LPS 105a/205a, PHIL 105a/205a, LSCI 145a: Set Theory and Mathematical Reasoning

Kenny Easwaran

December 2, 2025

Set theory is the branch of mathematics that makes precise our ordinary reasoning about collections of things. For much of the history of mathematics, there was very little formal study of sets — the relevant ideas were just treated intuitively. But starting in the 1870s, Georg Cantor began explicitly reasoning about the properties of infinite sets. His work radically reshaped philosophical and mathematical understandings of infinity in ways that were quite controversial for several decades.

In the early 20th century, set theory became an important branch of mathematics. While calculus, geometry, and linear algebra are the branches of mathematics that have had the most applicability in much of science and engineering, set theory (and related areas of logic) has been the most important branch of mathematics in philosophy and linguistics.

In mathematics, set theory has primarily become relevant as a framework within which all of the rest of mathematics is done. Sets are essential for the concept of number — what you are counting when you count is a *set* of things. Gottlob Frege's *Grundlagen der Arithmetik* (1884) (The Foundations of Arithmetic) and *Grundgesetze der Arithmetic* (vol. 1, 1893; vol. 2, 1903) (The Basic Laws of Arithmetic) showed how to define the counting numbers as certain sets of sets, and Richard Dedekind's *Was Sind und Was Sollen die Zahlen?* (1888) (What are the numbers, and what should they be?) showed how to define other types of numbers, like rational and real numbers, as sets of counting numbers. Although Bertrand Russell found a contradiction in Frege's set theory, Russell and Alfred North Whitehead developed a new set theory in their *Principia Mathematica* (vol. 1, 1910; vol. 2, 1912; vol. 3, 1913) and showed how to formulate most of the currently existing mathematics inside of it. Bartel Leendert van der Waerden's *Moderne Algebra* (vol. 1, 1930; vol. 2, 1931) reformulated all of abstract algebra in a set-theoretic language, and his use of the Axiom of Choice led to many new developments in mathematics, but also various controversies. But in recent decades, most mathematicians have found it useful to jump directly from some simple, intuitive set theory to the Axiom of Choice, and set theory is no longer considered as central and important as number theory, algebra, topology, geometry, and analysis.

In linguistics, set theory has been central to the development of semantics (the study of systematic linguistic meaning). It has also been relevant to pragmatics (the study of the purposes for which speakers use linguistic meaning) and syntax (the study of the systematic arrangement of words), but has been less important in morphology, phonology, and phonetics. This work began with Frege's *Begriffsschrift* (1879) (Concept Writing), in which he introduces modern quantifier logic, and uses sets as the meanings of concepts. Alfred Tarski then showed, in "Der Wahrheitsbegriff in den formalisierten Sprachen" (1935, Polish original in 1933) (The concept of truth in formalized languages) how to use a set-theoretic formulation to precisely define what it is for a sentence to be true. While Tarski's work concerned formal languages, Richard Montague showed how to use the same set-theoretic method to give a formal semantics for natural language in a series of three papers, "Universal Grammar" (1970), "English as a Formal Language" (1970), and "The Proper Treatment of Quantification in Ordinary English" (1973). Barbara Partee showed how to generalize this work, and set-theoretic concepts soon became relevant to many other aspects of meaning. Saul Kripke gave a basic treatment of modal verbs ("can", "may", "must", "should", etc.) in terms of sets of possible worlds, which Angelika Kratzer has turned into a fuller theory of conversational meaning, with sets representing not just the meanings of concepts and sets of possible worlds, but also information states, events, and other

important aspects of meaning. Irene Heim used set theory to represent conversational contexts, and showed how meanings can be understood in terms of changes in context rather than sets of worlds.

In philosophy, the importance of set theory primarily comes through its influence in mathematics and linguistics. The work of Frege, Russell, and Tarski resolved many ancient paradoxes around truth and infinity, but also introduced new paradoxes that have not fully been resolved. The role of set theory in providing a foundation for mathematics helped address some questions in the philosophy of mathematics, particularly about the existence of abstract objects, but also raised new ones. The role of set theory in semantics has also become important for philosophy of language, but also philosophy of mind and epistemology. The use of set theory in the understanding of infinity, and modal concepts like necessity and possibility, has taken on great importance in metaphysics, and thinking about the relationship of sets to their members has been relevant for the study of parts and wholes.

Each section of this document corresponds roughly to my plans for one week of the term. Each section will include some definitions and expository writing, some example questions with their answers worked out, and then starting with Week 1, some exercises. I will aim to have student volunteers to present some or all of the exercises in class, and students will be asked to turn in a written answer to one of these exercises each week. Feel free to choose a relatively easy question if you want practice being precise, or a relatively hard one if you want practice managing the steps in a more complex proof.

Part I does set theory informally, in the style of Cantor, leading up to the controversies around the Axiom of Choice, and the paradoxes that doomed Frege's approach to set theory. Although contradictions are possible within the informal set theory that we start with, Cantor was able to skillfully avoid them. In this first part of the class, we will sometimes point out these contradictions, but will carefully back away until it comes time to confront them.

Part II begins after we confront the paradoxes of informal set theory. This traces the development of modern axiomatic set theory, starting with the work of Ernst Zermelo in his 1908 paper, "Untersuchungen über die Grundlagen der Mengenlehre" (Investigations on the foundations of set theory) and its modification by Abraham Fraenkel. In the beginning of this section, we will mostly be re-proving the same results that we proved informally, but this time working within a formalized axiom system. Both formalized and informal proofs are parts of contemporary mathematical practice, so it will be helpful to spend some time with each. We will be able to see how the modern formalized system blocks the specific paradoxes that arose for the informal system (though some of the important results we will be getting close to at the end show that there can be no formal proof that contradictions are impossible, except by using even more questionable formal systems).

New terms that are being introduced will be given in *italics*.

Part I: Informal Set Theory

0 Sets as the meanings of words

Sets are collections of things. It's hard to be more precise about this until we get into axiomatic set theory, but we can work with examples. I'll do this by introducing a bit of terminology. The things that are in a set are called its *members* or *elements*. Everything that matters about a set is determined by which things are members of the set and which things are not members of the set. There are three ways that sets are often presented: entirely in words by saying "The set of ...", or explicitly by listing all the elements between two curly braces, or with the *set-builder notation* by using curly braces, a variable, a vertical line, and a description involving that variable that definitively specifies for each object whether it is a member of the set or not. For instance:

$$\begin{aligned}\text{The set of days of the week} &= \{\text{Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, Saturday}\}\\ &= \{x \mid x \text{ is a day of the week}\}.\end{aligned}$$

Although there are many other things that have members, like teams and committees and departments, these are importantly different from sets, in that they can be different even if they have the same members. Maybe there are two committees that both have Sue, John, and Tania as their only members — but if one of the committees has Sue as the chair and is in charge of determining the curriculum, while the other committee has John as the chair and is in charge of hiring, then they are different committees. But they have the same set of members. More clearly:

The set of days of the week whose name begins with "S" = The set of weekend days.

Note also that

$$\{\text{Saturday, Sunday}\} = \{\text{Sunday,Saturday}\} = \{\text{Saturday, Sunday, Saturday}\}.$$

Everything that is a member of one of them is a member of the others, so they are the same set, even though they are written differently.

The symbol "$\in$" is used to state that something is a member of a set (or with a slash to say it is not). For instance

$$\begin{aligned}\text{Ireland} &\in \text{The set of members of the European Union}\\ \text{Kenny} &\in \text{The set of faculty in the Department of Logic and Philosophy of Science}\\ \text{Canada} &\notin \text{The set of countries that have Spanish as an official language}\\ 5 &\notin \{x \mid x \text{ is even}\}\end{aligned}$$

There is one very significant set known as the *empty set*. This is the set for which nothing is a member. Although we could notate this set "{}", it's more standard to notate it as "$\emptyset$". Of course, this set also has other names, like "The set of faculty in LPS who are over 100 years old" and "The set of members of the European Union whose capital is in South America". (Why is there only one empty set?)

There are some important relations and operations among sets. If every member of A is also a member of B, we say that A is a *subset* of B or B is a *superset* of A, written as "$A \subseteq B$" or "$B \supseteq A$". Which of the following sets are subsets of each other?

- A =The set of even numbers
- $B = \{2, 6\}$
- $C = \{x \mid x < 7\}$
- D =The set of odd numbers

- E =The set of all positive integers
- $F = \emptyset$
- G =The set of all numbers that are not multiples of 4

Note that every set is a subset and a superset of itself. But if A is a subset of B and is not equal to B, we say A is a *proper subset* of B or B is a *proper superset* of A, and we usually write "$A \subset B$" or "$B \supset A$" with this meaning (like the difference between "$<$" and "$\leq$" or "$>$" and "$\geq$" — though there are some texts that use "$\subset$" to mean "$\subseteq$", which can be confusing).

For any two sets, x and y, we define the *intersection* of x and y ("$x \cap y$") to be the set of all things that are in both sets, and the *union* of x and y ("$x \cup y$") to be the set of all things that are in either of the sets. (You can remember which symbol is which because the $\cup$ is for Union and the $\cap$ is for iNtersection.) What is $A \cap C \cap E \cap G$? What is $B \cup C$? What is $A \cup D$? What is $B \cup F$? What is $B \cap F$?

We say that two sets are *disjoint* if their intersection is empty. Which pairs of sets from the ones above are disjoint?

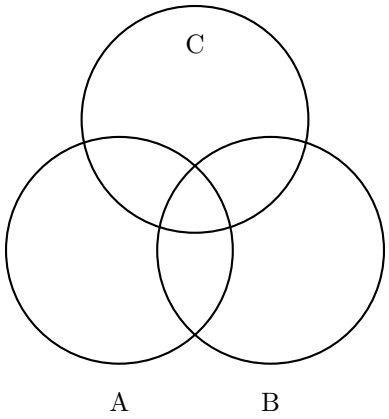

If we think of each letter as standing for the set of points inside the relevant circle, then which regions are $A \cup B$? $A \cup C$? $B \cup C$? $A \cap B$? $A \cap C$? $B \cap C$? $A \cap (B \cup C)$? $(A \cap B) \cup C$? $(A \cap B) \cup (A \cap C)$? $(A \cup C) \cap (B \cup C)$? $A \cap (B \cap C)$? $(A \cap B) \cap C$? $A \cup (B \cup C)$? $(A \cup B) \cup C$? (If you are familiar with the terms from algebra, you can think about commutativity, associativity, and distributivity.)

The important thing for a set is that for every object, it is either definitely a member of the set, or definitely not a member of the set. It doesn't matter if anyone knows what the members are, or even could know what the members are. But there must be a definite fact. Consider the set of students in this class who will some day win a Nobel prize or the set of cities that have un-discovered tyrannosaur fossils buried under them or the set of numbers that have never been written down in Arabic notation. These are perfectly fine sets. But some problematic definitions that don't work include definitions like "The set of tall people" or "The set of mountains". (Why are these problematic?)

Paradox 0.1 (The Sorites Paradox) Both of these premises seem plausible:

- A million grains of sand form a heap (σωρειτης).
- A heap of sand minus one grain is still a heap of sand.

But they together lead to the conclusion that there can be a heap of sand with no sand grains at all! (Why is that?)

A word that has some version of the sorites paradox is usually said to be *vague*. This is importantly different from an *ambiguous* word like "star", which has multiple distinct meanings (it can either refer to famous movie actors or to balls of hydrogen undergoing fusion in outer space — note that at least one of

these meanings is in fact vague). It is also importantly different from a *polysemous* word like "month", which has several related meanings (it can refer to one of the 12 subdivisions of the year, or to any period of about 30 days). It is common to replace a vague word with a *precisification*, some set that does have precise membership. For instance, we can replace the set of tall people with the set of people who are at least six feet tall, or the set of mountains with the set of geographic peaks that are over 5,000 feet high and not connected to any higher peak without going below that elevation.

Once we have precisified vague terms, we can think of one aspect of the meaning of most adjectives, nouns, and intransitive verbs as being the set of objects that they describe. We call this set the *extension* of the word. That is, the extension of "quick" is the set of quick things, the extension of "brown" is the set of brown things, the extension of "dog" is the set of dogs, the extension of "fox" is the set of foxes, the extension of "jumped" is the set of things that jumped, the extension of "sleeps" is the set of things that sleep.

The extensions of phrases composed out of adjectives and nouns can usually be thought of as the intersection of the extensions of the individual words. So the extension of "lazy fox" is the intersection of the set of lazy things and the set of foxes. The extension of "quick brown dog" is the intersection of the set of quick things, the set of brown things, and the set of dogs. (The extension of the word "the" is somewhat complicated, so we can't get to the extension of the "the quick brown dog" yet!)

