The Absolute Fundamentals of Math

A Primer for Curious Adults

For everyone who got through school maths but never quite understood what it was for.

Eight chapters covering the foundations school almost certainly skipped.

Contents

Introduction

Why this primer exists, who it's for, and how to use it.

Chapter 1: What Is Math, Actually?

Axioms, theorems, and the idea that math is a human construction.

Chapter 2: The Language of Math

Constants, variables, terms, expressions, and equations.

Chapter 3: Logic — The Engine Underneath

Propositions, connectives, proofs, and quantifiers.

Chapter 4: Sets — Collections of Things

Set notation, operations, subsets, and why sets matter.

Chapter 5: Numbers — A Tower Built in Layers

Natural numbers, integers, rationals, reals, and beyond.

Chapter 6: Properties and Structure — The Rules of the Game

Commutativity, associativity, distributivity, and algebraic structure.

Chapter 7: Functions — Machines That Transform

Domain, range, graphs, linear functions, and composition.

Chapter 8: Abstraction — The Superpower of Mathematics

How stripping away details reveals universal truth.

Conclusion

Where to go from here.

Introduction

There is a particular kind of frustration that comes from feeling like you missed something important. Maybe you did reasonably well in school maths — you could follow the procedures, pass the tests — but it never quite *clicked*. Nobody explained why. Why do we have negative numbers? What is a proof, really? Why does it matter whether something is commutative? The procedures were taught, but the foundations were skipped.

This primer exists to fill that gap. It is not a textbook, and it is not a drill book. You will not find pages of practice exercises for long division. What you will find is an honest attempt to explain what mathematics is actually about, at the level of ideas — the kind of explanation that makes everything else make sense retroactively.

Who This Is For

This primer is written for any intelligent adult who feels like they 'missed the foundations.' You do not need to remember anything from school. You need only curiosity and the willingness to read slowly and think. If you studied maths to age 15 or so — fractions, basic algebra, some geometry — that is more than enough background. The goal is not to teach you to calculate faster. It is to show you the conceptual bedrock that calculation rests on.

What This Covers (and What It Doesn't)

We cover eight foundational topics: the nature of mathematics itself, its language, the logic that underlies all of it, the concept of sets, the tower of number systems, the structural properties of arithmetic, functions, and the grand unifying idea of abstraction. Each chapter is self-contained but builds on what came before.

What this primer does *not* cover: calculation techniques, trigonometry, calculus, statistics, or any applied mathematics. Those are important, but they come after the foundations. Think of this as the basement of a building — not glamorous, but everything else sits on top of it.

How to Use This Primer

Read it linearly. Each chapter builds on ideas from earlier chapters, and the order is deliberate. If a section feels fuzzy, re-read it before moving on — the fuzziness usually signals something important. Each chapter ends with a set of exercises. Some are computational, but most are reflective: they ask you to explain a concept in your own words, or to apply an idea to a new situation. These are not optional extras — doing them is how understanding gets consolidated.

A note on pacing: this is not light reading. Some paragraphs contain a lot of ideas compressed into a small space. Give yourself permission to read slowly. The reward for doing so is genuine understanding — not the feeling of understanding, but the real thing.

Chapter 1

What Is Math, Actually?

Here is a question that sounds simple but almost nobody can answer satisfyingly: *what is mathematics?* Ask most people and they will say something like 'numbers' or 'calculations' or 'equations.' Ask a working mathematician and they will probably pause, look slightly pained, and then give you an answer that sounds more like philosophy than arithmetic. That gap — between the school version and the real thing — is exactly what this chapter is about. We are going to dismantle the popular conception and replace it with something much more interesting.

The Popular Misconception

The version of mathematics most people encounter in school is mostly *arithmetic*: adding, subtracting, multiplying, dividing. Then some algebra — manipulating symbols according to rules. Then perhaps some geometry, a little trigonometry, possibly some calculus at the end if you stayed in the system long enough. All of it is presented as a collection of procedures to be memorised and applied: do this when you see that; carry the one; move the variable to the other side; remember PEMDAS.

This procedural view of mathematics is not wrong, exactly. Procedures are part of the picture. But it is a bit like describing music as 'pressing keys in the right order.' Technically accurate in a very narrow sense, but entirely missing the point. A musician does not primarily think about key-pressing. They think about tension and resolution, phrase and silence, emotion and structure. The key-pressing is the physical substrate of something much richer.

Something similar is true of mathematics. The calculations are the surface. Underneath them is a vast, intricate structure of ideas — the study of *abstract structure, pattern, and logical consequence*. Mathematics is not primarily about numbers. Numbers are just one of many kinds of mathematical object. At its core, mathematics is about reasoning: starting from clearly stated assumptions and seeing, with absolute certainty, what must follow.

This is not just a philosophical distinction. It has practical consequences. If you understand mathematics as a collection of procedures, you are stuck — every new type of problem requires learning a new procedure. But if you understand mathematics as a system of reasoning, you can figure out new things. The goal of this primer is to give you enough of the underlying structure that mathematics starts to feel more like thinking clearly and less like following recipes.

A System Built on Agreements: Axioms

Here is the key insight about how mathematics actually works. Every mathematical system begins with a set of **axioms** — statements that we simply *agree to accept as true*, without proof. We do not prove axioms. We choose them. Then we apply the rules of logic rigorously and see what follows.

This might sound like cheating, but it is not. The choice of axioms is not arbitrary — we choose them to be as simple and obviously true as possible, and to capture the intuitions we care about. But once chosen, we are completely rigorous: we do not allow ourselves to sneak in any additional assumptions. Every new result must be derived from the axioms alone, using only the rules of logic.

The power of mathematics comes from what happens *after* you set up your axioms. From a small handful of agreed-upon starting points, an enormous and intricate structure can be derived — all of it logically guaranteed. Every theorem (a proven result) in mathematics is ultimately traceable back to the axioms. The connection might be indirect, passing through dozens or hundreds of intermediate results, but it is there.

A **theorem**, by contrast to an axiom, is a statement that has been *proved* — derived by a sequence of logical steps from the axioms (and possibly from other theorems proved earlier). Axioms are the seeds; theorems are the entire forest that grows from them. A **lemma** is a smaller, helper theorem, proved along the way to something bigger. A **corollary** is a result that follows immediately and easily from a theorem just proved.

A Concrete Example: Euclid's Geometry

The most famous axiomatic system in history is Euclid's geometry, set out around 300 BCE in a work called the *Elements*. Euclid began with five postulates (a slightly older word for axioms) — five statements so simple and obvious that no one could reasonably object to them. Here are the first three, loosely stated:

Postulate 1: You can draw a straight line between any two points.

Postulate 2: You can extend any straight line indefinitely in a straight line.

Postulate 3: You can draw a circle with any centre and any radius.

These seem almost embarrassingly obvious. Of course you can draw a line between two points. Why bother writing that down? The answer is: because once you write it down as a formal starting point, you can build on it rigorously. You know *exactly* what you are assuming, and therefore you know exactly where your theorems come from. Every geometric result you derive will trace back to these (and the other two) postulates.

From just five such postulates, Euclid derived hundreds of geometric results — the properties of triangles, the relationship between angles in parallel lines, the Pythagorean theorem, properties of circles, and much more. The *Elements* is one of the most influential books in the history of human knowledge. Not because of any particular result in it, but because of the method: the idea that you could take a small number of obvious truths and, through pure reasoning, derive an entire science.

What If We Change the Axioms?

Here is where things get genuinely mind-bending. Euclid's fifth postulate — usually stated as: through any point not on a given line, there is exactly *one* line parallel to the given line — was controversial from the start. It is more complex and less obviously 'obvious' than the other four. For nearly 2,000 years, mathematicians tried to prove it from the other four, suspecting it was actually a theorem rather than a true axiom. All attempts failed.

In the early 19th century, mathematicians tried something radical: what if you simply replaced that fifth postulate with a different one? What if you assumed that through a given external point there are *no* parallel lines to the given line? Or *infinitely many*? Would the result be self-contradictory nonsense?

The result was not contradiction — it was a completely new, internally consistent kind of geometry. With the 'no parallels' version, you get **spherical geometry** (the geometry of the surface of a sphere, where 'straight lines' are great circles, and yes, two 'parallel' great circles always meet). With the 'infinitely many parallels' version, you get **hyperbolic geometry** — a saddle-shaped space. Both are perfectly valid mathematical systems. And, strikingly, general relativity — our best theory of gravity and the structure of the universe — uses non-Euclidean geometry to describe curved spacetime. The abstract mathematical exploration turned out to describe physical reality.

The lesson is profound: different axiom sets lead to different but equally valid mathematical systems. Mathematics is not about uncovering a single fixed truth that exists out there in the world. It is about exploring what follows, with perfect logical rigour, from whatever starting points you choose.

Axioms in Other Parts of Mathematics

Euclid's geometry is not the only axiom system in mathematics. Modern mathematics is full of them. The study of arithmetic — the theory of natural numbers — can be built on the **Peano axioms**, a set of five axioms proposed by Giuseppe Peano in 1889. The Peano axioms define the natural numbers in terms of a starting point (zero, or one, depending on the version) and the concept of a 'successor' function that maps each number to the next one. From these five axioms, all of arithmetic can be derived.

Set theory — which we will cover in Chapter 4 — has its own axiom system, called ZFC (Zermelo-Fraenkel axioms plus the Axiom of Choice). This is the axiom system that most working mathematicians implicitly use when they talk about sets. It has nine axioms, each one capturing some aspect of how sets ought to behave. All of modern mathematics can (in principle) be derived from ZFC.

In practice, working mathematicians rarely think about axioms explicitly. When a number theorist proves a new result about prime numbers, they are not consciously checking each step against the Peano axioms. But the axiomatic foundation is always there, providing the ultimate guarantee of validity. It is like the constitution underlying a legal system: rarely consulted directly, but providing the ultimate source of legitimacy for all the rules above it.

Math as a Human Construction — But Not Arbitrary

This brings us to a philosophical point that surprises many people: mathematics is a *human construction*. Not in the sense that we make it up arbitrarily — the logical consequences of a given set of axioms are completely determined, not up to us — but in the sense that we choose what axioms to adopt, what objects to study, and what questions to ask.

The fact that the same mathematical structures (like non-Euclidean geometry) keep turning up as useful descriptions of physical reality does suggest that mathematics is somehow latching onto genuine features of the universe. The physicist Eugene Wigner wrote a famous essay about 'the unreasonable effectiveness of mathematics' — the mysterious fact that mathematical structures developed with no practical application in mind repeatedly turn out to be exactly what physics needs. Complex numbers (Chapter 5) were considered a pure abstraction until they turned out to be essential for quantum mechanics. Non-Euclidean geometry was pure curiosity until Einstein needed it.

