NOTES ON SET THEORY
18.510, FALL 2015

HENRY COHN

Contents

We will now axiomatize set theory. From this point on, we cannot use any previous knowledge of mathematics, except as intuition or motivation. Everyday words such as "finite" or "number" no longer have any technical meaning for us, except as specified by our axioms and definitions. These axioms and definitions will become our entire world, and we must reconstruct mathematics within it.

1. First axioms of set theory

Our development of set theory will differ in one way from informal mathematics: everything will be a set. In other words, there are no atomic elements, contained in sets but not actually sets themselves. Instead, all mathematical objects are built entirely out of sets. The only elements of sets are sets themselves. What are the elements of those sets? Still more sets. It's sets all the way down.

From an ordinary mathematical perspective, this feels counterintuitive. We're not used to thinking of numbers like 2 or π as sets. If a student asks "Is $2 \in \pi$?", it's natural to criticize it as a meaningless question that betrays a fundamental misunderstanding of what real numbers are. However, in set theory it will be a bizarre but technically meaningful question. (The answer may depend on exactly how one constructs the real numbers, for example by Cauchy sequences or Dedekind cuts.) Exotic questions like this are of little interest and it is generally best to ignore them. We will put up with these sorts of issues so that we can build mathematics on the simplest foundation. Ultimately, we will see in these notes that all mathematical objects can be built recursively as sets, starting just with the empty set.

We will use two basic binary relations, namely $=$ and $\in$. In other words, given sets x and y, it's meaningful to ask whether $x = y$ and whether $x \in y$. The intuitive interpretations are that "x equals y" and "x is an element of y," but we must not rely on intuition. So far we have not stated any axioms and therefore know nothing about how $=$ and $\in$ behave. We will implicitly assume the usual properties of equality (everything is equal to itself, and if $x = y$ then substituting y for some or all of the occurrences of x in an assertion preserves truth), which we will deal with more formally in the second half of the course.

Everything else must be defined in terms of them, together with logical connectives (and, or, not, implies, is equivalent), quantifiers (there exists, for all), and variables that stand for sets. For example, we define $A \subseteq B$ to mean that every element of A is an element of B. In symbols, $\forall x\,((x \in A) \to (x \in B))$. Note that quantifiers quantify over all sets. Sometimes we write $\forall x \in S \dots$ as shorthand for $\forall x\,((x \in S) \to \dots)$ (here $\dots$ stands for some subformula), but there is no need to restrict the values we are quantifying over to a given set.

In the second half of the course we will develop the logical formalism more explicitly. For now, we'll treat it somewhat informally, but still carefully. In particular, we cannot make use of any intuitive ideas that haven't been precisely defined. For example, we all know what it means for a set to be finite, and that for example the union of two finite sets is always finite, but we cannot use these intuitions until we have rigorously justified them. One rule of thumb is that any time you need to write an ellipsis, something needs to be defined more carefully.

Because everything in set theory is a set, all variables will stand for sets. The sort of letter used is sometimes meant to be suggestive. For example, when one set is an element of another, it is sometimes convenient to use a lowercase letter for the element and an uppercase letter for the set it is in. However, there is no way to do this consistently, and these conventions play no role in the logic.

The theory we'll develop is called ZFC, which stands for Zermelo-Fraenkel set theory with the Axiom of Choice. (ZF is the same theory without the Axiom of Choice.) ZFC is the consensus foundation for modern mathematics, and almost everything mathematicians do can be done in this framework. We'll see later that ZFC is far from complete, and that additional axioms are sometimes useful, but ZFC is surprisingly comprehensive in practice.

The first question we need to address is when two sets are the same.

Axiom (Extensionality)**.** *Two sets are equal if and only if they have the same elements. In symbols,* $\forall S\,\forall T\,((S = T) \leftrightarrow \forall x\,((x \in S) \leftrightarrow (x \in T)))$.

We'll frequently use words instead of symbols, because symbolic expressions are often more cumbersome yet aren't actually more precise, once you are used to translating between words and symbols.[1] Later we'll formalize the symbolic language and give precise rules for what constitutes a proof, but it's worth developing some set theory as an example before getting into abstractions.

Intuitively, the Axiom of Extensionality means that elements of a set do not have associated multiplicities or degrees of membership, do not come in any sort of order, etc. The only thing that matters is whether each potential element is in or out of the set. Note also that extensionality justifies the common technique of proving that $S = T$ by showing that $S \subseteq T$ and $T \subseteq S$.

[1] If you reach a point at which you are no longer sure precisely what some words mean, then you have become confused and need to retrace your steps.

It is common in mathematics to equate a "property" with the set of all things having that property. For example, consider the following informal statement of mathematical induction:

> Given any property P of natural numbers, if 0 has property P, and if whenever n has property P, so does $n + 1$, then every natural number has property P.

This statement can be made precise as follows:

> Given any set S of natural numbers, if $0 \in S$, and if whenever $n \in S$, it follows that $n + 1 \in S$, then every natural number is in S.

However, one must be slightly careful about language when one equates properties with sets. In ordinary language, properties do not satisfy extensionality: it is possible to have two different properties satisfied by exactly the same things. For example, being the day after Monday is a different property than being the day before Wednesday, even though both properties uniquely define Tuesday.

This explains the name "extensionality." The *extent* or *extension* of a property is the set of things that satisfy it, while the *intent* or *intension* is how it characterizes this set, i.e., the idea it expresses. (This is correctly spelled, and it's subtly different from intention, but that is close enough for our purposes.) The Axiom of Extensionality says that when we define a set, only the extent matters, not the intent.

The next question is how we can make sets. There's an obvious guess, namely that we can make them however we want:

False Axiom (Comprehension)**.** *For every mathematical property $\varphi(x)$, there exists a set S such that for all x, we have $x \in S$ iff $\varphi(x)$ holds.*

The set S is unique, by the Axiom of Extensionality. We write

$$S = \{x : \varphi(x)\}.$$

Note that by a property $\varphi(x)$, we mean any syntactically correct statement about x in our formal language. We'll be more explicit about exactly what this means later, when we formalize the language. In the meantime, it's enough to understand that $\varphi(x)$ can be any property of x that we can state precisely.

Comprehension is labeled "false axiom" above because it is not part of ZFC. It cannot be, because it leads immediately to *Russell's paradox*: let

$$R = \{x : x \notin x\}.$$

(Here $x \notin x$ is an abbreviation for not $x \in x$, the negation of $x \in x$.) Then $x \in R$ iff $x \notin x$, and taking $x = R$ shows that $R \in R$ iff $R \notin R$, which is a contradiction.

One of the difficulties of set theory is that it is sometimes not obvious whether an attempted definition really does define a set. For example, the definition of Russell's set R looks innocent enough: sure, "$x \notin x$" looks a little odd, but it's a perfectly reasonable property for a set to have. Indeed, we don't expect any set to be an element of itself. It is therefore a little unnerving that the set R cannot exist.

However, there is no reason to be worried about this. Merely proposing a definition of something does not ensure its existence, any more than saying "Let Skeffington be a purple unicorn" ensures that Skeffington exists. From this perspective, Russell's paradox is no worse than saying "Let n be the greatest integer. Then $n + 1$ is an even greater integer, which is a contradiction." This paradox is resolved by saying there is no greatest integer, and Russell's paradox is resolved by saying $\{x : x \notin x\}$ is not actually a set.

Instead of comprehension, set theory uses the weaker Axiom of Separation. The Axiom of Separation says that once we have a set, we can define a subset by any property we want:

Axiom (Separation)**.** *For every mathematical property $\varphi(x)$ and every set S,*

$$T = \{x \in S : \varphi(x)\}$$

is a set. In other words, there exists a set T such that for all x, we have $x \in T$ if and only if $x \in S$ and $\varphi(x)$ holds.

As above, there is a unique such set T by extensionality. The name "separation" refers to separating the elements $x \in S$ that satisfy $\varphi(x)$ from those that do not. By contrast, "comprehension" refers to creating a comprehensive set of all x satisfying $\varphi(x)$. The former is always possible, but the later is sometimes impossible.

Definition 1.1. A *universe*, or *universal set*, is a set that contains every set as an element.

By extensionality, there can be only one universe, if any, but it is not clear whether there is one. In fact, there cannot be:

Proposition 1.2. *There is no universe.*

Proof. Suppose U were a universal set. By the Axiom of Separation, we can define

$$R = \{x \in U : x \notin x\}.$$

(In other words, such a set R exists.) Because U is universal, $R \in U$, and we conclude from taking $x = R$ that $R \in R$ if and only if $R \notin R$. This is a contradiction, so R cannot exist and hence U cannot exist. □

Note that the non-existence of the universe is simply Russell's paradox, with the paradoxical aspect eliminated. Every apparent paradox is really a theorem in disguise, if one can identify the problematic assumption being made.

If there were a universe U, comprehension would be a special case of separation, because $\{x : \varphi(x)\} = \{x \in U : \varphi(x)\}$. Thus, although Proposition 1.2 may seem surprising, it is fundamental to set theory.

The basic difficulty is that U is too big to be a set. This phenomenon, called limitation of size, is the fundamental reason why attempted definitions can fail to define actual sets.

It's possible to work in a bigger theory, which includes not just sets but also objects called *classes*. These are set-like objects, but possibly much bigger, and a set is simply a class that can be an element of another class. Then there is a universe, i.e., a class consisting of all sets, but it is a proper class, not a set itself. This language is sometimes convenient, but never necessary, and we will not use it in our formal development of set theory.

So far, there is a worrisome gap in our axioms: are there any sets at all? Extensionality tells when two sets are equal, but it doesn't guarantee that any sets exist in the first place, and separation only lets us make further sets given a starting set.

Axiom (Empty Set)**.** *There is a set with no elements.*

The empty set is unique, by extensionality, and we denote it $\emptyset$ or $\{\}$. Note that we could replace the Empty Set Axiom by an axiom simply asserting the existence of some unspecified set, since given any set S, we can define $\emptyset = \{x \in S : x \neq x\}$.

Now we know there is at least one set, but our axioms so far allow the possibility that the empty set is the only set.

Axiom (Pairing)**.** *For all sets a and b, there exists a set S such that $a \in S$ and $b \in S$.*

The set S whose existence is guaranteed by the axiom could contain further elements, but we can cut it down by separation to form the set $\{x \in S : x = a \text{ or } x = b\}$, which is unique by extensionality (since its only elements are a and b). We write $\{a, b\}$ for that set. Note that we allow $a = b$, in which case we can also write $\{a\}$.

From $\emptyset$, we can now form a new set $\{\emptyset\}$. It is a different set: $\emptyset$ has no elements, while $\{\emptyset\}$ has one element, namely $\emptyset$. We can go on the form further sets, such as $\{\emptyset, \{\emptyset\}\}$ or $\{\{\emptyset\}\}$. However, using the axioms so far we cannot form a set with more than two elements.

Axiom (Union)**.** *For every set S, there is a set whose elements are exactly the elements of the elements of S.*

This set is unique by extensionality. We denote it by

$$\bigcup_{x \in S} x$$

and refer to it as the *union of S* or the *union of the elements of S* (which is more verbose but perhaps clearer). A set y is an element of $\bigcup_{x \in S} x$ if and only if $y \in x$ for some $x \in S$.

By combining the Union Axiom with the Pairing Axiom, we see that every pair of sets has a union, but of course the Union Axiom is much broader than this restricted statement. We write $a \cup b$ for the union of $\{a, b\}$.

Note that we do not need an intersection axiom. If S is a non-empty set, let x_0 be an element of S. Then

$$\bigcap_{x \in S} x = \{y \in x_0 : y \in x \text{ for all } x \in S\},$$

which defines a set by separation. The restriction that S must be non-empty is necessary, because the intersection of the elements of the empty set should be a universal set, which doesn't exist.[2]

Similarly, using separation we can define

$$S \setminus T = \{s \in S : s \notin T\}.$$

Definition 1.3. The *power set* $\mathcal{P}(S)$ of a set S is the set of all subsets of S.

The power set of S is unique if it exists, by extensionality, but it is not clear that every set has a power set (notation doesn't imply existence). Existence of power

[2] If you don't find it clear why the intersection of the elements of the empty set should be a universal set, it's worth thinking this through carefully. The general principle is that for a commutative and associative operation, applying the operation to no inputs should yield an identity element for the operation. For example, the sum of no elements is 0, and the product of no elements is 1. The union of no sets is the empty set, but the intersection of no sets would have to be a universal set, which doesn't exist.

sets doesn't follow from the axioms so far, because they are all satisfied by the class of countable sets, while the power set of a countably infinite set is not countable.[3]

Axiom (Power Set)**.** *Every set has a power set.*

Even with the Power Set Axiom, our list of axioms is far from complete. So far, they all hold for the class of finite sets, so we have no guarantee that there are any infinite sets.[4] However, we have enough axioms to start building up some of mathematical practice.

The key tool we need is the ordered pair. We will use Kuratowski's definition:

Definition 1.4. The *ordered pair* (x, y) is defined to be $\{\{x\}, \{x, y\}\}$.

This definition is completely ad hoc, and nobody actually thinks the ordered pair (x, y) means anything like $\{\{x\}, \{x, y\}\}$ intuitively. The only purpose of this definition is to offer a purely set-theoretic construction that satisfies the following lemma. Any definition with that property would work equally well.

Lemma 1.5. *For all u, v, x, y, we have $(u, v) = (x, y)$ if and only if $u = x$ and $v = y$.*

The intuition is simple: given $\{\{x\}, \{x, y\}\}$, we can reconstruct x as the only element of both $\{x\}$ and $\{x, y\}$, and then y is the remaining element of these two sets (or $y = x$ if there is no other element). Justifying this argument carefully using the axioms requires a depressingly large number of applications of extensionality, but it's not hard.

Proof. The nontrivial direction is the one starting with $(u, v) = (x, y)$, so suppose $(u, v) = (x, y)$. In other words, $\{\{u\}, \{u, v\}\} = \{\{x\}, \{x, y\}\}$. It will be convenient to break the proof up into cases according to whether $u = v$, $x = y$, or neither equation holds.

If $u = v$, then $\{\{u\}, \{u, v\}\} = \{\{u\}\}$. (We're using extensionality here: $\{u, v\} = \{u\}$ because both sets have the same elements, namely just u.) Then the only element of $\{\{x\}, \{x, y\}\}$ is $\{u\}$, and so $\{x, y\} = \{x\} = \{u\}$. By extensionality, $\{x\} = \{u\}$ implies $x = u$ and $\{x, y\} = \{x\}$ implies $x = y$. Thus, $x = y = u = v$.

Similarly, $x = y$ implies $x = y = u = v$.

The remaining case is $u \neq v$ and $x \neq y$. Then $\{u\} \neq \{u, v\}$ by extensionality (since v is in $\{u, v\}$ but not $\{u\}$), and $\{x\} \neq \{x, y\}$. Furthermore, $\{u\} \neq \{x, y\}$, since $\{x, y\}$ has distinct elements and $\{u\}$ does not, and similarly $\{x\} \neq \{u, v\}$. Now it follows from $\{\{u\}, \{u, v\}\} = \{\{x\}, \{x, y\}\}$ and extensionality that $\{u\} = \{x\}$ and $\{u, v\} = \{x, y\}$. The equation $\{u\} = \{x\}$ implies $u = x$ by extensionality, and then $\{u, v\} = \{x, y\}$ implies $v = y$. □

There are certainly many other ways to define ordered pairs for which this lemma would still be true. However, there are some subtleties. For example, one could try to use the definition $\{x, \{x, y\}\}$. In fact, given an additional axiom we will introduce later (the Axiom of Foundation), that definition can be proved to work, but one must deal with the following issue. If $u = \{x, y\}$ and $x = \{u, v\}$, then

$$\{x, \{x, y\}\} = \{u, \{u, v\}\}.$$

[3] Note that this sentence is just commentary about set theory using our prior knowledge, since of course we have not yet defined countability formally or proved Cantor's theorem.

[4] Again just commentary, since we have not defined finite or infinite sets.

This circular relationship between u and x may seem odd, and in fact it will be ruled out by the Axiom of Foundation, but nothing we have seen so far prohibits it. Kuratowski's definition looks slightly more complicated, but logically it's a little simpler since it avoids this issue.

Once we have defined ordered pairs, we can define the *Cartesian product* $S \times T$. Specifically,

$$S \times T = \{(s,t) \in \mathcal{P}(\mathcal{P}(S \cup T)) : s \in S, t \in T\}.$$

Slightly more formally,

$$S \times T = \{x \in \mathcal{P}(\mathcal{P}(S \cup T)) : \text{there exist } s \in S \text{ and } t \in T \text{ such that } x = (s,t)\}.$$

The one strange part of this definition is the occurrence of $\mathcal{P}(\mathcal{P}(S \cup T))$. In order to apply separation, we must have a set that is guaranteed to contain every ordered pair (s,t) with $s \in S$ and $t \in T$, but fortunately $\mathcal{P}(\mathcal{P}(S \cup T))$ can play this role: both $\{s\}$ and $\{s,t\}$ are elements of $\mathcal{P}(S \cup T)$, and so $\{\{s\},\{s,t\}\}$ is an element of $\mathcal{P}(\mathcal{P}(S \cup T))$. By contrast, simply defining $S \times T$ to be

$$\{(s,t) : s \in S, t \in T\}$$

would be unjustified, since it is not obvious that such a set even exists. In fact, the argument above shows that it does exist, but that requires a construction via separation.

Definition 1.6. A *relation* between S and T is a subset of $S \times T$. A *function* $f\colon S \to T$ is a subset $f \subseteq S \times T$ such that for all $s \in S$, there exists a unique $t \in T$ with $(s,t) \in f$. We write

$$\mathcal{F}(S,T)$$

for the set of all functions from S to T. We call S the *domain* and the T the *codomain*. The *identity function* from S to S is the function $\{(s,t) \in S \times S : s = t\}$.

Of course, if f is a function, we write $f(s) = t$ to mean $(s,t) \in f$. Note that the domain of a function f is uniquely determined by f; it equals

$$\left\{ s \in \bigcup_{x \in f} \bigcup_{y \in x} y : \text{there exists } t \text{ such that } (s,t) \in f \right\}.$$

The purpose of the strange union $\bigcup_{x \in f} \bigcup_{y \in x} y$ is to provide a set that is guaranteed to contain all the elements of the ordered pairs in f, so we can apply separation. (This set exists by two applications of the union axiom.) By contrast, functions do not have unique codomains: a function $f\colon S \to T$ can be viewed as a function from S to T' for any superset T' of T. Thus, if we care what the codomain is, we must specify it.