However, there are some adjectives that don't work this way. One set of counterexamples includes words like "former", "fake", and "imitation", and some similar words like "suspected" and "alleged". (How would you characterize the similarities and differences between these words?) Another set of counterexamples include words like "typical", "excellent", "bad", and "tall". (What's the difference between Alex, who is a tall six-year-old, a typical boy, an excellent student, and a bad basketball player, ad Charlie, who is a tall basketball player, an excellent boy, and a bad student? Could he be a typical six-year-old?)

Paradox 0.2 (Grelling's paradox) Note that some words are a member of their own extension: "word", "pentasyllabic", "English", "familiar". Note that other words are not a member of their own extension: "dog", "monosyllabic", "French", "microscopic". Kurt Grelling invented the words "autologous" for words that are members of their own extension, and "heterologous" for words that are not members of their own extension. Is the word "heterologous" heterologous? (Also, is the word "autologous" autologous?)

1 Sets and numbers

What are numbers?

The natural numbers are the things we use to say how many of something there are when there are finitely many of them. $\mathbb{N} = \{0, 1, 2, 3, \dots\}$. Note that we have to include 0, because that is sometimes how many of something there are — it is how many unicorns are in the room as I type this.

What does it mean to say how many of something there are? Well, when do we care how many of something there are?

Think about the set of students in this class and the set of chairs in the room. Do we have the same number of students as chairs, or does one of these sets have more? You don't need to count — just look at whether there are any chairs left over, or any students left over, when everyone is sitting in their own chair!

This is the basic idea of numbers. Two sets have the same number of things when there is a *one-to-one correspondence* between them. We will define it precisely below.

What is a one-to-one correspondence between two sets? The best way to think of it is as a special type of set of ordered pairs.

What's an ordered pair? We write it with ordinary parentheses, and its two members. The important thing is that even though $\{x, y\} = \{y, x\}$, $(x, y) \neq (y, x)$, unless $x = y$. You might be familiar with ordered pairs as the coordinates of points in the plane. Remember that the point $(1, 0)$ is 1 unit to the right of the origin, while $(0, 1)$ is 1 unit above the origin. Note that $(0, 0)$, $(1, 1)$, and so on are ordered pairs as well — because of the ordering, it makes a difference if we have the same element twice.

If X and Y are sets, we can consider some set S of ordered pairs whose first member is in X and whose second member is in Y.

Here are the four conditions for a one-to-one correspondence:

- For every $x \in X$, there is at most one $y \in Y$ with $(x, y) \in S$ (Function)
- For every $x \in X$, there is at least one $y \in Y$ with $(x, y) \in S$ (Total)
- For every $y \in Y$, there is at most one $x \in X$ with $(x, y) \in S$ (One-to-one)
- For every $y \in Y$, there is at least one $x \in X$ with $(x, y) \in S$ (Onto)

Why are these the conditions we want for the set of students and chairs? Are there any conditions we don't really care about?

Exercise 1.1 For every set, there is at least one correspondence between it and itself.

Exercise 1.2 For a set with two elements there are precisely two correspondences between it and itself. For a set with three elements there are precisely six correspondences between it and itself.

Note that although the existence of a correspondence proves that two sets have the same size, the fact that one attempted correspondence fails doesn't mean that they don't have the same size. Imagine someone came into the room early and claimed 15 chairs for themself, and let everyone else fight over the rest. There would be lots of pairs with the one person attached to different chairs, so the Function condition would fail. For other people, they might either not get a chair (so that the Total condition fails) or might share a chair (so the One-to-one condition fails). But that doesn't mean that there can't be a correspondence — just that this isn't it.

We don't always try to establish a correspondence to show that two sets have the same number of things. For finite sets, we rarely do. Instead, we count.

What do we do when we count? We say each number in sequence, while picking out one element of the set. The last number we say is the number of things in the set.

We have used the numbers in two different ways — as we are counting, we are picking out a first, second, third, etc. member of the set. These are called *ordinal* number, and they represent positions in an ordering. We use the last number of the counting to say how many members are in the set. This is called a *cardinal* number.

Exercise 1.3 If you can count two sets, and get the same cardinal number, then there is a correspondence between the two sets.

Exercise 1.4 Counting a set establishes a correspondence between it and the natural numbers below its cardinality.

Note that counting only works for finite sets, but correspondences can exist even for infinite sets. But with infinite sets, some strange things can happen.

Galileo observed that there is a correspondence between the natural numbers and the perfect squares, even though the perfect squares are a proper subset of the natural numbers. Galileo used this to suggest that there is no meaning of "same size" for infinite sets, but Cantor was a bit more careful. He treated this correspondence as meaning that these sets have the same size even though one is a proper subset of the other.

Note that this means that establishing a one-to-one correspondence between X and a subset of Y doesn't necessarily mean that Y is larger than X — it might still be the same size. This attempt would just be like the attempt above where one person hogged some of the chairs.

In fact, Dedekind defined an infinite set as one that has a one-to-one correspondence with a proper subset. (Modern set theorists don't exactly use this definition — it turns out that it gets caught up in the controversy around the Axiom of Choice.)

We can still show some results without using counting.

Exercise 1.5 If there is a correspondence between A and B, and a correspondence between B and C, then there is a correspondence between A and C.

Exercise 1.6 If there is a correspondence between A and B, then there is a correspondence between B and A.

Theorem 1.1 If a set's cardinality is a natural number, then it does not have a correspondence with any proper subset.

Proof. By induction.

Base case: Prove that a set with 0 elements does not have a correspondence with any proper subset. (Trivial, because there are no proper subsets!) Induction step: Prove that if a set with n elements has no correspondence with any proper subset, then a set with $n+1$ elements does not either. This step takes more work. We show that *if* there *were* a correspondence of a set with $n+1$ elements with some proper subset, *then* we could use that to build a correspondence of a set with n elements to a proper subset. Since that latter correspondence is assumed not to exist, neither can the first. We do this in two cases. Let x be the extra element in our set of $n+1$ elements.

First case: Assume the correspondence between a set with $n+1$ elements and a proper subset left out x in its range. Then that means that x in the domain must have been sent to some element a of the original set in the range. Since it's a correspondence, no other element of the domain can have been sent to a. So delete x from the domain, and now we have a correspondence of the original set to some proper subset of itself, leaving out a.

Second case: Assume the correspondence between a set with $n+1$ elements and a proper subset included x in its range. Then there must be some element a of the original set that is left out of the range. Just deleting x from the domain and range doesn't give a correspondence — some element b of the domain must have been sent to x in the range, and x in the domain must have been sent to some element c of the range. But if we delete x from the domain *and* range, and send b to c, then now we have a correspondence between the original set and a proper subset leaving out a. □

Once Cantor noticed that an infinite set can have a correspondence to a proper subset, he started showing they can also have correspondences to sets that seem bigger.

Let $\mathbb{Z} = \{\ldots, -3, -2, -1, 0, 1, 2, 3, \ldots\}$ be the set of all integers.

Consider the set of all ordered pairs of natural numbers.

Cantor's Enumeration of $\mathbb{N} \times \mathbb{N}$

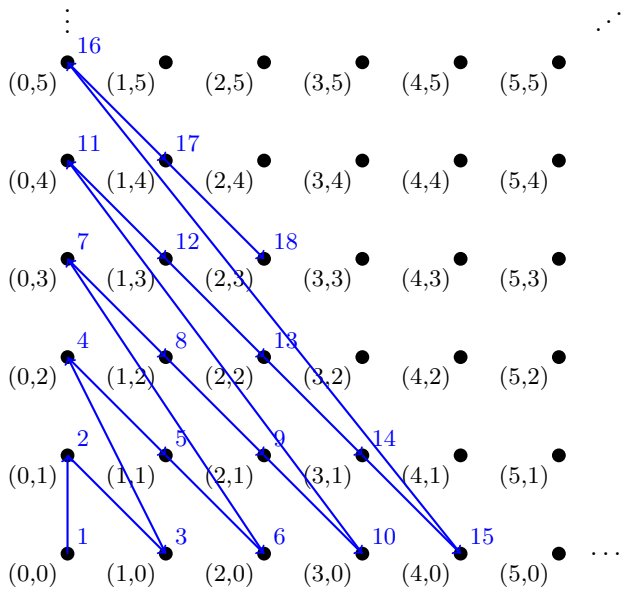

Each pair (n, m) gets a unique natural number following the blue path

2 Well-orderings, countable vs uncountable infinities

For natural numbers, we have already identified two different roles they play in counting. One is the *ordinal* number role, where they represent what position in the count an object has. When you count out all the elements of a set, you are effectively giving each element of that set an ordinal number as the first, second, third, etc. member of the set that you are identifying.

The other is the *cardinal* number role, which tells you how many objects there are in the set as a whole. This role is only played by the *last* number that you say when you count a set.

If you count up to 5 while counting the chairs in the room, and you count up to 5 when counting the people, then you know that the chairs and the people can be put into a correspondence. That's what it means to say the two sets have the same cardinal number. But the particular correspondence that we know must exist can be derived by using the ordinal numbers — the first person is assigned to the first chair; the second person is assigned to the second chair; and so on.

We have seen that for sets that can be put into correspondence with one of the natural numbers, there is no way to put that set into a correspondence with a proper subset. In particular, that means that no two natural numbers can be put in correspondence with each other. If a set has 1,352 elements, it is in correspondence with the set $\{1, 2, 3, \ldots, 1350, 1351, 1352\}$. But it can't also be put in correspondence with any proper subset of that set. That means that no matter how you start counting out elements of the set, there will always be some elements left if you haven't gone through all these numbers. And conversely, there will always be numbers from that set left if you haven't gone through all elements of the original set. Any way of giving out ordinal numbers up to the cardinal number will work out just fine.

However, this isn't true when we it comes to infinite sets. We have seen that the natural numbers can be put into correspondence with several proper subsets (like the even numbers and the squares) and also with some proper supersets (like the integers) and sets that feel like supersets (like the set of ordered pairs of natural numbers). If we're not careful when giving out ordinal numbers, we could accidentally give them all out before finishing the set we are trying to number. If we count the ordered pairs of natural numbers by just beginning with all the pairs whose second element is 0, we will run out of natural numbers before even getting to the second row of the chart! So we had to do our counting in a different order — either zigzag through the chart, like Cantor suggested, or be stingy when giving out natural numbers, only giving out the half of them at a time, using the odd numbers for the first row, half of the even numbers for the second row, half of the remaining numbers for the third row, and so on.

Every way of providing a correspondence to the natural numbers is a way of putting an infinite set in order so that it has a first element, a second element, a third element, and so on. For instance, these are the orderings of the evens, integers, and pairs of natural numbers under the correspondences to the natural numbers that we have come up with.

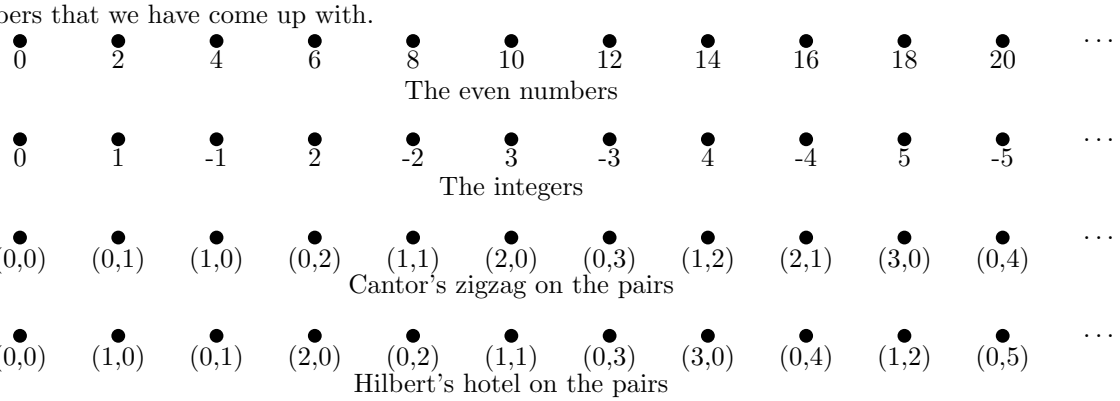

But there are other ways of ordering these infinite sets too. For instance, there is the standard ordering of the integers, and the "lexicographic" ordering of the ordered pairs of natural numbers. (We call it "lexicographic", because this is the principle that a dictionary is organized on — everything with the earliest first letter comes first, then everything with the second-earliest first letter, etc.)