This is one of the deepest puzzles about the nature of mathematics, and there is no consensus on the answer. But whatever your philosophical position, one thing is agreed upon: the reliability of mathematics comes from its logical rigour. A theorem is true not because of experiment or authority, but because of proof. This makes mathematics unique among human intellectual endeavours: its results are certain, and that certainty is timeless. The theorems Euclid proved in 300 BCE are still true. They will always be true.

The Role of Definitions

Alongside axioms, **definitions** play a crucial role in mathematics. A definition precisely specifies what a term means. When a mathematician says 'let n be an even integer,' the word 'even' has a precise definition: an integer is even if and only if it is divisible by 2, or equivalently, if it can be written in the form 2k for some integer k.

This precision is not pedantry — it is what makes proofs possible. You cannot reason rigorously about 'even numbers' unless you know exactly what 'even' means. In Chapter 3, when we proved that the square of an even number is even, the first step was 'if n is even, then n = 2k for some integer k.' That step only works because 'even' has been precisely defined.

One of the habits mathematical thinking develops is the instinct to ask 'but what exactly do you mean by that?' This is not being difficult — it is recognising that vague terms lead to uncertain reasoning. Almost every paradox and confusion in mathematics can be traced back to an undefined or ambiguous term.

Why This Matters for the Rest of This Primer

Understanding that mathematics is an axiomatic, logically rigorous system changes how you read everything that follows in this primer. When we define what a 'set' is in Chapter 4, or what a 'function' is in Chapter 7, we are not just introducing vocabulary for its own sake. We are laying down the precise definitions that allow us to reason rigorously about these objects. When we state properties like commutativity or distributivity in Chapter 6, we are identifying structural regularities that can be verified from first principles.

Mathematics is not a collection of facts to be memorised. It is a way of building certain, justified knowledge from precise starting points. Once you see it that way, it starts to feel less like an arbitrary collection of rules and more like an extraordinary human achievement — and a genuinely learnable one, because it all connects.

Exercises

1.1 In your own words, explain the difference between an axiom and a theorem. Why can't axioms be proved? What would happen if you tried to prove every statement?

1.2 Euclid's first postulate says you can draw a straight line between any two points. On the surface of a sphere, 'straight lines' are great circles (like the equator or lines of longitude). Can you draw a 'straight line' between any two points on a sphere? Is the result unique? What does this tell you about the axiom?

1.3 Can you think of an 'axiom' in everyday life — something that a community or institution simply agrees to accept as a starting point? (Examples might come from law, economics, or social norms.) What makes a good choice of axiom?

1.4 Why might it be useful to state seemingly obvious things (like 'you can draw a line between two points') formally and explicitly, even when everyone agrees they are true?

1.5 If two mathematical systems start from different axioms, can both be valid simultaneously? Use the example of Euclidean versus spherical geometry to argue that the answer is yes.

1.6 Wigner wrote about 'the unreasonable effectiveness of mathematics.' Think of one example (it can be from everyday life, not just physics) where an abstract mathematical idea turned out to be surprisingly useful. What might this tell us about the nature of mathematics?

1.7 Reflect: has your understanding of what mathematics 'is' changed after reading this chapter? What was your previous conception, and in what specific way has it shifted?

Chapter 2

The Language of Math

Terms, Expressions, and Equations

Every field of study develops its own language. Lawyers have 'tort' and 'estoppel.' Doctors have 'prognosis' and 'contraindication.' Biologists have 'mitosis' and 'phenotype.' Mathematicians have their own precise vocabulary too — and a great deal of confusion about mathematics comes not from the ideas being too hard, but from the words being misunderstood or never properly introduced.

Unlike natural language, mathematical language is designed for precision above all else. Every term has exactly one meaning. Ambiguity is not a feature — it is a bug, to be eliminated wherever it appears. This chapter introduces the basic vocabulary with care, because getting this right will make everything else in the primer much clearer.

Constants and Variables

Mathematics talks about two fundamentally different kinds of things: quantities that are fixed, and quantities that can vary. This distinction is so basic that it is easy to miss its importance.

Constants

A **constant** is a value that does not change. The number 7 is a constant. The number *pi* (approximately 3.14159...) is a constant — it always means the same thing, regardless of context. The number e (approximately 2.71828...) is a constant. When you see a specific number written in a mathematical expression, it is a constant.

Some constants have names because they appear so frequently and are so important. *Pi* (π) is the ratio of a circle's circumference to its diameter — the same ratio for every circle, everywhere. e is the base of natural logarithms, arising naturally in problems involving growth and decay. Both of these constants are irrational (as we'll see in Chapter 5), meaning their decimal expansions never terminate or repeat.

Variables

A **variable** is a placeholder — a letter or symbol that represents some quantity that could take different values. When we write x, n, t, or a, we are saying 'some number, the specific value of which either we do not know yet, or which is allowed to vary.'

Variables are extraordinarily powerful, because they let us state general truths rather than specific ones. Consider the difference between '3 plus 4 equals 4 plus 3' and 'a plus b equals b plus a.' The first is one specific numerical fact, true for exactly this pair of numbers. The second is a universal truth about addition — it holds for every pair of numbers in existence, whether those numbers are small integers, large fractions, irrational numbers, or complex numbers. A single statement with variables replaces an infinite family of specific statements.

The conventional choice of letters for variables has historical roots. x, y, z tend to be used for unknowns in equations. n, m, k tend to be used for integers (whole number counts). t often represents time. f, g, h often represent functions (Chapter 7). But these are conventions, not rules. In any given piece of mathematics, the author chooses their notation and defines it.

Terms, Expressions, and Equations

With constants and variables in hand, we can build up larger mathematical objects. The hierarchy is: terms combine to form expressions; expressions combine with '=' to form equations.

Terms: The Building Blocks

A **term** is the basic unit of a mathematical expression. A term is a single 'chunk' — it can be:

- A number on its own: 5, −3, 1/2
- A variable on its own: x, n
- A number multiplied by one or more variables: $3x$, $7ab$, $-2x^2$

The numerical part of a term (the number multiplied by the variables) is called the **coefficient**. In the term $3x$, the coefficient is 3. In the term $-5y^2$, the coefficient is −5. In the term x (with no explicit number), the coefficient is an implicit 1.

The key characteristic of a term: a term does not contain addition or subtraction. It is a single multiplicative unit. $3x^2$ is one term. $3x^2 + 2x$ is two terms combined. This might seem like a picky distinction, but it matters for operations like 'collecting like terms' and factoring.

Like Terms

Two terms are called **like terms** if they have identical variable parts (same variables raised to the same powers). $3x^2$ and $7x^2$ are like terms — both have x^2. $4xy$ and $-2xy$ are like terms. But $5x$ and $5x^2$ are NOT like terms — the variable parts differ.

Why does this matter? Because you can only add or subtract like terms. $3x^2 + 7x^2 = 10x^2$ (add the coefficients, keep the variable part). But $3x^2 + 7x$ cannot be simplified further — those are unlike terms, like trying to add metres and kilograms.

Expressions: Combinations of Terms

An **expression** is formed by combining terms using operations (addition, subtraction, multiplication, division). An expression represents a value — it computes to a number when specific values are substituted for the variables — but it does not make any claim about equality. It simply *is* a quantity.

Examples: $3x + 7$ is an expression with two terms. $x^2 - 4x + 1$ is an expression with three terms. $2(a + b)$ is an expression with two terms inside parentheses (equivalent to $2a + 2b$ by the distributive property — see Chapter 6).

If you know that $x = 2$, you can evaluate the expression $3x + 7$ to get $3(2) + 7 = 13$. But the expression itself does not say anything is 13. It is a recipe, not a statement.

Equations: Making a Claim

An **equation** is a statement of equality between two expressions. It uses the equals sign '=' to assert that the expression on the left and the expression on the right represent the same quantity. An equation is a claim — it may be true or false.

The equation $3x + 7 = 13$ is a claim that for some value of x, $3x + 7$ equals 13. Solving the equation means finding which value (or values) of x make the claim true. (In this case, $x = 2$.) The equation $a + b = b + a$ is a claim that is true for *all* values of a and b — this is the commutative law of addition. Not all equations are puzzles to be solved; some are universal truths.

The Real Meaning of '='

Here is a misconception that causes enormous confusion, especially for students who have learned mathematics through arithmetic: many people think '=' means 'the answer is.' They read $3 + 4 = 7$ as 'three plus four gives the answer seven,' and they think of the left side as a 'question' and the right side as the 'answer.' This is understandable — it is how arithmetic is often taught — but it is wrong.

The equals sign means: *the expression on the left and the expression on the right represent the same quantity*. It is a statement of identity, not a signal that a calculation has been performed. '3 + 4' and '7' are two different ways of writing the same number. The equals sign says they are the same.

This is why mathematicians can legitimately write $7 = 3 + 4$, or $7 = 7$, or $3 + 4 = 1 + 6$. All of these are perfectly valid equations. You would never read '3 + 4' as 'the answer is' and then write the answer on the left — but all three equations above are using '=' correctly. The equals sign is symmetric; neither side has priority.

This matters practically when you do algebra. Manipulating an equation means doing the same thing to both sides to preserve the balance. If $3x + 7 = 13$, you can subtract 7 from both sides (getting $3x = 6$) because doing the same thing to both sides of a true equation gives another true equation. You are maintaining the identity, not 'performing a calculation.'

Inequalities

Closely related to equations are **inequalities**. Where an equation says two things are equal, an inequality says one is larger or smaller than the other. We have four inequality symbols:

- $a < b$: 'a is strictly less than b' (e.g., $3 < 7$)
- $a > b$: 'a is strictly greater than b' (e.g., $7 > 3$)
- $a \leq b$: 'a is less than or equal to b'
- $a \geq b$: 'a is greater than or equal to b'

Inequalities describe ranges of values rather than exact ones. The inequality $x > 3$ does not pin down a single value of x — it describes the infinite set of all numbers greater than 3. This connects inequalities to the concept of sets (Chapter 4) in a natural way: the solutions to an inequality form a set.

One important property: if you multiply or divide both sides of an inequality by a *negative* number, the direction of the inequality flips. If $a < b$, then $-a > -b$. This is because multiplying by −1 reflects numbers across zero on the number line, reversing their order. This surprises many people and is a common source of errors.

Mathematical Notation as Technology

It is worth pausing to appreciate that mathematical notation is not natural — it was invented. The symbols +, −, ×, ÷, =, x, and the others were developed gradually over many centuries, with different cultures and mathematicians using different notations at different times.

The plus sign (+) and minus sign (−) only appeared in Europe in the late 15th century. The equals sign (=) was introduced by the Welsh mathematician Robert Recorde in 1557. He chose two parallel horizontal lines because, he wrote, 'no two things can be more equal.' The notation for algebra we use today was largely standardised by René Descartes in the 17th century — including the convention of using letters from the end of the alphabet (x, y, z) for unknowns and letters from the beginning (a, b, c) for constants.