We will write $f[U]$ to denote the *image* of the set $U \subseteq S$ under a function $f\colon S \to T$, i.e., the set $\{t \in T : t = f(u) \text{ for some } u \in U\}$. Often the notation $f(U)$ is used instead, but that can be ambiguous. For example, if f has domain $\{x, y, \{x,y\}\}$, then $f(\{x,y\})$ is very different from $f[\{x,y\}]$ (which is $\{f(x), f(y)\}$). This may seem like an arcane scenario, but nested sets like these come up frequently in set theory.

Given a function $f\colon S \to T$, its *restriction* $f|_{S'}$ to a subset S' of S is simply

$$\{(s,t) \in f : s \in S'\}.$$

Definition 1.7. A function $f\colon S \to T$ is *injective* (or *one-to-one*, or an *injection*) if for all $s_1, s_1 \in S$, the equation $f(s_1) = f(s_2)$ implies $s_1 = s_2$. It is *surjective* (or *onto*, or a *surjection*) if for every $t \in T$, there exists an $s \in S$ such that $f(s) = t$; equivalently, it is surjective if $f[S] = T$. It is *bijective* (or a *one-to-one correspondence*, or a *bijection*) if it is both injective and surjective.

Note that whether a function is surjective depends on its codomain.

An injective function $f\colon S \to T$ is a bijective function from S to $f[S]$, and we can then define its *inverse function* $f^{-1}\colon f[S] \to S$ by

$$f^{-1} = \{(t,s) \in T \times S : (s,t) \in f\}.$$

Definition 1.8. Given functions $f\colon T \to U$ and $g\colon S \to T$, their *composition* $f \circ g\colon S \to U$ is defined by $(f \circ g)(s) = f(g(s))$ for $s \in S$.

In terms of sets of ordered pairs,

$$f \circ g = \{(s,u) \in S \times U : \text{there exists } t \in T \text{ with } (s,t) \in g \text{ and } (t,u) \in f\}.$$

It is not hard to check that the composition of injective functions is injective, the composition of surjective functions is surjective, and the composition of bijective functions is bijective. Similarly, one can show that if f is a bijection from S to T, then f^{-1} is a bijection from T to S.

However, some intuitions cannot be justified with the axioms we have so far. For example, if there is a surjective map from S to T, must there be an injective map from T to S? The answer will turn out to be yes (as long as $S \neq \emptyset$), as shown in Lemma 7.8, but the proof will depend on the Axiom of Choice.

2. The natural numbers

One crucial step in building mathematics within set theory is constructing the natural numbers. After that, the rest of mathematics ($\mathbb{Z}$, $\mathbb{Q}$, $\mathbb{R}$, $\mathbb{C}$, and beyond) can be built by the usual approach from algebra and analysis classes. However, $\mathbb{N}$ has to come from somewhere, and it's not obvious how to construct the natural numbers out of sets.

As with defining ordered pairs, there is no canonical way to do this, but we will follow a particularly simple and beautiful approach due to von Neumann. Each natural number n will be a specific set with n elements, namely the set of all the previous natural numbers.

Thus, we define $0 = \emptyset$, $1 = \{0\}$, $2 = \{0,1\}$, $3 = \{0,1,2\}$, etc. Of course, the "etc." is hiding something, and we will need to carry this out more formally. We will do so below.

These sets look really confusing if you write them out in terms of braces: $0 = \{\}$, $1 = \{\{\}\}$, $2 = \{\{\},\{\{\}\}\}$, $3 = \{\{\},\{\{\}\},\{\{\},\{\{\}\}\}\}$, etc. However, the weirdness is misleading, and it's best not to focus too much on the nested braces. The point is just that 3 is a set containing three elements, specifically the three natural numbers that came before it.

We will need a way to describe how each natural number is built from the one before it. To see how, note that

$$\begin{aligned} 1 &= \{0\} = 0 \cup \{0\},\\ 2 &= \{0,1\} = 1 \cup \{1\},\\ 3 &= \{0,1,2\} = 2 \cup \{2\}, \end{aligned}$$

etc. We call this construction the successor:

Definition 2.1. The *successor* of a set x is the set $x^+ = x \cup \{x\}$. A set is *inductive* if it contains $\emptyset$ and it contains x^+ whenever it contains x.

Note that every set has a successor, by the pairing and union axioms.

The natural numbers should form an inductive set, and the Axiom of Infinity guarantees that such a set exists. The reason for the name "Axiom of Infinity" is that it is the first axiom that guarantees the existence of an infinite set.

Axiom (Infinity)**.** *There exists an inductive set.*

We should not expect being inductive to be enough to specify a set uniquely. Given any inductive set, one can expand it to include any desired element x as well, as long as we also add its iterated successors x^+, x^{++}, etc. (Strictly speaking our axioms and definitions so far do not guarantee that we can carry out such a construction, but we should expect it to work.)

To single out the natural numbers, we need to characterize them as being the smallest possible inductive set. In other words, they should contain only 0 and its iterated successors.

Definition 2.2. An inductive set is *minimal* if it is a subset of every inductive set.

Extensionality implies that there can be only one minimal inductive set, because given any two minimal inductive sets, each is a subset of the other.

Theorem 2.3. *There exists a minimal inductive set.*

One natural approach to proving this theorem would be to take the intersection of all the inductive sets and show that this intersection is inductive. However, there are too many inductive sets for there to be a set of all inductive sets, so one must tread carefully to avoid paradox. Fortunately, it is not hard to repair the proof, by fixing a single inductive set and working within it, so that we avoid the whole issue of whether there is a set of all inductive sets.

Proof. Let S be any inductive set, and define

$$\omega = \{x \in S : x \text{ is in every inductive set}\}$$

(which exists by separation). Then ω is certainly a subset of every inductive set, by construction. To see that it is a minimal inductive set, we must show that ω is inductive. To begin, $\emptyset \in \omega$, because the empty set is in every inductive set. Furthermore, if $x \in \omega$, then $x^+ \in S$ and x^+ is in every inductive set (by the definition of inductive), and thus $x^+ \in \omega$. It follows that ω is a minimal inductive set, as desired. □

Definition 2.4. The set ω of *natural numbers* is the minimal inductive set.

In set theory it's traditional to use the symbol ω for the set of natural numbers. We'll follow this tradition, partly because it fits nicely with the theory of ordinals, and partly to emphasize that ω is a formal construction within set theory, and we must not make illicit use of our intuitions about natural numbers when analyzing ω.

Defining ω to be the minimal inductive set enables us to give proofs by induction as follows. Suppose S is any subset of ω with the property that $0 \in S$ and if $n \in S$, then $n^+ \in S$. That just means S is an inductive set, and hence $\omega \subseteq S$ by the minimality of ω. On the other hand $S \subseteq \omega$ by assumption, and hence $S = \omega$. It's a little strange philosophically that Definition 2.4 turns the principle of mathematical induction into a definition (we're basically defining ω to be the unique set for which induction works, given our definitions of 0 and the successor). However, it's not vacuous since we have proved such a set exists.

Definition 2.5. A set is *finite* if there is a bijection from it to an element of ω. If $n \in \omega$, then a set has *size* (or *cardinality*) n if there is a bijection between it and the set n.

For example, a set has size 3 if there is a bijection between it and $\{0, 1, 2\}$, because $3 = \{0, 1, 2\}$. It is not hard to prove some simple consequences of Definition 2.5. For example, a set has size 0 iff it is the empty set, and a set has size 1 iff it is of the form $\{x\}$ for some set x.

However, other properties are less clear. For example, it is not obvious from our definitions and axioms that a finite set cannot have two different sizes. See Proposition 4.11 for a proof.

Proposition 2.6. *The union of two finite sets is always finite.*

The proof is an illustration of how the "minimal inductive set" characterization of ω enables us to carry out proofs by mathematical induction.

Proof. Let

$$S = \{n \in \omega : \text{the union of a set of size } n \text{ and a finite set is always finite}\}.$$

We will show that S is inductive, and thus $S = \omega$. Clearly, $\emptyset \in S$, because taking a union with the empty set changes nothing. Thus, we just need to show that S is closed under taking the successor n^+ of n.

Suppose $n \in S$, x has size n^+, and y is finite; we wish to show that $x \cup y$ is finite. Saying x has size n^+ means there is a bijection $f\colon x \to n^+ = n \cup \{n\}$. Suppose $z \in x$ is sent to $n \in n^+$ under f; i.e., $f(z) = n$. Then f restricts to a bijection from $x \setminus \{z\}$ to the set n. Because $n \in S$, we see that $(x \setminus \{z\}) \cup y$ is finite. In other words, there is a bijection $g\colon (x \setminus \{z\}) \cup y \to k$ for some $k \in \omega$. The set we want to understand is $x \cup y$, and we have $x \cup y = \big((x \setminus \{z\}) \cup y\big) \cup \{z\}$. If $z \in (x \setminus \{z\}) \cup y$ (i.e., $z \in y$), then $x \cup y = (x \setminus \{z\}) \cup y$ and we are done. Otherwise, we extend g to a function $h\colon x \cup y \to k \cup \{k\}$ by $h(w) = g(w)$ for $w \in (x \setminus \{z\}) \cup y$ and $h(z) = k$. Then h is a bijection from $x \cup y$ to the natural number k^+, so $x \cup y$ is finite, as desired. □

Proposition 2.7. *A subset of a finite set is always finite.*

Proof. As in the previous proof, we will use induction. Let

$$S = \{n \in \omega : \text{every subset of a set of size } n \text{ is finite}\}.$$

We will show that S is inductive, and thus $S = \omega$. Clearly, $\emptyset \in S$. Thus, we just need to show that S is closed under taking the successor n^+ of n.

Suppose $n \in S$. Because $n^+ = n \cup \{n\}$, every subset of n^+ is either a subset of n or the union of $\{n\}$ with a subset of n. Every subset of n is finite by hypothesis, and the union of $\{n\}$ with a subset of n is finite by Proposition 2.6 (or, more simply, if a subset T of n is in bijection with m, then $T \cup \{n\}$ is in bijection with m^+). Thus, $n^+ \in S$, as required. □

3. Partially, totally, and well-ordered sets

The set ω of natural numbers has given us a rigorous understanding of what cardinality means for finite sets, but it is much less clear what it means in the infinite case. Perhaps surprisingly, the right approach is to study orderings, and not just counting. In the finite case, cardinal numbers (one, two, three, etc., which are used for counting) and ordinal numbers (first, second, third, etc., which are used for ordering) are structurally identical, but in the infinite case we will see that there are major differences. The theory of ordinals is more comprehensive than the theory of cardinals, and we will derive cardinals as a special case, but this will actually be easier than trying to define the cardinals directly. Even though counting seems like a more basic notion than ordering, ordering is easier to pin down mathematically.

Definition 3.1. A *partially ordered set* (or *poset*) is a set S with a relation $R \subseteq S \times S$ satisfying the following three properties for all $x, y, z \in S$, where we write $x \le y$ to mean $(x, y) \in R$:

(1) Reflexivity: $x \le x$.
(2) Antisymmetry: if $x \le y$ and $y \le x$, then $x = y$.
(3) Transitivity: if $x \le y$ and $y \le z$, then $x \le z$.

Strictly speaking, it is an abuse of notation to refer to a set as a poset, because we must specify the ordering as well, and not just the set of elements. To be formally correct, we should define a poset to be an ordered pair (S, R) with $R \subseteq S \times S$, where R represents the ordering relation on S. However, this can be a little cumbersome, so we will typically not be so careful, except in cases where we are considering several different orderings on the same set.

The name "partial ordering" refers to the fact that incomparable elements are allowed, i.e., pairs x, y for which neither $x \le y$ nor $y \le x$.

Definition 3.2. Two elements in a poset are *comparable* if one is greater than or equal to the other, and the poset is *totally ordered* if every pair of elements is comparable. A totally ordered subset of a poset is sometimes called a *chain*.

One example of a total ordering is the usual ordering on $\mathbb{R}$, or on any subset of it. Indeed, every subset of a poset becomes a poset by restricting the ordering. However, most posets are not totally ordered. For example, the power set of any set is a poset under the ordering $\subseteq$, but most pairs of elements are not comparable (neither subset contains the other).

It is traditional to use weak orderings $\le$, allowing the possibility of equality, but for the theory of ordinals it will be convenient to use strict orderings $<$.

Definition 3.3. A *strict partial ordering* replaces the three properties above with the following two:

(1) Antisymmetry: if $x < y$, then $y \not< x$.
(2) Transitivity: if $x < y$ and $y < z$, then $x < z$.

It is not hard to check that this definition is equivalent to Definition 3.1 if we define $<$ to mean $\le$ but not equal, or conversely define $\le$ to mean $<$ or equal. In other words, one can pass back and forth between weak and strict partial orderings by changing whether equality is allowed.

Note that $y \not< x$ does not mean $y \ge x$, because x and y may be incomparable.

Definition 3.4. An *upper bound* for a subset S of a poset is an element x in the poset such that $x \geq y$ for all $y \in S$. It is the *least upper bound* for S if $x \leq x'$ for every upper bound x' for S. A *lower bound* and the *greatest lower bound* are defined identically, except with the inequalities reversed.

There can be at most one least upper bound for a given set, because if x and x' are both least upper bounds, then $x \leq x'$ and $x' \leq x$, which implies $x = x'$. Note also that every element of a poset is an upper bound for the empty set, so x is the least upper bound for $\emptyset$ if and only if $x \leq y$ for all y in the poset.

The least upper bound of a set may not be in the set. For example, $\sqrt{2}$ is the least upper bound of the interval $[0, \sqrt{2})$ in $\mathbb{R}$. Even if a set has an upper bound, there may be no least upper bound. For example, $[0, \sqrt{2}) \cap \mathbb{Q}$ has no least upper bound in $\mathbb{Q}$.

As pointed out above, for every set S, the power set $\mathcal{P}(S)$ is a poset under $\subseteq$. Every subset of $\mathcal{P}(S)$ has a least upper bound, namely the union of its elements.

Theorem 3.5 (Knaster-Tarski fixed point theorem)**.** *Let P be a poset in which every subset has a least upper bound, and let $f\colon P \to P$ be an ordering-preserving map (i.e., if $x \leq y$ then $f(x) \leq f(y)$). Then f has a fixed point. In other words, there exists $x \in P$ such that $f(x) = x$.*

Proof. Let

$$S = \{x \in P : x \leq f(x)\}.$$

Then f maps S to itself, because $x \leq f(x)$ implies $f(x) \leq f(f(x))$. Let y be the least upper bound of S. For every $x \in S$, we have $x \leq y$ and hence $f(x) \leq f(y)$; thus, because $x \leq f(x)$ for $x \in S$, we see that $f(y)$ is also an upper bound for S. Because y is the least upper bound, $y \leq f(y)$, so $y \in S$. However, that means $f(y) \in S$, because S is closed under f, and thus $f(y) \leq y$. It follows that $f(y) = y$, as desired. □

We can now prove a famous and fundamental result about comparing the sizes of sets. We will use the Knaster-Tarski fixed point theorem, but it is possible to give more straightforward proofs. See, for example, *Naive Set Theory* by Halmos. Besides the beauty of the Knaster-Tarski theorem, one motivation for choosing this proof is that it is pleasant to write down formally in axiomatic set theory.

Theorem 3.6 (Cantor-Schröder-Bernstein)**.** *If S and T are sets for which there are injective maps from S to T and from T to S, then there is a bijection between S and T.*

In fact, we will prove a slightly stronger result. Let $f\colon S \to T$ and $g\colon T \to S$ be injective. We will construction a bijection $h\colon S \to T$ such that $h(s) = t$ occurs only when $f(s) = t$ or $g(t) = s$. In other words, we will form h by piecing together f and g^{-1}. Specifically, we will choose a subset A of S and set $h|_A = f|_A$ and $h|_{S \setminus A} = g^{-1}|_{S \setminus A}$. In order for this to work, A must have the property that $f[A]$ and $g^{-1}[S \setminus A]$ are complements of each other in T. We will find such an A using the Knaster-Tarski fixed point theorem.

Proof. Let $f\colon S \to T$ and $g\colon T \to S$ be injective. Consider the partially ordered set $\mathcal{P}(S)$, and define $F\colon \mathcal{P}(S) \to \mathcal{P}(S)$ by

$$F(A) = S \setminus g[T \setminus f[A]].$$

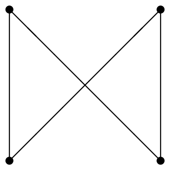

FIGURE 3.1. A poset with two minimal elements.

Recall that $f[A]$ denotes the image of the set A under f, i.e., $\{f(x) : x \in A\}$.

This function is order-preserving: taking the image under f or g preserves inequalities (if $A \subseteq B$ then $f[A] \subseteq f[B]$) while taking complements reverses them (if $A \subseteq B \subseteq S$ then $S \setminus B \subseteq S \setminus A$), so F reverses the order twice and thus preserves it.

Thus, the Knaster-Tarski fixed point theorem provides a subset A of S such that $F(A) = A$. This means A and $g[T \setminus f[A]]$ are complements of each other within S.

The function f restricts to a bijection from A to $f[A]$, because f is injective, and g restricts to a bijection from $T \setminus f[A]$ to $g[T \setminus f[A]]$. We can invert g to get a bijection g^{-1} from $g[T \setminus f[A]]$ to $T \setminus f[A]$, and now we have bijections from the complementary sets A and $g[T \setminus f[A]]$ to the complementary sets $f[A]$ and $T \setminus f[A]$. Thus, if we define

$$h(x) = \begin{cases} f(x) & \text{if } x \in A, \text{ and} \\ g^{-1}(x) & \text{if } x \in g[T \setminus f[A]], \end{cases}$$

then h is a bijection from S to T. □

Definition 3.7. If S is a subset of a poset, then $x \in S$ is a *minimal element* of S if there is no $y \in S$ with $y < x$ (and a *maximal element* is defined the same way but with the opposite inequality).

Note that we do not require that a minimal element be less than or equal to everything else in the set, just that nothing be strictly less than it. A poset can have several different minimal elements. For example, Figure 3.1 is a drawing of a four-element poset, in which the lines indicate comparability (higher is greater) and elements at the same horizontal level are incomparable. This poset has two minimal elements.

In a totally ordered set, a minimal element x of a subset S is an element of S such that $x \leq y$ for all $y \in S$. It follows that S can have at most one minimal element.

Definition 3.8. A *well-ordered set* is a totally ordered set in which every non-empty subset has a minimal element.

The standard example of a well-ordering is the usual ordering on the natural numbers. Most totally ordered sets are not well-ordered; for example, $\mathbb{Z}$ is not well-ordered under its usual ordering (it does not even have a least element), and neither is $\{x \in \mathbb{Q} : x \geq 0\}$ (it has a least element, but the subset $\{x \in \mathbb{Q} : x > 0\}$ does not). Note that every subset of a well-ordered set is itself well-ordered by the restriction of the ordering relation.

