈ ∪ i, A i) : ∃ i, x ∈ A i := h
example (h : x ∈ ∩ i, A i) : ∀ i, x ∈ A i := h
```

You can type \bigcap and \bigcup with `\I` and `\Un`, respectively. As with quantifiers, the notation $\bigcup i, A\ i$ and $\bigcap i, A\ i$ bind the variable i in the expression, and the scope extends as widely as possible. For example, if you write $\bigcap i, A\ i \cup B$, Lean assumes that the i th element of the sequence is $A\ i \cup B$. If you want to restrict the scope more narrowly, use parentheses.

The good news is that Lean's library does define indexed union and intersection, with this notation, and the definitions are made available with `import data.set`. The bad news is that it uses a different definition, so that $x \in \text{Inter } A$ and $x \in \text{Union } A$ are *not* definitionally equal to $\forall i, x \in A\ i$ and $\exists i, x \in A\ i$, as above. The good news is that Lean at least knows that they are equivalent:

```
import data.set
open set

variables {I U : Type}
variables {A B : I → set U}

theorem exists_of_mem_Union {x : U} (h : x ∈ ∪ i, A i) :
  ∃ i, x ∈ A i :=
by simp * at *

theorem mem_Union_of_exists {x : U} (h : ∃ i, x ∈ A i) :
  x ∈ ∪ i, A i :=
```

(continues on next page)

(continued from previous page)

```

by simp * at *

theorem forall_of_mem_Inter {x : U} (h : x ∈ ⋂ i, A i) :
  ∀ i, x ∈ A i :=
by simp * at *

theorem mem_Inter_of_forall {x : U} (h : ∀ i, x ∈ A i) :
  x ∈ ⋂ i, A i :=
by simp * at *
    
```

The command `simp * at *` calls upon Lean's automation to carry out the proofs. Here is an example of how these can be used:

```

example : (⋂ i, A i ∩ B i) = (⋂ i, A i) ∩ (⋂ i, B i) :=
ext $
assume x : U,
iff.intro
(assume h : x ∈ ⋂ i, A i ∩ B i,
  have h1 : ∀ i, x ∈ A i ∩ B i, from forall_of_mem_Inter h,
  have h2 : ∀ i, x ∈ A i, from assume i, and.left (h1 i),
  have h3 : ∀ i, x ∈ B i, from assume i, and.right (h1 i),
  have h4 : x ∈ ⋂ i, A i, from mem_Inter_of_forall h2,
  have h5 : x ∈ ⋂ i, B i, from mem_Inter_of_forall h3,
  and.intro h4 h5)
(assume h : x ∈ (⋂ i, A i) ∩ (⋂ i, B i),
  have h1 : ∀ i, x ∈ A i,
    from forall_of_mem_Inter (and.left h),
  have h2 : ∀ i, x ∈ B i,
    from forall_of_mem_Inter (and.right h),
  have h3 : ∀ i, x ∈ A i ∩ B i,
    from assume i, and.intro (h1 i) (h2 i),
  show x ∈ ⋂ i, A i ∩ B i, from mem_Inter_of_forall h3)
    
```

Even better, we can prove introduction and elimination rules for intersection and union:

```

import data.set
open set

variables {I U : Type}
variables {A : I → set U}

theorem Inter.intro {x : U} (h : ∀ i, x ∈ A i) : x ∈ ⋂ i, A i :=
by simp; assumption

@[elab_simple]
theorem Inter.elim {x : U} (h : x ∈ ⋂ i, A i) (i : I) : x ∈ A i :=
by simp at h; apply h

theorem Union.intro {x : U} (i : I) (h : x ∈ A i) :
  x ∈ ⋃ i, A i :=
by {simp, existsi i, exact h}
    
```

(continues on next page)

(continued from previous page)

```

theorem Union.elim {b : Prop} {x : U}
(h1 : x ∈ ⋃ i, A i) (h2 : ∀ (i : I), x ∈ A i → b) : b :=
by {simp at h1, cases h1 with i h, exact h2 i h}
    
```

Don't worry about what the proofs mean. What is important is how they can be used, which is what the following patterns illustrate:

```

example (x : U) : x ∈ ⋂ i, A i :=
Inter.intro $
assume i,
show x ∈ A i, from sorry

example (x : U) (i : I) (h : x ∈ ⋂ i, A i) : x ∈ A i :=
Inter.elim h i

example (x : U) (i : I) (h : x ∈ A i) : x ∈ ⋃ i, A i :=
Union.intro i h

example (C : Prop) (x : U) (h : x ∈ ⋃ i, A i) : C :=
Union.elim h $
assume i,
assume h : x ∈ A i,
show C, from sorry
    
```

Remember that the dollar sign saves us the trouble of having to put parentheses around the rest of the proof. Notice that with `Inter.intro` and `Inter.elim`, proofs using indexed intersections looks just like proofs using the universal quantifier. Similarly, `Union.intro` and `Union.elim` mirror the introduction and elimination rules for the existential quantifier. The following example provides one direction of an equivalence proved above:

```

variables {I U : Type}
variables (A : I → set U) (B : I → set U) (C : set U)

example : (⋂ i, A i ∩ B i) ⊆ (⋂ i, A i) ∩ (⋂ i, B i) :=
assume x : U,
assume h : x ∈ ⋂ i, A i ∩ B i,
have h1 : x ∈ ⋂ i, A i, from
  Inter.intro $
    assume i : I,
    have h2 : x ∈ A i ∩ B i, from Inter.elim h i,
    show x ∈ A i, from and.left h2,
have h2 : x ∈ ⋂ i, B i, from
  Inter.intro $
    assume i : I,
    have h2 : x ∈ A i ∩ B i, from Inter.elim h i,
    show x ∈ B i, from and.right h2,
show x ∈ (⋂ i, A i) ∩ (⋂ i, B i), from and.intro h1 h2
    
```

You are asked to prove the other direction in the exercises below. Here is an example that shows how to use the introduction and elimination rules for indexed union:

```

variables {I U : Type}
    
```

(continues on next page)

(continued from previous page)

```

variables (A : I → set U) (B : I → set U) (C : set U)

example : (⋃ i, C ∩ A i) ⊆ C ∩ (⋃ i, A i) :=
assume x,
assume h : x ∈ ⋃ i, C ∩ A i,
Union.elim h $
assume i,
assume h1 : x ∈ C ∩ A i,
have h2 : x ∈ C, from and.left h1,
have h3 : x ∈ A i, from and.right h1,
have h4 : x ∈ ⋃ i, A i, from Union.intro i h3,
show x ∈ C ∩ ⋃ i, A i, from and.intro h2 h4
    
```

Once again, we ask you to prove the other direction in the exercises below.

Sometimes we want to work with families $(A_{i,j})_{i \in I, j \in J}$ indexed by two variables. This is also easy to manage in Lean: if we declare $A : I \rightarrow J \rightarrow \text{set } U$, then given $i : I$ and $j : J$, we have that $A \ i \ j : \text{set } U$. (You should interpret the expression $I \rightarrow J \rightarrow \text{set } U$ as $I \rightarrow (J \rightarrow \text{set } U)$, so that $A \ i$ has type $J \rightarrow \text{set } U$, and then $A \ i \ j$ has type $\text{set } U$.) Here is an example of a proof involving a such a doubly-indexed family:

```

section
variables {I J U : Type}
variables (A : I → J → set U)

example : (⋃ i, ⋂ j, A i j) ⊆ (⋂ j, ⋃ i, A i j) :=
assume x,
assume h : x ∈ ⋃ i, ⋂ j, A i j,
Union.elim h $
assume i,
assume h1 : x ∈ ⋂ j, A i j,
show x ∈ ⋂ j, ⋃ i, A i j, from
  Inter.intro $
    assume j,
    have h2 : x ∈ A i j, from Inter.elim h1 j,
    Union.intro i h2
end
    
```

12.4 Power Sets

We can also define the power set in Lean:

```

variable {U : Type}

def powerset (A : set U) : set (set U) := {B : set U | B ⊆ A}

example (A B : set U) (h : B ∈ powerset A) : B ⊆ A :=
h
    
```

As the example shows, $B \in \text{powerset } A$ is then definitionally the same as $B \subseteq A$.

In fact, `powerset` is defined in Lean in exactly this way, and is available to you when you import `data.set` and open `set`. Here is an example of how it is used:

```
#check powerset A

example : A ∈ powerset (A ∪ B) :=
  assume x,
  assume : x ∈ A,
  show x ∈ A ∪ B, from or.inl <x ∈ A>
```

In essence, the example proves $A \subseteq A \cup B$. In the exercises below, we ask you to prove, formally, that for every $A B : \text{set } U$, we have $\text{powerset } A \subseteq \text{powerset } B$

12.5 Exercises

1. Fill in the sorry's.

```
section
  variable U : Type
  variables A B C : set U

  example : ∀ x, x ∈ A ∩ C → x ∈ A ∪ B :=
    sorry

  example : ∀ x, x ∈ -(A ∪ B) → x ∈ -A :=
    sorry
end
```

2. Fill in the sorry.

```
import data.set
open set

section
  variable {U : Type}

  /- defining "disjoint" -/

  def disj (A B : set U) : Prop := ∀ {x}, x ∈ A → x ∈ B → false

  example (A B : set U) (h : ∀ x, ¬ (x ∈ A ∧ x ∈ B)) :
    disj A B :=
  assume x,
  assume h1 : x ∈ A,
  assume h2 : x ∈ B,
  have h3 : x ∈ A ∧ x ∈ B, from and.intro h1 h2,
  show false, from h x h3

  -- notice that we do not have to mention x when applying
  -- h : disj A B
  example (A B : set U) (h1 : disj A B) (x : U)
    (h2 : x ∈ A) (h3 : x ∈ B) :
```

(continues on next page)

(continued from previous page)

```

false :=
h1 h2 h3

-- the same is true of  $\subseteq$ 
example (A B : set U) (x : U) (h : A  $\subseteq$  B) (h1 : x  $\in$  A) :
  x  $\in$  B :=
h h1

example (A B C D : set U) (h1 : disj A B) (h2 : C  $\subseteq$  A)
  (h3 : D  $\subseteq$  B) :
  disj C D :=
sorry
end
    
```

3. Prove the following facts about indexed unions and intersections, using the theorems `Inter.intro`, `Inter.elim`, `Union.intro`, and `Union.elim` listed above.

```

variables {I U : Type}
variables (A : I  $\rightarrow$  set U) (B : I  $\rightarrow$  set U) (C : set U)

example : ( $\bigcap$  i, A i)  $\cap$  ( $\bigcap$  i, B i)  $\subseteq$  ( $\bigcap$  i, A i  $\cap$  B i) :=
sorry

example : C  $\cap$  ( $\bigcup$  i, A i)  $\subseteq$   $\bigcup$  i, C  $\cap$  A i :=
sorry
    
```

4. Prove the following fact about power sets. You can use the theorems `subset.trans` and `subset.refl`

```

variable {U : Type}
variables A B C : set U

-- For this exercise these two facts are useful
example (h1 : A  $\subseteq$  B) (h2 : B  $\subseteq$  C) : A  $\subseteq$  C :=
subset.trans h1 h2

example : A  $\subseteq$  A :=
subset.refl A

example (h : A  $\subseteq$  B) : powerset A  $\subseteq$  powerset B :=
sorry

example (h : powerset A  $\subseteq$  powerset B) : A  $\subseteq$  B :=
sorry
    
```


RELATIONS

In [Chapter 7](#) we discussed the notion of a *relation symbol* in first-order logic, and in [Chapter 10](#) we saw how to interpret such a symbol in a model. In mathematics, we are generally interested in different sorts of relationships between mathematical objects, and so the notion of a relation is ubiquitous. In this chapter, we will consider some common kinds of relations.

In some axiomatic foundations, the notion of a relation is taken to be primitive, but in axiomatic set theory, a relation is taken to be a set of tuples of the corresponding arity. For example, we can take a binary relation on A to be a subset of $A \times A$, where $R(a, b)$ means that $(a, b) \in R$. The foundational definition is generally irrelevant to everyday mathematical practice; what is important is simply that we can write expressions like $R(a, b)$, and that they are true or false, depending on the values of a and b . In mathematics, we often use *infix* notation, writing aRb instead of $R(a, b)$.

13.1 Order Relations

We will start with a class of important binary relations in mathematics, namely, *partial orders*.

Definition. A binary relation \leq on a domain A is a *partial order* if it has the following three properties:

- *reflexivity*: $a \leq a$, for every a in A
 - *transitivity*: if $a \leq b$ and $b \leq c$, then $a \leq c$, for every a, b , and c in A
 - *antisymmetry*: if $a \leq b$ and $b \leq a$ then $a = b$, for every a and b in A
-

Notice the compact way of introducing the symbol \leq in the statement of the definition, and the fact that \leq is written as an infix symbol. Notice also that even though the relation is written with the symbol \leq , it is the only symbol occurring in the definition; mathematical practice favors natural language to describe its properties.

You now know enough, however, to recognize the universal quantifiers that are present in the three clauses. In symbolic logic, we would write them as follows:

- $\forall a (a \leq a)$
- $\forall a, b, c (a \leq b \wedge b \leq c \rightarrow a \leq c)$
- $\forall a, b (a \leq b \wedge b \leq a \rightarrow a = b)$

Here the variables a, b , and c implicitly range over the domain A .

The use of the symbol \leq is meant to be suggestive, and, indeed, the following are all examples of partial orders:

- \leq on the natural numbers
- \leq on the integers
- \leq on the rational numbers
- \leq on the real numbers

But keep in mind that \leq is only a symbol; it can have unexpected interpretations as well. For example, all of the following are also partial orders:

- \geq on the natural numbers
- \geq on the integers
- \geq on the rational numbers
- \geq on the real numbers

These are not fully representative of the class of partial orders, in that they all have an additional property:

Definition. A partial order \leq on a domain A is a *total order* (also called a *linear order*) if it also has the following property:

- for every a and b in A , either $a \leq b$ or $b \leq a$
-

You can check these these are two examples of partial orders that are not total orders:

- the divides relation, $x \mid y$, on the integers
- the subset relation, $x \subseteq y$, on sets of elements of some domain A

On the integers, we also have the strict order relation, $<$, which is not a partial order, since it is not reflexive. It is, rather, an instance of a *strict partial order*:

Definition. A binary relation $<$ on a domain A is a *strict partial order* if it satisfies the following:

- *irreflexivity*: $a \not< a$ for every a in A
- *transitivity*: $a < b$ and $b < c$ implies $a < c$, for every a , b , and c in A

A strict partial order is a *strict total order* (or *strict linear order*) if, in addition, we have the following property:

- *trichotomy*: $a < b$, $a = b$, or $a > b$ for every a and b in A
-

Here, $b \not< a$ means, of course, that it is not the case that $a < b$, and $a > b$ is alternative notation for $b < a$. To distinguish an ordinary partial order from a strict one, an ordinary partial order is sometimes called a *weak partial order*.

Proposition. A strict partial order $<$ on A is *asymmetric*: for every a and b , $a < b$ implies $b \not< a$.

Proof. Suppose $a < b$ and $b < a$. Then, by transitivity, $a < a$, contradicting irreflexivity.

On the integers, there are precise relationships between $<$ and \leq : $x \leq y$ if and only if $x < y$ or $x = y$, and $x < y$ if and only if $x \leq y$ and $x \neq y$. This illustrates a more general phenomenon.

Theorem. Suppose \leq is a partial order on a domain A . Define $a < b$ to mean that $a \leq b$ and $a \neq b$. Then $<$ is a strict partial order. Moreover, if \leq is total, so is $<$.

Theorem. Suppose $<$ is a strict partial order on a domain A . Define $a \leq b$ to mean $a < b$ or $a = b$. Then \leq is a partial order. Moreover, if $<$ is total, so is \leq .

We will prove the first here, and leave the second as an exercise. This proof is a nice illustration of how universal quantification, equality, and propositional reasoning are combined in a mathematical argument.

Proof. Suppose \leq is a partial order on A , and $<$ be defined as in the statement of the theorem. Irreflexivity is immediate, since $a < a$ implies $a \neq a$, which is a contradiction.

To show transitivity, suppose $a < b$ and $b < c$. Then we have $a \leq b$, $b \leq c$, $a \neq b$, and $b \neq c$. By the transitivity of \leq , we have $a \leq c$. To show $a < c$, we only have to show $a \neq c$. So suppose $a = c$. then, from the hypotheses, we have $c < b$ and $b < c$, violating asymmetry. So $a \neq c$, as required.

To establish the last claim in the theorem, suppose \leq is total, and let a and b be any elements of A . We need to show that $a < b$, $a = b$, or $a > b$. If $a = b$, we are done, so we can assume $a \neq b$. Since \leq is total, we have $a \leq b$ or $a \leq b$. Since $a \neq b$, in the first case we have $a < b$, and in the second case, we have $a > b$.

13.2 More on Orderings

Let \leq be a partial order on a domain, A , and let $<$ be the associated strict order, as defined in the last section. It is possible to show that if we go in the other direction, and define \leq' to be the partial order associated to $<$, then \leq and \leq' are the same, which is to say, for every a and b in A , $a \leq b$ if and only if $a \leq' b$. So we can think of every partial order as really being a pair, consisting of a weak partial order and an associated strict one. In other words, we can assume that $x < y$ holds if and only if $x \leq y$ and $x \neq y$, and we can assume $x \leq y$ holds if and only if $x < y$ or $x = y$.

We will henceforth adopt this convention. Given a partial order \leq and the associated strict order $<$, we leave it to you to show that if $x \leq y$ and $y < z$, then $x < z$, and, similarly, if $x < y$ and $y \leq z$, then $x < z$.

Consider the natural numbers with the less-than-or-equal relation. It has a least element, 0. We can express the fact that 0 is the least element in at least two ways:

- 0 is less than or equal to every natural number.
- There is no natural number that is less than 0.

In symbolic logic, we could formalize these statements as follows:

- $\forall x (0 \leq x)$
- $\forall x (x \not< 0)$

Using the existential quantifier, we could render the second statement more faithfully as follows:

- $\neg \exists x (x < 0)$

Notice that this more faithful statement is equivalent to the original, using deMorgan’s laws for quantifiers.

Are the two statements above equivalent? Say an element y is *minimum* for a partial order if it is less than or equal to any other element, that is, if it takes the place of 0 in the first statement. Say that an element y is *minimal* for a partial order if no element is less than it, that is, if it takes the place of 0 in the second statement. Two facts are immediate.

Theorem. Any minimum element is minimal.

Proof. Suppose x is minimum for \leq . We need to show that x is minimal, that is, for every y , it is not the case that $y < x$. Suppose $y < x$. Since x is minimum, we have $x \leq y$. From $y < x$ and $x \leq y$, we have $y < y$, contradicting the irreflexivity of $<$.

Theorem. If a partial order \leq has a minimum element, it is unique.

Proof. Suppose x_1 and x_2 are both minimum. Then $x_1 \leq x_2$ and $x_2 \leq x_1$. By antisymmetry, $x_1 = x_2$.

Notice that we have interpreted the second theorem as the statement that if x_1 and x_2 are both minimum, then $x_1 = x_2$. Indeed, this is exactly what we mean when we say that something is “unique.” When a partial order has a minimum element x , uniqueness is what justifies calling x *the* minimum element. Such an x is also called the *least* element or the *smallest* element, and the terms are generally interchangeable.

The converse to the second theorem – that is, the statement that every minimal element is minimum – is false. As an example, consider the nonempty subsets of the set $\{1, 2\}$ with the subset relation. In other words, consider the collection of sets $\{1\}$, $\{2\}$, and $\{1, 2\}$, where $\{1\} \subseteq \{1, 2\}$, $\{2\} \subseteq \{1, 2\}$, and, of course, every element is a subset of itself. Then you can check that $\{1\}$ and $\{2\}$ are each minimal, but neither is minimum. (One can also exhibit such a partial order by drawing a diagram, with dots labeled a , b , c , etc., and upwards edges between elements to indicate that one is less than or equal to the other.)

Notice that the statement “a minimal element of a partial order is not necessarily minimum” makes an “existential” assertion: it says that there is a partial order \leq , and an element x of the domain, such that x is minimal but not minimum. For a fixed partial order \leq , we can express the assertion that such an x exists as follows:

$$\exists x (\forall y (y \not< x) \wedge \forall y (x \leq y)).$$

The assertion that there exists a domain A , and a partial order \leq on that domain A , is more dramatic: it is a “higher order” existential assertion. But symbolic logic provides us with the means to make assertions like these as well, as we will see later on.

We can consider other properties of orders. An order is said to be *dense* if between any two distinct elements, there is another element. More precisely, an order is dense if, whenever $x < y$, there is an element z satisfying $x < z$ and $z < y$. For example, the rational numbers are dense with the usual \leq ordering, but not the integers. Saying that an order is dense is another example of an implicit use of existential quantification.

13.3 Equivalence Relations and Equality

In ordinary mathematical language, an *equivalence relation* is defined as follows.

Definition. A binary relation \equiv on some domain A is said to be an *equivalence relation* if it is reflexive, symmetric, and transitive. In other words, \equiv is an equivalent relation if it satisfies these three properties:

- *reflexivity*: $a \equiv a$, for every a in A

- *symmetry*: if $a \equiv b$, then $b \equiv a$, for every a and b in A
- *transitivity*: if $a \equiv b$ and $b \equiv c$, then $a \equiv c$, for every a, b , and c in A

We leave it to you to think about how you could write these statements in first-order logic. (Note the similarity to the rules for a partial order.) We will also leave you with an exercise: by a careful choice of how to instantiate the quantifiers, you can actually prove the three properties above from the following two:

- $\forall a (a \equiv a)$
- $\forall a, b, c (a \equiv b \wedge c \equiv b \rightarrow a \equiv c)$

Try to verify this using natural deduction or Lean.

These three properties alone are not strong enough to characterize equality. You should check that the following informal examples are all instances of equivalence relations:

- the relation on days on the calendar, given by “ x and y fall on the same day of the week”
- the relation on people currently alive on the planet, given by “ x and y have the same age”
- the relation on people currently alive on the planet, given by “ x and y have the same birthday”
- the relation on cities in the United States, given by “ x and y are in the same state”

Here are two common mathematical examples:

- the relation on lines in a plane, given by “ x and y are parallel”
- for any fixed natural number $m \geq 0$, the relation on natural numbers, given by “ x is congruent to y modulo m ” (see [Chapter 19](#))

Here, we say that x is congruent to y modulo m if they leave the same remainder when divided by m . Soon, you will be able to prove rigorously that this is equivalent to saying that $x - y$ is divisible by m .

Consider the equivalence relation on citizens of the United States, given by “ x and y have the same age.” There are some properties that respect that equivalence. For example, suppose I tell you that John and Susan have the same age, and I also tell you that John is old enough to vote. Then you can rightly infer that Susan is old enough to vote. On the other hand, if I tell you nothing more than the facts that John and Susan have the same age and John lives in South Dakota, you cannot infer that Susan lives in South Dakota. This little example illustrates what is special about the *equality* relation: if two things are equal, then they have exactly the same properties.

Let A be a set and let \equiv be an equivalence relation on A . There is an important mathematical construction known as forming the *quotient* of A under the equivalence relation. For every element a in A , let $[a]$ be the set of elements $\{c \mid c \equiv a\}$, that is, the set of elements of A that are equivalent to a . We call $[a]$ the equivalence class of A . The set A/\equiv , the *quotient of A by \equiv* , is defined to be the set $\{[a] : a \in A\}$, that is, the set of all the equivalence classes of elements in A . The exercises below ask you to show that if $[a]$ and $[b]$ are elements of such a quotient, then $[a] = [b]$ if and only if $a \equiv b$.

The motivation is as follows. Equivalence tries to capture a weak notion of equality: if two elements of A are equivalent, they are not necessarily the same, but they are similar in some way. Equivalence classes collect similar objects together, essentially glomming them into new objects. Thus A/\equiv is a version of the set A where similar elements have been compressed into a single element. For example, given the equivalence relation \equiv of congruence modulo 5 on the integers, \mathbb{N}/\equiv is the set $\{[0], [1], [2], [3], [4]\}$, where, for example, $[0]$ is the set of all multiples of 5.

13.4 Exercises

- Suppose $<$ is a strict partial order on a domain A , and define $a \leq b$ to mean that $a < b$ or $a = b$.
 - Show that \leq is a partial order.
 - Show that if $<$ is moreover a strict total order, then \leq is a total order.

(Above we proved the analogous theorem going in the other direction.)

- Suppose $<$ is a strict partial order on a domain A . (In other words, it is transitive and asymmetric.) Suppose that \leq is defined so that $a \leq b$ if and only if $a < b$ or $a = b$. We saw in class that \leq is a partial order on a domain A , i.e.~it is reflexive, transitive, and antisymmetric.

Prove that for every a and b in A , we have $a < b$ iff $a \leq b$ and $a \neq b$, using the facts above.

- An *ordered graph* is a collection of vertices (points), along with a collection of arrows between vertices. For each pair of vertices, there is at most one arrow between them: in other words, every pair of vertices is either unconnected, or one vertex is “directed” toward the other. Note that it is possible to have an arrow from a vertex to itself.

Define a relation \leq on the set of vertices, such that for two vertices a and b , $a \leq b$ means that there is an arrow from a pointing to b .

On an arbitrary graph, is \leq a partial order, a strict partial order, a total order, a strict total order, or none of the above? If possible, give examples of graphs where \leq fails to have these properties.

- Let \equiv be an equivalence relation on a set A . For every element a in A , let $[a]$ be the equivalence class of a : that is, the set of elements $\{c \mid c \equiv a\}$. Show that for every a and b , $[a] = [b]$ if and only if $a \equiv b$.

(Hints and notes:

- Remember that since you are proving an “if and only if” statement, there are two directions to prove.
 - Since that $[a]$ and $[b]$ are sets, $[a] = [b]$ means that for every element c , c is in $[a]$ if and only if c is in $[b]$.
 - By definition, an element c is in $[a]$ if and only if $c \equiv a$. In particular, a is in $[a]$.)
- Let the relation \sim on the natural numbers \mathbb{N} be defined as follows: if n is even, then $n \sim n + 1$, and if n is odd, then $n \sim n - 1$. Furthermore, for every n , $n \sim n$. Show that \sim is an equivalence relation. What is the equivalence class of the number 5? Describe the set of equivalence classes $\{[n] \mid n \in \mathbb{N}\}$.
 - Show that the relation on lines in the plane, given by “ l_1 and l_2 are parallel,” is an equivalence relation. What is the equivalence class of the x-axis? Describe the set of equivalence classes $\{[l] \mid l \text{ is a line in the plane}\}$.
 - A binary relation \leq on a domain A is said to be a *preorder* if it is reflexive and transitive. This is weaker than saying it is a partial order; we have removed the requirement that the relation is asymmetric. An example is the ordering on people currently alive on the planet defined by setting $x \leq y$ if and only if x ’s birth date is earlier than y ’s. Asymmetry fails, because different people can be born on the same day. But, prove that the following theorem holds:

Theorem. Let \leq be a preorder on a domain A . Define the relation \equiv , where $x \equiv y$ holds if and only if $x \leq y$ and $y \leq x$. Then \equiv is an equivalence relation on A .

RELATIONS IN LEAN

In the last chapter, we noted that set theorists think of a binary relation R on a set A as a set of ordered pairs, so that $R(a, b)$ really means $(a, b) \in R$. An alternative is to think of R as a function which, when applied to a and B , returns the proposition that $R(a, b)$ holds. This is the viewpoint adopted by Lean: a binary relation on a type A is a function $A \rightarrow A \rightarrow \text{Prop}$. Remember that the arrows associate to the right, so $A \rightarrow A \rightarrow \text{Prop}$ really means $A \rightarrow (A \rightarrow \text{Prop})$. So, given $a : A$, $R\ a$ is a predicate (the property of being related to A), and given $a\ b : A$, $R\ a\ b$ is a proposition.

14.1 Order Relations

With first-order logic, we can say what it means for a relation to be reflexive, symmetric, transitive, anti-symmetric, and so on:

```
namespace hidden

variable {A : Type}

def reflexive (R : A → A → Prop) : Prop :=
  ∀ x, R x x

def symmetric (R : A → A → Prop) : Prop :=
  ∀ x y, R x y → R y x

def transitive (R : A → A → Prop) : Prop :=
  ∀ x y z, R x y → R y z → R x z

def anti_symmetric (R : A → A → Prop) : Prop :=
  ∀ x y, R x y → R y x → x = y

end hidden
```

We can then use the notions freely. Notice that Lean will unfold the definitions when necessary, for example, treating `reflexive R` as $\forall x, R\ x\ x$.

```
variable R : A → A → Prop

example (h : reflexive R) (x : A) : R x x := h x

example (h : symmetric R) (x y : A) (h1 : R x y) : R y x :=
  h x y h1
```

(continues on next page)

(continued from previous page)

```
example (h : transitive R) (x y z : A) (h1 : R x y) (h2 : R y z) :
  R x z :=
h x y z h1 h2

example (h : anti_symmetric R) (x y : A) (h1 : R x y)
  (h2 : R y x) :
  x = y :=
h x y h1 h2
```

In the command `variable {A : Type}`, we put curly braces around `A` to indicate that it is an *implicit* argument, which is to say, you do not have to write it explicitly; Lean can infer it from the argument `R`. That is why we can write `reflexive R` rather than `reflexive A R`: Lean knows that `R` is a binary relation on `A`, so it can infer that we mean reflexivity for binary relations on `A`.

Given `h : transitive R`, `h1 : R x y`, and `h2 : R y z`, it is annoying to have to write `h x y z h1 h2` to prove `R x z`. After all, Lean should be able to infer that we are talking about transitivity at `x`, `y`, and `z`, from the fact that `h1` is `R x y` and `h2` is `R y z`. Indeed, we can replace that information by underscores:

```
variable R : A → A → Prop

example (h : transitive R) (x y z : A) (h1 : R x y)
  (h2 : R y z) :
  R x z :=
h _ _ _ h1 h2
```

But typing underscores is annoying, too. The best solution is to declare the arguments `x y z` to a transitivity hypothesis to be implicit as well:

```
variable {A : Type}

variable R : A → A → Prop

example (h : transitive R) (x y z : A) (h1 : R x y) (h2 : R y z) :
  R x z :=
h h1 h2
```

In fact, the notions `reflexive`, `symmetric`, `transitive`, and `anti_symmetric` are defined in Lean’s core library in exactly this way, so we are free to use them without defining them. That is why we put our temporary definitions of in a namespace `hidden`; that means that the full name of our version of `reflexive` is `hidden.reflexive`, which, therefore, doesn’t conflict with the one defined in the library.

In [Section 13.1](#) we showed that a strict partial order—that is, a binary relation that is transitive and irreflexive—is also asymmetric. Here is a proof of that fact in Lean.

```
variable A : Type
variable R : A → A → Prop

example (h1 : irreflexive R) (h2 : transitive R) :
  ∀ x y, R x y → ¬ R y x :=
assume x y,
assume h3 : R x y,
assume h4 : R y x,
```

(continues on next page)

(continued from previous page)

```

have h5 : R x x, from h2 h3 h4,
have h6 : ¬ R x x, from h1 x,
show false, from h6 h5
    
```

In mathematics, it is common to use infix notation and a symbol like \leq to denote a partial order. Lean supports this practice:

```

section
parameter A : Type
parameter R : A → A → Prop

local infix ≤ := R

example (h1 : irreflexive R) (h2 : transitive R) :
  ∀ x y, x ≤ y → ¬ y ≤ x :=
  assume x y,
  assume h3 : x ≤ y,
  assume h4 : y ≤ x,
  have h5 : x ≤ x, from h2 h3 h4,
  have h6 : ¬ x ≤ x, from h1 x,
  show false, from h6 h5

end
    
```

The `parameter` and `parameters` commands are similar to the `variable` and `variables` commands, except that parameters are fixed within a section. In other words, if you prove a theorem about R in the section above, you cannot apply that theorem to another relation, S , without closing the section. Since the parameter R is fixed, Lean allows us to define notation for R to be used locally in the section.

In the example below, having fixed a partial order, R , we define the corresponding strict partial order and prove that it is, indeed, a strict order.