Why does this history matter? Because it shows that notation is a technology — a tool humans invented to make thinking easier. And like all technology, better notation enables things that were previously impossible.

Consider Roman numerals versus Arabic numerals. Try multiplying XLVII by XXIII in Roman numerals. It is genuinely extremely cumbersome. The same calculation with Arabic numerals ($47 \times 23 = 1081$) is straightforward, because the notation carries structural information — the place value system — that makes multiplication algorithms possible. Better notation enables harder problems.

When you feel frustrated by mathematical notation, it helps to remember that the notation exists to make things easier, not harder. If a piece of notation feels confusing, the notation has usually not been explained well enough — because the notation itself, once understood, should be illuminating.

Polynomials: A First Encounter

Before leaving this chapter, it is worth briefly introducing **polynomials** — one of the most important types of expression you will encounter. A polynomial in a variable x is an expression formed by adding together terms of the form $a_n x^n$, where the powers are non-negative integers.

Examples of polynomials: $3x + 2$ (a polynomial of degree 1, i.e., linear); $x^2 - 5x + 6$ (degree 2, quadratic); $4x^3 + x - 7$ (degree 3, cubic). The **degree** of a polynomial is the highest power of the variable that appears.

Polynomials are important because they are the simplest kind of expression involving variables, they are easy to work with algebraically, and many real-world relationships can be modelled by them. When we study functions in Chapter 7, polynomials will be among the first examples.

Exercises

2.1 Identify each of the following as a constant, variable, term, expression, or equation. (Some may belong to more than one category.) (a) 5 (b) n (c) $3x + 2$ (d) $4y^2$ (e) $2a + 3b = 12$ (f) $7 = 7$

2.2 Write an equation that captures the statement: 'A number plus seven equals fifteen.' Then solve for the unknown number.

2.3 In the expression $4x^2 - 7x + 3$, identify the three terms and, for each, state the coefficient and the variable part. Which pairs of terms (if any) are like terms?

2.4 Explain in your own words why the equals sign does NOT mean 'the answer is.' Give a mathematical example that illustrates the difference — an equation that looks 'backwards' but is perfectly valid.

2.5 If you multiply both sides of $3 < 7$ by -2, what happens to the inequality sign? Write out the result and explain geometrically why it makes sense.

2.6 Robert Recorde chose two parallel lines for '=' because 'no two things can be more equal.' Do you think this is a good justification? Can you think of a better or alternative notation for equality?

2.7 Write down a polynomial of degree 2 in the variable t. What is its value when $t = 0$? When $t = 1$? When $t = -1$?

Chapter 3

Logic — The Engine Underneath

Mathematics is built on logic. Not the informal, intuitive logic of everyday reasoning ('well, it stands to reason...') but precise, formal logic — a system of rules for what counts as valid reasoning. These rules are not the same as 'common sense.' In fact, they sometimes produce conclusions that violate common sense, but they are reliable in a way that common sense never is. This chapter introduces the logical machinery that mathematics runs on.

Why does a primer on mathematics need a chapter on logic? Because proof — the defining activity of mathematics — is simply logic applied carefully. Understanding what a proof is, what makes an argument valid, and what the common logical errors are will change how you read mathematics, and arguably how you think in general.

Propositions: The Atoms of Logic

Logic works with **propositions** — statements that are either true or false (not both, and not neither). This 'either-or' requirement is strict: a proposition has a definite truth value, even if we do not currently know what it is.

Examples of propositions: '7 is a prime number' (true), '12 is prime' (false), 'the sum of any two even numbers is even' (true), 'every even number greater than 2 is the sum of two primes' (this is the famous Goldbach conjecture — believed to be true but unproved as of 2026).

Examples of things that are NOT propositions: 'Close the door' (a command, not a statement), 'Is the lecture interesting?' (a question), 'This is a nice proof' (a subjective opinion, not objectively true or false), 'Hello' (a greeting). In formal logic, we work only with propositions.

The Liar Paradox

Consider the sentence: 'This statement is false.' Is it a proposition? If it is true, then what it says must be accurate — so it is false. But if it is false, then the opposite of what it says holds — so it is true. It cannot be true or false without contradiction. This is the ancient **Liar Paradox**, and it is not a curiosity but a deep challenge: it shows that not all grammatically correct sentences are propositions. Self-referential statements — statements that refer to themselves — can violate the basic requirement that a proposition have a definite truth value. We simply exclude such sentences from our logic.

Logical Connectives: AND, OR, NOT

Single propositions are useful, but we need to combine them. We do this with **logical connectives** — words (and corresponding symbols) that join propositions to form more complex ones.

NOT (Negation)

NOT P (written ¬P) is the negation of P — it is true exactly when P is false, and false exactly when P is true. NOT simply flips the truth value. If P is 'it is raining,' then NOT P is 'it is not raining.'

Negation has a subtlety worth noticing: the negation of 'all swans are white' is not 'all swans are not white' (i.e., 'no swans are white'). The correct negation is 'there exists at least one swan that is not white.' Getting negations right, especially with quantifiers (see later in this chapter), is one of the skills that careful logical thinking develops.

AND (Conjunction)

P AND Q (written P ∧ Q) is true only when *both* P and Q are individually true. If either is false, the conjunction is false.

Everyday example: 'It is raining AND it is cold.' This compound statement is true only if both sub-statements are simultaneously true. If it is raining but warm, the statement is false. If it is cold but sunny, the statement is false.

OR (Disjunction)

P OR Q (written P ∨ Q) is true when *at least one* of P, Q is true. Crucially — and this differs from everyday usage — in mathematics OR is **inclusive**: if both P and Q are true, 'P OR Q' is still true.

In everyday speech, 'Do you want tea or coffee?' usually means one or the other, not both. This is *exclusive or*. In logic and mathematics, OR is always inclusive: 'A or B or both.' This distinction matters in proofs where you want to split into cases — 'in case A, or in case B (including when both hold).'

Truth Tables

A **truth table** is a systematic way to show how a compound proposition's truth value depends on the truth values of its components. Here is the complete truth table for NOT, AND, and OR:

P	Q	NOT P	P AND Q	P OR Q
True	True	False	True	True
True	False	False	False	True
False	True	True	False	True
False	False	True	False	False

If/Then: The Backbone of Mathematical Reasoning

The most important logical structure in all of mathematics is the **conditional statement**, also called an 'if/then' statement. We write it as: *If P, then Q*, or symbolically P ⇒ Q (read 'P implies Q'). P is the **hypothesis** and Q is the **conclusion**.

The conditional says: whenever P is true, Q must also be true. It does NOT say that P is true. It does NOT say that if P is false, Q must be false. It says only that the truth of P guarantees the truth of Q.

Everyday example: 'If it is raining, then I will carry an umbrella.' This is a promise. If it rains, I will have my umbrella. But the statement says nothing about what happens if it is not raining — maybe I carry an umbrella anyway. The hypothesis being false leaves the conclusion undetermined.

This is captured in the truth table for P ⇒ Q. The conditional is false in only one case: when P is true and Q is false. All other combinations — P false and Q anything, or both true — make the conditional true. A promise is only broken if the conditions were met and the conclusion failed.

P	**Q**	**P ⇒ Q**
True	True	True
True	False	False
False	True	True
False	False	True

The Converse: Tempting but Different

The **converse** of 'If P then Q' is 'If Q then P' — hypothesis and conclusion are swapped. This is *a different statement*, and critically, it may be false even when the original is true.

Original: 'If it is raining, then the ground is wet.' — TRUE. Converse: 'If the ground is wet, then it is raining.' — FALSE (a garden hose or sprinkler could have wetted the ground).

Original (mathematical): 'If n is divisible by 4, then n is even.' — TRUE. Converse: 'If n is even, then n is divisible by 4.' — FALSE (6 is even but not divisible by 4).

Confusing a statement with its converse is one of the most common logical errors in everyday reasoning. 'Successful people work hard' (if successful then hard-working) does not imply 'hard-working people are successful' (if hard-working then successful). Many career myths and self-help fallacies involve exactly this error.

The Contrapositive: Logically Equivalent

The **contrapositive** of 'If P then Q' is 'If NOT Q, then NOT P.' Unlike the converse, the contrapositive is always *logically equivalent* to the original — if one is true, the other must be too, and vice versa.

Original: 'If it is raining, then the ground is wet.' Contrapositive: 'If the ground is NOT wet, then it is NOT raining.' Both say the same thing in different words. You can verify this by checking the truth tables: they match in every row.

This equivalence is a practical proof tool. Sometimes it is hard to prove 'If P then Q' directly, but easy to prove its contrapositive 'If NOT Q then NOT P.' Because they are logically equivalent, proving either one suffices.

A Common Fallacy: Affirming the Consequent

Here is a classic logical fallacy: 'If it is raining, then the ground is wet. The ground is wet. Therefore, it is raining.' This looks valid. It is not. It is the fallacy of **affirming the consequent**: knowing that Q is true does not let you conclude P is true.

In formal terms: knowing $P \Rightarrow Q$ and Q does not give you P. What you can validly conclude from '$P \Rightarrow Q$ and Q' is... nothing about P. (You could additionally conclude P if you also knew $Q \Rightarrow P$ — the converse — but that needs to be proved separately.)

This error appears constantly in everyday reasoning. 'If the economy is growing, unemployment falls. Unemployment just fell. Therefore, the economy is growing.' Not necessarily — unemployment can fall for other reasons. Careful logical thinking flags this kind of inference as requiring additional justification.

Necessary and Sufficient Conditions

Two phrases that frequently appear in precise reasoning are 'necessary' and 'sufficient.' They are directly related to the conditional.

P is sufficient for Q: $P \Rightarrow Q$. Having P is enough to guarantee Q. Rain is sufficient for wet ground. Being divisible by 4 is sufficient for being even.

P is necessary for Q: $Q \Rightarrow P$ (equivalently, NOT $P \Rightarrow$ NOT Q). You cannot have Q without having P. Wet ground is necessary for rain (if it hasn't rained, the ground is not wet from rain). Being even is necessary for being divisible by 4.

P is necessary and sufficient for Q: $P \Rightarrow Q$ AND $Q \Rightarrow P$. P and Q always occur together and only together. We write this $P \Leftrightarrow Q$, read 'P if and only if Q' (often abbreviated 'iff'). An integer is even if and only if its last digit is 0, 2, 4, 6, or 8. The 'if and only if' in mathematical definitions is almost always this: necessary and sufficient.

What Is a Proof?

Now we can say precisely what a mathematical **proof** is. A proof is a finite sequence of statements where each statement is either an axiom, a previously proved theorem, or a logical consequence of earlier statements in the sequence — and the final statement is the theorem we wanted to prove.

A proof is not a plausibility argument, not a very strong intuition, not a long list of examples that all work out. It is a *logical guarantee*. A single counterexample refutes a universal claim ('all primes are odd' is refuted by 2). A million confirming examples do not prove one. This asymmetry is fundamental: proofs establish universal truths; examples can only establish particular ones.

