

Infinite Sets

Lecture Notes for NMAI074

Jakub Smolík

`smolikj@matfyz.cz`

Last Update:

March 19, 2026

Contents

Introduction	3
1 Review of Set Theory Basics	4
1.1 Sets and Classes	4
1.2 Axiom Schema of Replacement	4
1.3 Indexed Collections of Sets	5
1.4 Axiom of Choice	5
1.5 Natural Numbers and the Axiom of Infinity	8
1.6 Comparing Sizes of Sets	8
1.7 Well-Orderings and Initial Segments	9
2 Ordinal Numbers	11
2.1 Ordinals as a Generalization of Naturals	11
2.2 Ordinals as Types of Well-Ordered Sets	12
2.3 Transfinite Induction and Recursion	14
2.4 Well-Ordering Principle	16
2.5 Zorn's Lemma	17
2.6 Trichotomy Principle	20
3 Operations on Ordinals	20
3.1 Ordinal Functions	20
3.2 Ordinal Arithmetic	24
3.2.1 Definitions and Intuition	24
3.2.2 Properties of Ordinal Operations	26
3.2.3 Ordinal Equations and Power Expansions	29
3.3 Countable and Uncountable Ordinals	32
3.3.1 Epsilon Numbers	32
3.3.2 Veblen Hierarchy	35
3.3.3 Hartogs' Theorem	37
3.4 Peano Arithmetic	39
3.4.1 Peano Axioms	39
3.4.2 Models of Arithmetic	40
3.4.3 Gödel's Incompleteness Theorems	41
3.4.4 Consistency and the Connection with ε_0	43
3.4.5 Limits of Predicative Mathematics	44
3.5 Applications of Countable Ordinals	45
3.5.1 Goodstein Sequences	45
3.5.2 Kirby–Paris Hydra Game	50
3.5.3 Fundamental Sequences	51
3.5.4 Fast-Growing Hierarchy	52
4 Cardinal Numbers	56
4.1 Cardinals as Sizes of Well-Ordered Sets	56
4.2 Infinite Cardinals \aleph_α	58
4.3 Cofinality and Inaccessible Cardinals	62
4.4 Cardinal Arithmetic	66
4.4.1 Cardinality of the Continuum	67

4.4.2	Infinite Sums and Products	67
4.4.3	König's Inequality	70
4.4.4	Continuum Hypothesis	71
5	Large Cardinals	73
5.1	Cumulative Hierarchy of Sets V_α	73
5.2	Beth Numbers \beth_α	76
5.3	Constructible Universe L_α	77
5.4	Significance of the Axiom of Infinity	78
5.5	Large Cardinal Axioms	80
6	Topics in Pure Set Theory	82
6.1	Filters and Ultrafilters	82
6.2	Stationary Sets	87
6.3	Infinite Trees	87
7	Infinitary Combinatorics	90
7.1	Compactness Principles	90
7.1.1	Rado's Selection Principle	90
7.1.2	Infinite Hall's Theorem	92
7.1.3	De Bruijn–Erdős Theorem	92
7.1.4	Compactness in Ramsey Theory	93
7.2	Chromatic Numbers of Infinite Graphs	95
7.3	Partition Relations of Cardinals	98
7.3.1	Limits of the Partition Arrow	99
7.3.2	Ramsey's Theorem	100
7.3.3	An Unprovable Theorem	102
7.3.4	Transfinite Partition Relations	103
7.3.5	Partitions and Large Cardinals	105
7.4	Banach–Tarski Paradox	109
	Appendix	112
A.1	Group Theory	112
A.2	Measure Theory	113
	Bibliography	120

Introduction

This document serves as lecture notes for the course NMAI074 taught at MFF CUNI, but in some places it goes far beyond the scope of the lecture. A significant portion of the text follows parts of the second and third chapters of [1]. If you found any mistakes, please contact me at smolikj@matfyz.cz.

I would like to express my gratitude to Jan Kynčl for teaching the course and for giving me useful feedback and advice.

1 Review of Set Theory Basics

We will be working within Zermelo–Fraenkel (ZF) set theory; that is, Zermelo’s (1908) theory Z^1 augmented² by Fraenkel’s (1922) axiom schema of replacement and von Neumann’s (1925) axiom of foundation. If we further include the axiom of choice, we obtain the much stronger theory ZFC.

As for notation, I always use the symbol \subset for a proper subset (or subclass) and \subseteq for a general subset (or subclass). I use \subsetneq only when it is important that the two sets (or classes) are not equal. Concatenated expressions such as $a \in b \in c$ mean $a \in b \wedge b \in c$. I differentiate between the symbol for equality of two objects “=” and the symbol for the definition of an object “:=”. I use the following notation for defining functions:

- $f : A \rightarrow B$ is a function with domain A and codomain B ,
- $f : a \mapsto b$ denotes that f maps the set $a \in A$ to the set $b \in B$,
- $f = g \circ h$ means that $f(x) = h(g(x))$ for all suitable x .

I use the terms “function,” “map,” and “mapping” interchangeably. If f maps A into B , it means that $f[A] \subseteq B$, and if f maps A onto B , it means that $f[A] = B$.

1.1 Sets and Classes

Definition 1.1 (Class). If $\varphi(x)$ is a formula, then the expression $\{x \mid \varphi(x)\}$ is called a *class term*. It defines the “collection” of all sets x satisfying $\varphi(x)$. We call this collection the *class* determined by $\varphi(x)$.

Every set is a class, but not all classes are sets (consider the class of all sets). A class that is not a set is called a *proper class*. The major difference between sets and classes is that classes cannot be members of other classes or sets, while sets can. We can substitute class terms into logical formulas in place of free variables, but unlike sets, we cannot quantify them using \forall and \exists . It is not hard to show that for every formula with class terms (but without quantified class variables), there is an equivalent formula in the base language without class terms.

We will usually denote sets using small letters a, b, c, x, y, \dots and classes using capital letters A, B, C , etc. The exception to this are well-ordered sets, which will often be denoted as W . Finally, the class of all sets, also called the *universal class*, is denoted by \mathbf{V} .

1.2 Axiom Schema of Replacement

Axiom 1.2. When we take any (even a class) map F and a preimage set a , then the class of images $b = F[a]$ is also a set. Formally, if $\psi(x, y)$ is a formula without free variables y_1, y_2 and b , then

$$\begin{aligned} (\forall x)(\forall y_1, y_2)((\psi(x, y_1) \wedge \psi(x, y_2)) \Rightarrow y_1 = y_2) \Rightarrow \\ (\forall a)(\exists b) : (\forall y)(y \in b \Leftrightarrow (\exists x)(x \in a \wedge \psi(x, y))) \end{aligned}$$

¹Zermelo’s theory Z actually already contained AC, which he formulated in 1904.

²Replacement is necessary to prove certain important theorems and the existence of some key sets. Foundation is more of a “cleanup” axiom, as virtually all results in the branches of mathematics based on set theory hold even without it.

is an axiom. The formula $\psi(x, y_1)$, resp. $\psi(x, y_2)$ are created from $\psi(x, y)$ by substituting y_1 , resp. y_2 for y .

The first part of this axiom says that $\psi(x, y)$ should behave like a map $y = F(x)$. In the second part, a denotes the set of preimages and b the set of corresponding images.

1.3 Indexed Collections of Sets

Definition 1.3. A *collection of sets* $\langle X_i \mid i \in I \rangle$ is a map F with domain I , where X_i denotes the set $F(i)$. We say that I is the *index set* or *index class* of the collection. We define union, intersection, and Cartesian product as

$$\begin{aligned} \bigcup_{i \in I} X_i &:= \bigcup \text{Rng}(F), \\ \bigcap_{i \in I} X_i &:= \bigcap \text{Rng}(F), \\ \times_{i \in I} X_i &:= \left\{ f \mid f : I \rightarrow \bigcup_{i \in I} X_i \wedge (\forall i \in I)(f(i) \in X_i) \right\}. \end{aligned}$$

We define the Cartesian product only for collections indexed by a set, as the mappings f would otherwise be proper classes.

Exercise 1.4. Verify that if I is a set, then $\bigcup X_i$, $\bigcap X_i$, and $\times X_i$ are also sets.

1.4 Axiom of Choice

The *axiom of choice*, denoted AC, is one of the most important principles in modern mathematics, with profound implications in areas such as classical analysis or linear algebra. It states that for any collection of nonempty sets, it is possible to choose exactly one element from each set, even if the collection is infinite. When added to Zermelo–Fraenkel set theory, it yields the much more powerful ZFC. Many theorems that seem intuitively true, such as every vector space having a basis, depend on this axiom.

However, the axiom of choice is also controversial, as it leads to counter-intuitive results, such as the well-ordering principle, which claims that every set can be well-ordered, or the Banach–Tarski paradox, which provides a way to decompose a solid ball into finitely many pieces and reassemble them into two identical copies of the original.

Definition 1.5 (Choice function). A *choice function* (or a *selector*) on the set x is any function $f : x \rightarrow \bigcup x$ such that

$$(\forall t \in x)(t \neq \emptyset \Rightarrow f(t) \in t).$$

We can WLOG assume that the choice function is defined on $x \setminus \{\emptyset\}$ and all $t \in \text{Dom}(f)$ satisfy $f(t) \in t$.

One can prove in ZF via finite induction that every finite set has a choice function; that is, we are allowed to make finitely many choices (out of potentially

infinite sets). However, it can be shown that ZF cannot prove that every countable set has a choice function, and certainly not that *every* set has a choice function.

Since the assumption that *every* set has a choice function can lead to some paradoxical results (such as the Banach–Tarski paradox), we distinguish three different “power levels” of this axiom:

Axiom 1.6 (Axiom of Countable Choice AC_ω). Every countable set has a choice function; one can make only a countable number of choices.

Axiom 1.7 (Axiom of Dependent Choice DC). One can make a countable sequence of choices, where each choice may *depend* on the previous ones. Formally, for any nonempty set A and a binary relation $R \subseteq A \times A$ such that

$$(\forall x \in A)(\exists y \in A) x R y,$$

there exists an infinite sequence (x_n) satisfying $x_n R x_{n+1}$ for all $n \in \omega$.

Axiom 1.8 (Axiom of Choice AC). Every set has a choice function; the amount of choices is not limited.

Exercise 1.9. Show that $\text{AC} \implies \text{DC} \implies \text{AC}_\omega$.

Exercise 1.10. Show that $\text{AC} \iff$ the Cartesian product of an arbitrary collection of nonempty sets is nonempty.

We will now present a brief overview of some of the results that can be proven from each power level (but cannot be proven in ZF). More details and proofs for most of them can be found in [14].

Axiom of countable choice AC_ω

- \iff any Cartesian product of countably many nonempty sets is nonempty.
- \implies any union of countably many countable sets is countable. Without AC_ω , a countable union of finite sets might be uncountable.
- \implies ω_1 is a regular cardinal.
- \implies every infinite set x has a countably infinite subset; equivalently $\omega \preceq x$.
- \implies every set x is finite \iff it is Dedekind finite; that is $(\forall y)(y \subset x \implies y \prec x)$.
- \implies every function is continuous \iff it is Heine continuous; that is: for every convergent sequence (x_n) we have $\lim f(x_n) = f(\lim x_n)$.
- \implies König’s lemma: every infinite tree with finite degrees contains an infinite path.

Axiom of dependent choice DC

- \implies a partially ordered set is well-founded (every nonempty subset has a minimal element) \iff it contains no infinite strictly descending sequences.
- \implies most results of classical analysis and topology.

Axiom of choice AC

- \iff any Cartesian product of nonempty sets is nonempty.
- \implies there exists a mapping which assigns to each set x a set $|x|$ such that $x \approx |x|$ and $x \approx y \iff |x| = |y|$.
- \iff for any infinite set x it holds that $x \approx x \times x$.
- \iff every vector space (even of infinite dimension) has a basis.
- \iff the product of (even infinitely many) compact topological spaces is compact.
- \iff every surjection $f : X \rightarrow Y$ has a right inverse, i.e. a function $g : Y \rightarrow X$ such that $f(g(y)) = y$ for all $y \in Y$.
- \iff every connected (even infinite) graph has a spanning tree.
- \implies the Compactness Theorem in first-order logic: if every finite subset of a theory T has a model, then T has a model.
- \implies every infinite set has a non-principal ultrafilter.
- \implies it is possible to well-order the set of real numbers \mathbb{R} .
- \iff **the Well-Ordering Principle**: every set can be well-ordered.
- \iff **Zorn's Lemma**: every ordered set containing upper bounds for every chain necessarily contains at least one maximal element.
- \iff **the Trichotomy Principle**: for any sets x and y either $x \preceq y$, or $y \preceq x$.

We will show the equivalence³ of AC and the last three conditions in Sections 2.4, 2.5 and 2.6.

Paradoxical results implied by AC

- \implies the Banach–Tarski paradox: it is possible to decompose a solid ball into a few pieces and reassemble them into two identical copies of the original ball. Vsauce has a VIDEO with an intuitive explanation.
- \implies there exist subsets of \mathbb{R} that are not Lebesgue measurable. The most famous such set is probably the *Vitali set*. Veritasium has a great VIDEO on this topic (and the history of AC in general).
- \implies there exists a function $f : \mathbb{R} \rightarrow \mathbb{R}$ satisfying $f(x + y) = f(x) + f(y)$ (linearity) that is continuous nowhere, and the graph of f is dense in \mathbb{R}^2 . It is constructed via a Hamel basis: a basis of the vector space \mathbb{R} over \mathbb{Q} .

The axioms of countable and dependent choice are implicitly used in disciplines such as classical analysis all the time. The full power of the axiom of choice is rarely needed, and we try to avoid it when possible.

³“The Axiom of Choice is obviously true, the well-ordering principle obviously false, and who can tell about Zorn’s lemma?” — Jerry Bona

1.5 Natural Numbers and the Axiom of Infinity

We use *von Neumann ordinals*, meaning that natural numbers are defined as

$$0 := \emptyset, 1 := \{0\}, 2 := \{0, 1\}, \dots, n + 1 := \{0, 1, \dots, n\} = n \cup \{n\}.$$

Definition 1.11. The *successor function* is a mapping $S : \mathbf{V} \rightarrow \mathbf{V}$ defined as $v \mapsto v \cup \{v\}$. For convenience, we write $v + 1 := S(v) = v \cup \{v\}$.

Definition 1.12. A set w is *inductive* if $0 \in w$ and for all $n \in w$ also $n + 1 \in w$.

Axiom 1.13 (Axiom of Infinity). There exists an inductive set.

Definition 1.14. We define the *set of all natural numbers* as the \subseteq -smallest inductive set. Or, equivalently, as $\bigcap\{w \mid w \text{ is inductive}\}$. We denote it by ω .

1.6 Comparing Sizes of Sets

Definition 1.15. For sets x and y we define the relations

- (a) $x \approx y$, if there exists a bijection $x \rightarrow y$, ... x and y are *equinumerous*
- (b) $x \preceq y$, if there exists an injection $x \rightarrow y$, ... x is *subvalent* to y
- (c) $x \prec y$, if $x \preceq y$ and $x \not\approx y$.

Theorem 1.16 (Cantor–Bernstein theorem). $x \approx y \iff (x \preceq y \wedge y \preceq x)$.

Theorem 1.17 (Cantor’s theorem). $x \prec \mathcal{P}(x)$ for every set x .

Definition 1.18. We say that a set x is

- (a) *finite* if $x \approx n$ for some $n \in \omega$, and *infinite* otherwise,
- (b) *countable* if $x \preceq \omega$,
- (c) *countably infinite* if $x \approx \omega$,
- (d) *uncountable* if it is not countable ($x \not\preceq \omega$).

It is easy to see that finite sets are subvalent to infinite sets and that countably infinite and uncountable sets are infinite.

Exercise 1.19. Verify that a countable set that is not countably infinite is finite.

Hint. Consider what happens if a subset $A \subseteq \omega$ is bounded or unbounded, and use the result of Exercise 2.18.

Definition 1.20. A set x is *Dedekind-finite* if $y \prec x$ for each $y \subset x$; otherwise, it is *Dedekind-infinite*; that is, there exists an equinumerous subset.

Clearly, each finite set is Dedekind-finite, and thus every Dedekind-infinite set is infinite. However, the other implication does not hold in ZF.

Fact 1.21. *There exist models of ZF containing infinite but Dedekind-finite sets.*

Observation 1.22. *If a set x is infinite but Dedekind-finite, then it contains no countably infinite subset $y \subseteq x$; that is, $\omega \not\preceq x$.*

Proof. Suppose x has a countable subset $y = \{a_0, a_1, a_2, \dots\}$. Then we can construct a bijection $g : x \rightarrow x \setminus \{a_0\}$ as $g(a) = a$ if $a \notin y$, and $g(a_n) = a_{n+1}$. \square

Hence such sets x are uncountable, but it is not true that $\omega \prec x$.

Theorem 1.23 (AC_ω). *Every infinite set contains a countably infinite subset. Thus every infinite set is Dedekind-infinite, and every Dedekind-finite set is finite.*

Proof. Let x be an infinite set. Because $n \prec x$ for each $n \in \omega$, we can choose an injection $G_n : n \rightarrow x$ using AC_ω for each $n \in \omega$. Note that each G_n corresponds to a sequence a_0, a_1, \dots, a_{n-1} . Writing the elements of all the sequences $G_0, G_1, G_2, \dots, G_n, \dots$ one after another and erasing all occurrences of each member a of x other than its first occurrence yields an infinite sequence H of length ω of different members of x . H is indeed infinite since for every n the sequence G_n consists of n different terms. The formal definition of the sequence H is left to the reader. $\text{Rng}(H)$ is a countably infinite subset of x . \square

Corollary 1.24 (AC_ω). *A set x is uncountable $\iff \omega \prec x$*

Proof. “ \implies ”: x is an infinite set, so by the previous theorem $\omega \preceq x$. But $x \not\approx \omega$ since $x \not\preceq \omega$. “ \impliedby ”: if x were countable, that is $x \preceq \omega$, then by the Cantor–Bernstein theorem (since $\omega \preceq x$) we have $x \approx \omega$, a contradiction with $\omega \prec x$. \square

1.7 Well-Orderings and Initial Segments

Definition 1.25 (Ordering). A binary relation R on a class X is a

- (a) *trichotomy* if for all $x, y \in X$, either $x = y$, or $x R y$, or $y R x$,
- (b) *strict order* if it is anti-reflexive, strongly anti-symmetric, and transitive on X ; (note that strong anti-symmetry follows from the other two),
- (c) *(partial) order* if it is reflexive, weakly anti-symmetric, and transitive on X ,
- (d) *total (or linear) order* if it is a trichotomous partial order on X .

If R is an ordering, then instead of $x R y$ we write $x \leq_R y$ and we call (X, \leq_R) an *ordered class*. Similarly, if R is a strict ordering, then we write $x <_R y$ and we call $(X, <_R)$ a *strictly ordered class*.

Note that we can easily create a strict ordering $<_R$ from \leq_R and vice versa. For this reason, we will not define properties for both strict and non-strict orderings separately, because one implicitly defines the other.

Definition 1.26. We call an element of an ordered class (X, \leq) *minimal* if there is no smaller one, and we call it a *minimum* if it is smaller than all others. If a minimum exists, we denote it by $\min_{\leq}(X)$. The *supremum* of a subset $Y \subseteq X$ is the minimum of all its upper bounds. If it exists, we denote it by $\sup_{\leq}(Y)$

Observation 1.27. *Every minimum is minimal. Furthermore, if \leq_R is a total order, then there is at most one minimal element, and if it exists, then it is also the minimum. There is always at most one minimum.*

Definition 1.28 (Well-ordering). An ordered class (A, \leq_R) is

- (a) *well-founded* if every non-empty subset of A has a minimal element.
- (b) *well-ordered* if every non-empty subset of A has a minimum (least element).

Notice that every well-ordered class is totally ordered since we can take any two elements, and one of them has to be the minimum and is therefore smaller.

Observation 1.29. *Well-order \iff well-founded total order.*

Observation 1.30. *The well-ordering property is hereditary. That is, if X is well-ordered by \leq_R , then every $Y \subseteq X$ is also well-ordered by \leq_R .*

Observation 1.31. *Well-founded ordered sets contain no infinite strictly decreasing sequences, as such a sequence has no minimal element.*

Exercise 1.32. Prove that the reverse implication also holds, provided we accept the axiom of dependent choice (see Axiom 1.7).

Definition 1.33 (Lower part and subset). Let $(A, <_R)$ be a (strictly) ordered class. A subclass $X \subseteq A$ is a *lower part* of A if

$$(\forall x \in X)(\forall a \in A)(a <_R x \Rightarrow a \in X).$$

Additionally, if X is a set, we call it a *lower subset* of A , and if $X \neq A$, then we call it a *proper lower part*, or *proper lower subset* of A .

Lemma 1.34. *Let $(W, <_R)$ be a (strictly) well-ordered set, and suppose that X is a proper lower subset of W . Then there exists a unique $x \in W$ such that X is equal to the set $\{y \in W \mid y <_R x\}$. We denote this set as (\leftarrow, x) .*

Proof. We define x as the minimum of $W \setminus X$. Then every $y <_R x$ belongs to X , so $(\leftarrow, x) \subseteq X$. We also want the opposite inclusion. For contradiction, suppose there is a $y \in X$ such that $y \notin (\leftarrow, x)$. If $y \not<_R x$, then necessarily $x <_R y$ as $x \neq y$ since $x \notin X$. But this means that $x \in X$ because X is a lower subset and $y \in X$. But this is a contradiction since $x \notin X$. \square

Definition 1.35 (Initial segment). If $(W, <_R)$ is a (strictly) well-ordered set, then we call its proper lower subsets *initial segments* instead. We denote the unique initial segment of W determined by $x \in W$ as

$$(\leftarrow, x) := \{y \in W \mid y <_R x\}.$$

It contains all the elements of W from the minimum of W until x , but not x itself.

Observation 1.36. *Note that $x <_R y \iff (\leftarrow, x) \subset (\leftarrow, y)$.*

2 Ordinal Numbers

Informally, *ordinal numbers* are a way to generalize natural numbers. We will first do a quick recap of the basics of ordinal numbers and then prove a theorem that deeply links ordinals and well-ordered sets.

2.1 Ordinals as a Generalization of Naturals

Definition 2.1. A class X is called *transitive* if for all $x \in X$ we have $x \subseteq X$. Or equivalently, if for every x, y such that $y \in x \in X$ we have $y \in X$.

Theorem 2.2. *Every natural number and the set of all natural numbers ω is transitive and (strictly) well-ordered by the membership relation \in .*

From now on, we will denote the (strictly) well-ordered set (ω, \in) as $(\omega, <)$ instead and write $n < m$ instead of $n \in m$ when talking about natural numbers.

Definition 2.3 (Ordinal numbers). A set α is an *ordinal number* if it is transitive and (strictly) well-ordered by the membership relation \in . If α is infinite, we say that it is a *transfinite ordinal*. We denote the *class of all ordinal numbers* by On .

Theorem 2.4. *Finite ordinals are exactly the natural numbers, and ω is the smallest transfinite ordinal.*

Theorem 2.5. *The class On itself is transitive and (strictly) well-ordered by \in . This implies that it is not a set; otherwise, $\text{On} \in \text{On}$. Furthermore, any proper class X that is transitive and well-ordered by \in is identical to On .*

As for notation, we will usually denote ordinals using letters from the beginning of the Greek alphabet: $\alpha, \beta, \gamma, \delta \dots$. An exception to this is the letter λ , which we will reserve for limit ordinals. Furthermore, we compare ordinals using the symbol ' $<$ '. That is, we write $\beta < \alpha$ instead of $\beta \in \alpha$.

Observation 2.6. *If $\beta < \alpha$, then $\beta \subset \alpha$ and β is an initial segment of α . Additionally, $\alpha = (\leftarrow, \alpha)$.*

Definition 2.7. If α is an ordinal, then we call all $\beta < \alpha$ the *predecessors* of α . The *successor* of α is the ordinal $\alpha + 1 := \alpha \cup \{\alpha\}$. We say that α is the *direct predecessor* of $\alpha + 1$.

Remark. It is easy to show that $\alpha + 1$ is the smallest ordinal larger than α .

Definition 2.8. We say that an ordinal number α is an

- (a) *isolated* ordinal if $\alpha = 0$ or α has a direct predecessor,
- (b) a *limit* ordinal otherwise.

Isolated ordinals $\alpha > 0$ are also sometimes called *successor* ordinals.

Example. Every $n \in \omega$ is isolated, ω is limit, and $\omega + 1$ is isolated again.

2.2 Ordinals as Types of Well-Ordered Sets

The definition of ordinals presented above was formalized by John von Neumann in 1923. This elegant approach, however, came decades after Georg Cantor first introduced ordinals (around 1885) as *order types of well-ordered sets*. Cantor's intuition was that ordinals serve as labels for well-ordered sets: the smallest element is labeled 0, the next 1, and so on. The *order type* of the set is then the first label we did not have to use; it represents the "shape" of the ordering.