Well-ordered sets play a key role in set theory, because they generalize the fundamental property of ordering a finite set: there is always a next element. In a well-ordered set, given any subset S that is not the entire set, there is always a least element outside S. This means if we are building up a subset, until we fill up the entire set there is always a single element to consider next.

We will show that ω is well-ordered, and this is essentially equivalent to the validity of mathematical induction. Sometimes proofs can be expressed particularly elegantly using well-ordering directly; this technique is often called *infinite descent*. For example, stepping slightly outside our formal development of set theory, we can use well-ordering to prove the irrationality of $\sqrt{2}$. This is certainly not the simplest or best-motivated proof, but it is an entertaining illustration of infinite descent.

Suppose $\sqrt{2}$ were rational, and let $\sqrt{2} = p/q$, where p and q are positive integers with q as small as possible. Then

$$\sqrt{2} = \frac{2-\sqrt{2}}{\sqrt{2}-1} = \frac{2-p/q}{p/q-1} = \frac{2q-p}{p-q}.$$

However, $1 < \sqrt{2} < 2$ implies that $q < p < 2q$, and thus $0 < p - q < q$, so we have decreased the denominator of $\sqrt{2}$, which contradicts the minimality of q.

The name "infinite descent" describes how the rationality of $\sqrt{2}$ would lead to an infinite decreasing sequence of positive integers, namely the denominators of the expressions for $\sqrt{2}$. Such a sequence would have no least element, so the set of denominators for $\sqrt{2}$ must be empty.

Well-ordered sets are just on the threshold of what we can analyze: they can be subtle, but their structure is far more predictable than in an arbitrary poset. Let's think informally about what a well-ordered set looks like. The next two paragraphs should be considered merely intuitive, and not part of our formal development of set theory. (One can make them precise using ordinal arithmetic.)

The simplest well-ordered set is $\emptyset$. If a well-ordered set is non-empty, then it must have a least element, which we will label 0. The set could just be $\{0\}$, but if there are any further elements there must be a least one greater than 0, which we will label 1, and a least one greater than that (if the set doesn't end at $\{0, 1\}$), which we will label 2. Thus, in this labeling the well-ordered set must begin $0 < 1 < 2 < \dots$, and it either stops after finitely many steps or continues forever.

However, even continuing forever may not exhaust the set. If there are any further elements (greater than each of $0, 1, 2, \dots$), then there must be a least one, which we will label ω. Then we can continue with $\omega < \omega + 1 < \omega + 2 < \dots$, unless we run out of elements along the way. However, even that may not cover everything: there may be an element $\omega \cdot 2$ greater than all of them. We can continue with $\omega \cdot 2 < \omega \cdot 2 + 1 < \omega \cdot 2 + 2 < \dots$, and then $\omega \cdot 3$ greater than those, and $\omega \cdot 4$, etc. If we don't stop along the way, eventually we reach $\omega \cdot \omega$, also known as ω^2, and higher powers of ω. After still longer (infinite iterations of infinite iterations of infinite iterations), we may reach ω^ω. Beyond that is ω^{ω^ω}, and eventually even

$$\omega^{\omega^{\omega^{\cdot^{\cdot^{\cdot}}}}}.$$

Even this is just the beginning, and well-ordered sets can continue far beyond there. However, the crucial point is that we have no choice along the way: at each step of this process, we simply pass to the next element in the ordering. In Section 5, we'll see more precise versions of these intuitive phenomena.

One potentially confusing point is that the previous paragraph uses "ordinal arithmetic," which is different from the cardinal arithmetic we will study in Section 8 but uses the same notation. To try to minimize confusion, we will avoid using ordinal arithmetic elsewhere.

4. Ordinals and their basic properties

Recall the von Neumann definition of the natural numbers, with $0 = \emptyset$, $1 = \{0\}$, $2 = \{0, 1\}$, $3 = \{0, 1, 2\}$, etc. In other words, the successor $n + 1$ of n is given by $n \cup \{n\}$. In this setting, we can easily define the ordering by $n < m$ iff $n \in m$. Then each natural number equals the set of all natural numbers that come before it, and the natural numbers are strictly well-ordered by $\in$ (although we haven't proved this yet).

Ordinals generalize these two properties: they will be well-ordered by $\in$, and each ordinal will equal the set of all the ordinals that come before it. However, it's not immediately obvious how to use these properties to define the ordinals without circularity. Instead, we will characterize them using some consequences of the properties, and we will show that this definition then leads to the properties we wanted.

To start off, if the elements of an ordinal are supposed to be ordinals, and if all ordinals are supposed to be well-ordered by $\in$, then the elements of each ordinal should be well-ordered by $\in$. That will be the first defining property of an ordinal.

Furthermore, if α, β, and γ are ordinals satisfying $\alpha \in \beta$ and $\beta \in \gamma$, then the transitivity of the ordering should lead to $\alpha \in \gamma$. What that means is that each element β of an ordinal γ should be a subset of γ as well (since that is equivalent to saying every element of β is an element of γ).

Our formal definition of an ordinal simply requires these two properties, but we will see that they are enough to build up the whole theory.

Definition 4.1. An *ordinal* is a set α such that

(1) α is strictly well-ordered under $\in$, and
(2) α is *transitive*; i.e., for all $x \in \alpha$, we have $x \subseteq \alpha$.

Note that we do not directly assume, for example, that the elements of an ordinal must be ordinals themselves. It's difficult to formulate such an assumption as part of the definition without circularity.

It's traditional, although of course not mandatory, to use lowercase Greek letters for ordinals. We'll generally use other letters for sets that either aren't ordinals or aren't yet known to be ordinals, but keep in mind that the notation doesn't really imply anything.

Definition 4.2. The *successor* of an ordinal α is $\alpha^+ = \alpha \cup \{\alpha\}$.

This agrees with our earlier definition of x^+ for an arbitrary set x.

Lemma 4.3. *The empty set $\emptyset$ is an ordinal, and for each ordinal α, its successor α^+ is an ordinal.*

Proof. The empty set satisfies the definition of an ordinal vacuously, since it has no elements.

For the second part, let α be an ordinal. Then α is well-ordered by $\in$, and the only additional element of α^+ is α itself. Every element of α is less than α under the ordering $\in$, so we are simply adding another element above all the elements of α. Now we can see that every non-empty subset of α^+ has a minimal element: if the subset has non-empty intersection with α, then the least element of that intersection is the least element of the whole subset. The only non-empty subset of α^+ that doesn't intersect α is $\{\alpha\}$, and it also has a minimal element (namely, α). Thus, α^+ is well-ordered by $\in$.

Transitivity is also not hard to check. Let x be any element of $\alpha \cup \{\alpha\}$. If $x \in \alpha$, then $x \subseteq \alpha$ by the transitivity of α, so $x \subseteq \alpha^+$. Aside from elements of α, the only other element of $\alpha \cup \{\alpha\}$ is $x = \alpha$ itself, in which case $x \subseteq \alpha^+$ is still true. It follows that α^+ is both well-ordered by $\in$ and transitive, so it is an ordinal. □

Recall that ω denotes the set of natural numbers.

Corollary 4.4. *Every element of ω is an ordinal.*

Proof. Lemma 4.3 is all we need for a proof by induction, and the characterization of ω as the minimal inductive set enables us to carry out the proof. Let

$$S = \{\alpha \in \omega : \alpha \text{ is an ordinal}\}.$$

Then $S \subseteq \omega$, and Lemma 4.3 implies that S is inductive. It follows that $\omega \subseteq S$ (because ω is a minimal inductive set). Thus $S = \omega$, and so every element of ω is an ordinal. □

So far, we have recovered the motivating examples, namely the natural numbers. It will turn out that ω is itself an ordinal. We will also see that ω is the first limit ordinal, other than the empty set.

Definition 4.5. A *limit ordinal* is an ordinal that is not the successor of any ordinal.

Opinions differ as to whether 0 should be considered a limit ordinal: it satisfies the above definition, but it's conceptually different from other limit ordinals. We will allow it in our definition.

Definition 4.6. For ordinals α and β, we define $\alpha < \beta$ to mean $\alpha \in \beta$. Of course, $\alpha \leq \beta$ means $\alpha < \beta$ or $\alpha = \beta$, and $\beta > \alpha$ means $\alpha < \beta$, etc.

As an immediate consequence of the definitions so far, we have $\alpha < \alpha^+$ for all ordinals α.

Lemma 4.7. *For each ordinal α, either $\alpha = 0$ or $\alpha > 0$.*

Proof. Suppose $\alpha \neq 0$. In other words, $\alpha \neq \emptyset$, so because α is well-ordered under $\in$, there exists an $\in$-minimal element $x \in \alpha$. Minimality means no element of α is an element of x. However, by transitivity, $x \subseteq \alpha$. Thus, the fact that no element of α is in x means $x = \emptyset$. It follows that $\emptyset = x \in \alpha$, so $0 < \alpha$, as desired. □

Lemma 4.8. *No ordinal α satisfies $\alpha \in \alpha$.*

Proof. By assumption, α is strictly well-ordered under $\in$. The antisymmetry of this well-ordering implies that no element x of α satisfies $x \in x$. (Recall that antisymmetry for a strict well-ordering means $x < y$ implies $y \not< x$. Taking $y = x$ implies that $x < x$ is impossible, and for ordinals we are using $<$ to mean $\in$.) If $\alpha \in \alpha$, then taking $x = \alpha$ contradicts antisymmetry. □

In fact, in ZFC no set can be a member of itself, but we cannot prove that given our axioms so far. We need a new axiom:

Axiom (Foundation)**.** *Every non-empty set contains an element disjoint from it.*

In other words, every non-empty set S contains an element x such that no $y \in S$ satisfies $y \in x$. This means x is an $\in$-minimal element of S. Thus, the Axiom of Foundation implies that every set that is totally ordered under $\in$ is actually well-ordered under $\in$.

Corollary 4.9. *No set is an element of itself.*

Proof. Suppose $S \in S$. If we let $T = \{S\}$, then $T \cap S = \{S\}$. The only element of T is not disjoint from T, which contradicts the Axiom of Foundation. □

Corollary 4.10. *There do not exist sets $S_0, S_1, S_2, \dots$ such that $S_0 \ni S_1 \ni S_2 \ni \dots$. More formally, there do not a exist a set R and a function $f\colon \omega \to R$ such that $f(i^+) \in f(i)$ for all $i \in \omega$.*

Note that the formal version is needed because "…" is imprecise. In general, an infinite sequence simply means a function defined on ω. This sort of translation is routine once you are used to it.

Proof. The set $T = \{S_0, S_1, \dots\}$ contradicts the Axiom of Foundation. In formal terms, let T be the image $f[\omega]$. For every element $f(i)$ of T, we have $f(i^+) \in T \cap f(i)$, and so no element of T is disjoint from T. This contradicts the Axiom of Foundation. □

The Axiom of Foundation is the least useful axiom of set theory for everyday mathematics: its primary purpose to rule out pathological examples, but these examples can actually be convenient (see the book *Vicious Circles* by Barwise and Moss), and in any case they won't hurt us if we ignore them. One can define well-founded sets, which are the sets such that they, their elements, their elements' elements, etc. all satisfy the Axiom of Foundation. Then if one ever needs the Axiom of Foundation, instead of assuming it one can insert the words "well-founded" into the theorem statement. However, assuming the axiom is standard.

Intuitively, the Axiom of Foundation tells us that all sets are constructed iteratively, starting with the empty set. More precisely, if you start with any set, look at one of its elements, look at one of that set's elements, etc., you must eventually reach the empty set (by Corollary 4.10). This idea is elaborated on in the cumulative hierarchy, which we will briefly discuss later.

Proposition 4.11. *There is no bijection between two distinct elements of ω.*

Proof. Let

$$S = \{n \in \omega : \text{there exists no bijection from } n \text{ to } m \text{ with } m \in \omega \text{ and } m \neq n\}.$$

We will show that S is inductive, from which it follows that $S = \omega$.

Clearly, $\emptyset \in S$, since there is no bijection from it to any non-empty set. Now suppose $n \in S$. We must show that n^+ is not in bijection with any element of ω other than n^+. If it is, then that element must be non-empty and is therefore the successor m^+ of a natural number. (This is easy: the set of all natural numbers that are either 0 or a successor is inductive, so it contains all the natural numbers. Alternatively, if a natural number k were neither 0 nor a successor, then $\omega \setminus \{k\}$ would be inductive.)

Thus, suppose f is a bijection from n^+ to m^+. We will construct a bijection from n to m, and then the fact that $n \in S$ implies that $m = n$, as desired.

The function f maps $n \cup \{n\}$ to $m \cup \{m\}$. If $f(n) = m$, then f restricts to a bijection from n to m. Thus, suppose $f(n) \neq m$, and let ℓ be the unique element of n^+ (and, in particular, element of n) such that $f(\ell) = m$. Then we define a new function $g \colon n \to m$ by $g(x) = f(x)$ for $x \neq \ell$ and $g(\ell) = f(n)$. We will show that g is a bijection.

To check that g is surjective, note that because $g(x) = f(x)$ for $x \in n$ with $x \neq \ell$, the image of g includes every value in the image of f except for $f(\ell)$ (which we don't want anyway, since it equals m) and $f(n)$, which is covered by $g(\ell) = f(n)$. Thus, because f is surjective onto m^+, g is surjective onto m. To check that it is injective, suppose $g(x) = g(y)$ with $x, y \in n$. If neither x nor y equals ℓ, then $f(x) = f(y)$ and thus $x = y$ since f is injective. If $x = \ell \neq y$, then $f(n) = g(x) = g(y) = f(y)$, from which $y = n$ follows, and this contradicts $y \in n$ (by Lemma 4.8). Thus, $g(x) = g(y)$ implies $x = y$ as desired. □

Lemma 4.12. *For all ordinals α and β, their intersection $\alpha \cap \beta$ is an ordinal.*

Proof. Certainly $\alpha \cap \beta$ is well-ordered under $\in$, as a subset of α (or β). To prove transitivity, we begin by noting that if $x \in \alpha \cap \beta$, then $x \in \alpha$ and $x \in \beta$. Because α and β are transitive, it follows that $x \subseteq \alpha$ and $x \subseteq \beta$, and so $x \subseteq \alpha \cap \beta$. Thus, $\alpha \cap \beta$ is transitive, so it is an ordinal. □

Lemma 4.13. *If α and β are ordinals, then $\alpha \leq \beta$ iff $\alpha \subseteq \beta$.*

Proof. For the easy direction, suppose $\alpha \leq \beta$. This means $\alpha = \beta$ or $\alpha \in \beta$, and $\alpha \in \beta$ implies $\alpha \subseteq \beta$ by transitivity. Either way, $\alpha \subseteq \beta$.

For the other direction, suppose $\alpha \subseteq \beta$. Assume furthermore than $\alpha \neq \beta$, since the $\alpha = \beta$ case is trivial. Then let x be the least element of β that is not an element of α (which exists because β is well-ordered). We will show that $x = \alpha$, and hence $\alpha \in \beta$.

First, we prove $x \subseteq \alpha$. By the transitivity of β, we have $x \subseteq \beta$. All smaller elements of β than x are in α, by the definition of x. This means for all $y \in x$, we have $y \in \alpha$, or in other words $x \subseteq \alpha$.

Now we must prove $\alpha \subseteq x$. Suppose $y \in \alpha$. Because β is totally ordered by $\in$, we must have $x = y$, $x \in y$, or $y \in x$, and we want to show that the last possibility is the true one. If $x = y$, then $x \in \alpha$ because $y \in \alpha$, and $x \in \alpha$ contradicts the definition of x. If $x \in y$, then we apply the transitivity of the ordering of β by $\in$ to pass from $x \in y \in \alpha$ to $x \in \alpha$, which again contradicts the definition of x. The only remaining possibility is $y \in x$. This shows that $\alpha \subseteq x$, so it follows that $\alpha = x$ and hence $\alpha \in \beta$, as desired. □

Corollary 4.14. *If α and β are ordinals, then $\alpha < \beta$ if and only if $\alpha^+ \leq \beta$, $\alpha \leq \beta$ if and only if $\alpha < \beta^+$, and $\alpha < \beta$ if and only if $\alpha^+ < \beta^+$. If α is an ordinal and β is a limit ordinal, then $\alpha < \beta$ implies $\alpha^+ < \beta$.*

Proof. For the first equivalence, we begin by supposing that $\alpha < \beta$. Then $\alpha \in \beta$ and, by transitivity, $\alpha \subseteq \beta$. It follows that $\alpha^+ = \alpha \cup \{\alpha\} \subseteq \beta$, which implies $\alpha^+ \leq \beta$ by Lemma 4.13. Conversely, if $\alpha^+ \leq \beta$, then $\alpha \cup \{\alpha\} \subseteq \beta$ and hence $\alpha \in \beta$, as desired.

For the second equivalence, $\alpha < \beta^+$ means $\alpha \in \beta \cup \{\beta\}$, which holds iff $\alpha \in \beta$ or $\alpha = \beta$.

The third equivalence follows from the first two: $\alpha < \beta$ is equivalent to $\alpha^+ \leq \beta$ by the first equivalence, which is in turn equivalent to $\alpha^+ < \beta^+$ by the second.

Finally, suppose $\alpha < \beta$ and β is a limit ordinal. Then $\alpha^+ \leq \beta$, but $\alpha^+ \neq \beta$ since β is not a successor, and hence $\alpha^+ < \beta$. □

Proposition 4.15. *Every set of ordinals is (strictly) totally ordered by* $<$.

The reason we don't say "the set of all ordinals" is that, as we will see below, there are too many ordinals to form a set.

Proof. We must verify that $<$ is antisymmetric and transitive, and that all pairs of elements are comparable.

For antisymmetry, suppose α and β are ordinals satisfying $\alpha < \beta$ and $\beta < \alpha$, i.e., $\alpha \in \beta$ and $\beta \in \alpha$. By the transitivity property of ordinals, $\beta \in \alpha$ implies $\beta \subseteq \alpha$, and combining this with $\alpha \in \beta$ yields $\alpha \in \alpha$, which contradicts Lemma 4.8. Thus, $<$ is antisymmetric for ordinals.

Transitivity is similar: suppose α, β, and γ are ordinals satisfying $\alpha < \beta < \gamma$. Then $\alpha \in \beta$ and $\beta \in \gamma$, and the latter implies that $\beta \subseteq \gamma$, so we conclude that $\alpha \in \gamma$, as desired.