A Direct Proof Example

Let us prove: *if n is an even integer, then* n^2 *is an even integer.*

Proof: Suppose n is an even integer. By definition of 'even,' $n = 2k$ for some integer k. Then $n^2 = (2k)^2 = 4k^2 = 2 \times (2k^2)$. Since $2k^2$ is an integer (products of integers are integers), this shows n^2 is equal to 2 times an integer, i.e., n^2 is even. ■

Notice the structure: we started from the hypothesis (n is even), used the definition of 'even,' performed algebraic steps each of which is justified by arithmetic properties, and arrived at the conclusion (n^2 is even). Every step is explicit and justified.

Proof by Contradiction

Proof by contradiction (or *reductio ad absurdum*) is a powerful technique: assume the negation of what you want to prove, then derive a logical contradiction. Since the negation leads to nonsense, it must be false, so the original statement must be true.

The classic example: $\sqrt{2}$ is irrational. We want to prove it cannot be written as p/q with p, q integers. Assume for contradiction that $\sqrt{2} = p/q$ where p and q have no common factors (we can always reduce a fraction to lowest terms). Squaring: $2 = p^2/q^2$, so $p^2 = 2q^2$. Since p^2 is even, p itself must be even (by our direct proof above). Write $p = 2m$. Then $(2m)^2 = 2q^2$, so $4m^2 = 2q^2$, giving $q^2 = 2m^2$. So q is also even. But then p and q are both even — they share the factor 2 — contradicting our assumption that they have no common factors. Contradiction. Therefore $\sqrt{2}$ is irrational. ■

Quantifiers: For All and There Exists

Mathematical statements often range over entire collections of objects. We need precise language for this, and that is what **quantifiers** provide.

The **universal quantifier** ∀ ('for all,' 'for every') says that a property holds for every element of a collection. $\forall n \in \blacksquare, n + 0 = n$ says 'for every natural number n, $n + 0 = n$.' This is a claim about every natural number at once.

The **existential quantifier** ∃ ('there exists,' 'there is') says that at least one element has a property. $\exists x$ such that $x^2 = 4$ says 'there is a number x whose square equals 4.' (In fact $x = 2$ and $x = -2$ both work.)

The order of quantifiers matters enormously. 'For every person, there exists a person who is taller' means every person has someone taller than them — plausible. 'There exists a person who is taller than every person' means someone is taller than everyone — impossible. Same words, different order, completely different meaning.

Negating quantified statements follows specific rules. The negation of 'all swans are white' (∀ swans, it is white) is 'there exists a swan that is not white' (∃ a swan, it is not white). The negation of a universal is an existential, and vice versa. Getting this right is crucial for correct mathematical reasoning.

Exercises

3.1 Write the converse and contrapositive of: 'If a shape is a square, then it has four equal sides.' Is the converse true? Is the contrapositive true?

3.2 Identify the logical error: 'All elite athletes train hard. Jordan trains hard. Therefore Jordan is an elite athlete.'

3.3 Use the truth table for $P \Rightarrow Q$ to verify that the contrapositive (NOT Q ⇒ NOT P) always has the same truth value as $P \Rightarrow Q$.

3.4 Is 'This statement is false' a proposition? Try to assign it a truth value and explain what goes wrong.

3.5 Give an example from everyday life of: (a) a condition that is sufficient but not necessary, (b) a condition that is necessary but not sufficient, (c) a condition that is both necessary and sufficient.

3.6 In the proof that $\sqrt{2}$ is irrational, what is the role of the assumption that p/q has no common factors? What would go wrong if we did not make this assumption?

3.7 Negate the following statements, using the rule that $\neg\forall = \exists\neg$ and $\neg\exists = \forall\neg$: (a) Every prime number is odd. (b) There exists a real number x such that $x^2 < 0$.

Chapter 4

Sets — Collections of Things

If mathematics is a tower, sets are the ground floor. Almost every concept in modern mathematics — numbers, functions, geometric spaces, probability, logic itself — can be defined in terms of sets. Understanding sets is not just one topic among many; it is close to understanding the language in which the rest of mathematics is written. This chapter introduces sets carefully, because the concept repays careful attention.

What Is a Set?

A **set** is a collection of distinct objects, considered as a single entity. The objects in a set are called its **elements** or **members**. Two features of this definition deserve emphasis.

First, *distinct*: a set does not contain duplicates. {1, 2, 2, 3} and {1, 2, 3} are the same set. Repetitions are ignored. Second, *no inherent ordering*: {1, 2, 3} and {3, 1, 2} are also the same set. Unlike a list or a sequence, a set has no 'first' or 'second' element. What matters is solely which objects are members.

Sets can contain anything. The set {cat, 7, √2, Monday} is a perfectly valid mathematical set, even if it is not a natural or useful one. The set of all even natural numbers is a set. The set of all countries in the world is a set. The set of solutions to an equation is a set. Sets are completely general.

Why This Definition?

The definition of a set sounds almost trivially simple. So why is it so important? Because it is the simplest possible notion of a collection — and from this simple notion, combined with the membership relation (is this object in this set?), virtually all of mathematics can be constructed.

In the foundations of mathematics, this programme — building everything from sets — was largely carried out in the early 20th century. The natural numbers, the integers, the rationals, the reals, functions, relations, ordered pairs — all can be defined purely in terms of sets and the single relation 'is a member of.' The details are intricate and we will not go into them here, but the fact that it is possible is philosophically significant: mathematics reduces to one single primitive concept.

Notation

The standard notation for listing a set uses curly braces: {1, 2, 3} is the set whose elements are 1, 2, and 3. Sets are typically named with capital letters: A = {1, 2, 3}, B = {red, green, blue}, C = $\{x \mid x$ is a prime number$\}$.

Roster and Set-Builder Notation

The notation {1, 2, 3} — listing all elements explicitly — is called **roster notation** or **tabular form**. It works well for small, finite sets. For larger or infinite sets, we use **set-builder notation**: {*x* | property of *x*}, read as 'the set of all *x* such that [property].'

Examples of set-builder notation: {*n* | *n* is a natural number and *n* is even} = {2, 4, 6, 8, ...}. $\{x \mid x^2 = 9\} = \{3, -3\}$. {*x* ∈ ■ | 0 ≤ *x* ≤ 1} = the set of all real numbers between 0 and 1, inclusive.

Membership

To say an object is a member of a set, we use the symbol ∈ (read 'is an element of' or 'is in'). So 2 ∈ {1, 2, 3} means '2 is in {1, 2, 3}' — true. To say something is NOT in a set, we use ∉: 5 ∉ {1, 2, 3} — true.

The Empty Set

The **empty set** — written ∅ or {} — is the unique set containing no elements at all. It might seem like a pointless concept, but it is crucial. The empty set is to set theory roughly what zero is to arithmetic: a 'nothing' that is necessary for the system to be complete.

The empty set is unique: there is only one set with no elements. Why? Because two sets are equal when they have exactly the same elements. If two sets both have no elements, they have the same elements (vacuously), so they must be equal. There is exactly one empty set.

Subsets and Power Sets

Set A is a **subset** of set B (written A ⊆ B) if every element of A is also an element of B. We say A is 'contained in' B. If A ⊆ B and B ⊆ A, then A = B — the two sets have exactly the same elements. This gives us a way to prove two sets are equal: show that each is a subset of the other.

Some important subset facts: every set A is a subset of itself (A ⊆ A), because every element of A is trivially in A. The empty set ∅ is a subset of every set (∅ ⊆ A for any A), because there are no elements of ∅ that could fail to be in A.

A set A is a **proper subset** of B (written A ⊂ B) if A ⊆ B and A ≠ B — in other words, A is contained in B but is missing at least one element that B has.

The Power Set

The **power set** of A, written P(A) or 2^A, is the set of *all* subsets of A. This sounds abstract, but let's see a concrete example.

If A = {a, b}, then the subsets of A are: ∅ (the empty set), {a}, {b}, and {a, b} itself. So P({a, b}) = {∅, {a}, {b}, {a, b}} — a set with four elements. If A has 2 elements, its power set has $2^2 = 4$ elements. If A has 3 elements, its power set has $2^3 = 8$ elements. In general, if A has *n* elements, P(A) has 2^n elements — hence the notation 2^A.

Set Operations

Just as we combine numbers with arithmetic operations, we can combine sets with set operations. These operations are the tools of set theory, and they have natural connections to the logical connectives of Chapter 3.

Union (A ∪ B)

The **union** A ∪ B is the set of all elements that are in A, or in B, or in both. Formally: A ∪ B = {x | x ∈ A OR x ∈ B}.

Notice the connection to Chapter 3: union corresponds to logical OR. If A = {1, 2, 3} and B = {3, 4, 5}, then A ∪ B = {1, 2, 3, 4, 5}. The element 3 appears in both sets but only once in the union, because sets contain no duplicates.

Intersection (A ∩ B)

The **intersection** A ∩ B is the set of elements that are in *both* A and B. Formally: A ∩ B = {x | x ∈ A AND x ∈ B}.

Intersection corresponds to logical AND. With A = {1, 2, 3} and B = {3, 4, 5}: A ∩ B = {3}. Only 3 belongs to both sets. If two sets share no elements — like {1, 2} and {3, 4} — their intersection is ∅. Such sets are called **disjoint**.

Set Difference (A \ B)

The **set difference** A \ B (also written A − B) is the set of elements that are in A but NOT in B. Formally: A \ B = {x | x ∈ A AND x ∉ B}.

With A = {1, 2, 3} and B = {3, 4, 5}: A \ B = {1, 2}. We start with A and remove anything that B also contains. Note: A \ B is generally different from B \ A. Here, B \ A = {4, 5}.

Complement

When we are working within a fixed **universal set** U (the 'universe' of all objects under discussion), the **complement** of A — written A' or A^c — is U \ A: everything in U that is not in A.

For example, if U = {1, 2, 3, 4, 5} and A = {1, 3, 5}, then A' = {2, 4}. The complement of A contains exactly the elements of U that A misses. Note that (A')' = A — complementing twice gets you back where you started.

De Morgan's Laws

There are two beautiful laws connecting complement with union and intersection. They are called **De Morgan's Laws**:

- (A ∪ B)' = A' ∩ B'
- (A ∩ B)' = A' ∪ B'

The first says: the complement of a union is the intersection of the complements. In words: if something is NOT in (A or B), then it is neither in A nor in B. The second says: the complement of an intersection is the union of the complements. These are direct set-theoretic analogues of the logical laws from Chapter 3: NOT (P OR Q) = (NOT P) AND (NOT Q), and NOT (P AND Q) = (NOT P) OR (NOT Q). This connection is not coincidental — it reflects the deep relationship between set theory and logic.