Consider, for example, a set ordered as

$$a_0 < a_1 < a_2 < \overbrace{\dots}^{\infty} < b.$$

Here, there are countably infinitely many elements a_i , followed by one additional element b . leaving no finite label for b . This is precisely why we need transfinite ordinals: we assign the label ω to b . Hence, the order type of this ordering is $\omega + 1$.

It is important to realize that different orderings of the same sets can have different order types. This means that the ordinal numbers do not count the number of objects in the set; they only label them.




Figure 1: Different orderings of the same set can have different types [7].

Lemma 2.9. *Every proper lower part of $(\text{On}, <)$ is an ordinal number.*

Proof. Let X be a proper lower part of On . Then

- (i) X is transitive. Suppose $\alpha \in \beta \in X$, that is $\alpha < \beta \in X$. Because X is a lower part, we have $\alpha \in X$.
- (ii) X is well-ordered by \in because On is well-ordered by \in and $X \subseteq \text{On}$.

We also need to argue that X is a set. If it were a proper class, then by Theorem 2.5 it would be the entire On , but $X \subsetneq \text{On}$. \square

Definition 2.10 (Isomorphism). Let (A, \leq_R) and (B, \leq_S) be ordered classes. A bijection $F : A \rightarrow B$ is an *order-isomorphism* of (A, \leq_R) and (B, \leq_S) if

$$(\forall x, y \in A)(x \leq_R y \iff F(x) \leq_S F(y)).$$

Because we will not be dealing with other types of isomorphisms, we will usually simply say *isomorphism* instead of order-isomorphism.

Theorem 2.11 (About comparing well-orderings). *If (W_1, \leq_1) and (W_2, \leq_2) are well-ordered sets, then exactly one of the following holds:*

- (a) *either W_1 and W_2 are isomorphic, or*
- (b) *W_1 is isomorphic to an initial segment of W_2 , or*
- (c) *W_2 is isomorphic to an initial segment of W_1 .*

In each case, the isomorphism is unique.

Corollary 2.12. *No two distinct ordinal numbers can be isomorphic.*

Proof. Suppose $\alpha < \beta$, that is $\alpha \in \beta$ and $\alpha \subset \beta$. Clearly α is an initial segment of β . This means that we are in case (b) of the previous theorem. \square

Theorem 2.13 (About the type of well-ordering). *Every well-ordered set W is isomorphic to a unique ordinal number α , which is called the order type of W and is denoted by $\text{otp}(W)$.*

Proof (following [19, Thm. 3.1, Chap. 6]). Let $(W, <_R)$ be a well-ordered set. We want to show that there is a unique ordinal $(\alpha, <)$ isomorphic to it. Define X as the set of all $x \in W$ for which (\leftarrow, x) is isomorphic to an ordinal. As no two distinct ordinals are isomorphic, this ordinal is uniquely determined, and we denote it α_x ; we denote the isomorphism as $i_x : (\leftarrow, x) \rightarrow \alpha_x$.

Suppose that there exists a set S such that $S = \{\alpha_x \mid x \in X\} \subseteq \text{On}$. Because we assume that S is a set, $S \subsetneq \text{On}$. We claim that S is a proper lower part of $(\text{On}, <)$, and thus, by Lemma 2.9, it is an ordinal; let's call it α . Indeed, suppose $\beta < \alpha_x \in S$, we want $\beta \in S$. Note that β is an initial segment of α_x . This implies that $i_x^{-1}[\beta]$ is an initial segment of W . Because W is well-ordered, $i_x^{-1}[\beta]$ is equal to (\leftarrow, b) for some $b \in W$ (using Lemma 1.34). So $\beta = \alpha_b \in S$ by the definition of S . More precisely, $i_x \upharpoonright (\leftarrow, b)$ is an isomorphism between (\leftarrow, b) and β . We will argue why we can make the assumption that S is a set later.

A similar argument shows that X is a lower subset of W . To show this, suppose $x \in X$ and take $y \in W$ such that $y <_R x$. We want $y \in X$. We have $y <_R x$; therefore, (\leftarrow, y) is an initial segment of (\leftarrow, x) . Because isomorphisms conserve all ordering properties, $i_x \upharpoonright (\leftarrow, y)$ is an isomorphism between (\leftarrow, y) and an initial segment of α_x . By Lemma 2.9, this is an ordinal; by our previous notation, α_y . Therefore $y \in X$.

We conclude that either $X = W$ or $X = (\leftarrow, c) \subset W$ for some $c \in W$ (using Lemma 1.34). We now define a function $f : X \rightarrow S = \alpha$ by $f : x \mapsto \alpha_x$. From the definition of S and the fact that

$$x <_R y \iff (\leftarrow, x) \subset (\leftarrow, y) \iff \alpha_x \subset \alpha_y \iff \alpha_x < \alpha_y,$$

it is obvious that f is an isomorphism of $(X, <_R)$ and $(\alpha, <)$. If

- $X = (\leftarrow, c)$, then by the definition of the set X , $c \in X$ because (\leftarrow, c) is isomorphic to an ordinal $\alpha_c = \alpha$. But this is a contradiction because $c \notin (\leftarrow, c) = X$.
- Therefore, $X = W$ and α is the sought-after ordinal isomorphic to $(W, <_R)$.

The uniqueness of α follows from the simple observation that if W were isomorphic to two distinct α_1 and α_2 , then by the transitivity of isomorphisms, the ordinals α_1 and α_2 would be isomorphic, which is impossible by Corollary 2.12.

This would complete the proof if we were justified to make the assumption that the class S is a set and therefore an ordinal. In fact, we have to use the axiom of replacement to guarantee it. If we assume this axiom, then S is a set because it is the image of the set X under the map f . \square

Exercise 2.14. Is there a well-ordered proper class not isomorphic to $(\text{On}, <)$?

Hint. Try to modify On to contradict the property described in Lemma 2.9. If you run out of ideas, consult Section 3.2.

2.3 Transfinite Induction and Recursion

In mathematics, we often use induction on the natural numbers to prove statements, and we can use recursion, such as $f(0) = 1$ and $f(n) = n \cdot f(n - 1)$, to define functions. We will now show how to generalize this to all ordinals.

Theorem 2.15 (Transfinite Induction Principle). *Let $A \subseteq \text{On}$ be a class such that for all ordinals $\alpha \in \text{On}$, we have $\alpha \subseteq A \Rightarrow \alpha \in A$, or in other words*

$$(\forall \beta < \alpha)(\beta \in A) \implies (\alpha \in A). \quad (2.1)$$

Then $A = \text{On}$.

Equivalently, assume that $\varphi(x)$ is a property, and for all ordinals α :

If $\varphi(\beta)$ holds for all $\beta < \alpha$, then $\varphi(\alpha)$.

Then $\varphi(\alpha)$ holds for all ordinals α .

Proof. Suppose that $\gamma \in \text{On} \setminus A$ and let $S = \{\alpha \leq \gamma \mid \alpha \notin A\}$. Because ordinals are well-ordered, the set S has a minimum element α . Since every $\beta < \alpha$ is in A , it follows by (2.1) that $\alpha \in A$, which is a contradiction.

The equivalence can be easily seen by taking the class $A = \{x \mid \varphi(x)\}$ or the property $\varphi(x) = x \in A$. \square

We can also formulate the principle separately for isolated and limit ordinals, which allows us to use the transfinite induction principle in a form closer to the usual formulation of the induction principle for the naturals.

Theorem 2.16 (Transfinite Induction Principle II). *Let $A \subseteq \text{On}$ be a class satisfying*

(i) $0 \in A$,

(ii) $\alpha \in A \Rightarrow \alpha + 1 \in A$, *... this is just induction on ω*

(iii) *if α is a limit ordinal and $(\forall \beta < \alpha)(\beta \in A)$, then $\alpha \in A$.*

Then $A = \text{On}$. Note that we can again easily reformulate this in terms of a property $\varphi(x)$.

Proof. We need to show that these three assumptions imply (2.1). So let α be an ordinal such that $\beta \in A$ for all $\beta < \alpha$. If $\alpha = 0$, then $\alpha \in A$ by (i). If $\alpha \neq 0$ is isolated, that is if there is a $\beta < \alpha$ such that $\alpha = \beta + 1$, we know that $\beta \in A$, so $\alpha \in A$ by (ii). If α is a limit ordinal, we have $\alpha \in A$ by (iii). \square

We can use transfinite induction to prove properties of certain infinite structures. On the other hand, transfinite recursion—the technique described in the following theorem—allows us to construct various infinitely complex structures and define functions in a recurrent fashion.

Theorem 2.17 (About construction by transfinite recursion). *If $G : \mathbf{V} \rightarrow \mathbf{V}$ is a class map, then there is a unique class map $F : \text{On} \rightarrow \mathbf{V}$ satisfying*

$$F(\alpha) = G(F \upharpoonright \alpha). \quad (2.2)$$

So we define the image of the next ordinal using its predecessors and their images.

Remark. This should seem a bit suspicious because it looks like we are saying that for every class G , there exists a class F for which something holds. But we cannot quantify classes. Well, we can replace the quantification of G with a theorem *schema*, one for each G . And we aren't really quantifying F because the following proof explicitly constructs it.

Remark. The generality of defining $F(\alpha)$ based on the ordered pairs $(\beta, F(\beta))$ for all $\beta < \alpha$ allows us to define functions using many other recursions. For example:

- $F(\alpha) = G(F[\alpha]) = G(\{F(\beta) \mid \beta < \alpha\})$,
- $G : \text{On} \times \mathbf{V} \rightarrow \mathbf{V}$ and $F(\alpha) = G(\alpha, F \upharpoonright \alpha)$,
- $F(\alpha)$ is $G_1(F(\beta))$ if $\alpha = \beta + 1$ is isolated, and $G_2(F[\alpha])$ if α is limit. This is the form we will use most often.

Additionally, the theorem about construction by transfinite recursion is equivalent to the axiom of replacement, as shown in [18].

Proof. We define A as the class of “set approximations” of F . That is set mappings f , the domain of which is some ordinal number β , and that for all $\alpha < \beta$, we have $f(\alpha) = G(f \upharpoonright \alpha)$. Now we define F as $F := \bigcup A$. Clearly $F \subseteq \text{On} \times \mathbf{V}$. We will show that $F : \text{On} \rightarrow \mathbf{V}$ is the unique mapping satisfying (2.2).

First, we show that the approximations of F agree. Let $f, f' \in A$ and $\alpha \in \text{Dom}(f) \cap \text{Dom}(f')$. We claim that $f(\alpha) = f'(\alpha)$. Note that $\text{Dom}(f) \cap \text{Dom}(f')$ is an ordinal δ . For contradiction, suppose that $\alpha \in \delta$ is the smallest ordinal for which $f(\alpha) \neq f'(\alpha)$. Then $f \upharpoonright \alpha = f' \upharpoonright \alpha$ so $f(\alpha) = G(f \upharpoonright \alpha) = G(f' \upharpoonright \alpha) = f'(\alpha)$, a contradiction.

Second, we verify that F satisfies (2.2); that is, for all $\alpha \in \text{Dom}(F)$, we have $F(\alpha) = G(F \upharpoonright \alpha)$. So let $\alpha \in \text{Dom}(F)$. It is there due to some $f \in A$ satisfying $\alpha \in \text{Dom}(f)$ and $f(\alpha) = G(f \upharpoonright \alpha)$. Also, $F(\alpha) = f(\alpha)$ and $F \upharpoonright \alpha = f \upharpoonright \alpha$. Therefore, by combining these equalities $F(\alpha) = G(F \upharpoonright \alpha)$.

Next, we show that $\text{Dom}(F) = \text{On}$. First, we prove that $\text{Dom}(F)$ is a lower part of On . Suppose $\alpha \in \text{Dom}(F)$; then it is there thanks to some $f \in A$ with domain $\delta > \alpha$. If $\beta < \alpha$, then also $\beta \in \delta$, and thus $\beta \in \text{Dom}(F)$.

According to Lemma 2.9, either $\text{Dom}(F) = \text{On}$, which we want, or $\text{Dom}(F) = \gamma \in \text{On}$. Suppose, for contradiction, that $\text{Dom}(F) = \gamma$. Then F is a set because $\text{Dom}(F)$ is a set, $\text{Rng}(F)$ is a set using the axiom of replacement, and $F \subseteq \text{Dom}(f) \times \text{Rng}(f)$. This implies that $F \in A$ because its domain is an ordinal, and we have verified that it satisfies the recursive definition property.

Now that $F \in A$, we define a slightly “longer” function $F_1 := F \cup \{(\gamma, G(F))\}$; note that $F = F_1 \upharpoonright \gamma$. Notice that $F_1 \in A$ because $\text{Dom}(F_1) = \gamma + 1$ is an ordinal, and we defined it to satisfy the recursive definition property. Because $F = \bigcup A$, this implies $F_1 \subseteq F$, but then $\gamma \in \text{Dom}(F_1) \subseteq \text{Dom}(F) = \gamma$, which is a contradiction. We conclude that $\text{Dom}(F) = \text{On}$.

Finally, we prove the uniqueness of F . For contradiction, suppose that there is another mapping $F' \neq F$ satisfying this theorem. Because $(\text{On}, <)$ is well-ordered, we can take the smallest ordinal α where $F(\alpha) \neq F'(\alpha)$. Therefore $F \upharpoonright \alpha = F' \upharpoonright \alpha$ and so $F(\alpha) = G(F \upharpoonright \alpha) = G(F' \upharpoonright \alpha) = F'(\alpha)$, which is a contradiction. \square

Exercise 2.18. Prove by induction on ω that every infinite well-ordered set A , such that each initial segment (\leftarrow, a) is finite, is isomorphic to $(\omega, <)$.

Hint. Since each (\leftarrow, a) is finite, there is a unique $n_a \in \omega$ with the same cardinality. The isomorphism we are looking for is $f : A \rightarrow \omega$ defined by $f : a \mapsto n_a$.

Exercise 2.19. Prove by transfinite recursion that every well-ordered proper class W , such that each proper lower part (\leftarrow, a) is a set, is isomorphic to $(\text{On}, <)$.

Hint. Use transfinite recursion to define an order-isomorphism $F : \text{On} \rightarrow W$ using $G(x) = \min(W \setminus x)$ and $F(\alpha) = G(F[\alpha])$. Then $F(0) = G(\emptyset) = \min(W)$, and $F(\alpha)$ is smallest element of W that has not yet been used to define $F(\beta)$ for some smaller $\beta < \alpha$.

We will use transfinite recursion to prove the equivalence of AC to the well-ordering principle and Zorn’s lemma. But transfinite recursion, together with AC, can also be used to prove some wild geometrical claims, such as:

- \mathbb{R}^3 is a union of pair-wise disjoint unit circles, or that
- there is a set in \mathbb{R}^2 that intersects every line in exactly two points.

2.4 Well-Ordering Principle

The *well-ordering principle*—the statement that every set can be well-ordered—was a foundational belief of Georg Cantor, but he was unable to provide a proof for it. This challenge was famously solved by Ernst Zermelo in 1904. Zermelo was the first person to explicitly state the axiom of choice, which he identified as the principle Cantor (and many others) had been implicitly using in many proofs. He then demonstrated that AC and the well-ordering principle are equivalent, which is why the principle is now often called the “Well-Ordering Theorem” or “Zermelo’s Theorem.” Veritasium has a great video [38] on this topic.

Principle 2.20 (Well-Ordering Principle). Every set can be well-ordered.

Theorem 2.21. *The well-ordering principle is equivalent to the axiom of choice.*

Proof. **WO** \Rightarrow **AC**. Let $A \neq \emptyset$ be a set, without loss of generality $\emptyset \notin A$. We want to construct a selector $f : A \rightarrow \bigcup A$ such that for all $a \in A$ we have $f(a) \in a$. The well-ordering principle guarantees a well-ordering \leq on $\bigcup A$, and because every a is a nonempty subset of $\bigcup A$, it has a least element with respect to \leq . We chose this minimum as $f(a)$.

AC \Rightarrow **WO**. Let $A \neq \emptyset$ be a set. We will use transfinite recursion to label the elements of A by ordinal numbers and then use the well-order of the ordinals to define a well-order on A . Let $g : \mathcal{P}(A) \rightarrow A$ be a selector on $\mathcal{P}(A)$, assigning to each nonempty $B \subseteq A$ an element $b \in B$. We will want to use transfinite recursion based on g , so we should extend it to be a class map $G : \mathbf{V} \rightarrow \mathbf{V}$, for example, by defining it to be equal to \emptyset when g is not defined.

We can now use transfinite recursion to define the function $F : \text{On} \rightarrow A \cup \{\emptyset\}$ as $F(0) = G(A)$ and $F(\alpha) = G(A \setminus F[\alpha])$. This function assigns to each ordinal a unique element from A until they “run out” (when $F[\alpha] = A$), and then it assigns \emptyset to all larger ordinals.

Define W as the class of all ordinals α for which $F[\alpha] \subsetneq A$. Denote the restriction of F to W as $F_W : W \rightarrow A$. Plan: first, we show that W itself is an ordinal. From this, it will follow that F_W is a bijection between W and A , allowing us to denote the unique ordinal mapped to $a \in A$ as α_a . Once this is established, we define a well-ordering R of A as

$$a <_R b \iff \alpha_a < \alpha_b.$$

This is a well-ordering since $(A, <_R)$ is order-isomorphic to $(W, <)$, which is well-ordered (as W is an ordinal).

Firstly, we claim that W is a set. Indeed, because F_W is injective, it has an inverse F_W^{-1} that maps the set $\text{Rng}(F_W) \subseteq A$ onto W , which is therefore, using the axiom of replacement, a set. Now we claim that W is a lower subset of On , and so it is an ordinal (by Lemma 2.9). Suppose $\alpha \in W$, that is $F[\alpha] \subsetneq A$, and let $\beta < \alpha$. Then $\beta \subseteq \alpha$ and $F[\beta] \subseteq F[\alpha]$, so $\beta \in W$.

To complete the proof, we must show that $F_W : W \rightarrow A$ is a bijection. It is clearly injective. To show that it is surjective, suppose for contradiction that there exists some $b \in A \setminus F_W[W]$. Because W is an ordinal number γ , it satisfies the definition of W (thanks to b) and thus $W = \gamma \in W$, which is a contradiction. \square

2.5 Zorn’s Lemma

Zorn’s lemma is perhaps the most useful application of the axiom of choice outside of set theory. It is also known as the *maximality principle*, a name that dates back to the German mathematician Felix Hausdorff, who proved an earlier and equivalent version of the theorem in 1914 (see [41] for details). The formulation known today as Zorn’s lemma was introduced in 1935 by the German mathematician Max Zorn. However, it had already been independently proved in 1922 by the Polish mathematician Kazimierz Kuratowski, whom you might know for Kuratowski’s theorem—a forbidden-graph characterization of planar graphs.

Definition 2.22 (Chain). Let (a, \leq_R) be an ordered set. We call a subset $b \subseteq a$ a *chain* if b is totally ordered by \leq_R .

Principle 2.23 (Zorn’s Lemma). Every (partially) ordered set containing upper bounds for every chain necessarily contains at least one maximal element.

There is also a parameterized version of this statement.

Principle 2.24 (Parametrized Zorn's Lemma). Let A be a (partially) ordered set containing upper bounds for every chain. Then for every $a \in A$, there is a maximal element $b \in A$ such that $a \leq b$.

We can obtain the parameterized version from the unparameterized one by restricting ourselves to the elements above or equal to a . The other direction is obvious.

Remark. Zorn's lemma can be made slightly stronger by assuming that only well-ordered chains have upper bounds. The proof remains virtually unchanged.

Theorem 2.25. *The axiom of choice implies Zorn's lemma.*

Proof. Let $(A, <_R)$ be an ordered set containing upper bounds for each chain and for contradiction suppose that there is no maximal element. Note that this implies that every chain, in fact, has a *strict* upper bound. If a chain C had no strict upper bound, then the non-strict upper bound $b \in C$ would be a maximal element. We denote the set of strict upper bounds of C as $C^>$.

We take $f : \mathcal{P}(A) \rightarrow A$, a selector on $\mathcal{P}(A)$, and define a function g from the set of all chains in A as $g(C) := f(C^>)$. So g maps a chain to one of its strict upper bounds. Now pick an arbitrary $a \in A$ and define the mapping $H : \text{On} \rightarrow A$ by transfinite recursion as $H(0) = a$ and $H(\alpha + 1) = g(\{H(\alpha)\})$ for successor ordinals, and as $H(\delta) = g(H[\delta])$ for limit ordinals. We start with a and get larger and larger elements of A using successor ordinals, each time taking a strict upper bound of a single element chain. If an ordinal δ is limit, we notice that $H[\delta]$ is a chain (all the smaller elements that we picked previously are strict upper bounds of each other and are therefore comparable), and $H(\delta)$ is a strict upper bound of this chain.

Note that if we want to be rigorous about the construction by transfinite recursion, we should define g on the entire \mathbf{V} . But we can do this in any way, for example, by defining $G(x)$ as \emptyset if x is not a chain of A , and $g(x)$ otherwise.

Finally, observe that $H : \text{On} \rightarrow A$ is an increasing function (each value is a strictly larger upper bound than the previous one) and that it is injective. Thus, we obtain an injection from the proper class On into the set A , which is impossible. Indeed, taking the inverse mapping and applying the axiom of replacement would imply that On itself is a set, which is a contradiction. \square

Theorem 2.26. *Zorn's lemma implies the well-ordering principle.*

Proof. Let X be any set. We will find a well-ordering of it by considering all of its possible well-ordered subsets, picking the maximal one using Zorn's lemma, and showing that it orders the entire X . Consider the set:⁴

$$\mathcal{W} := \{(A, <_R) \mid <_R \text{ is a well-order on } A \subseteq X\},$$

and define a partial order $\prec_{\mathcal{W}}$ on it by $(A, <_R) \prec_{\mathcal{W}} (B, <_S)$ if B end-extends A . That is, if $A \subset B$, and $<_R$ is the restriction of $<_S$ to A , and A is an initial segment of B . We will apply Zorn's lemma to \mathcal{W} .

⁴Why is this a set?

First, we need to show that chains have upper bounds. Let $\mathcal{C} \subseteq \mathcal{W}$ be a chain. Define the set

$$M := \bigcup \{A \mid (A, <_R) \in \mathcal{C}\} \subseteq X,$$

and for $x, y \in M$ put $x <_M y$ if there exists some $(A, <_R) \in \mathcal{C}$ such that $x, y \in A$ and $x <_R y$. Because \mathcal{C} is a chain, this is well-defined: if x and y belong to two distinct orderings in \mathcal{C} , then one extends the other and hence they agree.

We claim that $(M, <_M)$ is well-ordered. Let $S \subseteq M$ be nonempty and pick some $s \in S$. Then $s \in A_s$ for some $(A_s, <_R) \in \mathcal{C}$. Note that $A_s \cap S$ is nonempty, and because A_s is well-ordered, there exists a minimum $m = \min_{<_R}(A_s \cap S)$. Notice that also $m = \min_{<_M}(S)$. Indeed, if there were a $t \in S \setminus A_s$ such that $t <_M m$, then it would be in S due to some $A_t \in \mathcal{C}$ containing t . Since both A_s and A_t are in the chain, either

- (a) $A_t \subseteq A_s$, which is impossible since then $t \in A_s$, or
- (b) $A_s \subset A_t$, meaning that A_s is an initial segment of A_t , and therefore $m \in A_s$ is smaller than $t \in A_t \setminus A_s$, which contradicts the assumption that $t <_M m$.

Therefore $(M, <_M)$ is well-ordered and thus an upper bound of \mathcal{C} in \mathcal{W} .

Because all chains are bounded, by Zorn's lemma, \mathcal{W} has a maximal element $(W, <_W)$. We claim that $W = X$ and so it is the sought-after well-ordering of X . For contradiction, suppose there exists some $x \in X \setminus W$ and extend the ordering $<_W$ to $W' := W \cup \{x\}$ by making each $y \in W$ smaller than x . Notice that this slightly "longer" order is a well-ordering of W' and therefore is in \mathcal{W} . Moreover, it end-extends $(W, <_W)$ which hence is not maximal in $(\mathcal{W}, \prec_{\mathcal{W}})$. We have arrived at a contradiction and can conclude that $W = X$. \square

Exercise 2.27. Would the proof still have worked if instead of end-extensions, we had simply used general extensions? Meaning that the smaller ordering doesn't need to be an initial segment of the larger one.

Hint. By defining the end-extension ordering, we have ensured that chains have a similar structure to chains of ordinals (larger ordinals end-extend the smaller ones). Thus, when proving that M is well-ordered, we could have used a similar strategy as when proving that the ordinals are well-ordered.

To demonstrate an application of Zorn's lemma, consider the following question. Does every connected graph have a spanning tree? Finding one in a finite graph is easy: simply remove the edges of cycles until there are no cycles left. But this process may not terminate for infinite graphs.

Proposition 2.28. *Every connected graph has a spanning tree.*

Proof sketch. The set of all sub-graphs that are trees is partially ordered by inclusion, and the union of a chain is its upper bound. Zorn's lemma states that a maximal tree must exist, which is a spanning tree since the graph is connected. \square

Remark. In general, suppose that we have a structure represented by a set X (a graph) with substructures $A \subseteq X$ (subgraphs that are trees), and we want to show that there is a maximal substructure. Then we simply need to check that the union of a chain of substructures is itself a substructure.