```

section
parameters {A : Type} (R : A → A → Prop)
parameter (reflR : reflexive R)
parameter (transR : transitive R)
parameter (antisymmR : ∀ {a b : A}, R a b → R b a → a = b)

local infix ≤ := R

definition R' (a b : A) : Prop := a ≤ b ∧ a ≠ b

local infix < := R'

theorem irreflR (a : A) : ¬ a < a :=
  assume : a < a,
  have a ≠ a, from and.right this,
  have a = a, from rfl,
  show false, from <a ≠ a> <a = a>

theorem transR {a b c : A} (h1 : a < b) (h2 : b < c) : a < c :=
  have a ≤ b, from and.left h1,
  have a ≠ b, from and.right h1,
    
```

(continues on next page)

(continued from previous page)

```

have b ≤ c, from and.left h₂,
have b ≠ c, from and.right h₂,
have a ≤ c, from transR <a ≤ b> <b ≤ c>,
have a ≠ c, from
  assume : a = c,
  have c ≤ b, from eq.subst <a = c> <a ≤ b>,
  have b = c, from antisymmR <b ≤ c> <c ≤ b>,
  show false, from <b ≠ c> <b = c>,
show a < c, from and.intro <a ≤ c> <a ≠ c>
end
    
```

Notice that we have used suggestive names `reflR`, `transR`, `antisymmR` instead of `h1`, `h2`, `h3` to help remember which hypothesis is which. The proof also uses anonymous `have` and `assume`, referring back to them with the French quotes, `\f< anf \f>`. Remember also that `eq.subst <a = c> <a ≤ b>` is a proof of the fact that amounts for substituting `c` for `a` in `a ≤ b`. You can also use the equivalent notation `<a = c> ► <a ≤ b>`, where the triangle is written `\t`.

In Section [Section 13.1](#), we also noted that you can define a (weak) partial order from a strict one. We ask you to do this formally in the exercises below.

Here is one more example. Suppose `R` is a binary relation on a type `A`, and we define `S x y` to mean that both `R x y` and `R y x` holds. Below we show that the resulting relation is reflexive and symmetric.

```

section
parameter A : Type
parameter R : A → A → Prop

variable h1 : transitive R
variable h2 : reflexive R

def S (x y : A) := R x y ∧ R y x

example : reflexive S :=
  assume x,
  have R x x, from h2 x,
  show S x x, from and.intro this this

example : symmetric S :=
  assume x y,
  assume h : S x y,
  have h1 : R x y, from h.left,
  have h2 : R y x, from h.right,
  show S y x, from ⟨h.right, h.left⟩

end
    
```

In the exercises below, we ask you to show that `S` is transitive as well.

In the first example, we use the anonymous `assume` and `have`, and then refer back to the `have` with the keyword `this`. In the second example, we abbreviate `and.left h` and `and.right h` as `h.left` and `h.right`, respectively. We also abbreviate `and.intro h.right h.left` with an anonymous constructor, writing `⟨h.right, h.left⟩`. Lean figures out that we are trying to prove a conjunction, and figures out that `and.intro` is the relevant introduction principle. You can type the corner brackets with `\<` and `\>`, respectively.

14.2 Orderings on Numbers

Conveniently, Lean has the normal orderings on the natural numbers, integers, and so on defined already.

```
open nat
variables n m : ℕ

#check 0 ≤ n
#check n < n + 1

example : 0 ≤ n := zero_le n
example : n < n + 1 := lt_succ_self n

example (h : n + 1 ≤ m) : n < m + 1 :=
have h1 : n < n + 1, from lt_succ_self n,
have h2 : n < m, from lt_of_lt_of_le h1 h,
have h3 : m < m + 1, from lt_succ_self m,
show n < m + 1, from lt.trans h2 h3
```

There are many theorems in Lean that are useful for proving facts about inequality relations. We list some common ones here.

```
variables (A : Type) [partial_order A]
variables a b c : A

#check (le_trans : a ≤ b → b ≤ c → a ≤ c)
#check (lt_trans : a < b → b < c → a < c)
#check (lt_of_lt_of_le : a < b → b ≤ c → a < c)
#check (lt_of_le_of_lt : a ≤ b → b < c → a < c)
#check (le_of_lt : a < b → a ≤ b)
```

Here the declaration at the top says that A has the structure of a partial order. There are also properties that are specific to some domains, like the natural numbers:

```
variable n : ℕ

#check (nat.zero_le : ∀ n : ℕ, 0 ≤ n)
#check (nat.lt_succ_self : ∀ n : ℕ, n < n + 1)
#check (nat.le_succ : ∀ n : ℕ, n ≤ n + 1)
```

14.3 Equivalence Relations

In Section 13.3 we saw that an *equivalence relation* is a binary relation on some domain A that is reflexive, symmetric, and transitive. We will see such relations in Lean in a moment, but first let's define another kind of relation called a *preorder*, which is a binary relation that is reflexive and transitive.

```
namespace hidden

variable {A : Type}

def preorder (R : A → A → Prop) : Prop :=
```

(continues on next page)

(continued from previous page)

```

reflexive R ∧ transitive R

end hidden

```

Lean’s library provides a different formulation of preorders, so, in order to use the same name, we have to put it in the `hidden` namespace. Lean’s library defines other properties of relations, such as these:

```

namespace hidden

variables {A : Type} (R : A → A → Prop)

def equivalence := reflexive R ∧ symmetric R ∧ transitive R

def total := ∀ x y, R x y ∨ R y x

def irreflexive := ∀ x, ¬ R x x

def anti_symmetric := ∀ {x y}, R x y → R y x → x = y

end hidden

```

You can ask Lean to print their definitions:

```

#print equivalence
#print total
#print irreflexive
#print anti_symmetric

```

Building on our previous definition of a preorder, we can describe a partial order as an antisymmetric preorder, and show that an equivalence relation as a symmetric preorder.

```

namespace hidden

variable {A : Type}

def preorder (R : A → A → Prop) : Prop :=
  reflexive R ∧ transitive R

def partial_order (R : A → A → Prop) : Prop :=
  preorder R ∧ anti_symmetric R

example (R : A → A → Prop):
  equivalence R ↔ preorder R ∧ symmetric R :=
  iff.intro
    (assume h1 : equivalence R,
     have h2 : reflexive R, from and.left h1,
     have h3 : symmetric R, from and.left (and.right h1),
     have h4 : transitive R, from and.right (and.right h1),
     show preorder R ∧ symmetric R,
     from and.intro (and.intro h2 h4) h3)
    (assume h1 : preorder R ∧ symmetric R,
     have h2 : preorder R, from and.left h1,

```

(continues on next page)

(continued from previous page)

```

show equivalence R,
  from and.intro (and.left h2)
    (and.intro (and.right h1) (and.right h2)))
end hidden

```

In Section 13.3 we claimed that there is yet another way to define an equivalence relation, namely, as a binary relation satisfying the following two properties:

- $\forall a (a \equiv a)$
- $\forall a, b, c (a \equiv b \wedge c \equiv b \rightarrow a \equiv c)$

Let's prove this in Lean. Remember that the `parameters` and `local infix` commands serve to fix a relation R and introduce the notation \approx to denote it. (You can type \approx as `\~.`) In the assumptions `reflexive` (\approx) and `symmetric` (\approx), the notation (\approx) denotes R .

```

section
parameters {A : Type} (R : A → A → Prop)
local infix ≈ := R

variable (h1 : reflexive (≈))
variable (h2 : ∀ {a b c}, a ≈ b ∧ c ≈ b → a ≈ c)

example : symmetric (≈) :=
assume a b (h : a ≈ b),
have b ≈ b ∧ a ≈ b, from and.intro (h1 b) h,
show b ≈ a, from h2 this

example : transitive (≈) :=
assume a b c (h3 : a ≈ b) (h4 : b ≈ c),
have c ≈ b, from h2 (and.intro (h1 c) h4),
have a ≈ b ∧ c ≈ b, from and.intro h3 this,
show a ≈ c, from h2 this

end

```

14.4 Exercises

1. Replace the `sorry` commands in the following proofs to show that we can create a partial order R' out of a strict partial order R .

```

section
parameters {A : Type} {R : A → A → Prop}
parameter (irreflR : irreflexive R)
parameter (transR : transitive R)

local infix < := R

def R' (a b : A) : Prop := R a b ∨ a = b
local infix ≤ := R'

```

(continues on next page)

(continued from previous page)

```

theorem reflR' (a : A) : a ≤ a := sorry

theorem transR' {a b c : A} (h1 : a ≤ b) (h2 : b ≤ c) :
  a ≤ c :=
sorry

theorem antisymmR' {a b : A} (h1 : a ≤ b) (h2 : b ≤ a) :
  a = b :=
sorry

end

```

2. Replace the sorry by a proof.

```

section
parameters {A : Type} {R : A → A → Prop}
parameter (reflR : reflexive R)
parameter (transR : transitive R)

def S (a b : A) : Prop := R a b ∧ R b a

example : transitive S :=
sorry

end

```

3. Only one of the following two theorems is provable. Figure out which one is true, and replace the sorry command with a complete proof.

```

section
parameters {A : Type} {a b c : A} {R : A → A → Prop}
parameter (Rab : R a b)
parameter (Rbc : R b c)
parameter (nRac : ¬ R a c)

-- Prove one of the following two theorems:

theorem R_is_strict_partial_order :
  irreflexive R ∧ transitive R :=
sorry

theorem R_is_not_strict_partial_order :
  ¬(irreflexive R ∧ transitive R) :=
sorry

end

```

4. Complete the following proof.

```

open nat

example : 1 ≤ 4 :=
sorry

```

FUNCTIONS

In the late nineteenth century, developments in a number of branches of mathematics pushed towards a uniform treatment of sets, functions, and relations. We have already considered sets and relations. In this chapter, we consider functions and their properties.

A function, f , is ordinarily understood as a mapping from a domain X to another domain Y . In set-theoretic foundations, X and Y are arbitrary sets. We have seen that in a type-based system like Lean, it is natural to distinguish between types and subsets of a type. In other words, we can consider a type X of elements, and a set A of elements of that type. Thus, in the type-theoretic formulation, it is natural to consider functions between types X and Y , and consider their behavior with respect to subsets of X and Y .

In everyday mathematics, however, set-theoretic language is common, and most mathematicians think of a function as a map between sets. When discussing functions from a mathematical standpoint, therefore, we will also adopt this language, and later switch to the type-theoretic representation when we talk about formalization in Lean.

15.1 The Function Concept

If X and Y are any sets, we write $f : X \rightarrow Y$ to express the fact that f is a function from X to Y . This means that f assigns a value $f(x)$ in Y to every element x of X . The set X is called the *domain* of f , and the set Y is called the *codomain*. (Some authors use the word “range” for the codomain, but today it is more common to use the word “range” for what we call the *image* of A below. We will avoid the ambiguity by avoiding the word range altogether.)

The simplest way to define a function is to give its value at every x with an explicit expression. For example, we can write any of the following:

- Let $f : \mathbb{N} \rightarrow \mathbb{N}$ be the function defined by $f(n) = n + 1$.
- Let $g : \mathbb{R} \rightarrow \mathbb{R}$ be the function defined by $g(x) = x^2$.
- Let $h : \mathbb{N} \rightarrow \mathbb{N}$ be the function defined by $h(n) = n^2$.
- Let $k : \mathbb{N} \rightarrow \{0, 1\}$ be the function defined by

$$k(n) = \begin{cases} 0 & \text{if } n \text{ is even} \\ 1 & \text{if } n \text{ is odd.} \end{cases}$$

The ability to define functions using an explicit expression raises the foundational question as to what counts as legitimate “expression.” For the moment, let us set that question aside, and simply note that modern mathematics is comfortable with all kinds of exotic definitions. For example, we can define a function $f : \mathbb{R} \rightarrow \{0, 1\}$ by

$$f(x) = \begin{cases} 0 & \text{if } x \text{ is rational} \\ 1 & \text{if } x \text{ is irrational.} \end{cases}$$

This is at odds with a view of functions as objects that are computable in some sense. It is not at all clear what it means to be presented with a real number as input, let alone whether it is possible to determine, algorithmically, whether such a number is rational or not. We will return to such issues in a later chapter.

Notice that the choice of the variables x and n in the definitions above are arbitrary. They are bound variables in that the functions being defined do not depend on x or n . The values remain the same under renaming, just as the truth values of “for every x , $P(x)$ ” and “for every y , $P(y)$ ” are the same. Given an expression $e(x)$ that depends on the variable x , logicians often use the notation $\lambda x e(x)$ to denote the function that maps x to $e(x)$. This is called “lambda notation,” for the obvious reason, and it is often quite handy. Instead of saying “let f be the function defined by $f(x) = x + 1$,” we can say “let $f = \lambda x(x + 1)$.” This is *not* common mathematical notation, and it is best to avoid it unless you are talking to logicians or computer scientists. We will see, however, that lambda notation is built in to Lean.

For any set X , we can define a function $i_X(x)$ by the equation $i_X(x) = x$. This function is called the *identity function*. More interestingly, let $f : X \rightarrow Y$ and $g : Y \rightarrow Z$. We can define a new function $k : X \rightarrow Z$ by $k(x) = g(f(x))$. The function k is called *the composition of f and g* or *f composed with g* and it is written $g \circ f$. The order is somewhat confusing; you just have to keep in mind that to evaluate the expression $g(f(x))$ you first evaluate f on input x , and then evaluate g .

We think of two functions $f, g : X \rightarrow Y$ as being equal, or the same function, when for they have the same values on every input; in other words, for every x in X , $f(x) = g(x)$. For example, if $f, g : \mathbb{R} \rightarrow \mathbb{R}$ are defined by $f(x) = x + 1$ and $g(x) = 1 + x$, then $f = g$. Notice that the statement that two functions are equal is a universal statement (that is, for the form “for every x , ...”).

Proposition. For every $f : X \rightarrow Y$, $f \circ i_X = f$ and $i_Y \circ f = f$.

Proof. Let x be any element of X . Then $(f \circ i_X)(x) = f(i_X(x)) = f(x)$, and $(i_Y \circ f)(x) = i_Y(f(x)) = x$.

Suppose $f : X \rightarrow Y$ and $g : Y \rightarrow X$ satisfy $g \circ f = i_X$. Remember that this means that $g(f(x)) = x$ for every x in X . In that case, g is said to be a *left inverse* to f , and f is said to be a *right inverse* to g . Here are some examples:

- Define $f, g : \mathbb{R} \rightarrow \mathbb{R}$ by $f(x) = x + 1$ and $g(x) = x - 1$. Then g is both a left and a right inverse to f , and vice-versa.
- Write $\mathbb{R}^{\geq 0}$ to denote the nonnegative reals. Define $f : \mathbb{R} \rightarrow \mathbb{R}^{\geq 0}$ by $f(x) = x^2$, and define $g : \mathbb{R}^{\geq 0} \rightarrow \mathbb{R}$ by $g(x) = \sqrt{x}$. Then $f(g(x)) = (\sqrt{x})^2 = x$ for every x in the domain of g , so f is a left inverse to g , and g is a right inverse to f . On the other hand, $g(f(x)) = \sqrt{x^2} = |x|$, which is not the same as x when x is negative. So g is not a left inverse to f , and f is not a right inverse to g .

The following fact is not at all obvious, even though the proof is short:

Proposition. Suppose $f : X \rightarrow Y$ has a left inverse, h , and a right inverse, k . Then $h = k$.

Proof. Let y be any element in Y . The idea is to compute $h(f(k(y)))$ in two different ways. Since h is a left inverse to f , we have $h(f(k(y))) = k(y)$. On the other hand, since k is a right inverse to f , $f(k(y)) = y$, and so $h(f(k(y))) = h(y)$. So $k(y) = h(y)$.

If g is both a right and left inverse to f , we say that g is simply the inverse of f . A function f may have more than one left or right inverse (we leave it to you to cook up examples), but it can have at most one inverse.

Proposition. Suppose $g_1, g_2 : Y \rightarrow X$ are both inverses to f . Then $g_1 = g_2$.

Proof. This follows from the previous proposition, since (say) g_1 is a left inverse to f , and g_2 is a right inverse.

When f has an inverse, g , this justifies calling g *the* inverse to f , and writing f^{-1} to denote g . Notice that if f^{-1} is an inverse to f , then f is an inverse to f^{-1} . So if f has an inverse, then so does f^{-1} , and $(f^{-1})^{-1} = f$. For any set A , clearly we have $i_X^{-1} = i_X$.

Proposition. Suppose $f : X \rightarrow Y$ and $g : Y \rightarrow Z$. If $h : Y \rightarrow X$ is a left inverse to f and $k : Z \rightarrow Y$ is a left inverse to g , then $h \circ k$ is a left inverse to $g \circ f$.

Proof. For every x in X ,

$$(h \circ k) \circ (g \circ f)(x) = h(k(g(f(x)))) = h(f(x)) = x.$$

Corollary. The previous proposition holds with “left” replaced by “right.”

Proof. Switch the role of f with h and g with k in the previous proposition.

Corollary. If $f : X \rightarrow Y$ and $g : Y \rightarrow Z$ both have inverses, then $(f \circ g)^{-1} = g^{-1} \circ f^{-1}$.

15.2 Injective, Surjective, and Bijective Functions

A function $f : X \rightarrow Y$ is said to be *injective*, or an *injection*, or *one-one*, if given any x_1 and x_2 in A , if $f(x_1) = f(x_2)$, then $x_1 = x_2$. Notice that the conclusion is equivalent to its contrapositive: if $x_1 \neq x_2$, then $f(x_1) \neq f(x_2)$. So f is injective if it maps distinct element of X to distinct elements of Y .

A function $f : X \rightarrow Y$ is said to be *surjective*, or a *surjection*, or *onto*, if for every element y of Y , there is an x in X such that $f(x) = y$. In other words, f is surjective if every element in the codomain is the value of f at some element in the domain.

A function $f : X \rightarrow Y$ is said to be *bijective*, or a *bijection*, or a *one-to-one correspondence*, if it is both injective and surjective. Intuitively, if there is a bijection between X and Y , then X and Y have the same size, since f makes each element of X correspond to exactly one element of Y and vice-versa. For example, it makes sense to interpret the statement that there were four Beatles as the statement that there is a bijection between the set $\{1, 2, 3, 4\}$ and the set $\{\text{John, Paul, George, Ringo}\}$. If we claimed that there were *five* Beatles, as evidenced by the function f which assigns 1 to John, 2 to Paul, 3 to George, 4 to Ringo, and 5 to John, you should object that we double-counted John—that is, f is not injective. If we claimed there were only three Beatles, as evidenced by the function f which assigns 1 to John, 2 to Paul, and 3 to George, you should object that we left out poor Ringo—that is, f is not surjective.

The next two propositions show that these notions can be cast in terms of the existence of inverses.

Proposition. Let $f : X \rightarrow Y$.

- If f has a left inverse, then f is injective.
- If f has a right inverse, then f is surjective.
- If f has an inverse, then it is f bijective.

Proof. For the first claim, suppose f has a left inverse g , and suppose $f(x_1) = f(x_2)$. Then $g(f(x_1)) = g(f(x_2))$, and so $x_1 = x_2$.

For the second claim, suppose f has a right inverse h . Let y be any element of Y , and let $x = g(y)$. Then $f(x) = f(g(y)) = y$.

The third claim follows from the first two.

The following proposition is more interesting, because it requires us to define new functions, given hypotheses on f .

Proposition. Let $f : X \rightarrow Y$.

- If X is nonempty and f is injective, then f has a left inverse.
- If f is surjective, then f has a right inverse.
- If f is bijective, then it has an inverse.

Proof. For the first claim, let \hat{x} be any element of X , and suppose f is injective. Define $g : Y \rightarrow X$ by setting $g(y)$ equal to any x such that $f(x) = y$, if there is one, and \hat{x} otherwise. Now, suppose $g(f(x)) = x'$. By the definition of g , x' has to have the property that $f(x) = f(x')$. Since f is injective, $x = x'$, so $g(f(x)) = x$.

For the second claim, because f is surjective, we know that for every y in Y there is any x such that $f(x) = y$. Define $h : Y \rightarrow X$ by again setting $h(y)$ equal to any such x . (In contrast to the previous paragraph, here we know that such an x exists, but it might not be unique.) Then, by the definition of h , we have $f(h(y)) = y$.

Notice that the definition of g in the first part of the proof requires the function to “decide” whether there is an x in X such that $f(x) = y$. There is nothing mathematically dubious about this definition, but in many situations, this cannot be done *algorithmically*; in other words, g might not be computable from the data. More interestingly, the definition of h in the second part of the proof requires the function to “choose” a suitable value of x from among potentially many candidates. We will see in [Section 23.3](#) that this is a version of the *axiom of choice*. In the early twentieth century, the use of the axiom of choice in mathematics was hotly debated, but today it is commonplace.

Using these equivalences and the results in the previous section, we can prove the following:

Proposition. Let $f : X \rightarrow B$ and $g : Y \rightarrow Z$.

- If f and g are injective, then so is $g \circ f$.
- If f and g are surjective, then so is $g \circ f$.

Proof. If f and g are injective, then they have left inverses h and k , respectively, in which case $h \circ k$ is a left inverse to $g \circ f$. The second statement is proved similarly.

We can prove these two statements, however, without mentioning inverses at all. We leave that to you as an exercise.

Notice that the expression $f(n) = 2n$ can be used to define infinitely many functions with domain \mathbb{N} , such as:

- a function $f : \mathbb{N} \rightarrow \mathbb{N}$
- a function $f : \mathbb{N} \rightarrow \mathbb{R}$

- a function $f : \mathbb{N} \rightarrow \{n \mid n \text{ is even}\}$

Only the third one is surjective. Thus a specification of the function's codomain as well as the domain is essential to making sense of whether a function is surjective.

15.3 Functions and Subsets of the Domain

Suppose f is a function from X to Y . We may wish to reason about the behavior of f on some subset A of X . For example, we can say that f is *injective on A* if for every x_1 and x_2 in A , if $f(x_1) = f(x_2)$, then $x_1 = x_2$.

If f is a function from X to Y and A is a subset of X , we write $f[A]$ to denote the *image of f on A* , defined by

$$f[A] = \{y \in Y \mid y = f(x) \text{ for some } x \text{ in } A\}.$$

In words, $f[A]$ is the set of elements of Y that are “hit” by elements of A under the mapping f . Notice that there is an implicit existential quantifier here, so that reasoning about images invariably involves the corresponding rules.

Proposition. Suppose $f : X \rightarrow Y$, and A is a subset of X . Then for any x in A , $f(x)$ is in $f[A]$.

Proof. By definition, $f(x)$ is in $f[A]$ if and only if there is some x' in A such that $f(x') = f(x)$. But that holds for $x' = x$.

Proposition. Suppose $f : X \rightarrow Y$ and $g : Y \rightarrow Z$. Let A be a subset of X . Then

$$(g \circ f)[A] = g[f[A]].$$

Proof. Suppose z is in $(g \circ f)[A]$. Then for some $x \in A$, $z = (g \circ f)(x) = g(f(x))$. By the previous proposition, $f(x)$ is in $f[A]$. Again by the previous proposition, $g(f(x))$ is in $g[f[A]]$.

Conversely, suppose z is in $g[f[A]]$. Then there is a y in $f[A]$ such that $f(y) = z$, and since y is in $f[A]$, there is an x in A such that $f(x) = y$. But then $(g \circ f)(x) = g(f(x)) = g(y) = z$, so z is in $(g \circ f)[A]$.

Notice that if f is a function from X to Y , then f is surjective if and only if $f[X] = Y$. So the previous proposition is a generalization of the fact that the composition of surjective functions is surjective.

Suppose f is a function from X to Y , and A is a subset of X . We can *view f* as a function from A to Y , by simply ignoring the behavior of f on elements outside of A . Properly speaking, this is another function, denoted $f \upharpoonright A$ and called “the restriction of f to A .” In other words, given $f : X \rightarrow Y$ and $A \subseteq X$, $f \upharpoonright A : A \rightarrow Y$ is the function defined by $(f \upharpoonright A)(x) = f(x)$ for every x in A . Notice that now “ f is injective on A ” means simply that the restriction of f to A is injective.

There is another important operation on functions, known as the *preimage*. If $f : X \rightarrow Y$ and $B \subseteq Y$, then the *preimage of B under f* , denoted $f^{-1}[B]$, is defined by

$$f^{-1}[B] = \{x \in X \mid f(x) \in B\},$$

that is, the set of elements of X that get mapped into B . Notice that this makes sense even if f does not have an inverse; for a given y in B , there may be no x 's with the property $f(x) \in B$, or there may be many. If f has an inverse, f^{-1} , then for every y in B there is exactly one $x \in X$ with the property $f(x) \in B$, in which case, $f^{-1}[B]$ means the same thing whether you interpret it as the image of B under f^{-1} or the preimage of B under f .

Proposition. Suppose $f : X \rightarrow Y$ and $g : Y \rightarrow Z$. Let C be a subset of Z . Then

$$(g \circ f)^{-1}[C] = f^{-1}[g^{-1}[C]].$$

Proof. For any y in C , y is in $(g \circ f)^{-1}[C]$ if and only if $g(f(y))$ is in C . This, in turn, happens if and only if $f(y)$ is in $g^{-1}[C]$, which in turn happens if and only if y is in $f^{-1}[g^{-1}[C]]$.

Here we give a long list of facts properties of images and preimages. Here, f denotes an arbitrary function from X to Y , A, A_1, A_2, \dots denote arbitrary subsets of X , and B, B_1, B_2, \dots denote arbitrary subsets of Y .

- $A \subseteq f^{-1}[f[A]]$, and if f is injective, $A = f^{-1}[f[A]]$.
- $f[f^{-1}[B]] \subseteq B$, and if f is surjective, $B = f[f^{-1}[B]]$.
- If $A_1 \subseteq A_2$, then $f[A_1] \subseteq f[A_2]$.
- If $B_1 \subseteq B_2$, then $f^{-1}[B_1] \subseteq f^{-1}[B_2]$.
- $f[A_1 \cup A_2] = f[A_1] \cup f[A_2]$.
- $f^{-1}[B_1 \cup B_2] = f^{-1}[B_1] \cup f^{-1}[B_2]$.
- $f[A_1 \cap A_2] \subseteq f[A_1] \cap f[A_2]$, and if f is injective, $f[A_1 \cap A_2] = f[A_1] \cap f[A_2]$.
- $f^{-1}[B_1 \cap B_2] = f^{-1}[B_1] \cap f^{-1}[B_2]$.
- $f[A] \setminus f[B] \subseteq f[A \setminus B]$.
- $f^{-1}[A] \setminus f^{-1}[B] \subseteq f^{-1}[A \setminus B]$.
- $f[A] \cap B = f[A \cap f^{-1}[B]]$.
- $f[A] \cup B \supseteq f[A \cup f^{-1}[B]]$.
- $A \cap f^{-1}[B] \subseteq f^{-1}[f[A] \cap B]$.
- $A \cup f^{-1}[B] \subseteq f^{-1}[f[A] \cup B]$.

Proving identities like this is typically a matter of unfolding definitions and using basic logical inferences. Here is an example.

Proposition. Let X and Y be sets, $f : X \rightarrow Y$, $A \subseteq X$, and $B \subseteq Y$. Then $f[A] \cap B = f[A \cap f^{-1}[B]]$.

Proof. Suppose $y \in f[A] \cap B$. Then $y \in B$, and for some $x \in A$, $f(x) = y$. But this means that x is in $f^{-1}[B]$, and so $x \in A \cap f^{-1}[B]$. Since $f(x) = y$, we have $y \in f[A \cap f^{-1}[B]]$, as needed.

Conversely, suppose $y \in f[A \cap f^{-1}[B]]$. Then for some $x \in A \cap f^{-1}[B]$, we have $f(x) = y$. For this x , have $x \in A$ and $f(x) \in B$. Since $f(x) = y$, we have $y \in B$, and since $x \in A$, we also have $y \in f[A]$, as required.

15.4 Functions and Relations

A binary relation $R(x, y)$ on A and B is *functional* if for every x in A there exists a unique y in B such that $R(x, y)$. If R is a functional relation, we can define a function $f_R : X \rightarrow B$ by setting $f_R(x)$ to be equal to the unique y in B such that $R(x, y)$. Conversely, it is not hard to see that if $f : X \rightarrow B$ is any function, the relation $R_f(x, y)$ defined by $f(x) = y$ is a functional relation. The relation $R_f(x, y)$ is known as the *graph* of f .

It is not hard to check that functions and relations travel in pairs: if f is the function associated with a functional relation R , then R is the functional relation associated the function f , and vice-versa. In set-theoretic foundations, a function is often defined to be a functional relation. Conversely, we have seen that in type-theoretic foundations like the one adopted by Lean, relations are often defined to be certain types of functions. We will discuss these matters later on, and in the meanwhile only remark that in everyday mathematical practice, the foundational details are not so important; what is important is simply that every function has a graph, and that any functional relation can be used to define a corresponding function.

So far, we have been focusing on functions that take a single argument. We can also consider functions $f(x, y)$ or $g(x, y, z)$ that take multiple arguments. For example, the addition function $f(x, y) = x + y$ on the integers takes two integers and returns an integer. Remember, we can consider binary functions, ternary functions, and so on, and the number of arguments to a function is called its “arity.” One easy way to make sense of functions with multiple arguments is to think of them as unary functions from a cartesian product. We can think of a function f which takes two arguments, one in A and one in B , and returns an argument in C as a unary function from $A \times B$ to C , whereby $f(a, b)$ abbreviates $f((a, b))$. We have seen that in dependent type theory (and in Lean) it is more convenient to think of such a function f as a function which takes an element of A and returns a function from $B \rightarrow C$, so that $f(a, b)$ abbreviates $(f(a))(b)$. Such a function f maps A to $B \rightarrow C$, where $B \rightarrow C$ is the set of functions from B to C .

We will return to these different ways of modeling functions of higher arity later on, when we consider set-theoretic and type-theoretic foundations. One again, we remark that in ordinary mathematics, the foundational details do not matter much. The two choices above are inter-translatable, and sanction the same principles for reasoning about functions informally.

In mathematics, we often also consider the notion of a *partial function* from X to Y , which is really a function from some subset of X to Y . The fact that f is a partial function from X to Y is sometimes written $f : X \dashrightarrow Y$, which should be interpreted as saying that $f : A \rightarrow Y$ for some subset A of X . Intuitively, we think of f as a function from $X \rightarrow Y$ which is simply “undefined” at some of its inputs; for example, we can think of $f : \mathbb{R} \dashrightarrow \mathbb{R}$ defined by $f(x) = 1/x$, which is undefined at $x = 0$, so that in reality $f : \mathbb{R} \setminus \{0\} \rightarrow \mathbb{R}$. The set A is sometimes called the *domain of f* , in which case, there is no good name for X ; others continue to call X the domain, and refer to A as the *domain of definition*. To indicate that a function f is defined at x , that is, that x is in the domain of definition of f , we sometimes write $f(x) \downarrow$. If f and g are two partial functions from X to Y , we write $f(x) \simeq g(x)$ to mean that either f and g are both defined at x and have the same value, or are both undefined at x . Notions of injectivity, surjectivity, and composition are extended to partial functions, generally as you would expect them to be.

In terms of relations, a partial function f corresponds to a relation $R_f(x, y)$ such that for every x there is at most one y such that $R_f(x, y)$ holds. Mathematicians also sometimes consider *multifunctions* from X to Y , which correspond to relations $R_f(x, y)$ such that for every x in X , there is *at least* one y such that $R_f(x, y)$ holds. There may be many such y ; you can think of these as functions which have more than one output value. If you think about it for a moment, you will see that a *partial multifunction* is essentially nothing more than an arbitrary relation.

15.5 Exercises

1. Let f be any function from X to Y , and let g be any function from Y to Z .

- Show that if $g \circ f$ is injective, then f is injective.
 - Give an example of functions f and g as above, such that that $g \circ f$ is injective, but g is not injective.
 - Show that if $g \circ f$ is injective and f is surjective, then g is injective.
2. Let f and g be as in the last problem. Suppose $g \circ f$ is surjective.
 - Is f necessarily surjective? Either prove that it is, or give a counterexample.
 - Is g necessarily surjective? Either prove that it is, or give a counterexample.
 3. A function f from \mathbb{R} to \mathbb{R} is said to be *strictly increasing* if whenever $x_1 < x_2$, $f(x_1) < f(x_2)$.
 - Show that if $f : \mathbb{R} \rightarrow \mathbb{R}$ is strictly increasing, then it is injective (and hence it has a left inverse).
 - Show that if $f : \mathbb{R} \rightarrow \mathbb{R}$ is strictly increasing, and g is a right inverse to f , then g is strictly increasing.
 4. Let $f : X \rightarrow Y$ be any function, and let A and B be subsets of X . Show that $f[A \cup B] = f[A] \cup f[B]$.
 5. Let $f : X \rightarrow Y$ be any function, and let A and B be any subsets of X . Show $f[A] \setminus f[B] \subseteq f[A \setminus B]$.
 6. Define notions of composition and inverse for binary relations that generalize the notions for functions.

FUNCTIONS IN LEAN

16.1 Functions and Symbolic Logic

Let us now consider functions in formal terms. Even though we have avoided the use of quantifiers and logical symbols in the definitions in the last chapter, by now you should be seeing them lurking beneath the surface. That fact that two functions $f, g : X \rightarrow Y$ are equal if and only if they take the same values at every input can be expressed as follows:

$$\forall x \in X (f(x) = g(x)) \leftrightarrow f = g.$$

This principle is known as *function extensionality*, analogous to the principle of extensionality for sets, discussed in Section 12.1. Recall that the notation $\forall x \in X P(x)$ abbreviates $\forall x (x \in X \rightarrow P(x))$, and $\exists x \in X P(x)$ abbreviates $\exists x (x \in X \wedge P(x))$, thereby relativizing the quantifiers to X .

We can avoid set-theoretic notation if we assume we are working in a logical formalism with basic types for X and Y , so that we can specify that x ranges over X . In that case, we will write instead

$$\forall x : X (f(x) = g(x) \leftrightarrow f = g)$$

to indicate that the quantification is over X . Henceforth, we will assume that all our variables range over some type, though we will sometimes omit the types in the quantifiers when they can be inferred from context.

The function f is injective if it satisfies

$$\forall x_1, x_2 : X (f(x_1) = f(x_2) \rightarrow x_1 = x_2),$$

and f is surjective if

$$\forall y : Y \exists x : X f(x) = y.$$

If $f : X \rightarrow Y$ and $g : Y \rightarrow X$, g is a left inverse to f if

$$\forall x : X g(f(x)) = x.$$

Notice that this is a universal statement, and it is equivalent to the statement that f is a right inverse to g .

Remember that in logic it is common to use lambda notation to define functions. We can denote the identity function by $\lambda x x$, or perhaps $\lambda x : X x$ to emphasize that the domain of the function is X . If $f : X \rightarrow Y$ and $g : Y \rightarrow Z$, we can define the composition $g \circ f$ by $g \circ f = \lambda x : X g(f(x))$.

Also remember that if $P(x)$ is any predicate, then in first order logic we can assert that there exists a unique x satisfying $P(x)$, written $\exists! x P(x)$, with the conjunction of the following two statements:

- $\exists x P(x)$

- $\forall x_1, x_2 (P(x_1) \wedge P(x_2) \rightarrow x_1 = x_2)$

Equivalently, we can write

$$\exists (P(x) \wedge \forall x' (P(x') \rightarrow x' = x)).$$

Assuming $\exists!x P(x)$, the following two statements are equivalent:

- $\exists x (P(x) \wedge Q(x))$
- $\forall x (P(x) \rightarrow Q(x))$

and both can be taken to assert that “the x satisfying P also satisfies Q .”

A binary relation R on X and Y is functional if it satisfies

$$\forall x \exists!y R(x, y).$$

In that case, a logician might use *iota notation*,

$$f(x) = \iota y R(x, y)$$

to define $f(x)$ to be equal to the unique y satisfying $R(x, y)$. If R satisfies the weaker property

$$\forall x \exists y R(x, y),$$

a logician might use the *Hilbert epsilon* to define a function

$$f(x) = \varepsilon y R(x, y)$$

to “choose” a value of y satisfying $R(x, y)$. As we have noted above, this is an implicit use of the axiom of choice.

16.2 Second- and Higher-Order Logic

In contrast to first-order logic, where we start with a fixed stock of function and relation symbols, the topics we have been considering in the last few chapters encourage us to consider a more expressive language with variables ranging over functions and relations as well. For example, saying that a function $f : X \rightarrow Y$ has a left-inverse implicitly involves a quantifying over functions,

$$\exists g \forall x g(f(x)) = x.$$

The theorem that asserts that if any function f from X to Y is injective then it has a left-inverse can be expressed as follows:

$$\forall x_1, x_2 (f(x_1) = f(x_2) \rightarrow x_1 = x_2) \rightarrow \exists g \forall x g(f(x)) = x.$$

Similarly, saying that two sets X and Y have a one-to-one correspondence asserts the existence of a function $f : X \rightarrow Y$ as well as an inverse to f . For another example, in [Section 15.4](#) we asserted that every functional relation gives rise to a corresponding function, and vice-versa.

What makes these statements interesting is that they involve quantification, both existential and universal, over functions and relations. This takes us outside the realm of first-order logic. One option is to develop a theory in the language of first-order logic in which the universe contains functions and relations as objects; we will see later that this is what axiomatic set theory does. An alternative is to extend first-order logic to involve new kinds of quantifiers and variables, to range over functions and relations. This is what higher-order logic does.

There are various ways to go about this. In view of the relationship between functions and relations described earlier, one can take relations as basic, and define functions in terms of them, or vice-versa. The following formulation of higher-order logic, due to the logician Alonzo Church, follows the latter approach. It is sometimes known as *simple type theory*.

Start with some basic types, X, Y, Z, \dots and a special type, Prop , of propositions. Add the following two rules to build new types:

- If U and V are types, so is $U \times V$.
- If U and V are types, so is $U \rightarrow V$.

The first is intended to denote the type of ordered pairs (u, v) , where u is in U and v is in V . The second is intended to denote the type of functions from U to V . Simple type theory now adds the following means of forming expressions:

- If u is of type U and v is of type V , (u, v) is of type $U \times V$.
- If p is of type $U \times V$, then $(p)_1$ is of type U and $(p)_2$ is of type V . (These are intended to denote the first and second element of the pair p .)
- If x is a variable of type U , and v is any expression of type V , then $\lambda x v$ is of type $U \rightarrow V$.
- If f is of type $U \rightarrow V$ and u is of type U , $f(u)$ is of type V .

In addition, simple type theory provides all the means we have in first-order logic—boolean connectives, quantifiers, and equality—to build propositions.

A function $f(x, y)$ which takes elements of X and Y to a type Z is viewed as an object of type $X \times Y \rightarrow Z$. Similarly, a binary relation $R(x, y)$ on X and Y is viewed as an object of type $X \times Y \rightarrow \text{Prop}$. What makes higher-order logic “higher order” is that we can iterate the function type operation indefinitely. For example, if \mathbb{N} is the type of natural numbers, $\mathbb{N} \rightarrow \mathbb{N}$ denotes the type of functions from the natural numbers to the natural numbers, and $(\mathbb{N} \rightarrow \mathbb{N}) \rightarrow \mathbb{N}$ denotes the type of functions $F(f)$ which take a function as argument, and return a natural number.

We have not specified the syntax and rules of higher-order logic very carefully. This is done in a number of more advanced logic textbooks. The fragment of higher-order logic which allows only functions and relations on the basic types (without iterating these constructions) is known as second-order logic.

These notions should seem familiar; we have been using these constructions, with similar notation, in Lean. Indeed, Lean’s logic is an even more elaborate and expressive system of logic, which fully subsumes all the notions of higher-order logic we have discussed here.

16.3 Functions in Lean

The fact that the notions we have been discussing have such a straightforward logical form means that it is easy to define them in Lean. The main difference between the formal representation in Lean and the informal representation above is that, in Lean, we distinguish between a type X and a subset $A : \text{set } X$ of that type.

In Lean’s library, composition and identity are defined as follows:

```
variables {X Y Z : Type}

def comp (f : Y → Z) (g : X → Y) : X → Z :=
λx, f (g x)

infixr ` ∘ ` := comp
```

(continues on next page)

(continued from previous page)

```
def id (x : X) : X :=
x
```

Ordinarily, we use `funext` (for “function extensionality”) to prove that two functions are equal.

```
example (f g : X → Y) (h : ∀ x, f x = g x) : f = g :=
funext h
```

But Lean can prove some basic identities by simply unfolding definitions and simplifying expressions, using reflexivity.

```
lemma left_id (f : X → Y) : id ∘ f = f := rfl

lemma right_id (f : X → Y) : f ∘ id = f := rfl

theorem comp.assoc (f : Z → W) (g : Y → Z) (h : X → Y) :
  (f ∘ g) ∘ h = f ∘ (g ∘ h) := rfl

theorem comp.left_id (f : X → Y) : id ∘ f = f := rfl

theorem comp.right_id (f : X → Y) : f ∘ id = f := rfl
```

We can define what it means for f to be injective, surjective, or bijective:

```
def injective (f : X → Y) : Prop := ∀ {x1 x2}, f x1 = f x2 → x1 = x2

def surjective (f : X → Y) : Prop := ∀ y, ∃ x, f x = y

def bijective (f : X → Y) := injective f ∧ surjective f
```

Marking the variables x_1 and x_2 implicit in the definition of `injective` means that we do not have to write them as often. Specifically, given $h : \text{injective } f$, and $h_1 : f\ x_1 = f\ x_2$, we write $h\ h_1$ rather than $h\ x_1\ x_2\ h_1$ to show $x_1 = x_2$.

We can then prove that the identity function is bijective:

```
theorem injective_id : injective (@id X) :=
assume x1 x2,
assume H : id x1 = id x2,
show x1 = x2, from H

theorem surjective_id : surjective (@id X) :=
assume y,
show ∃ x, id x = y, from exists.intro y rfl

theorem bijective_id : bijective (@id X) :=
and.intro injective_id surjective_id
```

More interestingly, we can prove that the composition of injective functions is injective, and so on.

```
theorem injective_comp {g : Y → Z} {f : X → Y}
  (Hg : injective g) (Hf : injective f) :
```

(continues on next page)

(continued from previous page)

```

    injective (g ∘ f) :=
  assume x₁ x₂,
  assume : (g ∘ f) x₁ = (g ∘ f) x₂,
  have f x₁ = f x₂, from Hg this,
  show x₁ = x₂, from Hf this

theorem surjective_comp {g : Y → Z} {f : X → Y}
  (hg : surjective g) (hf : surjective f) :
  surjective (g ∘ f) :=
  assume z,
  exists.elim (hg z) $
  assume y (hy : g y = z),
  exists.elim (hf y) $
  assume x (hx : f x = y),
  have g (f x) = z, from eq.subst (eq.symm hx) hy,
  show ∃ x, g (f x) = z, from exists.intro x this

theorem bijective_comp {g : Y → Z} {f : X → Y}
  (hg : bijective g) (hf : bijective f) :
  bijective (g ∘ f) :=
  have ginj : injective g, from hg.left,
  have gsurj : surjective g, from hg.right,
  have finj : injective f, from hf.left,
  have fsurj : surjective f, from hf.right,
  and.intro (injective_comp ginj finj)
  (surjective_comp gsurj fsurj)

```

The notions of left and right inverse are defined in the expected way.

```

-- g is a left inverse to f
def left_inverse (g : Y → X) (f : X → Y) : Prop := ∀ x, g (f x) = x

-- g is a right inverse to f
def right_inverse (g : Y → X) (f : X → Y) : Prop := left_inverse f g

```

In particular, composing with a left or right inverse yields the identity.

```

def id_of_left_inverse {g : Y → X} {f : X → Y} : left_inverse g f → g ∘ f = id :=
  assume H, funext H

def id_of_right_inverse {g : Y → X} {f : X → Y} : right_inverse g f → f ∘ g = id :=
  assume H, funext H

```

Notice that we need to use `funext` to show the equality of functions.

The following shows that if a function has a left inverse, then it is injective, and if it has a right inverse, then it is surjective.

```

theorem injective_of_left_inverse {g : Y → X} {f : X → Y} :
  left_inverse g f → injective f :=
  assume h, assume x₁ x₂, assume feq,
  calc x₁ = g (f x₁) : by rw h
      ... = g (f x₂) : by rw feq

```

(continues on next page)

(continued from previous page)

```

... = x₂      : by rw h

theorem surjective_of_right_inverse {g : Y → X} {f : X → Y} :
  right_inverse g f → surjective f :=
  assume h, assume y,
  let x : X := g y in
  have f x = y, from calc
  f x = (f (g y))    : rfl
  ... = y            : by rw [h y],
  show ∃ x, f x = y, from exists.intro x this

```

16.4 Defining the Inverse Classically

All the theorems listed in the previous section are found in the Lean library, and are available to you when you open the function namespace with `open function`:

```

open function

#check comp
#check left_inverse
#check has_right_inverse

```

Defining inverse functions, however, requires classical reasoning, which we get by opening the classical namespace:

```

open classical

section
  variables A B : Type
  variable P : A → Prop
  variable R : A → B → Prop

  example : (∀ x, ∃ y, R x y) → ∃ f : A → B, ∀ x, R x (f x) :=
  axiom_of_choice

  example (h : ∃ x, P x) : P (some h) :=
  some_spec h
end

```

The axiom of choice tells us that if, for every $x : X$, there is a $y : Y$ satisfying $R\ x\ y$, then there is a function $f : X \rightarrow Y$ which, for every x chooses such a y . In Lean, this “axiom” is proved using a classical construction, the `some` function (sometimes called “the indefinite description operator”) which, given that there is some x satisfying $P\ x$, returns such an x . With these constructions, the inverse function is defined as follows:

```

open classical function
local attribute [instance] prop_decidable

variables {X Y : Type}

```

(continues on next page)

(continued from previous page)

```

noncomputable def inverse (f : X → Y) (default : X) : Y → X :=
λ y, if h : ∃ x, f x = y then some h else default

```

Lean requires us to acknowledge that the definition is not computational, since, first, it may not be algorithmically possible to decide whether or not condition H holds, and even if it does, it may not be algorithmically possible to find a suitable value of x .

Below, the proposition `inverse_of_exists` asserts that `inverse` meets its specification, and the subsequent theorem shows that if f is injective, then the `inverse` function really is a left inverse.

```

theorem inverse_of_exists (f : X → Y) (default : X) (y : Y)
  (h : ∃ x, f x = y) :
f (inverse f default y) = y :=
have h1 : inverse f default y = some h, from dif_pos h,
have h2 : f (some h) = y, from some_spec h,
eq.subst (eq.symm h1) h2

theorem is_left_inverse_of_injective (f : X → Y) (default : X)
  (injf : injective f) :
left_inverse (inverse f default) f :=
let finv := (inverse f default) in
assume x,
have h1 : ∃ x', f x' = f x, from exists.intro x rfl,
have h2 : f (finv (f x)) = f x, from inverse_of_exists f default (f x) h1,
show finv (f x) = x, from injf h2

```

16.5 Functions and Sets in Lean

In [Section 7.4](#) we saw how to represent relativized universal and existential quantifiers when formalizing phrases like “every prime number greater than two is odd” and “some prime number is even.” In a similar way, we can relativize statements to sets. In symbolic logic, the expression $\exists x \in A P(x)$ abbreviates $\exists x (x \in A \wedge P(x))$, and $\forall x \in A P(x)$ abbreviates $\forall x (x \in A \rightarrow P(x))$.

Lean also defines notation for relativized quantifiers:

```

variables (X : Type) (A : set X) (P : X → Prop)

#check ∀ x ∈ A, P x
#check ∃ x ∈ A, P x

```

Here is an example of how to use the bounded universal quantifier:

```

example (h : ∀ x ∈ A, P x) (x : X) (h1 : x ∈ A) : P x := h x h1

```

Using bounded quantifiers, we can talk about the behavior of functions on particular sets:

```

import data.set
open set function

variables {X Y : Type}
variables (A : set X) (B : set Y)

```

(continues on next page)

(continued from previous page)

```
def maps_to (f : X → Y) (A : set X) (B : set Y) := ∀ x ∈ A, f x ∈ B

def inj_on (f : X → Y) (A : set X) := ∀ (x₁ ∈ A) (x₂ ∈ A), f x₁ = f x₂ → x₁ = x₂

def surj_on (f : X → Y) (A : set X) (B : set Y) := B ⊆ f '' A
```

The expression `maps_to f A B` asserts that `f` maps elements of the set `A` to the set `B`, and the expression `inj_on f A` asserts that `f` is injective on `A`. The expression `surj_on f A B` asserts that, viewed as a function defined on elements of `A`, the function `f` is surjective onto the set `B`. Here are examples of how they can be used:

```
variables (f : X → Y) (A : set X) (B : set Y)

example (h : maps_to f A B) (x : X) (h1 : x ∈ A) : f x ∈ B := h x h1

example (h : inj_on f A) (x₁ x₂ : X) (h1 : x₁ ∈ A) (h2 : x₂ ∈ A)
  (h3 : f x₁ = f x₂) : x₁ = x₂ :=
h x₁ h1 x₂ h2 h3

example (h : surj_on f A B) (y : Y) (h1 : y ∈ B) : ∃ x, x ∈ A ∧ f x = y :=
h h1
```

Actually, it is slightly more convenient to mark the variables in `maps_to` and `inj_on` as implicit:

```
def maps_to (f : X → Y) (A : set X) (B : set Y) := ∀ {x}, x ∈ A → f x ∈ B

def inj_on (f : X → Y) (A : set X) := ∀ {x₁ x₂}, x₁ ∈ A → x₂ ∈ A → f x₁ = f x₂ → x₁ = x₂
```

In that case, we can leave out some arguments:

```
variables (f : X → Y) (A : set X) (B : set Y)

example (h : maps_to f A B) (x : X) (h1 : x ∈ A) : f x ∈ B := h h1

example (h : inj_on f A) (x₁ x₂ : X) (h1 : x₁ ∈ A) (h2 : x₂ ∈ A)
  (h3 : f x₁ = f x₂) : x₁ = x₂ :=
h h1 h2 h3
```

In the examples below, we'll use the versions with implicit arguments. The expression `surj_on f A B` asserts that, viewed as a function defined on elements of `A`, the function `f` is surjective onto the set `B`.

With these notions in hand, we can prove that the composition of injective functions is injective. The proof is similar to the one above, though now we have to be more careful to relativize claims to `A` and `B`:

```
theorem inj_on_comp (fAB : maps_to f A B) (hg : inj_on g B) (hf : inj_on f A) :
  inj_on (g ∘ f) A :=
assume x1 x2 : X,
assume x1A : x1 ∈ A,
assume x2A : x2 ∈ A,
have fx1B : f x1 ∈ B, from fAB x1A,
have fx2B : f x2 ∈ B, from fAB x2A,
```

(continues on next page)

(continued from previous page)

```

assume h1 : g (f x1) = g (f x2),
have h2 : f x1 = f x2, from hg fx1B fx2B h1,
show x1 = x2, from hf x1A x2A h2
    
```

We can similarly prove that the composition of surjective functions is surjective:

```

theorem surj_on_comp (hg : surj_on g B C) (hf : surj_on f A B) :
  surj_on (g ∘ f) A C :=
assume z,
assume zc : z ∈ C,
exists.elim (hg zc) $
assume y (h1 : y ∈ B ∧ g y = z),
exists.elim (hf (and.left h1)) $
assume x (h2 : x ∈ A ∧ f x = y),
show ∃x, x ∈ A ∧ g (f x) = z, from
  exists.intro x
    (and.intro
      (and.left h2)
      (calc
        g (f x) = g y : by rw and.right h2
          ... = z : by rw and.right h1))
    
```

The following shows that the image of a union is the union of images:

```

theorem image_union : f '' (A1 ∪ A2) = f '' A1 ∪ f '' A2 :=
ext (assume y, iff.intro
  (assume h : y ∈ image f (A1 ∪ A2),
    exists.elim h $
      assume x h1,
      have xA1A2 : x ∈ A1 ∪ A2, from h1.left,
      have fxy : f x = y, from h1.right,
      or.elim xA1A2
        (assume xA1, or.inl ⟨x, xA1, fxy⟩)
        (assume xA2, or.inr ⟨x, xA2, fxy⟩))
  (assume h : y ∈ image f A1 ∪ image f A2,
    or.elim h
      (assume yifA1 : y ∈ image f A1,
        exists.elim yifA1 $
          assume x h1,
          have xA1 : x ∈ A1, from h1.left,
          have fxy : f x = y, from h1.right,
          ⟨x, or.inl xA1, fxy⟩)
      (assume yifA2 : y ∈ image f A2,
        exists.elim yifA2 $
          assume x h1,
          have xA2 : x ∈ A2, from h1.left,
          have fxy : f x = y, from h1.right,
          ⟨x, (or.inr xA2), fxy⟩)))
    
```

Note that the expression $y \in \text{image } f A_1$ expands to $\exists x, x \in A_1 \wedge f x = y$. We therefore need to provide three pieces of information: a value of x , a proof that $x \in A_1$, and a proof that $f x = y$. On the eighth line of this proof, after `or.inl`, we could justify the necessary information by writing `exists.intro x (and.intro xA1 fxy)` in parentheses. But in this case Lean’s “anonymous constructor” notation, that is,

the corner brackets entered with `\<` and `\>`, allow us to use the more compact expression $\langle x, xA_1, fxy \rangle$.

16.6 Exercises

1. Fill in the sorry's in the last three proofs below.

```

open function int algebra

def f (x : ℤ) : ℤ := x + 3
def g (x : ℤ) : ℤ := -x
def h (x : ℤ) : ℤ := 2 * x + 3

example : injective f :=
assume x1 x2,
assume h1 : x1 + 3 = x2 + 3, -- Lean knows this is the same as f x1 = f x2
show x1 = x2, from eq_of_add_eq_add_right h1

example : surjective f :=
assume y,
have h1 : f (y - 3) = y, from calc
  f (y - 3) = (y - 3) + 3 : rfl
  ... = y           : by rw sub_add_cancel,
show ∃ x, f x = y, from exists.intro (y - 3) h1

example (x y : ℤ) (h : 2 * x = 2 * y) : x = y :=
have h1 : 2 ≠ (0 : ℤ), from dec_trivial, -- this tells Lean to figure it out
→itself
show x = y, from eq_of_mul_eq_mul_left h1 h

example (x : ℤ) : -(-x) = x := neg_neg x

example (A B : Type) (u : A → B) (v : B → A) (h : left_inverse u v) :
  ∀ x, u (v x) = x :=
h

example (A B : Type) (u : A → B) (v : B → A) (h : left_inverse u v) :
  right_inverse v u :=
h

-- fill in the sorry's in the following proofs

example : injective h :=
sorry

example : surjective g :=
sorry

example (A B : Type) (u : A → B) (v1 : B → A) (v2 : B → A)
  (h1 : left_inverse v1 u) (h2 : right_inverse v2 u) : v1 = v2 :=
funext
  (assume x,
    calc

```

(continues on next page)

(continued from previous page)

```
v1 x = v1 (u (v2 x)) : sorry
... = v2 x           : sorry
```

2. Fill in the sorry in the proof below.

```
import data.set
open function set

variables {X Y : Type}
variable f : X → Y
variables A B : set X

example : f '' (A ∪ B) = f '' A ∪ f '' B :=
eq_of_subset_of_subset
  (assume y,
    assume h1 : y ∈ f '' (A ∪ B),
    exists.elim h1 $
      assume x h,
      have h2 : x ∈ A ∪ B, from h.left,
      have h3 : f x = y, from h.right,
      or.elim h2
        (assume h4 : x ∈ A,
          have h5 : y ∈ f '' A, from ⟨x, h4, h3⟩,
          show y ∈ f '' A ∪ f '' B, from or.inl h5)
        (assume h4 : x ∈ B,
          have h5 : y ∈ f '' B, from ⟨x, h4, h3⟩,
          show y ∈ f '' A ∪ f '' B, from or.inr h5))
    (assume y,
      assume h2 : y ∈ f '' A ∪ f '' B,
      or.elim h2
        (assume h3 : y ∈ f '' A,
          exists.elim h3 $
            assume x h,
            have h4 : x ∈ A, from h.left,
            have h5 : f x = y, from h.right,
            have h6 : x ∈ A ∪ B, from or.inl h4,
            show y ∈ f '' (A ∪ B), from ⟨x, h6, h5⟩)
        (assume h3 : y ∈ f '' B,
          exists.elim h3 $
            assume x h,
            have h4 : x ∈ B, from h.left,
            have h5 : f x = y, from h.right,
            have h6 : x ∈ A ∪ B, from or.inr h4,
            show y ∈ f '' (A ∪ B), from ⟨x, h6, h5⟩))

-- remember, x ∈ A ∩ B is the same as x ∈ A ∧ x ∈ B
example (x : X) (h1 : x ∈ A) (h2 : x ∈ B) : x ∈ A ∩ B :=
and.intro h1 h2

example (x : X) (h1 : x ∈ A ∩ B) : x ∈ A :=
and.left h1
```

(continues on next page)

(continued from previous page)

```
-- Fill in the proof below.  
-- (It should take about 8 lines.)  
  