Finally, comparability follows from Lemma 4.13. If α and β are ordinals, then so is $\alpha \cap \beta$ by Lemma 4.12. If it equals either of α or β, then one is a subset of the other and Lemma 4.13 implies that α and β are comparable. However, if it equals neither, then Lemma 4.13 implies that it is an element of both. Then, however, $\alpha \cap \beta \in \alpha$ and $\alpha \cap \beta \in \beta$ together imply that $\alpha \cap \beta \in \alpha \cap \beta$, which contradicts Lemma 4.8. □

Proposition 4.16. *Every non-empty set of ordinals has a minimal element.*

This follows immediately from the Axiom of Foundation, which says that every non-empty set contains an $\in$-minimal element. We'll give a proof that doesn't use that axiom, just to show that it can be done:

Proof. Let S be a non-empty set of ordinals. A minimal element of S is just an element $\alpha \in S$ such that there exists no $\beta \in S$ satisfying $\beta < \alpha$ (equivalently, $\beta \in \alpha$). In other words, it is an element $\alpha \in S$ satisfying $\alpha \cap S = \emptyset$.

Let α be any element of S. If $\alpha \cap S = \emptyset$, then α is minimal. Otherwise, suppose $\alpha \cap S \neq \emptyset$. Because α is well-ordered under $\in$ (by the definition of an ordinal), its non-empty subset $\alpha \cap S$ has a least element β. Then $\beta \in S$ and $\beta \cap S = \emptyset$: if $\beta \cap S \neq \emptyset$, then there is some $\gamma < \beta$ in S, but then $\gamma < \beta < \alpha$ implies $\gamma < \alpha$ (by Proposition 4.15). Thus, $\gamma \in \alpha \cap S$ and $\gamma < \beta$, which contradicts the minimality of β in $\alpha \cap S$. It follows that $\beta \cap S = \emptyset$, so β is the minimal element of S. □

Corollary 4.17. *Every set of ordinals is strictly well-ordered under* $<$.

This follows from Propositions 4.15 and 4.16.

Corollary 4.18. *The set* ω *of natural numbers is an ordinal, and in particular a limit ordinal.*

Proof. Corollary 4.17 implies that ω is well-ordered under $\in$. To see that it is transitive, we simply prove that $S = \{x \in \omega : x \subseteq \omega\}$ is inductive. This is straightforward: clearly $\emptyset \in S$, and if $x \in S$, then $x \subseteq \omega$ and $x \in \omega$ by assumption, so it follows that $x \cup \{x\} \subseteq \omega$ as well. Thus, $S = \omega$, so ω is transitive and is therefore an ordinal.

To see that it is a limit ordinal, suppose $\omega = \alpha^+$. Then $\alpha \in \omega$, because $\alpha \in \alpha^+$. However, ω is an inductive set, so it is closed under taking the successor. Thus, $\alpha \in \omega$ implies $\alpha^+ \in \omega$, so we conclude that $\omega \in \omega$, which contradicts Lemma 4.8. □

Proposition 4.19. *Every element of an ordinal is an ordinal.*

Proof. Let α be an ordinal and $x \in \alpha$. Then $x \subseteq \alpha$ by transitivity of α, and every subset of a well-ordered set is well-ordered. Thus, x is strictly well-ordered under $\in$.

For transitivity, suppose $y \in x$. We would like to conclude that $y \subseteq x$; in other words, we would like to show that for all $z \in y$, we have $z \in x$. Because α is transitive, $x \in \alpha$ implies $x \subseteq \alpha$, and then $y \in x$ implies $y \in \alpha$. Furthermore, the transitivity of α then implies that $y \subseteq \alpha$, so $z \in y$ implies $z \in \alpha$. Thus, x, y, and z are all elements of α. However, α is totally ordered under $\in$, so $y \in x$ and $z \in y$ imply $z \in x$, as desired. □

Theorem 4.20 (Burali-Forti paradox)**.** *There is no set of all ordinals.*

Despite the hyphen, Burali-Forti was one person.

Proof. Call the hypothetical set of all ordinals Ω. As a set of ordinals, Ω is well-ordered under $\in$ (Corollary 4.17). Furthermore, Ω is transitive: if $\alpha \in \Omega$, then $\alpha \subseteq \Omega$ since all elements of α are ordinals (Proposition 4.19). Thus, Ω is itself an ordinal. However, that means $\Omega \in \Omega$, which contradicts Lemma 4.8. □

From a modern perspective, this is not really a paradox, but simply a proof that the set of all ordinals cannot exist.

Lemma 4.21. *The union of any set of ordinals is an ordinal.*

Proof. Let S be any set of ordinals, and let $U = \bigcup_{\alpha \in S} \alpha$. Every element of an ordinal is an ordinal, so U is a set of ordinals and is therefore well-ordered under $\in$. Thus, to show that U satisfies the definition of an ordinal, we just need to show that it is transitive. Suppose $x \in U$. This means there exists some $\alpha \in S$ such that $x \in \alpha$. Then the transitivity of α implies that $x \subseteq \alpha$, and it follows that $x \subseteq U$, as desired. □

Corollary 4.22. *Let S be any set of ordinals. Then*

$$\bigcup_{\alpha \in S} \alpha$$

is the least upper bound for S (i.e., the smallest ordinal that is greater than or equal to each element of S).

Proof. Combine Lemmas 4.21 and 4.13. □

Lemma 4.23. *An ordinal is a limit ordinal if and only if it equals the union of all smaller ordinals.*

Proof. Let α be an ordinal. If $\alpha = \beta^+$, then β is an upper bound for all the ordinals less than α, by Corollary 4.14. Thus,

$$\bigcup_{\gamma < \alpha} \gamma = \beta < \alpha,$$

by Corollary 4.22.

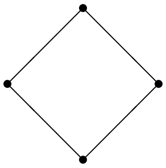

Figure 5.1. A poset with a nontrivial automorphism.

By contrast, suppose α is a limit ordinal. Corollary 4.22 implies that

$$\bigcup_{\beta<\alpha} \beta \le \alpha,$$

since α is an upper bound for all $\beta < \alpha$. To show that the union equals α, we must show that it contains each element γ of α. If $\gamma \in \alpha$, then $\gamma < \alpha$ and hence $\gamma^+ < \alpha$ since α is a limit ordinal (see Corollary 4.14). Thus, taking $\beta = \gamma^+$ shows that

$$\gamma \in \gamma^+ \subseteq \bigcup_{\beta<\alpha} \beta.$$

It follows that

$$\bigcup_{\beta<\alpha} \beta = \alpha,$$

as desired. □

5. Ordinals and well-ordered sets

Definition 5.1. An *isomorphism* between two posets S and T (with orderings $<_S$ and $<_T$, respectively) is a bijection $f\colon S \to T$ such that for all $x, y \in S$,

$$x <_S y \quad \Leftrightarrow \quad f(x) <_T f(y).$$

We say that S and T are *isomorphic* (written $S \cong T$) if there is an isomorphism between them. An *automorphism* of a poset is an isomorphism to itself.

Usually we will abuse notation slightly and denote both orderings by the same symbol. This shouldn't cause much confusion, since it's generally clear which ordering is meant just from context, but of course it's important to keep in mind that different posets will have different orderings.

Every poset has at least one automorphism, namely the identity function (called the trivial automorphism). It's easy to check that isomorphism is an equivalence relation: it is reflexive since every poset is isomorphic to itself via the identity function, symmetric since the inverse of an isomorphism is an isomorphism, and transitive since composing two isomorphisms yields an isomorphism.

For an example of an isomorphism, the map $x \mapsto \tan(\pi x/2)$ is an isomorphism from the interval $(-1, 1)$ to $\mathbb{R}$, where both sets are given their usual orderings. The posets $\mathbb{R}$ has many automorphisms, since as $x \mapsto x + 1$. For an example in a poset that is not totally ordered, consider the four-element poset shown in Figure 5.1, where the diagonal lines indicate covering relations (the higher-up element is greater than the lower) and the two middle elements are incomparable. This poset has two automorphisms, with the nontrivial one switching the two incomparable elements.

Proposition 5.2. *Well-ordered sets have no nontrivial automorphisms.*

Proof. Let S be an well-ordered set, and let f be any automorphism of S. Define

$$T = \{x \in S : f(x) \neq x\}.$$

If $T = \emptyset$, then f is the identity function, as desired. Otherwise, let x be the least element of T. Because f preserves the ordering, $f(x)$ must be the minimal element of the image $f[T]$ of T under f. However, $f[T] = T$, because f fixes the complement of T pointwise. It follows that $f(x)$ is the minimal element of T, and thus $f(x) = x$, which contradicts $x \in T$. Therefore $T = \emptyset$ and f must be the identity function. □

Corollary 5.3. *If two well-ordered sets are isomorphic, then there is a unique isomorphism in each direction between them.*

Proof. Suppose $f\colon S \to T$ and $g\colon S \to T$ are two isomorphisms between well-ordered sets. Then the composition $f^{-1} \circ g$ is an automorphism of S, so it must be the identity function, and hence $f = g$. □

Definition 5.4. If S is a well-ordered set and $x \in S$, the *initial segment* S_x is $\{y \in S : y < x\}$.

For example, if α and β are ordinals with $\alpha < \beta$, then α is an initial segment of β. Specifically, $\alpha \in \beta$, from which $\alpha \subseteq \beta$ follows, and therefore α is the initial segment β_α.

Proposition 5.5. *If S is well-ordered and $x \in S$, then $S \ncong S_x$.*

Proof. The proof is very much like that of Proposition 5.2. Suppose $f\colon S \to S_x$ is an isomorphism, and let

$$T = \{y \in S : f(y) \neq y\}.$$

If $T = \emptyset$, then $f(x) = x$ and so S_x must contain x, which contradicts the definition of S_x. Otherwise, let y be the least element of T. Because y is the least element of $S \setminus S_y$, it follows that $f(y)$ is the least element of $f[S] \setminus f[S_y]$. However, $f[S] = S_x$ by assumption, and every element of S_y is fixed by f by the definition of y. Thus, $f(y)$ must be the least element of $S_x \setminus S_y$, which is y (if indeed there is any element of $S_x \setminus S_y$). Thus, $f(y) = y$, which contradicts $y \in T$. □

Corollary 5.6. *Two distinct ordinals cannot be isomorphic as well-ordered sets.*

Proof. If α and β are ordinals satisfying $\alpha \neq \beta$, and without loss of generality $\alpha < \beta$, then $\alpha = \beta_\alpha$, so it follows from Proposition 5.5 that $\alpha \ncong \beta$. □

Corollary 5.7. *No two different initial segments of a well-ordered set are isomorphic.*

Proof. Suppose S is a well-ordered set and $S_x \cong S_y$ with $x \neq y$. Without loss of generality, $x < y$, and then $S_x = (S_y)_x$. However, then $(S_y)_x \cong S_y$, which contradicts Proposition 5.5. □

Definition 5.8. A subset S of a poset is *downwards closed* if whenever $x \in S$ and $y \leq x$, it follows that $y \in S$.

Lemma 5.9. *The only downwards-closed subsets of a well-ordered set are the set itself and its initial segments.*

Proof. Let T be a downwards closed subset of a well-ordered set S. If $T \neq S$, then let x be the least element of S not in T. Then by the definition of x, we have $S_x \subseteq T$, and $T \subseteq S_x$ holds because if T contained any greater element it would also contain x (being downwards closed). Thus, $T = S_x$, as desired. □

Theorem 5.10. *If S and T are well-ordered sets, then $S \cong T$, or $S_x \cong T$ for some $x \in S$, or $S \cong T_y$ for some $y \in T$.*

The intuition is that as we try to build an isomorphism from S to T, there is no flexibility or choice along the way. The minimal element of S must map to T, then the same must be true for the next least elements, etc. If we run out of elements in one set before the other, we end up with an isomorphism from it to an initial segment of the other. Otherwise, we get an isomorphism between S and T. One can make this intuition precise and use it to prove the theorem, but we will use a slicker approach. It essentially encodes the same idea, namely that a partial isomorphism can always be extended until one set runs out of elements, but it uses a particularly nice characterization of how to match up elements of S and T. Specifically, the isomorphism should map $x \in S$ to $y \in T$ if and only if $S_x \cong T_y$.

Proof. Let

$$f = \{(x, y) \in S \times T : S_x \cong T_y\}.$$

Then f is a function from some subset of S to T, because if $(x, y_1) \in f$ and $(x, y_2) \in f$, then $T_{y_2} \cong T_{y_1}$, which implies $y_1 = y_2$ by Corollary 5.7. Therefore we can use functional notation (writing $f(x) = y$ to mean $(x, y) \in f$). Furthermore, f is injective, because $f(x_1) = f(x_2)$ implies $S_{x_1} \cong S_{x_2}$ and thus $x_1 = x_2$. This means f is a bijection from its domain to its image.

Furthermore, f is a poset isomorphism from its domain to its image: if $x_1 < x_2$, then $T_{f(x_1)}$ is isomorphic to an initial segment of $T_{f(x_2)}$, because they are isomorphic to S_{x_1} and S_{x_2}, respectively. This implies that $f(x_1) < f(x_2)$, because if $f(x_1) \geq f(x_2)$, then $T_{f(x_1)}$ would be isomorphic to an initial segment of itself.

Both the domain and the image of f are downwards closed: if $S_x \cong T_y$, then every initial segment of S_x is isomorphic to an initial segment of T_y, and vice versa. Thus, the domain of f is either S or an initial segment of S, by Lemma 5.9, and the image is either T or an initial segment of T.

All that remains to be shown is that the domain is S or the image is T. If the domain of f is not S and the image is not T, let x be the least element of S not in the domain, and let y be the least element of T not in the image. Then f is an isomorphism from S_x to T_y, so $(x, y) \in f$ after all. □

Theorem 5.11. *Every well-ordered set is isomorphic to a unique ordinal.*

In other words, ordinals classify all possible well-ordered sets, by bringing them into a canonical form: each well-ordered set is isomorphic to an ordinal, and two well-ordered sets are isomorphic if and only if the corresponding ordinals are equal. Furthermore, this also canonically labels the elements of each well-ordered set, since the isomorphism from it to an ordinal is unique.

To prove Theorem 5.11, we will need the last axiom of ZF set theory (and the last axiom of ZFC except for the Axiom of Choice):

Axiom (Replacement)**.** *Suppose $\varphi(x, y)$ is a mathematical statement and S is a set such that for all $x \in S$, there is a unique set y such that $\varphi(x, y)$ holds. Then there is a set T such that for all y, $y \in T$ iff there exists $x \in S$ such that $\varphi(x, y)$.*

As with the previous axioms, we will make precise later what is meant by a "mathematical statement." The intuition here is that the only thing keeping T from being a set is that it might be too big. If we form it by replacing each element x in a set S with a unique y, then it will have the same size as S and should not be too big to be a set.

We can think of $\varphi(x, y)$ as defining a functional relationship between x and y: for each $x \in S$ it defines a unique y. However, this is a little bit tricky given the formalization of functions we have been using. Specifically, without assuming the Axiom of Replacement it is not clear that φ actually defines a function: we want a define the corresponding function as the set of ordered pairs in $S \times T$ consisting of exactly the pairs (x, y) satisfying $\varphi(x, y)$, and we can't even get off the ground without having the set T already!

A variant of replacement has "at most one set y" instead of "a unique set y." This looks like a slight generalization, but one can prove equivalence with the version above as follows. Suppose $\varphi(x, y)$ is satisfied by at most one y for each $x \in S$. Furthermore, suppose there is some $x \in S$ for which there is a y satisfying $\varphi(x, y)$ (otherwise we are building the empty set), and let y_0 be such a y. Then define

$$\varphi'(x, y) = \varphi(x, y) \vee \big((\neg \exists z\, \varphi(x, z)) \to (y = y_0)\big).$$

In other words, we allow $y = y_0$ when no y can satisfy the original formula. Then φ' is suitable for use in the Axiom of Replacement, while defining the same set.

Proof of Theorem 5.11. Let S be a well-ordered set, and define

$$T = \{x \in S : \text{there exists an ordinal } \alpha \text{ such that } \alpha \cong S_x\}.$$

For each $x \in S$, the ordinal α with $\alpha \cong S_x$ is unique if it exists, by Corollary 5.6. Thus, by the Axiom of Replacement,

$$\{\alpha : \alpha \text{ is an ordinal and there exists } x \in S \text{ with } S_x \cong \alpha\}$$

is a set. Call this set β.

First, we observe that β is an ordinal: it is well-ordered under $\in$ because it is a set of ordinals, and it is transitive because if $\alpha \in \beta$ and $\gamma \in \alpha$, we have an isomorphism $f \colon \alpha \to S_x$ for some $x \in S$, and then

$$\gamma = \alpha_\gamma \cong (S_x)_{f(\gamma)} = S_{f(\gamma)},$$

so $\gamma \in \beta$, as desired.

Now we can apply Theorem 5.10. If $\beta \cong S_x$ for some $x \in S$, then $\beta \in \beta$, which contradicts Lemma 4.8. If $\beta_\alpha \cong S$ for some $\alpha \in \beta$, then $\alpha \cong \beta_\alpha \cong S$, but $\alpha \in \beta$ means $\alpha \cong S_x$ for some $x \in S$, and thus $S \cong S_x$, which contradicts Proposition 5.5. Thus, the only remaining possibility is $\beta \cong S$, so S is isomorphic to an ordinal, which is unique by Corollary 5.6. □

Definition 5.12. The *order type* type(S) of a well-ordered set S is the unique ordinal isomorphic to it.

For several applications, it will be important to be able to define functions on well-ordered sets recursively in terms of their previous values. This technique is called *transfinite recursion*:

Theorem 5.13. *Let S be a well-ordered set, and let*

$$g \colon \{(x, h) : x \in S \text{ and } h \text{ is a function from } S_x \text{ to } T\} \to T$$

be a function. Then there is a unique function $f\colon S \to T$ such that for all $x \in S$,

$$f(x) = g(x, f|_{S_x}).$$

If x is the least element of S, then $S_x = \emptyset$ and $f(x) = g(x, \emptyset)$. Each further value $f(x)$ is then uniquely determined by what came before, i.e., by the restriction of f to S_x. Note that we cannot use a recurrence that relies just on the previous value, i.e., the value at the immediate predecessor of x, because there may not be such a value. For example, ω has no predecessor as an element of the ordinal ω^+.

Note also that the theorem statement builds the initial conditions into the recurrence. Specifically, if x is the least element of S, then $f(x) = g(x, f|_\emptyset) = g(x, \emptyset)$.

Proof. Call a function f admissible if its domain is a downwards-closed subset of S and it satisfies the recurrence

$$f(x) = g(x, f|_{S_x}).$$

There can be at most one admissible function on any downwards-closed subset of S: if there were two, say f and f', and if we let x be the first point at which they differ, then $f|_{S_x} = f'|_{S_x}$ and hence $f(x) = f'(x)$.