2.6 Trichotomy Principle

Principle 2.29 (Trichotomy Principle). The relation \preceq is trichotomous on \mathbf{V} . That is, for any sets x and y either $x \preceq y$, or $y \preceq x$.

Trichotomy seems only natural: how “weird” would sets x and y really have to be to be completely incomparable? We have mentioned in Section 1.6 that there exist models of \mathbf{ZF} containing infinite but Dedekind-finite sets, and that such sets have no countably infinite subset. Therefore, if we let x be such a set, then $\omega \not\preceq x$, but also clearly $x \not\preceq \omega$, otherwise x would be Dedekind-infinite.

Theorem 2.30. *Zorn’s lemma implies the trichotomy principle.*

Proof. Let x, y be arbitrary sets; we want to find an injection $x \rightarrow y$ or $y \rightarrow x$. Consider the set⁵

$$\mathcal{F} = \{f \mid f \text{ is an injection, } \text{Dom}(f) \subseteq x \text{ and } \text{Rng}(f) \subseteq y\}.$$

Notice that the ordered set (\mathcal{F}, \subseteq) satisfies the conditions of Zorn’s lemma since the union of a chain of injections is again an injection. Let g be a maximal element of \mathcal{F} . If both $x \setminus \text{Dom}(g)$ and $y \setminus \text{Rng}(g)$ were non-empty, then it would be possible to extend g by an extra pair, contradicting its maximality. Hence either $\text{Dom}(g) = x$ and then $x \preceq y$, or $\text{Rng}(g) = y$ and then $y \preceq x$. Here, we used the fact that the inverse of an injection is also an injection. \square

Later, (Theorem 3.57), we will show that the trichotomy principle implies the well-ordering principle and thus also the axiom of choice.

Theorem 2.31. *We conclude that the following statements are equivalent in \mathbf{ZF} :*

- (1) *the axiom of choice,*
- (2) *the well-ordering principle,*
- (3) *Zorn’s lemma,*
- (4) *the trichotomy principle.*

3 Operations on Ordinals

3.1 Ordinal Functions

Definition 3.1 (Ordinal function). We say that a mapping F is an *ordinal function* if its domain is a lower part of On , that is $\text{Dom}(F) \in \text{On}$ or $\text{Dom}(F) = \text{On}$, and $\text{Rng}(F) \subseteq \text{On}$. We say that F is

- (a) *increasing* if for all $\beta \in \text{Dom}(F)$ and $\alpha < \beta$ we have $F(\alpha) < F(\beta)$, and
- (b) *non-decreasing* if for all $\beta \in \text{Dom}(F)$ and $\alpha < \beta$ we have $F(\alpha) \leq F(\beta)$.

Remark. We do not define decreasing ordinal functions as they would not be very interesting — the well-ordering of On does not allow infinite decreasing sequences (Observation 1.31), hence F can be decreasing only when $\text{Dom}(F)$ is finite.

⁵Why is this a set?

Lemma 3.2. *Increasing ordinal functions grow at least as fast as the identity function. That is $F(\alpha) \geq \alpha$ for every $\alpha \in \text{Dom}(F)$ for increasing F .*

Proof. For contradiction, suppose that α is the least ordinal such that $F(\alpha) < \alpha$. This means that for every $\beta < \alpha$, we have $F(\beta) \geq \beta$ (note that $\beta \in \text{Dom}(F)$). Suppose $\beta = F(\alpha)$; then $F(\beta) \geq \beta = F(\alpha)$, which is a contradiction since F is increasing. \square

Lemma 3.3. *If α and β are the ordinal types of the well-ordered sets A and $B \subseteq A$, then $\beta \leq \alpha$. In other words $B \subseteq A \implies \text{otp}(B) \leq \text{otp}(A)$.*

Note that $B \subset A$ does not imply that $\beta < \alpha$; consider ω and $\omega \setminus \{\emptyset\}$.

Proof. Let $i_a : A \rightarrow \alpha$ and $i_b : B \rightarrow \beta$ be the isomorphisms of A and B with their types. Suppose $\beta > \alpha$ and define $f : \beta \rightarrow \alpha$ as $f = i_b^{-1} \circ i_a$. Notice that if $\gamma < \delta$, then $f(\gamma) < f(\delta)$ because both i_b^{-1} and i_a preserve order (they are order-isomorphisms), and thus f is increasing. Because $\alpha \in \text{Dom}(f)$, we have that $f(\alpha) \in \text{Rng}(f) \subseteq \alpha$, so $f(\alpha) < \alpha$. But this contradicts the previous lemma. \square

Recall what you know about metric spaces, namely about closed sets and continuous functions. A subset X of a metric space M is closed if for each convergent sequence $(a_n) \subset X$ we have that $\lim a_n \in X$.

You might also recall that a function f is continuous \iff the preimage $f^{-1}[Y]$ of every closed set Y is closed \iff for every sequence (a_n) we have that $f(\lim a_n) = \lim f(a_n)$. The first equivalence statement is utilized in topology to define continuous functions, and we could use it here as well. However, the second equivalence seems more natural, since $\lambda = \sup\{\alpha \mid \alpha < \lambda\}$ for any limit ordinal λ . Hence limit ordinals essentially represent limits of sequences.

Lemma 3.4. *If $A \subseteq \text{On}$ is a set, then $\bigcup A \in \text{On}$, and in fact, $\bigcup A = \sup(A)$. We say that $\sup(A)$ is the limit of the sequence of ordinals A .*

Remark. We use the term “sequence” informally here. What we really mean is that we can label the elements of A as $A = \{\alpha_\delta \mid \delta < \gamma\}$ for some ordinal γ . We use $(\alpha_\delta)_{\delta < \gamma}$ to denote the bijection $\delta \mapsto \alpha_\delta$.

Definition 3.5 (Closed class). A subclass $C \subseteq \text{On}$ is *closed* if, for every subset $Y \subseteq C$, we have $\sup(Y) \in C$. For an ordinal α , we say that a subset $C \subseteq \alpha$ is *closed in α* if, for every $Y \subseteq C$ satisfying $\sup(Y) < \alpha$, we have $\sup(Y) \in C$.

Observation 3.6. *If C is a closed set, then it has a maximum $\max(C) = \sup(C)$.*

Example. The ordinal ω is not closed as $\sup(\omega) = \omega \notin \omega$, but $\omega + 1$ is closed as $\sup(\omega + 1) = \omega \in \omega + 1$, and for any $y \subset \omega + 1$ we have $\sup(y) \leq \sup(\omega + 1)$.

Observation 3.7. *An ordinal number α is closed \iff it is isolated.*

Definition 3.8 (Normal function). An ordinal function F is *continuous* if for every limit ordinal $\lambda \in \text{Dom}(F)$ it holds that

$$F(\lambda) = \sup\{F(\alpha) \mid \alpha < \lambda\}.$$

We say that a function is *normal* if it is increasing and continuous.

Example. The simplest normal function is identity. But consider the (very innocent looking) function $F(\alpha) = \alpha + 1$. It is increasing but not continuous. It fails on limit ordinals, for example $F(\omega) = \omega + 1$, but

$$\sup\{F(\alpha) \mid \alpha < \omega\} = \sup\{\alpha + 1 \mid \alpha < \omega\} = \sup(\omega \setminus \{\emptyset\}) = \omega.$$

Observation 3.9. *If A is a set of ordinals and $\sup(A)$ is isolated, then $\sup(A) \in A$, so $\max(A)$ exists. Hence if $\max(A)$ does not exist, then $\sup(A)$ is limit.*

Observation 3.10. *If F is an increasing ordinal function and λ is a limit ordinal, then $\sup\{F(\alpha) \mid \alpha < \lambda\}$ is also a limit ordinal. Specifically, if F is normal and λ limit, then $F(\lambda)$ is limit as well.*

Proof. Let $A = \{F(\alpha) \mid \alpha < \lambda\}$ and suppose that $\sup(A)$ is isolated. By the previous observation, there exists some $\alpha < \lambda$ such that $F(\alpha) = \max(A)$. But this is a contradiction because F is increasing, so $F(\alpha + 1) > F(\alpha)$ even though $\alpha + 1 < \lambda$ since λ is limit. \square

Exercise 3.11. Show that the *topological definition of continuity* would make sense. Prove that an increasing function F is continuous \iff the preimage $F^{-1}[C]$ of every closed set $C \subset \text{On}$ is closed in $\text{Dom}(F)$.

Observation 3.12. *The composition $F \circ G$ of normal functions F and G is a normal function. This can be seen easily from the previous exercise.*

Lemma 3.13. *If F is a normal function, then for every ordinal β , such that $F(0) \leq \beta < \sup \text{Rng}(F)$, the maximum $\max\{\alpha \mid F(\alpha) \leq \beta\}$ exists.*

Intuition. For a natural number β and $F(n) = n^2$, we might consider the largest natural number α such that $F(\alpha) \leq \beta$. This α exists, it is in fact equal to $\lfloor \sqrt{\beta} \rfloor$.

Proof. Notice that the set $[0, \beta] := \{\alpha \mid \alpha \leq \beta\}$ is closed by Observation 3.7 because $[0, \beta] = \beta + 1$ is an isolated ordinal. We will use the topological definition of continuity (Exercise 3.11) and note that the preimage C of the closed set $[0, \beta]$ is closed in $\text{Dom}(F)$. We would like to say that C is closed (in general). But consider $F : \omega \rightarrow \text{On}$; then $\text{Dom}(F)$ is not closed in On .

The bound on β will save us. Notice that there is some $\gamma \in \text{Rng}(F)$ such that $\beta < \gamma \leq \sup \text{Rng}(F)$. Because F is increasing, the elements of C are bounded by $F^{-1}(\gamma)$, and therefore $\sup(C) \in \text{Dom}(F)$. Since C is closed in $\text{Dom}(F)$, it follows that $\sup(C) \in C$ and $\sup(C) = \max(C)$, which we will denote as α . Because F is increasing, α is the largest ordinal satisfying $F(\alpha) \leq \beta$. \square

Lemma 3.14. *If $K \subseteq \text{On}$ is a closed proper class, then there exists a unique bijective normal ordinal function $J : \text{On} \rightarrow K$ enumerating the elements of K . Equivalently we can say that J is a continuous order-isomorphism of On and K .*

Proof. The proper class K inherits from On its well-order and the property that every proper lower part $(\leftarrow, a) \subset K$ is a set (see Lemma 2.9). Exercise 2.19 claims that there is a unique isomorphism (bijective increasing function) $J : \text{On} \rightarrow K$.

To show that J is continuous, let λ be a limit ordinal. Because it is the first ordinal larger than all $\alpha < \lambda$, it will be mapped to the first $\kappa \in K$ larger than all $J(\alpha)$ for $\alpha < \lambda$. Because K is closed, $\delta := \sup\{J(\alpha) \mid \alpha < \lambda\} \in K$. We claim that $\kappa = \delta$ (hence J is continuous). If not, then δ was already used by some $\beta < \lambda$, that is $J(\beta) = \delta$. Because J is increasing, we have for $\beta + 1 < \lambda$ that $J(\beta + 1) > J(\beta) = \delta$, so δ is not the supremum, a contradiction. \square

Definition 3.15 (Fixed point). We call an ordinal ξ a *fixed point* F if $F(\xi) = \xi$.

Theorem 3.16 (About fixed points). *Let $F : \text{On} \rightarrow \text{On}$ be a normal function.*

- (i) *For every $\alpha \in \text{On}$, there exists $\beta \geq \alpha$ that is a fixed point of F .*
- (ii) *The first fixed point $\beta \geq \alpha$ is the limit (supremum) of the sequence $(\alpha_n)_{n < \omega}$ defined as $\alpha_0 = \alpha$ and $\alpha_{n+1} = F(\alpha_n)$.*
- (iii) *The class K of all fixed points of F is a closed proper class.*
- (iv) *There exists a unique bijective normal function $F' : \text{On} \rightarrow K$ called the derivative of F , which enumerates the fixed points of F .*

Proof. First, notice that (iv) is a direct consequence of (iii) by Lemma 3.14. To prove (i) and (ii), notice that $\alpha_{n+1} \geq \alpha_n$ since F grows at least as fast as the identity function, and that the supremum $\beta = \sup\{\alpha_n \mid n \in \omega\}$ is a fixed point:

- Consider the case when $\alpha_0 < \alpha_1 < \dots < \alpha_i = \alpha_{i+1}$ for some i ; then also $\alpha_{i+2} = F(\alpha_{i+1}) = F(\alpha_i) = \alpha_i$ and by induction $\alpha_n = \alpha_i$ for all $n \geq i$, thus $\beta = \alpha_i$ is a fixed point.
- Suppose the sequence never stabilizes; then Observation 3.9 implies that β is limit. Since F is continuous and non-decreasing, we have

$$F(\beta) = \sup\{F(\gamma) \mid \gamma < \beta\} = \sup\{F(\alpha_n) \mid n \in \omega\} = \sup\{\alpha_{n+1} \mid n \in \omega\} = \beta.$$

Second, we show that β is the smallest fixed point larger than α . If there were a fixed point $\xi \geq \alpha$ such that $\xi < \beta$, then there would exist an index n at which $\alpha_n \leq \xi < \alpha_{n+1}$. This is because the sequence is strictly increasing, and β is its supremum. We have $\xi < \alpha_{n+1} = F(\alpha_n) \leq F(\xi)$, so ξ is not a fixed point.

Finally, we prove (iii). We claim that K is closed. Let $C \subseteq K$ be a set; we need to show that the supremum $\beta = \sup(C)$ is a fixed point. Since F is continuous and non-decreasing, we have

$$F(\beta) = \sup\{F(\gamma) \mid \gamma < \beta\} = \sup\{F(\xi) \mid \xi \in C\} = \sup\{\xi \mid \xi \in C\} = \beta.$$

To complete the proof, we show that K is a proper class. If it were a set, then by Lemma 3.4, $\sup(K)$ would be an ordinal γ . We let the ordinal $\gamma + 1$ take the role of α in (i) and find a new fixed point of F , larger than all those in K . \square

Theorem 3.17 (About simultaneous fixed points). *Let $\langle F_i \mid i \in I \rangle$ be a collection of normal functions $F_i : \text{On} \rightarrow \text{On}$ indexed by a set I . We say that ξ is a simultaneous fixed point of $\langle F_i \mid i \in I \rangle$ if $F_i(\xi) = \xi$ for all $i \in I$.*

- (i) *For every $\alpha \in \text{On}$, there is $\beta \geq \alpha$, a simultaneous fixed point of $\langle F_i \mid i \in I \rangle$.*
- (ii) *The first simultaneous fixed point $\beta \geq \alpha$ is the limit of the sequence $(\alpha_n)_{n < \omega}$ defined as $\alpha_0 = \alpha$ and $\alpha_{n+1} = \sup\{F_i(\alpha_n) \mid i \in I\}$.*
- (iii) *The class K of all simultaneous fixed points is a closed proper class.*
- (iv) *There exists a unique bijective normal function $J : \text{On} \rightarrow K$ enumerating the simultaneous fixed points of $\langle F_i \mid i \in I \rangle$.*

Proof. Define an ordinal function $\overline{F} : \text{On} \rightarrow \text{On}$ as $\overline{F}(\alpha) := \sup\{F_i(\alpha) \mid i \in I\}$. This is well defined since we can use replacement to guarantee that $\{F_i(\alpha) \mid i \in I\}$ is a set. Notice that ξ is a fixed point of \overline{F} if and only if it is a simultaneous fixed point of $\langle F_i \mid i \in I \rangle$ since

$$\overline{F}(\xi) = \xi \implies (\forall i \in I) F_i(\xi) \leq \xi \iff (\forall i \in I) F_i(\xi) = \xi \implies \overline{F}(\xi) = \xi. \quad (3.1)$$

The equivalence holds because all F_i grow at least as fast as the identity function. Note that this also implies that \overline{F} itself grows at least as fast as the identity function. Moreover, \overline{F} is continuous. Indeed, if λ is a limit ordinal, then

$$\begin{aligned} \overline{F}(\lambda) &= \sup\{F_i(\lambda) \mid i \in I\} \\ &= \sup\{\sup\{F_i(\delta) \mid \delta < \lambda\} \mid i \in I\} && \dots \text{each } F_i \text{ is continuous} \\ &= \sup\{\sup\{F_i(\delta) \mid i \in I\} \mid \delta < \lambda\} \\ &= \sup\{\overline{F}(\delta) \mid \delta < \lambda\}. \end{aligned}$$

In general, \overline{F} is not normal, as it is not guaranteed to be increasing. However, notice that the proof of Theorem 3.16 only uses the fact that F is continuous, non-decreasing, and that it grows at least as fast as the identity. Hence, we can use the theorem on \overline{F} to obtain the claims (i)–(iv). \square

Proposition 3.18. *If $F, G : \text{On} \rightarrow \text{On}$ are normal functions, and $F(\alpha) \geq G(\alpha)$ for all α , then every fixed point of F is also a fixed point of G .*

Proof. This is a simple version of (3.1). \square

3.2 Ordinal Arithmetic

In this section we define operations such as ordinal addition and multiplication. Before proceeding further, I highly recommend watching the video [39] by Vsauce, which illustrates the concepts of constructing larger ordinals from earlier ones in a very illustrative and intuitive way.

3.2.1 Definitions and Intuition

Definition 3.19. Let α and β be ordinals. We define ordinal numbers

- (a) $\alpha + \beta$ as the order type of the set $(\{0\} \times \alpha) \cup (\{1\} \times \beta)$ when ordered lexicographically,
- (b) $\alpha \cdot \beta$ as the order type of the set $\beta \times \alpha$ when ordered lexicographically.

Using the popular “matchstick” representation of ordinals, $\alpha + \beta$ can be imagined as a pile of decreasing matchsticks labeled by α , followed by another pile of matchsticks labeled by β . Notice that our previous notation of denoting $\alpha \cup \{\alpha\}$ by $\alpha + 1$ is consistent with the above definition. We first use the elements of α to label the first pile, and we need one additional ordinal to label the second pile (which contains only a single matchstick).

Notice that we are using $\beta \times \alpha$ in the definition of $\alpha \cdot \beta$. The ordinal $\alpha \cdot \beta$ can be imagined as taking multiple piles of matchsticks labeled by α and arranging

them next to each other. How should the piles be arranged? In a way that we need β to label them.

With this intuition, it should not be surprising that ordinal addition and multiplication are generally not commutative. It is easy to see that $1 + \omega = \omega$ (label the first pile by 0 and the other pile by $\omega \setminus \{0\}$), but $\omega + 1 \neq \omega$. For multiplication, consider $2 \cdot \omega$, the order type of countably infinitely many copies of $\{0, 1\}$ stacked behind each other. This can be clearly labeled by ω , so $2 \cdot \omega = \omega$. But $\omega \cdot 2$ is the order type of two consecutive copies of ω . When we try to label them using ω , we use all $n \in \omega$ to label the first copy and need more ordinals for the second copy. Therefore $\omega \cdot 2 > \omega$.

Observation 3.20. For any ordinals α, β, γ and natural $n \in \omega$ it holds that

- (a) $\alpha + 0 = \alpha = 0 + \alpha, \quad \alpha \cdot 0 = 0 = 0 \cdot \alpha, \quad \alpha \cdot 1 = \alpha = 1 \cdot \alpha,$
- (b) $\alpha + (\beta + \gamma) = (\alpha + \beta) + \gamma, \quad \alpha \cdot (\beta \cdot \gamma) = (\alpha \cdot \beta) \cdot \gamma,$
- (c) $\alpha \cdot 2 = \alpha + \alpha, \quad \alpha \cdot 3 = \alpha + \alpha + \alpha, \quad \alpha \cdot (n + 1) = \alpha \cdot n + \alpha.$

Definition 3.21. For ordinal numbers α and β , we define α^β recursively as

- (i) $\alpha^0 := 1,$
- (ii) if $\beta = \gamma + 1$ is isolated, then $\alpha^\beta := \alpha^\gamma \cdot \alpha,$
- (iii) if β is a limit ordinal, then $\alpha^\beta := \sup\{\alpha^\gamma \mid 0 < \gamma < \beta\}.$

Remark. Formally, for every fixed α we use transfinite recursion to define an ordinal function $F_\alpha : \beta \mapsto \alpha^\beta$.

To get an intuition for ordinal powers, consider the ordinal $\omega^2 = \omega \cdot \omega$. It represents multiple copies of ω arranged in the same manner as ω .




Figure 2: A representation of the ordinal ω^2 . Each stick corresponds to an ordinal of the form $\omega \cdot m + n$ where m and n are natural numbers [17].

To construct $\omega^3 = (\omega \cdot \omega) \cdot \omega$, we take multiple copies of ω^2 and arrange them in a way that requires ω to label them. If we repeat this process countably infinitely many times, we arrive at ω^ω .

We can continue and arrive at larger and larger ordinals, such as $\omega^{(\omega^\omega)}$ or $\omega^{\omega^{(\omega^\omega)}}$. That is a lot of parentheses, so from now on, we will write ω^{ω^ω} instead of $\omega^{(\omega^\omega)}$ and use parentheses only when we mean to say $(\omega^\omega)^\omega$.




Figure 3: A spiral representation of ordinals up to ω^ω . One full turn corresponds to the mapping $f(\alpha) = \omega \cdot (1 + \alpha)$. Since ω^ω is the smallest fixed point of f , larger ordinals cannot be represented in this way [6].

Observation 3.22. *For any ordinals α and $\beta > 0$ it holds that*

- (a) $0^0 = 1, \quad 0^\beta = 0,$
- (b) $1^0 = 1, \quad 1^\beta = 1,$
- (c) $\alpha^0 = 1, \quad \alpha^1 = \alpha, \quad \alpha^2 = \alpha \cdot \alpha, \quad \alpha^3 = (\alpha \cdot \alpha) \cdot \alpha.$

3.2.2 Properties of Ordinal Operations

You should now have an intuition for how ordinal numbers constructed using these standard operations look. We continue by proving some of their basic properties.

Lemma 3.23 (Monotonicity of sum). *For any α, β and γ it holds that*

- (a) $\alpha < \beta \implies \gamma + \alpha < \gamma + \beta,$
- (b) $\alpha < \beta \implies \alpha + \gamma \leq \beta + \gamma.$

Proof. (a) From the definition of addition and order types, it is easy to see that $\gamma + \alpha$ is an initial segment of $\gamma + \beta$. (b) The set of ordered pairs that defines $\alpha + \gamma$ is a subset of the set of ordered pairs that defines $\beta + \gamma$. And Lemma 3.3 states that the order type (under lexicographic order) of the first is at most that of the second. \square

Lemma 3.24 (Monotonicity of product). *For any α, β and $\gamma > 0$ it holds that*

- (a) $\alpha < \beta \implies \gamma \cdot \alpha < \gamma \cdot \beta,$
- (b) $\alpha < \beta \implies \alpha \cdot \gamma \leq \beta \cdot \gamma. \quad \dots \text{for } \gamma = 0 \text{ also holds}$

Proof. (a) If $\alpha < \beta$, then $\alpha \times \gamma$ is an initial segment of $\beta \times \gamma$ when ordered lexicographically. (b) If $\alpha < \beta$, then $\gamma \times \alpha \subseteq \gamma \times \beta$, and the claim follows from Lemma 3.3. \square

Remark. Note that the second statement in the two preceding lemmas does not, in general, hold under strict inequality. For example, $1 < 2$, but

$$1 + \omega = 2 + \omega = \omega, \quad \text{and} \quad 1 \cdot \omega = 2 \cdot \omega = \omega.$$

In fact, for any natural $n \in \omega$, we have that $n + \omega = \omega$ and $n \cdot \omega = \omega$.

Lemma 3.25 (Distributivity). *For any ordinals α and β_1, β_2 , we have*

$$\alpha \cdot (\beta_1 + \beta_2) = \alpha \cdot \beta_1 + \alpha \cdot \beta_2.$$

That is, ordinal addition and multiplication are left-distributive. However, in general, they are not right-distributive. Meaning that for some α and β_1, β_2

$$(\beta_1 + \beta_2) \cdot \alpha \neq \beta_1 \cdot \alpha + \beta_2 \cdot \alpha$$

Proof. Left distributivity essentially states that if we arrange $\beta_1 + \beta_2$ copies of α next to each other, then it is the same as first arranging β_1 copies of α , followed by β_2 copies of α . This is obviously true from how we defined addition and multiplication. However, in general, these operations are not right-distributive. Consider $(1+1) \cdot \omega \neq \omega + \omega$. On the left, we have $2 \cdot \omega$, and on the right, $\omega \cdot 2$. \square

Theorem 3.26. *If m, n and k are natural numbers, then $m + n$, $m \cdot n$, and m^n are also natural numbers. Furthermore*

$$m + n = n + m, \quad m \cdot n = n \cdot m, \quad (m + n) \cdot k = m \cdot k + n \cdot k.$$

That is: addition and multiplication of natural numbers is commutative and right-distributive.

Proof. It is easy to see that $m + n$ and $m \cdot n$ are finite ordinals, and one can use induction on n to show that m^n is also finite.