example : f '' (A ∩ B) ⊆ f '' A ∩ f '' B :=  
assume y,  
assume h1 : y ∈ f '' (A ∩ B),  
show y ∈ f '' A ∩ f '' B, from sorry
```

THE NATURAL NUMBERS AND INDUCTION

This chapter marks a transition from the abstract to the concrete. Viewing the mathematical universe in terms of sets, relations, and functions gives us useful ways of thinking about mathematical objects and structures and the relationships between them. At some point, however, we need to start thinking about *particular* mathematical objects and structures, and the natural numbers are a good place to start. The nineteenth century mathematician Leopold Kronecker once proclaimed “God created the whole numbers; everything else is the work of man.” By this he meant that the natural numbers (and the integers, which we will also discuss below) are a fundamental component of the mathematical universe, and that many other objects and structures of interest can be constructed from these.

In this chapter, we will consider the natural numbers and the basic principles that govern them. In [Chapter 18](#) we will see that even basic operations like addition and multiplication can be defined using means described here, and their properties derived from these basic principles. Our presentation in this chapter will remain informal, however. In [Chapter 19](#), we will see how these principles play out in number theory, one of the oldest and most venerable branches of mathematics.

17.1 The Principle of Induction

The set of natural numbers is the set

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

In the past, opinions have differed as to whether the set of natural numbers should start with 0 or 1, but these days most mathematicians take them to start with 0. Logicians often call the function $s(n) = n + 1$ the *successor* function, since it maps each natural number, n , to the one that follows it. What makes the natural numbers special is that they are *generated* by the number zero and the successor function, which is to say, the only way to construct a natural number is to start with 0 and apply the successor function finitely many times. From a foundational standpoint, we are in danger of running into a circularity here, because it is not clear how we can explain what it means to apply a function “finitely many times” without talking about the natural numbers themselves. But the following principle, known as the *principle of induction*, describes this essential property of the natural numbers in a non-circular way.

Principle of Induction. Let P be any property of natural numbers. Suppose P holds of zero, and whenever P holds of a natural number n , then it holds of its successor, $n + 1$. Then P holds of every natural number.

This reflects the image of the natural numbers as being generated by zero and the successor operation: by covering the zero and successor cases, we take care of all the natural numbers.

The principle of induction provides a recipe for proving that every natural number has a certain property: to show that P holds of every natural number, show that it holds of 0, and show that whenever it holds of

some number n , it holds of $n + 1$. This form of proof is called a *proof by induction*. The first required task is called the *base case*, and the second required task is called the *induction step*. The induction step requires temporarily fixing a natural number n , assuming that P holds of n , and then showing that P holds of $n + 1$. In this context, the assumption that P holds of n is called the *inductive hypothesis*.

You can visualize proof by induction as a method of knocking down an infinite stream of dominoes, all at once. We set the mechanism in place and knock down domino 0 (the base case), and every domino knocks down the next domino (the induction step). So domino 0 knocks down domino 1; that knocks down domino 2, and so on.

Here is an example of a proof by induction.

Theorem. For every natural number n ,

$$1 + 2 + \dots + 2^n = 2^{n+1} - 1.$$

Proof. We prove this by induction on n . In the base case, when $n = 0$, we have $1 = 2^{0+1} - 1$, as required.

For the induction step, fix n , and assume the *induction hypothesis*

$$1 + 2 + \dots + 2^n = 2^{n+1} - 1.$$

We need to show that this same claim holds with n replaced by $n + 1$. But this is just a calculation:

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

In the notation of first-order logic, if we write $P(n)$ to mean that P holds of n , we could express the principle of induction as follows:

$$P(0) \wedge \forall n (P(n) \rightarrow P(n + 1)) \rightarrow \forall n P(n).$$

But notice that the principle of induction says that the axiom holds *for every property* P , which means that we should properly use a universal quantifier for that, too:

$$\forall P (P(0) \wedge \forall n (P(n) \rightarrow P(n + 1)) \rightarrow \forall n P(n)).$$

Quantifying over properties takes us out of the realm of first-order logic; induction is therefore a second-order principle.

The pattern for a proof by induction is expressed even more naturally by the following natural deduction rule:

$$\frac{\begin{array}{c} \overline{P(n)}^1 \\ \vdots \\ P(0) \quad P(n + 1) \end{array}}{\forall n P(n)}$$

You should think about how some of the proofs in this chapter could be represented formally using natural deduction.

For another example of a proof by induction, let us derive a formula that, given any finite set S , determines the number of subsets of S . For example, there are four subsets of the two-element set $\{1, 2\}$, namely \emptyset , $\{1\}$, $\{2\}$, and $\{1, 2\}$. You should convince yourself that there are eight subsets of the set $\{1, 2, 3\}$. The following theorem establishes the general pattern.

Theorem. For any finite set S , if S has n elements, then there are 2^n subsets of S .

Proof. We use induction on n . In the base case, there is only one set with 0 elements, the empty set, and there is exactly one subset of the empty set, as required.

In the inductive case, suppose S has $n + 1$ elements. Let a be any element of S , and let S' be the set containing the remaining n elements. In order to count the subsets of S , we divide them into two groups.

First, we consider the subsets of S that don't contain a . These are exactly the subsets of S' , and by the inductive hypothesis, there are 2^n of those.

Next we consider the subsets of S that *do* contain a . Each of these is obtained by choosing a subset of S' and adding a . Since there are 2^n subsets of S' , there are 2^n subsets of S that contain a .

Taken together, then, there are $2^n + 2^n = 2^{n+1}$ subsets of S , as required.

We have seen that there is a correspondence between properties of a domain and subsets of a domain. For every property P of natural numbers, we can consider the set S of natural numbers with that property, and for every set of natural numbers, we can consider the property of being in that set. For example, we can talk about the property of being even, or talk about the set of even numbers. Under this correspondence, the principle of induction can be cast as follows:

Principle of Induction. Let S be any set of natural numbers that contains 0 and is closed under the successor operation. Then $S = \mathbb{N}$.

Here, saying that S is “closed under the successor operation” means that whenever a number n is in S , so is $n + 1$.

17.2 Variants of Induction

In this section, we will consider variations on the principle of induction that are often useful. It is important to recognize that each of these can be justified using the principle of induction as stated in the last section, so they need not be taken as fundamental.

The first one is no great shakes: instead of starting from 0, we can start from any natural number, m .

Principle of Induction from a Starting Point. Let P be any property of natural numbers, and let m be any natural number. Suppose P holds of m , and whenever P holds of a natural number n greater than or equal to m , then it holds of its successor, $n + 1$. Then P holds of every natural number greater than or equal to m .

Assuming the hypotheses of this last principle, if we let $P'(n)$ be the property “ P holds of $m + n$,” we can prove that P' holds of every n by the ordinary principle of induction. But this means that P holds of every number greater than or equal to m .

Here is one example of a proof using this variant of induction.

Theorem. For every natural number $n \geq 5$, $2^n > n^2$.

Proof. By induction on n . When $n = 5$, we have $2^n = 32 > 25 = n^2$, as required.

For the induction step, suppose $n \geq 5$ and $2^n > n^2$. Since n is greater than or equal to 5, we have $2n + 1 \leq 3n \leq n^2$, and so

$$\begin{aligned}(n + 1)^2 &= n^2 + 2n + 1 \\ &\leq n^2 + n^2 \\ &< 2^n + 2^n \\ &= 2^{n+1}.\end{aligned}$$

For another example, let us derive a formula for the sum total of the angles in a convex polygon. A polygon is said to be *convex* if every line between two vertices stays inside the polygon. We will accept without proof the visually obvious fact that one can subdivide any convex polygon with more than three sides into a triangle and a convex polygon with one fewer side, namely, by closing off any two consecutive sides to form a triangle. We will also accept, without proof, the basic geometric fact that the sum of the angles of any triangle is 180 degrees.

Theorem. For any $n \geq 3$, the sum of the angles of any convex n -gon is $180(n - 2)$.

Proof. In the base case, when $n = 3$, this reduces to the statement that the sum of the angles in any triangle is 180 degrees.

For the induction step, suppose $n \geq 3$, and let P be a convex $(n + 1)$ -gon. Divide P into a triangle and an n -gon. By the inductive hypotheses, the sum of the angles of the n -gon is $180(n - 2)$ degrees, and the sum of the angles of the triangle is 180 degrees. The measures of these angles taken together make up the sum of the measures of the angles of P , for a total of $180(n - 2) + 180 = 180(n - 1)$ degrees.

For our second example, we will consider the principle of *complete induction*, also sometimes known as *total induction*.

Principle of Complete Induction. Let P be any property that satisfies the following: for any natural number n , whenever P holds of every number less than n , it also holds of n . Then P holds of every natural number.

Notice that there is no need to break out a special case for zero: for any property P , P holds of all the natural numbers less than zero, for the trivial reason that there aren't any! So, in particular, any such property automatically holds of zero.

Notice also that if such a property P holds of every number less than n , then it also holds of every number less than $n + 1$ (why?). So, for such a P , the ordinary principle of induction implies that for every natural

number n , P holds of every natural number less than n . But this is just a roundabout way of saying that P holds of every natural number. In other words, we have justified the principle of complete induction using ordinary induction.

To use the principle of complete induction we merely have to let n be any natural number and show that P holds of n , assuming that it holds of every smaller number. Compare this to the ordinary principle of induction, which requires us to show $P(n+1)$ assuming only $P(n)$. The following example of the use of this principle is taken verbatim from the introduction to this book:

Theorem. Every natural number greater than or equal to 2 can be written as a product of primes.

Proof. We proceed by induction on n . Let n be any natural number greater than 2. If n is prime, we are done; we can consider n itself as a product with one factor. Otherwise, n is composite, and we can write $n = m \cdot k$ where m and k are smaller than n and greater than 1. By the inductive hypothesis, each of m and k can be written as a product of primes:

$$\begin{aligned} m &= p_1 \cdot p_2 \cdot \dots \cdot p_u \\ k &= q_1 \cdot q_2 \cdot \dots \cdot q_v. \end{aligned}$$

But then we have

$$n = m \cdot k = p_1 \cdot p_2 \cdot \dots \cdot p_u \cdot q_1 \cdot q_2 \cdot \dots \cdot q_v.$$

We see that n is a product of primes, as required.

Finally, we will consider another formulation of induction, known as the least element principle.

The Least Element Principle. Suppose P is some property of natural numbers, and suppose P holds of some n . Then there is a smallest value of n for which P holds.

In fact, using classical reasoning, this is equivalent to the principle of complete induction. To see this, consider the contrapositive of the statement above: “if there is no smallest value for which P holds, then P doesn’t hold of any natural number.” Let $Q(n)$ be the property “ P does *not* hold of n .” Saying that there is no smallest value for which P holds means that, for every n , if P holds at n , then it holds of some number smaller than n ; and this is equivalent to saying that, for every n , if Q doesn’t hold at n , then there is a smaller value for which Q doesn’t hold. And *that* is equivalent to saying that if Q holds for every number less than n , it holds for n as well. Similarly, saying that P doesn’t hold of any natural number is equivalent to saying that Q holds of every natural number. In other words, replacing the least element principle by its contrapositive, and replacing P by “not Q ,” we have the principle of complete induction. Since every statement is equivalent to its contrapositive, and every predicate has its negated version, the two principles are the same.

It is not surprising, then, that the least element principle can be used in much the same way as the principle of complete induction. Here, for example, is a formulation of the previous proof in these terms. Notice that it is phrased as a proof by contradiction.

Theorem. Every natural number greater than equal to 2 can be written as a product of primes.

Proof. Suppose, to the contrary, some natural number greater than or equal to 2 cannot be written as a product of primes. By the least element principle, there is a smallest such element; call it n . Then n is not prime, and since it is greater than or equal to 2, it must be composite. Hence we can write $n = m \cdot k$ where

m and k are smaller than n and greater than 1. By the assumption on n , each of m and k can be written as a product of primes:

$$\begin{aligned}m &= p_1 \cdot p_2 \cdot \dots \cdot p_u \\k &= q_1 \cdot q_2 \cdot \dots \cdot q_v.\end{aligned}$$

But then we have

$$n = m \cdot k = p_1 \cdot p_2 \cdot \dots \cdot p_u \cdot q_1 \cdot q_2 \cdot \dots \cdot q_v.$$

We see that n is a product of primes, contradicting the fact that n cannot be written as a product of primes.

Here is another example:

Theorem. Every natural number is interesting.

Proof. Suppose, to the contrary, some natural number is uninteresting. Then there is a smallest one, n . In other words, n is the smallest uninteresting number. But that is really interesting! Contradiction.

17.3 Recursive Definitions

Suppose I tell you that I have a function $f : \mathbb{N} \rightarrow \mathbb{N}$ in mind, satisfying the following properties:

$$\begin{aligned}f(0) &= 1 \\f(n+1) &= 2 \cdot f(n)\end{aligned}$$

What can you infer about f ? Try calculating a few values:

$$\begin{aligned}f(1) &= f(0+1) = 2 \cdot f(0) = 2 \\f(2) &= f(1+1) = 2 \cdot f(1) = 4 \\f(3) &= f(2+1) = 2 \cdot f(2) = 8\end{aligned}$$

It soon becomes apparent that for every n , $f(n) = 2^n$.

What is more interesting is that the two conditions above specify *all* the values of f , which is to say, there is exactly one function meeting the specification above. In fact, it does not matter that f takes values in the natural numbers; it could take values in any other domain. All that is needed is a value of $f(0)$ and a way to compute the value of $f(n+1)$ in terms of n and $f(n)$. This is what the principle of definition by recursion asserts:

Principle of Definition by Recursion. Let A be any set, and suppose a is in A , and $g : \mathbb{N} \times A \rightarrow A$. Then there is a unique function f satisfying the following two clauses:

$$\begin{aligned}f(0) &= a \\f(n+1) &= g(n, f(n)).\end{aligned}$$

The principle of recursive definition makes two claims at once: first, that there is a function f satisfying the clauses above, and, second, that any two functions f_1 and f_2 satisfying those clauses are equal, which is to

say, they have the same values for every input. In the example with which we began this section, A is just \mathbb{N} and $g(n, f(n)) = 2 \cdot f(n)$.

In some axiomatic frameworks, the principle of recursive definition can be justified using the principle of induction. In others, the principle of induction can be viewed as a special case of the principle of recursive definition. For now, we will simply take both to be fundamental properties of the natural numbers.

As another example of a recursive definition, consider the function $g : \mathbb{N} \rightarrow \mathbb{N}$ defined recursively by the following clauses:

$$\begin{aligned} g(0) &= 1 \\ g(n+1) &= (n+1) \cdot g(n) \end{aligned}$$

Try calculating the first few values. Unwrapping the definition, we see that $g(n) = 1 \cdot 2 \cdot 3 \cdot \dots \cdot (n-1) \cdot n$ for every n ; indeed, definition by recursion is usually the proper way to make expressions using “...” precise. The value $g(n)$ is read “ n factorial,” and written $n!$.

Indeed, summation notation

$$\sum_{i < n} f(i) = f(0) + f(1) + \dots + f(n-1)$$

and product notation

$$\prod_{i < n} f(i) = f(0) \cdot f(1) \cdot \dots \cdot f(n-1)$$

can also be made precise using recursive definitions. For example, the function $k(n) = \sum_{i < n} f(i)$ can be defined recursively as follows:

$$\begin{aligned} k(0) &= 0 \\ k(n+1) &= k(n) + f(n) \end{aligned}$$

Induction and recursion are complementary principles, and typically the way to prove something about a recursively defined function is to use the principle of induction. For example, the following theorem provides a formula for the sum $1 + 2 + \dots + n$, in terms of n .

Theorem. For every n , $\sum_{i < n+1} i = n(n+1)/2$.

Proof. In the base case, when $n = 0$, both sides are equal to 0.

In the inductive step, we have

$$\begin{aligned} \sum_{i < n+2} i &= \left(\sum_{i < n+1} i \right) + (n+1) \\ &= n(n+1)/2 + n+1 \\ &= \frac{n^2 + n}{2} + \frac{2n+2}{2} \\ &= \frac{n^2 + 3n + 2}{2} \\ &= \frac{(n+1)(n+2)}{2}. \end{aligned}$$

There are just as many variations on the principle of recursive definition as there are on the principle of induction. For example, in analogy to the principle of complete induction, we can specify a value of $f(n)$

in terms of the values that f takes at all inputs smaller than n . When $n \geq 2$, for example, the following definition specifies the value of a function $\text{fib}(n)$ in terms of its two predecessors:

$$\begin{aligned}\text{fib}(0) &= 0 \\ \text{fib}(1) &= 1 \\ \text{fib}(n+2) &= \text{fib}(n+1) + \text{fib}(n)\end{aligned}$$

Calculating the values of fib on $0, 1, 2, \dots$ we obtain

$$0, 1, 1, 2, 3, 5, 8, 13, 21, \dots$$

Here, after the second number, each successive number is the sum of the two values preceding it. This is known as the *Fibonacci sequence*, and the corresponding numbers are known as the *Fibonacci numbers*. An ordinary mathematical presentation would write F_n instead of $\text{fib}(n)$ and specify the sequence with the following equations:

$$F_0 = 0, \quad F_1 = 1, \quad F_{n+2} = F_{n+1} + F_n$$

But you can now recognize such a specification as an implicit appeal to the principle of definition by recursion. We ask you to prove some facts about the Fibonacci sequence in the exercises below.

17.4 Defining Arithmetic Operations

In fact, we can even use the principle of recursive definition to define the most basic operations on the natural numbers and show that they have the properties we expect them to have. From a foundational standpoint, we can characterize the natural numbers as a set, \mathbb{N} , with a distinguished element 0 and a function, $\text{succ}(m)$, which, for every natural number m , returns its *successor*. These satisfy the following:

- $0 \neq \text{succ}(m)$ for any m in \mathbb{N} .
- For every m and n in \mathbb{N} , if $m \neq n$, then $\text{succ}(m) \neq \text{succ}(n)$. In other words, succ is *injective*.
- If A is any subset of \mathbb{N} with the property that 0 is in A and whenever n is in A then $\text{succ}(n)$ is in A , then $A = \mathbb{N}$.

The last clause can be reformulated as the principle of induction:

Suppose $P(n)$ is any property of natural numbers, such that P holds of 0, and for every n , $P(n)$ implies $P(\text{succ}(n))$. Then every P holds of every natural number.

Remember that this principle can be used to justify the principle of definition by recursion:

Let A be any set, a be any element of A , and let $g(n, m)$ be any function from $\mathbb{N} \times A$ to A . Then there is a unique function $f : \mathbb{N} \rightarrow A$ satisfying the following two clauses:

- $f(0) = a$
- $f(\text{succ}(n)) = g(n, f(n))$ for every n in \mathbb{N}

We can use the principle of recursive definition to define addition with the following two clauses:

$$\begin{aligned}m + 0 &= m \\ m + \text{succ}(n) &= \text{succ}(m + n)\end{aligned}$$

Note that we are fixing m , and viewing this as a function of n . If we write $1 = \text{succ}(0)$, $2 = \text{succ}(1)$, and so on, it is easy to prove $n + 1 = \text{succ}(n)$ from the definition of addition.

We can proceed to define multiplication using the following two clauses:

$$\begin{aligned} m \cdot 0 &= 0 \\ m \cdot \text{succ}(n) &= m \cdot n + m \end{aligned}$$

We can also define a predecessor function by

$$\begin{aligned} \text{pred}(0) &= 0 \\ \text{pred}(\text{succ}(n)) &= n \end{aligned}$$

We can define *truncated subtraction* by

$$\begin{aligned} m \dot{-} 0 &= m \\ m \dot{-} (\text{succ}(n)) &= \text{pred}(m \dot{-} n) \end{aligned}$$

With these definitions and the induction principle, one can prove all the following identities:

- $n \neq 0$ implies $\text{succ}(\text{pred}(n)) = n$
- $0 + n = n$
- $\text{succ}(m) + n = \text{succ}(m + n)$
- $(m + n) + k = m + (n + k)$
- $m + n = n + m$
- $m(n + k) = mn + mk$
- $0 \cdot n = 0$
- $1 \cdot n = n$
- $(mn)k = m(nk)$
- $mn = nm$

We will do the first five here, and leave the remaining ones as exercises.

Proposition. For every natural number n , if $n \neq 0$ then $\text{succ}(\text{pred}(n)) = n$.

Proof. By induction on n . We have ruled out the case where n is 0, so we only need to show that the claim holds for $\text{succ}(n)$. But in that case, we have $\text{succ}(\text{pred}(\text{succ}(n))) = \text{succ}(n)$ by the second defining clause of the predecessor function.

Proposition. For every n , $0 + n = n$.

Proof. By induction on n . We have $0 + 0 = 0$ by the first defining clause for addition. And assuming $0 + n = n$, we have $0 + \text{succ}(n) = \text{succ}(0 + n) = n$, using the second defining clause for addition.

Proposition. For every m and n , $\text{succ}(m) + n = \text{succ}(m + n)$.

Proof. Fix m and use induction on n . Then $n = 0$, we have $\text{succ}(m) + 0 = \text{succ}(m) = \text{succ}(m + 0)$, using the first defining clause for addition. Assuming the claim holds for n , we have

$$\begin{aligned} \text{succ}(m) + \text{succ}(n) &= \text{succ}(\text{succ}(m) + n) \\ &= \text{succ}(\text{succ}(m + n)) \\ &= \text{succ}(m + \text{succ}(n)) \end{aligned}$$

using the inductive hypothesis and the second defining clause for addition.

Proposition. For every m , n , and k , $(m + n) + k = m + (n + k)$.

Proof. By induction on k . The case where $k = 0$ is easy, and in the induction step we have

$$\begin{aligned}(m + n) + \text{succ}(k) &= \text{succ}((m + n) + k) \\ &= \text{succ}(m + (n + k)) \\ &= m + \text{succ}(n + k) \\ &= m + (n + \text{succ}(k))\end{aligned}$$

using the inductive hypothesis and the definition of addition.

Proposition. For every pair of natural numbers m and n , $m + n = n + m$.

Proof. By induction on n . The base case is easy using the second proposition above. In the inductive step, we have

$$\begin{aligned}m + \text{succ}(n) &= \text{succ}(m + n) \\ &= \text{succ}(n + m) \\ &= \text{succ}(n) + m\end{aligned}$$

using the third proposition above.

17.5 Arithmetic on the Natural Numbers

Continuing as in the last section, we can establish all the basic properties of the natural numbers that play a role in day-to-day mathematics. We summarize the main ones here:

$$\begin{aligned}m + n &= n + m && \text{(commutativity of addition)} \\ m + (n + k) &= (m + n) + k && \text{(associativity of addition)} \\ n + 0 &= n && \text{(0 is a neutral element for addition)} \\ n \cdot m &= m \cdot n && \text{(commutativity of multiplication)} \\ m \cdot (n \cdot k) &= (m \cdot n) \cdot k && \text{(associativity of multiplication)} \\ n \cdot 1 &= n && \text{(1 is a neutral element for multiplication)} \\ n \cdot (m + k) &= n \cdot m + n \cdot k && \text{(distributivity)} \\ n \cdot 0 &= 0 && \text{(0 is an absorbing element for multiplication)}\end{aligned}$$

In an ordinary mathematical argument or calculation, they can be used without explicit justification. We also have the following properties:

- $n + 1 \neq 0$
- if $n + k = m + k$ then $n = m$
- if $n \cdot k = m \cdot k$ and $k \neq 0$ then $n = m$

We can define $m \leq n$, “ m is less than or equal to n ,” to mean that there exists a k such that $m + k = n$. If we do that, it is not hard to show that the less-than-or-equal-to relation satisfies all the following properties, for every n , m , and k :

- $n \leq n$ (*reflexivity*)
- if $n \leq m$ and $m \leq k$ then $n \leq k$ (*transitivity*)
- if $n \leq m$ and $m \leq n$ then $n = m$ (*antisymmetry*)
- for all n and m , either $n \leq m$ or $m \leq n$ is true (*totality*)

- if $n \leq m$ then $n + k \leq m + k$
- if $n + k \leq m + k$ then $n \leq m$
- if $n \leq m$ then $nk \leq mk$
- if $m \geq n$ then $m = n$ or $m \geq n + 1$
- $0 \leq n$

Remember from [Chapter 13](#) that the first four items assert that \leq is a linear order. Note that when we write $m \geq n$, we mean $n \leq m$.

As usual, then, we can define $m < n$ to mean that $m \leq n$ and $m \neq n$. In that case, we have that $m \leq n$ holds if and only if $m < n$ or $m = n$.

Proposition. For every m , $m + 1 \not\leq 0$.

Proof. Otherwise, we would have $(m + 1) + k = (m + k) + 1 = 0$ for some k .

In particular, taking $m = 0$, we have $1 \not\leq 0$.

Proposition. We have $m < n$ iff and only if $m + 1 \leq n$.

Proof. Suppose $m < n$. Then $m \leq n$ and $m \neq n$. So there is a k such that $m + k = n$, and since $m \neq n$, we have $k \neq 0$. Then $k = u + 1$ for some u , which means we have $m + (u + 1) = m + 1 + u = n$, so $m \leq n$, as required.

In the other direction, suppose $m + 1 \leq n$. Then $m \leq n$. We also have $m \neq n$, since if $m = n$, we would have $m + 1 \leq m + 0$ and hence $1 \leq 0$, a contradiction.

In a similar way, we can show that $m < n$ if and only if $m \leq n$ and $m \neq n$. In fact, we can demonstrate all of the following from these properties and the properties of \leq :

- $n < n$ is never true (*irreflexivity*)
- if $n < m$ and $m < k$ then $n < k$ (*transitivity*)
- for all n and m , either $n < m$, $n = m$ or $m < n$ is true (*trichotomy*)
- if $n < m$ then $n + k < m + k$
- if $k > 0$ and $n < m$ then $nk < mk$
- if $m > n$ then $m = n + 1$ or $m > n + 1$
- for all n , $n = 0$ or $n > 0$

The first three items mean that $<$ is a strict linear order, and the properties above means that \leq is the associated linear order, in the sense described in [Section 13.1](#).

Proof. We will prove some of these properties using the previous characterization of the less-than relation.

The first property is straightforward: we know $n \leq n + 1$, and if we had $n + 1 \leq n$, we should have $n = n + 1$, a contradiction.

For the second property, assume $n < m$ and $m < k$. Then $n + 1 \leq m \leq m + 1 \leq k$, which implies $n < k$.

For the third, we know that either $n \leq m$ or $m \leq n$. If $m = n$, we are done, and otherwise we have either $n < m$ or $m < n$.

For the fourth, if $n + 1 \leq m$, we have $n + 1 + k = (n + k) + 1 \leq m + k$, as required.

For the fifth, suppose $k > 0$, which is to say, $k \geq 1$. If $n < m$, then $n + 1 \leq m$, and so $nk + 1 \leq nk + k \leq mk$. But this implies $nk < mk$, as required.

The rest of the remaining proofs are left as an exercise to the reader.

Here are some additional properties of $<$ and \leq :

- $n < m$ and $m < n$ cannot both hold (*asymmetry*)
- $n + 1 > n$
- if $n < m$ and $m \leq k$ then $n < k$
- if $n \leq m$ and $m < k$ then $n < k$
- if $m > n$ then $m \geq n + 1$
- if $m \geq n$ then $m + 1 > n$
- if $n + k < m + k$ then $n < m$
- if $nk < mk$ then $k > 0$ and $n < m$

These can be proved from the ones above. Moreover, the collection of principles we have just seen can be used to justify basic facts about the natural numbers, which are again typically taken for granted in informal mathematical arguments.

Proposition. If n and m are natural numbers such that $n + m = 0$, then $n = m = 0$.

Proof. We first prove that $m = 0$. We know that $m = 0$ or $m > 0$. Suppose that $m > 0$. Then $n + m > n + 0 = n$. Since $n \geq 0$, we conclude that $n + m > 0$, which contradicts the fact that $n + m = 0$. Since $m > 0$ leads to a contradiction, we must have $m = 0$.

Now we can easily conclude that $n = 0$, since $n = n + 0 = n + m = 0$. Hence $n = m = 0$.

Proposition. If n is a natural number such that $n < 3$, then $n = 0$, $n = 1$ or $n = 2$.

Proof. In this proof we repeatedly use the property that if $m > n$ then $m = n + 1$ or $m > n + 1$. Since $2 + 1 = 3 > n$, we conclude that either $2 + 1 = n + 1$ or $2 + 1 > n + 1$. In the first case we conclude $n = 2$, and we are done. In the second case we conclude $2 > n$, which implies that either $2 = n + 1$, or $2 > n + 1$. In the first case, we conclude $n = 1$, and we are done. In the second case, we conclude $1 > n$, and appeal one last time to the general principle presented above to conclude that either $1 = n + 1$ or $1 > n + 1$. In the first case, we conclude $n = 0$, and we are once again done. In the second case, we conclude that $0 > n$. This leads to a contradiction, since now $0 > n \geq 0$, hence $0 > 0$, which contradicts the irreflexivity of $>$.

17.6 The Integers

The natural numbers are designed for counting discrete quantities, but they suffer an annoying drawback: it is possible to subtract n from m if n is less than or equal to m , but not if m is greater than n . The set of

integers, \mathbb{Z} , extends the natural numbers with negative values, to make it possible to carry out subtraction in full:

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

We will see in a later chapter that the integers can be extended to the *rational numbers*, the *real numbers*, and the *complex numbers*, each of which serves useful purposes. For dealing with discrete quantities, however, the integers will get us pretty far.

You can think of the integers as consisting of two copies of the natural numbers, a positive one and a negative one, sharing a common zero. Conversely, once we have the integers, you can think of the natural numbers as consisting of the nonnegative integers, that is, the integers that are greater than or equal to 0. Most mathematicians blur the distinction between the two, though we will see that in Lean, for example, the natural numbers and the integers represent two different data types.

Most of the properties of the natural numbers that were enumerated in the last section hold of the integers as well, but not all. For example, it is no longer the case that $n + 1 \neq 0$ for every n , since the claim is false for $n = -1$. For another example, it is not the case that every integer is either equal to 0 or greater than 0, since this fails to hold of the negative integers.

The key property that the integers enjoy, which sets them apart from the natural numbers, is that for every integer n there is a value $-n$ with the property that $n + (-n) = 0$. The value $-n$ is called the *negation* of n . We define subtraction $n - m$ to be $n + (-m)$. For any integer n , we also define the *absolute value* of n , written $|n|$, to be n if $n \geq 0$, and $-n$ otherwise.

We can no longer use proof by induction on the integers, because induction does not cover the negative numbers. But we can use induction to show that a property holds of every nonnegative integer, for example. Moreover, we know that every negative integer is the negation of a positive one. As a result, proofs involving the integers often break down into two cases, where one case covers the nonnegative integers, and the other case covers the negative ones.

17.7 Exercises

1. Write the principle of complete induction using the notation of symbolic logic. Also write the least element principle this way, and use logical manipulations to show that the two are equivalent.
2. Show that for every n , $0^2 + 1^2 + 2^2 + \dots + n^2 = \frac{1}{6}n(n+1)(n+2)$.
3. Show that for every n , $0^3 + 1^3 + \dots + n^3 = \frac{1}{4}n^2(n+1)^2$.
4. Given the definition of the Fibonacci numbers in [Section 17.3](#), prove Cassini's identity: for every n , $F_{n+1}^2 - F_{n+2}F_n = (-1)^n$. Hint: in the induction step, write F_{n+2}^2 as $F_{n+2}(F_{n+1} + F_n)$.
5. Prove $\sum_{i < n} F_{2i+1} = F_{2n}$.
6. Prove the following two identities:
 - $F_{2n+1} = F_{n+1}^2 + F_n^2$
 - $F_{2n+2} = F_{n+2}^2 - F_n^2$

Hint: use induction on n , and prove them both at once. In the induction step, expand $F_{2n+3} = F_{2n+2} + F_{2n+1}$, and similarly for F_{2n+4} . Proving the second equation is especially tricky. Use the inductive hypothesis and the first identity to simplify the left-hand side, and repeatedly unfold the Fibonacci number with the highest index and simplify the equation you need to prove. (When you have worked out a solution, write a clear equational proof, calculating in the “forward” direction.)

7. Prove that every natural number can be written as a sum of *distinct* powers of 2. For this problem, $1 = 2^0$ is counted as power of 2.

8. Let V be a non-empty set of integers such that the following two properties hold:

- If $x, y \in V$, then $x - y \in V$.
- If $x \in V$, then every multiple of x is an element of V .

Prove that there is some $d \in V$, such that V is equal to the set of multiples of d . Hint: use the least element principle.

9. Give an informal but detailed proof that for every natural number n , $1 \cdot n = n$, using a proof by induction, the definition of multiplication, and the theorems proved in [Section 17.4](#).

10. Show that multiplication distributes over addition. In other words, prove that for natural numbers m , n , and k , $m(n + k) = mn + mk$. You should use the definitions of addition and multiplication and facts proved in [Section 17.4](#) (but nothing more).

11. Prove the multiplication is associative, in the same way. You can use any of the facts proved in [Section 17.4](#) and the previous exercise.

12. Prove that multiplication is commutative.

13. Prove $(m^n)^k = m^{nk}$.

14. Following the example in [Section 17.5](#), prove that if n is a natural number and $n < 5$, then n is one of the values 0, 1, 2, 3, or 4.

15. Prove that if n and m are natural numbers and $nm = 1$, then $n = m = 1$, using only properties listed in [Section 17.5](#).

This is tricky. First show that n and m are greater than 0, and hence greater than or equal to 1. Then show that if either one of them is greater than 1, then $nm > 1$.

16. Prove any of the other claims in [Section 17.5](#) that were stated without proof.

17. Prove the following properties of negation and subtraction on the integers, using only the properties of negation and subtraction given in [Section 17.6](#).

- If $n + m = 0$ then $m = -n$.
- $-0 = 0$.
- If $-n = -m$ then $n = m$.
- $m + (n - m) = n$.
- $-(n + m) = -n - m$.
- If $m < n$ then $n - m > 0$.
- If $m < n$ then $-m > -n$.
- $n \cdot (-m) = -nm$.
- $n(m - k) = nm - nk$.
- If $n < m$ then $n - k < m - k$.

18. Suppose you have an infinite chessboard with a natural number written in each square. The value in each square is the average of the values of the four neighboring squares. Prove that all the values on the chessboard are equal.

19. Prove that every natural number can be written as a sum of *distinct non-consecutive* Fibonacci numbers. For example, $22 = 1 + 3 + 5 + 13$ is not allowed, since 3 and 5 are consecutive Fibonacci numbers, but $22 = 1 + 21$ is allowed.

THE NATURAL NUMBERS AND INDUCTION IN LEAN

18.1 Induction and Recursion in Lean

Internally, in Lean, the natural numbers are defined as a type generated inductively from an axiomatically declared `zero` and `succ` operation:

```
inductive nat : Type
| zero : nat
| succ : nat → nat
```

If you click the button that copies this text into the editor in the online version of this textbook, you will see that we wrap it with the phrases `namespace hidden` and `end hidden`. This puts the definition into a new “namespace,” so that the identifiers that are defined are `hidden.nat`, `hidden.nat.zero` and `hidden.nat.succ`, to avoid conflicting with the one that is in the Lean library. Below, we will do that in a number of places where our examples duplicate objects defined in the library. The unicode symbol \mathbb{N} , entered with `\N` or `\nat`, is a synonym for `nat`.

Declaring `nat` as an inductively defined type means that we can define functions by recursion, and prove theorems by induction. For example, these are the first two recursive definitions presented in the last chapter:

```
open nat

def two_pow : ℕ → ℕ
| 0       := 1
| (succ n) := 2 * two_pow n

def fact : ℕ → ℕ
| 0       := 1
| (succ n) := (succ n) * fact n
```

Addition and numerals are defined in such a way that Lean recognizes `succ n` and `n + 1` as essentially the same, so we could instead write these definitions as follows:

```
def two_pow : ℕ → ℕ
| 0       := 1
| (n + 1) := 2 * two_pow n

def fact : ℕ → ℕ
| 0       := 1
| (n + 1) := (n + 1) * fact n
```

If we wanted to define the function m^n , we would do that by fixing m , and writing doing the recursion on the second argument:

```
def pow : ℕ → ℕ → ℕ
| m 0      := 1
| m (n + 1) := pow m n * m
```

In fact, this is how the power function on the natural numbers, `nat.pow`, is defined in Lean's library.

Lean is also smart enough to interpret more complicated forms of recursion, like this one:

```
def fib : ℕ → ℕ
| 0      := 0
| 1      := 1
| (n + 2) := fib (n + 1) + fib n
```

In addition to defining functions by recursion, we can prove theorems by induction. In Lean, each clause of a recursive definition results in a new identity. For example, the two clauses in the definition of `pow` above give rise to the following two theorems:

```
open nat

theorem pow_zero (n : ℕ) : pow n 0 = 1 := rfl
theorem pow_succ (m n : ℕ) : pow m (n+1) = pow m n * m := rfl
```

Lean defines the usual notation for exponentiation:

```
open nat

theorem pow_zero (n : ℕ) : n^0 = 1 := rfl
theorem pow_succ (m n : ℕ) : m^(n+1) = m^n * m := rfl
```

Notice that we could alternatively have used `m * pow m n` in the second clause of the definition of `pow`. Of course, we can prove that the two definitions are equivalent using the commutativity of multiplication, but, using a proof by induction, we can also prove it using only the associativity of multiplication, and the properties $1 * m = m$ and $m * 1 = m$. This is useful, because the power function is also often used in situations where multiplication is not commutative, such as with matrix multiplication. The theorem can be proved in Lean as follows:

```
open nat

theorem pow_succ' (m n : ℕ) : m^(succ n) = m * m^n :=
nat.rec_on n
  (show m^(succ 0) = m * m^0, from calc
    m^(succ 0) = m^0 * m : by rw pow_succ
    ... = 1 * m       : by rw pow_zero
    ... = m           : by rw one_mul
    ... = m * 1       : by rw mul_one
    ... = m * m^0     : by rw pow_zero)
  (assume n,
    assume ih : m^(succ n) = m * m^n,
    show m^(succ (succ n)) = m * m^(succ n), from calc
      m^(succ (succ n)) = m^(succ n) * m : by rw pow_succ
      ... = (m * m^n) * m : by rw ih
```

(continues on next page)

(continued from previous page)

```

... = m * (m^n * m)      : by rw mul_assoc
... = m * m^(succ n)    : by rw pow_succ)

```

This is a typical proof by induction in Lean. It begins with the phrase `nat.rec_on n`, and is followed by the base case and the inductive hypothesis. (The phrase `open nat` allows us to write `pow` instead of `nat.pow`. The proof can be shortened using `rewrite`:

```

open nat

theorem pow_succ' (m n : ℕ) : m^(succ n) = m * (m^n) :=
nat.rec_on n
  (show m^(succ 0) = m * m^0,
   by rw [pow_succ, pow_zero, mul_one, one_mul])
  (assume n,
   assume ih : m^(succ n) = m * m^n,
   show m^(succ (succ n)) = m * m^(succ n),
   by rw [pow_succ, ih, mul_assoc, mul_comm (m^n)])

```

Remember that you can write a `rewrite` proof incrementally, checking the error messages to make sure things are working so far, and to see how far Lean got.

As another example of a proof by induction, here is a proof of the identity $m^{(n + k)} = m^n * m^k$.