Venn Diagrams

The operations above have beautiful visual representations called **Venn diagrams**, invented by the English logician John Venn in 1880. In a Venn diagram, sets are represented as overlapping circles inside a rectangle (the rectangle represents the universal set U).

The overlapping region of two circles represents the intersection. The entire combined region of both circles represents the union. The region inside circle A but outside circle B represents A \ B. The region outside circle A (but inside the rectangle) represents A'.

Venn diagrams are excellent for building intuition about set operations. They are also useful for checking whether two expressions are always equal (shade both expressions and see if you get the same region) or for solving logic puzzles. They are less useful for proofs — for rigorous mathematics, you use the formal definitions.

Why Sets Matter: A Deeper Look

Let us be more concrete about why sets are the foundation of mathematics. Consider the natural numbers. In Peano's axiom system (Chapter 1), natural numbers are defined by axioms. But in set theory, we can actually *construct* them. We define 0 = ∅ (the empty set). We define 1 = {∅} = {0}. We define 2 = {∅, {∅}} = {0, 1}. In general, each natural number is defined as the set containing all previous natural numbers.

This might look like a trick, but it shows something remarkable: you do not need to assume numbers exist. Given only set theory, you can construct them. Similarly, functions can be defined as sets of ordered pairs — we will return to this in Chapter 7. Geometric points can be defined as pairs of real numbers. Everything mathematical can, in principle, be built from sets.

Cardinality: The Size of Sets

The **cardinality** of a set is its 'size' — how many elements it contains. We write |A| for the cardinality of A. |{a, b, c}| = 3. |∅| = 0. For finite sets, cardinality is straightforward: just count the elements.

Infinite sets are more interesting. The cardinality of the natural numbers ■ is written ■$_0$ (aleph-null). A set is called **countably infinite** if it has the same cardinality as ■ — meaning its elements can be listed in a sequence: first element, second element, third element, and so on, hitting every element eventually.

Surprisingly, the integers ■ are also countably infinite, despite appearing twice as large: we can list them as 0, 1, −1, 2, −2, 3, −3, ... Every integer appears eventually. The rational numbers ■ are also countably infinite — this can be shown by a clever zigzag argument for listing all fractions. This is counterintuitive: the rationals are dense (between any two rationals is another rational), yet they can still be listed in a sequence.

But Georg Cantor proved in 1874 that the real numbers ■ are *not* countably infinite. There is no way to list all real numbers in a sequence — any attempt will always miss some. The cardinality of ■ is strictly larger than the cardinality of ■. There are, in a precise sense, more real numbers than natural numbers, even though both are infinite. This was one of the most shocking results in the history of mathematics, and Cantor faced considerable hostility from contemporaries who could not accept it. Today it is a cornerstone of mathematics.

Exercises

4.1 Let $A = \{1, 2, 3, 4\}$ and $B = \{3, 4, 5, 6\}$. Find: (a) $A \cup B$ (b) $A \cap B$ (c) $A \setminus B$ (d) $B \setminus A$ (e) $|A \cup B|$

4.2 List all subsets of $\{a, b, c\}$. How many are there? Does the formula 2^n correctly predict this number? What are all subsets of the empty set?

4.3 Use set-builder notation to describe: (a) all integers divisible by 3, (b) all real numbers strictly between 0 and 1, (c) all positive real numbers x satisfying $x^2 < 4$.

4.4 Verify De Morgan's first law with a specific example. Let $U = \{1,2,3,4,5,6\}$, $A = \{1,2,3\}$, $B = \{3,4,5\}$. Compute $(A \cup B)'$ and $A' \cap B'$ and check they are equal.

4.5 Two sets are disjoint if their intersection is $\varnothing$. Give an example of three sets A, B, C where A and B are disjoint, B and C are disjoint, but A and C are not disjoint.

4.6 The natural numbers ■ and the integers ■ are both infinite. Cantor would say they have the same cardinality (both countably infinite). What is your intuitive reaction to this? Does 'size' mean the same thing for infinite sets as for finite ones?

Chapter 5

Numbers — A Tower Built in Layers

Ask someone what a number is and you will most likely get a puzzled look. Numbers seem so fundamental, so obvious, that the question appears almost silly. But the history of mathematics is, in large part, the history of numbers being invented and reinvented — each time to solve a problem the previous system could not handle. Numbers are not a single, fixed thing. They are a family of increasingly powerful systems, each one built to fill a gap that its predecessor left.

The Natural Numbers (N)

The **natural numbers** are the counting numbers: 1, 2, 3, 4, 5, ... (Some conventions include 0 in the natural numbers; we'll discuss 0 separately.) They are denoted N (from the Latin *naturales* or German *natürliche Zahlen*). These are the numbers you use to count objects, measure discrete quantities, and keep score. They are the oldest mathematical objects, present in every known human culture.

With natural numbers, you can add and multiply freely: the sum or product of two natural numbers is always another natural number. We say the naturals are **closed** under addition and multiplication. But subtraction causes problems. 5 – 3 = 2 — fine. But 3 – 5 = ? There is no natural number that answers this question. The naturals are not closed under subtraction. To fix this, we need to extend the system.

The natural numbers were axiomatised by Giuseppe Peano in 1889, using five elegant axioms. The key ideas: there is a starting number (usually 1 or 0); every natural number has a unique successor (the 'next' number); and the principle of induction (if a property holds for 1 and for n+1 whenever it holds for n, then it holds for all natural numbers). Mathematical induction — proof by showing a domino effect propagates through all naturals — follows directly from Peano's axioms.

Zero: A Surprisingly Late Arrival

Before we get to negative numbers, a word about zero. Zero as a number — not just a placeholder in positional notation — was a profound conceptual step. The Babylonians had a placeholder symbol (like our '0' in 302, indicating an absent column) but did not treat it as a number you could do arithmetic with.

The systematic treatment of zero as a number first appears in Indian mathematics, particularly in the work of Brahmagupta in the 7th century CE. He gave rules for arithmetic with zero: 'The sum of zero and zero is zero.' 'A debt minus zero is a debt. A fortune minus zero is a fortune.' However, even Brahmagupta struggled with division by zero — he incorrectly gave 0/0 = 0, a confusion that took centuries to resolve.

Zero is fascinating mathematically: it is the additive identity ($a + 0 = a$ for all a), it is even, its square is zero, it has no multiplicative inverse, and dividing by it is undefined. These properties make zero simultaneously simple and special.

The Integers (Z)

The **integers** extend the natural numbers by including zero and the negative numbers: ..., −3, −2, −1, 0, 1, 2, 3, ... They are denoted Z (from the German *Zahlen*, 'numbers'). Now subtraction always works: $3 - 5 = -2$.

The inclusion of negative numbers took a long time historically. Chinese and Indian mathematicians accepted them by around the 7th century CE, representing positive numbers as one colour and negative numbers as another. European mathematicians were more resistant — many in the 16th and 17th centuries called negative numbers 'fictitious,' 'absurd,' or 'nonsensical.' Even the great mathematician Blaise Pascal wrote in 1659 that nothing is more absurd than a number less than nothing.

The resistance faded as the practical utility became undeniable — debts, temperatures below zero, coordinates to the left of an origin — and as mathematicians discovered that negative numbers followed the same algebraic rules as positive ones. The rules of arithmetic extend consistently and cleanly. There was nothing 'absurd' about them; they were just unfamiliar.

Properties of the Integers

The integers are closed under addition, subtraction, and multiplication: any integer plus, minus, or times any integer gives another integer. But division still causes problems: $7 \div 2$ is not an integer. For that, we need to go further.

An important concept with integers is **divisibility**. We say a divides b (written $a \mid b$) if there is an integer k such that $b = ak$. For example, $3 \mid 12$ because $12 = 3 \times 4$. But 3 does not divide 13 (since 13/3 is not an integer).

The study of divisibility properties of integers — which numbers divide which, how to factor integers, properties of prime numbers — is called **number theory**. It is one of the oldest and richest branches of mathematics, with deep connections to cryptography (the mathematics of secure communication depends crucially on properties of large prime numbers).

The Rational Numbers (Q)

The **rational numbers** (denoted Q from 'quotient') are all numbers that can be expressed as a ratio p/q where p and q are integers and $q \neq 0$. Rationals include all integers (since $n = n/1$), all proper fractions (1/2, 3/4, 7/8), and all mixed numbers.

A key fact: a real number is rational if and only if its decimal expansion either terminates (comes to an end) or eventually repeats in a cycle. $1/4 = 0.25$ (terminates). $1/3 = 0.333...$ (repeating). $1/7 = 0.142857142857...$ (repeating cycle of length 6). These decimal representations are familiar from school, but the *why* — the connection to the fraction form — is worth understanding: when you perform long division with a divisor q, there are only q possible remainders (0 through $q-1$).

Eventually a remainder must repeat, at which point the decimal pattern cycles.

Why Can't We Divide by Zero?

Division by zero is undefined — not just 'really difficult' or 'infinity,' but simply has no meaning. Here is why. Division is defined as the inverse of multiplication: $a \div b = c$ means $a = b \times c$.

So $5 \div 0 = c$ would mean $5 = 0 \times c$. But $0 \times c = 0$ for any number c. There is no number that, multiplied by zero, gives 5. Division by zero has no answer. It is not that the answer is 'infinity' — infinity is not a number in ordinary arithmetic. It simply has no answer.

What about $0 \div 0$? This would mean $0 = 0 \times c$, which is satisfied by every number c. So it has infinitely many 'answers,' which is just as bad. Division requires a unique answer; 0/0 has none or all, neither of which is acceptable.

The Irrational Numbers

Despite the fact that the rational numbers are densely packed on the number line (between any two rationals, there is always another rational), they still do not fill the entire line. There are 'gaps' — points that no fraction reaches. These are the **irrational numbers**: numbers that cannot be expressed as any ratio of integers.

We proved in Chapter 3 that $\sqrt{2}$ is irrational. In fact, $\sqrt{p}$ is irrational for any prime p. But the most famous irrational numbers are:

Pi (π): Approximately 3.14159265358979..., defined as the ratio of a circle's circumference to its diameter. This ratio is the same for all circles. Pi was shown to be irrational by Johann Lambert in 1768. It was later shown to be *transcendental* — not only irrational but not satisfying any polynomial equation with rational coefficients. Computers have calculated over 100 trillion decimal places, and no repeating pattern has emerged (nor can it, by proof).

Euler's number (e): Approximately 2.71828182845904..., arising in exponential growth, compound interest, and calculus. Euler's number satisfies the remarkable property that the function e^x is its own derivative — it is the unique function equal to its own rate of change. Like π, e is both irrational and transcendental.

The golden ratio (ϕ): Approximately 1.61803398874..., defined as $(1 + \sqrt{5})/2$. It appears in the proportions of pentagons, in Fibonacci sequences, and (despite popular claims) moderately in art and architecture. It is irrational but algebraic (it satisfies $\phi^2 = \phi + 1$).