To show that $m \cdot n = n \cdot m$, take any bijection $f : n \times m \rightarrow m \times n$ and use it along with the lexicographic order on $n \times m$ to define an isomorphic linear order $<_f$ on $m \times n$. It is well known (and it is proved in the basic set theory course) that any two linear orders on a finite set are isomorphic, so $<_f$ is isomorphic to the lexicographic order of $m \times n$. Since isomorphisms are transitive, the two lexicographically ordered sets $n \times m$ and $m \times n$ are isomorphic and thus have the same order types. Hence $m \cdot n = n \cdot m$. One can similarly show that addition commutes as well.

Right-distributivity is implied by commutativity and left-distributivity. \square

Lemma 3.27 (Existence of the right difference). *If $\alpha \leq \beta$, then there is a unique ordinal ϱ such that $\alpha + \varrho = \beta$. We denote ϱ by $\beta \dot{-} \alpha$.*

Intuition. Any ordinal can be extended by a specific amount to reach any larger ordinal. Furthermore, when restricted to natural numbers, $a \dot{-} b$ is the standard subtraction operation.

Proof. If $\alpha \leq \beta$, then $\alpha = (\leftarrow, \alpha)$ is an initial segment of β , and its complement, $\beta \setminus \alpha$, is what we might denote as $[\alpha, \rightarrow)$. If ϱ is the order type of $\beta \setminus \alpha$, then clearly $\alpha + \varrho = \beta$. The uniqueness of ϱ follows from Lemma 3.23 (a): suppose there were $\varrho_1 < \varrho_2$ satisfying $\alpha + \varrho_1 = \beta = \alpha + \varrho_2$. But since $\varrho_1 < \varrho_2$, we have that $\alpha + \varrho_1 < \alpha + \varrho_2$. \square

Lemma 3.28 (Division with remainder). *If $\beta > 0$, then for every ordinal α there are unique ordinals $\delta \leq \alpha$ and $\varrho < \beta$ such that $\alpha = \beta \cdot \delta + \varrho$.*

Intuition. Any ordinal α can be created by arranging multiple copies of β in a specific way, and following this with a short tail ϱ .

Proof. Since $\beta \geq 1$, we have $\alpha \leq \beta \cdot \alpha$. If $\alpha = \beta \cdot \alpha$ (for example $\omega = 3 \cdot \omega$), choose $\delta := \alpha$ and $\varrho := 0$. It is not hard to show that the monotonicity of sum and product, together with left-distributivity, implies uniqueness; we will skip it.

If $\alpha < \beta \cdot \alpha$, let j be the isomorphism of the lexicographically ordered set $\alpha \times \beta$ and the ordinal $\beta \cdot \alpha$. Let $(\delta, \varrho) \in \alpha \times \beta$ (so $\delta < \alpha$ and $\varrho < \beta$) be the (unique) pair mapped by j to $\alpha \in \beta \cdot \alpha$. Since $\alpha \times \beta$ is ordered lexicographically, it is easy to see that $\alpha = \beta \cdot \delta + \varrho$. \square

Lemma 3.29 (Monotonicity of power). *For any α, β, γ and $\rho > 1$, it holds that*

$$(a) \quad \alpha < \beta \implies \alpha^\gamma \leq \beta^\gamma,$$

$$(b) \quad \alpha < \beta \implies \rho^\alpha < \rho^\beta.$$

Proof. (a) Using transfinite induction on γ . If $\gamma = 0$, then $\alpha^\gamma = \beta^\gamma = 1$. If $\gamma = \delta + 1$ and $\alpha^\delta \leq \beta^\delta$, from the monotonicity of product we have that

$$\alpha^\gamma = \alpha^\delta \cdot \alpha \leq \beta^\delta \cdot \beta = \beta^\gamma.$$

If γ is a limit ordinal and for every $\delta < \gamma$ already $\alpha^\delta \leq \beta^\delta$, then also

$$\alpha^\gamma = \sup\{\alpha^\delta \mid 0 < \delta < \gamma\} \leq \sup\{\beta^\delta \mid 0 < \delta < \gamma\} = \beta^\gamma.$$

(b) Suppose that $\rho > 1$. It is easy to show using transfinite induction on δ that for every $\delta > 1$ it holds that $\rho^\alpha < \rho^{\alpha+\delta}$. If $\alpha < \beta$, then according to Lemma 3.27 there is a unique $\delta > 0$ satisfying $\beta = \alpha + \delta$. \square

Remark. Note that the first statement in the previous lemma does not, in general, hold under the strict inequality, even if $\gamma > 0$. For example, $2 < 3$, but $2^\omega = 3^\omega = \omega$. In general, if $n \in \omega$, then $n^\omega = \omega$.

Lemma 3.30 (Continuity in the second argument). *The ordinal function*

$$(a) \quad F(\xi) = \alpha + \xi \text{ is normal for every } \alpha \geq 0,$$

$$(b) \quad F(\xi) = \alpha \cdot \xi \text{ is normal for every } \alpha > 0,$$

$$(c) \quad F(\xi) = \alpha^\xi \text{ is normal for every } \alpha > 1.$$

Proof. All of the functions mentioned above are increasing for the specified α , since the respective operations are monotonic. We claim that they are also continuous; that is $F(\lambda) = \sup\{F(\xi) \mid \xi < \lambda\}$ for limit λ . Notice that this holds for (c) as this is simply the definition of α^λ .

Because F is increasing, $F(\lambda)$ is an upper bound of the values $F(\xi)$ for $\xi < \lambda$. Assume for contradiction that $F(\lambda)$ is not the smallest upper bound, meaning that there exists $\sigma < F(\lambda)$ such that $\sigma \geq F(\xi)$ for all $\xi < \lambda$.

(a) Lemma 3.27 claims that there is a unique ϱ such that $\sigma = \alpha + \varrho = F(\varrho)$. Because $\sigma < F(\lambda)$ we have that $\alpha + \varrho < \alpha + \lambda$ thus by monotonicity $\varrho < \lambda$ and also $\varrho + 1 < \lambda$. It should hold that $\sigma \geq F(\varrho + 1)$, but $\sigma = F(\varrho) < F(\varrho + 1)$.

(b) Lemma 3.28 claims the existence of unique ordinals δ and $\varrho < \alpha$ such that $\sigma = \alpha \cdot \delta + \varrho$. Since $\sigma < \alpha \cdot \lambda$ we have from monotonicity that $\delta < \lambda$ and also $\delta + 1 < \lambda$. It should hold that $\sigma \geq F(\delta + 1)$, but from monotonicity we have

$$\sigma = \alpha \cdot \delta + \varrho < \alpha \cdot \delta + \alpha = \alpha \cdot (\delta + 1) = F(\delta + 1). \quad \square$$

Lemma 3.31 (Properties of exponents). *For any α, β and γ , it holds that*

(a) $\alpha^{\beta+\gamma} = \alpha^\beta \cdot \alpha^\gamma$,

(b) $(\alpha^\beta)^\gamma = \alpha^{\beta \cdot \gamma}$

Proof. (a) Trivially holds for $\alpha \leq 1$. Suppose $\alpha > 1$, we will use transfinite induction on γ . If $\gamma = 0$, there is nothing to prove. If $\gamma = \delta + 1$ is isolated, then $\beta + \gamma = (\beta + \delta) + 1$, and the statement follows from the induction hypothesis:

$$\alpha^{(\beta+\delta)+1} = \alpha^{(\beta+\delta)} \cdot \alpha^1 = \alpha^\beta \cdot \alpha^\delta \cdot \alpha = \alpha^\beta \cdot \alpha^{\delta+1} = \alpha^\beta \cdot \alpha^\gamma.$$

Finally, if γ is a limit ordinal and the statement holds for all $\delta < \gamma$, then

$$\begin{aligned} \alpha^\beta \cdot \alpha^\gamma &= \sup\{\alpha^\beta \cdot \xi \mid \xi < \alpha^\gamma\} && \dots F(\xi) = \alpha^\beta \cdot \xi \text{ is normal} \\ &= \sup\{\alpha^\beta \cdot \alpha^\delta \mid 0 < \delta < \gamma\} && \dots \alpha^\gamma = \sup\{\alpha^\delta \mid 0 < \delta < \gamma\} \\ &= \sup\{\alpha^{\beta+\delta} \mid 0 < \delta < \gamma\} && \dots \text{induction hypothesis} \\ &= \sup\{\alpha^\varepsilon \mid 0 < \varepsilon < \beta + \gamma\} && \dots F(\xi) = \beta + \xi \text{ is normal} \\ &= \alpha^{\beta+\gamma}. \end{aligned}$$

The last equality holds because $\beta + \gamma$ is a limit ordinal from Observation 3.10.

(b) Suppose that $\beta, \gamma \neq 0$ and $\alpha > 1$, otherwise, it trivially holds. We will again use transfinite induction on γ . If $\gamma = 0$, then it holds. If $\gamma = \delta + 1$ is isolated, then the statement follows from (a) and the induction hypothesis:

$$(\alpha^\beta)^{\delta+1} = (\alpha^\beta)^\delta \cdot (\alpha^\beta)^1 = \alpha^{\beta \cdot \delta} \cdot \alpha^\beta = \alpha^{\beta \cdot \delta + \beta} = \alpha^{\beta \cdot (\delta+1)} = \alpha^{\beta \cdot \gamma}.$$

Finally, if γ is a limit ordinal, then $\beta \cdot \gamma$ is also limit and

$$\begin{aligned} (\alpha^\beta)^\gamma &= \sup\{(\alpha^\beta)^\delta \mid 0 < \delta < \gamma\} && \dots \gamma \text{ is limit} \\ &= \sup\{\alpha^{\beta \cdot \delta} \mid 0 < \delta < \gamma\} && \dots \text{induction hypothesis} \\ &= \sup\{\alpha^\varepsilon \mid 0 < \varepsilon < \beta \cdot \gamma\} && \dots F(\xi) = \beta \cdot \xi \text{ is normal} \\ &= \alpha^{\beta \cdot \gamma}. \end{aligned}$$

The last equality holds because $\beta \cdot \gamma$ is a limit ordinal. \square

3.2.3 Ordinal Equations and Power Expansions

Example. Suppose we want to find all ξ and β satisfying $\xi + \beta = \omega$. Lemma 3.27 claims that $\xi \leq \omega$ and $\beta = \omega \dot{-} \xi$. Suppose $\xi = \omega$, then $\beta = 0$. If $\xi = n$ is a natural number, then $\beta = \omega \dot{-} n = \omega$. We conclude that β can attain only two different values.

Proposition 3.32. *Let α be an ordinal and consider the equation $\xi + \beta = \alpha$. The set of solutions (ξ, β) contains only finitely many distinct values of β .*

Proof. Suppose that for some α , there are infinitely many distinct values of β in the solution set. Let $(\xi_n, \beta_n)_{n \in \omega}$ be a sequence of solutions such that $\beta_n < \beta_{n+1}$ for all n . Since $\xi_n + \beta_n = \xi_{n+1} + \beta_{n+1}$, from the monotonicity of sum we have that $\xi_n > \xi_{n+1}$ for all $n \in \omega$. We have constructed an infinite strictly decreasing sequence, which is impossible since On is well-ordered. \square

We are able to express any natural number n as an expansion of powers of any base $b > 1$. We will prove that a similar statement holds for ordinal numbers too. A base of special importance is ω (as it is the first transfinite ordinal), and the expansion of α over ω is called its *Cantor normal form*; however, an expansion is possible over any base $\beta > 1$.

Lemma 3.33. *If k, m_0, m_1, \dots, m_k are natural numbers and $\delta > \gamma_0, \gamma_1, \dots, \gamma_k$ are ordinals, then*

$$\omega^\delta > \omega^{\gamma_0} \cdot m_0 + \omega^{\gamma_1} \cdot m_1 + \dots + \omega^{\gamma_k} \cdot m_k.$$

Proof. Let m be the largest among all m_i , and γ be the largest among all γ_i . Then $\omega^\gamma \cdot m \cdot k$ is an upper bound of the sum on the right side of the equation. We assumed that $\delta \geq \gamma + 1$, so $\omega^\delta \geq \omega^{\gamma+1} > \omega^\gamma \cdot m \cdot k$. \square

Theorem 3.34 (Expansion over ω). *For any $\alpha > 0$ there are unique natural numbers $k, m_0, m_1, \dots, m_k \neq 0$ and ordinals $\gamma_0 > \gamma_1 > \dots > \gamma_k$ which satisfy*

$$\alpha = \omega^{\gamma_0} \cdot m_0 + \omega^{\gamma_1} \cdot m_1 + \dots + \omega^{\gamma_k} \cdot m_k. \quad (3.2)$$

The sum on the right side of the equation is called the Cantor normal form of α . Furthermore, if

$$\beta = \omega^{\delta_0} \cdot n_0 + \omega^{\delta_1} \cdot n_1 + \dots + \omega^{\delta_l} \cdot n_l \quad (3.3)$$

is the Cantor normal form of an ordinal β , then $\beta > \alpha$ if and only if one of the two following cases occurs:

- (a) $l > k$ and the first k terms of β are identical to those of α . For example: $\alpha = \omega^2 + \omega \cdot 2$ and $\beta = \omega^2 + \omega \cdot 2 + 3$.
- (b) there exists an index $i \leq \min(k, l)$ at which (γ_i, m_i) and (δ_i, n_i) differ, and for the smallest such index i either $\delta_i > \gamma_i$, or $\delta_i = \gamma_i$ and $n_i > m_i$. For example $\alpha = \omega^2 + \omega \cdot 2$ and $\beta = \omega^2 \cdot 5 + \omega \cdot 2$.

Proof. We prove the first part by transfinite induction on α . The CNF of $\alpha = 1$ is $\alpha = \omega^0 \cdot 1$. Suppose $\alpha > 1$ and that every nonzero $\beta < \alpha$ has a unique CNF. The ordinal function $\gamma \mapsto \omega^\gamma$ is normal, so according to Lemma 3.13, there exists a maximal ordinal γ such that $\omega^\gamma \leq \alpha$. Similarly, from the normality of product in the second argument follows the existence of a maximal ordinal δ such that $\omega^\gamma \cdot \delta \leq \alpha$. Also, $\delta < \omega$, since $\omega^{\gamma+1} = \omega^\gamma \cdot \omega > \alpha$, which contradicts the choice of γ . If $\omega^\gamma \cdot \delta = \alpha$, then the uniqueness of this expansion follows from Lemma 3.33.

If $\omega^\gamma \cdot \delta < \alpha$, then there exists a unique ordinal $\beta = \alpha \dot{-} \omega^\gamma \cdot \delta$ such that $\omega^\gamma \cdot \delta + \beta = \alpha$. Note that $\beta < \omega^\gamma$; otherwise, we get $\omega^\gamma \cdot \delta + \beta \geq \omega^\gamma \cdot (\delta + 1)$, which contradicts the choice of δ . To find the CNF of α , let

$$\beta = \omega^{\gamma_1} \cdot m_1 + \omega^{\gamma_2} \cdot m_2 + \cdots + \omega^{\gamma_k} \cdot m_k$$

be the CNF of β . Define $\gamma_0 := \gamma$ and $m_0 := \delta$. Then $\gamma_0 > \gamma_1$, and (3.2) is the CNF of α . The uniqueness of this expansion follows from Lemma 3.33 and the unique choice of β .

Next, we prove the second part of the theorem. Suppose that the ordinals α and β have Cantor normal forms (3.2) and (3.3). If (a) holds, then $\beta > \alpha$ because the trailing terms in the expansion of β are nonzero. Suppose that (b) holds and that i is the least index at which the two expansions differ. If $\delta_i > \gamma_i$, then $\beta > \alpha$ from Lemma 3.33. If $\delta_i = \gamma_i$ and $n_i > m_i$, then $n_i \geq m_i + 1$ and

$$\omega^{\delta_i} \cdot n_i \geq \omega^{\gamma_i} \cdot m_i + \omega^{\gamma_i}.$$

Lemma 3.33 claims that the second summand on the right (ω^{γ_i}) is a strict upper bound of the remaining summands in (3.2), the expansion of α ; thus $\beta > \alpha$.

All that remains is to prove the reverse implication. If $\beta > \alpha$, then their Cantor normal forms (3.2) and (3.3) have to differ. Either the CNF of one of the ordinals is the same as the beginning of the CNF of the other, or there exists an index at which they differ. We can use the already proven implication to show that the only two possible cases are (a) and (b). \square

Corollary 3.35 (Alternative expansions). *For any $\alpha > 0$, it holds that*

- (a) *there is a unique natural number $l > 0$ and unique ordinals $\gamma_0 \geq \gamma_1 \geq \cdots \geq \gamma_l$ which satisfy*

$$\alpha = \omega^{\gamma_0} + \omega^{\gamma_1} + \cdots + \omega^{\gamma_l},$$

- (b) *there are unique ordinals β and γ such that*

$$\alpha = \omega^\gamma \cdot (\beta + 1).$$

Proof. (a) For any ordinal γ and natural m , is the ordinal number $\omega^\gamma \cdot m$ equal to the sum of m summands of the form ω^γ . We obtain the expansion in (a) by expressing each term in the CNF of α in this expanded form.

(b) If α has CNF (3.2), we let $\gamma = \gamma_k$. Then for all $i \leq k$ is $\gamma_i = \gamma + \delta_i$ for $\delta_i = \gamma_i \dot{-} \gamma$. From the properties of exponents and left-distributivity, we get

$$\alpha = \omega^\gamma \cdot (\omega^{\delta_0} \cdot m_0 + \omega^{\delta_1} \cdot m_1 + \cdots + \omega^0 \cdot m_k).$$

The parentheses on the right contain an isolated ordinal $\beta + 1$, because m_k is a nonzero natural number. The uniqueness of the ordinals γ and β follows from the uniqueness of the CNF of α . \square

Theorem 3.36 (Expansion over any base). *The choice of ω as a base in Theorem 3.34 was arbitrary; the same holds for any ordinal base $\beta > 1$. We just need to restrict the nonzero coefficients m_0, m_1, \dots, m_k to be smaller than β .*

Proof. We did not use any special properties of ω in the proof of Theorem 3.34, so we only need to modify Lemma 3.33. If we slightly change its claim, only for decreasing exponents $\gamma_0 > \gamma_1 > \dots > \gamma_k$ and coefficients $m_i < \beta$, we can prove it by transfinite induction on γ_0 . \square

Remark. If we restrict ourselves only to natural numbers, we obtain the familiar theorem about expanding natural numbers using powers of a base $b > 1$.

3.3 Countable and Uncountable Ordinals

We saw earlier that for all natural numbers n , it holds that $n + \omega = n \cdot \omega = n^\omega = \omega$. We also proved that the functions corresponding to these basic operations,

$$A_n(\xi) = n + \xi, \quad M_n(\xi) = n \cdot \xi, \quad E_n(\xi) = n^\xi,$$

are normal. Theorem 3.16 claims that each of them has infinitely many fixed points. It is easy to see that no (nonzero) natural number is a fixed point, and above we have observed that ω is a fixed point of all of them. It is, in fact, the smallest (nonzero) fixed point. Notice that this makes intuitive sense. When restricted to natural numbers, these are all fast growing functions ($A \ll M \ll E$), so we need a new concept (countable infinity) to find a fixed point.

Now consider what would happen if we replaced n with ω and tried to find a (nonzero) fixed point of these new ω -functions. Theorem 3.16 claims that the smallest such fixed points are:

- $F_A = \sup\{0, \omega, \omega + \omega, \omega + \omega + \omega, \omega \cdot 4, \omega \cdot 5, \dots\} = \omega \cdot \omega$,
- $F_M = \sup\{1, \omega, \omega \cdot \omega, \omega \cdot \omega \cdot \omega, \omega^4, \omega^5, \dots\} = \omega^\omega$,
- $F_E = \sup\{1, \omega, \omega^\omega, \omega^{\omega^\omega}, \omega^{\omega^{\omega^\omega}}, \dots\}$, and we no longer have notation to describe this number; we will denote it as ε_0 .

The question is: did we escape the countable infinity represented by ω ? No, we will soon see that all of these numbers are, in fact, still countable. Nonetheless, we have stumbled upon something important. The last number, $\varepsilon_0 = \omega^{\varepsilon_0}$, is closely connected to Peano arithmetic, and we will also use it when proving Goodstein's theorem. It also gives rise to an entire class of ordinals called the *epsilon numbers*.

3.3.1 Epsilon Numbers

Definition 3.37. An ordinal ξ is an *epsilon number* if it is a fixed point of the normal function $\xi \mapsto \omega^\xi$. That is, if $\xi = \omega^\xi$. Theorem 3.16 (iv) asserts the existence of a bijective normal function $\varepsilon : \text{On} \rightarrow \{\xi \mid \xi = \omega^\xi\}$ enumerating the epsilon numbers. We denote by ε_β the ordinal $\varepsilon(\beta)$.

Proposition 3.38. *For any ordinal β , it holds that*

- (i) $\varepsilon_0 = \sup\{1, \omega, \omega^\omega, \omega^{\omega^\omega}, \omega^{\omega^{\omega^\omega}}, \dots\}$,
- (ii) $\varepsilon_{\beta+1} = \sup\{1, \varepsilon_\beta, \varepsilon_\beta^{\varepsilon_\beta}, \varepsilon_\beta^{\varepsilon_\beta^{\varepsilon_\beta}}, \varepsilon_\beta^{\varepsilon_\beta^{\varepsilon_\beta^{\varepsilon_\beta}}}, \dots\}$,

(iii) $\varepsilon_\beta = \sup\{\varepsilon_\delta \mid \delta < \beta\}$, whenever β is a limit ordinal.

Proof. We prove the theorem by transfinite induction on β . (i) This is the definition of ε_0 . (iii) Holds because the epsilon function is normal. (ii) Following Theorem 3.16, we know that $\varepsilon_{\beta+1}$ is the limit of the sequence

$$\varepsilon_\beta + 1, \omega^{\varepsilon_\beta+1}, \omega^{\omega^{\varepsilon_\beta+1}}, \omega^{\omega^{\omega^{\varepsilon_\beta+1}}}, \dots$$

Let α_n denote the element with index $n \in \omega$. Define a different sequence for $n \geq 2$ as $\alpha'_2 := \varepsilon_\beta^\omega$ and $\alpha'_{n+1} := \varepsilon_\beta^{\alpha'_n}$. Clearly,

$$\sup\{\alpha'_n \mid n \geq 2\} = \sup\left\{\varepsilon_\beta^\omega, \varepsilon_\beta^{\varepsilon_\beta^\omega}, \varepsilon_\beta^{\varepsilon_\beta^{\varepsilon_\beta^\omega}}, \dots\right\} = \sup\left\{1, \varepsilon_\beta, \varepsilon_\beta^{\varepsilon_\beta}, \varepsilon_\beta^{\varepsilon_\beta^{\varepsilon_\beta}}, \dots\right\}.$$

We will use induction on n to show $\alpha_n = \alpha'_n$ for all $n \geq 2$:

$$\alpha_1 = \omega^{\varepsilon_\beta+1} = \omega^{\varepsilon_\beta} \cdot \omega = \varepsilon_\beta \cdot \omega$$

$$\alpha_2 = \omega^{\omega^{\varepsilon_\beta+1}} = \omega^{(\varepsilon_\beta \cdot \omega)} = (\omega^{\varepsilon_\beta})^\omega = \varepsilon_\beta^\omega = \alpha'_2$$

$$\alpha_3 = \omega^{\omega^{\omega^{\varepsilon_\beta+1}}} = \omega^{\varepsilon_\beta^\omega} = \omega^{\varepsilon_\beta^{1+\omega}} = \omega^{\varepsilon_\beta \cdot \varepsilon_\beta^\omega} = (\omega^{\varepsilon_\beta})^{\varepsilon_\beta^\omega} = \varepsilon_\beta^{\varepsilon_\beta^\omega} = \alpha'_3$$

$$\alpha_{n+2} = \omega^{\alpha_{n+1}} = \omega^{\alpha'_{n+1}} = \omega^{\varepsilon_\beta^{\alpha'_n}} = \omega^{\varepsilon_\beta^{1+\alpha'_n}} = \omega^{\varepsilon_\beta \cdot \varepsilon_\beta^{\alpha'_n}} = (\omega^{\varepsilon_\beta})^{\varepsilon_\beta^{\alpha'_n}} = \varepsilon_\beta^{\varepsilon_\beta^{\alpha'_n}} = \alpha'_{n+2}$$

Hence $\varepsilon_\beta = \sup\{\alpha_n \mid n < \omega\} = \sup\{\alpha'_n \mid 2 \leq n < \omega\}$. \square

Lemma 3.39 (AC_ω). *A countable union of countable sets is countable. Specifically, if β and γ_α for $\alpha < \beta$ are countable ordinals, then $\gamma = \sup\{\gamma_\alpha \mid \alpha < \beta\}$ is also a countable ordinal.*

Remark. AC_ω denotes the *axiom of countable choice*, which states that every countable set has a choice function. Without AC_ω , a countable union of countable (even finite) sets might be uncountable.