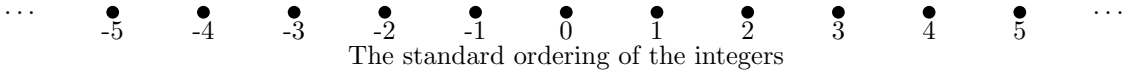

The standard ordering of the integers

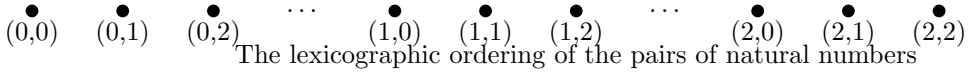

The lexicographic ordering of the pairs of natural numbers

Even though neither of these orderings is the same as the standard ordering of the natural numbers, Cantor noticed an important difference between the two orderings. The ordering on the integers is in a sense "slippery" — if you slide everything over, it still looks the same. But the lexicographic ordering on the pairs of natural numbers is not—even if you delete the names in the positions, you can figure out which position is which. There is still a first, and a second, and a third, and so on. But even after you run out of natural numbers, there's a first position beyond all those (where the pair $(1, 0)$ is going to go) and then a second, and then a third, and so on.

Just as Cantor generalized the concept of a cardinal number to infinite sets by seeing when they have correspondences, he generalized the concept of an ordinal number to positions in infinite orderings that aren't "slippery" in this way. To make this precise, Cantor defined the concept of a *well-ordering*: an ordering called a well-ordering when every non-empty subset has a first element. Intuitively, we sometimes think that we can't "identify positions in an ordering" if you can't count up to there. But Cantor's idea was that to define an ordering, all you need is a set of elements, and a rule that determines, for any two elements, which comes first. (Just like with membership in a set, this rule doesn't need to let us *know* which of two elements comes first—all that matters is that there *is a fact* for any two elements about which one comes first.) For this rule to be an ordering, it needs to be transitive—if x is before y and y is before z, then x is before z. And it should be anti-symmetric—nothing should come before itself.

Exercise 2.1 Show that any relation that is transitive (if x is before y and y is before z, then x is before z) and antisymmetric (there is no x with x coming before x) must be irreflexive (if x is before y then y is *not* before x).

We will use the symbol "$<$" for any ordering. Let's be precise and define the *lexicographic* ordering on pairs of natural numbers by saying $(x, y) < (z, w)$ iff $x < z$ or $x = z$ and $y < w$.

Consider the following sets of ordered pairs. What is the earliest element of each of these sets, in the lexicographic ordering?

- $\{(1, 3), (5, 2), (2, 7), (8, 0), (0, 4), (1, 5)\}$
- $\{2, 4), (2, 6), (2, 8), (2, 10), (1, 11), (1, 13), (3, 15)\}$
- The set of all pairs (m, n) where $m = 2n$
- $\{(3, 5), (3, 4), (3, 2), (2, 7), (2, 5), (2, 1), (5, 9), (4, 3), (5, 10)\}$
- $\{(10, 0), (9, 1), (8, 2), (7, 3), (6, 4)\}$

Think about how you found the earliest element of each of these sets.

Exercise 2.2 Show that the lexicographic ordering on ordered pairs of natural numbers is a well-ordering.

We can see that no well-ordering is "slippery" in the sense we've talked about—if there *were* some position that could move, then the set of positions that could move would be non-empty. By the well-ordering property, there must be a *first* such position. But this position can't move earlier in the ordering (since all earlier positions are fixed) and it can't move later in the ordering (or else there's nothing left to fill its position).

Cantor then came up with names for each of the positions in the lexicographic ordering of the pairs of natural numbers, and called them *ordinal numbers*.

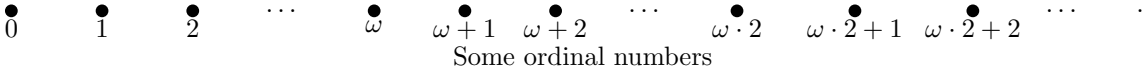

Some ordinal numbers

Let's try to explain what these names mean. The first step is just that he uses the letter "ω" to stand for the first infinite position in the well-ordering. But to explain the rest, we need to explain what the symbols "$+$" and "$\cdot$" mean—we are familiar with them for the natural numbers, but we need to explain what they mean in order to see how they generalize to the infinite ordinals.

The way we understand what "$m + n$" means, when m and n are both finite ordinals is by counting m spots, and then counting n more, and seeing what position you get to.

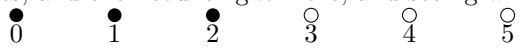

Illustration that $2 + 3 = 5$

And we can do this when one of the terms is infinite. When we take an infinite ordinal and add a finite ordinal to it, it works as we might expect.

$$(\omega + 1) + 3 = \omega + 4$$

But when we do it in the opposite order, some surprising things happen.

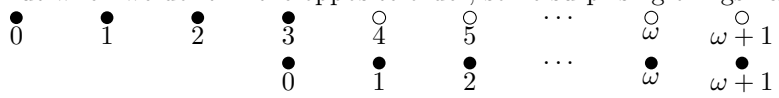

$$3 + (\omega + 1) = \omega + 1$$

Effectively, adding something infinite to the end of something finite "absorbs" the finite thing. If you add something infinite at the end of something infinite with a finite remainder, that finite remainder gets absorbed too.

Exercise 2.3 What is $(\omega + 1) + (\omega + 3)$? What is $(\omega + 2) + (\omega + 2)$? What is $(\omega + 3) + (\omega + 3)$? What is $(\omega + 6) + \omega$? Which, if any, of these are equal?

Multiplication of well-orderings

$\omega \cdot 2 \neq 2 \cdot \omega$

For any ordinal, there is a larger ordinal — just add one element at the end.

An *initial segment* of a well-ordering is a subset such that any element of the well-ordering that is before something in the initial segment is also in the initial segment.

For any two initial segments of well-orderings, if there is a correspondence between them, then that correspondence is unique.

Proof. Call the two initial segments A and B. If the correspondence between A and B is not unique, then there must be some element in one initial segment (let's say A) that can correspond to two different elements in the other. Because A is well-ordered, there must be some least element in A that corresponds to two different elements in B. Let's call this least element a, and let's call two of the elements it can correspond to in B b_1 and b_2, with $b_1 < b_2$. Because every element before a corresponds uniquely to some element in B, they must all correspond to elements less than b_1. So in the correspondence where a corresponds to b_2, there can be nothing that corresponds to b_1 — everything less than a corresponds to something less than b_1, while everything greater than a must correspond to something greater than b_2. Thus, this second "correspondence" is not a correspondence at all, so there is a contradiction. □

For any two well-orderings A and B, one of them corresponds to an initial segment of the other (since every well-ordering is an initial segment of itself, this could mean that they correspond to each other).

Proof. Consider the set of correspondences between initial segments of A and B. From the above result, we can see that any element in A that is in any such correspondence must correspond to the same element of B in all such correspondences. Thus, we can put all these correspondences in initial segments together to get a maximal correspondence between initial segments. If every element of A is in one of these correspondences, this means that A corresponds to an initial segment of B. If every element of B is in one of these correspondences, this means that B corresponds to an initial segment of A. So the only way the result could fail is if there are elements of both A and B that aren't in any initial segment with a correspondence. Since A and B are both well-ordered, we can take the least non-corresponding element a of A and the least non-corresponding element b of B. Every element before a and b is corresponded by the maximal correspondence. But adding a and b to these maximal initial segments gets an even larger initial segment with a correspondence. This contradicts the maximality, so these must not both exist, so one of the two well-orderings must correspond to an initial segment of the other. □

Theorem 2.1 For any non-empty set of ordinals, there is a least ordinal in the set.

Proof. Let S be a set of ordinals. Let α be any element of this set. If no elements of α are in S, then α is the least ordinal in the set. If there are some elements of α that are in S, then since α is well-ordered, there must be a least element of α that is in S. This must then be the least ordinal in S. □

Thus, the ordinals are themselves well-ordered.

Paradox 2.1 (Burali-Forti Paradox) Consider the set Ω of all ordinals. Because Ω is a well-ordered set of ordinals, it must itself be an ordinal. But now consider $\Omega+1$. This is itself an ordinal, so it must be a proper initial segment of Ω. But Ω is a proper initial segment of it. This would mean that Ω is in correspondence with a proper initial segment of itself. But this is a contradiction!

2.1 Power sets

One more method we have of generating larger sets is the operation of taking the "power set". The *power set* of a set is the set of all subsets of a set. Here I list out all the elements of the power set of $\{0, 1, 2\}$.

$\{\quad\quad\quad\quad\quad\quad\}=\emptyset$
$\{\quad\quad\quad\quad\quad 2\quad\}$
$\{\quad\quad\quad 1\quad\quad\quad\}$
$\{\quad\quad\quad 1,\quad 2\quad\}$
$\{\quad 0\quad\quad\quad\quad\quad\}$
$\{\quad 0,\quad\quad\quad 2\quad\}$
$\{\quad 0,\quad 1\quad\quad\quad\}$
$\{\quad 0,\quad 1,\quad 2\quad\}$

One of the things to notice—every time you add one more element to a set, you *double* the number of elements in the power set. Basically, the power set of the new set will contain all the subsets of the old set (none of which have the new element), and it will also contain one set corresponding to each of those subsets of the old set, with the new element added. For each subset of the old set, there are two subsets of the new set—one with and one without the new element.

Now we will show that even for infinite sets, the power set is a bigger set than the original set.

Theorem 2.2 (Cantor's Theorem) There is no correspondence between the set $\mathbb{N}$ of all natural numbers, and the power set of the natural numbers, $\mathcal{P}(\mathbb{N})$.

Proof. First, let's show this for the power set of the set of natural numbers. The way we will do this is by showing that no function f from $\mathbb{N}$ to $\mathcal{P}(\mathbb{N})$ can be a correspondence. In particular, we will show that no such function can have the onto property. So here is some particular function f from $\mathbb{N}$ to its power set:

$$
\begin{array}{lccccccccc}
f(0) = & \{ & 0, & 1, & 2, & 3, & 4, & \ldots & \} \\
f(1) = & \{ & & & 2, & 3, & & \ldots & \} \\
f(2) = & \{ & 0, & & 2, & & 4, & \ldots & \} \\
f(3) = & \{ & 0, & 1, & & & 4, & \ldots & \} \\
f(4) = & \{ & & 1, & & 3, & & \ldots & \} \\
\vdots & & & & & \vdots & & & \\
\end{array}
$$

In order to show that this set is not onto, we need to find a particular element of $\mathcal{P}(\mathbb{N})$ that is not in the range of this function. We do this by going down the diagonal, and reversing everything, to generate this "diagonal' set D:

$$
\begin{array}{lccccccccc}
f(0) = & \{ & \boxed{0,} & 1, & 2, & 3, & 4, & \ldots & \} \\
f(1) = & \{ & & \boxed{} & 2, & 3, & & \ldots & \} \\
f(2) = & \{ & 0, & & \boxed{2,} & & 4, & \ldots & \} \\
f(3) = & \{ & 0, & 1, & & \boxed{} & 4, & \ldots & \} \\
f(4) = & \{ & & 1, & & 3, & \boxed{} & \ldots & \} \\
\vdots & & & & & \vdots & & \ddots & \\
D = & \{ & & 1, & & 3, & 4, & \ldots & \}
\end{array}
$$

We can see that D is not in the range of f. If it were, then it would appear in some row of the table. But every row intersects the diagonal, and D has the *opposite* on the diagonal, making clear that it *can't* be the same as the set on any row of the diagram. □

Exercise 2.4 Show that this same argument works for *any* set, not just the set of natural numbers. That is, let f be a function from S to $\mathcal{P}(S)$, let $D = \{x \in S \mid x \notin f(x)\}$, and show that D isn't in the range of f.

It is straightforward to see that there is a correspondence between $\mathbb{N}$ and a *subset* of $\mathcal{P}(\mathbb{N})$ — in particular, let $f(n) = \{n\}$, the set containing just that natural number. It is straightforward to see that this is a correspondence. (Again, this works for any set, not just $\mathbb{N}$.) Thus, we can see that the power set of any infinite set is larger than that set.

Paradox 2.2 Let V be the set of all sets. We can see that $\mathcal{P}(V)$ is the set of all sets of sets, and thus all its members are sets, so it is a subset of V. But Cantor's theorem says that $\mathcal{P}(V)$ cannot be in correspondence with any subset of V.

Note that the correspondence between $\mathcal{P}(V)$ and a subset of V is based on the identity function $f(x) = x$. Bertrand Russell plugged this into Cantor's proof and noted the more specific paradox.