The Real Numbers (R)

The **real numbers** (denoted R) are formed by taking all the rational numbers and all the irrational numbers together. Geometrically, the real numbers correspond exactly to all the points on a continuous number line — no gaps, no holes.

The real numbers have a crucial property that the rationals lack: **completeness**. Informally, the real numbers have no gaps. Any sequence of real numbers that is converging — getting closer and closer to some value — actually converges to a real number. This is not true of the rationals: you can construct a sequence of rationals (by computing better and better decimal approximations to √2) that converges towards √2, but √2 is not rational. The sequence converges 'to a gap' in the rationals. The real numbers fill all such gaps.

Completeness is what makes calculus rigorous. The formal definition of a limit — the foundation of calculus — requires completeness: you need to know that the value a sequence is converging to actually exists as a number. This is why calculus is built on the real numbers, not the rationals.

Complex Numbers: One More Extension

The real numbers solve all the obvious arithmetic problems, but one equation has no real solution: $x^2 + 1 = 0$. No real number squared gives -1, because squaring any real number gives a non-negative result.

Rather than declaring this impossible, mathematicians extended the number system one more time. We introduce the **imaginary unit** i, defined by the single rule $i^2 = -1$. A **complex number** has the form $a + bi$, where a and b are real numbers. When $b = 0$, we get an ordinary real number — so the reals are a subset of the complex numbers.

Complex numbers are not just an abstract trick. They are essential in physics (quantum mechanics uses complex amplitudes), engineering (electrical circuit analysis uses complex impedance), signal processing (Fourier analysis uses complex exponentials), and many branches of pure mathematics. The equation $e^{i\pi} + 1 = 0$ (Euler's identity) connects the five most fundamental constants of mathematics in one elegant formula, and it only makes sense in the complex numbers.

The Tower in Summary

The tower of number systems forms a chain of inclusions:

■ ⊆ ■ ⊆ ■ ⊆ ■ ⊆ ■

Each step extends the previous system to solve a problem the previous one could not: the integers solve $3 - 5$; the rationals solve $7 \div 2$; the reals provide completeness; the complex numbers solve $x^2 = -1$. Does the tower stop at complex numbers? For most of mathematics, yes — the complex numbers are 'algebraically closed': every polynomial has roots there. But extensions to quaternions, octonions, and beyond have been studied, each sacrificing some algebraic property (quaternions lose commutativity; octonions lose associativity) in exchange for new structure.

Exercises

5.1 Classify each of the following as natural, integer, rational, irrational, or complex — and state whether it belongs to each category in the tower: 7, −3, 4/5, √2, 0, π, √9, −7/2, $2 + 3i$.

5.2 Give an example of a calculation impossible with natural numbers but possible with integers. Then one possible with rationals but not integers. Then one possible only with reals (or complex numbers).

5.3 Why can't you divide by zero? Give two explanations: (a) using the concept of 'inverse of multiplication,' (b) by considering what happens to 1/x as x gets closer and closer to zero from both positive and negative sides.

5.4 The decimal expansion of 1/7 = 0.142857142857... has a repeating block of length 6. What is the decimal expansion of 2/7? Of 3/7? What pattern do you notice?

5.5 Without a calculator, explain why √9 is rational but √11 is not. Then explain why √9 is both rational and (trivially) a natural number.

5.6 Euler's identity $e^{i\pi} + 1 = 0$ is often called 'the most beautiful equation in mathematics.' Look up what each symbol means, and explain in plain words what this equation is saying about the relationship between exponential growth and circular motion.

5.7 The real numbers are 'complete': every sequence converging to a value actually converges to a real number. The rationals are not complete. Give an intuitive example of a sequence of rationals that 'converges to a gap' in the rationals.

Chapter 6

Properties and Structure — The Rules of the Game

Numbers are interesting, but what makes them truly powerful is how they *behave* — the structural properties that govern the operations we perform. These properties are not arbitrary rules that someone decided to impose. They are consequences of how addition and multiplication are defined, and they carry enormous amounts of information about why algebra works the way it does. This chapter unpacks those properties carefully, and ends by gesturing at the beautiful generalisations they suggest.

Commutativity

An operation is **commutative** if swapping the order of the inputs does not change the output. For addition and multiplication of numbers:

$$a + b = b + a$$

$$a \times b = b \times a$$

These hold for all real (and complex) numbers. $3 + 5 = 5 + 3 = 8$. $4 \times 7 = 7 \times 4 = 28$. When you shuffle the terms of a sum, the total is unchanged. When you rearrange the factors in a product, the result is unchanged. Seems obvious — but subtraction and division fail:

$5 - 3 = 2$, but $3 - 5 = -2 \neq 2$

$8 \div 4 = 2$, but $4 \div 8 = 1/2 \neq 2$

This is why the order in which you write a subtraction or division problem matters. 'The bus travels from A to B' is not the same journey as 'the bus travels from B to A.' Subtraction has direction. In a commutative world (addition and multiplication), direction does not matter.

Commutativity is also important in less obvious contexts. Matrix multiplication — an operation on arrays of numbers important in engineering and data science — is generally NOT commutative: the order of multiplication matters. This is a significant difference from scalar multiplication, and it is one of the first genuinely surprising things students encounter in linear algebra.

Associativity

An operation is **associative** if the way you group repeated applications does not change the result. For addition and multiplication:

$$(a + b) + c = a + (b + c)$$

$$(a \times b) \times c = a \times (b \times c)$$

For addition: (2 + 3) + 4 = 5 + 4 = 9, and 2 + (3 + 4) = 2 + 7 = 9. The grouping of parentheses does not matter; you get 9 either way. This is why we can write 2 + 3 + 4 without any parentheses: the result is unambiguous.

Subtraction is not associative: (8 − 3) − 2 = 5 − 2 = 3, but 8 − (3 − 2) = 8 − 1 = 7. The parentheses matter. When we write 8 − 3 − 2 without parentheses, we follow the convention of evaluating left to right, getting 3.

Associativity is what allows us to extend operations to arbitrarily many operands without specifying a grouping. 1 + 2 + 3 + 4 + 5 has a definite value (15) regardless of the order of operations, because addition is both commutative and associative. This seems obvious until you encounter an operation that is not associative.

Distributivity

Distributivity is the property that connects multiplication to addition. There are two forms (both hold for numbers):

$a \times (b + c) = a \times b + a \times c$ (left distributivity)

$(a + b) \times c = a \times c + b \times c$ (right distributivity)

Concrete example: 4 × (10 + 3). We could evaluate 10 + 3 = 13 first, then multiply: 4 × 13 = 52. Or we could 'distribute' the 4 over the sum: 4 × 10 + 4 × 3 = 40 + 12 = 52. Same answer. This is particularly useful when one of the additions is inconvenient and the multiplications are easier.

Distributivity is, arguably, the single most important property for algebra. Expanding brackets: $3(x + 4) = 3x + 12$. Factoring: $6x + 10 = 2(3x + 5)$. Multiplying out (FOIL): $(x + 2)(x + 3) = x^2 + 3x + 2x + 6 = x^2 + 5x + 6$. All of these use distributivity. Without it, algebra would not work.

Why does multiplication distribute over addition? At an intuitive level: 5 groups of (3 + 4) is the same as (5 groups of 3) + (5 groups of 4). You can distribute the groups among the parts of the sum. This intuition can be formalised into a rigorous proof from the axioms of arithmetic.

Identity Elements

An **identity element** for an operation is a special value that, when combined with any other value, leaves that value unchanged. It is the 'do nothing' element.

For addition: the identity element is **0**. $a + 0 = a$ and $0 + a = a$ for every number a. Adding zero does nothing — the value is unchanged.

For multiplication: the identity element is **1**. $a \times 1 = a$ and $1 \times a = a$ for every number a. Multiplying by 1 does nothing.

For subtraction, there is a right identity (a − 0 = a) but no left identity (0 − a = −a ≠ a in general). For division, there is a right identity (a ÷ 1 = a) but no left identity. This asymmetry is why subtraction and division are considered 'weaker' operations — they lack the full symmetry of addition and multiplication.

Inverse Elements

For every number *a*, there is an **additive inverse** $-a$ such that $a + (-a) = 0$ (the additive identity). The additive inverse 'undoes' addition — adding a and then adding −a gets you back to where you started.

Subtraction can be understood as adding the additive inverse: $a - b = a + (-b)$. This reframes subtraction as a special case of addition, which explains why addition is the 'more fundamental' operation.

For multiplication: every *nonzero* number *a* has a **multiplicative inverse** $1/a$ such that $a \times (1/a) = 1$. Division is multiplication by the inverse: $a \div b = a \times (1/b)$. The number 0 has no multiplicative inverse: there is no number 1/0, which is exactly why division by zero is undefined.

Order of Operations

When we write 2 + 3 × 4, there is a potential ambiguity: do we add first (giving 5 × 4 = 20) or multiply first (giving 2 + 12 = 14)? The standard convention — usually remembered as PEMDAS or BODMAS — says: Parentheses/Brackets first, then Exponents/Orders, then Multiplication and Division (left to right), then Addition and Subtraction (left to right). So 2 + 3 × 4 = 14.

This convention is not a mathematical law — it is a human agreement, like driving on the right side of the road. But the choice is not arbitrary. Treating multiplication as 'stronger' than addition reflects the distributive property: expressions like $ax + bx + c$ appear naturally in algebraic manipulation (they come from expanding polynomials), and writing them without parentheses at every multiplication is far more readable. The convention makes the most common expressions unambiguous without clutter.

Algebraic Structures: The Bigger Picture

Here is the remarkable payoff of thinking about these properties abstractly. The properties we have been discussing — closure, commutativity, associativity, distributivity, identity elements, and inverses — each capture something about the *structure* of an operation, independent of what the underlying objects are.

A **group** is any set with a single operation that is associative, has an identity element, and has inverses for all elements. The integers under addition form a group. The non-zero rationals under multiplication form a group. The symmetries of a square (rotations and reflections that map the square to itself) form a group — a fact that has no apparent connection to numbers but uses exactly the same abstract structure.

A **ring** is a set with two operations (like + and ×) where the set is an abelian (commutative) group under the first operation, the second operation is associative and distributes over the first, and there is a multiplicative identity. The integers form a ring. Polynomials with integer coefficients form a ring. So do 2×2 matrices with real entries (though matrix multiplication is not commutative in this ring).

A **field** is a ring where every nonzero element has a multiplicative inverse. The rationals, the reals, and the complex numbers are all fields. The integers are not — 2 has no multiplicative inverse in the integers, because 1/2 is not an integer.

Why does this matter? Because any theorem you prove about groups applies to every group — numbers, symmetries, cryptographic systems, all at once. This is abstraction in action: by identifying the common structure, you get enormously general results for free. We will return to this in Chapter 8.

Exercises

6.1 Show with a specific numerical example that subtraction is not commutative and not associative. For associativity, use three numbers.