Proof. Let $A = \langle A_n \mid n \in I \rangle$ be a countable collection of sets, WLOG $I = \omega$, such that all A_n are countable, and denote $S := \bigcup A$. We will define an injection $g : S \rightarrow \omega \times \omega$ (here, $\omega \times \omega$ is countable⁶). Since each A_n is countable, it injects into ω , and we can *choose* an injection $j_n : A_n \rightarrow \omega$ (and because there are only countably many sets A_n , we are making only countably many choices). For an element $a \in S$, define

$$n_a := \min\{n \in \omega \mid a \in A_n\}.$$

This number indicates in which A_n does a first appear in. Notice that more elements $a \in S$ can have the same number n_a , but that $j_{n_a}(a)$ uniquely identifies a among these elements (since j_{n_a} is injective). Hence, we can define an injection $g : a \mapsto (n_a, j_{n_a}(a))$. \square

From now on, we will generally assume AC_ω .

Lemma 3.40 (AC_ω). *The ordinal ω^α is countable $\iff \alpha$ is countable.*

⁶Prove that the Cartesian product of finitely many countable sets is countable. Hint: prime numbers might help. Then try proving it without the use of primes.

Proof. We first prove ‘ \Leftarrow ’ by transfinite induction on α . If $\alpha = 0$, it holds. Now suppose the claim holds for a countable α and consider the ordinal $\omega^{\alpha+1} = \omega^\alpha \cdot \omega$. This is countable because it is the order type of the set $\omega \times \omega^\alpha$, which is countable since it is the cartesian product of two countable sets. Finally, if α is a countable limit ordinal, then

$$\omega^\alpha = \sup\{\omega^\delta \mid \delta < \alpha\} = \bigcup\{\omega^\delta \mid \delta < \alpha\}.$$

Because α is countable and all ω^δ are countable ordinals (induction hypothesis), Lemma 3.39 implies that ω^α is countable as well.

We prove ‘ \Rightarrow ’ by contraposition. Suppose that α is uncountable; since $\xi \mapsto \omega^\xi$ is a normal function, it grows at least as fast as the identity function, and so $\omega^\alpha \geq \alpha$. We can now easily define an injection $\alpha \rightarrow \omega^\alpha$, showing that ω^α is uncountable. \square

Lemma 3.41. *The epsilon number ε_0 is countable.*

Proof. By definition, ε_0 is the limit of the sequence $\alpha_0 = 1$ and $\alpha_{n+1} = \omega^{\alpha_n}$. By induction on n and using the previous lemma, one can show that all α_n are, in fact, countable. This implies that ε_0 is a countable union of countable ordinals and is therefore countable as well. \square

Note on ordinal notations We have proven the previous lemma via a statement that requires AC_ω , because it made our job easier (and will continue to do so greatly down the line). However, we could have shown that ε_0 is countable by realizing that every $\alpha < \varepsilon_0$ has a finite *hereditary* Cantor normal form; therefore, we can use prime numbers to encode α as a unique natural number, hence constructing an injection $\varepsilon_0 \rightarrow \omega$. Hereditary CNF simply means that if any of the powers γ_i are ordinals larger than ω , then we express them in CNF as well, and we repeat the process inductively. For example:

$$\alpha = \omega^{\omega^{\omega+1} + \omega^2 \cdot 3 + 5} + \omega^{\omega \cdot 2 + 1} + \omega \cdot 2 + 7.$$

But ε_0 cannot be represented by a finite hereditary CNF, since $\varepsilon_0 = \omega^{\varepsilon_0}$ is its CNF, which is self-referential.

This concept can be generalized: if λ is a large ordinal, and every $\alpha < \lambda$ can be represented as a finite syntactic structure over some finite alphabet, then we can encode α as a unique natural number, constructing an injection $\lambda \rightarrow \omega$. This concept is called *ordinal notations*, and it does not require any kind of choice. However, the larger the ordinal λ , the more complicated the finite structure and encoding become. It is crucial to understand that when we use AC_ω to prove the countability of a large ordinal (like ε_0 or Γ_0 later), we are using it to simplify the process, not because we *have to* use it.

However, general statements like Lemma 3.40 or the following theorem do require the use of countable choice and cannot be proved in bare **ZF**.

Theorem 3.42 (AC_ω). *The epsilon number ε_β is countable $\iff \beta$ is countable.*

Proof. We first prove ‘ \Leftarrow ’ by transfinite induction on β . The base case $\beta = 0$ has been verified by the previous lemma. Suppose the claim holds for a countable

ordinal β ; that is, ε_β is countable, and we want to show that $\varepsilon_{\beta+1}$ is countable as well. Theorem 3.16 claims that $\varepsilon_{\beta+1}$ is the limit of the sequence $\alpha_0 = \varepsilon_\beta + 1$ and $\alpha_{n+1} = \omega^{\alpha_n}$. Note that α_0 is countable since ε_β is countable. One can now use induction on n and Lemma 3.40 to show that all α_n are in fact countable. This implies that $\varepsilon_{\beta+1}$ is a countable union of countable ordinals and is therefore also countable.

Finally, assume that β is a countable limit ordinal. Because the epsilon function $\beta \mapsto \varepsilon_\beta$ is normal, it holds that

$$\varepsilon_\beta = \sup\{\varepsilon_\delta \mid \delta < \beta\} = \bigcup\{\varepsilon_\delta \mid \delta < \beta\}.$$

Since β is countable, this is a countable union of countable ordinals (induction hypothesis), so it is countable as well.

One can prove ‘ \Rightarrow ’ in the same manner as in Lemma 3.40. \square

The theorem we have just proven places us in a difficult position. Does an uncountable ordinal even exist? If we assume the axiom of choice, then it is fairly easy to find one: just well-order the uncountable set $\mathcal{P}(\omega)$ and take its order type. Finding one in $\text{ZF} + \text{AC}_\omega$ seems to be much more difficult.

3.3.2 Veblen Hierarchy

We know that combining countable ordinals using the standard operations defined above produces more countable ordinals. The best tool for constructing large ordinals we currently have are the epsilon numbers (and, moreover, Theorem 3.16 implies that for any ordinal, there is a larger epsilon number), but it seems like they will not help us either. Consider the sequence

$$\gamma_0 = \varepsilon_0, \gamma_{n+1} = \varepsilon_{\gamma_n} \quad \longrightarrow \quad \varepsilon_0, \varepsilon_{\varepsilon_0}, \varepsilon_{\varepsilon_{\varepsilon_0}}, \varepsilon_{\varepsilon_{\varepsilon_{\varepsilon_0}}}, \dots$$

The largest number we can currently construct is the limit of this sequence; the first fixed point of the epsilon function $\xi \mapsto \varepsilon_\xi$, a number denoted as ζ_0 . However, this number, while enormously large, is still countable. The reason is that all the terms γ_n are countable (by induction and the previous theorem); hence, ζ_0 is a countable union of countable ordinals and is thus also countable.

We could define *zeta* (ζ) *numbers* in a similar fashion to how we defined epsilon numbers; however, for the same reasons that epsilon numbers with countable indices are countable, we would arrive at the conclusion that any zeta number with a countable index is still countable.

We could even create an entire hierarchy of these special fixed-point numbers. The bottom tier would be $\varphi_0(\beta) = \omega^\beta$; the second tier would be the epsilon numbers $\varphi_1(\beta) = \varepsilon_\beta$; the third tier would be the zeta numbers $\varphi_2(\beta) = \zeta_\beta$; the next one would be the so-called *eta numbers* $\varphi_3(\beta) = \eta_\beta$, and so on. The tiers are defined in such a way that the values of φ_{n+1} are the fixed points of φ_n . We could now consider the ordinal

$$\Lambda := \sup\{\varphi_n(0) \mid n < \omega\}.$$

However, this ordinal is *still countable*, as it is a countable union of countable ordinals.

The functions φ_n we have just constructed are called the *Veblen functions*, and they can be generalized for arbitrary ordinal indices.

Definition 3.43 (Veblen hierarchy, 1908). The functions $\varphi_\alpha : \text{On} \rightarrow \text{On}$ are defined for all ordinals $\alpha \geq 0$ recursively as:

- (i) $\varphi_0(\beta) = \omega^\beta$,
- (ii) $\varphi_{\alpha+1}(\beta)$ is the β th fixed point of φ_α , starting at $\beta = 0$.
- (iii) when α is a limit ordinal, we define $\varphi_\alpha(\beta)$ as the β th simultaneous fixed point of all the φ_δ for $\delta < \alpha$, also starting at $\beta = 0$.

Observation 3.44. *The function φ_α is normal for every α .*

Proof. By transfinite induction on α . It holds for $\alpha = 0$ because exponentiation is normal in the second argument. If $\alpha = \gamma + 1$ is isolated, then φ_α enumerates the fixed points of the normal function φ_γ and is by Theorem 3.16 (iv) normal. If α is isolated, it enumerates the simultaneous fixed points of the normal functions φ_δ for $\delta < \alpha$ and so φ_α is normal by Theorem 3.17 (iv). \square

Exercise 3.45. Show that if $\gamma > \alpha$, then $\varphi_\alpha(\varphi_\gamma(\beta)) = \varphi_\gamma(\beta)$ for any β .

This demonstrates that the values of $\varphi_{\alpha+1}$ are not only fixed points of φ_α , but they are fixed points of all φ_δ for $\delta \leq \alpha$. This means that we could have used condition (iii) to define φ_α for all ordinals α , not only for limit ordinals.

Exercise 3.46. Show that the ordinal function $\alpha \mapsto \varphi_\alpha(0)$ is normal.

Definition 3.47. An ordinal worth noting is the *Feferman–Schütte ordinal* Γ_0 , defined as the first fixed point of the function $\alpha \mapsto \varphi_\alpha(0)$, or equivalently, as the limit of the sequence $\gamma_0 = \varphi_0(0)$, $\gamma_{n+1} = \varphi_{\gamma_n}(0)$, that is

$$\gamma_0 = 1, \gamma_1 = \varphi_1(0) = \varepsilon_0, \gamma_2 = \varphi_{\varepsilon_0}(0) = \varphi_{\varphi_{\varphi_0(0)}(0)}(0), \gamma_3 = \varphi_{\varphi_{\varphi_{\varphi_0(0)}(0)}(0)}(0), \dots$$

It is one of the most famous ordinals in logic, and we will attempt to provide an explanation of why in Section 3.4.5.

Exercise 3.48. Show that Γ_0 is the first ordinal $\gamma > 0$ closed under φ . That is, the least $\gamma > 0$ such that for all $\alpha, \beta < \gamma$ we have $\varphi_\alpha(\beta) < \gamma$.

In other words, Γ_0 is the first ordinal that cannot be reached from below via repeated application of the Veblen functions. Does that mean that we have finally found an uncountable ordinal?

Exercise 3.49. Show that for all ordinals $\alpha < \Gamma_0$ it holds that $\varphi_\alpha(\Gamma_0) = \Gamma_0$.

That is, all φ_α attain the value Γ_0 at the same time! Even the very “slow” growing $\varphi_0(\beta) = \omega^\beta$ catches up to functions like $\varphi_{\varepsilon_0}(\beta)$ or $\varphi_{\varphi_{\varepsilon_0}}(\beta)$, and they all momentarily synchronize at $\beta = \Gamma_0$.

Theorem 3.50 (AC_ω). *The ordinal $\varphi_\alpha(\beta)$ is countable $\iff \alpha, \beta$ are countable.*

Proof. For the ‘ \Leftarrow ’ direction, we use transfinite induction on α . Notice that Lemma 3.40 is our base case (when $\alpha = 0$). If $\alpha = \gamma + 1$ is isolated, the proof is essentially the same as the proof of Theorem 3.42. If α is a countable limit ordinal, then we prove the claim by transfinite induction on β .

- If $\beta = 0$, then $\varphi_\alpha(0)$ is by Theorem 3.17 (ii) equal to the limit of the sequence

$$\gamma_0 = 0, \quad \gamma_{n+1} = \sup\{\varphi_\delta(\gamma_n) \mid \delta < \alpha\}.$$

Using our outer induction hypothesis (for α), one can show by induction on n that all γ_n are countable. Thus $\varphi_\alpha(0)$ is a countable union of countable ordinals and is therefore also countable.

- If $\beta = \beta' + 1$ is isolated, then the proof is almost identical to the case when $\beta = 0$. One just starts with $\gamma_0 = \varphi_\alpha(\beta') + 1$.
- If β is a limit ordinal, then since φ_α is normal we have that $\varphi_\alpha(\beta) = \sup\{\varphi_\alpha(\delta) \mid \delta < \beta\}$. Because we assume that β is countable, $\varphi_\alpha(\beta)$ is thus a countable union of countable ordinals (inner induction hypothesis) and is countable as well.

For the ‘ \Rightarrow ’ direction, we need to show that $\varphi_\alpha(\beta) \geq \alpha, \beta$; so if either α or β is uncountable, then $\varphi_\alpha(\beta)$ is also uncountable. Clearly $\varphi_\alpha(\beta) \geq \beta$ because φ_α is a normal function. Exercise 3.46 implies that $\varphi_\alpha(0) \geq \alpha$, and from here we have that $\varphi_\alpha(\beta) \geq \varphi_\alpha(0) \geq \alpha$. \square

Corollary 3.51. *The Feferman–Schütte ordinal Γ_0 is countable, as it is a countable union of countable ordinals.*

Remark. As mentioned in Section 3.3.1, the axiom of countable choice is not required to show that Γ_0 is countable, but it makes the task easier. If we wanted to prove it in ZF, we would show that the Veblen functions $\varphi_\alpha(\beta)$ provide a way of expressing every $\gamma < \Gamma_0$ as a unique finite syntactic structure. For more details about the Veblen hierarchy, see Sections 6.5, 7, and 8 of [13].

What we have shown in this section demonstrates an important concept in set theory when assuming AC_ω : one cannot reach uncountable infinity by starting from ω and applying ordinal operations such as addition, multiplication, exponentiation, finding fixed points of normal functions, and taking suprema — all the while utilizing only the ordinals we have already constructed along the way.

3.3.3 Hartogs’ Theorem

Does that mean that all hope is lost and there are no uncountable ordinals? Thankfully, no. The following theorem gives us a way out.

Theorem 3.52 (Hartogs, 1915). *For any set x , there exists an ordinal η such that there is no injection $\eta \rightarrow x$. The least such η is called the Hartogs number of x .*

Proof (cf. [10]). Consider the set (why is this a set?)

$$\mathcal{W} = \{(A, <_R) \mid A \subseteq x \text{ and } <_R \text{ is a well-ordering of } A\}.$$

We can use replacement to construct the set

$$S = \{\alpha \in \text{On} \mid \text{there exists } (A, <_R) \in \mathcal{W} \text{ order-isomorphic to } \alpha\}$$

by assigning to each $(A, <_R)$ its order type.

But this set is exactly the Hartogs number of x . Notice that S is transitive: if $\alpha \in S$ and $\gamma < \alpha$, then $\alpha \in S$ as well. A transitive set of ordinals is again an ordinal (Lemma 2.9), so S is an ordinal number η . Furthermore, there is no injection from η into x , because if there were, then we would get the contradiction that $\eta \in \eta$. And finally, η is the least such ordinal. If $\alpha < \eta$, then also $\alpha \in \eta$, and there is an injection $\alpha \rightarrow x$. \square

Remark. It is crucial to note that the theorem we just proved, which gives us the Hartogs number as a von Neumann ordinal, is more powerful than Hartogs' original 1915 result. Hartogs, working in \mathbf{Z} (proposed in 1908 by Zermelo, containing the axioms of \mathbf{ZF} except replacement and foundation), only proved the existence of a well-ordered set that could not be injected into x ; but he did not—and could not—show it was a von Neumann ordinal. The general theorem that “every well-ordered set is isomorphic to a unique von Neumann ordinal” is itself not provable in \mathbf{Z} and requires replacement (see proof of Theorem 2.13). We will now use this modern, replacement-based construction to construct an uncountable ordinal as the Hartogs number of ω . It is this very step, guaranteeing that the collection of all countable ordinals is a set, that fails in \mathbf{Z} and was one of the motivations for Fraenkel and Skolem to propose the axiom of replacement in 1922.

Notice that the theorem does not say that $x \prec \eta$, because this does not in general hold without \mathbf{AC} . However, if x is well-ordered, then it has an order type α , and we can compare α with η .

This allows us to access an uncountable ordinal. Let ω_1 be the Hartogs number of ω . That is, the first ordinal with the property $\omega_1 \not\preceq \omega$, or in other words, the *first uncountable ordinal*,⁷ and we can write

$$\omega_1 = \{\alpha \in \text{On} \mid \alpha \preceq \omega\}$$

Exercise 3.53. Show that an ordinal α is uncountable ($\alpha \not\preceq \omega$) $\iff \omega \prec \alpha$. Note that this is not true for general sets in \mathbf{ZF} , we would need to accept \mathbf{AC}_ω .

It is almost impossible to grasp just how unfathomably large ω_1 is. The entire vast, complex, mind-boggling hierarchy of ordinals described by the Veblen functions up to Γ_0 is still just a tiny, countable speck at the absolute “bottom” of the ordinal line from the perspective of ω_1 .

Exercise 3.54. Show that for any countable ordinal α , it holds that $\varphi_\alpha(\omega_1) = \omega_1$, specifically $\omega^{\omega_1} = \omega_1$ and $\varepsilon_{\omega_1} = \omega_1$. Furthermore show that $\varphi_{\omega_1}(0) = \omega_1$.

Realize that there was nothing special about the choice of ω . We can apply the same process to ω_1 to get ω_2 , and continue doing this to construct larger and larger ordinals (in the sense of cardinality).

Definition 3.55. For an ordinal α we define ω_α recursively as

- (i) $\omega_0 := \omega$,
- (ii) $\omega_{\alpha+1}$ is the Hartogs number of ω_α ,
- (iii) $\omega_\lambda := \sup\{\omega_\alpha \mid \alpha < \lambda\}$ for limit ordinals λ .

⁷The ordinal ω_1 is also commonly denoted as Ω .

Observation 3.56. *The number ω_α is the first ordinal that is larger (in the sense of cardinality) than all previous ω -numbers.*

This definition foreshadows the section about cardinal numbers, where we will encounter these omega numbers again and explore their properties in depth.

Hartogs numbers also allow us to finally prove that the trichotomy principle implies AC. In fact, this was the original motivation behind Hartogs' theorem.

Theorem 3.57. *The trichotomy principle implies the well-ordering principle.*

Proof. Let x be an arbitrary set, and let η be its Hartogs number. Apply the trichotomy principle to x and η . One of the following holds:

- (a) $x \preceq \eta$, there is an injection $x \rightarrow \eta$, or
- (b) $\eta \preceq x$, there is an injection $\eta \rightarrow x$.

The second case is impossible due to the defining property of η . Hence, there exists an injection $f : x \rightarrow \eta$. We can now well-order x by inheriting the order of η by f . □

3.4 Peano Arithmetic

To understand this section, the reader should be familiar with the basic notions of logic, including concepts such as language, theory, model, etc. Explanations of these concepts can be found in the lecture notes [5] for the course NAIL062.

3.4.1 Peano Axioms

Peano Arithmetic, denoted PA, is the standard axiomatic theory of the natural numbers. In ZFC, we have encountered the set of natural numbers, ω , constructed as the set of finite von Neumann ordinals. This is no coincidence; the set ω , together with the restrictions of operations of ordinal arithmetic to ω , serves as the *standard model* for PA, denoted by \mathcal{N} .

Our study of ordinal arithmetic in Section 3.2 has already established that these operations, when restricted to finite ordinals, are commutative and satisfy all the familiar properties of elementary arithmetic. The axioms of PA can therefore be seen as a precise, first-order logic attempt to capture the properties of this standard model.

Definition 3.58 (PA, [5]). The language of PA is $\mathcal{L}_{PA} = \langle 0, S, +, \times, \leq \rangle$ with equality. The base axioms of PA are the following formulas:

$$\begin{array}{ll}
 \neg Sx = 0 & x \times 0 = 0 \\
 Sx = Sy \implies x = y & x \times Sy = x \times y + x \\
 x + 0 = x & \neg x = 0 \implies (\exists y)(x = Sy) \\
 x + Sy = S(x + y) & x \leq y \iff (\exists z)(z + x = y)
 \end{array}$$

These axioms alone yield the much weaker *Robinson Arithmetic* (Q). It cannot prove, for example, the commutativity or associativity of addition or multiplication, or the transitivity of order. To obtain PA, we need to add the *Axiom*

Schema of Induction. That is, for each \mathcal{L}_{PA} -formula $\varphi(x, \vec{y})$, the following axiom is added:

$$(\varphi(0, \vec{y}) \wedge (\forall x)(\varphi(x, \vec{y}) \Rightarrow \varphi(Sx, \vec{y}))) \Longrightarrow (\forall x)\varphi(x, \vec{y}) \quad (3.4)$$

Remark. The last axiom schema should seem similar to the induction principle on ω from set theory:

$$(\forall X \subseteq \omega) \left((0 \in X \wedge (\forall x)(x \in X \Rightarrow x \cup \{x\} \in X)) \Longrightarrow X = \omega \right).$$

However, the axiom schema of induction is a weaker version, as it is a first-order logic attempt to simulate a second-order logic axiom with an axiom schema. The familiar induction principle could be expressed with the following second-order \mathcal{L}_{PA} -formula

$$(\forall X) \left((X(0) \wedge (\forall x)(X(x) \Rightarrow X(Sx))) \Longrightarrow X = (\forall x)X(x) \right).$$

By adding it to **PA**, we would obtain the much stronger second-order theory **PA**₂.

Here X represents (any) unary relation; that is, a subset of the universe. The important distinction is that (3.4) provides an infinite collection of axioms, one for each subset of the universe that is *definable* by a \mathcal{L}_{PA} -formula φ .

This restriction is the source of **PA**'s most profound properties and limitations. For example, **PA**₂ is categorical; that is, it has only one model (up to isomorphism) — the standard model \mathcal{N} . On the other hand, **PA** allows the existence of other “nonstandard” models.

3.4.2 Models of Arithmetic

We have already mentioned that the *standard model* of **PA** is the \mathcal{L}_{PA} -structure $\mathcal{N} = (\omega, 0^{\mathcal{N}}, S^{\mathcal{N}}, +^{\mathcal{N}}, \times^{\mathcal{N}}, \leq^{\mathcal{N}})$, where the domain is the set ω , the interpretation of the symbol ‘0’ is $0^{\mathcal{N}} = \emptyset$, the successor of x is $S^{\mathcal{N}}(x) = x \cup \{x\}$, and $+^{\mathcal{N}}$, $\times^{\mathcal{N}}$ and $\leq^{\mathcal{N}}$ are the operations of ordinal arithmetic restricted to ω . Any model that is not isomorphic to \mathcal{N} is called *nonstandard*.

Theorem 3.59. *There exist countable nonstandard models of **PA**.*

Proof sketch. By the *compactness theorem*. We extend \mathcal{L}_{PA} with a new constant symbol c . Consider the theory $T = \mathbf{PA} \cup \{c > \bar{n} \mid n \in \omega\}$, where \bar{n} is the \mathcal{L}_{PA} -term $S(S(\dots S(0)\dots))$ (n times). Any finite subset $T_0 \subset T$ is satisfiable: we take \mathcal{N} as the model and interpret c as a standard natural number larger than any numeral \bar{n} explicitly mentioned in T_0 . By the compactness theorem, T has a model \mathcal{M} . This \mathcal{M} must be a model of **PA**, but the interpretation of c is a “nonstandard” number, an element larger than all standard elements \bar{n} . Thus, $\mathcal{M} \not\cong \mathcal{N}$. \square

All countable nonstandard models \mathcal{M} share a common structure: they begin with an initial segment isomorphic to ω (the standard part), which is then followed by a collection of “blocks” of nonstandard numbers. This “pathology” of **PA** is not merely set-theoretic, but also computational:

Theorem 3.60 (Tennenbaum, 1959). *No countable nonstandard model of **PA** is recursive.*

Intuition. If $\mathcal{M} = (M, 0^{\mathcal{M}}, S^{\mathcal{M}}, +^{\mathcal{M}}, \times^{\mathcal{M}}, \leq^{\mathcal{M}})$ is a nonstandard model with a countable domain M , then for every bijection $j : \omega \rightarrow M$, the functions $\oplus, \otimes : \omega \times \omega \rightarrow \omega$ defined as

$$\oplus(x, y) := j^{-1}(j(x) +^{\mathcal{M}} j(y)), \quad \otimes(x, y) := j^{-1}(j(x) \times^{\mathcal{M}} j(y))$$

are not recursive: there is no algorithm (Turing machine) that would compute them for every input. See [34] for more details and a proof.

The induction schema, while syntactically “weaker” than its second-order counterpart, thus imposes enormous computational complexity on any nonstandard structure that satisfies it, effectively isolating the standard model as the only computationally tractable one.

3.4.3 Gödel’s Incompleteness Theorems

When working with a formal theory, it is natural to ask what statements we can prove from its axioms. If a theory T can prove the sentence ψ , we write $T \vdash \psi$. A theory is *consistent* if it is free from contradictions, meaning it is impossible to prove both a statement ψ and its negation $\neg\psi$ from its axioms; or equivalently, if it has a model. A consistent theory is *complete* if it has an “opinion” on every statement, meaning for every sentence ψ in its language, the theory can prove either ψ or $\neg\psi$. If it cannot do either, it is said to be *incomplete*, and ψ is said to be *independent* in T . Equivalently, ψ is independent in T , if it holds in some models of the theory but does not hold in others.