Paradox 2.3 (Russell's Paradox) Let $D = \{x \in V \mid x \notin x\}$. That is, D is the set of all sets that are not members of themselves. Either D is a member of itself or it is not. If it is, then it does not satisfy the definition, so it should not be a member of itself. But if it is not, then it does satisfy the definition, so it should be. Either way, there is a contradiction.

To distinguish the sizes of infinity, we say that infinite sets that *can* be put in correspondence with the natural numbers are *countable* infinite sets, while the ones that can't are called *uncountable* infinite sets. However, it's important to note that although all countable infinite sets are the same size (because they can all be put in correspondence with each other), there are many different sizes of uncountable set, since the power set of one uncountable set is always an even larger uncountable set.

But this raises a natural question — is there any set whose size is somewhere in between that of $\mathbb{N}$ and $\mathcal{P}(\mathbb{N})$? Or in general, for any infinite set S, is there a set whose size is somewhere in between that of S and $\mathcal{P}(S)$? This turns out to be a deep and difficult question.

3 Many countable infinities

We have already shown that the set $\mathbb{N} \times \mathbb{N}$ of ordered pairs of natural numbers can be put in correspondence with the set $\mathbb{N}$ of natural numbers, and in fact we have shown this in two ways. One way involved writing them out in a grid, and numbering them in zigzag fashion through the grid, so that every pair was assigned a natural number. The other way involved assigning the odd numbers to the pairs beginning with 0 (leaving the infinitely many even numbers in reserve), and then assigning every other even number to the pairs beginning with 1 (leaving the infinitely many multiples of 4 in reserve), and then assigning every other multiple of 4 to the pairs beginning with 2 (leaving the infinitely many multiples of 8 in reserve), and so on. On this method, the pair (m, n) is represented by the natural number $2^m \cdot (2n + 1)$. It's clear that each pair is represented by a distinct natural number.

These methods naturally generalize to the ordered triples of natural numbers. Tilt into a diagonal pyramid and do a three-dimensional zigzag. Powers of 2, powers of 3, powers of 5. Group the triples as a pair followed by a natural.

Although we know that uncountable sets exist, we will talk about many countable sets

Ordered pairs of naturals

Rational numbers

But define the real numbers.

Consider a type of measurement like distance, time, volume, mass. (I won't consider temperature, for a reason you'll see in a moment.)

What does it mean to measure?

For any two amounts, either one is more than the other, or they are the same amount. For any two amounts, there is an amount you get by sticking them together, *concatenating* them, which I will write as "$x \circ y$".

If $x > y$, then $x \circ a > y \circ a$.

$x \circ y = y \circ x$ — you can see this by turning the objects around.

For any x, there is a y such that $x = y \circ y$ — you can divide it in two.

If $x < y$ then there is a z with $x \circ z = y$.

Continuity property — if there are two sequences with $x_0 < x_1 < x_2 < \cdots < x_n < \cdots < y_n < \cdots < y_2 < y_1 < y_0$, then there is a z that is greater than every x_i and less than every y_i. This is important to ensure that everything can be measured.

Archimedean property — for any $x < y$, either $x \circ x > y$ or $x \circ x \circ x > y$ or $x \circ x \circ x \circ x > y$ or ....

Now we can prove that the set of possible measurements is not countable.

Ordered triples

Ordered n-tuples

Exercise 3.1 Any subset of a countable set is countable.

Exercise 3.2 The union of two countable sets is countable.

Exercise 3.3 The union of countably many sets that are each countable is countable.

The set of all finite sets of natural numbers

The set of all co-finite sets of natural numbers

The set of sentences in English

The set of all computer programs

The set of all definitions of sets of natural numbers

Paradox 3.1 (Richard's Paradox) Consider the set of all definable sets of natural numbers. This set must be countable, because we just argued that the set of definitions of sets of natural numbers is countable. So this set can be enumerated. Now let us use the diagonal argument to define a set that is not on that list. This appears to be a definable set of natural numbers that is not in the set of all definable sets of natural numbers.

4 Comparing the real numbers and the power set of the natural numbers

Explain what the real numbers are

Measurement theory - not decimal notation

Real number measurement axioms.

There is a consistent way to say when one measurement is larger than another.

Measurements can be concatenated.

For any x, y, $x \circ y = y \circ x$ and $x \circ y > x$.

For any x, there is a y with $y \circ y = x$.

For any $x > y$ there is a z with $x = y \circ z$.

Continuity: If $x_0 < x_1 < x_2 < \cdots < x_n < \cdots < y_n < \cdots < y_2 < y_1 < y_0$, then there is a z in between them.

Archimedean: For any two measurements, there is a finite number of copies of one that is greater than the other.

These mean that everything can be measured.

Every measurement can be represented by a sequence of approximations by fractions. Decimal notation is one particular sequence of approximations.

Exercise 4.1 The Continuity axiom for measurement states that for every sequence of nested intervals there is a measurement that lies inside all of these intervals. Consider a sequence of nested intervals defined by decimal numbers that differ only by one in their last digit, with one more digit in the endpoints defining each interval (like the intervals $[3, 4], [3.1, 3.2], [3.14, 3.15], [3.141, 3.142], [3.1415, 3.1416], \ldots$ or the intervals $[0, 1], [0.9, 1.0], [0.99, 1.00], [0.999, 1.000], [0.9999, 1.0000], \ldots$). Use the Archimedean axiom to show that there cannot be two distinct measurements that are both contained in all of these intervals. (As a byproduct, this will show that $0.999\cdots = 1$.)

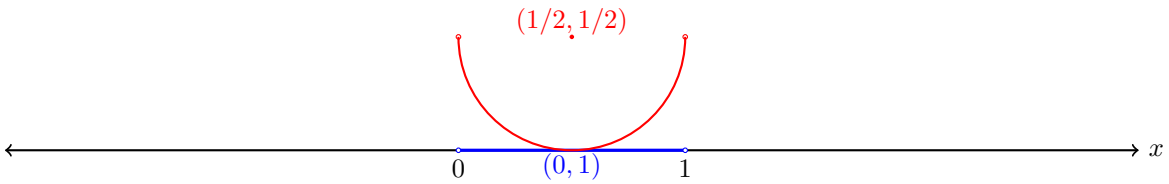

Exercise 4.2 Show that there is a correspondence between the set of all real numbers between 0 and 1 (blue in the above diagram) and the set of points on the semicircle of radius $1/2$, centered at $(1/2, 1/2)$ (red in the above diagram), not including the endpoints of either the interval or the semicircle. Make sure to show that all four properties of a correspondence hold.

Exercise 4.3 Show that there is a correspondence between the set of points on the semicircle of radius $1/2$, centered at $(1/2, 1/2)$ (not including the endpoints) and the set of *all* real numbers. Make sure to show that all four properties of a correspondence hold.

Exercise 4.4 Show that there is a correspondence between the set of points in any particular finite interval of real numbers (not including endpoints) and the set of *all* real numbers.

Explain binary notation. Show that binary and decimal notation are two equally good ways of representing real number measurements (and in fact, any finite base will also work equally well).

Exercise 4.5 Using binary notation, show that there is a function from the real numbers in the interval from 0 to 1 (including endpoints) to $\mathcal{P}(\mathbb{N})$, the power set of the natural numbers, that is total, and one-to-one, but *not* onto. (Remember that in binary notation, $0.011111\cdots = 0.1$.)

Exercise 4.6 Using notation in any base greater than 2, show that there is a function from $\mathcal{P}(\mathbb{N})$ to the set of real numbers in the interval from 0 to 1 (including endpoints), that is total, and one-to-one, but *not* onto.

By the next theorem, we can prove that there is in fact a correspondence between these two sets (and thus by earlier exercises, between $\mathbb{R}$ and $\mathcal{P}(\mathbb{N})$).

Theorem 4.1 (Cantor-Schroeder-Bernstein Theorem) If f is a function from X to Y that is total, and one-to-one, but not onto, while g is a function from Y to X that is total, and one-to-one, but not onto, then there is a correspondence between X and Y.

Proof. Consider the following diagram, illustrating the features given in the statement of the theorem. The fact that f sends X in a total, one-to-one, but not onto way into Y is illustrated by the fact that X's shape is duplicated inside Y (but not covering it), and conversely for g. Every region of a given shape and shading on one side is sent to a region of the same shape and shading on the other — but note that f sends the outermost dotted square of X to the dotted square just inside Y, while g sends that dotted square just inside Y to the *second* dotted square in X.

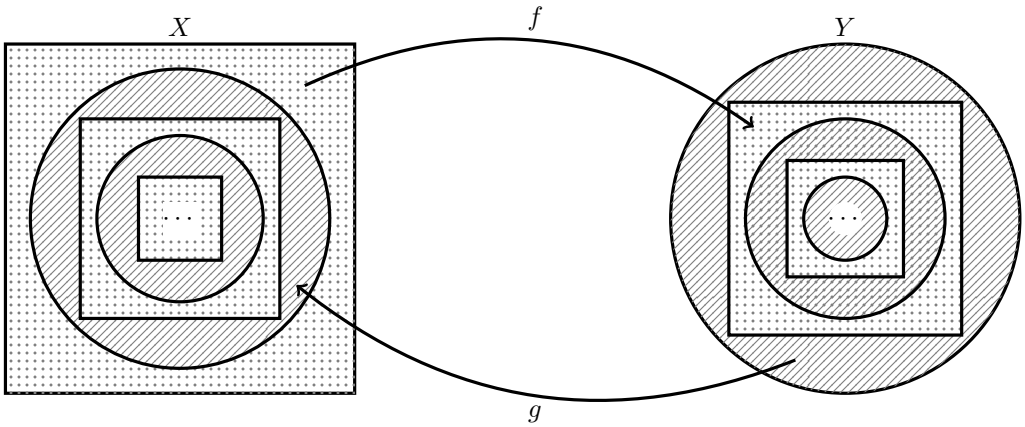

We will construct the correspondence between all of X and all of Y by putting these things together. We define a function h by letting it behave like f for the dotted regions of X, and letting it behave like g^{-1} for the striped regions of X. This turns out to be the correspondence.

To make this precise, and not rely on the visual aid, define subsets in the following way. Let

$$X_0 = \{x \in X \mid \neg\exists y{\in}Y(g(y) = x)\}, Y_0 = \{y \in Y \mid \neg\exists x{\in}X(f(x) = y)\}.$$

Then for every natural number n, define

$$X_{n+1} = \{x \in X \mid \exists y{\in}Y_n(g(y) = x)\}, Y_{n+1} = \{y \in Y \mid \exists x{\in}X(f(x) = y)\}.$$

Finally, define

$$X_\infty = \{x \in X \mid \neg\exists n(x \in X_n)\}, Y_\infty = \{y \in Y \mid \neg\exists n(y \in Y_n)\}.$$

Now we make some observations. Observe that if $i \neq j$, then X_i is disjoint from X_j, and X_∞ (and similarly for Y) — proving this is a tedious exercise in proof by induction. Observe also that for each i, f is a correspondence between X_i and Y_{i+1}, while g is a correspondence between Y_i and X_{i+1} — this is actually straightforward to prove (f and g are both functions that are total and one-to-one on these subsets, because they are functions that are total and one-to-one on the whole sets X and Y, and they are onto because of how X_{i+1} and Y_{i+1} are defined). Finally, note that, for the same reason, f is a correspondence between X_∞ and Y_∞, and g is a correspondence between Y_∞ and X_∞ (though they don't necessarily behave in any similar way on these sets).

At last, we construct the correspondence. Let $h(x) = \begin{cases} f(x) & \text{if } x \in X_n \text{ with } n \text{ even} \\ g^{-1}(x) & \text{if } x \in X_n \text{ with } n \text{ odd} \\ x & \text{if } x \in X_\infty. \end{cases}$
The observations we have already made ensure that h is a correspondence between X and Y.

□

If you don't want to use the Cantor-Schroeder-Bernstein theorem to prove that there is a correspondence between $\mathbb{R}$ and $\mathcal{P}(\mathbb{N})$, we can also do it more directly.

Proof. We say that a set of natural numbers is *cofinite* if there are only finitely many natural numbers left out. Recall that we have shown that the set of finite sets of natural numbers is countable. It is also straightforward to see that the set of cofinite sets of natural numbers is countable — there is a correspondence between these two sets by taking any finite set S to the cofinite set $\bar{S} = \{n \in \mathbb{N} \mid n \notin S\}$.