6.2 Use the distributive property to: (a) expand $4 \times (10 + 3)$ and verify it equals 52; (b) expand $(x + 2)(x + 5)$; (c) factor $12x + 18y$.

6.3 What is the additive identity for the integers? The multiplicative identity? What is the additive inverse of 7? Of -3? What is the multiplicative inverse of 4? Of 1/3?

6.4 Why is there no multiplicative inverse for 0? Give an algebraic argument.

6.5 Evaluate each of the following using the correct order of operations: (a) $3 + 4 \times 2 - 1$ (b) $(3 + 4) \times (2 - 1)$ (c) $2^3 + 4 \times 2$ (d) $10 - 2 \times 3 + 1$

6.6 The set of all even integers forms a group under addition. Check: (a) is it closed? (b) is it associative? (c) does it have an identity? (d) does every element have an inverse? Does the set of all odd integers form a group under addition? Why or why not?

Chapter 7

Functions — Machines That Transform

If you had to identify the single most important concept in all of applied mathematics, functions would be a very strong candidate. Functions appear everywhere — in physics (position as a function of time), economics (price as a function of quantity), biology (population as a function of time), computing (a program is, abstractly, a function from inputs to outputs). Every equation in science is a function in disguise. Understanding functions — really understanding them, not just manipulating symbols — is the gateway to almost all of applied mathematics.

What Is a Function? The Precise Definition

A **function** f from a set A to a set B (written f: A → B) is a rule that assigns to each element of A *exactly one* element of B. A is called the **domain** and B is called the **codomain**. The element assigned by f to an input $x \in$ A is called the **value** of f at x, written $f(x)$.

The phrase 'exactly one' is the heart of the definition. Every input in the domain must produce an output (the function is total — defined everywhere on A). And each input produces only *one* output (the function is well-defined — no ambiguity). Different inputs can produce the same output — that is allowed. But one input cannot produce two different outputs — that is not allowed.

Think of a function as a machine: put in an input, get out an output. The machine is deterministic: the same input always gives the same output. Put in 3 today and get 7; put in 3 tomorrow and get 7 again. Functions are reliable transformations.

Functions as Sets

In Chapter 4 we said that all mathematical objects can be defined using sets. Functions are no exception. A function f: A → B can be formally defined as a *set of ordered pairs* $\{(x, f(x)) \mid x \in \text{A}\}$, with the property that no two pairs have the same first element (since each input has only one output). This set is called the **graph** of f.

This formal definition might seem overly abstract, but it is important: it shows that functions are mathematical objects in their own right, not just 'processes' or 'algorithms.' A function is a specific set of input-output pairs. Two functions are equal if and only if they produce the same output for every input.

Domain, Codomain, and Range

Three sets are associated with every function f: A → B:

Domain (A): The set of all valid inputs. Every element of A is assigned an output by f. Sometimes the domain is given explicitly; sometimes it is implied to be all inputs for which the rule makes sense.

Codomain (B): The set into which the function maps. The codomain is not necessarily hit by every element; it is the 'target' set, specifying what kind of output the function can produce in principle.

Range (or image): The set of outputs actually produced: $\{f(x) \mid x \in A\}$. The range is always a subset of the codomain, but may be smaller. For $f(x) = x^2$ with domain all real numbers and codomain all real numbers: the range is all non-negative reals, since no real number squared is negative.

The distinction between codomain and range is subtle but important in more advanced mathematics. For now, the key point is: the domain constrains valid inputs, and the range is what you actually get out.

When a function's rule involves a formula, the natural domain is often implied: it is all values for which the formula is defined. $f(x) = 1/x$: domain is all nonzero reals (division by zero undefined). $g(x) = \sqrt{x}$: domain is all non-negative reals (square root of negative number is not real). $h(x) = \log(x)$: domain is all positive reals (logarithm of zero or negative is undefined).

Examples of Functions

Functions appear throughout everyday life, long before anyone calls them 'functions.'

A vending machine: Input: a button code (B2, C3, etc.). Output: a snack. Each code gives exactly one item. This is a function. Domain: the set of valid codes. Codomain: the set of items in the machine.

A temperature conversion: $F(C) = (9/5)C + 32$ converts Celsius to Fahrenheit. Input: temperature in Celsius. Output: temperature in Fahrenheit. Each Celsius temperature gives exactly one Fahrenheit temperature. Domain: all real numbers (temperatures from absolute zero to as high as you like). Range: all real numbers.

A phone book (historically): Input: a name. Output: a phone number. But a phone book is only a function if each name maps to exactly one number — if two people have the same name, the 'phone book function' breaks down. Real phone books are not functions; they are relations.

Graphs of Functions

For functions from real numbers to real numbers, we can create a visual representation: the **graph**. The graph of f is the set of all points $(x, f(x))$ in the coordinate plane. The horizontal axis represents inputs and the vertical axis represents outputs. Each point on the graph records: 'when the input is x, the output is y.'

The coordinate plane has two perpendicular axes meeting at the origin (0, 0). Any point is described by its x-coordinate (horizontal position) and y-coordinate (vertical position). The point (3, 7) is 3 units right and 7 units up from the origin.

The **vertical line test**: a curve in the plane represents a function if and only if every vertical line crosses the curve at most once. If a vertical line crosses at two points, the corresponding input x would have two outputs — violating the definition. A circle fails the vertical line test (at most x

values, a vertical line crosses the circle at two points), confirming that the equation of a circle does not define a function from x to y.

Linear Functions

The simplest non-trivial functions are **linear functions**: $f(x) = mx + b$. Their graphs are straight lines.

m is the **slope** or **gradient**: the rate of change of the output per unit increase in input. If $m = 2$, then every time x increases by 1, $f(x)$ increases by 2. If $m = -3$, the function decreases by 3 for each unit increase in input. If $m = 0$, the output never changes — a horizontal line.

b is the **y-intercept**: the output when the input is zero, i.e., $f(0) = b$. It determines where the line crosses the vertical axis.

Linear functions model any relationship with a constant rate of change: distance as a function of time at constant speed, cost as a function of quantity at a fixed price, temperature conversion formulas, simple interest on a loan. They are the first functions you should reach for when modelling a real-world relationship — if plotting the data gives roughly a straight line, a linear function is appropriate.

Composition of Functions

Composition is the operation of chaining functions together: the output of one becomes the input of the next. If f: A → B and g: B → C, then the composition $g \blacksquare f$: A → C is defined by $(g \blacksquare f)(x) = g(f(x))$ — first apply f, then apply g to the result.

Example: $f(x) = 2x + 1$ and $g(x) = x^2$. $g(f(3)) = g(7) = 49$. $f(g(3)) = f(9) = 19$. The order matters — composition is generally not commutative. $g \blacksquare f \neq f \blacksquare g$ in general.

Composition is how complex transformations are built from simple ones. In computing, it is how function pipelines work: each function processes the output of the previous one. In calculus, the chain rule tells you how to differentiate composed functions. In abstract algebra, composition of functions is the fundamental operation that makes functions themselves form algebraic structures.

Injective, Surjective, Bijective

Three important properties of functions classify how inputs and outputs relate:

A function is **injective** (or one-to-one) if different inputs always give different outputs: if $f(a) = f(b)$, then $a = b$. Equivalently, no two inputs map to the same output. Example: $f(x) = 2x$ is injective. $g(x) = x^2$ is NOT injective (since $g(2) = g(-2) = 4$).

A function is **surjective** (or onto) if every element of the codomain is the output for at least one input: the range equals the codomain. Example: f: R → R defined by $f(x) = x^3$ is surjective (every real number is the cube of some real number). g: R → R defined by $g(x) = x^2$ is NOT surjective (no real number squares to give -4, for example).

A function that is both injective and surjective is called **bijective** (or a one-to-one correspondence). Bijections pair up the elements of the domain and codomain exactly: each element of the domain maps to a unique element of the codomain, and every element of the codomain is hit. Bijections have inverses — this leads to the concept of inverse functions.

Inverse Functions

A function f: A → B has an **inverse function** f^{-1}: B → A if and only if f is bijective. The inverse 'undoes' f: $f^{-1}(f(x)) = x$ for all $x \in$ A, and $f(f^{-1}(y)) = y$ for all $y \in$ B.

If $f(x) = 2x + 3$, then $f^{-1}(y) = (y - 3)/2$. Check: $f^{-1}(f(x)) = f^{-1}(2x + 3) = ((2x + 3) - 3)/2 = 2x/2 = x$. Correct.

The graphical interpretation: the graph of f^{-1} is the reflection of the graph of f across the line $y = x$. This is because if (a, b) is on the graph of f (meaning $f(a) = b$), then (b, a) is on the graph of f^{-1} (meaning $f^{-1}(b) = a$).

Exercises

7.1 If $f(x) = 3x - 2$, find $f(0)$, $f(1)$, $f(5)$, and $f(-3)$. What is the domain and range of f?

7.2 A taxi charges £3.50 base fare plus £2.20 per kilometre. Write this as a function $C(d)$. What is the domain? What does $C(0)$ represent? What distance corresponds to a fare of £16.10?

7.3 If $f(x) = 2x + 1$ and $g(x) = x^2$, compute (a) $f(g(3))$ (b) $g(f(3))$ (c) $f(f(2))$.

7.4 Explain the vertical line test in your own words. Which of the following are functions from R to R: (a) a straight line; (b) a parabola opening upward; (c) a full circle?

7.5 For each function, determine whether it is injective, surjective, or bijective (assume domain and codomain are both all real numbers R): (a) $f(x) = 3x + 1$ (b) $g(x) = x^2$ (c) $h(x) = x^3$.

7.6 Find the inverse of $f(x) = 5x - 3$. Verify that $f^{-1}(f(x)) = x$. What is the graphical relationship between f and f^{-1}?

7.7 Give an example of a real-world 'function' that is not injective (one input can give different outputs, or two inputs give the same output). What goes wrong with defining an inverse in that case?

Chapter 8

Abstraction — The Superpower of Mathematics

We have now covered seven chapters of foundations: the nature of mathematics, its language, its logical machinery, sets, numbers, structural properties, and functions. Each of these is important on its own. But this final chapter is about the single idea that ties them all together — the idea that, once understood, transforms how you think about mathematics, about science, and about problems in general. That idea is **abstraction**.

What Abstraction Means

Abstraction means stripping away specific, concrete details to reveal the underlying structure. When you abstract, you ignore what is *particular* about a situation and focus on what is *general* — on the pattern, the structure, the skeleton — rather than the specific instance.

This sounds simple but it is one of the most powerful intellectual moves available to a human mind. Almost every major advance in mathematics has come from abstraction: from seeing that two apparently different problems have the same underlying structure, and exploiting that commonality.