Probably the most influential result linking these concepts together with PA are the famous Incompleteness Theorems, published by Kurt Gödel⁸ [15] in 1931. Veritasium has an amazing video [37] that provides an intuitive explanation of this topic. We provide only a simplified explanation of these profound results; for more details and proofs, refer to [5].

Despite its limitations, PA is a remarkably powerful theory. Its expressive power is sufficient to represent all computable (recursive) functions. This strength is the key to PA’s own undoing. It allows for the *arithmetization of syntax* (Gödel numbering), whereby the syntax of \mathcal{L}_{PA} (terms, formulas, proofs) can be uniquely encoded as natural numbers. Syntactic operations (like substitution) and relations (like “is a proof of”) become recursive functions and relations on these numbers. Crucially, this allows for the creation of a provability predicate.

Definition 3.61. There exists an \mathcal{L}_{PA} -formula $\text{Prov}_{PA}(x)$ such that for any sentence ϕ it holds that $(PA \vdash \phi) \iff \text{Prov}_{PA}(\ulcorner \phi \urcorner)$. Here, $\ulcorner \phi \urcorner$ denotes the Gödel number of ϕ . The formula $\text{Prov}_{PA}(\ulcorner \phi \urcorner)$ is: “there exists x such that x is the Gödel number of a proof of the sentence with Gödel number $\ulcorner \phi \urcorner$.”

⁸The life of Kurt Gödel (1906–1978) is a fascinating story. Born in Brno, he left for Vienna at the age of eighteen to study mathematics and logic. At twenty-four, he proved his incompleteness theorems, which formed the basis of his doctoral dissertation. He later emigrated to the United States following the rise of Nazism. Albert Einstein regarded Gödel as the greatest logician since Aristotle and once remarked that the only reason he went to his office was to have the privilege of walking home with Gödel. Yet Gödel’s life was not without darkness: he struggled with psychological illness throughout adulthood and ultimately died of self-starvation, driven by the paranoid belief that someone was trying to poison him. Perhaps the most detailed account of Gödel’s life (as of the writing of this text) can be found in [3].

This predicate allows the theory to “talk about” its own provability, leading directly to sentences that self-reference and assert their own unprovability.

Theorem 3.62 (Gödel’s First Incompleteness Theorem, 1931). *If PA is consistent, then it is incomplete.*

Proof sketch. Consider a sentence \mathbf{g} (the *Gödel sentence*) saying: “there is no x such that x is the Gödel number of a proof of the sentence with Gödel number $\ulcorner \mathbf{g} \urcorner$.” Notice that $(\text{PA} \vdash \mathbf{g}) \iff \neg \text{Prov}_{\text{PA}}(\ulcorner \mathbf{g} \urcorner)$. Hence if $\text{PA} \vdash \mathbf{g}$, then PA is inconsistent. Therefore, if PA is consistent, then $\text{PA} \not\vdash \mathbf{g}$, and it is incomplete. \square

As a corollary of this theorem, Gödel achieved his second result.

Theorem 3.63 (Gödel’s Second Incompleteness Theorem, 1931). *PA cannot prove its own consistency.*

Proof sketch. Let $\text{Con}(\text{PA})$ be the \mathcal{L}_{PA} -sentence $\neg \text{Prov}_{\text{PA}}(\ulcorner \perp \urcorner)$ (where \perp is a contradiction, e.g., $0 = S0$). That is, $\text{Con}(\text{PA})$ is true if and only if PA is consistent. Because Gödel formalized the entire proof of the previous theorem in PA (using Gödel numbers), his first theorem can be expressed as

$$\text{PA} \vdash (\text{Con}(\text{PA}) \implies \neg \text{Prov}_{\text{PA}}(\ulcorner \mathbf{g} \urcorner)).$$

This together with the equivalence $(\text{PA} \vdash \mathbf{g}) \iff \neg \text{Prov}_{\text{PA}}(\ulcorner \mathbf{g} \urcorner)$ gives

$$\text{PA} \vdash (\text{Con}(\text{PA}) \implies \mathbf{g}).$$

Now, suppose for contradiction that PA could prove its own consistency. Combining this with the last formula gives $\text{PA} \vdash \mathbf{g}$, but this is a contradiction, since (the end of the previous proof) if PA is consistent, then $\text{PA} \not\vdash \mathbf{g}$. \square

Gödel’s original formulation of these theorems did not, in fact, talk about PA, but about a system he called P, a close relative of PA. Gödel then had to make a philosophical assumption. He argued that any other system “related” to, and at least as strong as P (and therefore capable of arithmetic), would also be capable of producing a Gödel sentence \mathbf{g} ; thus, his incompleteness theorems would apply to this system as well. This was a strong, intuitive argument, but he could not formally prove it.

The missing piece was provided in 1936 by Alan Turing [36], who formalized the notion of computability using the Turing machine, which made a formal proof of Gödel’s conjecture possible.

Theorem 3.64 (Generalized Gödel’s Incompleteness Theorems). *For any consistent, recursively axiomatized theory T, it holds that:*

- (1) *If T is an extension of Robinson arithmetic Q, then T is incomplete.*
- (2) *If T is an extension of Peano arithmetic PA, then T cannot prove its own consistency.*

Remark. Recursively axiomatized means that there is an algorithm (Turing machine) that, for every input formula φ , halts and answers whether φ is an axiom of T. The condition that T is an extension of Q (or PA) essentially means that T is at least as powerful as Q (or PA). For example, PA is an extension of Q.

Corollary 3.65. *It is impossible to prove the consistency of ZFC inside ZFC.*

An example of an independent statement in ZFC is the famous continuum hypothesis CH, claiming that there is no set x such that $\omega \prec x \prec \mathcal{P}(\omega)$. Similarly, AC can be shown to be independent in ZF, meaning that if ZF is consistent, then ZFC is as well.

In 1940, Gödel showed that neither AC can be disproved from ZF, nor CH from ZFC, by constructing the *constructible universe*, a model of ZF in which both AC and CH hold. This model begins with the empty set and adds only those sets that are definable from previous ones, thus forming the minimal universe compatible with the axioms. Later, in 1963, Paul Cohen showed that CH cannot be proved from ZFC by developing the method of *forcing*, which allowed him to construct a model of ZFC in which CH fails. Through a different forcing argument, he likewise obtained a model of ZF that violates AC.

3.4.4 Consistency and the Connection with ε_0

Gödel’s second theorem seems to place us in a difficult position: a consistency proof for PA must employ principles that transcend PA itself. While ZFC is far stronger than PA and easily proves $\text{Con}(\text{PA})$ (by exhibiting the model \mathcal{N}), this isn’t a very “unilluminating” result. PA is a “finitary” theory, while ZFC is a wildly “infinitary” theory (it assumes the existence of various vast infinities). By proving the consistency of PA in ZFC, we base our proof on the assumption that ZFC is consistent. It would be better to find a weaker system that is still capable of proving $\text{Con}(\text{PA})$.

Theorem 3.66 (Gentzen’s Consistency Proof,⁹ 1936). *The consistency of PA is provable in Primitive Recursive Arithmetic PRA (which by itself is weaker than PA), augmented with a schema for transfinite induction up to the ordinal ε_0 .*

Gentzen’s proof precisely identified the principle transcending PA that is required to prove its consistency. Recall from Section 3.3.1 that ε_0 is the first fixed point of the ordinal function $\alpha \mapsto \omega^\alpha$, the limit of the sequence $\omega, \omega^\omega, \omega^{\omega^\omega}, \dots$. Gentzen’s result, $\text{PRA} + \text{TI}(\varepsilon_0) \vdash \text{Con}(\text{PA})$, thus establishes two facts:

- (a) Proving the consistency of PA does not require the full power of ZFC; only transfinite induction up to a countable ordinal. That is, the assumption that ε_0 contains no infinite decreasing chains.
- (b) The principle of transfinite induction up to ε_0 , $\text{TI}(\varepsilon_0)$, must be unprovable in PA (lest PA prove its own consistency).

Gentzen also showed that using any smaller ordinal $\alpha < \varepsilon_0$ is not enough. This calibrates the strength of PA with extraordinary precision. The collected strength of PA’s infinite induction schema is exactly equivalent to the single principle of transfinite induction up to (but not including) ε_0 . This is formalized in the concept of the *proof-theoretic ordinal*.

Theorem 3.67. *The proof-theoretic ordinal of PA is $|\text{PA}| = \varepsilon_0$.*

⁹See [28] for a modern version of the proof. Moreover, [8] provides an alternative view on this result, and talks about the consistency of PA in general.

This theorem has a twofold meaning that we can understand intuitively:

- (a) What PA *can* prove: PA is strong enough to prove the well-foundedness of any recursive well-ordering $<_R$ on ω with order-type $\alpha < \varepsilon_0$.
- (b) What PA *cannot* prove: PA is *not* strong enough to prove the well-foundedness of any recursive well-ordering $<_R$ on ω with order-type $\alpha \geq \varepsilon_0$.

Here, “recursive” means that there exists an algorithm that can answer whether $x <_R y$ or $y <_R x$ for all $x, y \in \omega$. Well-foundedness is the arithmetical statement that every nonempty subset of ω has a minimal element. Proving well-foundedness is thus equivalent to proving that transfinite induction “works” for that ordering (as there cannot be any infinite decreasing chains).

A point of confusion here might be the fact that any well-ordering is well-founded. But PA does not know that $<_R$ is a well-ordering; it only receives an “object,” $<_R$, together with the instructions: “prove that the ordering you received is well-founded.”

Therefore, PA can formalize proofs by transfinite induction up to any ordinal $\alpha < \varepsilon_0$, but it cannot justify the principle of transfinite induction up to ε_0 itself.

3.4.5 Limits of Predicative Mathematics

The discovery of set-theoretic paradoxes (such as Russell’s and Burali-Forti’s) in the early 20th century triggered the so-called *Grundlagenkrise*, or foundational crisis, in mathematics. The naive assumption that any property $\phi(x)$ could define a set $\{x \mid \phi(x)\}$ was shown to lead to contradictions. This prompted a range of philosophical responses, but most mathematicians eventually turned to the axiomatic framework of ZFC, which resolved the paradoxes by carefully restricting what counts as a set.

However, a significant objection came from mathematicians like Poincaré and, most notably, Hermann Weyl. They argued that the core problem was the use of *impredicative definitions*—definitions that define an object S by quantifying over a totality T that already includes S . This went directly against ZFC, as it uses impredicative definitions all the time. For example, the supremum of a set is defined as its least upper bound. In this case, the totality is the set of all upper bounds, and since the supremum is itself an upper bound, it is a member of that totality.

Weyl argued that such definitions were circular and potentially dangerous. He decided to rebuild mathematical analysis on a “safe” *predicative* basis, starting with his 1918 paper “The Continuum.” He succeeded in developing a significant portion of classical analysis, but was unable to replicate everything. The big question became: “How much of mathematics can we *actually* recover using only predicative methods?” In the 1960s, Solomon Feferman and Kurt Schütte independently found the precise answer. We provide an intuitive interpretation of their result.

Theorem 3.68 (Feferman–Schütte, c. 1965). *The proof-theoretic ordinal of “predicative mathematics” is the Feferman–Schütte ordinal, denoted Γ_0 . It is the first fixed point of the function $\alpha \mapsto \varphi_\alpha(0)$ (see Section 3.3.2).*

This shows that Γ_0 is the first ordinal that cannot be proven to be well-founded by predicative means, just as ε_0 is the first ordinal that cannot be proven well-founded by the “finitistic” means of PA.

Remark. It should be noted that not all mathematicians agree on what “predicative mathematics” means exactly. It would be more accurate to say that Feferman and Schütte showed that Γ_0 is the first ordinal that cannot be proven well-founded by *certain* predicative means, and most people agree that those means are a reasonable interpretation of predicativity. This was later formalized in the 1970s by Friedman and Simpson in a formal system called ATR_0 (Arithmetical Transfinite Recursion). So what the theorem above really says is that the proof theoretic ordinal of ATR_0 is Γ_0 .

The discipline of finding the proof-theoretic ordinals of theories is called *Ordinal Analysis*. For an introduction to the topic, I recommend [28]. Other notable sources are [9], [32] and [26].

3.5 Applications of Countable Ordinals

Gödel’s independent sentences, \mathfrak{g} and $\text{Con}(\text{PA})$, are meta-mathematical statements, not “natural” theorems of number theory. For decades, it was an open question whether any “ordinary” theorem of arithmetic or combinatorics was unprovable in PA, leading to speculation that Gödel’s Incompleteness Theorem would not have meaningful implications to practical mathematics.

However, in 1977, Paris and Harrington [27] showed that a very natural variation of Ramsey’s Theorem was true but not provable in PA. Five years later, in 1982, Kirby and Paris [25] showed that Goodstein’s theorem, a statement purely about sequences of natural numbers, cannot be proven in PA either.

The second theorem presented in Kirby and Paris’s 1982 paper establishes an analogous result, this time showing that a statement about the Hydra game is unprovable in PA.

3.5.1 Goodstein Sequences

Definition 3.69. The *hereditary base- n representation* of a natural number m is achieved by first writing m in base n , and then applying the procedure inductively to each exponent until there are no numbers larger than n .

Example. 100 in hereditary base-2 is $64+32+4 = 2^6+2^5+2^2 = 2^{2^2+2}+2^{2^2+1}+2^2$.

The Goodstein sequence of a natural number $m > 0$ is generated as follows.

- Start with $m_0 = m$.
- To get m_1 , write m_0 in hereditary base-2 representation, replace all 2s with 3s, and then subtract 1 from the result.
- To get m_{n+1} from m_n , write m_n in hereditary base- $(n+2)$ representation, replace each occurrence of $(n+2)$ with $(n+3)$, and subtract 1.
- If ever $m_n = 0$, then $m_{n+1} = 0$.

For example, when we start with $m_0 = 3$, we get the sequence:

$$\begin{aligned}
m_0 &= 2^1 + 1 &= 3 \\
m_1 &= 3^1 + 1 - 1 = 3^1 &= 3 \\
m_2 &= 4^1 - 1 = 3 &= 3 \\
m_3 &= 3 - 1 = 2 &= 2 \\
m_4 &= 2 - 1 = 1 &= 1 \\
m_5 &= 1 - 1 = 0 &= 0
\end{aligned}$$

Notice that at step $m_2 \rightarrow m_3$, the base used (4 at m_2) exceeded the value of $m_2 = 3$, which caused the sequence to start decreasing and eventually terminate. Does it always terminate? Let's try again, this time with $m_0 = 29$.

$$\begin{aligned}
m_0 &= 2^{2^2} + 2^{2^{2^1}} + 2^2 + 1 &= 29 \\
m_1 &= 3^{3^3} + 3^{3^{3^1}} + 3^3 &\sim 8 \cdot 10^{12} \\
m_2 &= 4^{4^4} + 4^{4^{4^1}} + 4^4 - 1 = &\sim 10^{154} \\
&= 4^{4^4} + 4^{4^{4^1}} + 3 \cdot 4^3 + 3 \cdot 4^2 + 3 \cdot 4 + 3 \\
m_3 &= 5^{5^5} + 5^{5^{5^1}} + 3 \cdot 5^3 + 3 \cdot 5^2 + 3 \cdot 5 + 2 &\sim 10^{2200} \\
m_4 &= 6^{6^6} + 6^{6^{6^1}} + 3 \cdot 6^3 + 3 \cdot 6^2 + 3 \cdot 6 + 1 &\sim 10^{36305}
\end{aligned}$$

This does not look like it will terminate. Let's try it one more time, this time with $m_0 = 4$, to really understand what is going on.

$$\begin{aligned}
m_0 &= 2^2 &= 4 \\
m_1 &= 3^3 - 1 = 2 \cdot 3^2 + 2 \cdot 3 + 2 &= 26 \\
m_2 &= 2 \cdot 4^2 + 2 \cdot 4 + 1 &= 41 \\
m_3 &= 2 \cdot 5^2 + 2 \cdot 5 &= 60 \\
&\vdots &\vdots \\
m_9 &= 2 \cdot 11^2 + 11 &= 253 \\
m_{10} &= 2 \cdot 12^2 + 12 - 1 = 2 \cdot 12^2 + 11 &= 299 \\
&\vdots &\vdots \\
m_{22} &= 2 \cdot 24^2 - 1 = 24^2 + 23 \cdot 24 + 23 &= 1151 \\
&\vdots &\vdots \\
m_{B-2} &= 2 \cdot B^1 &= 2B > 10^{10^8} \\
m_{B-1} &= 2 \cdot (B+1)^1 - 1 = (B+1)^1 + B &= 2B + 1 \\
m_B &= (B+2)^1 + B - 1 &= 2B + 1 \\
&\vdots &\vdots \\
m_{B+k} &= (B+k+2)^1 + B - (k+1) &= 2B + 1 \\
&\vdots &\vdots \\
m_{2B-2} &= (2B)^1 + 1 &= 2B + 1 \\
m_{2B-1} &= (2B+1)^1 &= 2B + 1
\end{aligned}$$

$$\begin{array}{rcl}
m_{2B} & = & (2B + 2)^1 - 1 = 2B + 1 & = & 2B + 1 \\
m_{2B+1} & = & 2B + 1 - 1 & = & 2B \\
& & \vdots & & \vdots \\
m_{2B+k} & = & 2B + 1 - k & = & 2B - (k - 1) \\
& & \vdots & & \vdots \\
m_{3B} & = & 2B + 1 - B = 1 & = & B + 1 \\
& & \vdots & & \vdots \\
m_{4B} & = & 2B + 1 - 2B = 1 & = & 1 \\
& & \vdots & & \vdots \\
m_{4B+1} & = & 0 & = & 0
\end{array}$$

This demonstrates that when we start at $m_0 = 4$, the sequence first rises to index $B - 2$, where it attains the value $2 \cdot B^1$ (current base is B). In the next step, this is decomposed to $(B + 1)^1 + B$ (this is also the maximum value). The sequence lingers here for the next B steps (each base change raises the value by 1, and this is immediately subtracted), and afterward, it begins its long descent. Finally, at index $4B + 1$, it reaches zero. The exact value of B is $3 \cdot 2^{402653209} - 1$, and the length of this sequence is $4B + 2 = 3 \cdot 2^{402653211} - 2$.

Goodstein's theorem claims that a similar fate awaits every Goodstein sequence, no matter how large the starting value might be.

Theorem 3.70 (Goodstein, 1944). *For every natural number m , there exists a natural number n such that $m_n = 0$.*

Definition 3.71 (Change of base function). For a natural number $b \geq 2$ we define the functions

- (a) *change of base function* $R_b : \omega \rightarrow \omega$ as the function that takes a natural number n and replaces each b in the hereditary base- b representation of n with $b + 1$.
- (b) *ω -change of base function* $R_b^\omega : \omega \rightarrow \varepsilon_0$ as the function that takes a natural number n and replaces each b in the hereditary base- b representation of n with ω .

Formally, we define $R_b(0) = R_b^\omega(0) = 0$ and for $n > 0$ expressed in base- b as

$$n = \sum_{i=0}^k b^i \cdot p_i$$

we let

$$R_b(n) = \sum_{i=0}^k (b + 1)^{R_b(i)} \cdot p_i, \quad R_b^\omega(n) = \sum_{i=0}^k \omega^{R_b^\omega(i)} \cdot p_i.$$

Example. For example:

- $R_2(29) = R_2(2^2 + 2^{2+1} + 2^2 + 1) = 3^3 + 3^{3+1} + 3^3 + 1,$
- $R_2^\omega(29) = R_2^\omega(2^2 + 2^{2+1} + 2^2 + 1) = \omega^{\omega^\omega} + \omega^{\omega+1} + \omega^\omega + 1$

Remark. Since each $n \in \omega$ is finite, the hereditary base- b representation of n is finite, and thus $R_b^\omega(n)$ also contains only finitely many occurrences of ω . Therefore $R_b^\omega(n) < \varepsilon_0$ for each $b \geq 2$ and $n \in \omega$.

Observation 3.72. *The terms of a Goodstein sequence starting in m could be defined as $m_{n+1} = R_{n+2}(m_n) - 1$.*

Lemma 3.73. *For every $b \geq 2$ and $n \geq 0$ it holds that $R_b^\omega(n+1) > R_b^\omega(n)$.*

Intuition. This should seem obvious, for example when $b = 2$ and $n = 13$ we get

$$\begin{aligned} R_2^\omega(13) &= R_2^\omega(2^{2+1} + 2^2 + 1) = \omega^{\omega+1} + \omega^\omega + 1 \\ R_2^\omega(13+1) &= R_2^\omega(2^{2+1} + 2^2 + 2) = \omega^{\omega+1} + \omega^\omega + \omega, \end{aligned}$$

Proof. Let $b \geq 2$ be given; we prove the claim using induction on n . If $n = 0$, then we have $R_b^\omega(1) = 1 > 0 = R_b^\omega(0)$. If $n > 0$, then let the following be the base- b expansions of n and $n+1$:

$$\begin{aligned} n &= b^{c_0} \cdot p_0 + b^{c_1} \cdot p_1 + \cdots + b^{c_k} \cdot p_k \\ n+1 &= b^{d_0} \cdot q_0 + b^{d_1} \cdot q_1 + \cdots + b^{d_l} \cdot q_l \end{aligned}$$

where all p_i and q_j are nonzero, and $c_0 > c_1 > \cdots > c_k$ and $d_0 > d_1 > \cdots > d_l$. Theorem 3.36, or more precisely, part two of Theorem 3.34, describes the two sufficient and necessary conditions that these expansions must satisfy in order for one of them to be greater than the other ($n+1 > n$). Notice that the claim already holds for all exponents c_i and d_j from the induction hypothesis. Thus, when we apply R_b^ω to these expansions—change the base b to ω and apply R_b^ω to the exponents—the sufficient and necessary conditions will not be affected; therefore $R_b^\omega(n+1) > R_b^\omega(n)$. \square

We can now prove Goodstein's theorem.

Proof of Theorem 3.70 (cf. [33]). Let $m = m_0$ be given. We will define a sequence of ordinals $\mu_n < \varepsilon_0$ satisfying

$$\mu_n > 0 \implies \mu_{n+1} < \mu_n \quad \text{and} \quad \mu_n > 0 \iff m_n > 0.$$

Since ε_0 is well-founded, it does not admit infinite strictly decreasing sequences; thus, there exists an index k at which $\mu_k = 0$, and therefore also $m_k = 0$.

Define $\mu_n := R_{n+2}^\omega(m_n)$. Clearly $\mu_n > 0 \iff m_n > 0$, so it remains to show that μ_n is decreasing. If $\mu_n > 0$, then

$$\begin{aligned} \mu_{n+1} &= R_{n+3}^\omega(m_{n+1}) \\ &= R_{n+3}^\omega(R_{n+2}(m_n) - 1) \\ &< R_{n+3}^\omega(R_{n+2}(m_n)) \\ &= R_{n+2}^\omega(m_n) = \mu_n. \end{aligned}$$

The inequality holds by Lemma 3.73 since if $\mu_n > 0$, then $m_n > 0$ and so $R_{n+2}(m_n) \geq m_n > 0$. The equality after the inequality is a trivial property of the base change functions ($n+2 \rightarrow n+3 \rightarrow \omega$ is the same as $n+2 \rightarrow \omega$). \square

Extended Goodstein’s theorem In the version of Goodstein sequences presented above, we started with base $b_0 = 2$ for m_0 , then changed it to $b_1 = 3$ for m_1 , and in general worked with base $b_n = n + 2$ for m_n . This can be generalized by considering any non-decreasing sequence $2 \leq b_0 \leq b_1 \leq \dots$ of bases and defining the term m_{n+1} from m_n by expressing m_n in hereditary base- b_n notation, replacing each occurrence of b_n with b_{n+1} , and subtracting one. It is not hard to modify the proof above to show that this sequence still terminates.

This extended version is, in fact, the one Goodstein originally considered in [16], and he proved that it is equivalent to the claim that ε_0 is well-founded. Recall that we mentioned in Section 3.4.4 that PA cannot justify the well-foundedness of ε_0 . Does that mean that Goodstein showed that the extended Goodstein’s theorem cannot be proved in PA? Not quite; the theorem cannot even be stated in PA as it is impossible to formalize. The theory lacks the expressive power to represent arbitrary infinite sequences.

The simple version we considered can however be formalized in PA easily (as we are working with only one specific sequence), and it seems much “tamer” than the extended version. For a long time it was unknown whether the simple version could be proved without using tools beyond the reach of PA (such as the well-foundedness of ε_0). Almost forty years later, Kirby and Paris [25] showed that no such finitary proof is possible.