```

open nat

theorem pow_add (m n k : ℕ) : m^(n + k) = m^n * m^k :=
nat.rec_on k
  (show m^(n + 0) = m^n * m^0, from calc
   m^(n + 0) = m^n      : by rw add_zero
   ... = m^n * 1      : by rw mul_one
   ... = m^n * m^0    : by rw pow_zero)
  (assume k,
   assume ih : m^(n + k) = m^n * m^k,
   show m^(n + succ k) = m^n * m^(succ k), from calc
   m^(n + succ k) = m^(succ (n + k)) : by rw nat.add_succ
   ... = m^(n + k) * m      : by rw pow_succ
   ... = m^n * m^k * m      : by rw ih
   ... = m^n * (m^k * m)    : by rw mul_assoc
   ... = m^n * m^(succ k)  : by rw pow_succ)

```

Notice the same pattern. This time, we do induction on `k`, and the base case and inductive step are routine. Once again, with a bit of cleverness, we can shorten the proof with `rewrite`:

```

open nat

theorem pow_add (m n k : ℕ) : m^(n + k) = m^n * m^k :=
nat.rec_on k
  (show m^(n + 0) = m^n * m^0,
   by rw [add_zero, pow_zero, mul_one])
  (assume k,
   assume ih : m^(n + k) = m^n * m^k,
   show m^(n + succ k) = m^n * m^(succ k),
   by rw [nat.add_succ, pow_succ, ih, mul_assoc, pow_succ])

```

You should not hesitate to use `calc`, however, to make the proofs more explicit. Remember that you can also use `calc` and `rewrite` together, using `calc` to structure the calculational proof, and using `rewrite` to fill in each justification step.

18.2 Defining the Arithmetic Operations in Lean

In fact, addition and multiplication are defined in Lean essentially as described in [Section 17.4](#). The defining equations for addition hold by reflexivity, but they are also named `add_zero` and `add_succ`:

```
open nat

variables m n : ℕ

example : m + 0 = m := add_zero m
example : m + succ n = succ (m + n) := add_succ m n
```

Similarly, we have the defining equations for the predecessor function and multiplication:

```
open nat

#check @pred_zero
#check @pred_succ
#check @mul_zero
#check @mul_succ
```

Here are the five propositions proved in [Section 17.4](#).

```
theorem succ_pred (n : ℕ) : n ≠ 0 → succ (pred n) = n :=
nat.rec_on n
  (assume H : 0 ≠ 0,
   show succ (pred 0) = 0, from absurd rfl H)
  (assume n,
   assume ih,
   assume H : succ n ≠ 0,
   show succ (pred (succ n)) = succ n,
   by rewrite pred_succ)

theorem zero_add (n : nat) : 0 + n = n :=
nat.rec_on n
  (show 0 + 0 = 0, from rfl)
  (assume n,
   assume ih : 0 + n = n,
   show 0 + succ n = succ n, from
    calc
    0 + succ n = succ (0 + n) : rfl
    ... = succ n : by rw ih)

theorem succ_add (m n : nat) : succ m + n = succ (m + n) :=
nat.rec_on n
  (show succ m + 0 = succ (m + 0), from rfl)
  (assume n,
   assume ih : succ m + n = succ (m + n),
```

(continues on next page)

(continued from previous page)

```

show succ m + succ n = succ (m + succ n), from
  calc
succ m + succ n = succ (succ m + n) : rfl
... = succ (succ (m + n)) : by rw ih
... = succ (m + succ n) : rfl

theorem add_assoc (m n k : nat) : m + n + k = m + (n + k) :=
nat.rec_on k
  (show m + n + 0 = m + (n + 0), by rw [add_zero, add_zero])
  (assume k,
    assume ih : m + n + k = m + (n + k),
    show m + n + succ k = m + (n + (succ k)), from calc
      m + n + succ k = succ (m + n + k) : by rw add_succ
      ... = succ (m + (n + k)) : by rw ih
      ... = m + (n + succ k) : by rw add_succ)