Consider a simple example. You observe that '3 + 5 = 5 + 3.' A specific fact. You observe it again: '17 + 42 = 42 + 17.' Another specific fact. You abstract: '$a + b = b + a$ for all numbers.' A general law, captured by variables. Then you go further: you notice this holds not just for numbers but for vector addition, for function composition in some cases, for the union and intersection of sets. The property — commutativity — is even more abstract than its numerical instance. It names a structural feature that can appear in any system with a binary operation.

Why Abstraction Is Powerful: Transferability

Here is the central payoff of abstraction: *prove something about the abstract structure, and it holds for every specific instance of that structure.* You do not need to re-prove it for each case.

Suppose you prove a theorem about groups: 'In any group, the identity element is unique.' This is a theorem about an abstract algebraic structure (see Chapter 6). It instantly applies to: the integers under addition, the nonzero rationals under multiplication, the symmetries of any geometric shape, the set of all invertible matrices under multiplication, the set of all permutations of a finite set, and every other group that has ever been or will ever be discovered — all with a single proof.

This is an extraordinary economy. Instead of proving the same thing over and over again for different systems, you prove it once at the abstract level and inherit the result in every specific case. The more abstract the theorem, the broader its reach.

Abstraction Across the Chapters

Looking back over this primer, we can see abstraction at work at every level.

From Numbers to Variables to Functions

In Chapter 2, we moved from specific numbers (3, 7) to variables (x, n). This is one step of abstraction: instead of one specific number, a variable represents any number. Then in Chapter 7, we moved from variables to functions $f(x) = 2x + 1$ — capturing a relationship between any input and its corresponding output. Then to properties of functions: linearity, injectivity, bijectivity. Each step is more general, more abstract, and more powerful.

From Specific Proofs to Proof Techniques

In Chapter 3, we proved specific results: that $\sqrt{2}$ is irrational, that the square of an even number is even. But we also introduced *proof by contradiction* as a general technique. This is abstraction applied to proof itself: instead of a specific argument, a template that can be applied to many different propositions. Learning the technique gives you a tool, not just a result.

From Arithmetic to Algebraic Structures

In Chapter 6, we identified specific properties of arithmetic: commutativity, associativity, distributivity. Then we noted these same properties appear in many other systems. The abstraction is to define a *group* or *ring* as any system having these properties — divorcing the properties from any particular implementation. The real numbers are one example of a field; the complex numbers are another; and there are many finite fields used in cryptography and coding theory.

From Specific Sets to Set Theory

Chapter 4 showed that we can abstract from specific collections (the set of prime numbers less than 100, the set of all students in a classroom) to the general concept of a set: any collection of distinct objects. At this abstract level, we can prove results that hold for all sets — De Morgan's laws, facts about cardinality — without worrying about what the elements are.

The Same Structure in Different Places

Here is one of the most remarkable phenomena in mathematics: the same abstract structure keeps appearing in completely different real-world contexts. When mathematicians see this, it is a signal that something deep is happening.

Exponential growth: The differential equation $dy/dt = ky$ — 'the rate of change of y is proportional to y' — governs bacterial population growth, radioactive decay, compound interest, the discharge of a capacitor, and the cooling of an object (Newton's law of cooling). These are bacteria, atoms, money, electricity, and heat — utterly different physical systems. Yet they share one mathematical structure, and understanding that structure tells you about all of them at once.

Symmetry: The mathematical study of symmetry — what transformations leave an object unchanged — uses group theory (Chapter 6). The symmetries of a snowflake, the symmetries of a crystal lattice, the symmetry principles of particle physics (which predict the existence of new particles before they are observed), and the symmetries of equations — all are instances of the same abstract group-theoretic structure.

Optimisation: Finding the maximum or minimum of a function (a problem in calculus) has the same abstract structure whether you are maximising a company's profit, minimising the error in a machine learning model, or finding the path of least time for a beam of light. The mathematics is identical; only the interpretation changes.

Abstraction and the Unreasonable Effectiveness of Mathematics

In 1960, the physicist Eugene Wigner published a famous essay called 'The Unreasonable Effectiveness of Mathematics in the Natural Sciences.' He observed that mathematical structures developed with no physical application in mind — often as pure intellectual exercises — repeatedly turned out to be exactly what physics needed, sometimes decades or centuries later.

Non-Euclidean geometry was developed as a purely abstract exercise in the early 19th century. Sixty years later, Einstein needed it for general relativity. Complex numbers were considered a convenient fiction for centuries. Quantum mechanics revealed that they describe the fundamental structure of reality. Matrix algebra was developed in the 19th century for abstract reasons. It turned out to be the language of quantum mechanics, computer graphics, and machine learning.

Why does this happen? Nobody fully knows. One view: mathematics is not arbitrary — we are exploring structures that are in some sense latent in the logic of consistent thought, and the physical universe also follows consistent logical principles, so the same structures appear in both. Another view: we tend to notice and develop mathematics that is useful, so the apparent mystery is partly selection bias. Both views have merit. The honest answer is that this is one of the deepest puzzles in all of philosophy.

Abstraction as a Way of Thinking

Here is perhaps the most important thing to take away from this entire primer. The real value of learning mathematics is not being able to calculate faster — computers do that incomparably better than any human. The real value is developing the cognitive habit of thinking abstractly: asking 'what is the underlying structure here?' rather than being trapped in the surface details.

This habit is transferable. A lawyer who thinks abstractly asks: 'What is the structural principle this case establishes, independent of its specific facts?' A scientist asks: 'What is the simplest model that captures the essential structure of this phenomenon?' A software engineer asks: 'What is the abstract interface here, separated from any particular implementation?' An economist asks: 'What is the underlying incentive structure of this situation?' All of these are the same intellectual move, learned through mathematics.

Mathematics also trains a complementary habit: precision. Abstraction without precision is vague hand-waving. What makes mathematical abstraction powerful is that it is precise: the concept of a group is not a vague notion of 'something like symmetry' — it is a specific list of properties with a specific definition. This precision is what allows proofs to be rigorous and conclusions to be certain.

Together — abstraction and precision — these are arguably the most powerful thinking tools a human mind can develop. And they are the core of what mathematics, at its deepest level, is about.

Where to Go from Here

This primer has laid the foundations. With them in place, the natural next territories are:

Algebra: Building on Chapters 2 and 6, algebra systematically develops the tools for solving equations and working with polynomial, rational, exponential, and logarithmic functions. It answers the question: given an equation, how do we find all its solutions, and what do those solutions tell us?

Calculus: The mathematics of continuous change, developed independently by Newton and Leibniz in the 17th century. The derivative captures the instantaneous rate of change of a function; the integral captures accumulation over a range. Together they form a coherent theory that describes motion, growth, decay, optimisation, and area. Built directly on the real number system and function theory developed in Chapters 5 and 7.

Linear Algebra: The study of vectors, matrices, and linear transformations. It generalises the linear functions of Chapter 7 to multiple dimensions, and turns out to be the mathematical backbone of machine learning, computer graphics, quantum mechanics, and data analysis. It is also where the algebraic structures of Chapter 6 — particularly fields — come into full play.

Probability and Statistics: The mathematics of uncertainty. Built directly on set theory (events are sets of outcomes, Chapter 4) and functions (probability is a function on sets). Statistics provides the tools to draw valid inferences from data — arguably one of the most important practical skills in the modern world.

Discrete Mathematics and Number Theory: The study of integers, combinatorics, graph theory, and algorithms. Building on the foundations of logic (Chapter 3) and sets (Chapter 4), and connecting to the computational world. Number theory in particular underlies all of modern cryptography.

Exercises

8.1 Think of two real-world situations that appear completely different on the surface but might share the same underlying mathematical structure. Describe the structure they share.

8.2 In what sense is a map of a city an 'abstraction'? What details does it preserve, what does it discard, and why? What would happen if a map tried to include everything?

8.3 The chapter describes the 'unreasonable effectiveness of mathematics.' Choose one historical example (from physics, engineering, or another field) where pure mathematics turned out to be unexpectedly applicable, and describe what happened.

8.4 Pick any concept from this primer (set, function, proof, axiom, etc.) and explain it in your own words to an imaginary 12-year-old, without using any technical vocabulary. What does this exercise reveal about your own understanding?

8.5 Consider the abstract concept of a 'group' (a set with an associative binary operation, an identity, and inverses). Find three very different real-world or mathematical systems that are groups. What does this tell you about the value of abstraction?

8.6 Of the five areas listed in 'Where to Go from Here' (algebra, calculus, linear algebra, probability/statistics, discrete mathematics), which interests you most right now, and what specific question would you most want it to help you answer?

8.7 Reflect on the primer as a whole: which chapter or concept felt most surprising or most changed your view of mathematics? Why?

Conclusion

You have now covered the foundations that school mathematics almost universally skips: what mathematics actually is (an axiomatic, logically rigorous system, not just arithmetic), its language (constants, variables, expressions, equations), the logic that underlies all of it (propositions, connectives, proofs), sets (the bedrock on which everything else is built), numbers (a tower of expanding systems, each motivated by a specific need), the structural properties of arithmetic (commutativity, associativity, distributivity, and the glimpse they give of abstract algebra), functions (the central concept of applied mathematics), and abstraction (the superpower that makes all of it both generalisable and deep).

Here is what that means in practical terms: you now understand why the rules of algebra work, not just what they are. You understand why negative numbers exist, why you can't divide by zero, what a proof actually is and why it matters, and what is really going on when someone talks about a function. These are not small things. Most people who 'did well' in school mathematics never got here.

A Final Word on Understanding

Understanding takes time. Do not be discouraged if some sections of this primer feel fuzzy even after a second reading. That fuzziness is informative — it points to exactly where your understanding is incomplete. Return to those sections. Try to explain the idea out loud. Try to apply it to a new example. Understanding is not an event; it is a process.

Mathematics has a reputation for being the hardest subject and the province of special minds. That reputation is mostly wrong, and largely a product of bad teaching. Given clear explanations, precise definitions, and time to think, mathematics is accessible to any patient, curious adult. The fact that you have read this far is good evidence that you are one.

Suggested Resources for Going Further

3Blue1Brown (YouTube channel, 3b1b.co): Exceptional visual explanations of mathematics, including the 'Essence of Calculus' and 'Essence of Linear Algebra' series. Perhaps the best introduction to these subjects available anywhere, and completely free.

'How to Prove It' by Daniel J. Velleman: The natural next step after this primer if you want to go deeper into logic and proof. Clear, well-paced, and beautifully written. It develops the tools introduced in Chapter 3 into a full toolkit for mathematical reasoning.

Brilliant.org: An interactive learning platform with excellent courses on mathematical thinking, calculus, linear algebra, probability, and more. Particularly good for building intuition through interactive problem-solving.

Khan Academy (khanacademy.org): If you want to rebuild procedural fluency — the ability to actually do calculations and solve standard exercises — Khan Academy is unmatched for this. It is not primarily conceptual, but the practice it provides is valuable once you have the conceptual foundations in place.

Mathematics is the art of giving the same name to different things.
— Henri Poincaré