Theorem 3.74 (Kirby–Paris, 1982). *Goodstein’s theorem is true, but unprovable in Peano arithmetic.*

Intuition. Even though for every fixed natural number m we have

$$\text{PA} \vdash (\exists n)(\overline{m}_n = 0),$$

where $\overline{m} = S(S(\dots S(0)\dots))$ repeated m times, (for each fixed m , PA can verify the finite descent of the Goodstein sequence by explicit computation, showing that it terminates after some finite number n of steps), it also holds that

$$\text{PA} \not\vdash (\forall m)(\exists n)(m_n = 0).$$

Proof sketch. Kirby and Paris started with a statement of the form

$$(\forall a)(\forall b)(\exists c) \varphi(a, b, c)$$

that was known to be independent¹⁰ in PA. Then they constructed a nonstandard model M of PA containing a nonstandard $b_0 \in M$ such that

$$M \models \neg(\exists y) \varphi(1, b_0, y).$$

That is, the above-mentioned independent statement does not hold in this model. To finish the proof, they showed that if Goodstein’s theorem could be proved in PA, then there would exist a (very large) number $e \in M$ (satisfying $m_e = 0$ for some carefully chosen m defined using the nonstandard b_0) such that $\varphi(1, b_0, e)$ holds (inside M), which is a contradiction. \square

¹⁰The formula $\varphi(a, b, c)$ talks about certain large ordinals $\alpha < \varepsilon_0$, utilizing the fact that these numbers have a finite hereditary Cantor normal form, and can therefore be formalized in PA.

This result should be surprising. It can be shown that PA is equivalent¹¹ to the theory of finite sets; that is, ZFC with the axiom of infinity replaced by the axiom “there are no limit ordinals.” From this, one can prove that all sets are finite. The Kirby–Paris theorem asserts that accepting an axiom about infinite sets changes what we can prove about finite sets.

3.5.2 Kirby–Paris Hydra Game

In Greek mythology, Hercules is a hero who must complete twelve difficult tasks known as the Labours of Hercules. His second labour is to destroy the Lernaean Hydra, a nine-headed serpent living in the lake of Lerna. Raised by the goddess Hera to defeat Hercules, the Hydra is a particularly challenging enemy, as whenever one of its heads is cut off, two new heads grow back in its place. Realizing he cannot win by force alone, Hercules calls upon his nephew Iolaus for help. They devise a strategy where Hercules cuts off a head and Iolaus immediately burns the stump with a torch, preventing the regeneration. This allows them to defeat the creature, with Hercules burying the final, immortal head under a massive rock.

Mathematicians have adapted this tale into the *Hydra game*, a problem where one iteratively removes leaves from a rooted tree that grows in response to being “attacked,” with the goal of reducing the tree to a single vertex. There are several versions of this game, and we will focus on the version introduced by Kirby and Paris in 1982 [25].

A *hydra* is a finite rooted tree, usually drawn with the root at the bottom. A *head* of the hydra is a leaf together with its attached edge. A *battle* between Hercules and a given hydra is divided into stages, starting at stage one. During stage n , Hercules chops off one head of the hydra. The hydra then grows n new “head segments” in the following manner:

- From the node that used to be attached to the head which was just chopped off, move along the length of one edge towards the root; that is, move to the grandparent of the chopped off node.
- From this node, sprout n replicas of that part of the hydra (after decapitation), which is “above” the edge just traveled.
- If the head just chopped off was attached to the root, no new head is grown.

Hercules *wins* if after some finite number of stages, the hydra is dead; that is, nothing is left of the hydra but its root. A *strategy* is a function which determines for Hercules which head to chop off at each stage of any battle. A strategy is *winning* if it ensures that Hercules wins against every hydra. It is not hard to show that a winning strategy exists (for example, always targeting one of the highest positioned heads). More surprisingly, Hercules cannot help winning:

Theorem 3.75. *Every strategy is a winning strategy.*

¹¹A theorem is provable in PA if and only if an equivalent theorem is provable in the theory of finite sets. See [33] for an intuition why this could be true and for more references.




Figure 4: The Hydra game, as presented by Kirby and Paris [25], if at each stage Hercules decides to chop off the head marked with an arrow.

Proof sketch. The idea of the proof is very similar to that of Goodstein’s theorem. We create a sequence of ordinal numbers $\mu_n < \varepsilon_0$ that is strictly decreasing, and its elements are positive if and only if the hydra is still alive. Since ε_0 is well-founded, it admits no infinite decreasing chains; thus, the sequence must eventually terminate, and the hydra will die with it.

Suppose we assign to each hydra an ordinal $\alpha < \varepsilon_0$. For any strategy σ , we can define a function $H_\sigma(\alpha, n) : \varepsilon_0 \times \omega \rightarrow \varepsilon_0$ that maps the ordinal of the hydra represented by α , together with a stage number n , to the ordinal of the hydra that results from Hercules chopping off the head of α as specified by the strategy σ at stage $n - 1$ (that is, n new “heads” are grown).

To prove the theorem, one only has to show that for any strategy σ , any $0 < \alpha < \varepsilon_0$, and any $n \in \omega$, it holds that $H_\sigma(\alpha, n) < \alpha$. \square

Exercise 3.76. Finish the proof by assigning to each hydra an ordinal α such that $\alpha > 0$ if and only if the hydra is still alive, and $H_\sigma(\alpha, n) < \alpha$.

Kirby and Paris then proved a second independence statement.

Theorem 3.77 (Kirby–Paris, 1982). *The statement “every recursive strategy is a winning strategy” is not provable in Peano arithmetic.*

Remark. “Recursive strategy” means any strategy that can be implemented by an algorithm. The reason why the theorem does not simply say “every strategy” is that it is not possible to represent an arbitrary infinite object in PA. The restriction to recursive strategies ensures that each one of them can be formalized by a (finite) Turing machine.

3.5.3 Fundamental Sequences

Definition 3.78. The *fundamental sequence* of a countable limit ordinal α is an increasing sequence $\alpha[0] < \alpha[1] < \dots$ such that $\alpha = \sup\{\alpha[n] \mid n < \omega\}$.

It is clear from the definition that the same ordinal α may have multiple fundamental sequences, but there is usually a “standard” one.

Definition 3.79. The following is a common assignment of fundamental sequences to limit ordinals $\alpha \leq \varepsilon_0$ (the first fixed point of $\xi \mapsto \omega^\xi$). Corollary 3.35 claims that every $\alpha < \varepsilon_0$ can be expressed as

$$\alpha = \omega^{\gamma_0} + \omega^{\gamma_1} + \cdots + \omega^{\gamma_k},$$

where $\gamma_0 \geq \gamma_1 \geq \cdots \geq \gamma_k$ (by expanding the Cantor normal form of α). Since $\alpha < \varepsilon_0$, each exponent γ_i satisfies $\gamma_i < \alpha$. Also, note that since α is limit, we have that $\gamma_k \neq 0$. We define $\alpha[n]$ inductively as

- (i) $\omega^{\gamma+1}[n] := \omega^\gamma \cdot (n + 1)$,
- (ii) $\omega^\gamma[n] := \omega^{\gamma[n]}$ for limit ordinals γ
- (iii) $(\omega^{\gamma_0} + \cdots + \omega^{\gamma_k})[n] := \omega^{\gamma_0} + \cdots + (\omega^{\gamma_k}[n])$ for $\gamma_0 \geq \cdots \geq \gamma_k$,
- (iv) $\varepsilon_0[n] = \gamma_n$, where $\gamma_0 = 1$ and $\gamma_{n+1} = \omega^{\gamma_n}$. That is, $\gamma_n = \omega \uparrow n$ is a ω -tower of height n .

Example. Some simple fundamental sequences are

- $\omega[n] = \omega^0 \cdot (n + 1) = n + 1$,
- $\omega^\omega[n] = \omega^{\omega[n]} = \omega^{n+1}$,
- $\omega^{\omega+5}[n] = \omega^{\omega+4} \cdot (n + 1)$,
- $(\omega^{\omega+2} + \omega^\omega \cdot 5)[n] = \omega^{\omega+2} + \omega^\omega \cdot 4 + (\omega^\omega[n]) = \omega^{\omega+2} + \omega^\omega \cdot 4 + \omega^{n+1}$

Fundamental sequences can also be assigned to limit ordinals larger than ε_0 , but this becomes much more complicated. The next common approach after Cantor normal form is to utilize the Veblen function described in Section 3.3.2. One can show that all ordinals $\alpha < \Gamma_0$ have a unique normal form in terms of Veblen functions; for details, see Section 8 of [13]. This normal form can then be utilized to define fundamental sequences for all limit ordinals $\alpha < \Gamma_0$. If you are interested in how the formulas look, see [40].

3.5.4 Fast-Growing Hierarchy

One of the use cases of fundamental sequences is the definition of hierarchies of functions $f_\alpha : \omega \rightarrow \omega$, where each function grows faster than the previous one.

Definition 3.80 (Fast-growing hierarchy¹²). For ordinals $\alpha \leq \varepsilon_0$, we define functions from natural numbers to natural numbers $f_\alpha : \omega \rightarrow \omega$ as follows:

- (i) $f_0(n) := n + 1$,
- (ii) $f_{\alpha+1}(n) := f_\alpha^n(n) = f_\alpha(f_\alpha(\dots f_\alpha(n)\dots))$, where f_α is composed n times,

¹²Numberphile has a VIDEO where they explore the growth rates of some extremely fast-growing functions (Graham’s iteration and the TREE sequence).

(iii) $f_\alpha(n) := f_{\alpha[n]}(n)$ for limit ordinals α .

Remark. Martin Löb and Stanley Wainer introduced this hierarchy in 1970s as a generalization of the Grzegorzcyk hierarchy, which only considered $\alpha < \omega$.

Remark. The ordinal ε_0 is not important in the definition; we could use any other large countable ordinal μ if we had fundamental sequences for all limit ordinals $\alpha \leq \mu$. Also, note that the values of the functions f_α may differ based on the chosen fundamental sequences for limit ordinals. The idea is that the values will be asymptotically the same.

Example. Some fast-growing hierarchy functions are

- $f_1(n) = 2n, \quad f_2(n) = 2^n \cdot n,$
- $f_3(n) > 2 \uparrow n$ is a 2-tower of height n ; this is called *tetration*.
- $f_\omega(n) > A(n, n)$, where A is the *Ackermann function*.¹³

Definition 3.81. A function $f : \omega \rightarrow \omega$ is

- (a) *total* if it is defined for every $n \in \omega$; *partial* otherwise,
- (b) *recursive* (or *computable*) if there exists an algorithm that for any given input n halts precisely when $f(n)$ is defined and outputs $f(n)$,
- (c) *primitive recursive* if there exists an algorithm that does not use recursion (it only uses loops and conditions) and for any given input n halts precisely when $f(n)$ is defined and outputs $f(n)$.

It is important to realize that all the functions f_α defined above are recursive, as for each of them, there is a straightforward algorithm to compute the value $f_\alpha(n)$ for any n by following α down to smaller ordinals using the fundamental sequences $\alpha[n]$. Even if the ordinal α is infinite, a Turing machine set to compute f_α would eventually find $f_\alpha(n)$ in a finite amount of time.

We have already mentioned that we can define fundamental sequences for ordinals $\alpha \leq \Gamma_0$ (and there are ways to go beyond), and the functions f_α would remain recursive as long as we have a well-defined, recursive method for choosing the sequence $\alpha[n]$. This raises the question: “How far can we go?” Eventually, it must become impossible to choose $\alpha[n]$ in a recursive manner because there are only countably many recursive functions (each corresponds to a Turing machine, which can be represented as a finite sequence of natural numbers, and there are only countably many of those), while there are uncountably many countable ordinals, as we have shown in Section 3.3.3.

Therefore, at some point, we will reach a countable “non-recursive” ordinal Λ for which f_Λ can no longer be recursive. The first such ordinal is called the *Church–Kleene ordinal* and it is denoted by ω_1^{CK} . All ordinals $\alpha < \omega_1^{\text{CK}}$ are *recursive*; that means there exists a recursive well-ordering $<_\alpha$ of ω with type α . How big is the Church–Kleene ordinal? It is certainly much, much larger than Γ_0 , but it is still countable, so nothing compared to ω_1 .

¹³Computerphile has a VIDEO about the Ackermann function.

Ackermann Function

The previously mentioned Ackermann function is a famous total recursive function that is not primitive recursive. It is defined as follows.

Algorithm 1: Ackermann function $A(m,n)$

Function Ackermann(m,n):

```
  if  $m = 0$  then
    | return  $n + 1$ ;
  else
    | if  $n = 0$  then
    | | return Ackermann( $m - 1, 1$ );
    | else
    | | return Ackermann( $m - 1, \text{Ackermann}(m, n - 1)$ );
```

Exercise 3.82. Convince yourself that the algorithm above always halts.

Definition 3.83. We say that a function $f : \omega \rightarrow \omega$ dominates a function $g : \omega \rightarrow \omega$ if, for all sufficiently large n , we have $f(n) > g(n)$.

The fast-growing hierarchy is sometimes referred to as the Grzegorzcyk hierarchy due to the following theorem.

Fact 3.84 (Grzegorzcyk, 1953). *Any primitive recursive function is eventually dominated by some f_n for $n \in \omega$.*

This means that the functions f_n for $n \in \omega$ measure the growth rates of all primitive recursive functions.

Corollary 3.85. *Since the Ackermann function is not primitive recursive, $A(n,n)$ dominates all of these. It can also be shown that $A(n,n)$ is dominated by $f_\omega(n)$, so the Ackermann function lies at the very edge between primitive and non-primitive recursiveness.*

Hardy Hierarchy

Definition 3.86 (Hardy hierarchy, 1904). For ordinals $\alpha \leq \varepsilon_0$, we define functions from natural numbers to natural numbers $H_\alpha : \omega \rightarrow \omega$ as follows:

- (i) $H_0(n) := n$,
- (ii) $H_{\alpha+1}(n) := H_\alpha(n + 1)$,
- (iii) $H_\alpha(n) := H_{\alpha[n]}(n)$ for limit ordinals α .

Example. Some simple Hardy functions are

- $H_1(n) = H_0(n + 1) = n + 1$,
- $H_k(n) = n + k$,
- $H_\omega(n) = H_{\omega[n]}(n) = H_{n+1}(n) = n + (n + 1) = 2n + 1$.

Exercise 3.87. Try computing $H_{\omega+\omega}(n)$, $H_{\omega \cdot k}(n)$ and $H_{\omega \cdot \omega}$.

The Hardy hierarchy seems to be much slower than the fast-growing hierarchy; they are in fact related by $f_\alpha \sim H_{\omega^\alpha}$ for all $\alpha \leq \varepsilon_0$. However, notice that the Hardy hierarchy “catches up” at $\alpha = \varepsilon_0$ (since $\varepsilon_0 = \omega^{\varepsilon_0}$) in the sense that

$$f_{\varepsilon_0}(n-1) \leq H_{\varepsilon_0}(n) \leq f_{\varepsilon_0}(n+1).$$

This means that the two hierarchies can often be treated as equal.

Fact 3.88 (Schwichtenberg–Wainer, c. 1972). *The total recursive functions that can be proved total by Peano arithmetic are exactly those that are eventually dominated by some f_α (or equivalently some H_α) for $\alpha < \varepsilon_0$.*

Remark. By “PA can prove that f is total,” we mean that PA can prove that the algorithm defining f always terminates.

Corollary 3.89. *PA cannot prove that neither f_{ε_0} nor H_{ε_0} are total.*

This is a “computational” counterpart of Gentzen’s consistency proof we saw in Section 3.4.4. Gentzen showed that PA cannot prove that ε_0 is well-founded. Intuitively, it should not be able to prove that the recursive definition of H_{ε_0} is total (that every recursive call eventually reaches the base case H_0).

Connection to Goodstein’s Theorem

Definition 3.90 (Goodstein function). The function $\mathcal{G} : \omega \rightarrow \omega$ mapping each natural number m to the length of the Goodstein sequence starting in m is called the *Goodstein function*. That is, if n is the first index where $m_n = 0$, then $\mathcal{G}(m) = n + 1$ (since we index from zero).

Example. The first few values of the Goodstein function are

- $\mathcal{G}(1) = 2, \quad \mathcal{G}(2) = 4, \quad \mathcal{G}(3) = 6,$
- $\mathcal{G}(4) = 3 \cdot 2^{402653211} - 2 > 10^{10^8} = 10^{100000000},$
- $\mathcal{G}(5) > 10^{10^{10^{10000}}}, \quad \mathcal{G}(12) > \text{Graham’s number}.$ ¹⁴

Observation 3.91. *The Goodstein function \mathcal{G} dominates f_α for every $\alpha < \varepsilon_0$.*

Proof. Notice that we could reformulate Goodstein’s theorem as follows: the function $\mathcal{G} : \omega \rightarrow \omega$ is total. Now, if \mathcal{G} was dominated by some f_α , then PA could prove that \mathcal{G} is total, thus proving Goodstein’s theorem. \square

Corollary 3.92. *The Goodstein function is not primitive recursive since it dominates the Ackerman function $A(n,n) \sim f_\omega(n)$.*

Fact 3.93 (Cichon, 1983). *The Goodstein function is dominated by H_{ε_0} (and thus also by f_{ε_0}). More precisely,*

$$\mathcal{G}(n) = H_{R_2^\omega(n+1)}(1) - 1,$$

where $R_2^\omega(n)$ is the result of writing n in hereditary base-2 notation and then replacing all 2s with ω (as we did when proving Goodstein’s theorem).

¹⁴Numberphile has a VIDEO where Ron Graham himself explains the Graham’s number.

It might not be immediately obvious that this is dominated by H_{ε_0} , but consider what happens when for example $n = 15 = 2^{2^2} - 1$. Then

$$\mathcal{G}(15) = H_{R_2^{\omega}(15+1)}(1) - 1 = H_{\omega^{\omega}}(1) - 1, \quad H_{\varepsilon_0}(15) = H_{\omega \uparrow 15}(15),$$

where $\omega \uparrow 15$ denotes the ω -tower of height 15.

Fact 3.94 (Caicedo, 2007). *If $n = 2^{m_1} + 2^{m_2} + \dots + 2^{m_k}$ is the base-2 representation of n with $m_1 > m_2 > \dots > m_k$, then*

$$\mathcal{G}(n) = f_{\alpha_1}(f_{\alpha_2}(\dots f_{\alpha_k}(3)\dots)) - 2,$$

where $\alpha_i = R_2^{\omega}(m_i)$.

4 Cardinal Numbers

Similar to how ordinal numbers represent order types of well-ordered sets, *cardinal numbers* represent *sizes* of well-ordered sets. Historically, these were two different classes of abstraction, but for us, cardinal numbers will be a special type of ordinal numbers.

4.1 Cardinals as Sizes of Well-Ordered Sets

We know that sizes of sets can be compared using injective mappings and that the relation $x \approx y$ is an equivalence on the class of all sets. We would now like to take the step from only *comparing* the sizes of sets to *quantifying* them. The question is, how do we represent classes

$$\{y \mid y \approx x\}$$

of all sets with the same size as x ? We would like to define a mapping that assigns to each set x a set $|x|$ so that for any two sets x and y , we would have

$$x \approx y \iff |x| = |y|. \tag{4.1}$$

It would be ideal if, in addition, $x \approx |x|$. If such a mapping exists, then we call the set $|x|$ the *cardinality* of the set x .

It is fairly easy to define the cardinalities of certain classes of sets. For example, if x is finite, then there is a unique natural number n such that $x \approx n$, and we can set $|x| := n$. Similarly, if x is countable, then $|x| := \omega$. We are beginning to see a pattern: if x can be well-ordered, choose $|x|$ as the order type of one of its orderings; but which one? In order for (4.1) to hold, we have to choose $|x|$ as the *least* such order type. This is the key to defining cardinal numbers.

Definition 4.1 (Cardinal numbers). An ordinal number κ is a *cardinal number* if for all ordinals $\alpha < \kappa$ there is no injection $\kappa \rightarrow \alpha$. Equivalently, if

$$\alpha < \kappa \implies \alpha \prec \kappa.$$

We denote the *class of all cardinal numbers* by Cn . ($\text{Cn} \subset \text{On}$).

We say that a cardinal number κ is the *cardinality* of a set x , and we write $|x| = \kappa$, if there exists a bijection $x \rightarrow \kappa$; in other words, if $x \approx \kappa$.

Observation 4.2. For cardinal numbers κ and λ it holds that

(a) $\kappa < \lambda \iff \kappa \prec \lambda$,

(b) $\kappa = \lambda \iff \kappa \approx \lambda$,

Observation 4.3. If the sets x and y have their cardinalities defined, then

(a) $x \approx y \iff |x| = |y|$,

(b) $x \approx |x|$.

Observation 4.4. The cardinality $|x|$ is defined $\iff x$ can be well-ordered.

Proof. Our earlier discussion and the way we defined cardinals show that if x can be well-ordered, then $|x|$ is defined (as the least order type of the well-orderings of x). On the other hand, if $|x| = \kappa$ is defined, then we can well-order x by inheriting the order of κ . \square

This demonstrates that if we want to have $|x|$ defined for every set x , then we need to accept the axiom of choice (which allows us to well-order any set).

Remark. There actually is a way to define a map assigning to every set x a set $|x|$ such that (a) holds *without* using the axiom of choice. This is achieved via the axiom of foundation,¹⁵ but we lose property (b). A. Lévy (1969) showed that without either the axiom of foundation or the axiom of choice, it is impossible to define a map assigning each set x a set $|x|$ such that (a) holds. D. Pincus (1974) showed that without the axiom of choice, there exists no mapping satisfying both (a) and (b) for all sets x .

Example. Some basic properties of Cn are:

- every $n \in \omega$ and ω are cardinal numbers,
- if $\alpha \geq \omega$, then $\alpha + 1$ is not a cardinal number,
- thus every cardinal number $\kappa \geq \omega$ is a limit ordinal number,
- but not every limit ordinal number is a cardinal number — for example, we have $\omega + \omega > \omega$, but $\omega + \omega \approx \omega$, so $\omega + \omega$ is not a cardinal number,
- by the same logic, if $\alpha > \omega$ is countable, then it is not a cardinal number,
- the first uncountable ordinal ω_1 we encountered in Section 3.3.3 is a cardinal number since every $\alpha < \omega_1$ is countable.

Lemma 4.5. If $A \subseteq \text{Cn}$ is a set of cardinal numbers, then $\sup(A) = \bigcup A$ is also a cardinal number. In other words, the class $\text{Cn} \subset \text{On}$ is closed.

¹⁵Briefly, this is done by showing that the axiom of foundation implies that the relation \in is well-founded on \mathbf{V} , allowing us to define a hierarchy of sets $V_0 := \emptyset$, $V_{\alpha+1} := \mathcal{P}(V_\alpha)$ and $V_\lambda := \bigcup_{\alpha < \lambda} V_\alpha$ for limit λ , such that $\mathbf{V} = \bigcup \{V_\alpha \mid \alpha \in \text{On}\}$. Now, for any set x denote $\varrho(x) := \min\{\alpha \mid x \subseteq V_\alpha\}$ and define $|x|$ as the set $\{y \mid y \approx x \wedge (\forall z)(z \approx x \Rightarrow \varrho(y) \leq \varrho(z))\}$.

Proof. Since A is a set of ordinals, according to Lemma 3.4, $\sigma := \sup(A) = \bigcup A$ is an ordinal number. We need to show the implication

$$\alpha < \sigma \implies \alpha \prec \sigma.$$

Let $\alpha < \sigma$. Then $\alpha \in \sigma = \bigcup A$, so $\alpha \in \kappa$ for some cardinal $\kappa \in A$. Because κ is a cardinal number and $\alpha < \kappa$, it must be that $\alpha \prec \kappa$. Furthermore, because $\kappa \in A$, we have $\kappa \subseteq \bigcup A = \sigma$ and so $\kappa \preceq \sigma$. The Cantor–Bernstein theorem together with $\alpha \prec \kappa \preceq \sigma$ implies that $\alpha \prec \sigma$. \square

Theorem 4.6. *For every cardinal there exists a larger cardinal.*

Proof. Suppose for contradiction that κ is the largest cardinal. Then, for every ordinal $\alpha \geq \kappa$, there exists a bijection $\alpha \rightarrow \kappa$, so κ can be well-ordered according to the type α . Corollary 2.12 implies that if r_α, r_β are well-orderings of κ with types $\alpha \neq \beta$, then $r_\alpha \neq r_\beta$ because α and β are not isomorphic.

Notice that every well-ordering of κ is a subset of $\kappa \times \kappa$, so an element of $\mathcal{P}(\kappa \times \kappa)$. For $\alpha \geq \kappa$, denote by $R_\alpha \in \mathcal{P}(\mathcal{P}(\kappa \times \kappa))$ the set of all well-orderings of κ with type α . Finally, notice that for $\alpha \neq \beta$ we have $R_\alpha \neq R_\beta$, so we can construct an injection $\alpha \mapsto R_\alpha$ that maps the proper class $\text{On} \setminus \kappa$ into the set $\mathcal{P}(\mathcal{P}(\kappa \times \kappa))$, which contradicts the axiom schema of replacement. \square

Remark. Alternatively, notice that if κ is a cardinal, then the Hartogs number (see Theorem 3.52) of κ is also a cardinal, and it is larger than κ .