theorem add_comm (m n : nat) : m + n = n + m :=
nat.rec_on n
  (show m + 0 = 0 + m, by rewrite [add_zero, zero_add])
  (assume n,
    assume ih : m + n = n + m,
    show m + succ n = succ n + m, from calc
      m + succ n = succ (m + n) : by rw add_succ
      ... = succ (n + m) : by rw ih
      ... = succ n + m : by rw succ_add)

```

18.3 Exercises

1. Consider the theorems in [Section 17.4](#) and [Section 17.5](#). Formalize as many of them as you can in Lean.

ELEMENTARY NUMBER THEORY

In the last two chapters, we saw that the natural numbers are characterized by the fact that they support *proof by induction* and *definition by recursion*. Moreover, with these components, we can actually define $+$, \times , and $<$ in a suitable axiomatic foundation, and prove that they have the relevant properties. In [Section 17.6](#) we also discussed the integers, which include negative numbers and support the operation of subtraction.

The natural numbers and the integers are the central components of *number theory*, a branch of mathematics dating back to the ancients. In this chapter, we will discuss some of the rudiments of this subject.

19.1 The Quotient-Remainder Theorem

A key property of the integers that we will use here is the quotient-remainder theorem:

Theorem. Let n and m be integers with $m > 0$. Then there are integers q and r satisfying $n = mq + r$ and $0 \leq r < m$.

Proof. First we prove this in the case where n is a natural number, in which case use complete induction on n . Let n be any natural number. If $n < m$, then we can take $q = 0$ and $r = n$, and we indeed have $n = mq + r$ and $0 \leq r < m$. Otherwise, we have $n \geq m$. In this case $n - m$ is a natural number smaller than n . By induction hypothesis, we know that we can find q' and r' such that $n - m = mq' + r'$ and $0 \leq r' < m$. Then we can choose $q = q' + 1$ and $r = r'$, and we obtain $n = mq + r$ and $0 \leq r < m$, as desired.

If n is negative, then $-(n+1)$ is a natural number, hence we can use the previous part for $-(n+1)$ to obtain q' and r' such that $-(n+1) = mq' + r'$ and $0 \leq r' < m$. Now let $q = -(q' + 1)$ and $r = m - r' - 1$. Then we can compute

$$\begin{aligned}mq + r &= -m(q' + 1) + m - r' - 1 \\ &= -(mq' + r') - m + m - 1 \\ &= -(-(n + 1)) - 1 \\ &= n + 1 - 1 \\ &= n.\end{aligned}$$

Also, since $r' \geq 0$ we have $r < m$ and since $r' < m$ we have $r \geq 0$. This completes the proof.

Intuitively, q is the integer *quotient* when you divide n by m and r is the *remainder*. Remember that using the word “the” presupposes that there are unique values meeting that description. That is, in fact, the case:

Proposition. If n and m are as above, $n = mq + r$ and $n = mq' + r'$ with both r and r' less than m , then $q = q'$ and $r = r'$.

Proof. By assumption, we have $mq + r = mq' + r'$. It suffices to show that $q = q'$, because then $mq = mq'$, and hence $r = r'$.

Suppose $q \neq q'$. Then either $q < q'$ or $q' < q$. Suppose without loss of generality that $q < q'$. (The other case is symmetric.) Then $mq < mq'$, so we can subtract mq from both sides of the equality $mq + r = mq' + r'$ to obtain

$$r = mq' + r' - mq = m(q - q') + r'.$$

But since $q' < q$, we have $q - q' \geq 1$, which means

$$m(q - q') + r' \geq m + r' \geq m,$$

which contradicts the fact that $r < m$.

19.2 Divisibility

We can define divisibility on the integers as follows.

Definition. Given two integers m and n , we say that m is a divisor of n , written $m \mid n$, if there exists some integer k such that $m \cdot k = n$. We also say that n is divisible by m or that m divides n . We write $m \nmid n$ to say that m is not a divisor of n .

We can now prove the following:

Theorem. The relation \mid is reflexive and transitive. Also, if $n \mid m$ and $m \mid n$, then $m = \pm n$. This means that restricted to the natural numbers, this relation is a partial order.

Proof. Reflexivity is immediate, because $n \cdot 1 = n$, hence $n \mid n$.

For transitivity, suppose $m \mid n$ and $n \mid r$. Then there are k, ℓ such that $m \cdot k = n$ and $n \cdot \ell = r$. Now we compute

$$\begin{aligned} m \cdot (k \cdot \ell) &= (m \cdot k) \cdot \ell \\ &= n \cdot \ell \\ &= r. \end{aligned}$$

Suppose that n and m are integers such that $n \mid m$ and $m \mid n$. Then there exist k and ℓ such that $n \cdot k = m$ and $m \cdot \ell = n$. We distinguish two cases. If $n = 0$, then we have $m = n \cdot k = 0 = n$, so we are done. If $n \neq 0$, then we use the equations to get $n \cdot k \cdot \ell = m \cdot \ell = n$, and we can cancel n on both sides to get $k \cdot \ell = 1$. We conclude that $k = \ell = \pm 1$, hence we get $m = n \cdot k = \pm n$.

Note that this means that if n and m are both natural numbers, then $n = m$, which means that \mid is antisymmetric, and hence a partial order, on the natural numbers.

See Exercise 1 for some basic properties of divisibility.

An integer is *even* if it is divisible by 2, in other words, n is even if $2 \mid n$. An integer is *odd* if it is not even. Of course, odd numbers are of the form $2k + 1$ for some k , and we can prove this now.

Theorem. If n is an odd integer, then $n = 2k + 1$ for some integer k .

Proof. By the quotient-remainder theorem, we can write $n = 2k + r$ for some integers k and r with $0 \leq r < 2$. The last condition means that $r = 0$ or $r = 1$. In the first case, we have $n = 2k$, hence $2 \mid n$, contradicting that n is odd. So we have $r = 1$, which means that $n = 2k + 1$.

Theorem. Every sequence of k consecutive numbers contains a number divisible by k .

Proof. Denote the largest element of the sequence by n . This means that the sequence is $n - (k - 1), \dots, n - 1, n$. By the quotient-remainder theorem, we have $n = qk + r$ for some integers q and r with $0 \leq r < k$. From these inequalities we conclude that $n - r$ is in our sequence, and $n - r = qk$, hence divisible by k .

Definition. Given two integers m and n such that either $m \neq 0$ or $n \neq 0$, we define the *greatest common divisor* $\gcd(m, n)$ of m and n to be the largest integer d which is both a divisor of m and n , that is $d \mid m$ and $d \mid n$.

This largest integer exists, because there is at least one common divisor, but only finitely many. There is at least one, since 1 is a common divisor of any two integers, and there are finitely many, since a nonzero number has only finitely many divisors.

If $n = m = 0$, then we define $\gcd(0, 0) = 0$.

The greatest common divisor of two numbers is always a natural number, since 1 is always a common divisor of two numbers. As an example, let us compute the greatest common divisor of 6 and 28. The positive divisors of 6 are $\{1, 2, 3, 6\}$ and the positive divisors of 28 are $\{1, 2, 4, 7, 14, 28\}$. The largest number in both these sets is 2, which is the greatest common divisor of 6 and 28.

However, computing the greatest common divisor of two numbers by listing all the divisors of both numbers is a lot of work, so we will now consider a method to compute the greatest common divisor more efficiently.

Lemma. For all integers n , m and k we have $\gcd(n, m) = \gcd(m, n - km)$.

Proof. Let $d = \gcd(n, m)$ and $r = n - km$. If $n = m = 0$, then $d = 0 = \gcd(m, r)$, and we're done.

In the other case we first show that the set of common divisors of n and m is the same as the set of the common divisors of m and r . To see this, let $d' \mid m$ and $d' \mid n$. Then also $d' \mid n - km$ by Exercise 1 below. Hence d' is a common divisor of m and r . On the other hand, if d' is a divisor of m and r , then $d' \mid r + km$, hence $d' \mid n$, hence d' is a common divisor of n and m .

Since the sets of common divisors are the same, the largest element in each set is also the same, hence $\gcd(n, m) = \gcd(m, n - km)$.

Lemma. For all integers n we have $\gcd(n, 0) = |n|$.

Proof. Every number is a divisor of 0, hence the greatest common divisor of n and 0 is just the greatest divisor of n , which is the absolute value of n .

These two lemmas give us a quick way to compute the greatest common divisor of two numbers. This is called the *Euclidean Algorithm*. Suppose we want to compute $\gcd(a, b)$.

- We let $r_0 = a$ and $r_1 = b$.

- Given r_n and r_{n+1} we compute r_{n+2} as the remainder of r_n when divided by r_{n+1} .
- Once $r_n = 0$, we stop, and $\gcd(a, b) = |r_{n-1}|$.

This works, because by the lemmas above, we have $\gcd(r_k, r_{k+1}) = \gcd(r_{k+1}, r_{k+2})$, since $r_{k+2} = r_k - qr_{k+1}$ for some q . Hence if $r_n = 0$ we have

$$\gcd(a, b) = \gcd(r_0, r_1) = \gcd(r_{n-1}, r_n) = \gcd(r_{n-1}, 0) = |r_{n-1}|.$$

For example, suppose we want to compute the greatest common divisor of 1311 and 5757. We compute the following remainders:

$$\begin{aligned} 5757 &= 4 \times 1311 + 513 \\ 1311 &= 2 \times 513 + 285 \\ 513 &= 1 \times 285 + 228 \\ 285 &= 1 \times 228 + 57 \\ 228 &= 4 \times 57 + 0. \end{aligned}$$

Hence $\gcd(1311, 5757) = 57$. This is much quicker than computing all the divisors of both 1311 and 5757.

Here is an important result about greatest common divisors. It is only called a “lemma” for historical reasons.

Theorem (Bézout’s Lemma). Let s and t be integers. Then there are integers a and b such that $as + bt = \gcd(s, t)$.

Proof. We compute $\gcd(s, t)$ by the Euclidean Algorithm given above, and during the algorithm we get the intermediate values r_0, r_1, \dots, r_n where $r_n = 0$. Now by induction on k we prove that we can write $r_k = a_k s + b_k t$ for some integers a_k and b_k . Indeed: $r_0 = 1 \cdot s + 0 \cdot t$ and $r_1 = 0 \cdot s + 1 \cdot t$. Now if we assume that $r_k = a_k s + b_k t$ and $r_{k+1} = a_{k+1} s + b_{k+1} t$, we know that $r_{k+2} = r_k - q \cdot r_{k+1}$, where q is the quotient of r_k when divided by r_{k+1} . These equations together give

$$r_{k+2} = (a_k - qa_{k+1})s + (b_k - qb_{k+1})t.$$

This completes the induction. In particular, $r_{n-1} = a_{n-1}s + b_{n-1}t$, and since $\gcd(s, t) = \pm r_{n-1}$ we can write $\gcd(s, t)$ as $as + bt$ for some a and b .

Corollary. If c is any common divisor of n and m , then $c \mid \gcd(n, m)$.

Proof. By Bézout’s Lemma, there are a and b such that $\gcd(n, m) = an + bm$. Since c divides both n and m , c divides $an + bm$ by Exercise 1 below, and hence also $\gcd(n, m)$.

Of special interest are pairs of integers which have no divisors in common, except 1 and -1 .

Definition. Two integers n and m are *coprime* if $\gcd(n, m) = 1$.

Proposition. Let n , m and k be integers such that n and k are coprime. If $k \mid nm$ then $k \mid m$.

Proof. By Bézout’s Lemma, there are a and b such that $an + bk = 1$. Multiplying by m gives $anm + bkm = m$. Since k divides nm , k divides the left-hand side of the equation, hence $k \mid m$.

19.3 Prime Numbers

In this section we consider properties of prime numbers.

Definition. An integer $p \geq 2$ is called *prime* if the only positive divisors of p are 1 and p . An integer $n \geq 2$ which is not prime is called *composite*.

An equivalent definition of a prime number is a positive number with exactly 2 positive divisors.

Recall from [Chapter 17](#) that every natural number greater than 1 can be written as the product of primes. In particular, every natural number greater than 1 is divisible by some prime number.

We now prove some other properties about prime numbers.

Theorem. There are infinitely many primes.

Proof. Suppose for the sake of contradiction that there are only finitely many primes p_1, p_2, \dots, p_k . Let $n = p_1 \times p_2 \times \dots \times p_k$. Since n is divisible by p_i for all $i \leq k$ we know that $n + 1$ is not divisible by p_i for any i . However, we assumed that these are all primes, contradicting the fact that every number is divisible by a prime number.

Lemma. If n is an integer and p is a prime number, then either n and p are coprime or $p \mid n$.

Proof. Let $d = \gcd(n, p)$. Since d is a positive divisor of p , either $d = 1$ or $d = p$. In the first case, n and p are coprime by definition, and in the second case we have $p \mid n$.

Proposition. If n and m are integers and p is a prime number such that $p \mid nm$ then either $p \mid n$ or $p \mid m$.

Proof. Suppose that $p \nmid n$. By the previous lemma, this means that p and n are coprime. From this we can conclude that $p \mid m$.

The last result in this section captures that the primes are the “building blocks” of the positive integers for multiplication: all other integers can be written as a product of primes in an essentially unique way.

Theorem (Fundamental Theorem of Arithmetic). Let $n > 0$ be an integer. Then there are primes p_1, \dots, p_k such that $n = p_1 \times \dots \times p_k$. Moreover, these primes are unique up to reordering. That means that if there are prime numbers q_1, \dots, q_ℓ such that $q_1 \times \dots \times q_\ell = n$, then the q_i are a reordering of the p_i . To be completely precise, this means that there is a bijection $\sigma : \{1, \dots, k\} \rightarrow \{1, \dots, \ell\}$ such that $q_i = p_{\sigma(i)}$.

Remark. 1 can be written as the product of zero prime numbers. The *empty product* is defined to be 1.

Proof. We have already seen that every number can be written as the product of primes, so we only need to prove the uniqueness up to reordering. Suppose this is not true, and by the least element principle, let n be the smallest positive integers such that n can be written as the product of primes in two ways: $n = p_1 \times \dots \times p_k = q_1 \times \dots \times q_\ell$.

Since 1 can be written as product of primes *only* as empty product, we have $n > 1$, hence $k \geq 1$. Since p_k is prime, we must have $p_k \mid q_j$ for some $j \leq \ell$. By swapping q_j and q_ℓ , we may assume that $j = \ell$. Since q_ℓ is also prime, we have $p_k = q_\ell$.

Now we have $p_1 \times \dots \times p_{k-1} = q_1 \times \dots \times q_{\ell-1}$. This product is smaller than n , but can be written as product of primes in two different ways. But we assumed n was the smallest such number. Contradiction!

19.4 Modular Arithmetic

In the discussion of equivalence relations in Section 13.3 we considered the example of the relation of modular equivalence on the integers. This is sometimes thought of as “clock arithmetic.” Suppose you have a 12-hour clock without a minute hand, so it only has an hour hand which can point to the hours 12, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11 and then it wraps to 12 again. We can do arithmetic with this clock.

- If the hand currently points to 10, then 5 hours later it will point to 3.
- If the hand points to 7, then 23 hours before that, it pointed to 8.
- If the hand points to 9, and we work for a 8 hours, then when we are done the hand will point to 5. If we worked twice as long, starting at 9, the hand will point to 1.

We want to write these statements using mathematical notation, so that we can reason about them more easily. We cannot write $10 + 5 = 3$ for the first expression, because that would be false, so instead we use the notation $10 + 5 \equiv 3 \pmod{12}$. The notation $\pmod{12}$ indicates that we forget about multiples of 12, and we use the “congruence” symbol with three horizontal lines to remind us that these values are not exactly equal, but only equal up to multiples of 12. The other two lines can be formulated as $7 - 23 \equiv 8 \pmod{12}$ and $9 + 2 \cdot 8 \equiv 1 \pmod{12}$.

Here are some more examples:

- $6 + 7 \equiv 1 \pmod{12}$
- $6 \cdot 7 \equiv 42 \equiv 6 \pmod{12}$
- $7 \cdot 5 \equiv 35 \equiv -1 \pmod{12}$

The last example shows that we can use negative numbers as well.

We now give a precise definition.

Definition. For integers a , b and n we say that a and b are *congruent modulo n* if $n \mid a - b$. This is written $a \equiv b \pmod{n}$. The number n is called the *modulus*.

Typically we only use this definition when the modulus n is positive.

Theorem. Congruence modulo n is an equivalence relation.

Proof. We have to show that congruence modulo n is reflexive, symmetric and transitive.

It is reflexive, because $a - a = 0$, so $n \mid a - a$, and hence $a \equiv a \pmod{n}$.

To show that it is symmetric, suppose that $a \equiv b \pmod{n}$. Then by definition, $n \mid a - b$. So $n \mid (-1) \cdot (a - b)$, which means that $n \mid b - a$. This means by definition that $b \equiv a \pmod{n}$.

To show that it is transitive, suppose that $a \equiv b \pmod{n}$ and $b \equiv c \pmod{n}$. Then we have $n \mid a - b$ and $n \mid b - c$. Hence we have $n \mid (a - b) + (b - c)$ which means that $n \mid a - c$. So $a \equiv c \pmod{n}$.

This theorem justifies the “chaining” notation we used above when we wrote $7 \cdot 5 \equiv 35 \equiv -1 \pmod{12}$. Since congruence modulo 12 is transitive, we can now actually conclude that $7 \cdot 5 \equiv -1 \pmod{12}$.

Theorem. Suppose that $a \equiv b \pmod{n}$ and $c \equiv d \pmod{n}$. Then $a + c \equiv b + d \pmod{n}$ and $a \cdot c \equiv b \cdot d \pmod{n}$.

Moreover, if $a \equiv b \pmod{n}$ then $a^k \equiv b^k \pmod{n}$ for all natural numbers k .

Proof. We know that $n \mid a - b$ and $n \mid c - d$. For the first statement, we can calculate that $(a + c) - (b + d) = (a - b) + (c - d)$, so we can conclude that $n \mid (a + c) - (b + d)$ hence that $a + c \equiv b + d \pmod{n}$.

For the second statement, we want to show that $n \mid a \cdot c - b \cdot d$. We can factor $a \cdot c - b \cdot d = (a - b) \cdot c + b \cdot (c - d)$. Now n divides both summands on the right, hence n divides $a \cdot c - b \cdot d$, which means that $a \cdot c \equiv b \cdot d \pmod{n}$.

The last statement follows by induction on k . If $k = 0$, then $1 \equiv 1 \pmod{n}$, and for the induction step, suppose that $a^k \equiv b^k \pmod{n}$, then we have $a^{k+1} = a \cdot a^k \equiv b \cdot b^k = b^{k+1} \pmod{n}$

This theorem is useful for carrying out computations modulo n . Here are some examples.

- Suppose we want to compute $77 \cdot 123$ modulo 12. We know that $77 \equiv 5 \pmod{12}$ and $123 \equiv 3 \pmod{12}$, so $77 \cdot 123 \equiv 5 \cdot 3 \equiv 15 \equiv 3 \pmod{12}$
- Suppose we want to compute $99 \cdot 998$ modulo 10. We know that $99 \equiv -1 \pmod{10}$ and $998 \equiv -2 \pmod{10}$, hence $99 \cdot 998 \equiv (-1) \cdot (-2) \equiv 2 \pmod{10}$.
- Suppose we want to know the last digit of 101^{101} . Notice that the last digit of a number n is congruent to n modulo 10, so we can just compute $101^{101} \equiv 1^{101} \equiv 1 \pmod{10}$. So the last digit of 101^{101} is 1.

Warning. You cannot do all computations you might expect with modular arithmetic:

- Modular arithmetic does not respect division. For example $12 \equiv 16 \pmod{4}$, but we cannot divide both sides of the equation by 2, because $6 \not\equiv 8 \pmod{4}$.
- Exponents also do not respect modular arithmetic. For example $8 \equiv 3 \pmod{5}$, but $2^8 \not\equiv 2^3 \pmod{5}$. To see this: $2^8 = 256 \equiv 1 \pmod{5}$, but $2^3 = 8 \equiv 3 \pmod{5}$.

Recall the quotient-remainder theorem: if $n > 0$, then any integer a can be expressed as $a = nq + r$, where $0 \leq r < n$. In the language of modular arithmetic this means that $a \equiv r \pmod{n}$. So if $n > 0$, then every integer is congruent to a number between 0 and $n - 1$ (inclusive). So there “are only n different numbers” when working modulo n . This can be used to prove many statements about the natural numbers.

Proposition. For every integer k , $k^2 + 1$ is not divisible by 3.

Proof. Translating this problem to modular arithmetic, we have to show that $k^2 + 1 \not\equiv 0 \pmod{3}$ or in other words that $k^2 \not\equiv 2 \pmod{3}$ for all k . By the quotient-remainder theorem, we know that k is either congruent to 0, 1 or 2, modulo 3. In the first case, $k^2 \equiv 0^2 \equiv 0 \pmod{3}$. In the second case, $k^2 \equiv 1^2 \equiv 1 \pmod{3}$, and in the last case we have $k^2 \equiv 2^2 \equiv 4 \equiv 1 \pmod{3}$. In all of those cases, $k^2 \not\equiv 2 \pmod{3}$. So $k^2 + 1$ is never divisible by 3.

Proposition. For all integers a and b , $a^2 + b^2 - 3$ is not divisible by 4.

Proof. We first compute the squares modulo 4. We compute

$$0^2 \equiv 0 \pmod{4}$$

$$1^2 \equiv 1 \pmod{4}$$

$$2^2 \equiv 0 \pmod{4}$$

$$3^2 \equiv 1 \pmod{4}.$$

Since every number is congruent to 0, 1, 2 or 3 modulo 4, we know that every square is congruent to 0 or 1 modulo 4. This means that there are only four possibilities for $a^2 + b^2 \pmod{4}$. It can be congruent to $0 + 0$, $1 + 0$, $0 + 1$ or $0 + 0$. In all those cases, $a^2 + b^2 \not\equiv 3 \pmod{4}$ Hence $4 \nmid a^2 + b^2 - 3$, proving the proposition.

Recall that we warned you about dividing in modular arithmetic. This doesn't always work, but often it does. For example, suppose we want to solve $2n \equiv 1 \pmod{5}$. We cannot solve this by saying that $n \equiv \frac{1}{2} \pmod{5}$, because we cannot work with fractions in modular arithmetic. However, we can still solve it by multiplying both sides with 3. Then we get $6n \equiv 3 \pmod{5}$, and since $6 \equiv 1 \pmod{5}$ we get $n \equiv 3 \pmod{5}$. So instead of dividing by 2 we could multiply by 3 to get the answer. The reason this worked is because $2 \times 3 \equiv 1 \pmod{5}$.

Definition. Let n and a be integers. A *multiplicative inverse of a modulo n* is an integer b such that $ab \equiv 1 \pmod{n}$.

For example, 3 is a multiplicative inverse of 5 modulo 7, since $3 \times 5 \equiv 1 \pmod{7}$. But 2 has no multiplicative inverse modulo 6. Indeed, suppose that $2b \equiv 1 \pmod{6}$, then $6 \mid 2b - 1$. However, $2b - 1$ is odd, and cannot be divisible by an even number. We can use multiplicative inverses to solve equations. If we want to solve $ax \equiv c \pmod{n}$ for x and we know that b is a multiplicative inverse of a , the solution is $x \equiv bc \pmod{n}$ which we can see by multiplying both sides by b .

Lemma Let n and a be integers. a has at most one multiplicative inverse modulo n . That is, if b and b' are both multiplicative inverses of a modulo n , then $b \equiv b' \pmod{n}$.

Proof. Suppose that $ab \equiv 1 \equiv ab' \pmod{n}$. Then we can compute bab' in two ways: $b \equiv b(ab') = (ba)b' \equiv b' \pmod{n}$.

Proposition. Let n and a be integers. a has a multiplicative inverse modulo n if and only if n and a are coprime.

Proof. Suppose b is a multiplicative inverse of a modulo n . Then $n \mid ab - 1$. Let $d = \gcd(a, b)$. Since $d \mid n$ we have $d \mid ab - 1$. But since d is a divisor of ab , we have $d \mid ab - (ab - 1) = 1$. Since $d \geq 0$ we have $d = 1$. Hence n and a are coprime.

On the other hand, suppose that n and a are coprime. By Bézout's Lemma we know that there are integers b and c such that $cn + ba = \gcd(n, a) = 1$. We can rewrite this to $ab - 1 = (-c)n$, hence $n \mid ab - 1$, which means by definition $ab \equiv 1 \pmod{n}$. This means that b is a multiplicative inverse of a modulo n .

Note that if p is a prime number and a is a integer not divisible by p , then a and p are coprime, hence a has a multiplicative inverse.

19.5 Properties of Squares

Mathematicians from ancient times have been interested in the question as to which integers can be written as a sum of two squares. For example, we can write $2 = 1^1 + 1^1$, $5 = 2^2 + 1^2$, $13 = 3^2 + 2^2$. If we make a sufficiently long list of these, an interesting pattern emerges: if two numbers can be written as a sum of two squares, then so can their product. For example, $10 = 5 \cdot 2$, and we can write $10 = 3^2 + 1^2$. Or $65 = 13 \cdot 5$, and we can write $65 = 8^2 + 1^2$.

At first, one might wonder whether this is just a coincidence. The following provides a proof of the fact that it is not.

Theorem. Let x and y be any two integers. If x and y are both sums of two squares, then so is xy .

Proof. Suppose $x = a^2 + b^2$, and suppose $y = c^2 + d^2$. I claim that

$$xy = (ac - bd)^2 + (ad + bc)^2.$$

To show this, notice that on the one hand we have

$$xy = (a^2 + b^2)(c^2 + d^2) = a^2c^2 + a^2d^2 + b^2c^2 + b^2d^2.$$

On the other hand, we have

$$\begin{aligned} (ac - bd)^2 + (ad + bc)^2 &= (a^2c^2 - 2abcd + b^2d^2) + (a^2d^2 + 2abcd + b^2c^2) \\ &= a^2c^2 + b^2d^2 + a^2d^2 + b^2c^2. \end{aligned}$$

Up to the order of summands, the two right-hand sides are the same.

We will now prove that $\sqrt{2}$ is not a fraction of two integers.

Theorem. There are no integers a and b such that $\frac{a}{b} = \sqrt{2}$.

Proof. Suppose that $\frac{a}{b} = \sqrt{2}$ for some integers a and b . By canceling common factors, we may assume that a and b are coprime. By squaring both sides, we get $\frac{a^2}{b^2} = 2$, and multiplying both sides by b^2 gives $a^2 = 2b^2$. Since $2b^2$ is even, we know that a^2 is even, and since odd squares are odd, we conclude that a is even. Hence we can write $a = 2c$ for some integer c . This means that $(2c)^2 = 2b^2$, hence $2c^2 = b^2$. The same reasoning shows that b is even. But we assumed that a and b are coprime, which contradicts the fact that they are both even.

Hence there are no integers a and b such that $\frac{a}{b} = \sqrt{2}$.

19.6 Exercises

- Prove the following properties about divisibility (for any integers a , b and c):
 - If $a \mid b$ and $a \mid c$ then $a \mid b + c$ and $a \mid b - c$.
 - If $a \mid b$ then $a \mid bc$.
 - $a \mid 0$;
 - If $0 \mid a$ then $a = 0$.
 - If $a \neq 0$ then the statements $b \mid c$ and $ab \mid ac$ are equivalent.
 - If $a \mid b$ and $b \neq 0$ then $|a| \leq |b|$.
- Prove that for any integer n , n^2 leaves a remainder of 0 or 1 when you divide it by 4. Conclude that $n^2 + 2$ is never divisible by 4.
- Prove that if n is odd, $n^2 - 1$ is divisible by 8.
- Prove that if m and n are odd, then $m^2 + n^2$ is even but not divisible by 4.
- Say that two integers “have the same parity” if they are both even or both odd. Prove that if m and n are any two integers, then $m + n$ and $m - n$ have the same parity.

6. Write 11160 as product of primes.
7. List all the divisors of 42 and 198, and find the greatest common divisor by looking at the largest number in both lists. Also compute the greatest common divisor of the numbers by the Euclidean Algorithm.
8. Compute $\gcd(15, 55)$, $\gcd(12345, 54321)$ and $\gcd(-77, 110)$
9. Show by induction on n that for every pair of integers x and y , $x - y$ divides $x^n - y^n$. (Hint: in the induction step, write $x^{n+1} - y^{n+1}$ as $x^n(x - y) + x^n y - y^{n+1}$.)
10. Compute $2^{12} \pmod{13}$. Use this to compute $2^{1212004} \pmod{13}$.
11. Find the last digit of 99^{99} . Can you also find the last two digits of this number?
12. Prove that $50^{22} - 22^{50}$ is divisible by 7.
13. Check whether the following multiplicative inverses exist, and if so, find them.
 - the multiplicative inverse of 5 modulo 7
 - the multiplicative inverse of 17 modulo 21
 - the multiplicative inverse of 4 modulo 14
 - the multiplicative inverse of -2 modulo 9
14. Find all integers x such that $75x \equiv 45 \pmod{8}$.
15. Show that for every integer n the number n^4 is congruent to 0 or 1 modulo 5. Hint: to simplify the computation, use that $4^4 \equiv (-1)^4 \pmod{5}$.
16. Prove that the equation $n^4 + m^4 = k^4 + 3$ has no solutions in the integers. (Hint: use the previous exercise.)
17. Suppose p is a prime number such that $p \nmid k$. Show that if $kn \equiv km \pmod{p}$ then $n \equiv m \pmod{p}$.
18. Let n , m and c be given integers. Use Bézout's Lemma to prove that the equation $an + bm = c$ has a solution for integers a and b if and only if $\gcd(n, m) \mid c$.
19. Suppose that $a \mid n$ and $a \mid m$ and let $d = \gcd(n, m)$. Prove that $\gcd(\frac{n}{a}, \frac{m}{a}) = \frac{d}{a}$. Conclude that for any two integers n and m with greatest common divisor d the numbers $\frac{n}{d}$ and $\frac{m}{d}$ are coprime.

COMBINATORICS

Combinatorics is the art of counting without counting. It is a fundamental mathematical task to determine how many things there are in a given collection, and when the collection is large, it can be tedious or infeasible to count the elements individually. Moreover, when the collection is described in terms of a changing parameter (say, a natural number, n), we would like a formula that tells us how the number of objects depends on that parameter. In this chapter we will set up a foundation for achieving this goal, and learn some of the tricks of the trade.

20.1 Finite Sets and Cardinality

It will be helpful, for every natural number n , to have a canonical set of elements of size n . To that end, we will choose the set

$$[n] = \{m \mid m < n\} = \{0, 1, \dots, n-1\}.$$

We used the same notation, $[n]$, to describe equivalence classes with respect to an equivalence relation, but hopefully our intended meaning will always be clear from the context.

A set A of elements is said to be *finite* if there is a bijection from $[n]$ to A for some n . In that case, we would like to say that A has n elements, or that the set A has cardinality n , and write $|A| = n$. But to do so, we need to know that when A is finite, there is a unique n with the property above.

Suppose there are bijections from both $[m]$ and $[n]$ to A . Composing the first bijection with the inverse of the second, we get a bijection from $[m]$ to $[n]$. It seems intuitively clear that this implies $m = n$, but our goal is to prove this from the fundamental properties of sets, functions, and the natural numbers.

So suppose, for the sake of contradiction, $m \neq n$. Without loss of generality, we can assume $m > n$ (why?). In particular, there is an injective function f from $[m]$ to $[n]$. Since $m > n$, $m \geq n+1$, and so we can restrict f to get an injective function from $[n+1]$ to $[n]$. The next theorem shows that this cannot happen.

Theorem. For any natural number n , there is no injective function from $[n+1]$ to $[n]$.

Proof. By induction on n . The theorem is clear when $n = 0$, because $[1] = \{0\}$ and $[0] = \emptyset$. If f were an injective function from $[1]$ to $[0]$, we would have $f(0) \in \emptyset$, which is impossible.

So suppose the claim is true for n , and suppose f is an injective function from $[n+2]$ to $[n+1]$. We consider two cases.

In the first case, suppose n is not in the image of f . Then f maps $[n+2]$ to $[n]$, and restricting the domain, we have an injective function from $[n+1]$ to $[n]$, contradicting the inductive hypothesis.

In the second case, there is some $m < n+2$ such that $f(m) = n$. The idea is to alter f slightly to get an injective function from $[n+1]$ to $[n]$, again contradicting the inductive hypothesis. If $m = n+1$, which is

to say it is the last element of $[n + 2]$ that is mapped to the last element of $[n + 1]$, we can just restrict f to $[n + 1]$. The fact that f was injective implies that all the elements in $[n + 1]$ are mapped to n .

Otherwise, define $f' : [n + 1] \rightarrow [n]$ by

$$f'(i) = \begin{cases} f(i) & \text{if } i \neq m \\ f(n + 1) & \text{if } i = m. \end{cases}$$

In other words, we map m to the value that $n + 1$ was mapped to. Since f is injective, $f(n + 1) \neq f(m)$, and so $f(n + 1) < n$, as required. It is not hard to check that f' is injective, so we have the contradiction we were after.

This theorem is known as the “pigeonhole principle.” It implies that if $n + 1$ pigeons inhabit n holes, then at least one hole has more than one pigeon. The principle implies that for every finite set A , there is a unique n such that there is a bijection from $[n]$ to A , and we can define the cardinality of A to be that n .

20.2 Counting Principles

Here is a basic counting principle.

Theorem. Let A and B be disjoint finite sets. Then $|A \cup B| = |A| + |B|$.

Proof. Suppose $f : [m] \rightarrow A$ and $g : [n] \rightarrow B$ are bijections. Define $h : [m + n] \rightarrow A \cup B$ by

$$h(i) = \begin{cases} f(i) & \text{if } i < m \\ g(i - m) & \text{if } m \leq i < m + n. \end{cases}$$

To see that h is surjective, note that every k in $A \cup B$ can be written as either $k = f(i)$ for some $i \in [m]$ or $k = g(j)$ for some $j \in [n]$. In the first case, $k = f(i) = h(i)$, and in the second case, $k = g(j) = h(m + j)$.

It is not hard to show that h is also injective. Suppose $h(i) = h(j)$. If $h(i)$ is in A , then it is not in the range of g , and so we must have $h(i) = f(i)$ and $h(j) = f(j)$. Then $f(i) = f(j)$, the injectivity of f implies that $i = j$. If $h(i)$ is instead in B , the argument is similar.

The proof only spells out our basic intuitions: if you want to list all of the elements of $A \cup B$, you can list all the elements of A and then all the elements of B . And if A and B have no elements in common, then to count the elements of $A \cup B$, you can count the elements of A and then continue counting the elements of B . Once you are comfortable translating the intuitive argument into a precise mathematical proof (and mathematicians generally are), you can use the more intuitive descriptions (and mathematicians generally do).

Here is another basic counting principle:

Theorem. Let A and B be finite sets. Then $|A \times B| = |A| \cdot |B|$.

Notice that this time we are counting the number of ordered pairs (a, b) with $a \in A$ and $b \in B$. The exercises ask you to give a detailed proof of this theorem. There are at least two ways to go about it. The first is to start with bijections $f : [m] \rightarrow A$ and $g : [n] \rightarrow B$ and describe an explicit bijection $h : [m \cdot n] \rightarrow A \times B$. The

second is to fix m , say, and use induction on n and the previous counting principle. Notice that if U and V are any sets and w is not in V , we have

$$U \times (V \cup \{w\}) = (U \times V) \cup (U \times \{w\}),$$

and the two sets in this union are disjoint.

Just as we have notions of union $\bigcup_{i \in I} A_i$ and intersection $\bigcap_{i \in I} A_i$ for indexed families of sets, it is useful to have a notion of a product $\prod_{i \in I} A_i$. We can think of an element a of this product as a function which, for each element $i \in I$, returns an element $a_i \in A_i$. For example, when $I = \{1, 2, 3\}$, an element of $\prod_{i \in I} A_i$ is just a triple a_1, a_2, a_3 with $a_1 \in A_1$, $a_2 \in A_2$, and $a_3 \in A_3$. This is essentially the same as $A_1 \times A_2 \times A_3$, up to the fiddly details as to whether we represent a triple as a function or with iterated pairing $(a_1, (a_2, a_3))$.

Theorem. Let I be a finite index set, and let $(A_i)_{i \in I}$ be a family of finite sets. Then:

- If each pair of sets A_i, A_j are disjoint, then $|\bigcup_{i \in I} A_i| = \sum_{i \in I} |A_i|$.
- $|\prod_{i \in I} A_i| = \prod_{i \in I} |A_i|$.

Proof. By induction on $|I|$, using the previous counting principles.

We can already use these principles to carry out basic calculations.

Example. The dessert menu at a restaurant has four flavors of ice cream, two kinds of cake, and three kinds of pie. How many dessert choices are there?

Solution. $4 + 2 + 3 = 9$, the cardinality of the union of the three disjoint sets.

Example. The menu at a diner has 6 choices of appetizers, 7 choices of entrée, and 5 choices of dessert. How many choices of three-course dinners are there?

Solution. A three-course dinner is a triple consisting of an appetizer, an entrée, and a dessert. There are therefore $6 \cdot 7 \cdot 5 = 210$ options.

A special case of the previous counting principles arises when all the sets have the same size. If I has cardinality k and each A_i has cardinality n , then the cardinality of $\bigcup_{i \in I} A_i$ is $k \cdot n$ if the sets are pairwise disjoint, and the cardinality of $\prod_{i \in I} A_i$ is n^k .

Example. A deck of playing cards has four suits (diamonds, hearts, spades, and clubs) and 13 cards in each suit, for a total of $4 \cdot 13 = 52$.

Example. A binary string of length n is a sequence of n many 0's and 1's. We can think of this as an element of

$$\{0, 1\}^n = \prod_{i < n} \{0, 1\},$$

so there are 2^n many binary strings of length n .

There is another counting principle that is almost too obvious to mention: if A is a finite set and there is a bijection between A and B , then B is also finite, and $|A| = |B|$.

Example. Consider the power set of $[n]$, that is, the collection of all subsets of $\{0, 1, 2, \dots, n - 1\}$. There is a one-to-one correspondence between subsets and binary strings of length n , where element i of the string is 1 if i is in the set and 0 otherwise. As a result, we have $|\mathcal{P}([n])| = 2^n$.

20.3 Ordered Selections

Let S be a finite set, which we will think of as being a set of options, such as items on a menu or books that can be selected from a shelf. We now turn to a family of problems in combinatorics that involves making repeated selections from that set of options. In each case, there are finitely many selections, and the order counts: there is a first choice, a second one, a third one, and so on.

In the first variant of the problem, you are allowed to repeat a choice. For example, if you are choosing 3 flavors from a list of 31 ice cream flavors, you can choose “chocolate, vanilla, chocolate.” This is known as *ordered selection with repetition*. If you are making k choices from among n options in S , such a selection is essentially a tuple $(a_0, a_1, \dots, a_{k-1})$, where each a_i is one of the n elements in S . In other words, the set of ways of making k selections from S with repetition is the set S^k , and we have seen in the last section that if S has cardinality n , the set S^k has cardinality n^k .

Theorem. Let S be a set of n elements. Then the number of ways of making k selections from S with repetition allowed is n^k .

Example. How many three-letter strings (like “xyz,” “qqa,” ...) can be formed using the twenty-six letters of the alphabet?

Solution. We have to make three selections from a set of 26 elements, for a total of $26^3 = 17,576$ possibilities.

Suppose instead we wish to make k ordered selections, but we are not allowed to repeat ourselves. This would arise, from example, if a museum had 26 paintings in its storeroom, and has to select three of them to put on display, ordered from left to right along a wall. There are 26 choices for the first position. Once we have made that choice, 25 remain for the second position, and then 24 remain for the third. So it seems clear that there are $26 \cdot 25 \cdot 24$ arrangements overall.

Let us try to frame the problem in mathematical terms. We can think of an ordered selection of k elements from a set S without repetition as being an *injective function* f from $[k]$ to S . The element $f(0)$ is the first choice; $f(1)$ is the second choice, which has to be distinct from $f(0)$; $f(2)$ is the third choice, which has to be distinct from $f(0)$ and $f(1)$; and so on.

Theorem. Let A and B be finite sets, with $|A| = k$ and $|B| = n$, and $k \leq n$. The number of injective functions from A to B is $n \cdot (n - 1) \cdot \dots \cdot (n - k + 1)$.

Proof. Using induction on k , we will show that for every A , B , and $n \geq k$, the claim holds. When $k = 0$ there is only one injective function, namely the function with empty domain. Suppose A has cardinality $k + 1$, let a_0 be any element of A . Then any injective function from A to B can be obtained by choosing an element b_0 for the image of a_0 , and then choosing an injective function from $A \setminus \{a_0\}$ to $B \setminus \{b_0\}$. There are n choices of b_0 , and since $|A \setminus \{a_0\}| = n - 1$ and $|B \setminus \{b_0\}| = k - 1$, there are $(n - 1) \cdot \dots \cdot (n - k + 1)$ choices of the injective function, by the inductive hypothesis.

Theorem. Let S be a finite set, with $|S| = n$. Then the number of ways of making k selections from S without repetition allowed is $n \cdot (n - 1) \cdot \dots \cdot (n - k + 1)$.

Proof. This is just a restatement of the previous theorem, where $A = [k]$ and $B = S$.

If A is a finite set, a bijection f from A to A is also called a *permutation* of A . The previous theorem shows that if $|A| = n$ then the number of permutations of A is $n \cdot (n - 1) \cdot \dots \cdot 1$. This quantity comes up so often that it has a name, *n factorial*, and a special notation, $n!$. If we think of the elements of A listed in some order, a permutation of A is essentially an ordered selection of n elements from A without repetition: we choose where to map the first element, then the second element, and so on. It is a useful convention to take $0!$ to be equal to 1.

The more general case where we are choosing only k elements from a set A is called a k -permutation of A . The theorem above says that the number of k -permutations of an n -element set is equal to $n!/(n - k)!$, because if you expand the numerator and denominator into products and cancel, you get exactly the $n \cdot (n - 1) \cdot \dots \cdot (n - k + 1)$. This number is often denoted $P(n, k)$ or P_k^n , or some similar variant. So we have $P(n, k) = n!/(n - k)!$. Notice that the expression on the right side of the equality provides an efficient way of writing the value of $P(n, k)$, but an inefficient way of calculating it.

20.4 Combinations and Binomial Coefficients

In the last section, we calculated the number of ways in which a museum could arrange three paintings along a wall, chosen from among 26 paintings in its storeroom. By the final observation in the previous section, we can write this number as $26!/23!$.

Suppose now we want to calculate the number of ways that a museum can choose three paintings from its storeroom to put on display, where we do not care about the order. In other words, if a , b , and c are paintings, we do not want to distinguish between choosing a then b then c and choosing c then b then a . When we were arranging paintings along all wall, it made sense to consider these two different arrangements, but if we only care about the *set* of elements we end up with at the end, the order that we choose them does not matter.

The problem is that each set of three paintings will be counted multiple times. In fact, each one will be counted six times: there are $3! = 6$ permutations of the set $\{a, b, c\}$, for example. So to count the number of outcomes we simply need to divide by 6. In other words, the number we want is $\frac{26!}{3! \cdot 23!}$.

There is nothing special about the numbers 26 and 3. The same formula holds for what we will call *unordered selections of k elements from a set of n elements*, or *k -combinations from an n -element set*. Our goal is once again to describe the situation in precise mathematical terms, at which point we will be able to state the formula as a theorem.