Consider the function that takes a set of natural numbers and returns a real number by saying $f(S) = \sum_{n \in S} \frac{1}{2^n}$. This takes a set of natural numbers to the real number whose notation in binary has a 1 in every position corresponding to a member of the set, and a 0 in every position corresponding to a number that is not in the set. Observe that f is a function from $\mathcal{P}(\mathbb{N})$ to the set of real numbers between 0 and 1 (including endpoints), that is total, and onto, but it is not one-to-one.

Note that whenever S is finite, there is a cofinite set T that is mapped to the same real number, and vice versa. Note also that the set of real numbers that have two sets mapped to them in this way is the set of real numbers strictly between 0 and 1 that can be written as fractions whose denominator is a power of 2. However, we have already shown that the set of finite sets of natural numbers is countable, and so is the set of cofinite sets, so the union of these two sets is also countable. But the set of real numbers strictly between 0 and 1 that can be written as fractions whose denominator is a power of 2 is a subset of $\mathbb{Q}$, which we have already shown is countable. Thus, there must exist a function g that is a correspondence between the set of finite or cofinite sets of natural numbers and this set of real numbers.

Using f on the set of sets of natural numbers that are *not* finite or cofinite, and g on the set of sets of natural numbers that *are* finite or cofinite, thus gives a correspondence between the set of *all* sets of natural numbers, and the entire interval from 0 to 1. □

Now that we have shown that there is a correspondence between the set $\mathbb{R}$ of real numbers, and the power set of the natural numbers, $\mathcal{P}(\mathbb{N})$, we can return to the question, is there any other infinity that is between the size of $\mathbb{N}$ and this one? Because $\mathbb{R}$ is the "continuous" set of real numbers, Cantor called this question "the continuum problem". "The Continuum Hypothesis" is the hypothesis that there is no infinity in between these two, and its negation is the hypothesis that there is an infinity in between these two.

5 The Axiom of Choice, the paradoxes

Axiom (Axiom of Choice) If X is a set of sets, and for every $S \in X$, S is non-empty, then there exists a function f such that for each $S \in X$, $f(S) \in S$.

Cantor used the Axiom of Choice to prove that every set can be well-ordered.

Theorem 5.1 (The Well-Ordering Theorem) Every set has a well-ordering.

Proof. Let S be any set. We will construct a function g from the ordinals to S that is one-to-one and onto. This will yield a correspondence between S and some ordinal, and this correspondence defines a well-ordering of S.

To do this, start by letting X be the set of all non-empty subsets of S. By the Axiom of Choice, there is a function f that "chooses" one element of each such set. We will use this function to define g. Intuitively, we will use f to "pick" elements of S one-by-one, "until we have picked them all", and this will generate the well-ordering. Of course, this is just a metaphor — there is no actual *process* here, but thinking of it in terms of a process that runs possibly infinitely long helps understand what is going on.

Let us consider a particular set W of functions whose domain is some set of ordinals, and whose range is some subset of S. Not every such function will be in the set, but just the functions satisfying certain conditions. In particular a function h will be a member of W if it satisfies two conditions. First, for every ordinal α, if $h(\alpha)$ is defined, then so is $h(\beta)$ for all $\beta < \alpha$. Second, for every ordinal α, consider the set $S_\alpha = \{g(\beta) \mid \beta < \alpha\}$. That is, S_α is the set of all elements of S picked in the first α steps. Let $\bar{S_\alpha} = \{x \in S \mid x \notin S_\alpha\}$ be the remainder, the set of all elements of S that haven't been picked in the first α steps. The second condition on h is that for every ordinal α, if $h(\alpha)$ is defined, then $h(\alpha) = f(\bar{S_\alpha})$.

Now notice a few things. First, each $h \in W$ is one-to-one. This is because each $h(\alpha)$ is the result of using f to pick some element that hadn't occurred earlier in the function.

Second, if $h_1 \in W$ and $h_2 \in W$, then for any ordinal α for which $h_1(\alpha)$ and $h_2(\alpha)$ are both defined, it must be the case that $h_1(\alpha) = h_2(\alpha)$. (If this weren't the case, then there would have to be some *earliest* α for which these two functions differ, but then since $h_1(\beta) = h_2(\beta)$ for all $\beta < \alpha$, this would mean that the two functions would have to have the same S_α, and so both $h_1(\alpha) = h_2(\alpha) = f(\bar{S_\alpha})$.)

If we think of each function as a set of ordered pairs, then the union of all these sets of ordered pairs is also a function, h_∞ whose domain is some set of ordinals, and whose range is some subset of S. This function h_∞ must satisfy the defining property, and thus $h_\infty \in W$. That is h_∞ is the unique largest element of W.

Now let us consider the domain of h_∞. This domain must in fact be a set, since we got it from a set by just applying a function and union operations. This domain can't be the set of all ordinals, because there is no set containing all ordinals. (This was the result of the Burali-Forti paradox.)

Thus, there must be some ordinal α such that $h_\infty(\alpha)$ is undefined. Consider the set $S_\alpha = \{h_\infty(\beta) \mid \beta < \alpha\}$, as above. If $S_\alpha \neq S$, then $\bar{S_\alpha}$ is non-empty, so there would be a possibility of extending h_∞ by defining $h_{\infty+1}(\alpha) = f(\bar{S_\alpha})$. But then this $h_{\infty+1}$ would have to be in W, contradicting the definition of h_∞ as the largest element of W.

So in order for $h_\infty(\alpha)$ to be undefined we must have $S_\alpha = S$. Since h_∞ is in W, it is one-to-one, and since $S_\alpha = S$, this means h_∞ is onto. Thus, we can set $g = h_\infty$, and see that there is in fact a function from the ordinals to S that is one-to-one and onto. This defines a well-ordering of S. □

Exercise 5.1 Show (using the Axiom of Choice) that every infinite set has a subset that has a correspondence with the natural numbers.

Recall that any two ordinals can be compared — that is, for every two ordinals, one of them appears within the order type of the other. Thus, since every set is in correspondence with an ordinal, we can see that for any two sets, one of them is in correspondence with a subset of the other. Thus, any two sets can be compared.

Furthermore, for any set, this theorem implies that there is at least one ordinal that is in correspondence with this set. Because the ordinals are well-ordered, this means there must be a *least* ordinal that is in

correspondence with the set. Just as we identify the finite ordinals with the finite cardinals (both are the natural numbers), we can identify the infinite cardinals with these lowest infinite ordinals that have a given correspondence. This then gives us a well-ordering of all the ordinals.

Cantor used the Hebrew letter Aleph, "ℵ", to stand for these ordinals. (Most of the Greek and Roman letters were already used for other concepts in mathematics, including Cantor's use of ω, and since Cantor was Jewish, he was familiar with the Hebrew letters.) He let $\aleph_0$ be the smallest infinite cardinal number (which is equal to ω), and then let $\aleph_1$ be the next cardinal number, and then $\aleph_2$ the next one, and so on. Because each of these is well-defined, there must be a set $\{\aleph_n \mid n \in \mathbb{N}\}$, and the union of this set would have to be an ordinal larger than any of them, which must be $\aleph_\omega$. But this is not the largest infinite cardinality (since the power set of any set has a larger cardinality than that set), so the next one must be $\aleph_{\omega+1}$, and then $\aleph_{\omega+2}$, and so on. A similar argument shows that in fact, for every ordinal α, there is a cardinal $\aleph_\alpha$.

Returning to Cantor's continuum problem, we see that the Continuum Hypothesis says that the cardinality of $\mathbb{R}$ is $\aleph_1$. But if the Continuum Hypothesis is false, the cardinality of $\mathbb{R}$ could be $\aleph_2$ or $\aleph_3$ or $\aleph_{17}$ or $\aleph_{\omega+3}$. In fact, Paul Cohen proved in 1963 that modern axiomatic set theory is not enough to answer the continuum problem. In fact, in 1970s, Easton's Theorem showed that the size of the continuum could be nearly any $\aleph$. It can't be precisely $\aleph_\omega$, or $\aleph_{\omega+\omega}$, or anything else like that (in a sense that can be made precise), but it can be anything else.

5.1 Games with hats

Let's take a break from infinite sets for a bit.

Let's consider a game, where n people stand in a circle. I will either put a black hat or a white hat on every person's head. Everyone can see the hat on everyone else's head, but not the hat on their own head. Now the players will go around the circle, and take turns guessing the color of the hat on their own head.

If everyone is just guessing at random, you can see that on average, about half of the people are going to guess incorrectly. But it turns out that if they plan ahead, they can do better. In fact, there's a strategy they can take that will guarantee that at most one person will get their guess wrong!

The important thing to think about here is that not only can everyone see everyone else's hat, they can also *hear the guesses* of everyone who guesses before them. So early people can use their guesses to communicate information to later people.

The basic idea is that the first person should look out at all the hats they can see, and should guess "white" if they see an odd number of white hats, and "black" if they see an even number of white hats. (The reverse would also work, but the players need to have agreed in advance which of these strategies they will use.) Once the first player has made their "guess", everyone else can look around at what hats they can see. They ignore the first person's hat, since they know that person can't see their own hat, but counting the rest, they can either see an odd number of white hats or an even number of white hats. If the first person saw an odd number of white hats, and this person sees an odd number of white hats on heads other than the first person's, then they know their own hat must be black. But if the first person saw an odd number of white hats, and this person sees an even number of white hats on heads other than the first person's, then they know their own hat must be white. (Similar reasoning works if the first person saw an even number of white hats.)

Thus, if they agree on one of these strategies, then as soon as the first person makes their "guess", everyone else can look around and know the color of their own hat. Thus, everyone but the first player can guess their own hat color correctly — and even the first player will sometimes get their own color right just by chance.

Exercise 5.2 Show that for any n, these (i.e., first player guesses white if they see an odd number of white hats; first player guesses black if they see an odd number of white hats) are the *only* two strategies that can guarantee that everyone but the first player guesses correctly.

However, if they all have to guess at once (using just the information about hats other than their own), we can prove that no matter what strategy they use, there must be some possibilities where at least half of

them get their guess incorrect. If there are any possibilities where a majority of them get their guess correct, then there must also be possibilities where a majority of them get their guess incorrect.

To see this, we will count the possible number of correct or incorrect guesses in two different ways. There are 2^n ways of assigning hat colors to the players, and each of the n players makes a guess on each such assignment, so there are $n \cdot 2^n$ total possible guesses. For a given player, we can see that regardless of the strategy, they must get their own hat right in exactly half of the different possible assignments. That's because, if we keep everyone else's hat fixed, we must keep a player's guess fixed, and there is exactly one possibility where that results in a correct guess for the player, and exactly one possibility where that results in an incorrect guess for the player. Thus, the total number of correct guesses across all 2^n possibilities must be precisely $\frac{n \cdot 2^n}{2}$.

If it were true that, in every possibility, a majority of players got their guess correct, then there would have to be more than $n/2$ correct guesses in each of the 2^n possibilities. Thus, there must be some possibilities where at least half of players get their guess incorrect. Furthermore, if there are *any* possibilities where more than $n/2$ guesses are correct, there must be some other possibilities where *less* than $n/2$ guesses are correct, in order that the total can be precisely $\frac{n \cdot 2^n}{2}$. Depending on their strategy, the players can affect how they cluster these correct and incorrect guesses — they can ensure that *usually* a majority will get it right, but then sometimes they will all get it wrong, and they will never all be right.

But look what happens if there are infinitely many players. Assume that there are $\aleph_0$ many players, i.e., one player for each natural number, and that each can see the hat on everyone's head but their own, and they all have to guess at once. Using the Axiom of Choice, we can prove that there is a strategy that guarantees that *always*, not only a *majority* of players get their guess right, but in fact *only finitely many* get it wrong!

The way they do this is the following. Consider the set of all possible assignments of hats to the players. Say that two of these assignments are "similar" if there are only finitely many differences between the assignments. For any assignment x, let S_x be the set of all assignments similar to x. We can see that if x is similar to y, then $S_x = S_y$ (anything similar to x is similar to y), while if x is not similar to y, then $S_x \cap S_y = \emptyset$ (nothing is both similar to x and similar to y). Call these sets the "similarity classes".

Now consider the set X of all similarity classes of assignments. Each of these similarity classes is non-empty, so the Axiom of Choice tells us there is a function f such that $f(S) \in S$ for any $S \in X$.

Here is the strategy that each player will use. When a player looks out at all the other hats, they can't tell precisely which assignment x is actual. However, they can tell which similarity class S_x the actual assignment is in. Then $f(S_x)$ is some assignment within this similarity class. Each player can then guess the color of their own hat in the assignment $f(S_x)$. Although it is unlikely that every player guesses their own hat correctly, we can see that the true assignment is in the same similarity class as the guess. Which means that the actual hat colors must only differ from the guessed hat colors for finitely many players. That is, all but finitely many of them get it right.