Now let f be the union of all the admissible functions. Then f is a function defined on a downwards-closed subset of S, because no two admissible functions can ever disagree at any point, and it too satisfies

$$f(x) = g(x, f|_{S_x})$$

and is thus admissible. The domain of f must be S; otherwise, if we let x be the least point outside the domain, then we could extend f by defining $f(x) = g(x, f|_{S_x})$, which would contradict the containment of all admissible functions within f. □

Proposition 5.14. *Let T be a well-ordered set, and give $S \subseteq T$ the induced well-ordering. Then* $\operatorname{type}(S) \le \operatorname{type}(T)$.

Proof. Let t be the least element of T (if $T = \emptyset$ then the proposition is trivial). By transfinite recursion, there is a unique function $f\colon S \to T$ satisfying

$$f(x) = \begin{cases} \text{least element of } T \setminus f[S_x] & \text{if } f[S_x] \ne T \text{, and} \\ t & \text{otherwise.} \end{cases}$$

First, we verify that $f(x) \le x$ for all $x \in S$. If not, then let x be the least x such that $f(x) > x$. Then $f[S_x] \subseteq S_x$ by the minimality of x. However, $x \in T \setminus S_x$, so the least element of $T \setminus S_x$ is no larger than x, and this contradicts $f(x) > x$.

Now the inequality $f(x) \le x$ means $f[S_x] \subseteq S_x$ for all x. In particular, $f[S_x] \ne T$ since $x \in T \setminus f[S_x]$, so f always satisfies

$$f(x) = \text{least element of } T \setminus f[S_x].$$

(The only purpose of the second clause in the definition of f was that Theorem 5.13 requires $g(x, f|_{S_x})$ to be defined under all circumstances. However, one can define it arbitrarily in cases that will never arise.)

If $x \le y$, then $T \setminus S_y \subseteq T \setminus S_x$ and hence $f(x) \le f(y)$. Thus, f is order-preserving, and it is injective because $f(x) \notin f[S_x]$. Furthermore, $f[S]$ is downwards-closed, because every element less than $f(x)$ must be in $f[S_x]$. Therefore f is an isomorphism from S to T or to an initial segment of T, so the order types satisfy $\operatorname{type}(S) \le \operatorname{type}(T)$, as desired. □

Proposition 5.15. *For every set S, there exists an ordinal α such that there is no injection from α to S.*

Proof. We would like to define

$$\alpha = \{\beta : \beta \text{ is an ordinal and there is an injective map from } \beta \text{ to } S\}.$$

However, this is a little tricky, because it is not clear that any set contains all these ordinals, so we cannot just use separation. First, we'll assume there is such a set and go on to complete the proof, and then we'll justify it using replacement.

Given that α exists, it is an ordinal: it is well-ordered under $\in$ because it is a set of ordinals, and it is transitive because if $\beta \in \alpha$ and $\gamma \in \beta$, then $\gamma \subseteq \beta$ by the transitivity of β, and the injection from β to S restricts to one from γ to S, so $\gamma \in \alpha$, as desired. However, α cannot have any injective map to S, since if it did, then $\alpha \in \alpha$.

Thus, all we need to do is to show that there is a set α as defined above. We start with the set of pairs consisting of a subset of S and a well-ordering of that subset. Let

$$U = \{(T, R) \in \mathcal{P}(S) \times \mathcal{P}(S \times S) : R \subseteq T \times T \text{ and } R \text{ is a well-ordering of } T\}.$$

For every $(T, R) \in U$, there is a unique ordinal β that is isomorphic to T with the ordering R, and these ordinals are exactly the ordinals that have injective functions into S: the isomorphism from β to T is such a function, and conversely an injective map from β to S defines a well-ordering on the image of the map and thus leads to an element of U. Now applying the Axiom of Replacement to replace (T, R) with the ordinal β shows that

$$\{\beta : \beta \text{ is an ordinal and there is an injective map from } \beta \text{ to } S\}$$

is indeed a set. □

Proposition 5.15 shows that ordinals can in a certain sense be arbitrarily large. That's just one step removed from saying that every set can be well-ordered, but for that we'll need the final axiom of ZFC.

6. The Axiom of Choice

Definition 6.1. A *choice function* for a set S is a function $f \colon S \to \bigcup_{x \in S} x$ such that $f(x) \in x$ for every $x \in S$.

In other words, if we view S as a collection of sets, a choice function chooses an element $f(x)$ in each set $x \in S$. Obviously, there cannot be a choice function for S if $\emptyset \in S$, but the Axiom of Choice says that is the only restriction.

Axiom (Choice)**.** *Every set of non-empty sets has a choice function.*

This is the final axiom of ZFC (the Zermelo-Fraenkel axioms with choice), while ZF consists of just the previous axioms.

The Axiom of Choice is radically different in character from ZF, because it in no way specifies what the choice function should be. By contrast, the previous axioms asserted the existence of sets that could be uniquely defined by some formula. The difficulty with the Axiom of Choice is that it is not clear there exists a function that describes infinitely many choices at once. Even aside from the philosophical question about whether it is possible to make infinitely many choices at once, set theory might be missing some functions we would naively expect to be there (the same way we might imagine there would be a universal set, or that $\sqrt{2}$ would be a rational number).[5]

However, we needn't worry about running into trouble with the Axiom of Choice. In 1938, Gödel proved that it is consistent with the other axioms of set theory: if one can derive a contradiction using choice, then the remaining axioms are already contradictory even without choice. Thus, the axiom is harmless, but mathematicians still wondered whether it was redundant. This was settled in 1963, when Cohen proved that in fact the Axiom of Choice is independent of the other axioms (i.e., its negation is also consistent with them).

It's worth briefly reviewing which sort of choices require the Axiom of Choice and which do not. (We'll formalize this later when we study first-order logic.) If you know something exists, you are always allowed to choose an example and assign it a variable name. For example, if you know a set S is non-empty, you can say "let s be an element of S." You can do this even if there are multiple elements of S. Similarly, if you know that $S \cong T$, then you are allowed to choose an isomorphism $f\colon S \to T$, because the hypothesis $S \cong T$ ensures that such an isomorphism exists.

On the other hand, you can't say "let f be a choice function" without some reason why choice functions exist in your set theory. The fact that you can imagine making infinitely many choices does not imply that set theory necessarily contains a function that makes those choices.

Constructing a choice function is only problematic for infinite sets, since ZF alone suffices to prove the following proposition:

Proposition 6.2. *Every finite set of non-empty sets has a choice function.*

Proof. Let

$$T = \{n \in \omega : \text{every set of } n \text{ non-empty sets has a choice function}\}.$$

We will show that T is inductive, from which it follows that $T = \omega$.

First, note that $0 \in T$ because the unique set of size 0 is the empty set, which has the empty choice function. Now suppose $n \in T$. Given a set S of size n^+, let $g\colon n^+ \to S$ be a bijection. Then $g|_n$ is a bijection from n to $S \setminus \{g(n)\}$, and so $S \setminus \{g(n)\}$ has a choice function f because $n \in T$. The set $g(n)$ is non-empty and so contains some element x. Then $f \cup \{(g(n), x)\}$ is a choice function on S, as desired. □

For the rest of this section, we will not assume the Axiom of Choice, except as indicated in the theorem statements. Instead, we will focus on describing its consequences and equivalent forms, assuming only the ZF axioms. Specifically, we will prove the equivalence of the following assertions:

(1) The Axiom of Choice.
(2) The well-ordering theorem: every set can be well-ordered.
(3) Trichotomy: for all sets S and T, there exists an injective function from S to T or one from T to S.
(4) Zorn's lemma (stated below).

[5] For comparison, consider set theory without the Axiom of Infinity. We can form the empty set 0, the set $1 = 0^+$, the set $2 = 1^+$, etc. Each of these sets is guaranteed to exist by the other axioms. However, that doesn't ensure that there is a single set ω that contains all of them. Similarly, each individual choice in a choice function could certainly be made, but we need the Axiom of Choice to guarantee that they can be combined to form a choice function.

It may seem odd that trichotomy involves only two possibilities, and perhaps it should be called dichotomy. The motivation for the name is that by the Cantor-Schröder-Bernstein theorem, it implies that $|S| = |T|$, $|S| < |T|$, or $|T| < |S|$ (although we will not define this notation until Section 7).

Proposition 6.3. *If the well-ordering theorem holds, then so does trichotomy.*

Proof. If we well-order S and T, then by Theorem 5.10, we have $S \cong T$, or $S_x \cong T$ for some $x \in S$, or $S \cong T_y$ for some $y \in T$. In the first case we get a bijection between S and T, in the second an injection from T to S, and in the third an injection from S to T. Thus, in each case there is an injection in one direction or the other. □

This proof may seem outrageous, since it imposes profound structure to deduce a simple, intuitive conclusion. However, that structure is in fact equivalent to the conclusion:

Proposition 6.4. *If trichotomy holds, then every set can be well-ordered.*

Proof. Given a set S, by Proposition 5.15 there is an ordinal α from which there is no injective map to S. By trichotomy, there must be an injective map from S to α, and it is a bijection from S to a subset of α. Every subset of a well-ordered set is well-ordered, and the bijection transfers this structure to S. □

Thus, the well-ordering theorem and trichotomy are equivalent.

Proposition 6.5. *If the well-ordering theorem holds, then so does the Axiom of Choice.*

Proof. Given a set S of non-empty sets, well-order

$$\bigcup_{x \in S} x$$

via an ordering we will denote by $\leq$. Now we can define a choice function $f\colon S \to \bigcup_{x\in S} x$ by letting $f(x)$ be the least element of x. More formally, let

$$f = \left\{ (x,y) \in S \times \bigcup_{x \in S} x : y \text{ is the least element of } x \text{ under } \leq \right\}.$$

This set exists by separation and defines a choice function. □

Intuitively, a well-ordering lets you make infinitely many choices in a systematic way, by always taking the minimal option. By contrast, we can't avoid choice by saying "y is some element of x" in place of "y is the least element of x under $\leq$," since then every $y \in x$ would work and f would not be single valued. To make f by separation, we need a mathematical statement that singles out a specific element in each element x of S.

To complete the proof that the Axiom of Choice, the well-ordering theorem, and trichotomy are equivalent, all we need to do is to prove that the Axiom of Choice implies the well-ordering theorem. We will prove this theorem twice, once using ordinals and once by a more elementary approach.

Theorem 6.6. *The Axiom of Choice implies that every set can be well-ordered.*

The intuition behind both proofs is simple. Given a set S, we will pick a choice function f for the non-empty subsets of S. Then $x_0 = f(S)$ will be the least element of S, $x_1 = f(S \setminus \{x_0\})$ will be the second-least element, $x_2 = f(S \setminus \{x_0, x_1\})$ will be the third-least, etc. We will continue until we fill the entire set. However, that is a little tricky: $\{x_i : i \in \omega\}$ may not fill S, in which case we must continue with $x_\omega = f(S \setminus \{x_i : i \in \omega\})$, etc. It is not obvious that this process can actually be completed, so there is genuinely something to prove.

First proof. Let S be a set, and let f be a choice function for $\mathcal{P}(S) \setminus \{\emptyset\}$. For each ordinal α, by transfinite recursion there is a unique function $g\colon \alpha \to S$ satisfying the recurrence

$$g(\beta) = \begin{cases} f(S \setminus \{g(\gamma) : \gamma < \beta\}) & \text{if } \{g(\gamma) : \gamma < \beta\} \neq S, \text{ and} \\ f(S) & \text{otherwise.} \end{cases}$$

The first line of this definition is the important part, while the value $f(S)$ in the second line is just an arbitrary choice.

Note that we can write $\{g(\gamma) : \gamma < \beta\}$ as $g[\beta]$. Suppose $g[\beta] = S$ for some $\beta \in \alpha$, and let β denote the first point at which that occurs. Then $g|_\beta$ is injective, because for $\gamma_1 < \gamma_2 < \beta$, we have

$$g(\gamma_2) = f(S \setminus g[\gamma_2]) \in S \setminus g[\gamma_2]$$

and $g(\gamma_1) \in g[\gamma_2]$, so $g(\gamma_1) \neq g(\gamma_2)$. Thus, in this case $g|_\beta$ is a bijection between the ordinal β and S, so S can be well-ordered.

If no such β exists, then the same argument shows that g is an injective function from α to S. However, Proposition 5.15 shows that there are ordinals α with no injective maps to S. For such an α, there must exist $\beta \in \alpha$ for which $g|_\beta$ is a bijection between the ordinal β and S, so we see that S can indeed be well-ordered. □

One can also prove the well-ordering theorem without even mentioning ordinals:

Second proof. Let S be a set, and let f be a choice function for $\mathcal{P}(S) \setminus \{\emptyset\}$. Define an *$f$-ordered subset* of S to be a subset A with a well-ordering $<_A$ on A, such that for all $x \in A$,

$$x = f(S \setminus A_x).$$

(Recall that A_x is the initial segment $\{y \in A : y <_A x\}$.)

First, we will show that every isomorphism $g\colon A \to B$ between f-ordered subsets is the identity function. If not, let x be the least element of A for which $g(x) \neq x$. Then

$$\begin{aligned} x &= f(S \setminus \{y \in A : y <_A x\}) \\ &= f(S \setminus \{g(y) : y \in A \text{ and } y <_A x\}) \\ &= f(S \setminus \{g(y) : y \in A \text{ and } g(y) <_B g(x)\}) \\ &= f(S \setminus \{z \in B : z <_B g(x)\}) \\ &= g(x), \end{aligned}$$

which contradicts $g(x) \neq x$. It now follows from Theorem 5.10 that for every pair of f-ordered subsets, either they are equal or one is an initial segment of the other. It follows that the orderings on two f-ordered subsets agree on their overlap.

We would like to define a maximal f-ordered subset T, i.e., one that contains all the others. Maximality implies that $T = S$, since otherwise we could define an

f-ordering on $T \cup \{f(S \setminus T)\}$ by making $f(S \setminus T)$ greater than every element of T. Thus, the existence of a maximal f-ordered subset implies that S can itself be well-ordered.

To construct T, we will take the union of all the f-ordered subsets of S. Define the ordering $<_T$ by $x <_T y$ if and only if $x <_A y$ for some f-ordered subset containing both x and y. Equivalently, $x <_A y$ for every f-ordered subset containing both x and y, because the orderings of the f-ordered subsets agree.

First, we must verify that T is totally ordered by $<_T$. If x and y are elements of T, then by the definition of T they are contained in f-ordered subsets A and B. Either those subsets are the same, or one is an initial segment of the other, and either way both x and y are contained in a single f-ordered subset (whichever of A and B is larger). Thus, they are comparable, and antisymmetry holds. Similarly, any three elements of T are all contained in a single f-ordered subset, so $<_T$ is transitive.

Next, we check that $<_T$ strictly well-orders T. If U is a non-empty subset of T, let u be an element of U. Then u is in some f-ordered subset A, and all smaller elements of T must be in A as well (since every other f-ordered subset is either an initial segment of A or has A as one of its initial segments). Thus, $\{t \in T : t \leq_T u\}$ is contained in A, and it has a least element in A. However, the orderings $<_T$ and $<_A$ agree on the elements of A, so there is a least element in T as well.

Finally, we verify that T is f-ordered. If $x \in T$, then $x \in A$ for some f-ordered A, and then

$$\begin{aligned} x &= f(S \setminus \{y \in A : y <_A x\}) \\ &= f(S \setminus \{y \in T : y <_T x\}), \end{aligned}$$

as desired.

Thus, there exists a maximal f-ordered subset of S, which must be S itself, and therefore S can be well-ordered. □

These proofs show how to turn a choice function for $\mathcal{P}(S) \setminus \{\emptyset\}$ into a well-ordering of S. It's far from obvious how to well-order $\mathbb{R}$, and it's also not obvious how to describe a choice function on $\mathcal{P}(\mathbb{R}) \setminus \{\emptyset\}$. (This would amount to specifying a single real number in each non-empty set of real numbers, and how would you give any sort of rule for doing that?) However, the Axiom of Choice tells us that such a choice function exists. It's just not something we can easily describe.

So far, we have shown that the Axiom of Choice, the well-ordering theorem, and trichotomy are equivalent. The final form will be *Zorn's lemma*: if every chain in a poset has an upper bound, then the poset has a maximal element. Recall from Section 3 that a chain is a totally ordered subset, an upper bound for a chain is an element of the poset that is greater than or equal to everything in the chain, and a maximal element of the poset is an element such that nothing is strictly greater than it.

Unlike trichotomy or the well-ordering theorem, Zorn's lemma is somewhat technical, and it is not obvious why one should care. However, constructing maximal elements of posets turns out to be very important. For example, Zorn's lemma implies that every vector space has a basis, as we prove below. Those unfamiliar with linear algebra can skip this example.

Proposition 6.7. *Zorn's lemma implies that every vector space has a basis.*

We will work over $\mathbb{R}$, although one could replace $\mathbb{R}$ with any field. Recall that a subset S of a vector space V is *linearly independent* if there do not exist vector $v_1, \dots, v_n \in S$ and coefficients $c_1, \dots, c_n \in \mathbb{R}$ with $c_1 v_1 + \dots + c_n v_n = 0$, unless $c_1 = \dots = c_n = 0$. A subset S *spans* V if for every $v \in V$, there exist $v_1, \dots, v_n \in S$ and $c_1, \dots, c_n \in \mathbb{R}$ such that $v = c_1 v_1 + \dots + c_n v_n$. A *basis* is a linearly independent subset that spans V.

Strictly speaking, we are sweeping one minor issue under the rug. What's a basis of the zero-dimensional vector space $\{0\}$? The only reasonable answer is the empty set, since a zero-dimensional vector space should have a zero-element basis, but how does the empty set span the vector 0? The zero vector is the sum of no elements (i.e., we take $n = 0$ above). More generally, for any associative operation with an identity element, applying that operation to no elements should result in the identity element.

Proof. Let V be a vector space, and let P be the set of all linearly independent subsets of V. Then P is a poset under $\subseteq$. Given any chain C in P, let

$$S = \bigcup_{T \in C} T.$$

Then S is linearly independent: if $v_1, \dots, v_n \in S$, then there exist $T_1, \dots, T_n \in C$ such that $v_i \in T_i$. Because C is a chain, $T_1, \dots, T_n$ are comparable and one of them therefore contains the others as subsets; call it T_j. Then $v_1, \dots, v_n \in T_j$, and because T_j is linearly independent, $c_1 v_1 + \dots + c_n v_n$ cannot vanish unless $c_1 = \dots = c_n = 0$. Thus, S is linearly independent, so $S \in P$. It is an upper bound for the chain C, so we have verified that every chain in P has an upper bound.