In fact, describing the situation in more mathematical terms is quite easy to do. If S is a set of n elements, an unordered selection of k elements from S is just a subset of S that has cardinality k .

Theorem. Let S be any set with cardinality n , and let $k \leq n$. Then the number of subsets of S of cardinality k is $\frac{n!}{k!(n-k)!}$.

Proof. Let U be the set of unordered selections of k elements from S , let V be the set of permutations of $[k]$, and let W be the set of *ordered* selections of k elements from S . There is a bijection between $U \times V$ and W , as follows. Suppose we assign to every k -element subset $\{a_0, \dots, a_{k-1}\}$ of S some way of listing the elements, as shown. Then given any such set and any permutation f of $[k]$, we get an ordered the ordered selection $(a_{f(0)}, a_{f(1)}, \dots, a_{f(k-1)})$. Any ordered selection arises from such a subset and a suitable permutation, so the

mapping is surjective. And a different set or a different permutation results in a different ordered selection, so the mapping is injective.

By the counting principles, we have

$$P(n, k) = |W| = |U \times V| = |U| \cdot |V| = |U| \cdot k!,$$

so we have $|U| = P(n, k)/k! = \frac{n!}{k!(n-k)!}$.

Example. Someone is going on vacation and wants to choose three outfits from ten in their closet to pack in their suitcase. How many choices do they have?

Solution. $\frac{10!}{3!7!} = \frac{10 \cdot 9 \cdot 8}{3 \cdot 2 \cdot 1} = 120$.

The number of unordered selections of k elements from a set of size n , or, equivalently, the number of k -combinations from an n -element set, is typically denoted by $\binom{n}{k}$, $C(n, k)$, C_k^n , or something similar. We will use the first notation, because it is most common. Notice that $\binom{n}{0} = 1$ for every n ; this makes sense, because there is exactly one subset of any n -element set of cardinality 0.

Here is one important property of this function.

Theorem. For every n and $k \leq n$, we have $\binom{n}{k} = \binom{n}{n-k}$.

Proof. This is an easy calculation:

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

But it is also easy to see from the combinatorial interpretation: choosing k outfits from n to take on vacation is the same task as choosing $n - k$ outfits to leave home.

Here is another important property.

Theorem. For every n and k , if $k + 1 \leq n$, then

$$\binom{n+1}{k+1} = \binom{n}{k+1} + \binom{n}{k}.$$

Proof. One way to understand this theorem is in terms of the combinatorial interpretation. Suppose you want to choose $k + 1$ outfits out of $n + 1$. Set aside one outfit, say, the blue one. Then you have two choices: you can either choose $k + 1$ outfits from the remaining ones, with $\binom{n}{k+1}$ possibilities; or you can take the blue one, and choose k outfits from the remaining ones.

The theorem can also be proved by direct calculation. We can express the left-hand side of the equation as follows:

$$\begin{aligned} \binom{n+1}{k+1} &= \frac{(n+1)!}{(k+1)!((n+1)-(k+1))!} \\ &= \frac{(n+1)!}{(k+1)!(n-k)!}. \end{aligned}$$

we would like to know how many ways there are of making k choices, where we can choose elements of S repeatedly, but we only care about the number of times each element was chosen, and not the order. We have the following:

The number of unordered selections of k elements from an n -element set, with repetition, is $\binom{n+k-1}{k}$.

A proof of this is outlined in the exercises.

20.5 The Inclusion-Exclusion Principle

Let A and B be any two subsets of some domain, U . Then $A = A \setminus B \cup (A \cap B)$, and the two sets in the union are disjoint, so we have $|A| = |A \setminus B| + |A \cap B|$. This means $|A \setminus B| = |A| - |A \cap B|$. Intuitively, this makes sense: we can count the elements of $A \setminus B$ by counting the elements in A , and then subtracting the number of elements that are in both A and B .

Similarly, we have $A \cup B = A \cup (B \setminus A)$, and the two sets on the right-hand side of this equation are disjoint, so we have

$$|A \cup B| = |A| + |B \setminus A| = |A| + |B| - |A \cap B|.$$

If we draw a Venn diagram, this makes sense: to count the elements in $A \cup B$, we can add the number of elements in A to the number of elements in B , but then we have to subtract the number of elements of both.

What happen when there are three sets? To compute $|A \cup B \cup C|$, we can start by adding the number of elements in each, and then subtracting the number of elements of $|A \cap B|$, $|A \cap C|$, and $|B \cap C|$, each of which have been double-counted. But thinking about the Venn diagram should help us realize that then we have over-corrected: each element of $A \cap B \cap C$ was counted three times in the original sum, and the subtracted three times. So we need to add them back in:

$$|A \cup B \cup C| = |A| + |B| + |C| - |A \cap B| - |A \cap C| - |B \cap C| + |A \cap B \cap C|.$$

This generalizes to any number of sets. To state the general result, suppose the sets are numbered A_0, \dots, A_{n-1} . For each nonempty subset I of $\{0, \dots, n-1\}$, consider $\bigcap_{i \in I} A_i$. If $|I|$ is odd (that is, equal to 1, 3, 5, ...) we want to add the cardinality of the intersection; if it is even we want to subtract it. This recipe is expressed compactly by the following formula:

$$\left| \bigcup_{i < n} A_i \right| = \sum_{\emptyset \neq I \subseteq [n]} (-1)^{|I|} \left| \bigcap_{i \in I} A_i \right|.$$

You are invited to try proving this as an exercise, if you are ambitious. The following example illustrates its use:

Example. Among a group of college Freshmen, 30 are taking Logic, 25 are taking History, and 20 are taking French. Moreover, 11 are taking Logic and History, 10 are taking Logic and French, 7 are taking History and French, and 3 are taking all three. How many students are taking at least one of the three classes?

Solution. Letting L , H , and F denote the sets of students taking Logic, History, and French, respectively, we have

$$|L \cup H \cup F| = 30 + 25 + 20 - 11 - 10 - 7 + 3 = 50.$$

20.6 Exercises

- Suppose that, at a party, every two people either know each other or don't. In other words, " x knows y " is symmetric. Also, let us ignore the complex question of whether we always know ourselves by restricting attention to the relation between distinct people; in other words, for this problem, take " x knows y " to be irreflexive as well.

Use the pigeonhole principle (and an additional insight) to show that there must be two people who know exactly the same number of people.

- Show that in any set of $n + 1$ integers, two of them are equivalent modulo n .
- Spell out in detail a proof of the second counting principle in [Section 20.2](#).
- An ice cream parlor has 31 flavors of ice cream.
 - Determine how many three-flavor ice-cream cones are possible, if we care about the order and repetitions are allowed. (So choosing chocolate-chocolate-vanilla scoops, from bottom to top, is different from choosing chocolate-vanilla-chocolate.)
 - Determine how many three flavor ice-cream cones are possible, if we care about the order, but repetitions are not allowed.
 - Determine how many three flavor ice-cream cones are possible, if we don't care about the order, but repetitions are not allowed.
- A club of 10 people has to elect a president, vice president, and secretary. How many possibilities are there:
 - if no person can hold more than one office?
 - if anyone can hold any number of those offices?
 - if anyone can hold up to two offices?
 - if the president cannot hold another office, but the vice president and secretary may or may not be the same person?
- How many 7 digit phone numbers are there, if any 7 digits can be used? How many are there is the first digit cannot be 0?
- In a class of 20 kindergarten students, two are twins. How many ways are there of lining up the students, so that the twins are standing together?
- A woman has 8 murder mysteries sitting on her shelf, and wants to take three of them on a vacation. How many ways can she do this?
- In poker, a "full house" is a hand with three of one rank and two of another (for example, three kings and two fives). Determine the number of full houses in poker.
- We saw in [Section 20.4](#) that

$$\binom{n+1}{k+1} = \binom{n}{k+1} + \binom{n}{k}.$$

Replacing $k + 1$ by k , whenever $1 \leq k \leq n$, we have

$$\binom{n+1}{k} = \binom{n}{k} + \binom{n}{k-1}.$$

Use this to show, by induction on n , that for every $k \leq n$, that if S is any set of n elements, $\binom{n}{k}$ is the number of subsets of S with k elements.

11. How many distinct arrangements are there of the letters in the word MISSISSIPPI?

(Hint: this is tricky. First, suppose all the S's, I's, and P's were painted different colors. Then determine how many distinct arrangements of the letters there would be. In the absence of distinguishing colors, determine how many times each configuration appeared in the first count, and divide by that number.

12. Prove the inclusion exclusion principle.

13. Use the inclusion exclusion principle to determine the number of integers less than 100 that are divisible by 2, 3, or 5.

14. Show that the number of *unordered* selections of k elements from an n -element set is $\binom{n+k-1}{k}$.

Hint: consider $[n]$. We need to choose some number i_0 of 0's, some number i_1 of 1's, and so on, so that $i_0 + i_1 + \dots + i_{n-1} = k$. Suppose we assign to each such tuple a the following binary sequence: we write down i_0 0's, then a 1, then i_1 0's, then a 1, then i_2 0's, and so on. The result is a binary sequence of length $n + k - 1$ with exactly k 1's, and such binary sequence arises from a unique tuple in this way.

THE REAL NUMBERS

21.1 The Number Systems

We have already come across some of the fundamental number systems: the natural numbers, \mathbb{N} , the integers, \mathbb{Z} , and the rationals, \mathbb{Q} . In a sense, each subsequent element of the list was designed to remedy defects in the previous system. We can subtract any integer from any other integer and end up with another integer, and we can divide any rational number by a nonzero rational number and end up with a rational number.

The integers satisfy all of the following properties:

- Addition is associative and commutative.
- There is an additive identity, 0, and every element x has an additive inverse, $-x$.
- Multiplication is associative and commutative.
- There is a multiplicative identity, 1.
- Multiplication distributes over addition: for every x , y , and z , we have $x(y + z) = xy + xz$.
- The ordering \leq is a total order.
- For any elements x , y , and z , if $x \leq y$ then $x + z \leq y + z$.
- For any elements x and y , if $0 \leq x$ and $0 \leq y$ then $0 \leq xy$.

The first five clauses say that with \times , $+$, 0, and 1, the integers form a *commutative ring*, and the last three say that together with \leq , the structure is an *ordered ring*. The natural numbers lack additive inverses, so they satisfy a slightly weaker set of axioms that make them an *ordered semiring*. On the other hand, the rational numbers also form an ordered ring, satisfying the following additional property:

- Every nonzero element has a multiplicative inverse, x^{-1} .

This makes them an instance of an *ordered field*.

It is worth knowing that once we have the natural numbers, it is possible to *construct* the integers and rational numbers, using set-theoretic constructions you have already seen. For example, we can take an integer to be a pair (i, n) of natural numbers where i is either 0 or 1, with the intention that $(0, n)$ represents the positive integer n , and $(1, n)$ represents the negative integer $-(n + 1)$. (We use $-(n + 1)$ instead of $-n$ to avoid having two representations of 0.) With this definition, the integers are simply $\{0, 1\} \times \mathbb{N}$. We can then go on to define the operations of addition and multiplication, the additive inverse, and the order relation, and prove they have the desired properties.

This construction has the side effect that the natural numbers themselves are not integers; for example, we have to distinguish between the natural number 2 and the integer 2. This is the case in Lean. In ordinary mathematics, it is common to think of the natural numbers as a subset of the integers. Once we construct the integers, however, we can throw away the old version of the natural numbers, and afterwards identify the natural numbers as nonnegative integers.

We can do the same for the rationals, defining them to be the set of pairs (a, b) in $\mathbb{Z} \times \mathbb{N}$, where either $a = 0$ and $b = 1$, or $b > 0$ and a and b have no common divisor (other than 1 and -1). The idea is that (a, b) represents a/b . With this definition, the rationals are really a subset of $\mathbb{Z} \times \mathbb{N}$, and we can then define all the operations accordingly.

In the next section, we will define a more sophisticated approach, one which will scale to a definition of the real numbers. And in a later chapter, we will show how to construct the natural numbers from the axioms of set theory. This shows that we can construct all the number systems from the bottom up.

But first, let us pause for a moment to consider why the real numbers are needed. We have seen that 2 has no rational square root. This means, in a sense, that there is a “gap” in the rationals: there are rationals whose squares are arbitrarily close to 2, but there is no rational x with the property that $x^2 = 2$. But it seems intuitively clear that there should be some *number* with that property: $\sqrt{2}$ is the length of the diagonal of a square with side length 1. Similarly, π , the area of a circle with radius 1, is missing from the rationals. These are the kinds of defects that the real numbers are designed to repair.

You may be used to thinking of real numbers as (potentially) infinite decimals: for example, $\sqrt{2} = 1.41421356\dots$ and $\pi = 3.14159265\dots$. A central goal of this chapter is to make the “...” precise. The idea is that we can take an infinite decimal to represent a sequence of rational approximations. For example, we can approximate the square root of 2 with the sequence 1, 1.4, 1.41, 1.414, \dots . We would like to define $\sqrt{2}$ to be the “limit” of that sequence, but we have seen that the sequence does not have a limit in the rationals. So we have to construct new objects, the real numbers, to serve that purpose.

In fact, we will define the real numbers, more or less, to *be* such sequences of rational approximations. But we will have to deal with the fact that, for example, there are *lots* of ways of approximating the square root of two. For example, we can just as well approach it from above, 2, 1.5, 1.42, \dots , or by oscillating above and below. The next section will show us how to “glue” all these sequences together and treat them as a single object.

21.2 Quotient Constructions

Let A be any set, and let \equiv be any equivalence relation on A . Recall from [Section 13.3](#) that we can assign to every element a of A the equivalence class $[a]$, where $b \in [a]$ means $b \equiv a$. This assignment has the property that for every a and b , $a \equiv b$ if and only if $[a] = [b]$.

Given any set A and equivalence relation \equiv , define A/\equiv to be the set $\{[a] \mid a \in A\}$ of *equivalence classes* of A modulo \equiv . This set is called “ A modulo \equiv ,” or the *quotient* of A by \equiv . You can think of this as the set A where equivalent elements are “glued together” to make a coarser set.

For example, if we consider the integers \mathbb{Z} with \equiv denoting equivalence modulo 5 (as in [Section 19.4](#)), then \mathbb{Z}/\equiv is just $\{[0], [1], [2], [3], [4]\}$. We can define addition on \mathbb{Z}/\equiv by $[a] + [b] = [a + b]$. For this definition to make sense, it is important to know that the right-hand side does not depend on which representatives of $[a]$ and $[b]$ we choose. In other words, we need to know that whenever $[a] = [a']$ and $[b] = [b']$, then $[a + b] = [a' + b']$. This, in turn, is equivalent to saying that if $a \equiv a'$ and $b \equiv b'$, then $a + b \equiv a' + b'$. In other words, we require that the operation of addition *respects* the equivalence relation, and we saw in [Section 19.4](#) that this is in fact the case.

This general strategy for transferring a function defined on a set to a function defined on a quotient of that set is given by the following theorem.

Theorem. Let A and B be any sets, let \equiv be any equivalence relation defined on A , and let $f : A \rightarrow B$. Suppose f respects the equivalence relation, which is to say, for every a and a' in A , if $a \equiv a'$, then $f(a) = f(a')$. Then there is a unique function $\bar{f} : A/\equiv \rightarrow B$, defined by $\bar{f}([a]) = f(a)$ for every a in A .

Proof. We have defined the value of \bar{f} on an equivalence class x by writing $x = [a]$, and setting $\bar{f}(x) = f(a)$. In other words, we say that $\bar{f}(x) = y$ if and only if there is an a such that $x = [a]$, and $f(a) = y$. What is dubious about the definition is that, a priori, it might depend on how we express x in that form; in other words, we need to show that there is a *unique* y meeting this description. Specifically, we need to know that if $x = [a] = [a']$, then $f(a) = f(a')$. But since $[a] = [a']$ is equivalent to $a \equiv a'$, this amounts to saying that f respects the equivalence relation, which is exactly what we have assumed.

Mathematicians often “define” \bar{f} by the equation $\bar{f}([a]) = f(a)$, and then express the proof above as a proof that “ \bar{f} is well defined.” This is confusing. What they really mean is what the theorem says, namely, that there is a unique function meeting that description.

To construct the integers, start with $\mathbb{N} \times \mathbb{N}$. Think of the pair of natural numbers (m, n) as representing $m - n$, where the subtraction takes place in the integers (which we haven’t constructed yet!). For example, both $(2, 5)$ and $(6, 9)$ represent the integer -3 . Intuitively, the pairs (m, n) and (m', n') will represent the same integer when $m - n = m' - n'$, but we cannot say this yet, because we have not yet defined the appropriate notion of subtraction. But the equation is equivalent to $m + n' = m' + n$, and *this* makes sense with addition on the natural numbers.

Definition. Define the relation \equiv on $\mathbb{N} \times \mathbb{N}$ by $(m, n) \equiv (m', n')$ if and only if $m + n' = m' + n$.

Proposition. \equiv is an equivalence relation.

Proof. For reflexivity, it is clear that $(m, n) \equiv (m, n)$, since $m + n = m + n$.

For symmetry, suppose $(m, n) \equiv (m', n')$. This means $m + n' = m' + n$. But the symmetry of equality implies $(m', n') \equiv (m, n)$, as required.

For transitivity, suppose $(m, n) \equiv (m', n')$, and $(m', n') \equiv (m'', n'')$. Then we have $m + n' = m' + n$ and $m' + n'' = n' + m''$. Adding these equations, we get

$$m + n' + m' + n'' = m' + n + n' + m''.$$

Subtracting $m' + n'$ from both sides, we get $m + n'' = n + m''$, which is equivalent to $(m, n) \equiv (m'', n'')$, as required.

We can now define the integers to be $\mathbb{N} \times \mathbb{N} / \equiv$. How should we define addition? If $[(m, n)]$ represents $m - n$, and $[(u, v)]$ represents $u - v$, then $[(m, n)] + [(u, v)]$ should represent $(m + u) - (n + v)$. Thus, it makes sense to define $[(m, n)] + [(u, v)]$ to be $[(m + u) - (n + v)]$. For this to work, we need to know that the operation which sends (m, n) and (u, v) to $(m + u, n + v)$ respects the equivalence relation.

Proposition. If $(m, n) \equiv (m', n')$ and $(u, v) \equiv (u', v')$, then $(m + u, n + v) \equiv (m' + u', n' + v')$.

Proof. The first equivalence means $m + n' = m' + n$, and the second means $u + v' = u' + v$. Adding the two equations, we get $(m + u) + (n' + v') \equiv (m' + u') + (n + v)$, which is exactly the same as saying $(m + u, n + v) \equiv (m' + u', n' + v')$.

Every natural number n can be represented by the integer $[(n, 0)]$, and, in particular, 0 is represented by $[(0, 0)]$. Moreover, if $[(m, n)]$ is any integer, we can define its negation to be $[(n, m)]$, since $[(m, n)] + [(n, m)] = [(m + n, n + m)] = [(0, 0)]$, since $(m + n, n + m) \equiv (0, 0)$. In short, we have “invented” the negative numbers!

We could go on this way to define multiplication and the ordering on the integers, and prove that they have the desired properties. We could also carry out a similar construction for the rational numbers. Here, we

would start with the set $\mathbb{Z} \times \mathbb{Z}^{>0}$, where $\mathbb{Z}^{>0}$ denotes the strictly positive integers. The idea, of course, is that (a, b) represents (a/b) . With that in mind, it makes sense to define $(a, b) \equiv (c, d)$ if $ad = bc$. We could go on to define addition, multiplication, and the ordering there, too. The details are tedious, however, and not very illuminating. So we turn, instead, to a construction of the real numbers.

21.3 Constructing the Real Numbers

The problem we face is that the sequence $1, 1.4, 1.41, 1.414, 1.4142, \dots$ of rational numbers seems to approach a value that *would* be the square root of 2, but there is no rational number that can play that role. The next definition captures the notion that this sequence of numbers “seems to approach a value,” without referring to a value that it is approaching.

Definition. A sequence of rational numbers $(q_i)_{i \in \mathbb{N}}$ is *Cauchy* if for every rational number $\varepsilon > 0$, there is some natural number $N \in \mathbb{N}$ such that for all $i, j \geq N$, we have that $|q_i - q_j| < \varepsilon$.

Roughly speaking, a Cauchy sequence is one where the elements become arbitrarily close, not just to their successors but to all following elements. It is common in mathematics to use ε to represent a quantity that is intended to denote something small; you should read the phrase “for every $\varepsilon > 0$ ” as saying “no matter how small ε is.” So a sequence is Cauchy if, for any $\varepsilon > 0$, no matter how small, there is some point N , beyond which the elements stay within a distance of ε of one another.

Cauchy sequences can be used to describe these gaps in the rationals, but, as noted above, many Cauchy sequences can be used to describe the same gap. At this stage, it is slightly misleading to say that they “approach the same point,” since there is no rational point that they approach; a more precise statement is that the sequences eventually become arbitrarily close.

Definition. Two Cauchy sequences $p = (p_i)_{i \in \mathbb{N}}$ and $q = (q_i)_{i \in \mathbb{N}}$ are *equivalent* if for every rational number $\varepsilon > 0$, there is some natural number $N \in \mathbb{N}$ such that for all $i \geq N$, we have that $|p_i - q_i| < \varepsilon$. We will write $p \equiv q$ to express that p is equivalent to q .

Proposition. \equiv is an equivalence relation on Cauchy sequences.

Proof. Reflexivity and symmetry are easy, so let us prove transitivity. Suppose $(p_i) \equiv (q_i)$ and $(q_i) \equiv (r_i)$. We want to show that the sequence (q_i) is equivalent to (r_i) . So, given any $\varepsilon > 0$, choose N_0 large enough such that for every $i \geq N_0$, $|p_i - q_i| < \varepsilon/2$. Choose another number, N_1 , so that for every $i \geq N_1$, $|q_i - r_i| < \varepsilon/2$. Let $N = \max(N_0, N_1)$. Then for every $i \geq N$, we have

$$|p_i - r_i| = |(p_i - q_i) + (q_i - r_i)| < |p_i - q_i| + |q_i - r_i| \leq \varepsilon/2 + \varepsilon/2 = \varepsilon,$$

as required.

Notice that the proof uses the *triangle inequality*, which states for any rational numbers a and b , $|a + b| \leq |a| + |b|$. If we define $|a|$ to be the maximum of a and $-a$, the triangle inequality in fact holds for any ordered ring:

Theorem. Let a and b be elements of any ordered ring. Then $|a + b| \leq |a| + |b|$.

Proof. By the definition of absolute value, it suffices to show that $a + b \leq |a| + |b|$ and $-(a + b) \leq |a| + |b|$. The first claim follows from the fact that $a \leq |a|$ and $b \leq |b|$. For the second claim, we similarly have $-a \leq |a|$ and $-b \leq |b|$, so $-(a + b) = -a + -b \leq |a| + |b|$.

In the theorem above, if we let $a = x - y$ and $b = y - z$, we get $|x - z| \leq |x - y| + |y - z|$. The fact that $|x - y|$ represents the distance between x and y on the number line explains the name: for any three “points” x , y , and z , the distance from x to z can’t be any greater than the distance from x to y plus the distance from y to z .

We now let A be the set of Cauchy sequences of rationals, and define the real numbers, \mathbb{R} , to be A/\equiv . In other words, the real numbers are the set of Cauchy sequence of rationals, modulo the equivalence relation we just defined.

Having the set \mathbb{R} by itself is not enough: we also would like to know how to add, subtract, multiply, and divide real numbers. As with the integers, we need to define operations on the underlying set, and then show that they respect the equivalence relation. For example, we will say how to add Cauchy sequences of rationals, and then show that if $p_1 \equiv p_2$ and $q_1 \equiv q_2$, then $p_1 + q_1 \equiv p_2 + q_2$. We can then lift this definition to \mathbb{R} by defining $[p] + [q]$ to be $[p + q]$.

Luckily, it is easy to define addition, subtraction, and multiplication on Cauchy sequences. If $p = (p_i)_{i \in \mathbb{N}}$ and $q = (q_i)_{i \in \mathbb{N}}$ are Cauchy sequences, let $p + q = (p_i + q_i)_{i \in \mathbb{N}}$, and similarly for subtraction and multiplication. It is trickier to show that these sequences are Cauchy themselves, and to show that the operations have the appropriate algebraic properties. We ask you to prove some of these properties in the exercises.

We can identify each rational number q with the constant Cauchy sequence q, q, q, \dots , so the real numbers include all the rationals. The next step is to abstract away the details of the particular construction we have chosen, so that henceforth we can work with the real numbers abstractly, and no longer think of them as given by equivalence classes of Cauchy sequences of rationals.

21.4 The Completeness of the Real Numbers

We constructed the real numbers to fill in the gaps in the rationals. How do we know that we have got them all? Perhaps we need to construct even more numbers, using Cauchy sequences of reals? The next theorem tells us that, on the contrary, there is no need to extend the reals any further in this way.

Definition. Let r be a real number. A sequence $(r_i)_{i \in \mathbb{N}}$ of real numbers *converges* to r if, for every $\varepsilon > 0$, there is an N such that for every $i \geq N$, $|r_i - r| < \varepsilon$.

Definition. A sequence $(r_i)_{i \in \mathbb{N}}$ *converges* if it converges to some r .

Theorem. Every Cauchy sequence of real numbers converges.

The statement of the theorem is often expressed by saying that the real numbers are *complete*. Roughly, it says that everywhere you look for a real number, you are bound to find one. Here is a similar principle.

Definition. An element $u \in \mathbb{R}$ is said to be an *upper bound* to a subset $S \subseteq \mathbb{R}$ if everything in S is less than or equal to u . S is said to be *bounded* if there is an upper bound to S . An element u is said to be a *least upper bound* to S if it is an upper bound to S , and nothing smaller than u is an upper bound to S .

Theorem. Let S be a bounded, nonempty subset of \mathbb{R} . Then S has a least upper bound.

The rational numbers do not have this property: if we set $S = \{x \in \mathbb{Q} \mid x^2 < 2\}$, then the rational number 2 is an upper bound for S , but S has no least upper bound in \mathbb{Q} .

It is a fundamental theorem that the real numbers are characterized exactly by the property that they are a complete ordered field, such that every real number r is less than or equal to some natural number N . Any two models that meet these requirements must behave in exactly the same way, at least insofar as the constants 0 and 1, the operations $+$ and $*$, and the relation \leq are concerned. This fact is extremely powerful because it allows us to avoid thinking about the Cauchy sequence construction in normal mathematics. Once we have shown that our construction meets these requirements, we can take \mathbb{R} to be “the” unique complete totally ordered field and ignore any implementation details. We are also free to implement \mathbb{R} in any way we choose, and as long as it meets this interface, and as long as they do not refer to the underlying representations, any theorems we prove about the reals will hold equally well for all constructions.

21.5 An Alternative Construction

Many sources use an alternative construction of the reals, taking them instead to be *Dedekind cuts*. A Dedekind cut is an ordered pair (A, B) of sets of rational numbers with the following properties:

- Every rational number q is in either A or B .
- Each $a \in A$ is less than every $b \in B$.
- There is no greatest element of A .
- A and B are both nonempty.

The first two properties show why we call this pair a “cut.” The set A contains all of the rational numbers to the left of some mark on the number line, and B all of the points to the right. The third property tells us something about what happens exactly at that mark. But there are two possibilities: either B has a least element, or it doesn’t. Picturing the situation where A has no greatest element and B has no least element may be tricky, but consider the example $A = \{x \in \mathbb{Q} \mid x^2 < 2\}$ and $B = \{x \in \mathbb{Q} \mid x^2 > 2\}$. There is no rational number q such that $q^2 = 2$, but there are rational numbers on either side that are arbitrarily close; thus neither A nor B contains an endpoint.

We can define \mathbb{R} to be the set of Dedekind cuts. A Dedekind cut (A, B) corresponds to a rational number q if q is the least element of B , and to an irrational number if B has no least element. It is straightforward to define addition on \mathbb{R} :

$$(A_1, B_1) + (A_2, B_2) = (\{a_1 + a_2 \mid a_1 \in A_1, a_2 \in A_2\}, \{b_1 + b_2 \mid b_1 \in B_1, b_2 \in B_2\}).$$

Some authors prefer this construction to the Cauchy sequence construction because it avoids taking the quotient of a set, and thus removes the complication of showing that arithmetic operations respect equivalence. Others prefer Cauchy sequences since they provide a clearer notion of approximation: if a real number r is given by a Cauchy sequence $(q_i)_{i \in \mathbb{N}}$, then an arbitrarily close rational approximation of r is given by q_N for a sufficiently large N .

For most mathematicians most of the time, though, the difference is immaterial. Both constructions create complete linear ordered fields, and in a certain sense, they create the *same* complete linear ordered field. Strictly speaking, the set of Cauchy reals is not equal to the set of Dedekind reals, since one consists of equivalence classes of rational Cauchy sequences and one consists of pairs of sets of rationals. But there is a bijection between the two sets that preserves the field properties. That is, there is a bijection f from the Cauchy reals to the Dedekind reals such that

- $f(0) = 0$
- $f(1) = 1$
- $f(x + y) = f(x) + f(y)$

- $f(x \cdot y) = f(x) \cdot f(y)$
- $f(-x) = -f(x)$
- $f(x^{-1}) = f(x)^{-1}$
- $f(x) \leq f(y) \iff x \leq y$

We say that the two constructions are *isomorphic*, and that the function f is an *isomorphism*. Since we often only care about the real numbers in regard to their status as a complete ordered field, and the two constructions are indistinguishable as ordered fields, it makes no difference which construction is used.

21.6 Exercises

1. Show that addition for the integers, as defined in Section 21.2, is commutative and associative.
2. Show from the construction of the integers in Section 21.2 that $a + 0 = a$ for every integer a .
3. Define subtraction for the integers by $a - b = a + (-b)$, and show that $a - b + b = a$ for every pair of integers a and b .
4. Define multiplication for the integers, by first defining it on the underlying representation and then showing that the operation respects the equivalence relation.
5. Show that every Cauchy sequence is bounded: that is, if $(q_i)_{i \in \mathbb{N}}$ is Cauchy, there is some rational M such that $|q_i| \leq M$ for all i . Hint: try letting $\varepsilon = 1$.
6. Let $p = (p_i)_{i \in \mathbb{N}}$ and $q = (q_i)_{i \in \mathbb{N}}$ be Cauchy sequences. Define $p + q = (p_i + q_i)_{i \in \mathbb{N}}$ and $pq = (p_i q_i)_{i \in \mathbb{N}}$.
 - a. Show that $p + q$ is Cauchy. That is, for arbitrary $\varepsilon > 0$, show that there exists an N such that for all $i, j \geq N$, $|(p_i + q_i) - (p_j + q_j)| < \varepsilon$.
 - b. Show that pq is Cauchy. In addition to the triangle inequality, you will find the previous exercise useful.
7. These two parts show that addition of Cauchy sequences respects equivalence.
 - a. Show that if p, p', q are Cauchy sequences and $p \equiv p'$, then $p + q \equiv p' + q$.
 - b. Using the first part of this problem, show that if p, p', q, q' are Cauchy sequences, $p \equiv p'$, and $q \equiv q'$, then $p + q \equiv p' + q'$. You can use the fact that addition on the real numbers is commutative.
8. Show that if (A_1, B_1) and (A_2, B_2) are Dedekind cuts, then $(A_1, B_1) + (A_2, B_2)$ is also a Dedekind cut.

22.1 Equinumerosity

Remember that in [Chapter 20](#) we defined, for each natural number n , the set $[n] = \{0, 1, \dots, n - 1\}$. We then said that a set A is *finite* if there is a bijection between A and $[n]$ for some n . A set is said to be *infinite* if it is not finite.

If A and B are two finite sets, then they have the same cardinality if and only if there is a bijection between them. It turns out that the same notion of “having the same cardinality” makes sense even if A and B are not finite.

Definition. Two sets A and B are said to be *equinumerous*, written $A \approx B$, if there is a bijection between them. Equivalently, we say that A and B *have the same cardinality*.

At this stage, saying that A and B have the same cardinality may sound strange, because it is not clear that there is any object, “the cardinality of A ,” that they both “have.” It turns out that, in set-theoretic foundations, there are certain objects—generalizations of the natural numbers—that one can use to measure the size of an infinite set. There are known as the “cardinal numbers” or “cardinals.” But they are irrelevant to our purposes here. For the rest of this chapter, when we say that A and B have the same cardinality, we mean neither more nor less than the fact that there is a bijection between them.

The following theorem says, essentially, that equinumerosity is an equivalence relation. (The caveat is that so far we have spoke only of relations between sets, and the collection of all sets is not itself a set.)

Proposition. Let A , B , and C be any sets.

- $A \approx A$.
 - If $A \approx B$, then $B \approx A$.
 - If $A \approx B$ and $B \approx C$ then $A \approx C$.
-

The proof is left as an exercise.

22.2 Countably Infinite Sets

The set of natural numbers, \mathbb{N} , is a prototypical example of an infinite set. To see that it is infinite, suppose, on the other hand, that it is finite. This means that there is a bijection f between \mathbb{N} and $[n]$ for some natural number n . We can restrict to the subset $[n + 1]$ of \mathbb{N} , and thereby obtain an injective map from $[n + 1]$ to $[n]$. But this violates the pigeonhole principle, proved in Chapter 20.

Definition. A set is said to be *countably infinite* if it is equinumerous with \mathbb{N} . A set is said to be *countable* if it is finite or countably infinite.

Since the identity map $id(x) = x$ is a bijection on any set, every set is equinumerous with itself, and thus \mathbb{N} itself is countably infinite.

The term “countably infinite” is meant to be evocative. Suppose A is a countable set. By definition, there is a bijection $f : \mathbb{N} \rightarrow A$. So A has a “first” element $f(0)$, a “second” element $f(1)$, a “third” element $f(2)$, and so on. Since f is a bijection, for every element a of A , a is the n th element enumerated in this way, for a unique value of n . That is, each element of A is “counted” at some finite stage.

With this definition in hand, it is natural to wonder which of our favorite sets are countable. Is the set of integers \mathbb{Z} countable? How about the set of rationals \mathbb{Q} , or the set of reals \mathbb{R} ? At this point, you should reflect on the logical form of the statement “ A is countable,” and think about what is required to show that a set A does or does not have this property.

Theorem. The set of integers, \mathbb{Z} , is countable.

Proof. We need to show that there exists a bijection between \mathbb{N} and \mathbb{Z} . Define $f : \mathbb{N} \rightarrow \mathbb{Z}$ as follows:

$$f(n) = \begin{cases} n/2 & \text{if } n \text{ is even} \\ -(n+1)/2 & \text{if } n \text{ is odd.} \end{cases}$$

We claim that f is a bijection. To see that it is injective, suppose $f(m) = f(n)$. If $f(m)$ (and hence also $f(n)$) is nonnegative, then m and n are even, in which case $m/2 = n/2$ implies $m = n$. Otherwise, m and n are odd, and again $-(m+1)/2 = -(n+1)/2$ implies $m = n$.

To see that f is surjective, suppose a is any integer. If a is nonnegative, then $a = f(2a)$. If a is strictly negative, then $2a - 1$ is also strictly negative, and hence $-(2a - 1)$ is an odd natural number. In that case, it is not hard to check that $a = f(-(2a - 1))$.

We will now build up an arsenal of theorems that we can use to show that various sets are countable.

Theorem. A set A is countable if and only if A is empty or there is a surjective function $f : \mathbb{N} \rightarrow A$.

Proof. For the forward direction, suppose A is countable. Then it is either finite or countably infinite. If A is countably infinite, there is a bijection from \mathbb{N} to A , and we are done. Suppose, then, that A is finite. If A is empty, we are done. Otherwise, for some n , there is a bijection $f : [n] \rightarrow A$, with $n \geq 1$. Define a function $g : \mathbb{N} \rightarrow A$ as follows:

$$g(i) = \begin{cases} f(i) & \text{if } i < n \\ f(0) & \text{otherwise.} \end{cases}$$

In other words, g enumerates the elements of A by using f first, and then repeating the element $f(0)$. Clearly f is surjective, as required.

In the other direction, if A is finite, then it is countable, and we are done. So suppose A is not finite. Then it is not empty, and so there is a surjective function $f : \mathbb{N} \rightarrow A$. We need to turn f into a *bijective* function. The problem is that f may not be injective, which is to say, elements in A may be enumerated more than once. The solution is to define a function, g , which eliminates all the duplicates. The idea is that g should enumerate the elements $f(0), f(1), f(2), \dots$, but skip over the ones that have already been enumerated.

To be precise, the function g is defined recursively as follows: $g(0) = f(0)$, and for every i , $g(i + 1) = f(j)$, where j is the least natural number such that $f(j)$ is not among $\{g(0), g(1), g(2), \dots, g(i)\}$. The assumption that A is infinite and f is surjective guarantees that some such j always exists.

We only need to check that g is a bijection. By definition, for every i , $g(i + 1)$ is different from $g(0), \dots, g(i)$. This implies that g is injective. But we can also show by induction that for every i , $\{g(0), \dots, g(i)\} \supseteq \{f(0), \dots, f(i)\}$. Since f is surjective, g is too.

In a manner similar to the way we proved that the integers are countable, we can prove the following:

Theorem. If A and B are countably infinite, then so is $A \cup B$.

Proof. Suppose $f : \mathbb{N} \rightarrow A$ and $g : \mathbb{N} \rightarrow B$ are surjective. Then we can define a function $h : \mathbb{N} \rightarrow A \cup B$:

$$h(n) = \begin{cases} f(n/2) & \text{if } n \text{ is even} \\ g((n-1)/2) & \text{if } n \text{ is odd.} \end{cases}$$

It is not hard to show that h is surjective.

Intuitively, if $A = \{f(0), f(1), f(2), \dots\}$ and $B = \{g(0), g(1), g(2), \dots\}$, then we can enumerate $A \cup B$ as $\{f(0), g(0), f(1), g(1), f(2), g(2), \dots\}$.

The next two theorems are also helpful. The first says that to show that a set B is countable, it is enough to “cover” it with a surjective function from a countable set. The second says that to show that a set A is countable, then it is enough to embed it in a countable set.

Theorem. If A is countable and $f : A \rightarrow B$ is surjective, then B is countable.

Proof. If A is countable, then there is a surjective function $g : \mathbb{N} \rightarrow A$, and $f \circ g$ is a surjective function from $\mathbb{N} \rightarrow B$.

Theorem. If B is countable and $f : A \rightarrow B$ is injective, then A is countable.

Proof. Assuming $f : A \rightarrow B$ is injective, it has a left inverse, $g : B \rightarrow A$. Since g has a right inverse, f , we know that g is surjective, and we can apply the previous theorem.

Corollary. If B is countable and $A \subseteq B$, then A is countable.

Proof. The function $f : A \rightarrow B$ defined by $f(x) = x$ is injective.

Remember that $\mathbb{N} \times \mathbb{N}$ is the set of ordered pairs (i, j) where i and j are natural numbers.

Theorem. $\mathbb{N} \times \mathbb{N}$ is countable.

Proof. Enumerate the elements as follows:

$$(0, 0), (1, 0), (0, 1), (2, 0), (1, 1), (1, 2), (3, 0), (2, 1), (1, 2), (0, 3), \dots$$

If you think of the pairs as coordinates in the x - y plane, the pairs are enumerated along diagonals: first the diagonal with pairs whose elements sum to 0, then the diagonal with pairs whose elements sum to 1, and so on. This is often called a “dovetailing” argument, because if you imagine drawing a line that weaves back and forth through the pairs enumerated this way, it will be analogous to the a carpenter’s practice of using a dovetail to join two pieces of wood. (And that term, in turn, comes from the similarity to a dove’s tail.)