5.2 More troubling results

The Banach-Tarski Paradox

"Russell's socks"

Part II: Formal axiomatic set theory

6 The formal language, and the "small" axioms: Extensionality, Pairing, Union, and Empty Set

In order to avoid the paradoxes, and ensure that we know precisely what is provable, we will make every assumption about set existence explicit.

Along the way, we will need to talk about a formal language that shows what we can express precisely in the language of set theory. This will be made more explicit as we get to that, but for those of you familiar with first-order logic, this formal language will include the relations "$\in$" and "$=$", and everything definable from these using the logical connectives "and" ($\wedge$), "or" ($\vee$), "not" ($\neg$), and the quantifiers "there exist" ($\exists$) and "for all" ($\forall$). We can also use the logical connective "if φ then ψ", defined as "$\neg(\varphi \wedge \neg\psi)$" or "$(\neg\varphi \vee \psi)$".

We have the intuitive idea that "$\in$" means "is a member of", but this is only the guiding intuition. We will try to make the meaning more precise by stating *axioms*, which are the basic facts that tell us everything we can say formally about sets and membership. But because the language with just "*in*" and "$=$" is somewhat restricted, we will also supplement this language with many *definitions* — each definition introduces a new symbol, and tells us everything there is to know about when that symbol applies. Anything we express using the new symbol could be expressed without it

Definition "$x \subseteq y$" is an abbreviation for "every element of x is an element of y" (or, in the formal language, $\forall z(z \in x \rightarrow z \in y)$).

Exercise 6.1 If $x \subseteq y$ and $y \subseteq z$ then $x \subseteq z$.

Axiom (Pure Set Extensionality) $x = y$ if x and y have the same elements.

Within the formal language, we could state this as $x = y \leftrightarrow \forall z(z \in x \leftrightarrow z \in y)$, and if we make use of the definition of "$\subseteq$", we could even state it as $x = y \leftrightarrow (x \subseteq y \& y \subseteq x)$. This axiom is just another way of stating something that we already stated back in the beginning, that two sets are actually the same set if they have the same elements, but now we are making explicit that this is one of the axioms of Zermelo-Fraenkel set theory, that tells us some of the basic facts about what sets are.

Notice that, among other things, it implies that there can be at most one object with no elements. (Any two objects that both have no elements would have the same elements, and would thus have to be the same object.) This means that, in this context, everything we are talking about is a set — the members of sets are themselves sets.

Axiom (Ur-Element Extensionality) If x and y are sets, then $x = y$ if they have the same elements.

Axiom (Set of Ur-Elements) There is a set that contains everything that is not a set.

There is an axiom that states that there is in fact an object that has no elements. This axiom turns out to be implied by axioms we will discuss later, so it is sometimes left out of some versions of the Zermelo-Fraenkel axioms, but it is useful to include it.

Axiom (Empty Set) There is a set that has no elements.

Definition "$\emptyset$" is a name for the set that has no elements.

Axiom (Pairing) For any x and y, there is a set that contains x, and contains y, such that any element of that set is either x or y.

Definition "$\{x, y\}$" is a name for the set that contains x, and contains y, such that any element of that set is either x or y.

Notice that this axiom (and the associated definition) doesn't require that x and y be distinct objects. Thus, among the implications of this axiom is the claim that for every object x, there is a set whose *only* element is x. Such a set is called a *singleton*.

Definition "$\{x\}$" is a name for the set that contains x, and nothing else.

It is straightforward to use these two axioms (and the relevant definitions) to see that $\{x, y\} = \{y, x\}$.

If x and y exist, then the above axioms and definitions let us conclude that $\{x\}$ and $\{x, y\}$ exist, and so the Pairing axiom lets us conclude that $\{\{x\}, \{x, y\}\}$ exists as well.

Definition (Kuratowski's definition) "$\langle x, y\rangle$" is an abbreviation for "$\{\{x\}, \{x, y\}\}$".

This is how we make precise the concept of an ordered pair, which we just took for granted before. One thing to note, which is slightly strange, is that if $x = y$, then $\langle x, y\rangle = \{\{x\}\}$.

Exercise 6.2 $\langle x, y\rangle = \langle z, w\rangle$ iff $x = z$ and $y = w$.

Axiom (Union) For any X, there is a set that contains all and only the elements of any sets that are elements of X.

Definition "$\bigcup X$" is a name for the set that contains all and only the elements of any sets that are elements of X.

Definition "$x \cup y$" is an abbreviation for "$\bigcup\{x, y\}$".

The relation between the symbols "$\bigcup$" and "$\cup$" is a lot like the relation between the mathematical symbols "$\sum$" and "$+$".

Notice that if $x \in X$ then $x \subseteq \bigcup X$.

Using the Pairing axiom, we can show that for any x, y, z, there are sets $\{x, y\}$ and $\{z\}$, and thus $\{\{x, y\}, \{z\}\}$. Then by the Union axiom, there is a set $\bigcup\{\{x, y\}, \{z\}\}$, and we can see that the elements of this set are all and only x, y, z. We can similarly do this for any other finite number of objects.

Definition For any finite list of objects $x_1, x_2, \ldots, x_n$, "$\{x_1, x_2, \ldots, x_n\}$" is a name for the set that contains these objects as members, and nothing else.

Thus, although Pairing only directly puts two elements together into a set, we can use Pairing and Union together to create sets with any specific finite set of elements.

Definition "$S(x)$" is an abbreviation for "$x \cup \{x\}$".

This lets us generate a particular sequence of sets that we will use to stand for the natural numbers. In particular, we will use the numeral "0" to stand for $\emptyset$, and for any natural number n, we will use the numeral $n + 1$ to stand for $S(n)$.

In order for this definition to make sense, we need to make sure that every set that arises in this sequence is different from all the earlier sets.

First, observe that if x has no element that is an element of itself, then x is not an element of itself. (After all, if x were an element of itself, then it would have an element that is an element of itself, namely x!)

Now note that if x has no element that is an element of itself, then $S(x) \neq x$, because x is an element of $S(x)$, but is not an element of x. Furthermore, $S(x)$ is not equal to any element of x, because it contains every element of x, and we assumed that no element of x contained itself. Finally, $S(x)$ also has no element that is an element of itself, because its elements are just the elements of x, and x itself, none of which is a member of itself.

Now, it's clear that 0 has no element that is an element of itself (because it has no elements at all!) so the above results apply. Thus $S(0) = 1$ is different from 0, and has no element that is an element of itself. Thus $S(1) = 2$ is different from 1, and from every element of 1, and has no element that is an element of itself. And similarly for $3, 4, 5, \ldots$. Each new natural number is different from each previous natural number, and has no element that is an element of itself.

Exercise 6.3 Using these definitions, show that each n is the set of all natural numbers strictly less than n. (This will be a proof by induction. The base case is to show that this is true when $n = 0$, where there is no natural number less than n, and the induction step involves showing that if it's true for n, then it's also true for $n + 1$.)

Note that this shows that each natural number n has exactly n members. These results also mean that there are infinitely many distinct sets — though we can't yet prove that there is a set that itself has infinitely many distinct elements.

7 The "medium" axioms: Power set, (Foundation), and Separation

7.1 Power set

Axiom (Power Set) For any set x, there is a set whose elements are all and only the subsets of x.

Definition "$\mathcal{P}(x)$" is a name for the set whose elements are all and only the subsets of x.

Exercise 7.1 Find $\mathcal{P}(\emptyset)$, $\mathcal{P}(\mathcal{P}(\emptyset))$, and $\mathcal{P}(\mathcal{P}(\mathcal{P}(\emptyset)))$ and write them explicitly using curly braces.

Note that for any x, both $\emptyset \in \mathcal{P}(x)$ and $x \in \mathcal{P}(x)$, but you should have already shown in the above exercise that not every power set has at least two elements. (Can you identify all the sets whose power set has fewer than two elements?)

Exercise 7.2 If $x \subseteq y$ then $\mathcal{P}(x) \subseteq \mathcal{P}(y)$.

Exercise 7.3 $x = \bigcup \mathcal{P}(x)$ and $x \subseteq \mathcal{P}(\bigcup x)$, but there exists a set x such that $x \neq \mathcal{P}(\bigcup x)$.

We will eventually be able to prove Cantor's result that the power set of a set is always bigger than the set, but in the axiomatic context, it will take some time to set up the precise formal statement of what that means.

7.2 Foundation

The next axiom is usually included for convenience, but isn't essential to the main mathematical results. Other than Extensionality, it is the only axiom that doesn't take the form of constructing a new set out of an old set.

Axiom (Foundation) For any non-empty set x, there is a member of x that has no elements in common with x.

This axiom is somewhat strange, but the most important implications of this axiom can be expressed more directly.

First, this axiom ensures that there is no set S with $S \in S$. If there were, we could then consider the set $x = \{S\}$. This set has only one element, namely S. But S has an element in common with x, namely S itself!

Second, this axiom ensures that there is no finite sequence $S_0 \in S_1 \in \cdots \in S_n \in S_0$. Again, if there were, then our earlier results with Pairing and Union would let us construct a set $x = \{S_0, S_1, \ldots, S_n\}$. But every element of this set has at least one element in common with this set, namely the earlier element.

Finally, this axiom also entails that there is no set of all sets — if there were such a set, it would contain itself. But we will use other means to show that there is no set of all sets.

7.3 Separation

Axiom (Separation) For any set x, and any formula $\phi(w)$ in the formal language, there is a set consisting of all and only the elements of x that satisfy the formula $\phi(w)$.

Definition "$\{w \in x \mid \phi(w)\}$" is a name for this set.

This axiom justifies *some* uses of the "set-builder" notation we defined early on.

Notice that since there is a set of all ur-elements, we can use this axiom to define a set of all ur-elements but one, or all ur-elements but two.

However, the axiom of Separation doesn't let us define a set of *all* objects satisfying a property, just the objects satisfying this property that come from a set that already exists. This restriction is one of the

main ways that the Zermelo-Fraenkel axioms avoid most of the paradoxes. Several of the paradoxes involved collecting together *all* objects of a given sort, like the Burali-Forti Paradox and Russell's Paradox.

The fact that $\phi(w)$ has to be expressible in the formal language solves the other paradoxes. The Sorites paradox involved vague language that isn't part of the language of set theory, and many of the paradoxes involved the concept of "definability" or "meaning", which are not themselves definable within the formal language, like Grelling's paradox (the word "heterologous"), and Richard's paradox (the definable diagonal of the definable sets of natural numbers).

Separation lets us construct various other sets. For instance, we can construct $x \setminus y$, the set of all elements of x that are not in y, and we can construct $x \cap y$, the intersection of x and y. In fact, we can construct $\bigcap X$, the intersection of all the sets in X, and even $\bigcap \varphi(w)$, the intersection of all sets that satisfy $\varphi(w)$! That's because being in every set satisfying $\varphi(w)$ is a property that can be expressed in the formal language, and if we start with any one set satisfying this property, it will contain all the elements that are needed.

Exercise 7.4 Use the Separation axiom to prove that the empty set exists, using only the assumption that at least one set exists.

Exercise 7.5 Use the Separation axiom to prove that there is no set of all sets. (Hint: Use Russell's paradox, letting $\varphi(w)$ be "$\neg(w \in w)$".)

7.4 Developing Cantor's theory of sizes of infinity

Once we have this we can develop all the machinery for Cantor's theory of infinity, and be pretty confident we are safe from the paradoxes.

We need to say what it is for sets to have the same size - this means we need to be able to talk about the existence of correspondences between the sets, so we have to think of correspondences as sets of a particular sort.

Correspondences are relations that satisfy four properties, so we need to be able to think of relations as sets of a given sort, describe these properties in the formal language, and show that the relevant ones exist.

Relations are going to be thought of as sets of ordered pairs. We have already shown that every particular ordered pair exists, so what we will do in the rest of this section is show that the relevant sets of ordered pairs exist, and then show that we can talk about these properties, and eventually prove Cantor's theorem that the power set of a set is larger than the set. We don't yet have the existence of an infinite set, but once we add that, we will have the whole theory.

We can also use Separation to construct lots of other interesting sets. The set of all singletons from X can be constructed from the power set of X, using the formula $\varphi(w)$ given by "$\exists x \forall y(y \in w \leftrightarrow y = x)$", which says there is an object x such that x is an element of w and all elements of w are x. The set of all ordered pairs of elements of X can be constructed from the power set of the power set of X, using the formula $\varphi(w)$ stating that w is an ordered pair of two elements of X (it's worth checking that this can be expressed in the formal language). The set of all ordered pairs of an element of X with an element of Y, which we write as $X \times Y$, the *cartesian product* of X and Y, can be constructed from the power set of the power set of the union of X and Y.