By Zorn's lemma, there is a maximal element S in P. That means for every $v \in V \setminus S$, the union $S \cup \{v\}$ is not linearly independent, so there exist $v_1, \dots, v_n \in S$ and $c_0, \dots, c_n \in \mathbb{R}$ such that

$$c_0 v + c_1 v_1 + \dots c_n v_n = 0,$$

with not all of $c_0, \dots, c_n$ equal to 0. Because S is linearly independent, we must have $c_0 \neq 0$, and hence

$$v = -\frac{c_1}{c_0} v_1 - \dots - \frac{c_n}{c_0} v_n.$$

Thus, S spans V, so it is a basis of V. □

We will deduce Zorn's lemma from a further lemma called *Hausdorff's maximal principle*: every poset has a maximal chain (i.e., a chain that is not a subset of any bigger chain).

Proposition 6.8. *Hausdorff's maximal principle implies Zorn's lemma.*

Proof. Let P be a poset in which every chain has an upper bound. By Hausdorff's maximal principle, there is a maximal chain C in P. Let u be an upper bound for C. Then u is maximal in P, because if $v > u$, then $C \cup \{v\}$ is an even larger chain. (Note that u must be in C, because otherwise $C \cup \{u\}$ would already be a larger chain than C, but v cannot be in C because it is strictly greater than each element of C.) □

Proposition 6.9. *The well-ordering theorem implies Hausdorff's maximal principle.*

Proof. Let P be a poset with ordering $\leq$. By the well-ordering theorem, there exists a well-ordering $<_W$ on P. Note that it needn't have any particular relationship to the original ordering $\leq$.

We will build a maximal chain by examining the elements of P successively and deciding whether to include each one based on whether it is comparable with all the previous elements of the chain. To formalize this process, we will use transfinite recursion.

Let p be the $<_W$-least element of P, and define by transfinite recursion a function $f\colon P \to P$ by

$$f(x) = \begin{cases} x & \text{if } \{x\} \cup \{f(y) : y <_W x\} \text{ is a chain in } P\text{, and} \\ p & \text{otherwise.} \end{cases}$$

Now consider the image $f[P]$ of f. We will show that it is a chain in P. To start, it always includes p, because $f(p) = p$. Then each element x is included in the image only if it is comparable with everything added before it (otherwise $f(x) = p$, which does not change the image since it already contains p). This means all elements of the image are comparable: if x_1 and x_2 are in $f[P]$ with $x_1 <_W x_2$, then the equation $f(x_2) = x_2$ tells us that x_2 must be comparable with $f(x_1)$, which is x_1.

Furthermore, anything not in the image was left out precisely because it was incomparable with something already in the image, so no further elements can be added to the chain. Thus, the image of f is a maximal chain in P. □

Finally, we will complete the circle by showing that Zorn's lemma implies the Axiom of Choice.

Proposition 6.10. *Zorn's lemma implies the Axiom of Choice.*

Proof. Let S be a set of non-empty sets, and let

$$P = \{(T, f) : T \subseteq S \text{ and } f \text{ is a choice function on } T\}.$$

Then P is a poset under the ordering $\leq$ defined by $(T_1, f_1) \leq (T_2, f_2)$ iff $T_1 \subseteq T_2$ and the restriction of f_2 to T_1 is f_1. Think of the elements of P as partial choice functions, defined only on subsets of S.

If there is a maximal element (T, f) in P, then $T = S$, because otherwise we could take $s \in S \setminus T$ and extend f to $T \cup \{s\}$ by letting $f(s)$ be any element of s. More formally, if $T \neq S$ and we take $s \in S \setminus T$ and $s_0 \in s$, then $(T \cup \{s\}, f \cup \{(s, s_0)\})$ is an element of P that is strictly greater than (T, f). Thus, by Zorn's lemma, all we need to check is that every chain in P has an upper bound.

Given a chain C in P, let T be the union of the domains of the partial choice functions in C. Because the elements of C agree wherever more than one is defined (this follows from comparability in P), we can define a partial choice function f on T by letting $f(t)$ be the unique element of t chosen by the elements of C whose domains include t. Then (T, f) is an upper bound in P for C. Thus, there is a maximal element of P, and it is a choice function for S, as desired. □

7. Cardinals

From this point on, we assume the Axiom of Choice. Thus, by the well-ordering theorem, every set is in bijection with some ordinal.

Definition 7.1. An ordinal is a *cardinal* if it is not in bijection with any smaller ordinal.

Note that we impose no restriction on being in bijection with larger ordinals. (Indeed, every infinite ordinal α is in bijection with α^+.)

By Proposition 4.11, every natural number is a cardinal. Furthermore, ω is a cardinal as well:

Lemma 7.2. *There is no bijection between ω and any natural number, and thus ω is a cardinal.*

Proof. Suppose there were a bijection $f\colon \omega \to n$ with $n \in \omega$. Then there would be an injective map from ω to n^+, namely the composition of f with the inclusion $n \subseteq n^+$, and an injective map from n^+ to ω, namely the inclusion $n^+ \subseteq \omega$. Thus, the Cantor-Schröder-Bernstein theorem would give a bijection between ω and n^+. However, there is no bijection between n and n^+, by Proposition 4.11. □

On the other hand, ω^+ is not a cardinal, because it is in bijection with ω.

Lemma 7.3. *Every set is in bijection with a unique cardinal.*

Proof. Let S be a set. By the well-ordering theorem, we can well-order S, and then Theorem 5.11 says it is isomorphic to α for some ordinal α. Since every set of ordinals is well-ordered, let κ be the least ordinal that is at most α and in bijection with S. (Note that the "at most α" is to guarantee that we are taking the least element of an actual set.) Then κ is a cardinal because, by construction, no smaller ordinal is in bijection with κ (equivalently, S). For uniqueness, note that no two different cardinals can be in bijection, since one of them would have to be smaller than the other. □

Definition 7.4. The *cardinality* $|S|$ of a set S is the unique cardinal (i.e., the smallest ordinal) that is in bijection with S.

Every cardinal is the cardinality of some set (for example, itself), and the cardinality of a set is always a cardinal. Of course, $|\alpha| \le \alpha$ for every ordinal α.

There are two natural ways to compare cardinals: using the ordering on ordinals and using the existence of injective functions between them. The next lemma shows that these orderings are the same.

Lemma 7.5. *Let κ and λ be cardinals. Then $\kappa \le \lambda$ as ordinals if and only if there is an injective map from κ to λ.*

Proof. If $\kappa \le \lambda$, then $\kappa \subseteq \lambda$ and the inclusion of κ in λ is injective. If $\kappa > \lambda$, then there cannot be an injective map from κ to λ, since otherwise combining it with the inclusion from λ to κ would yield a bijection between κ and λ by the Cantor-Schröder-Bernstein theorem, which would contradict the fact that κ is a cardinal. □

This proof used the Cantor-Schröder-Bernstein theorem, but we could also prove it without that theorem:

Proof. If $\kappa \le \lambda$, then $\kappa \subseteq \lambda$ and the inclusion of κ in λ is injective. Conversely, if there is an injective map from κ to λ, then κ is isomorphic to a subset of λ as a well-ordered set and Proposition 5.14 implies that $\kappa \le \lambda$. □

The advantage of this second proof is that the Cantor-Schröder-Bernstein theorem is an immediate corollary: if there are injections from S to T and vice versa, then

$|S| \leq |T|$ and $|T| \leq |S|$, from which it follows that $|S| = |T|$. On the other hand, our previous proof required far less machinery and did not use the Axiom of Choice.

Corollary 7.6. *For all sets S and T, $|S| \leq |T|$ iff there is an injective function from S to T, and $|S| = |T|$ iff there is a bijection between S and T.*

Corollary 7.7. *If α and β are ordinals and $\alpha \leq \beta$, then $|\alpha| \leq |\beta|$.*

Proof. If $\alpha \leq \beta$, then $\alpha \subseteq \beta$. □

On the other hand, $\alpha < \beta$ does not imply $|\alpha| < |\beta|$ (take $\alpha = \omega$ and $\beta = \omega^+$).

In terms of cardinality, the trichotomy theorem says that for all sets S and T, $|S| \leq |T|$ or $|T| \leq |S|$. (Note that we do not need the trichotomy theorem to reach this conclusion, because cardinals are totally ordered.)

The ordering on cardinals can also be defined using surjective maps, but one must be careful about the empty set: for $\kappa \neq 0$, the following lemma shows that $\kappa \leq \lambda$ if and only if there is a surjective map from λ to κ, but there is no surjective map from a non-empty set to the empty set so the assumption that $\kappa \neq 0$ is needed.

Lemma 7.8. *Let S and T be sets, with $S \neq \emptyset$. Then there is an injective map from S to T if and only if there is a surjective map from T to S.*

Proof. Given an injection $f\colon S \to T$, we can define a surjection from T to S by mapping the elements of $f[S]$ to their pre-images and the rest of T to some arbitrary element of S (which is why S must be non-empty). However, the other direction depends on the Axiom of Choice: given a surjection $g\colon T \to S$, we can define an injection $f\colon S \to T$ by choosing an element $f(s)$ from the pre-image of s under g. More formally, let $g^{-1}(s) = \{t \in T : g(t) = s\}$, which is non-empty for each $s \in S$ because g is surjective, and let $\widetilde{f}$ be a choice function for $\{x \in \mathcal{P}(T) : x = g^{-1}(s) \text{ for some } s \in S\}$. Then the function $f\colon S \to T$ defined by $f(s) = \widetilde{f}(g^{-1}(s))$ is injective. To see why, note that $\widetilde{f}(g^{-1}(s)) \in g^{-1}(s)$, which implies $g(f(s)) = s$; thus if $f(s_1) = f(s_2)$, then $g(f(s_1)) = g(f(s_2))$ and hence $s_1 = s_2$. □

It is not know whether Lemma 7.8 (sometimes called the Partition Principle) is equivalent to the Axiom of Choice.

We define $\aleph_0$ ("aleph naught") to be ω, thought of as a cardinal. Of course, there is no logical need to have two different names for the same set, but it is sometimes convenient because it makes it clear whether we are thinking about ordinals in general or cardinals specifically. The letter $\aleph$ is the Hebrew letter aleph.

Note that $\aleph_0$ is the smallest infinite cardinal (since all its elements are finite by definition), which means S is infinite iff $|S| \geq \aleph_0$. Expressed differently, a set S is infinite iff there is an injection from ω to S. As a side comment, it's not hard to turn this into another characterization of finite sets:

Proposition 7.9. *A set is infinite if and only if it has a bijection to one of its proper subsets.*

Proof. Suppose S is infinite, and let $f\colon \omega \to S$ be injective. Then there is a bijection from S to $S \setminus \{f(0)\}$, for example by mapping $f(i)$ to $f(i^+)$ and fixing $S \setminus f[\omega]$.

For the converse, we just need to check that for each natural number n, there is no bijection from n to one of its proper subsets. Suppose $f\colon n \to S$ is a bijection with $S \subsetneq n$. Since n is not the empty set, we must have $n = m^+$ for some $m \in \omega$, and thus $S \subsetneq m \cup \{m\}$. Without loss of generality we can assume that $S \subseteq m$, because otherwise $m \in S$ but some $i \in m$ must be missing from S (since S is a proper subset of $m \cup \{m\}$), and swapping i with m transforms f into a bijection with a subset of m. Now we have an injective map from n to m with $m < n$. Since the inclusion map from m to n is also injective, this means there is a bijection between n and m, by the Cantor-Schröder-Bernstein theorem, and that contradicts Proposition 4.11. □

Theorem 7.10. *Every set of cardinals is well-ordered.*

Of course, this theorem is an immediate corollary of Corollary 4.17, because all cardinals are ordinals. However, it is an important enough result to be worth stating as a theorem.

It follows from Proposition 5.15 that there is no greatest cardinal. However, there is an even simpler way to see this:

Theorem 7.11 (Cantor)**.** *For every set S,*

$$|S| < |\mathcal{P}(S)|.$$

It is clear that $|S| \le |\mathcal{P}(S)|$, for example because of the injective function mapping $s \in S$ to $\{s\} \in \mathcal{P}(S)$, but the strict inequality is not as obvious. Cantor's theorem is one of the most important theorems of set theory, and indeed of mathematics in general. The proof is a beautiful example of diagonalization.

Proof. We will show that there cannot be a surjective function from S to $\mathcal{P}(S)$. Given a function $f \colon S \to \mathcal{P}(S)$, let

$$T = \{s \in S : s \notin f(s)\}.$$

Then $T \in \mathcal{P}(S)$, and we will show that it cannot be in the image of f. If $T = f(t)$ with $t \in S$, then $t \in T$ if and only if $t \notin f(t)$ by the definition of T, and that means $t \in T$ if and only if $t \notin T$, which is a contradiction. Thus, f is not surjective. □

The motivation behind this diagonal argument is that for each $s \in S$, we want to rule out $T = f(s)$. We do so by arranging for them to differ in whether they contain s. That means we want $s \in T$ iff $s \notin f(s)$, which is the definition of T.

Proposition 7.12. *There is no set of all cardinals.*

Proof. For every ordinal α, there is a larger cardinal, for example $|\mathcal{P}(\alpha)|$, and $\alpha \in |\mathcal{P}(\alpha)|$. (To see why, note that if $|\mathcal{P}(\alpha)| \le \alpha$ instead, then $|\mathcal{P}(\alpha)| \le |\alpha|$, which would violate Cantor's theorem.) Thus, if the set S of all cardinals existed, then

$$\bigcup_{\kappa \in S} \kappa$$

would be the set of all ordinals, and we would run into the Burali-Forti paradox. □

Because infinite cardinals are well-ordered, we can index them by ordinals: $\aleph_1$ is the least cardinal greater than $\aleph_0$, $\aleph_2$ is the least cardinal greater than $\aleph_1$, etc. In fact, for every ordinal α there will be an infinite cardinal $\aleph_\alpha$, and every infinite cardinal will be of this form, but that requires proof.

To carry out this construction, we will measure the size of an infinite cardinal by looking at the set of infinite cardinals that come before it. Specifically, given any infinite cardinal κ, consider the set

$$I_\kappa = \{\lambda : \lambda \text{ is an infinite cardinal and } \lambda < \kappa\}.$$

(Note that this set exists by separation, because every such λ is in fact an element of κ.) This is a well-ordered set and is therefore isomorphic to a unique ordinal. Furthermore, if $\kappa < \lambda$ then I_κ is an initial segment of I_λ, so the corresponding ordinals are unequal, and every initial segment of I_κ corresponds to some infinite cardinal.

Lemma 7.13. *For every ordinal α, there is an infinite cardinal κ for which I_κ is isomorphic to α.*

Proof. If not, let α be the least ordinal not isomorphic to I_κ for any infinite cardinal κ. If α is a successor ordinal, with $\alpha = \beta^+$, then let β be isomorphic to I_λ. If we let κ be the least cardinal greater than λ, then I_κ is isomorphic to α, which is a contradiction. On the other hand, if α is a limit ordinal, then

$$\alpha = \bigcup_{\beta<\alpha} \beta.$$

Taking the union of the sets I_κ corresponding to the ordinals $\beta < \alpha$ gives a downwards closed set S of infinite cardinals that is isomorphic to $\bigcup_{\beta<\alpha} \beta$ and hence to α. It cannot contain all cardinals, by Proposition 7.12. If λ is the smallest infinite cardinal not contained in S, then $S = I_\lambda$, as desired. □

Definition 7.14. If α is an ordinal, then $\aleph_\alpha$ is the unique infinite cardinal such that

$$\{\kappa : \kappa \text{ is an infinite cardinal and } \kappa < \aleph_\alpha\}$$

is isomorphic to α as a well-ordered set. We say $\aleph_\alpha$ is a *successor cardinal* if α is a successor ordinal, and a *limit cardinal* otherwise.

Note that every infinite cardinal is always a limit ordinal (an infinite successor ordinal α^+ is in bijection with α and thus not a cardinal). Thus, limit ordinals and limit cardinals are not at all the same thing.

This definition indexes the infinite cardinals by the ordinals. People sometimes write ω_α for $\aleph_\alpha$ when they wish to think of it more as an ordinal than as a cardinal, but of course there's no mathematical distinction here, just a psychological distinction.

We know that $|\mathcal{P}(\aleph_0)| > \aleph_0$, and it is natural to ask how large it is. Cantor conjectured that $|\mathcal{P}(\aleph_0)| = \aleph_1$, and this conjecture is known as the *continuum hypothesis*, because $|\mathcal{P}(\aleph_0)|$ is the cardinality of the continuum $\mathbb{R}$. The *generalized continuum hypothesis* is that $|\mathcal{P}(\aleph_\alpha)| = \aleph_{\alpha^+}$ for all ordinals α. In other words, it says that the power set is as small as possible, subject to Cantor's theorem.

Gödel showed in 1938 that the generalized continuum hypothesis is consistent with ZFC (of course assuming ZFC is itself consistent), and Cohen showed in 1963 that the negation of the continuum hypothesis is also consistent with ZFC. Thus, the very natural question of how large $|\mathcal{P}(\aleph_0)|$ is cannot be settled using the current axioms of set theory. Of course, we could introduce new axioms to settle it; for example, we could even assume the generalized continuum hypothesis itself as an axiom. However, nobody has been able to propose an intuitively compelling axiom that resolves the continuum hypothesis.

Incidentally, the word "continuum" refers to the real numbers. One can show that $|\mathbb{R}| = |\mathcal{P}(\aleph_0)|$, and thus the continuum hypothesis asks whether $|\mathbb{R}| = \aleph_1$. It's remarkable that the answer is not known and may not be knowable.

8. Cardinal arithmetic

We can define arithmetic on cardinals by imitating the ideas behind elementary-school arithmetic. Recall that $\mathcal{F}(S,T)$ denotes the set of functions from S to T.

Definition 8.1. Let κ and λ be cardinals. Then we define

$$\kappa+\lambda = |(\kappa\times\{0\})\cup(\lambda\times\{1\})|,$$
$$\kappa\lambda = |\kappa\times\lambda|,$$

and

$$\kappa^\lambda = |\mathcal{F}(\lambda,\kappa)|.$$

The definition of multiplication is exactly what one would expect, but addition and exponentiation may look a little odd. For addition, the purpose of using $\kappa\times\{0\}$ and $\lambda\times\{1\}$ instead of κ and λ is to ensure that the sets are disjoint. (Just using κ and λ does not work, since the smaller one is a subset of the other.) For exponentiation, one might naively expect to use functions going in the other direction, but that would give the wrong answer if κ and λ are finite: a function from λ to κ is given by choosing one of κ values at each of λ points, and there are κ^λ ways to do that.

Given these definitions, we can easily apply the operations to the cardinalities of arbitrary sets:

Lemma 8.2. *For arbitrary sets S and T, we have $|S\cup T|+|S\cap T| = |S|+|T|$, $|S\times T| = |S||T|$, and*

$$|\mathcal{F}(T,S)| = |S|^{|T|}.$$

Note that we cannot write the first equation in the lemma as $|S\cup T| = |S| + |T| - |S\cap T|$, because there is no operation of subtraction on cardinals. For example, $\aleph_0+0=\aleph_0+1$, but we cannot subtract $\aleph_0$ to conclude that $0=1$. Similarly, there is no division, because $1\aleph_0 = 2\aleph_0$.