As far as proofs go, the informal description above and the associated diagram are perfectly compelling. It is possible to describe a bijection between $\mathbb{N} \times \mathbb{N}$ explicitly, however, in algebraic terms. You are asked to do this in the exercises.

The previous theorem has a number of interesting consequences.

Theorem. If A and B are countable, then so is $A \times B$.

Proof. If p is any element of $\mathbb{N} \times \mathbb{N}$, write p_0 and p_1 to denote the two components. Let $f : \mathbb{N} \rightarrow \mathbb{N} \times \mathbb{N}$ be a surjection, as guaranteed by the previous theorem. Suppose $g : \mathbb{N} \rightarrow A$ and $h : \mathbb{N} \rightarrow B$ be surjective. Then the function $k(i) = (g(f(i)_0), h(f(i)_1))$ is a surjective function from \mathbb{N} to $A \times B$.

Theorem. The set of rational numbers, \mathbb{Q} , is countable.

Proof. By the previous theorem, we know that $\mathbb{Z} \times \mathbb{Z}$ is countable. Define $f : \mathbb{Z} \times \mathbb{Z} \rightarrow \mathbb{Q}$ by

$$f(i, j) = \begin{cases} i/j & \text{if } j \neq 0 \\ 0 & \text{otherwise.} \end{cases}$$

Since every element of \mathbb{Q} can be written as i/j for some i and j in \mathbb{Z} , f is surjective.

Theorem. Suppose that A is countable. For each n , the set A^n is countable.

Proof. Remember that we can identify the set of n -tuples of elements from A with $A \times \dots \times A$, where there are n copies of A in the product. The result follows using induction on n .

Theorem. Let $(A_i)_{i \in \mathbb{N}}$ be a family of sets indexed by the natural numbers, and suppose that each A_i is countable. Then $\bigcup_i A_i$ is countable.

Proof. Suppose for each i , f_i is a surjective function from \mathbb{N} to A_i . Then the function $g(i, j) = f_i(j)$ is a surjective function from $\mathbb{N} \times \mathbb{N}$ to $\bigcup_i A_i$.

Theorem. Suppose that A is countable. Then the set of finite sequences of elements of A is countable.

Proof. The set of finite sequences of elements of A is equal to $\bigcup_i A^i$, and we can apply the previous two theorems.

Notice that the set of all alphanumeric characters and punctuation (say, represented as the set of all ASCII characters) is finite. Together with the last theorem, this implies that there are only countably many sentences in the English language (and, indeed, any language in which sentences are represented by finite sequences of symbols, chosen from any countable stock).

At this stage, it might seem as though everything is countable. In the next section, we will see that this is not the case: the set of real numbers, \mathbb{R} , is not countable, and if A is any set (finite or infinite), the powerset of A , $\mathcal{P}(A)$, is not equinumerous with A .

22.3 Cantor's Theorem

A set A is *uncountable* if it is not countable. Our goal is to prove the following theorem, due to Georg Cantor.

Theorem. The set of real numbers is uncountable.

Proof. Remember that $[0, 1]$ denotes the closed interval $\{r \in \mathbb{R} \mid 0 \leq r \leq 1\}$. It suffices to show that there is no surjective function $f : \mathbb{N} \rightarrow [0, 1]$, since if \mathbb{R} were countable, $[0, 1]$ would be countable too.

Recall that every real number $r \in [0, 1]$ has a decimal expansion of the form $r = 0.r_0r_1r_2r_3r_4\dots$, where each r_i is a digit in $\{0, 1, \dots, 9\}$. More formally, we can write $r = \sum_{i=0}^{\infty} \frac{r_i}{10^i}$ for each $r \in \mathbb{R}$ with $0 \leq r \leq 1$.

(Notice that 1 can be written 0.9999... In general every other rational number in $[0, 1]$ will have two representations of this form; for example, $0.5 = 0.5000\dots = 0.4999\dots$. For concreteness, for these numbers we can choose the representation that ends with zeros.)

As a result, we can write

- $f(0) = 0.r_0^0r_1^0r_2^0r_3^0r_4^0\dots$
- $f(1) = 0.r_0^1r_1^1r_2^1r_3^1r_4^1\dots$
- $f(2) = 0.r_0^2r_1^2r_2^2r_3^2r_4^2\dots$
- $f(3) = 0.r_0^3r_1^3r_2^3r_3^3r_4^3\dots$
- $f(4) = 0.r_0^4r_1^4r_2^4r_3^4r_4^4\dots$
- ...

(We use superscripts, r^i , to denote the digits of $f(i)$. The superscripts do not mean the “ i th power.”)

Our goal is to show that f is not surjective. To that end, define a new sequence of digits $(r_i)_{i \in \mathbb{N}}$ by

$$r_i = \begin{cases} 7 & \text{if } r_i^i \neq 7 \\ 3 & \text{otherwise.} \end{cases}$$

Then define the real number $r = 0.r_0r_1r_2r_3\dots$. Then, for each i , r differs from $f(i)$ in the i th digit. But this means that for every i , $f(i) \neq r$. Since r is not in the range of f , we see that f is not surjective. Since f was arbitrary, there is no surjective function from \mathbb{N} to $[0, 1]$.

(We chose the digits 3 and 7 only to avoid 0 and 9, to avoid the case where, for example, $f(0) = 0.5000\dots$ and $r = 0.4999\dots$. Since there are no zeros or nines in r , since the i th digit of r differs from $f(i)$, it really is a different real number.)

This remarkable proof is known as a “diagonalization argument.” We are trying to construct a real number with a certain property, namely, that it is not in the range of f . We make a table of digits, in which the rows represent infinitely many constraints we have to satisfy (namely, that for each i , $f(i) \neq r$), and the columns represent opportunities to satisfy that constraint (namely, by choosing the i th digit of r appropriately). Then we complete the construction by stepping along the diagonal, using the i th opportunity to satisfy the i th constraint. This technique is used often in logic and computability theory.

The following provides another example of an uncountable set.

Theorem. The power set of the natural numbers, $\mathcal{P}(\mathbb{N})$, is uncountable.

Proof. Let $f : \mathbb{N} \rightarrow \mathcal{P}(\mathbb{N})$ be any function. Once again, our goal is to show that f is not surjective. Let S be the set of natural numbers, defined as follows:

$$S = \{n \in \mathbb{N} \mid n \notin f(n)\}.$$

In words, for every natural number, n , n is in S if and only if it is not in $f(n)$. Then clearly for every n , $f(n) \neq S$. So f is not surjective.

We can also view this as a diagonalization argument: draw a table with rows and columns indexed by the natural numbers, where the entry in the i th row and j th column is “yes” if j is an element of $f(i)$, and “no” otherwise. The set S is constructed by switching “yes” and “no” entries along the diagonal.

In fact, exactly the same argument yields the following:

Theorem. For every set A , there is no surjective function from A to $\mathcal{P}(A)$.

Proof. As above, if f is any function from A to $\mathcal{P}(A)$, the set $S = \{a \in A \mid a \notin f(a)\}$ is not in the range of f .

This shows that there is an endless hierarchy of infinities. For example, in the sequence $\mathbb{N}, \mathcal{P}(\mathbb{N}), \mathcal{P}(\mathcal{P}(\mathbb{N})), \dots$, there is an injective function mapping each set into the next, but no surjective function. The union of all those sets is even larger still, and then we can take the power set of *that*, and so on. Set theorists are still today investigating the structure within this hierarchy.

22.4 An Alternative Definition of Finiteness

One thing that distinguishes the infinite from the finite is that an infinite set can have the same size as a proper subset of itself. For example, the natural numbers, the set of even numbers, and the set of perfect squares are all equinumerous, even though the latter two are strictly contained among the natural numbers.

In the nineteenth century, the mathematician Richard Dedekind used this curious property to *define* what it means to be finite. We can show that his definition is equivalent to ours, but the proof requires the axiom of choice.

Definition. A set is *Dedekind infinite* if A is equinumerous with a proper subset of itself, and *Dedekind finite* otherwise.

Theorem. A set is Dedekind infinite if and only if it is infinite.

Proof. Suppose A is Dedekind infinite. We need to show it is not finite; suppose, to the contrary, it is bijective with $[n]$ for some n . Composing bijections, we have that $[n]$ is bijective with a proper subset of itself. This means that there is an injective function f from $[n]$ to a proper subset of n . Modifying f , we can get an injective function from $[n]$ into $[n - 1]$, contradicting the pigeonhole principle.

Suppose, on the other hand, that A is infinite. We need to show that there is an injective function f from A to a proper subset of itself (because then f is a bijection between A and the range of f). Choose a sequence of distinct element a_0, a_1, a_2, \dots of A . Let f map each a_i to a_{i+1} , but leave every other element of A fixed. Then f is injective, but a_0 is not in the range of f , as required.

22.5 The Cantor-Bernstein Theorem

Saying that A and B are equinumerous means, intuitively, that A and B have the same size. There is also a natural way of saying that A is not larger than B :

Definition. For two sets A and B , we say the cardinality of A is less than or equal to the cardinality of B , written $A \preceq B$, when there is an injection $f : A \rightarrow B$.

As an exercise, we ask you to show that \preceq is a *preorder*, which is to say, it is reflexive and transitive. Here is a natural question: does $A \preceq B$ and $B \preceq A$ imply $A \approx B$? In other words, assuming there are injective functions $f : A \rightarrow B$ and $g : B \rightarrow A$, is there necessarily a bijection from A to B ?

The answer is “yes,” but the proof is tricky. The result is known as the *Cantor-Bernstein Theorem*, and we state it without proof.

Theorem. For any sets A and B , if $A \preceq B$ and $B \preceq A$, then $A \approx B$.

22.6 Exercises

1. Show that equinumerosity is reflexive, symmetric, and transitive.
2. Show that the function $f(x) = x/(1 - x)$ is a bijection between the interval $[0, 1)$ and $\mathbb{R}^{\geq 0}$.
3. Show that the $g(x) = x/(1 - |x|)$ gives a bijection between $(-1, 1)$ and \mathbb{R} .
4. Define a function $J : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ by $J(i, j) = \frac{(i+j)(i+j+1)}{2} + i$. This goal of this problem is to show that J is a bijection from $\mathbb{N} \times \mathbb{N}$ to \mathbb{N} .
 - a. Draw a picture indicating which pairs are sent to $0, 1, 2, \dots$
 - b. Let $n = i + j$. Show that $J(i, j)$ is equal the number of pairs (u, v) such that either $u + v < n$, or $u + v = n$ and $u < i$. (Use the fact that $1 + 2 + \dots + n = n(n + 1)/2$.)
 - c. Conclude that J is surjective: to find i and j such that $J(i, j) = k$, it suffices to find the largest n such that $n(n + 1)/2 \leq k$, let $i = k - n(n + 1)/2$, and let $j = n - i$.
 - d. Conclude that J is injective: if $J(i, j) = J(i', j')$, let $n = i + j$ and $n' = i' + j'$. Argue that $n = n'$, and so $i = i'$ and $j = j'$.
5. Let S be the set of functions from \mathbb{N} to $\{0, 1\}$. Use a diagonal argument to show that S is uncountable. (Notice that you can think of a function $f : \mathbb{N} \rightarrow \{0, 1\}$ as an infinite sequence of 0's and 1's, given by $f(0), f(1), f(2), \dots$. So, given a function $F(n)$ which, for each natural number n , returns an infinite sequence of 0's and 1's, you need to find a sequence that is not in the image of F .)
6. If f and g are functions from \mathbb{N} to \mathbb{N} , say that g *eventually dominates* f if there is some n such that for every $m \geq n$, $g(m) > f(m)$. In other words, from some point on, g is bigger than f .
 Show that if f_0, f_1, f_2, \dots is any sequence of functions from \mathbb{N} to \mathbb{N} , indexed by the natural numbers, then there is a function g that eventually dominates each f_i . (Hint: construct g so that for each i , $g(n) > f_i(n)$ for every $n \geq i$.)
7. Show that the relation \preceq defined in Section 22.5 is reflexive and transitive.

AXIOMATIC FOUNDATIONS

In this final chapter, our story comes full circle. We started our journey with symbolic logic, using the propositional connectives to model logical terms like “and,” “or,” “not,” and “implies.” To that we added the quantifiers and function and relation symbols of first-order logic. From there, we moved to sets, functions, and relations, which are ubiquitous in modern mathematics; the natural numbers and induction; and then topics such as number theory, combinatorics, the real numbers, and the infinite. Here we return to symbolic logic, and see how it can be used to provide a formal foundation for all of mathematics.

Specifically, we will consider an axiomatic framework known as *Zermelo-Fraenkel set theory*, which was introduced early in the twentieth century. In the set-theoretic view of mathematics, every mathematical object is a set. The axioms assert the existence of sets with various properties. From the collection of all sets, we carve out the usual inhabitants of the mathematical universe, not just the various number systems we have considered, but also pairs, finite sequences, relations, functions, and so on. This provides us with an idealized foundation for everything we have done since [Chapter 11](#).

At the end of this chapter, we will briefly describe another axiomatic framework, *dependent type theory*, which is the one used by Lean. We will see that it provides an alternative perspective on mathematical objects and constructions, but one which is nonetheless inter-interpretable with the set-theoretic point of view.

23.1 Basic Axioms for Sets

The axioms of set theory are expressed in first-order logic, for a language with a single binary relation symbol, \in . We think of the entire mathematical universe as consisting of nothing but sets; if x and y are sets, we can express that x is an element of y by writing $x \in y$. The first axiom says that two sets are equal if and only if they have the same elements.

$$\text{Extensionality: } \forall x, y (x = y \leftrightarrow \forall z (z \in x \leftrightarrow z \in y))$$

The next axiom tells us that there is at least one interesting set in the universe, namely, the set with no element.

$$\text{Empty set: } \exists x \forall y y \notin x$$

Here, of course, $x \notin y$ abbreviates $\neg(x \in y)$. By the axiom of extensionality, the set asserted to exist by this axiom is unique: in other words, if x_1 and x_2 each have no elements, then, vacuously, any element is in one if and only if it is in the other, so $x_1 = x_2$. This justifies using the word *the* in the phrase *the empty set*. Given this fact, it should seem harmless to introduce a new symbol, \emptyset , to denote the set matching that description. Indeed, one can show that this is case: in a precise sense, such expansions to a first-order language can be viewed as nothing more than a convenient manner of expression, and statements in the bigger language can be translated to the original language in a way that justifies all the expected inferences. We will not go into

the details here, and, rather, take this fact for granted. Using the new symbol, the empty set axiom tells us the empty set satisfies the property $\forall y y \notin \emptyset$.

The third axiom tells us that given two sets x and y , we can form a new set z whose elements are exactly x and y .

$$\text{Pairing: } \forall x, y \exists z \forall w (w \in z \leftrightarrow w = x \vee w = y)$$

There is a stealth usage of this axiom lurking nearby. The axiom does not require that x and y are different, so, for example, we can take them both to be the empty set. This tells us that the set $\{\emptyset\}$, whose only element is the empty set, exists. More generally, the axiom tells us that for any x , we have the set $\{x\}$ whose only element is x , and for any x and y , we have $\{x, y\}$, as described above. Once again, the axiom of extensionality tells us that the sets meeting these descriptions are unique, so it is fair to use the corresponding notation. We are now off and running! We now have all of the following sets, and more:

$$\emptyset, \{\emptyset\}, \{\{\emptyset\}\}, \{\emptyset, \{\emptyset\}\}, \{\{\{\emptyset\}\}\}, \dots$$

Still, we can never form a set with more than two elements in this way. To that end, it would be reasonable to add an axiom that asserts for every x and y , the set $x \cup y$ exists. But we can do better. Remember that if x is any set, $\bigcup x$ denotes the union of all the sets in x . In other words, for any set z , z is an element of $\bigcup x$ if and only if z is in w for some set w in x . The following axiom asserts that this set exists.

$$\text{Union: } \forall x \exists y \forall z (z \in y \leftrightarrow \exists w (w \in x \wedge z \in w))$$

Once again, this justifies the use of the \bigcup notation. We get the ordinary binary union using this axiom together with pairing, since we have $x \cup y = \bigcup\{x, y\}$.

At this stage, it will be useful to invoke some additional notation that was first introduced in our informal presentation of sets. If A is any first-order formula in the language of set theory, $\forall x \in y A$ abbreviates $\forall x (x \in y \rightarrow A)$ and $\exists x \in y A$ abbreviates $\exists x (x \in y \wedge A)$, relativizing the quantifiers as described in Section 7.4. The expression $x \subseteq y$ abbreviates $\forall z \in x (z \in y)$, as you would expect.

The next axiom asserts that for every set x , the power set, $\mathcal{P}(x)$ exists.

$$\text{Power Set: } \forall x \exists y \forall z (z \in y \leftrightarrow z \subseteq x)$$

We have begun to populate the universe with basic set constructions. It is the next axiom, however, that gives set theory its remarkable flexibility. Properly speaking, it is not a single axiom, but a *schema*, an infinite family of axioms given by a single template. The schema is meant to justify set-builder notation $\{w \mid \dots\}$ that was ubiquitous in Chapter 11. The first question we need to address is what we are allowed to write in place of the ellipsis. In our informal presentation of set theory, we said that one can define a set using any property, but that only prompts the question here as to what counts as a “property.” Axiomatic set theory provides a simple but powerful answer: we can use any first-order formula in the language of set theory.

Another concern centers around Russell’s paradox, as discussed in Section 11.1. Any theory that allows us to define the set $\{w \mid w \notin w\}$ is inconsistent, since if we call this set z , we can show $z \in z$ if and only if $z \notin z$, which is a contradiction. Once again, set theory offers a simple and elegant solution: for any formula $A(z)$ and set y , we can instead form the set $\{w \in y \mid A(w)\}$, consisting of the elements of y that satisfy A . In other words, we have to first use the other axioms of set theory to form a set y that is big enough to include all the elements that we want to consider, and then use the formula A to pick out the ones we want.

The axiom schema we want is called *separation*, because we use it to separate the elements we want from those in a bigger collection.

$$\text{Separation: } \forall x_1, x_2, \dots, x_n, y \exists z \forall w (w \in z \leftrightarrow w \in y \wedge A(w, x_1, x_2, \dots, x_n))$$

Here, A can be any formula, and the list of variables x_1, \dots, x_n that are shown indicate that the formula A can have some parameters, in which case the set we form depends on these values. For example, in ordinary

mathematics, given a number m we can form the set $\{n \in \mathbb{N} \mid \text{prime}(n) \wedge n > m\}$. In this example, the description involves m and n , and the set so defined depends on m .

We could use the separation axiom to simplify the previous axioms. For example, as long as we know that *any* set x exists, we can define the empty set as $\{y \in x \mid \perp\}$. Similarly, in the pairing axiom, it is enough to assert that there is a set that contains x and y as elements, because then we can use separation to carve out the set whose elements are exactly x and y .

These are only the first six axioms of set theory; we have four more to go. But these axioms alone provide a foundation for reasoning about sets, relations, and functions, as we did in [Chapter 11](#), [Chapter 13](#), and [Chapter 15](#). For example, we have already defined the union operation, and we can define set intersection $x \cap y$ as $\{z \in x \cup y \mid z \in x \wedge z \in y\}$. We cannot define arbitrary set complements; for example, the exercises ask you to show that in set theory we can prove that there is no set that contains all sets, and so the complement of the empty set does not exist. But given any two sets x and y , we can define their difference $x \setminus y$ as $\{z \in x \mid z \notin y\}$. The exercises below ask you to show that we can also define indexed unions and intersections, once we have developed the notion of a function.

We would like to define a binary relation between two sets x and y to be a subset of $x \times y$, but we first have to define the cartesian product $x \times y$. Remember that in [Section 11.4](#) we defined the ordered pair (u, v) to be the set $\{\{u\}, \{u, v\}\}$. As a result, we can use the separation axiom to define

$$x \times y = \{z \in \dots \mid \exists u \in x \exists v \in y (z = (u, v))\}$$

provided we can prove the existence of a set big enough to fill the “...” In the exercises below, we ask you to show that the set $\mathcal{P}(\mathcal{P}(x \cup y))$ contains all the relevant ordered pairs. A binary relation r on x and y is then just a subset of $x \times y$, where we interpret $r(u, v)$ as $(u, v) \in r$. We can think of ordered triples from the sets x, y, z as elements of $x \times (y \times z)$ and so on. This gives us ternary relations, four-place relations, and so on.

Now we can say that a function $f : x \rightarrow y$ is really a binary relation satisfying $\forall u \in x \exists! v \in y f(u, v)$, and we write $f(u) = v$ when v is the unique element satisfying $f(u, v)$. A function f taking arguments from sets x, y , and z and returning an element of w can be interpreted as a function $f : x \times y \times z \rightarrow w$, and so on.

With sets, relations, and functions, we have the basic infrastructure we need to do mathematics. All we are missing at this point are some interesting sets and structures to work with. For example, it would be nice to have a set of natural numbers, \mathbb{N} , with all the properties we expect it to have. So let us turn to that next.

23.2 The Axiom of Infinity

With the axioms we have so far, we can form lots of finite sets, starting with \emptyset and iterating pairing, union, powerset, and separation constructions. This will give us sets like

$$\emptyset, \{\emptyset\}, \{\{\emptyset\}\}, \{\emptyset, \{\emptyset\}\}, \{\{\{\emptyset\}\}\}, \dots$$

But the axioms so far do not allow us to define sets that are more interesting than these. In particular, none of the axioms gives us an infinite set. So we need a further axiom to tell us that such a set exists.

Remember that in [Chapter 17](#) we characterized the natural numbers as a set with a distinguished element, 0, and an injective operation *succ*, satisfying the principles of induction and recursive definition. In set theory, everything is a set, so if we want to represent the natural numbers in that framework, we need to identify them with particular sets. There is a natural choice for 0, namely, the empty set, \emptyset . For a successor operation, we will use the function *succ* defined by $\text{succ}(x) = x \cup \{x\}$. The choice is a bit of a hack; the best justification for the definition is that it works. With this definition, the first few natural numbers are as follows:

$$0 = \emptyset, \quad 1 = \{\emptyset\}, \quad 2 = \{\emptyset, \{\emptyset\}\}, \quad 3 = \{\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}\}, \quad \dots$$

It is more perspicuous to write them as follows:

$$0 = \emptyset, \quad 1 = \{0\}, \quad 2 = \{0, 1\}, \quad 3 = \{0, 1, 2\}, \quad 4 = \{0, 1, 2, 3\}, \quad \dots$$

In general, $n + 1$ is represented by the set $\{0, 1, \dots, n\}$, in which case, $m \in n$ is the same as $m < n$. This is just an incidental property of our encoding, but it is a rather charming one.

Recall from Chapter 17 that we can characterize the set of natural numbers as follows:

- There is an element $0 \in \mathbb{N}$ and there is an injective function $\text{succ} : \mathbb{N} \rightarrow \mathbb{N}$, with the additional property that $\text{succ}(x) \neq 0$ for any x in \mathbb{N} .
- The set \mathbb{N} satisfies the principle of induction: if x is a subset of \mathbb{N} that contains 0 and is closed under succ (that is, whenever z is in \mathbb{N} , so is $\text{succ}(z)$), then $x = \mathbb{N}$.

We have already settled on the definitions of 0 and succ , but we don't yet have any set that contains the first and is closed under applying the second. The axiom of infinity asserts precisely that there exists such a set.

$$\text{Infinity: } \exists x (\emptyset \in x \wedge \forall y (y \in x \rightarrow y \cup \{y\} \in x))$$

Say a set x is *inductive* if it satisfies the property after the existential quantifier, namely, that it contains the empty set and is closed under our successor operation. Notice that the set of natural numbers, which we are still trying to define formally, has this property. The axiom of infinity asserts the existence of *some* inductive set, but not necessarily the natural numbers themselves; an inductive set can have other things in it as well. In a sense, the principle of induction says that the natural numbers is the *smallest* inductive set. So we need a way to separate that set from the one asserted to exist by the axiom of infinity.

Let x be any inductive set, as asserted to exist by the axiom of infinity. Let

$$y = \bigcap \{z \subseteq x \mid z \text{ is inductive}\}.$$

Here $z \subseteq x$ can also be written $z \in \mathcal{P}(x)$, so the inside set exists by the separation axiom. According to this definition, y is the intersection of every inductive subset of x , so an element w is in y if and only if w is in every inductive subset of x . We claim that y itself is inductive. First, we have $\emptyset \in y$, since the empty set is an element of every inductive set. Next, suppose w is in y . Then w is in every inductive subset of x . But since every inductive set is closed under successor, $\text{succ}(w)$ is in every inductive subset of x . So $\text{succ}(w)$ is in the intersection of all inductive subsets of x —which is y !

It quickly follows that y is a subset of *every* inductive set. To see this, suppose that z is inductive. You can check that $z \cap x$ is inductive, and thus $y \subseteq z \cap x \subseteq z$.

The more interesting point is that y also satisfies the principle of induction. To see this, suppose $u \subseteq y$ contains the empty set and is closed under succ . Then u is inductive, and since y is a subset of every inductive set, we have $y \subseteq u$. Since we assumed $u \subseteq y$, we have $u = y$, which is what we want.

To summarize, then, we have proved the existence of a set that contains 0 and is closed under a successor operation and satisfies the induction axiom. Moreover, there is only one such set: if y_1 and y_2 both have this property, then so does $y_1 \cap y_2$, and by the induction principle, this intersection has to be equal to both y_1 and y_2 , in which case y_1 and y_2 are equal. It then makes sense to call the unique set with these properties the *natural numbers*, and denote it by the symbol \mathbb{N} .

There is only one piece of the puzzle missing. It is clear from the definition that 0 is not the successor of any number, but it is not clear that the successor function is injective. We can prove that by first noticing that the natural numbers, as we have defined them, have a peculiar property: if z is a natural number, y is an element of z , and x is an element of y , then x is an element of z . This says exactly that the \in relation is transitive on natural numbers, which is not surprising, since we have noted that \in on the natural numbers, under our representation, coincides with $<$. To prove this claim formally, say that a set z is *transitive* if it has the property just mentioned, namely, that every element of an element of z is an element of z . This is equivalent to saying that for every $y \in z$, we have $y \subseteq z$.

Lemma. Every natural number is transitive.

Proof. By induction on the natural numbers. Clearly, \emptyset is transitive. Suppose x is transitive, and suppose $y \in \text{succ}(x)$ and $z \in y$. Since $\text{succ}(x) = x \cup \{x\}$, we have $y \in x$ or $y \in \{x\}$. If $y \in x$, then by the inductive hypothesis, we have $z \in x$, and hence $z \in \text{succ}(x)$. Otherwise, we have $y \in \{x\}$, and so $y = x$. In that case, again we have $z \in x$, and hence $z \in \text{succ}(x)$.

The next lemma shows that, on transitive sets, union acts like the predecessor operation.

Lemma. If x is transitive, then $\bigcup \text{succ}(x) = x$.

Proof. Suppose y is in $\bigcup \text{succ}(x) = \bigcup (x \cup \{x\})$. Then either $y \in z$ for some $z \in x$, or $y \in x$. In the first case, also have $y \in x$, since x is transitive.

Conversely, suppose y is in x . Then y is in $\bigcup \text{succ}(x)$, since we have $x \in \text{succ}(x)$.

Theorem. succ is injective on \mathbb{N} .

Proof. Suppose x and y are in \mathbb{N} , and $\text{succ}(x) = \text{succ}(y)$. Then x and y are both transitive, and we have $x = \bigcup \text{succ}(x) = \bigcup \text{succ}(y) = y$.

With that, we are off and running. Although we will not present the details here, using the principle of induction we can justify the principle of recursive definition. We can then go on to define the basic operations of arithmetic and derive their properties, as done in [Chapter 17](#). We can go on to define the integers, the rational numbers, and the real numbers, as described in [Chapter 21](#), and to develop subjects like number theory and combinatorics, as described in [Chapter 19](#) and [Chapter 20](#). In fact, it seems that any reasonable branch of mathematics can be developed formally on the basis of axiomatic set theory. There are pitfalls, for example, having to do with large collections: for example, just as it is inconsistent to postulate the existence of a set of all sets, in the same way, there is no collection of all partial orders, or all groups. So when interpreting some mathematical claims, care has to be taken in some cases to restrict to sufficiently large collections of such objects. But this rarely amounts to more than careful bookkeeping, and it is a remarkable fact that, for the most part, the axioms of set theory are flexible and powerful enough to justify most ordinary mathematical constructions.

23.3 The Remaining Axioms

The seven axioms we have seen are quite powerful, and suffice to represent large portions of mathematics. We discuss the remaining axioms of Zermelo-Fraenkel set theory here.

So far, none of the axioms we have seen rule out the possibility that a set x can be an element of itself, that is, that we can have $x \in x$. The following axiom precludes that.

$$\text{Foundation } \forall x (\exists y y \in x \rightarrow \exists y \in x \forall z \in x z \notin y))$$

The axiom says that if x is a nonempty set, there is an element y of x with the property that no element of y is again an element of x . This implies we cannot have a descending chain of sets, each one an element of the one before:

$$x_1 \ni x_2 \ni x_3 \ni \dots$$

If we apply the axiom of foundation to the set $\{x_1, x_2, x_3, \dots\}$, we find that some element x_i does not contain any others, which is only possible if the sequence has terminated with x_i . In other words, the axiom implies (and is in fact equivalent to) the statement that the elementhood relation is *well founded*, which explains the name.

The axioms listed in the previous section tell a story of how sets come to be: we start with the empty set, and keep applying constructions like power set, union, and separation, to build more sets. Set theorists often imagine the hierarchy of sets as forming a big V , with the empty set at the bottom and a set at any higher level comprising, as its elements, sets that appear in levels below. In a precise sense (which we will not spell out here), the axiom of foundation says that every set arises in such a way.

Now consider the following sequence of sets:

$$\mathbb{N}, \mathcal{P}(\mathbb{N}), \mathcal{P}(\mathcal{P}(\mathbb{N})), \mathcal{P}(\mathcal{P}(\mathcal{P}(\mathbb{N}))), \dots$$

It is consistent with all the axioms we have seen so far that every set in the mathematical universe is an element of one of these. That still gives us a lot of sets, but, since we have described that sequence, we can just as well imagine a set that contains all of them:

$$\{\mathbb{N}, \mathcal{P}(\mathbb{N}), \mathcal{P}(\mathcal{P}(\mathbb{N})), \mathcal{P}(\mathcal{P}(\mathcal{P}(\mathbb{N}))), \dots\}.$$

The following axiom implies the existence of such a set.

$$\begin{aligned} \text{Replacement: } \forall x, y_1, \dots, y_n (\forall z \in x \exists! w A(z, w, y_1, \dots, y_n) \rightarrow \\ \exists u \forall w (w \in u \leftrightarrow \exists z \in x A(z, w, y_1, \dots, y_n))) \end{aligned}$$

Like the axiom of separation, this axiom is really a schema, which is to say, a separate axiom for each formula A . Here, too, the variables y_1, y_2, \dots, y_n are free variables that can occur in A . To understand the axiom, it is easiest to think of them as parameters that are fixed in the background, and then ignore them. The axioms says that if, for every z in x there is a unique w satisfying $A(z, w)$, then there is a single set, u , that consists of the w values corresponding to every such z . In other words, if you think of A as a function whose domain is x , the axiom asserts that the range of that function exists. In the example above, x is the natural numbers, and $A(z, w)$ says that w is the z -fold iterate of the power set of the natural numbers.

The nine axioms we have listed so far comprise what is known as *Zermelo-Fraenkel Set Theory*. There is on additional axiom, the axiom of choice, which is usually listed separately for historical reasons: it was once considered controversial, and in the early days, mathematicians considered it important to keep track of whether the axiom was actually used in a proof. There are many equivalent formulations, but this one is one of the most straightforward.

$$\text{Choice: } \forall x (\emptyset \notin x \rightarrow \exists f : x \rightarrow \bigcup x \forall y \in x f(y) \in y)$$

The axiom says that for any collection x of nonempty sets, there is a function f that selects an element from each one. We used this axiom, informally, in [Section 15.2](#) to show that every surjective function has a right inverse. In fact, this last statement can be shown to be equivalent to the axiom of choice on the basis of the other axioms.

To summarize, then, the axioms of Zermelo-Fraenkel Set Theory with the axiom of choice are as follows:

1. Extensionality:

$$\forall x, y (x = y \leftrightarrow \forall z (z \in x \leftrightarrow z \in y))$$

2. Empty set:

$$\exists x \forall y y \notin x$$

3. Pairing:

$$\forall x, y \exists z \forall w (w \in z \leftrightarrow w = x \vee w = y)$$

4. Union:

$$\forall x \exists y \forall z (z \in y \leftrightarrow \exists w (w \in x \wedge z \in w))$$

5. Power set:

$$\forall x \exists y \forall z (z \in y \leftrightarrow z \subseteq x)$$

6. Separation:

$$\forall x_1, x_2, \dots, x_n, y \exists z \forall w (w \in z \leftrightarrow w \in y \wedge A(w, x_1, x_2, \dots, x_n))$$

7. Infinity:

$$\exists x (\emptyset \in x \wedge \forall y (y \in x \rightarrow y \cup \{y\} \in x))$$

8. Foundation:

$$\forall x (\exists y y \in x \rightarrow \exists y \in x \forall z \in x z \notin y)$$

9. Replacement:

$$\forall x, y_1, \dots, y_n (\forall z \in x \exists! w A(z, w, y_1, \dots, y_n) \rightarrow \exists u \forall w (w \in u \leftrightarrow \exists z \in x A(z, w, y_1, \dots, y_n)))$$

10. Choice:

$$\forall x (\emptyset \notin x \rightarrow \exists f : x \rightarrow \bigcup x \forall y \in x f(y) \in y)$$

23.4 Type Theory

As a foundation for mathematics, Zermelo-Fraenkel set theory is appealing. The underlying logic, first-order logic, provides the basic logical framework for quantifiers and the logical connectives. On top of that, the theory describes a single, intuitively natural concept, that of a set of elements. The axioms are plausible

eminently reasonable. It is remarkable that virtually all of modern mathematics can be reduced to such simple terms.

There are other foundations on offer, however. These tend to be largely inter-interpretable with set theory. After all, set-theoretic language is now ubiquitous in everyday mathematics, so any reasonable foundation should be able to make sense of such language. On the other hand, we have already noted that set theory is remarkably expressive and robust, and so it should not be surprising that other foundational approaches can often be understood in set-theoretic terms.

This is, in particular, true of *dependent type theory*, which is the basis of the Lean theorem prover. The syntax of type theory is more complicated than that of set theory. In set theory, there is only one kind of object; officially, everything is a set. In contrast, in type theory, every well-formed expression in Lean has a *type*, and there is a rich vocabulary of defining types.

In fact, Lean is based on a version of an axiomatic framework known as the *Calculus of Inductive Constructions*, which provides all of the following:

- A hierarchy of *type universes*, **Type 0**, **Type 1**, **Type 2**, ... and a special type **Prop**. The expression **Type** abbreviates **Type 0**, and saying $T : \text{Type}$ can be interpreted as saying that T is a datatype. The type **Prop** is the type of propositions.
- *Dependent function types* $\Pi x : A, B\ x$. An element f of this type is a function which maps any element a of type A to an element $f\ a$ of type $B\ a$. The fact that the type of the output depends on the type of the input is what makes the function “dependent.” In the case where the output type does not depend on the input, we have the simple function type $A \rightarrow B$.
- *Inductive types*, like the natural numbers, specified by *constructors*, like zero and successor. Each such type comes with principles of induction and recursion.

These constructions account for both the underlying logic of assertions (that is, the propositions) as well as the objects of the universe, which are elements of the ordinary types.

It is straightforward to interpret type theory in set theory, since we can view each type as a set. The type universes are simply large collections of sets, and dependent function types and inductive types can be explained in terms of set-theoretic constructions. We can view **Prop** as the set $\{\top, \perp\}$ of truth values, just as we did when we described truth-table semantics for propositional logic.

Given this last fact, why not just use set theory instead of type theory for interactive theorem proving? Some interactive theorem provers do just that. But type theory has some advantages:

- The fact that the rules for forming expressions are so rigid makes it easier for the system to recognize typographical errors and provide useful feedback. In type theory, if f has type $\mathbb{N} \rightarrow \mathbb{N}$ it can be applied only to a natural number, and a theorem prover can flag an error if the argument has the wrong type. In set theory, anything can be applied to anything, whether or not doing so really makes sense.
- Again, because the rules for forming expressions are so rigid, the system can infer useful information from the components of an expression, whereas set theory would require us to make such information explicit. For example, with f as above, a theorem prover can infer that a variable x in $f\ x$ should have type \mathbb{N} , and that the resulting expression again has type \mathbb{N} . In set theory, $x \in \mathbb{N}$ has to be stated as an explicit hypothesis, and $f(x) \in \mathbb{N}$ is then a theorem.
- By encoding propositions as certain kinds of types, we can use the same language for defining mathematical objects and writing mathematical proofs. For example, we can apply a function to an argument in the same way we apply a theorem to some hypotheses.
- Expressions in a sufficiently pure part of dependent type theory have a computational interpretation, so, for example, the logical framework tells us how to evaluate the factorial function, given its definition. In set theory, the computational interpretation is specified independently, after the fact.

These facts hark back to the separation of concerns that we raised in [Chapter 1](#): different axiomatic foundations provide different idealized descriptions of mathematical activity, and can be designed to serve different purposes. If you want a clean, simple theory that accounts for the vast majority of mathematical proof, set theory is hard to beat. If you are looking for a foundation that makes computation central or takes the notion of a function rather than a set as basic, various flavors of type theory have their charms. For interactive theorem proving, pragmatic issues regarding implementation and usability come into play. What is important to recognize is that what all these idealized descriptions have in common is that they are all designed to model important aspects of mathematical language and proof. Our goal here has been to help you reflect on those features of mathematical language and proof that give mathematics its special character, and to help you better understand how they work.

23.5 Exercises

1. Use an argument similar Russell's paradox to show that there is no "set of all sets," that is, there is no set that contains every other set as an element.
2. Suppose x is a nonempty set, say, containing an element y . Use the axiom of separation to show that the set $\bigcap x$ exists. (Remember that something is an element of $\bigcap x$ if it is an element of every element of x .)
3. Justify the claim in [Section 23.1](#) that every element of $x \times y$ is an element of $\mathcal{P}(\mathcal{P}(x \cup y))$.
4. Given a set x and a function $A : x \rightarrow y$, use the axioms of set theory to prove the existence of $\bigcup_{i \in x} A(i)$.