Any subset of this cartesian product is called a binary relation from X to Y. For any $R \subseteq X \times Y$, we sometimes write xRy for $\langle x, y\rangle \in R$.

Again using Separation, for any X, we can construct the identity relation on X, $\{\langle x, x\rangle \mid x \in X\}$.

For any set of ordered pairs, we can use Union follows by Separation to construct the domain and range of this set, so that every set of ordered pairs can be thought of as a relation.

Putting the cartesian product in reverse, we can then use Separation to create the inverse of a relation.

If R is a relation from X to Y and S is a relation from Y to Z, then we can generate the composition $R \circ S$ as a relation from X to Z.

Say that R is a *function* if for every $x \in X$, if $\langle x, y_1\rangle \in R$ and $\langle x, y_2\rangle \in R$, then $y_1 = y_2$.

Say that R is *total* if for every $x \in X$, there exists a y with $\langle x, y\rangle \in R$.

Say that R is *one-to-one* if for every $y \in Y$, if $\langle x_1, y\rangle \in R$ and $\langle x_2, y\rangle \in R$, then $x_1 = x_2$.

Say that R is *onto* if for every $y \in Y$, there exists an x with $\langle x, y\rangle \in R$.
If $R \subseteq X \times Y$ and $S \subseteq Y \times Z$, then we define $R \circ S \subseteq X \times Z$ by $R \circ S = \{\langle x, z\rangle \mid \exists y(\langle x, y\rangle \in R \& \langle y, z\rangle \in S)\}$.

Exercise 7.6 If R is a relation from X to Y, and S is a relation from Y to Z, and R and S are both (functions/total/one-to-one/onto) then so is $R \circ S$ as a relation from X to Z.

By taking $\mathcal{P}(X \times Y)$ and using Separation, we can construct the set ${}^{X}Y$ of all total functions from X to Y.

Exercise 7.7 (Cantor's Theorem) Show that for any set X, there is no relation with domain X and range $\mathcal{P}(X)$ that is both a function and onto.

In the proof of Cantor's Theorem, you will need to be explicit about how, for any relation with domain X and range $\mathcal{P}(X)$ that is a function, you can use Separation to define a diagonal set, which is not part of the relation.

8 Infinity, formalized ordinals, and the well-ordering principle

Recall that we defined the successor function by $S(x) = x \cup \{x\}$. We say that a set z is *inductive* if $\emptyset \in z$ and $\forall x(x \in z \to S(x) \in z)$. That is, a set is inductive if it contains the empty set and contains the successor of anything it contains.

Axiom (Infinity) There is an inductive set.

Recall that the Separation axiom allows us to take the intersection of any definable collection of sets that has at least one member. Thus, the Axiom of Infinity allows the following definitions:

Definition ω is the intersection of all inductive sets.

Definition For x to be a *natural number* is for it to be a member of ω.

Notice that these definitions ensure that proof by induction works to establish that all natural numbers have a given property. If $\varphi(w)$ is some property, Separation allows us to consider the set of natural numbers that satisfy that property, $\{w \in \omega \mid \varphi(w)\}$. Call this set X. It is clear that $X \subseteq \omega$. But if we can show that $\varphi(\emptyset)$ and $\forall w(\varphi(w) \to \varphi(S(w)))$, then we can see that X must in fact be an inductive set. But because ω is defined to be the intersection of all inductive sets, this means that $\omega \subseteq X$. And thus, by Extensionality, since ω and X are both subsets of each other, they must be the same set. So we can conclude that every natural number has property $\varphi(w)$.

When we first discussed proof by induction, we just noted it as a method that intuitively should work. But now we have defined the natural numbers in a way that ensures that this method works, by definition.

Exercise 8.1 No natural number has a correspondence to a proper subset of itself. (Hint: The best way to prove this in the formal axiomatic context is to show that the set of natural numbers that don't have a correspondence to a proper subset is inductive.)

Exercise 8.2 ω does have a correspondence to a proper subset of itself.

Definition For x to be *finite* is for it to have a correspondence to some natural number.

Because no natural number has a correspondence with a subset of itself, it's clear that there is no more than one natural number with a correspondence to any given set.

8.1 Addition of natural numbers

We can define addition, multiplication, and exponentiation on the natural numbers in two different ways. Both ways are useful, and we can prove that they yield the same definition. Both methods show that these relations are definable in the formal language of set theory, but they have different advantages.

One method is *combinatorial* — we show some operation on sets that corresponds to the intuitive meaning of the operation, and identify the natural number that has a correspondence to the resulting set. The other method is *recursive* — we define how the operation works for the number 0, and then define how the operation works on $S(n)$ in terms of how it works on $S(n)$.

For addition, the combinatorial definition involves the union of disjoint sets: when you visualize what "13+15" means, you probably visualize something like a set of 13 students, and a separate set of 15 students being combined. It's important that these sets are disjoint — if there's one class with 13 students, and another class with 15 students, there may be less than 28 students when you combine them, if there are some students who are in both classes. To define the sum of m and n, we need to define two disjoint sets of these sizes. In order to do this, we will take advantage of the fact that ordered pairs with different second elements are always different. We will combinatorially define $m + n$ as the natural number that has a correspondence with $m \times \{0\} \cup n \times \{1\}$.

For addition, the recursive definition is as follows: we define $m + 0 = m$, and we define $m + S(n) = S(m + n)$. As I will illustrate more later, we can think of this sort of definition as saying that to add n to

m, you start with m (because $m + 0 = m$) and then apply the successor operation n times (because each time we add one more S to n, we add one more S to the output).

The first thing to note about the recursive definition is that it has two parts — there is one part that tells us how to add 0 to any number, and another part that tells us how to add the successor of n to any number, if we already know how to add n to that number. The second part of that definition might *feel* circular, but it isn't — it shows how addition of a bigger number can be defined in terms of addition of a smaller number, and if we follow these definitions back, we can eventually get down to the explicit definition with 0. It is straightforward to prove by induction that this recursive definition gives a unique answer for any $m + n$ — if $n = 0$, then obviously only the first clause applies so the definition is unique; and if $m + n$ is defined uniquely, then for $m + S(n)$, only the second clause applies, and the answer will be defined, and unique.[1]

Once we have given both the combinatorial and the recursive definition of addition, we can prove that they do in fact define the same function. To do this, we work by induction. We first prove the base case, that the combinatorial definition of $m + 0$ yields the value m. (Think about what the combinatorial definition of $m + 0$ is.) And then we prove the induction step, that the number that is the combinatorial definition of $m + S(n)$ is the successor of the number that is the combinatorial definition of $m + n$.

One feature of the recursive definition is that we can use it to algorithmically reduce any expression involving 0, the successor operation, and addition, to one just involving 0 and the successor operation. For instance:

$$\begin{aligned} S(S(0)) + S(S(S(0))) &= S\Big(S(S(0)) + S(S(0))\Big) \\ &= S(S\Big(S(S(0)) + S(0)\Big)) \\ &= S(S(S\Big(S(S(0)) + 0\Big))) \\ &= S(S(S(S(S(0))))) \end{aligned}$$

The first three steps in that calculation use the fact that the term on the right side of the addition operation begins with a S (which can then be moved outside the addition), and the last step uses the fact that an addition involving 0 is trivial. The same can be done when there are multiple addition operations in the expression.

Exercise 8.3 Use a similar sequence of operations to show that $S(S(0)) + (S(0) + S(0)) = S(S(S(S(0))))$ (that is, that $2 + (1 + 1) = 4$).

We can naturally think of the combinatorial definition as giving the "meaning" of addition, its semantic content, while the recursive definition gives a syntactic process for calculating the outcome. It's useful to compare this to the way that we all learned to add numbers using Arabic numerals — think about how you figure out that $176 + 247 = 423$ ($6 + 7 = 13$, write down the 3, carry the 1, $1 + 7 + 4 = 12$, write down the 2, carry the 1, $1 + 1 + 2 = 4$, write that down, see 423). You don't have to know what these numbers mean, or that addition corresponds to taking two disjoint sets of objects and finding their union, and you can even arrange for a small set of electrical circuits, or even mechanical switches, to do this automatically without

[1]This footnote outlines the proof of the general theorem for recursive definition of functions of natural numbers. A similar proof allows us to prove this for arbitrary functions of ordinals.

We can in fact show a general theorem about recursive definitions of this sort. Consider the set of all ordered triples of natural numbers, $(\omega \times \omega) \times \omega$. Take the set of all subsets of this set. Use Separation to remove all the sets that contain a triple whose second member is 0 and whose third member is non-zero. Use Separation again to remove all the sets that contain a triple of the form $(m, S(n), \ell)$ that don't also contain a triple of the form (m, n, k), where $\ell = S(k)$. We can prove that the remaining sets of triples will all define partial functions that satisfy the recursive definition of addition. We can prove, by induction, that for any two such partial functions, if one of them contains a triple of the form (m, n, k) and the other contains a triple of the form (m, n, ℓ), then $\ell = k$. The union of all these sets will then define a function that satisfies these conditions as well. Finally, we prove by induction that for every m, n, this union has a triple of the form (m, n, k).

This general proof strategy works for a recursive definition of *any* function $f(m, n)$, if there is some one-place function g that we define $f(m, 0) = g(m)$, and some three-place function h that we define $f(m, S(n)) = h(m, n, f(m, n))$. We will use this for multiplication and exponentiation as well.

any understanding, as in the earliest mechanical calculators built in the 1700s, and widely available as cash registers in shops by the 1890s. It's a lot more complicated to write out the rules for doing addition with Arabic numerals than to write out the two simple rules that $m + 0 = m$ and $m + S(n) = S(m + n)$ — but it's much easier for humans to tell the difference between "6" and "7" than between "$S(S(S(S(S(S(0))))))$" and "$S(S(S(S(S(S(S(0)))))))$".

The combinatorial definition of addition allows us to straightforwardly prove some general properties that addition has. For each of the following exercises, you should recall that "$m + n$" is defined as the natural number that has a correspondence with $m \times \{0\} \cup n \times \{1\}$, and show that there is a correspondence between the relevant sets in these exercises, which has the consequence that whatever natural number has a correspondence with one must have a correspondence with the other.

Exercise 8.4 (Commutativity of addition) For all natural numbers m, n, $m + n = n + m$.

Exercise 8.5 (Associativity of addition) For all natural numbers ℓ, m, n, $\ell + (m + n) = (\ell + m) + n$.

We can also prove these facts using the recursive definition, but it turns out to be much harder, and needs to be done by a double (or possibly triple) induction. Don't worry about following this proof (you can skip straight ahead to the section on multiplication instead), but just notice how much longer and less pleasant this is. It's much nicer to prove that the recursive definition agrees with the combinatorial definition, and then prove that the combinatorial definition satisfies commutativity.

Theorem 8.1 (Commutativity of addition, based on the recursive definition) For all natural numbers m,, $m + n = n + m$.

Proof. This will be a proof by induction on n. We need to show a base case (that 0 commutes with all numbers, i.e., for all m, $m + 0 = 0 + m$) and an inductive step (that if n commutes with all numbers, then $S(n)$ does as well).

Base case: We need to show that 0 commutes with all natural numbers, i.e., for all natural numbers m, $m + 0 = 0 + m$.

> *Proof.* The recursive definition of addition tells us that $m+0 = m$. So to prove this base case, we need to show that $0+m = m$. You might have hoped that this directly follows from the recursive definition, but unfortunately the recursive definition only tells us what to do with numbers on the *right* side of the "+" sign, not the left. So to prove that $0 + m = m$, we need to work by induction, this time on m.
>
> > Base case: We need to show that $0 + m = m$ when $m = 0$, i.e., we need to show that $0 + 0 = 0$. This does follow directly from the recursive definition of addition.
> >
> > Inductive step: Now we need to show that if $0 + m = m$, then $0 + S(m) = S(m)$. By the recursive definition of addition, $0 + S(m) = S(0 + m)$. So if $0 + m = m$, then we can replace this inside the parentheses, so that $0 + S(m) = S(m)$.
>
> Thus, we have shown that for all natural numbers m, $0 + m = m$. Since $m + 0 = m$, we now know that for all m, $m + 0 = 0 + m$, so 0 commutes with all natural numbers. □

Inductive step: We need to show that if n commutes with all numbers, then $S(n)$ does as well.