Proof. The equations for multiplication and exponentiation are simply based on the fact that every set is in bijection with a cardinal. Addition is a little more subtle because of the disjointness issue, but it is still straightforward. We can form a bijection f from $(S\times\{0\})\cup(T\times\{1\})$ to $((S\cup T)\times\{0\})\cup((S\cap T)\times\{1\})$ by setting $f(s,0)=(s,0)$ for $s\in S$, $f(t,1)=(t,0)$ for $t\in T\setminus S$, and $f(t,1)=(t,1)$ for $t\in T\cap S$. □

Lemma 8.3. *Cardinal arithmetic satisfies the following identities for all cardinals κ, λ, and μ:*

1. $\kappa+\lambda=\lambda+\kappa$ *and* $\kappa\lambda=\lambda\kappa$*,*
2. $(\kappa+\lambda)+\mu=\kappa+(\lambda+\mu)$ *and* $(\kappa\lambda)\mu=\kappa(\lambda\mu)$*,*
3. $\kappa+0=\kappa$ *and* $1\kappa=\kappa$*,*
4. $\kappa(\lambda+\mu)=\kappa\lambda+\kappa\mu$*,*
5. $\kappa^{\lambda+\mu}=\kappa^\lambda\kappa^\mu$*,*
6. $(\kappa\lambda)^\mu=\kappa^\mu\lambda^\mu$*, and*
7. $(\kappa^\lambda)^\mu=\kappa^{\lambda\mu}$*.*

Sketch of proof. Most of these identities are straightforward. For example, $\kappa\lambda=\lambda\kappa$ holds because switching the coordinates gives a bijection between $\kappa\times\lambda$ and $\lambda\times\kappa$. It is very much worth going over this list carefully and checking that each identity is indeed provable, but we will omit the details here.

The most subtle case is the last one, namely $(\kappa^\lambda)^\mu = \kappa^{\lambda\mu}$. To prove it, we must find a bijection between $\mathcal{F}(\mu, \mathcal{F}(\lambda, \kappa))$ and $\mathcal{F}(\lambda \times \mu, \kappa)$. The bijection is sometimes called *currying*. Given $f \colon \lambda \times \mu \to \kappa$, and given $y \in \mu$, let $g(y)$ be the function from λ to κ defined by $(g(y))(x) = f(x, y)$. Then g is a function from μ to functions from λ to κ; in other words, it is an element of $\mathcal{F}(\mu, \mathcal{F}(\lambda, \kappa))$. Conversely, given such a function g, we can define f by $f(x, y) = (g(y))(x)$, and this gives the desired bijection. □

It also follows easily from the definitions that $\kappa + \lambda$, $\kappa\lambda$, and κ^λ are (weakly) increasing in both κ and λ. On the other hand, they are not strictly increasing: $0 < 1$, but it is not the case that $0 + \aleph_0 < 1 + \aleph_0$.

Lemma 8.4. *If κ and λ are finite, then so are $\kappa + \lambda$, $\kappa\lambda$, and κ^λ.*

For $\kappa + \lambda$, this follows immediately from Proposition 2.6. The remaining two cases can be proved by a similar induction, with each case depending on the previous one.

Addition and multiplication turn out to be as simple as possible for infinite cardinals: they just give the maximum of their arguments. This means cardinal arithmetic lacks the richness of number theory, but on the positive side it means it is often quite easy to compute cardinalities of infinite sets, because taking unions or Cartesian products never complicates anything.

Theorem 8.5. *Let κ and λ be cardinals, with $\kappa \le \lambda$ and λ infinite. Then $\kappa + \lambda = \lambda$ and if $\kappa \ne 0$, then $\kappa\lambda = \lambda$.*

Proof. We will prove that $\lambda^2 = \lambda$. Then the desired results follow from

$$\lambda \le \kappa + \lambda \le \lambda + \lambda = 2\lambda \le \lambda^2 = \lambda$$

and

$$\lambda \le \kappa\lambda \le \lambda^2 = \lambda.$$

Suppose λ is the first infinite cardinal such that $\lambda^2 \ne \lambda$. To compute λ^2, we will start by well-ordering $\lambda \times \lambda$. Specifically, we set $(\alpha_1, \beta_1) < (\alpha_2, \beta_2)$ if

(1) $\max(\alpha_1, \beta_1) < \max(\alpha_2, \beta_2)$, or
(2) $\max(\alpha_1, \beta_1) = \max(\alpha_2, \beta_2)$, and $\alpha_1 < \alpha_2$, or
(3) $\max(\alpha_1, \beta_1) = \max(\alpha_2, \beta_2)$, $\alpha_1 = \alpha_2$, and $\beta_1 < \beta_2$.

In other words, we order first by the maximum of the two coordinates, and then refine that ordering lexicographically. It is straightforward to check that $\lambda \times \lambda$ is then well-ordered, using the fact that λ is itself well-ordered.

Now there exists a unique ordinal δ such that $\delta \cong \lambda \times \lambda$. We will prove that $\delta \le \lambda$, from which it follows that $|\lambda \times \lambda| = |\delta| \le |\lambda| = \lambda$ (note that $|\lambda| = \lambda$ since λ is a cardinal). On the other hand, $\lambda \le \lambda^2$ (multiply both sides of $1 \le \lambda$ by λ), and so $\lambda^2 \le \lambda$ is enough to conclude that $\lambda^2 = \lambda$.

Suppose, by way of contradiction, that $\delta > \lambda$. Then λ is an initial segment of δ, so this means there is an initial segment of $\lambda \times \lambda$ that has cardinality λ. To rule this out, we will compute the sizes of the initial segments of $\lambda \times \lambda$.

Consider the initial segment consisting of all the elements of $\lambda \times \lambda$ that are less than (α, β), where $\alpha, \beta \in \lambda$. Without loss of generality we can assume $\beta = \alpha$, because replacing the lesser of the two coordinates with the greater will only make the initial segment larger. Under our ordering on $\lambda \times \lambda$, the initial segment cut out by (α, α) consists of all pairs (β, γ) with $\beta, \gamma \le \alpha$ and $(\beta, \gamma) \ne (\alpha, \alpha)$. In other

9	10	11	15
4	5	8	14
1	3	7	13
0	2	6	12

FIGURE 8.1. A well-ordering of $\aleph_0 \times \aleph_0$.

words, it is the set $(\alpha^+ \times \alpha^+) \setminus \{(\alpha, \alpha)\}$. Thus, the initial segment has size at most $|\alpha^+|^2$.

Because λ is a cardinal and α is an ordinal with $\alpha < \lambda$ (i.e., $\alpha \in \lambda$), we have $|\alpha| < \lambda$. If α is finite, then so is $|\alpha^+|^2$. On the other hand, if α is infinite, then $|\alpha^+| = |\alpha|$ and $|\alpha|^2 = |\alpha|$ since we assumed λ was the smallest infinite cardinal such that $\lambda^2 \neq \lambda$. Either way, $|\alpha^+|^2 < \lambda$. Thus, every initial segment of $\lambda \times \lambda$ has cardinality strictly less than λ, and we conclude that $\delta \leq \lambda$ and $\lambda^2 = \lambda$. Thus, every infinite cardinal equals its own square, as desired. □

As a side comment about the proof, note that δ cannot be strictly less than λ, because λ is not in bijection with any previous ordinal. Thus, the well-ordering in the proof gives $\lambda \times \lambda$ the order type λ. See Figure 8.1 for a picture of how it well-orders $\aleph_0 \times \aleph_0$.

On the other hand, cardinal exponentiation is much more subtle than addition or multiplication. For every set S, we have

$$|\mathcal{P}(S)| = 2^{|S|},$$

because specifying a function from S to the set $2 = \{0, 1\}$ is the same as specifying the subset of S on which it takes the value 0. The fact that ZFC does not even determine whether $2^{\aleph_0} = \aleph_1$ means we will not be able to compute many exponentials explicitly. However, we can say something.

Lemma 8.6. *If κ, λ, and μ are cardinals satisfying $\lambda \leq \kappa \leq \lambda^\mu$ and μ is infinite, then $\kappa^\mu = \lambda^\mu$.*

Proof. Because $\lambda \leq \kappa \leq \lambda^\mu$, we have

$$\lambda^\mu \leq \kappa^\mu \leq (\lambda^\mu)^\mu.$$

Because μ is infinite,

$$(\lambda^\mu)^\mu = \lambda^{\mu^2} = \lambda^\mu.$$

Thus, $\kappa^\mu = \lambda^\mu$, as desired. □

For example, $\kappa^\kappa = 2^\kappa$ whenever κ is infinite. In other words, for every infinite set S, there are the same number of functions from S to S as there are subsets of S.

We can naturally extend cardinal arithmetic to infinite sums and products. For sums we simply take infinite unions; for products, we must define infinite Cartesian products.

Definition 8.7. Given a family of sets S_i indexed by $i \in I$, their *Cartesian product* is defined by

$$\bigtimes_{i\in I} S_i = \left\{ \text{functions } f \text{ from } I \text{ to } \bigcup_{i\in I} S_i \text{ such that } f(i) \in S_i \text{ for all } i \in I \right\}.$$

Strictly speaking, when $|I| = 2$ this definition conflicts with our earlier definition of the Cartesian product, but it does not matter because the new definition still satisfies Lemma 1.5. Of course, we could not have used the new definition earlier, because we had to define Cartesian products before we could define functions.

The intuition is that a k-tuple $(a_0, \dots, a_{k-1})$ is the same as a function f defined on the set k by $f(i) = a_i$. More generally, even if I is infinite we can still view functions defined on I as I-tuples.

We will often define *projection operators* $\pi_j \colon \bigtimes_{i\in I} S_i \to S_j$, which just return the value in the j-th coordinate. In terms of the definition we have given for the Cartesian product, $\pi_j(f) = f(j)$. Note that we do not really need the π_j notation, since $f(j)$ is even easier to write. However, obtaining the j-th coordinate by function evaluation looks a little weird, and it's often more natural to think of it as an operator we apply to I-tuples.

Definition 8.8. Given cardinals κ_i for $i \in I$, we define

$$\sum_{i\in I} \kappa_i = \left| \bigcup_{i\in I} \kappa_i \times \{i\} \right|$$

and

$$\prod_{i\in I} \kappa_i = \left| \bigtimes_{i\in I} \kappa_i \right|.$$

Note that the reason for the unconventional symbol $\bigtimes$ for the Cartesian product, instead of $\prod$, is that the latter would lead to the cryptic definition

$$\prod_{i\in I} \kappa_i = \left| \prod_{i\in I} \kappa_i \right|,$$

in which the two products would mean different things.

The Axiom of Choice is equivalent to the fact that an infinite product of nonempty sets is always nonempty: an element of $\bigtimes_{i\in I} S_i$ encodes the same information as a choice function for $\{S_i : i \in I\}$.

As in the case of finite sums and products, it is easy to justify various arithmetic laws. For example, if κ is a cardinal, then

$$\sum_{i\in I} \kappa = \kappa|I|$$

and

$$\prod_{i\in I} \kappa = \kappa^{|I|}.$$

A slightly less trivial identity is

$$\sum_{i\in\omega} \aleph_i = \aleph_\omega.$$

To see why it is true, note that

$$\sum_{i\in\omega} \aleph_i \geq \aleph_j$$

for all $j \in \omega$, and this implies

$$\sum_{i\in\omega} \aleph_i \geq \aleph_\omega.$$

On the other hand,

$$\sum_{i\in\omega} \aleph_i \leq \sum_{i\in\omega} \aleph_\omega = \aleph_\omega \aleph_0 = \aleph_\omega.$$

More generally, infinite sums are not so difficult to compute:

Lemma 8.9. *Suppose κ_i is a nonzero cardinal for each $i \in I$. If $|I| \geq \aleph_0$, then*

$$\sum_{i\in I} \kappa_i = |I| \sup_{i\in I} \kappa_i.$$

Here $\sup_{i\in I} \kappa_i$ denotes the least cardinal greater than or equal to each κ_i. Such a cardinal exists because there are upper bounds (such as $\sum_{i\in I} \kappa_i$) and every set of cardinals is well-ordered, so there must be a least upper bound.

Proof. Clearly,

$$\sum_{i\in I} \kappa_i \leq |I| \sup_{i\in I} \kappa_i.$$

On the other hand,

$$|I| \leq \sum_{i\in I} \kappa_i$$

because $\kappa_i \neq 0$ for all i, and

$$\sup_{i\in I} \kappa_i \leq \sum_{i\in I} \kappa_i$$

because the right side is an upper bound for the cardinals κ_i. Thus, since $|I|$ is infinite it follows from Theorem 8.5 that

$$|I| \sup_{i\in I} \kappa_i \leq \sum_{i\in I} \kappa_i,$$

as desired. □

Infinite products are quite a bit more subtle. To address one potential misconception, note that they cannot be obtained as "limits of partial products" or the like. For example,

$$\prod_{i\in\omega} 2 = 2^{\aleph_0} > \aleph_0 = \sup_{n\in\omega} \prod_{i\in n} 2,$$

so the partial products

$$\prod_{i\in n} 2$$

in no way converge to the infinite product.

One of the most important results about infinite sums and products is König's theorem:

Theorem 8.10 (König). *Let κ_i and λ_i be cardinals for $i \in I$, with $\kappa_i < \lambda_i$ for all i. Then*

$$\sum_{i \in I} \kappa_i < \prod_{i \in I} \lambda_i.$$

Part of why König's theorem is remarkable is the strict inequality in the conclusion. By contrast, $\kappa_i < \lambda_i$ immediately implies that

$$\sum_{i \in I} \kappa_i \le \sum_{i \in I} \lambda_i,$$

but strict inequality does not always hold. For example, if I is infinite, $\kappa_i = 1$, and $\lambda_i = 2$, then we get $|I|$ on both sides of the equation.

Proof. Let S_i and T_i be any sets with $|S_i| = \kappa_i$ and $|T_i| = \lambda_i$. We will prove that there is no surjective function from

$$\bigcup_{i \in I} S_i$$

to

$$\bigtimes_{i \in I} T_i,$$

which implies that

$$\left| \bigcup_{i \in I} S_i \right| < \left| \bigtimes_{i \in I} T_i \right|$$

because the Cartesian product is non-empty. König's theorem is simply the special case in which the sets S_i are disjoint, but of course allowing them to overlap does not really strengthen the result.

Let $f \colon \bigcup_{i \in I} S_i \to \bigtimes_{i \in I} T_i$ be any function. We want to construct an element x of $\bigtimes_{i \in I} T_i$ that is guaranteed not to be in the image of f. For each $i \in I$, we will choose the i-th coordinate of x to ensure that it is not in $f[S_i]$.

Specifically, let $\pi_i \colon \bigtimes_{j \in I} T_j \to T_i$ be the projection onto the i-th coordinate. Then

$$|\pi_i[f[S_i]]| \le |S_i| < |T_i|,$$

so there exists an element $x_i \in T_i \setminus \pi_i[f[S_i]]$. This means x_i cannot occur as the i-th coordinate of any element of $f[S_i]$.

Now define $x \in \bigtimes_{i \in I} T_i$ to have such an element as its i-th coordinate x_i for all i (we're using the Axiom of Choice here). Then for all i, x cannot be in $f[S_i]$, so it is not in the image of f at all. Thus, f is not surjective, as desired. □

Cantor's theorem is an immediate corollary of König's theorem: if we take $\kappa_i = 1$ and $\lambda_i = 2$ for all $i \in I$, then we find that

$$|I| = \sum_{i \in I} 1 < \prod_{i \in I} 2 = 2^{|I|}.$$

However, there are even deeper consequences. For example, König's theorem implies that $\aleph_\omega$ is not the cardinality of any power set:

Proposition 8.11. *For all cardinals κ,*

$$2^\kappa \ne \aleph_\omega.$$

Note that $\aleph_0$ has the same property for a simpler reason: $2^\kappa < \aleph_0$ if κ is finite and $2^\kappa \ge 2^{\aleph_0} > \aleph_0$ if κ is infinite.

Proof. By König's theorem,

$$\sum_{i\in\omega} \aleph_i < \prod_{i\in\omega} \aleph_\omega,$$

and thus

$$\aleph_\omega < \aleph_\omega^{\aleph_0}.$$

It suffices to prove the proposition when κ is infinite. Then the previous equation implies that

$$\aleph_\omega^\kappa > \aleph_\omega.$$

If $2^\kappa = \aleph_\omega$, then $\aleph_\omega^\kappa = 2^{\kappa^2} = 2^\kappa = \aleph_\omega$, which is a contradiction. □

As a special case, this remarkable argument tells us $2^{\aleph_0} \neq \aleph_\omega$, but it does not specify which is larger, and in fact both possibilities are consistent with ZFC.

Thus, there is some property of $\aleph_\omega$ other than its size that is incompatible with being a power of 2. The existence of such properties is not so surprising in principle; for example, it is obvious that

$$2^{128} \neq 340282366920938463453374607431768211457$$

because the right side is odd, although it is not obvious which side is larger. However, the relevant property of $\aleph_\omega$ is more subtle than being odd. In the proof of Proposition 8.11, all we needed was that $\aleph_\omega^{\aleph_0} > \aleph_\omega$. This inequality fits naturally into the theory of cofinality, which we will examine next.

9. Cofinality

Definition 9.1. A subset S of a partially ordered set P is *cofinal* if for all $x \in P$, there exists an $s \in S$ such that $x \leq s$.

In other words, no matter how long P continues, there is always an element of S coming up. The name "cofinal" captures this idea by saying that S is "equally final" compared with P.

We will be particularly interested in measuring the sizes of cofinal subsets of ordinals, by their order types or cardinalities (we will see in Corollary 9.5 that these two approaches are equivalent). Recall that the order type of a well-ordered set is the unique ordinal isomorphic to it.

Definition 9.2. The *cofinality* $\mathrm{cf}(\alpha)$ of an ordinal α is the smallest possible order type of a cofinal subset of α.

Note that every subset of α is well ordered and hence has an order type, and we can compare these order types because they are ordinals. Because ordinals are well-ordered, there is a smallest possible order type of a cofinal subset, and hence cofinality is well-defined. We'll see several equivalent definitions of cofinality, but this one is the strongest, which makes it a good starting point.

Clearly, $\mathrm{cf}(\alpha) \leq \alpha$, because α is a cofinal subset of itself, but it can be much smaller. For example, $\mathrm{cf}(\alpha^+) = 1$, because the subset $\{\alpha\}$ is cofinal in α^+. Cofinality also satisfies $\mathrm{cf}(\mathrm{cf}(\alpha)) = \mathrm{cf}(\alpha)$, because if $\mathrm{cf}(\mathrm{cf}(\alpha)) < \mathrm{cf}(\alpha)$, then one could replace a cofinal subset S of α that has order type $\mathrm{cf}(\alpha)$ with a cofinal subset of S that has order type $\mathrm{cf}(\mathrm{cf}(\alpha))$ to lower the cofinality of α (a cofinal subset of a cofinal subset is itself cofinal).