> *Proof.* We want to show that $S(n)$ commutes with all numbers, i.e., that for all m, $m + S(n) = S(n) + m$. We will do this by yet another induction, this time on m. That is, we need to show a base case (that if n commutes with all numbers, then 0 commutes with $S(n)$) and an inductive step (that if n commutes with all numbers, and m commutes with $S(n)$, then $S(m)$ commutes with $S(n)$).
>
> The base case just requires us to show that 0 commutes with $S(n)$. Fortunately, above we showed that 0 commutes with all numbers, so we don't need to prove this again.

For the inductive step, we need to show that if n commutes with all numbers, and m commutes with $S(n)$, then $S(m) + S(n) = S(n) + S(m)$.

By the recursive definition of addition, we know that $S(m) + S(n) = S\Big(S(m) + n\Big)$. By the assumption that n commutes with all numbers, this equals $S\Big(n + S(m)\Big)$, and by the recursive definition of addition, this equals $S(S\Big(n+m\Big))$. Because n commutes with all numbers, $S(S(n+m)) = S(S(m+n))$, and by the recursive definition of addition, this equals $S(m+S(n))$. By the assumption that m commutes with $S(n)$, this equals $S(S(n)+m)$, and by the recursive definition of addition, this equals $S(n) + S(m)$, and we are done with this inductive step.

Thus, we have shown that if n commutes with all numbers, then so does $S(n)$. □

Thus, we have shown that 0 commutes with all numbers, and that if n commutes with all numbers, then $S(n)$ commutes with all numbers. So this completes the big induction, allowing us to conclude that all numbers commute. □

8.2 Multiplication of natural numbers

For multiplication, the combinatorial definition involves the cartesian product: the set of ordered pairs. When you visualize what "$4 \cdot 6$" means, you probably visualize something like the number of squares in a rectangular grid with 4 rows and 6 columns. Each row can be numbered with a natural number below 4, and each column can be numbered with a natural number below 6, and each square can be labeled with the corresponding ordered pair.

$\langle 0,0\rangle$	$\langle 0,1\rangle$	$\langle 0,2\rangle$	$\langle 0,3\rangle$	$\langle 0,4\rangle$	$\langle 0,5\rangle$
$\langle 1,0\rangle$	$\langle 1,1\rangle$	$\langle 1,2\rangle$	$\langle 1,3\rangle$	$\langle 1,4\rangle$	$\langle 1,5\rangle$
$\langle 2,0\rangle$	$\langle 2,1\rangle$	$\langle 2,2\rangle$	$\langle 2,3\rangle$	$\langle 2,4\rangle$	$\langle 2,5\rangle$
$\langle 3,0\rangle$	$\langle 3,1\rangle$	$\langle 3,2\rangle$	$\langle 3,3\rangle$	$\langle 3,4\rangle$	$\langle 3,5\rangle$

Since m is the set of natural numbers below m, and n is the set of natural numbers below n, we combinatorially define $m \cdot n$ as the natural number that has a correspondence with $m \times n$, the set of ordered pairs from m and n.

One convenient fact to notice — the set used to define $m + m$ is the same as the set used to define $m \cdot 2$. Both of them include every ordered pair whose first element is a member of m, and whose second element is either 0 or 1.

For multiplication, the recursive definition works as follow. We define $m \cdot 0 = 0$, and we define $m \cdot S(n) = m \cdot n + m$. The same argument as above shows that this definition is sufficient to give a unique answer. We can also see that this definition basically tells us that multiplying m by n means starting with 0, and then, n times, applying the operation of adding m.

Exercise 8.6 Prove (by induction on n) that the combinatorial definition of multiplication and the recursive definition agree for all m, n. (The base case will be very straightforward; for the induction step, there is a combinatorial definition of $m \cdot S(n)$, and there is a combinatorial definition of $m \cdot n + m$, and you will want to demonstrate a correspondence between the relevant sets in this combinatorial definition.)

Again, we can use this recursive definition to first eliminate each S on the right side of a multiplication and replace it by additions, and then replace a multiplication by 0 with a 0, and then pull out each S on

the right side of an addition, and cancel 0's on the right side of addition:

$$\begin{aligned}
S(S(0))) \cdot S(S(0)) &= \Big(S(S(0))) \cdot S(0)\Big) + S(S(0))) \\
&= \Big(\big(S(S(0))) \cdot 0\big) + S(S(0)))\Big) + S(S(0))) \\
&= \Big(0 + S(S(0))\Big) + S(S(0)) \\
&= S\Big(0 + S(0)\Big) + S(S(0)) \\
&= S(S\Big(0 + 0\Big) + S(S(0)) \\
&= S(S(0)) + S(S(0)) \\
&= S\Big(S(S(0) + S(0)\Big) \\
&= S(S\Big(S(S(0)) + 0\Big)) \\
&= S(S(S(S(0))))
\end{aligned}$$

Again, we can use the combinatorial definitions to easily prove some features of multiplication that are very difficult to prove using the recursive definitions. For both of these exercises, you should find a correspondence between the relevant sets.

Exercise 8.7 Show that for all natural numbers m, n, $m \cdot n = n \cdot m$.

Exercise 8.8 Show that for all natural numbers ℓ, m, n, $\ell \cdot (m \cdot n) = (\ell \cdot m) \cdot n$.

8.3 Exponentiation of natural numbers

Exponentiation is an operation that is usually not taught in elementary school, and when it is taught, it tends to be taught only in terms of the recursive definition. However, there is a combinatorial definition as well.

The recursive definition of exponentiation says that $m^0 = 1$ and $m^{S(n)} = m^n \cdot m$. That is, we think of exponentiation as starting with 1, and then multiplying it by m n times. This is probably the way you were taught to understand exponentiation, if at all.

The combinatorial definition again identifies m^n as the unique natural number that has a correspondence to a particular set. This set is the set of all total functions from n to m. Recall, a total function f from n to m is a set of ordered pairs, such that for every $k \in n$, there is exactly one member of f whose first element is k, and such that the second element of any member of f is a member of m. Members of m can appear zero, one, or many times as second elements of ordered pairs in f, but members of n must appear exactly once as first elements.

Exercise 8.9 Show, using the combinatorial definitions, that $m^2 = m \cdot m$. (This involves finding a correspondence between things like $f = \{\langle 0, k\rangle, \langle 1, \ell\rangle\}$ and ordered pairs of elements of m.)

Exercise 8.10 Show, using the combinatorial definition, that the power set of m has 2^m members. (This involves finding a correspondence between functions from m to $\{0, 1\}$ and subsets of m.)

Exercise 8.11 Show, by induction, that the recursive and combinatorial definitions of exponentiation agree.

8.4 Well-orderings

Recall that an ordering R on X is defined as a subset of $X \times X$ that is transitive and irreflexive. That is if $R \subseteq X \times X$ is a relation on X, then (trasitivity) if $\langle x, y\rangle \in R$ and $\langle y, z\rangle \in R$, then $\langle x, z\rangle \in R$, and (irreflexivity) that $\langle x, x\rangle \notin R$.

An ordering R on X is *total* if for every $x, y \in X$, either $x = y$ or $\langle x, y\rangle \in R$ or $\langle y, x\rangle \in R$.

Finally, R is a *well-ordering* if for every non-empty $S \subseteq X$, there is an R-least element of S. That is, there is some $x \in S$ such that for every $y \in S$, if $y \neq x$ then $\langle x, y\rangle \in R$. Once we prove that some relation is a well-ordering, we usually write "$m < n$" rather than "$\langle m, n\rangle \in R$", if it is clear which relation we are using.

Let $R \subseteq \omega \times \omega$ such that $R = \{\langle m, n\rangle \mid m \in n\}$. That is, R is the relation of being a member. We can prove that R is a well-ordering on ω. We have already proved that R is irreflexive (no natural number contains itself). Transitivity and totality are proved by induction.

To prove that it is a well-ordering, we show that if $X \subseteq \omega$ does *not* have an R-least element, then the complement of X contains all the natural numbers. We actually won't do this by proving directly that the complement of X is inductive — instead, we define a new set $X^* = \{n \mid \forall m \leq n(m \notin X)\}$, and we will prove that *this* set is inductive. If X does not have an R-least element, then it is clear that $0 \notin X$ (because it would be R-least in any set that contains it), so $0 \in X^*$ (since every natural number less than or equal to it is not in X). Furthermore, if $n \in X^*$, then neither n nor any smaller natural number is in X. So if $S(n)$ were in X, it would be R-least. So $S(n)$ must not be in X, and thus it must be in X^*. Thus X^* is inductive, and thus contains all natural numbers, so X itself must be empty.

Thus, every non-empty subset of the natural numbers has an R-least element, so this is a well-ordering.

If X and Y are both well-orderings, then we can define a well-ordering on $X \times Y$ by saying that $\langle x_1, y_1\rangle < \langle x_2, y_2\rangle$ if either $y_1 < y_2$ or $y_1 = y_2$ and $x_1 < x_2$. (Note: This is lexicographic ordering where the *second* element in the ordered pair is more significant!) To show that this is a well-ordering, we consider any subset of $X \times Y$ and show how to find the least element of that subset — just look at the set of y that form the second elements of pairs in this set, and find the least such y, and the look at the set of pairs that have this y, and find the least x that forms a pair with it, and the resulting pair must be the least in the set.

Thus, we now have many well-orderings beyond ω, such as $\omega \times \omega$ and $\omega \times (\omega \times \omega)$, each ordered lexicographically.

8.5 Replacement

Axiom (Replacement) If $\varphi(x, y)$ is a formula that defines a function (that is, for every x there is at most one object y satisfying $\varphi(x, y)$), and X is any set, then there is a set $\{y \mid \exists x{\in}X(\varphi(x, y))\}$.

Effectively, this axiom says that for any definable function, we can apply that function to every element of an existing set to get a new set. This axiom wasn't actually part of Zermelo's formulation, but was added by Fraenkel. This axiom expresses some of the idea that as long as a set isn't "too big" it won't lead to paradox — because we just replace the elements of an existing set one for one, the new set can't be "too big" if the old one wasn't.

Using Replacement, Pairing, and Union, we can construct cartesian products without actually using the Power Set axiom. Recall that for any two objects x, y, the Pairing axiom ensures that $\{x\}$, $\{x, y\}$, and $\{\{x\}, \{x, y\}\} = \langle x, y\rangle$ all exist. Note that if we fix a given object y, we can write down a formula stating that one object is the ordered pair consisting of the other object followed by y. This is a formula that defines a function, since there is only one such ordered pair for any first object. We can apply this function to each element $x \in X$ to get a set $X \otimes y = \{\langle x, y\rangle \mid x \in X\}$. Now the relation that fixes X and associates each y to the set $X \otimes y$ is itself another function, and we can use it to replace each element of Y, to get the set $\{X \otimes y \mid y \in Y\}$. Applying the Union axiom to this, we get $X \times Y$.

We can also use Power Set, Empty Set, and Replacement to prove Pairing, without having it as a separate axiom. Using the Empty Set axiom, we know there is a set $\emptyset$. Using Power Set, we know there is a set that contains every subset of this set — and since $\emptyset$ is its own only subset, this means there is a set $\{\emptyset\}$. Using Power Set again, we know there is a set that contains all the subsets of this set, which are just itself and the empty set, so there is a set $\{\emptyset, \{\emptyset\}\}$. Now for any objects z, w, we can write down the formula $\varphi(x, y)$ that states "$(x = \emptyset \& y = z)$ or $(x = \{\emptyset\} \& y = w)$". It is clear that this formula defines a function — for any value of x there is at most one value of y that makes this formula true. We can then apply this function to replace

the elements of $\{\emptyset, \{\emptyset\}\}$, getting $\{z, w\}$. Since z and w were arbitrary, this means that every instance of the Pairing axiom is true.

8.6 Axiomatic ordinals

Well-ordering and Replacement are what we need to define arbitrary functions recursively.

We can define Mostowski collapse. If x is a member of a well-ordering, we define $M(x) = \{M(y) \mid y < x\}$. We can show that this is enough to ensure that $M(x)$ is defined for all members of a well-ordering. If there are any members for which it is *not* defined, we can find the least element of that set. There is a set consisting of all the earlier members, and M is defined for each of those. Thus, we can use the Axiom of Replacement to replace all these earlier members by the output of M, and that gives us a set that is the output of M for the element that was supposed to be the least element for which it wasn't defined.

For addition and multiplication of ordinals, we use the same sets as we used earlier for the natural numbers, and apply Mostowski collapse the lexicographic ordering.

This shows that $2 \cdot \omega \neq \omega \cdot 2$. There is a correspondence, but it is not the same order.

9 The Continuum Hypothesis and further topics