Lemma 9.3. *Let α and β be ordinals. If there is a function from α to β with cofinal image, then* $\mathrm{cf}(\beta) \leq \alpha$.

We can in fact go further: there is always a function from $\mathrm{cf}(\alpha)$ to α with cofinal image, and composing it with the function in the lemma statement gives one from $\mathrm{cf}(\alpha)$ to β. Thus, we arrive at the stronger corollary that $\mathrm{cf}(\beta) \le \mathrm{cf}(\alpha)$.

Proof. Given a function $f\colon \alpha \to \beta$ with cofinal image, define

$$S = \{\gamma \in \alpha : f(\gamma) > f(\delta) \text{ for all } \delta < \gamma\}.$$

In other words, S is the set of places where f sets a new size record compared with everything that came before. Then $S \subseteq \alpha$ and $f|_S$ is an isomorphism from S to $f[S]$, because S is chosen so that $f|_S$ is strictly order-preserving.

Furthermore, $f[S]$ is cofinal in $f[\alpha]$, because for each element $f(\delta)$ in $f[\alpha]$, the first $\gamma \in \alpha$ such that $f(\gamma) \ge f(\delta)$ is in S. Thus, $f[S]$ is cofinal in β, so there is a cofinal subset of β with order type $\mathrm{type}(S)$. However, $\mathrm{type}(S) \le \mathrm{type}(\alpha) = \alpha$, by Proposition 5.14. Thus, $\mathrm{cf}(\beta) \le \alpha$, as desired. □

Corollary 9.4. *For every ordinal α, its cofinality* $\mathrm{cf}(\alpha)$ *is a cardinal.*

Proof. We must show that $\mathrm{cf}(\alpha)$ is not in bijection with any lesser ordinal. Suppose β is any ordinal in bijection with $\mathrm{cf}(\alpha)$. Then the bijection is a function from β to $\mathrm{cf}(\alpha)$ with cofinal image, and hence $\mathrm{cf}(\mathrm{cf}(\alpha)) \le \beta$ by Lemma 9.3. This means $\beta \ge \mathrm{cf}(\alpha)$, as desired. □

This means we can think of $\mathrm{cf}(\alpha)$ as measuring the size of the smallest cofinal subset of α in either of two ways, by order type or by cardinality, and they give the same answer:

Corollary 9.5. *For every ordinal α,* $\mathrm{cf}(\alpha)$ *is the smallest cardinality $|S|$ of a cofinal subset S of α.*

Definition 9.6. An infinite cardinal κ is *regular* if $\kappa = \mathrm{cf}(\kappa)$, and *singular* otherwise.

This definition is not interesting for finite cardinals, so we will not apply it to them.

As a first example, $\aleph_0$ is regular, since every finite subset of $\aleph_0$ has a maximal element and is therefore not cofinal. In fact, the same is true of every successor cardinal, which justifies the implication of the word "regular" that most cardinals are regular:

Proposition 9.7. *For every ordinal α, the successor cardinal $\aleph_{\alpha^+}$ is regular.*

Proof. If $S \subseteq \aleph_{\alpha^+}$ is cofinal, then

$$\aleph_{\alpha^+} = \bigcup_{\beta \in S} \beta,$$

by Corollary 4.22. It follows that

$$\aleph_{\alpha^+} \le \sum_{\beta \in S} |\beta| \le \sum_{\beta \in S} \aleph_\alpha = \aleph_\alpha |S|,$$

because every element β of $\aleph_{\alpha^+}$ satisfies $|\beta| \le \aleph_\alpha$. Thus, $|S| \ge \aleph_{\alpha^+}$, since otherwise we would have $|S| \le \aleph_\alpha$ and hence

$$\aleph_{\alpha^+} \le \aleph_\alpha |S| \le \aleph_\alpha^2 = \aleph_\alpha.$$ □

However, not every infinite cardinal is regular. For example, $\aleph_\omega$ is a singular cardinal, because the subset $\{\aleph_i : i \in \omega\}$ is cofinal.

Proposition 9.8. *If α is a nonzero limit ordinal, then* $\mathrm{cf}(\aleph_\alpha) = \mathrm{cf}(\alpha)$.

Note that $\mathrm{cf}(\aleph_0) = \aleph_0 \neq 0 = \mathrm{cf}(0)$, so we must make an exception for $\alpha = 0$. Furthermore,

$$\mathrm{cf}(\aleph_{\alpha^+}) = \aleph_{\alpha^+} \neq 1 = \mathrm{cf}(\alpha^+),$$

which means the analogue of Proposition 9.8 is absolutely false for successor cardinals.

Proof. Because α is a nonzero limit ordinal, we have $\aleph_\alpha > \aleph_0$, and $\aleph_\beta < \aleph_\alpha$ implies $\aleph_{\beta^+} < \aleph_\alpha$. It follows that for every element γ of $\aleph_\alpha$, there exists an infinite cardinal κ in $\aleph_\alpha$ such that $\gamma < \kappa$. (Because $\aleph_\alpha$ is a cardinal, $|\gamma| < \aleph_\alpha$. If $|\gamma| = \aleph_\beta$, then we can take $\kappa = \aleph_{\beta^+}$.)

Given any cofinal subset S of $\aleph_\alpha$, without loss of generality we can assume that S consists only of infinite cardinals. Specifically, we can replace each element of S with the smallest infinite cardinal greater than it; this preserves cofinality and cannot increase the size of S.

Thus, we can assume S is of the form $\{\aleph_\beta : \beta \in T\}$ for some subset T of α. Then S is cofinal in $\aleph_\alpha$ if and only if T is cofinal in α, and hence $\mathrm{cf}(\aleph_\alpha) = \mathrm{cf}(\alpha)$. □

With the sole exception of $\aleph_0$, all the limit cardinals one can think of are singular, like $\aleph_\omega$. Is every regular cardinal beyond $\aleph_0$ a successor cardinal? Counterexamples are called weakly inaccessible cardinals:

Definition 9.9. A *weakly inaccessible* cardinal is an uncountable cardinal that is both regular and a limit cardinal.

ZFC cannot prove that weakly inaccessible cardinals exist, unless it is itself inconsistent. It cannot even prove relative consistency (that if ZFC is consistent, then so is ZFC plus existence of a weakly inaccessible cardinal). However, the existence of inaccessibles is a natural way to extend ZFC and the simplest example of a *large cardinal axiom*. We will return to this topic once we have covered Gödel's second incompleteness theorem.

Weakly inaccessible cardinals are huge, vastly bigger than the cardinals ZFC can prove must exist. Think of $\aleph_0$ for comparison, which is far bigger than any finite cardinal. We will explore this in more detail later, but one can get a rough impression as follows. Suppose $\aleph_\alpha$ is weakly inaccessible. Then $\mathrm{cf}(\aleph_\alpha) = \aleph_\alpha$ by regularity, but $\mathrm{cf}(\aleph_\alpha) = \mathrm{cf}(\alpha)$ by Proposition 9.8. Thus, $\aleph_\alpha = \mathrm{cf}(\aleph_\alpha) = \mathrm{cf}(\alpha) \leq \alpha$. We always have $\alpha \leq \aleph_\alpha$ because $\aleph_\alpha$ contains a subset of order type α (namely, $\{\aleph_\beta : \beta < \alpha\}$), and so every weakly inaccessible cardinal satisfies $\aleph_\alpha = \alpha$. It follows that

$$\aleph_\alpha = \aleph_{\aleph_\alpha} = \aleph_{\aleph_{\aleph_\alpha}} = \dots,$$

which makes $\aleph_\alpha$ rather large. (As a side comment, although ZFC cannot prove the existence of weakly inaccessible cardinals, it can prove the existence of fixed points $\aleph_\alpha = \alpha$, such as the union of $\aleph_0, \aleph_{\aleph_0}, \aleph_{\aleph_{\aleph_0}}, \dots$, so this is not as impressive as it sounds. Weakly inaccessible cardinals are much larger than even this calculation suggests.)

Proposition 9.10. *Let κ be an infinite cardinal. Then* $\mathrm{cf}(\kappa)$ *is the least cardinality $|I|$ of a set I for which there are cardinals λ_i for $i \in I$ such that $\lambda_i < \kappa$ and*

$$\sum_{i \in I} \lambda_i = \kappa.$$

In other words, the cofinality of an infinite cardinal κ is the smallest number of summands less than κ that can add up to κ.

Proof. First, we will show that we can write κ as a sum of $\mathrm{cf}(\kappa)$ cardinals that are less than κ. Let $\{\alpha_i : i \in \mathrm{cf}(\kappa)\}$ be a cofinal subset of κ. Then

$$\kappa = \bigcup_{i \in \mathrm{cf}(\kappa)} \alpha_i,$$

by Corollary 4.22. Because $\alpha_i \in \kappa$ and κ is a cardinal, $|\alpha_i| < \kappa$. Thus,

$$\kappa \le \sum_{i \in \mathrm{cf}(\kappa)} |\alpha_i| \le \mathrm{cf}(\kappa)\kappa = \kappa,$$

so

$$\sum_{i \in \mathrm{cf}(\kappa)} |\alpha_i| = \kappa$$

with $|\alpha_i| < \kappa$, as desired.

For the other direction, suppose $|I| < \kappa$ (otherwise, clearly $|I| \ge \mathrm{cf}(\kappa)$) and

$$\kappa = \sum_{i \in I} \lambda_i$$

for some cardinals λ_i with with $\lambda_i < \kappa$. If $\{\lambda_i : i \in I\}$ is not cofinal in κ, then there exists an ordinal $\alpha \in \kappa$ with $\lambda_i \le \alpha$ for all i, and hence $\lambda_i = |\lambda_i| \le |\alpha| < \kappa$. Then

$$\kappa = \sum_{i \in I} \lambda_i \le |I|\,|\alpha| < \kappa,$$

which is a contradiction. Thus, $\{\lambda_i : i \in I\}$ must be cofinal in κ, and so

$$|I| \ge |\{\lambda_i : i \in I\}| \ge \mathrm{cf}(\kappa). \qquad \square$$

We can now formulate two of the most important consequences of König's theorem:

Theorem 9.11. *For every infinite cardinal κ and cardinal $\lambda \ge 2$,*

$$\kappa^{\mathrm{cf}(\kappa)} > \kappa$$

and

$$\mathrm{cf}(\lambda^\kappa) > \kappa.$$

Note that the second part is a strengthening of Cantor's theorem, since it implies that $\lambda^\kappa > \kappa$.

Proof. By Proposition 9.10, we can write

$$\kappa = \sum_{i \in \mathrm{cf}(\kappa)} \lambda_i$$

with $\lambda_i < \kappa$. By König's theorem,

$$\sum_{i \in \mathrm{cf}(\kappa)} \lambda_i < \prod_{i \in \mathrm{cf}(\kappa)} \kappa,$$

which amounts to $\kappa < \kappa^{\mathrm{cf}(\kappa)}$.

For the second part, note that λ^κ is infinite because $\lambda \ge 2$. If $\mathrm{cf}(\lambda^\kappa) \le \kappa$, then

$$\lambda^\kappa < (\lambda^\kappa)^{\mathrm{cf}(\lambda^\kappa)} \le (\lambda^\kappa)^\kappa = \lambda^{\kappa^2} = \lambda^\kappa,$$

which is a contradiction. Thus, $\mathrm{cf}(\lambda^\kappa) > \kappa$. $\square$

Note that the heart of Theorem 9.11 is $\kappa^{\mathrm{cf}(\kappa)} > \kappa$. It's worth giving a second proof of that inequality:

Proof. We'll show that if S is cofinal in κ, then $|\mathcal{F}(S,\kappa)| > \kappa$. Suppose we are given κ elements f_α of $\mathcal{F}(S,\kappa)$, with $\alpha \in \kappa$. We will construct an element g of $\mathcal{F}(S,\kappa)$ that is not of the form f_α. Specifically, we will construct g so that for each $s \in S$, its value $g(s)$ rules out the functions f_α with $\alpha < s$.

For $s \in S$, let $g(s)$ be the least element of $\kappa \setminus \{f_\alpha(s) : \alpha < s\}$, which is a nonempty set since $|\{f_\alpha(s) : \alpha < s\}| \le |s| < \kappa$ (note that $|s| < \kappa$ since κ is a cardinal). Because κ is an infinite cardinal, it is a limit ordinal, and thus for all $\alpha \in \kappa$, there exists an $s \in S$ such that $s > \alpha$. (Cofinality yields $s \ge \alpha$, which we can strengthen in a limit ordinal because it has no maximal element.) In that case $g(s) \ne f_\alpha(s)$, so $g \notin \{f_\alpha : \alpha \in \kappa\}$. It follows that $|\mathcal{F}(S,\kappa)| > \kappa$. □

Theorem 9.11 generalizes Proposition 8.11, because $\mathrm{cf}(\aleph_\omega) = \aleph_0$ but $\mathrm{cf}(2^\kappa) > \kappa \ge \aleph_0$ when κ is infinite.

It turns out that Theorem 9.11 tells us everything ZFC has to say about 2^κ when κ is regular (besides the obvious monotonicity). Stated somewhat informally, we have the following theorem:

Informal Theorem 9.12 (Easton)**.** *Under ZFC, the only constraints on 2^κ for regular infinite cardinals κ and λ are*

$$\kappa < \mathrm{cf}(2^\kappa)$$

and

$$\kappa \le \lambda \quad \Rightarrow \quad 2^\kappa \le 2^\lambda.$$

We will not state this theorem precisely (let alone prove it). Roughly, the idea is that if we define any function f on a set S of regular infinite cardinals, then it is consistent with ZFC that $2^\kappa = f(\kappa)$ for all $\kappa \in S$, provided that f is weakly increasing and $\mathrm{cf}(f(\kappa)) > \kappa$. However, this rough version is not quite true, because one can give circular definitions of f. For example, suppose we try to define $f(\aleph_0)$ to be the first cardinal after $2^{\aleph_0}$ (which is a successor cardinal and thus has cofinality greater than $\aleph_0$). Then it is impossible to have $2^{\aleph_0} = f(\aleph_0)$, but only because we made a contradictory definition of $f(\aleph_0)$ in terms of $2^{\aleph_0}$.

I do not know of an elementary way to make the statement rigorous. The way Easton did it was that he started with a model of ZFC satisfying GCH, so 2^κ is as small as possible in this model. Then given any function f as above, he showed how to extend the model to a bigger model with the same cardinals and cofinalities but with $2^\kappa = f(\kappa)$ for all $\kappa \in S$. (Roughly, you are adding more subsets of certain sets, so 2^κ can get bigger, and you very carefully control how much bigger it gets. You need Cohen's forcing technique, both to obtain the initial model and to extend it.) This makes the circularity issue irrelevant: even if you define f in terms of 2^κ in the initial model, when you extend the model you change the value of 2^κ while keeping f the same, and there is no contradiction.

One consequence of Easton's theorem is that $2^{\aleph_0} = 2^{\aleph_1}$ is consistent with ZFC. In other words, the weak monotonicity in the theorem needn't be strict. It's weird to think that sets of different sizes could have the same number of subsets, but ZFC does not rule it out.

In fact, arbitrarily weird things can happen. For example, it is consistent with ZFC that $2^{\aleph_i} = \aleph_{i+1}$ for $0 \le i \le 10$, $2^{\aleph_i} = \aleph_{592}$ for $11 \le i \le 591$, and $2^{\aleph_i} = \aleph_{i+1}$ for $i \ge 592$, and of course there are far weirder possibilities than this.

It is consistent with ZFC that $2^{\aleph_0} = \aleph_{\omega^+}$, in which case $2^{\aleph_0} > \aleph_\omega$, or that $2^{\aleph_0} = \aleph_1$, in which case $2^{\aleph_0} < \aleph_\omega$. Thus, the cofinality constraint that rules out $2^{\aleph_0} = \aleph_\omega$ is not simply a matter of size.

Easton's theorem tells us nothing about 2^κ when κ is singular. It's tempting to guess that this is just a defect of the proof, and that in fact the regularity assumption could be dropped, but Silver showed that it cannot:

Theorem 9.13 (Silver)**.** *If $\aleph_\alpha$ is a singular cardinal with uncountable cofinality, and if $2^{\aleph_\beta} = \aleph_{\beta^+}$ for all $\beta < \alpha$, then $2^{\aleph_\alpha} = \aleph_{\alpha^+}$.*

In other words, a singular cardinal with uncountable cofinality cannot be the first counterexample to the generalized continuum hypothesis, so ZFC imposes some additional constraints on 2^κ with κ singular. These constraints are still not fully understood, but there are some remarkable results, such as this theorem of Shelah:

Theorem 9.14 (Shelah)**.** *If $2^{\aleph_i} < \aleph_\omega$ for all $i \in \omega$, then*

$$2^{\aleph_\omega} < \aleph_{\aleph_4}.$$

The upper bound of $\aleph_{\aleph_4}$ is truly enormous ($\aleph_4$ is vastly larger than ω), but it is noteworthy that there is any upper bound at all. Of course, we need some hypothesis along the lines given in Shelah's theorem. For example, it is consistent with ZFC that $2^{\aleph_0} > \aleph_{\aleph_4}$, in which case of course $2^{\aleph_\omega} > \aleph_{\aleph_4}$. Nobody knows whether the bound in Shelah's theorem can be improved to $\aleph_{\aleph_3}$.

We won't prove these theorems, but it is not hard to show that there are some additional constraints on singular cardinals:

Proposition 9.15. *If κ is a singular infinite cardinal and $2^\lambda = \mu$ for all infinite cardinals $\lambda < \kappa$, then $2^\kappa = \mu$ as well.*

For example, if $2^{\aleph_i} = \aleph_{\omega^+}$ for all $i \in \omega$ (which happens in some models of ZFC, by Easton's theorem), then $2^{\aleph_\omega} = \aleph_{\omega^+}$.

Proof. Since κ is singular, we can write

$$\kappa = \sum_{i \in I} \lambda_i$$

with $|I| < \kappa$ and $\lambda_i < \kappa$ (and we can assume without loss of generality that I and λ_i are both infinite). Then

$$2^\kappa = \prod_{i \in I} 2^{\lambda_i} = \prod_{i \in I} \mu = \mu^{|I|}.$$

However,

$$\mu^{|I|} = \left(2^{|I|}\right)^{|I|} = 2^{|I|^2} = 2^{|I|} = \mu.$$

Thus, $2^\kappa = \mu$. □