Corollary 4.7. *The class of all cardinal numbers Cn is a proper class.*

Proof. If it were a set, then by Lemma 4.5, its supremum would be the largest cardinal number, which is impossible. \square

Definition 4.8. The *successor* of a cardinal κ is the smallest cardinal larger than κ , and we denote it by κ^+ . Furthermore, we say that κ is the *predecessor* of κ^+ . Finally, $\lambda > 0$ is a *limit* cardinal if it has no predecessor.

4.2 Infinite Cardinals \aleph_α

Finite cardinals are not very interesting since they are the same as the natural numbers. We are interested in the cardinalities of infinite sets. We have shown that the class $\text{Cn} \setminus \omega$ is closed and proper. By Lemma 3.14 there exists a unique bijective normal ordinal function $\aleph : \text{On} \rightarrow \text{Cn} \setminus \omega$ enumerating the cardinals that measure the sizes of infinite (well-orderable) sets. Cantor introduced the symbol \aleph (“aleph”), the first letter of the Hebrew alphabet, to denote this function.

Definition 4.9. The unique normal function mapping On onto the class of all infinite cardinals is denoted by \aleph , and its values $\aleph(\alpha)$ are denoted by \aleph_α .

The smallest infinite cardinal number ω (the size of countable sets) is denoted by \aleph_0 and read as “aleph null.”

Observation 4.10. *These aleph numbers are exactly the omega numbers we have discovered in Section 3.3.3. That is, $\aleph_\alpha = \omega_\alpha$ and*

$$\aleph_\alpha = \{\beta \in \text{On} \mid |\beta| < \aleph_\alpha\}.$$

The cardinal \aleph_α is the first ordinal number with cardinality \aleph_α .

Remark. Historically, ordinals and cardinals were not concrete sets but abstract concepts: ordinals described well-ordering types, while cardinals measured size. This distinction led to the development of two parallel notation systems, ω_α and \aleph_α . Von Neumann’s 1923 set-theoretic definition of ordinals unified these ideas by providing canonical representatives for ordinal types. Today, we often write ω_α when considering the cardinal \aleph_α viewed as an ordinal with its well-order.

Observation 4.11. $\aleph_0 = \omega$ is a limit cardinal, and for $\alpha > 0$ we have that

$$\aleph_\alpha \text{ is a limit cardinal} \iff \alpha \text{ is a limit ordinal.}$$

Proof. We use the fact that a cardinal is limit when it is not isolated and prove the statement by contraposition. The claim follows from the simple observation that when $\alpha = \beta + 1$ is an isolated ordinal, then $\aleph_\alpha = \aleph_{\beta+1} = \aleph_\beta^+$. And when $\aleph_\alpha = \aleph_\beta^+$ is an isolated cardinal, then $\alpha = \beta + 1$. \square

Observation 4.12. If α is an ordinal and ξ is a limit ordinal, then

- (a) $\alpha \leq \aleph_\alpha$, ... \aleph is a normal function
- (b) there exist ordinals α such that $\alpha = \aleph_\alpha$, ... see Theorem 3.16
- (c) ω_α is a limit ordinal, ... it is an infinite cardinal
- (d) $\aleph_\xi = \sup\{\aleph_\alpha \mid \alpha < \xi\}$... \aleph is a normal function

Theorem 4.13. For every ordinal α , it holds that $|\aleph_\alpha \times \aleph_\alpha| = \aleph_\alpha$.

To prove this theorem, we first define a suitable well-ordering of $\aleph_\alpha \times \aleph_\alpha$.

Definition 4.14. We define the *maximo-lexicographic* order of $\text{On} \times \text{On}$ as

$$(\alpha_1, \beta_1) \sqsubset (\alpha_2, \beta_2) \iff \begin{cases} \max\{\alpha_1, \beta_1\} < \max\{\alpha_2, \beta_2\}, \text{ or} \\ \max\{\alpha_1, \beta_1\} = \max\{\alpha_2, \beta_2\} \wedge \alpha_1 < \alpha_2, \text{ or} \\ \max\{\alpha_1, \beta_1\} = \max\{\alpha_2, \beta_2\} \wedge \alpha_1 = \alpha_2 \wedge \beta_1 < \beta_2. \end{cases}$$

Intuition. Picture the product $\text{On} \times \text{On}$ as a grid whose horizontal axis is indexed by α and whose vertical axis is indexed by β . Every point (α_i, β_j) sits in a “right-angle” band determined by the value $\max\{\alpha_i, \beta_j\}$. These bands consist of all points whose coordinates share the same maximum, and the bands themselves move outward from the origin. The ordering \sqsubset simply compares points by the outward distance of the bands they belong to. Once two points lie in the same band, we break ties lexicographically. First, compare the α -coordinates. If those agree, compare the β -coordinates. Thus, inside each strip, the ordering “runs along the top” from left to right, and only then “climbs upward” on the right edge. It might be helpful to draw this on a piece of paper.

Exercise 4.15. Prove that the ordering defined above is a well-ordering.

Exercise 2.19 implies that the well-ordered class $(\text{On} \times \text{On}, \sqsubset)$ is order-isomorphic to On . Under this well-ordering, each $\alpha \times \alpha$ is an initial segment of $\text{On} \times \text{On}$; the induced well-ordering of $\alpha \times \alpha$ is called the *canonical well-ordering* of $\alpha \times \alpha$.

Exercise 4.16. Prove that $\omega \times \omega \approx \omega$ by showing that $(\omega \times \omega, \sqsubset)$ and $(\omega, <)$ are order-isomorphic, using the result of Exercise 2.18. Although one can also show that $\omega \times \omega \approx \omega$ using prime-number encoding, this set-theoretic approach is conceptually cleaner, as it does not require arithmetic.

Proof of Theorem 4.13. By transfinite induction on α . If $\alpha = 0$, then we have the countable case $\omega \times \omega \approx \omega$, which holds by the previous exercise. Suppose that $\alpha > 0$ and consider the canonical well-ordering of $\omega_\alpha \times \omega_\alpha$. Since this is a well-ordering, it is isomorphic to a unique ordinal η , and we claim that $\eta = \omega_\alpha$. Clearly $\omega_\alpha \leq \omega_\alpha \times \omega_\alpha \approx \eta$. This implies that $\omega_\alpha \leq \eta$ since ω_α is a cardinal number. Suppose for contradiction that $\omega_\alpha < \eta$; that is, $\omega_\alpha = (\leftarrow, \omega_\alpha) \subset \eta$ is isomorphic to an initial segment $(\leftarrow, (\gamma, \delta))$ of $(\omega_\alpha \times \omega_\alpha, \sqsubset)$, where $(\gamma, \delta) \in \omega_\alpha \times \omega_\alpha$.

Let $\xi = \max\{\gamma, \delta\} + 1$ and notice that $(\leftarrow, (\gamma, \delta)) \subseteq \xi \times \xi$. Since ω_α is a cardinal and $\xi < \omega_\alpha$, there exists $\beta < \alpha$ such that $|\xi| = \omega_\beta$. By the induction hypothesis,

$$|\xi \times \xi| = |\omega_\beta \times \omega_\beta| = \omega_\beta < \omega_\alpha.$$

This is a contradiction since ω_α is isomorphic to an initial segment of $\xi \times \xi$, and we would have $\omega_\alpha < \omega_\alpha$. \square

Corollary 4.17. *If x can be well-ordered, then $x \times x \approx x$. By induction, this also holds for any finite Cartesian product $x \times \cdots \times x \approx x$.*

Theorem 4.18. $\text{AC} \iff$ *For every infinite set x , we have $x \times x \approx x$.*

Proof sketch. The direction ‘ \Rightarrow ’ is easy; just well-order x and use Theorem 4.13.

The reverse implication is harder, and we only sketch the proof. We show that if A is an infinite set satisfying $A \times A \approx A$, then we can well-order A , implying the well-ordering principle. Let $j : A \times A \rightarrow A$ be a bijection, and for each $a \in A$ define $C_a := \{j(a, t) \mid t \in A\}$. Notice that the family $\{C_a \mid a \in A\}$ partitions A into “ A many” copies of A . In other words, the sets C_a are pairwise disjoint and satisfy $C_a \approx A$ and $\bigcup_{a \in A} C_a = A$.

Let η be the Hartogs number (see Theorem 3.52) of A ; that is, the first ordinal that does not inject into A . The idea of the proof is to try to use transfinite recursion to build a family of well-ordered subsets $A_\alpha \subseteq A$ for each $\alpha < \eta$ such that the order type of A_α is α , and the sets A_α are pairwise disjoint. Informally, we do this by placing each A_α into the slot reserved by some C_a . Since $C_a \approx A$, the set A_α can always “fit” into C_a , and the sets C_a are pairwise disjoint. Suppose the transfinite recursion succeeds, and we define all pairwise disjoint A_α with isomorphisms $j_\alpha : \alpha \rightarrow A_\alpha$. Then we can define an injection $\eta \setminus \{\emptyset\} \rightarrow A$ as $\alpha \mapsto j_\alpha(0)$, contradicting the choice of η .

Hence, the transfinite recursion must “fail,” and at some step $\alpha < \eta$, we will have already used up all slots C_a . That is, for every $\beta < \alpha$, there is a slot C_{a_β} containing A_β and $\bigcup_{\beta < \alpha} C_{a_\beta} = A$. Then we define a bijection $\alpha \rightarrow A$ as $\beta \mapsto a_\beta$, allowing us to well-order A . \square

Exercise 4.19. Given an infinite set X with cardinality κ , is it possible to construct a system of nested subsets $\langle X_i \mid i \in \kappa \rangle$ such that $X_0 = X$ and if $i > j$ then $X_i \subseteq X_j$, with all $|X_i| = \kappa$ but $\bigcap_{i \in \kappa} X_i = \emptyset$?

Theorem 4.20. *For any infinite cardinal κ it holds that*

- (a) $\kappa \prec \{x \mid x \subseteq \kappa\} = \mathcal{P}(\kappa)$,
- (b) $\kappa \approx \{x \mid x \subseteq \kappa \text{ is finite}\}$.

Proof. The first claim follows from Cantor’s theorem. We will prove the second claim in Theorem 4.62 while assuming AC. A proof without the use of AC can be found in §4 Chapter II of [1]. \square

Definition 4.21. Let κ and λ be cardinals. We define cardinal numbers

- (a) $\kappa + \lambda := |(\{0\} \times \kappa) \cup (\{1\} \times \lambda)|$,
- (b) $\kappa \cdot \lambda = |\lambda \times \kappa|$.

In other words, $\kappa + \lambda$ and $\kappa \cdot \lambda$ are cardinal numbers, which represent the size of the set on the right side of the equation, in contrast to ordinal addition and multiplication, which express the order type of the same set when ordered lexicographically. If we want to highlight this difference, we talk about *cardinal* addition and multiplication.

Observation 4.22. If A and B are disjoint sets with $|A|$ and $|B|$ defined, then

$$|A \cup B| = |A| + |B|, \quad |A \times B| = |A| \cdot |B|.$$

Observation 4.23. Cardinal addition and multiplication are associative, commutative, and distributive. Additionally, when restricted to ω , they are the same as the corresponding ordinal operations.

Recall that ordinal addition and multiplication are associative; however, in general, they are not commutative or right-distributive. This is because they have to “keep track” of the underlying orderings.

Lemma 4.24. If κ and λ are cardinals, and at least one of them is infinite, then $\kappa + \lambda = \max\{\kappa, \lambda\}$. If, in addition, they are nonzero, then $\kappa \cdot \lambda = \max\{\kappa, \lambda\}$.

Proof. Denote $\mu := \max\{\kappa, \lambda\}$. To show addition, consider

$$\begin{aligned} \mu \preceq \kappa + \lambda &\approx (\{0\} \times \kappa) \cup (\{1\} \times \lambda) \\ &\preceq (\{0\} \times \mu) \cup (\{1\} \times \mu) = 2 \times \mu \preceq \mu \times \mu \approx \mu. \end{aligned}$$

where the last ‘ \approx ’ follows from Theorem 4.13. The Cantor–Bernstein theorem now implies that $\kappa + \lambda \approx \mu$, so $\kappa + \lambda = \mu$ (because they are cardinals).

Similarly for multiplication: $\mu \preceq \kappa \cdot \lambda \approx \lambda \times \kappa \preceq \mu \times \mu \approx \mu$. \square

Corollary 4.25. For any ordinals α and β , it holds that

$$\aleph_\alpha + \aleph_\beta = \aleph_\alpha \cdot \aleph_\beta = \max\{\aleph_\alpha, \aleph_\beta\}.$$

Corollary 4.26. If A is an infinite set and $|B| < |A|$, then $|A \setminus B| = |A|$.

Proof. WLOG assume $B \subseteq A$. Since the sets $A \setminus B$ and B are disjoint, we have

$$|A| = |A \setminus B| + |B|.$$

Since $|A| > |B|$ and $|A| = \max\{|A \setminus B|, |B|\}$, we must have $|A \setminus B| = |A|$. \square

Lemma 4.27 (AC). For any set S , we have $|\bigcup S| \leq |S| \cdot \sup\{|A| \mid A \in S\}$.

Proof. Let $\kappa = |S|$ and $\lambda = \sup\{|A| \mid A \in S\}$. Since $\kappa \approx S$, we can enumerate the elements of S as $S = \{A_\alpha \mid \alpha < \kappa\}$. Moreover, each A_α injects into λ ; hence, we can choose an injection $j_\alpha : A_\alpha \rightarrow \lambda$. For an element $a \in \bigcup S$, define

$$\alpha_a := \min\{\alpha < \kappa \mid a \in A_\alpha\}.$$

This number indicates in which A_α does a first appear in. Notice that more elements $a \in \bigcup S$ can have the same number α_a , but that $j_{\alpha_a}(a)$ uniquely identifies a among these elements (since j_{α_a} is injective). This allows us to define an injection $g : \bigcup S \rightarrow \kappa \times \lambda$ as $a \mapsto (\alpha_a, j_{\alpha_a}(a))$. \square

Corollary 4.28 (AC). The union of a collection of \aleph_α sets, each of cardinality at most \aleph_α , has cardinality at most \aleph_α . If, in addition, they are non-empty and disjoint, then the union has cardinality exactly \aleph_α .

4.3 Cofinality and Inaccessible Cardinals

Recall the pigeonhole principle for $\omega = \aleph_0$, which says that ω cannot be partitioned into a finite number of finite sets. Thus, if $A = \bigcup\{A_i \mid i \in I\}$ is countably infinite, then either I or one of A_i has to be countably infinite. How does this generalize to higher cardinals?

Definition 4.29. Let (X, \leq_R) be a partially ordered set. We say that $Y \subseteq X$ is *cofinal* in X (or is a *cofinal subset* of X) with respect to \leq_R if every $x \in X$ is bounded by some $y \in Y$; (that is, $x \leq y$).

Observation 4.30. If Y is cofinal in X , then it contains all maximal elements of X . Moreover, the relation “to be cofinal in” is transitive.

Example. If X has a maximum x , then $\{x\}$ is the smallest cofinal subset of X .

We are usually interested in cofinality in the context of limit ordinals.

Observation 4.31. If α is an ordinal, $A \subseteq \alpha$, and

- (i) $\alpha = \beta + 1$ is isolated, then A is cofinal in $\alpha \iff \beta \in A$,
- (ii) α is limit, then A is cofinal in $\alpha \iff \sup(A) = \alpha$.

Definition 4.32 (Cofinality). The *cofinality* of a limit ordinal α , denoted by $\text{cf}(\alpha)$, is the least ordinal β that is the order type of some $A \subseteq \alpha$ cofinal in α :

$$\text{cf}(\alpha) := \min\{\text{otp}(A) \mid A \subseteq \alpha \wedge \sup(A) = \alpha\}.$$

Observation 4.33. $\text{cf}(\alpha)$ is the least ordinal β such that there is an increasing β -sequence $\langle \alpha_\xi \mid \xi < \beta \rangle$ with limit α . Hence $\text{cf}(\alpha)$ is always a limit ordinal.

Lemma 4.34. For every (limit) ordinal α , we have

$$\text{cf}(\text{cf}(\alpha)) = \text{cf}(\alpha).$$

Proof. Let $\beta = \text{cf}(\alpha)$ and $\gamma = \text{cf}(\text{cf}(\alpha))$, so $\gamma \leq \beta$. From the definition of $\text{cf}(\alpha)$, there exists a cofinal subset $A \subseteq \alpha$ and an increasing function (isomorphism) $f : \beta \rightarrow A$, and a cofinal subset $B \subseteq \beta$ and an increasing function (isomorphism) $g : \gamma \rightarrow B$. Notice that the map $g \circ f : \gamma \rightarrow A$ is an isomorphism of γ and a cofinal subset of α , thus (by the definition of $\text{cf}(\alpha)$), $\gamma \geq \text{cf}(\alpha) = \beta$. \square

Lemma 4.35. *The cofinality of any (limit) ordinal α is always an infinite cardinal. More precisely, it is the length of the shortest sequence with limit α :*

$$\text{cf}(\alpha) = \min\{|A| \mid A \subseteq \alpha \wedge \sup(A) = \alpha\}.$$

Proof. We know that $\text{cf}(\alpha)$ is an infinite ordinal. If $\text{cf}(\alpha)$ were not a cardinal number, then there would be a cardinal $\kappa < \text{cf}(\alpha)$ such that $\kappa = |\text{cf}(\alpha)|$, and a bijection $f : \kappa \rightarrow \text{cf}(\alpha)$. Since $\text{cf}(\alpha)$ is the cofinality of α , there exists a cofinal subset $A \subseteq \alpha$ with order type $\text{cf}(\alpha)$, and an isomorphism $h : \text{cf}(\alpha) \rightarrow A$. The idea is to use f to skip some terms of the $\text{cf}(\alpha)$ -sequence defined by h .

We define a non-decreasing function $g : \kappa \rightarrow A$ as

$$g : \beta \mapsto h(\sup\{f(\delta) \mid \delta < \beta\}).$$

Notice that this is well-defined: suppose there is some $\beta \in \kappa$ such that

$$\sup\{f(\delta) \mid \delta < \beta\} = \sup(f[\beta]) = \text{cf}(\alpha),$$

then $f[\beta]$ is a cofinal subset of $\text{cf}(\alpha)$, and according to the previous lemma, $f[\beta]$ has order type $\text{cf}(\alpha)$. Hence, there is a bijection between $\beta < \kappa$ and $\text{cf}(\alpha)$, contradicting the fact that $\kappa = |\text{cf}(\alpha)|$ is a cardinal.

Clearly $\text{Rng}(g)$ is a cofinal subset of A (and therefore also of α). We claim that it has order type $\gamma \leq \kappa < \text{cf}(\alpha)$, which is a contradiction. Indeed, the function $i : \text{Rng}(g) \rightarrow \kappa$ defined as

$$a \mapsto \min\{\beta \in \kappa \mid g(\beta) = a\}$$

is injective and increasing (since g is non-decreasing), and thus an isomorphism of $\text{Rng}(g)$ and $x := \text{Rng}(i) \subseteq \kappa$. Therefore $\text{Rng}(g)$ and x have the same order type γ , and Lemma 3.3 states that $\gamma \leq \kappa$ because $x \subseteq \kappa$. \square

Corollary 4.36. *For any (limit) ordinal α it holds that $\omega \leq \text{cf}(\alpha) \leq |\alpha|$.*

Example. Some cofinalities we already know:

- $\text{cf}(\omega) = \text{cf}(\omega + \omega) = \text{cf}(\omega \cdot \omega) = \text{cf}(\omega^\omega) = \text{cf}(\varepsilon_0) = \text{cf}(\Gamma_0) = \omega$,
- in general, $\text{cf}(\alpha) = \omega$ for countable (limit) α , because $\omega \leq \text{cf}(\alpha) \leq |\alpha| = \omega$,
- $\text{cf}(\aleph_\omega) = \omega$, since $\aleph_\omega = \sup\{\aleph_\alpha \mid \alpha < \omega\}$,
- $\text{AC}_\omega \implies \text{cf}(\omega_1) = \omega_1$, as Lemma 3.39 implies $\text{cf}(\omega_1) \geq \omega_1$.

Definition 4.37. An infinite cardinal number $\kappa = \aleph_\alpha$ is a

- (a) *regular cardinal* if $\text{cf}(\kappa) = \kappa$,
- (b) *singular cardinal* if $\text{cf}(\kappa) < \kappa$.

Intuition. If ω_α is a regular cardinal, then it is *almost* closed on taking suprema. As long as the length of the sequence is less than \aleph_α , the limit will never reach ω_α . However, if ω_α is singular, then it is possible to reach it from below via a shorter sequence.

We have already seen that $\aleph_0 = \omega$ is regular and \aleph_ω singular. Furthermore, Lemma 4.34 implies that the cofinality of any (limit) ordinal α is always a regular cardinal number. The question is: are there any regular cardinals besides \aleph_0 ? If we assume the axiom of countable choice, then ω_1 is also regular. But what if we are working in bare ZF?

M. Gitik (1979) showed that the statement “ \aleph_0 is the only regular cardinal” cannot be disproved in ZF. In other words, there exists a model of ZF where *every* cardinal κ has a cofinal subset of size less than κ . In particular, this model contains a countable sequence of countable ordinals whose limit is uncountable, and hence Lemma 3.39 cannot be proved in ZF.

Theorem 4.38. *An infinite cardinal number κ is singular \iff there exists a set X such that $\kappa = \bigcup X$, where $|X| < \kappa$ and $|x| < \kappa$ for all $x \in X$.*

Proof. ‘ \implies ’ Let $X \subseteq \kappa$ be a cofinal subset of κ with order type $\text{cf}(\kappa) < \kappa$. Then $\kappa = \bigcup X$ because $\kappa = \sup(X)$, and $\sup(X) = \bigcup X$ since X is a set of ordinals. Additionally, all $\alpha \in X$ satisfy $|\alpha| < \kappa$ because κ is a cardinal.

‘ \impliedby ’ Assume that $\kappa = \bigcup X$ where $|X| < \kappa$ and all $x \in X$ have $|x| < \kappa$. Note that the cardinalities of all these sets are well-defined since they inherit a well-order from κ . If there exists $x \in X$ such that x is cofinal in κ , then $\text{cf}(\kappa) \leq |x| < \kappa$. If none of $x \in X$ are cofinal in κ (and thus $\sup x < \kappa$) define $S := \{\sup x \mid x \in X\}$, and notice that S is cofinal in $\bigcup X = \kappa$. Furthermore, S is a set of ordinals, so it can be well-ordered and has cardinality $|S| \leq |X|$. From this, we have that $\text{cf}(\kappa) \leq |S| \leq |X| < \kappa$. \square

Corollary 4.39 (Pigeonhole principle for cardinals). *If κ is a regular cardinal and $\kappa = \bigcup X$, where $|X| < \kappa$, then there exists $x \in X$ such that $|x| = \kappa$.*

If we further assume the axiom of choice, every set will have a defined cardinality, and we can make a more general claim.

Corollary 4.40 (AC, Pigeonhole principle for infinite sets). *If S is an infinite set with regular cardinality $|S|$, and $S = \bigcup X$, where $|X| < |S|$, then there exists $x \in X$ such that $|x| = |S|$.*

Theorem 4.41 (AC). *Every infinite isolated cardinal $\aleph_{\alpha+1}$ is regular. Thus all infinite singular cardinals are limit.*

Proof. If it were singular, then by Theorem 4.38 it would be the supremum of at most \aleph_α sets of cardinality at most \aleph_α . But Corollary 4.28 implies that such a supremum must have cardinality at most \aleph_α (since suprema and unions of ordinals are the same). \square

Exercise 4.42 (AC). Every infinite singular cardinal κ is the supremum of $\text{cf}(\kappa)$ regular cardinals.

Theorem 4.43. *If $\aleph_\alpha > \aleph_0$ is a limit cardinal (α a limit ordinal), then*

$$\text{cf}(\aleph_\alpha) = \text{cf}(\alpha).$$

Proof. The cardinal \aleph_α is defined as the limit of the sequence $\{\aleph_\beta \mid \beta < \alpha\}$. The claim follows from the observation that the set $\{\aleph_\beta \mid \beta \in A\}$ is cofinal in \aleph_α if and only if A is a cofinal subset of α . This allows us to skip some terms of the sequence; in fact, we only need $\text{cf}(\alpha)$ many terms. \square

We know that in ZFC, every infinite isolated cardinal is regular, but we have not seen any regular limit cardinals besides \aleph_0 . Let's try different limit ordinal indices α and check whether \aleph_α is regular or not.

- $\text{cf}(\aleph_\omega) = \text{cf}(\aleph_{\varepsilon_0}) = \text{cf}(\aleph_{\Gamma_0}) = \omega$,
- $\text{cf}(\aleph_\alpha) = \omega$ for any countable (limit) α ,
- $\text{cf}(\aleph_{\omega_1}) = \text{cf}(\omega_1) = \aleph_1$,
- $\text{cf}(\aleph_{\omega_{\alpha+1}}) = \text{cf}(\omega_{\alpha+1}) = \aleph_{\alpha+1}$ for any isolated cardinal $\omega_{\alpha+1}$.

This clearly is not working. Notice that Theorem 4.43 implies that if $\aleph_\alpha > \aleph_0$ is a regular limit cardinal, then

$$\text{cf}(\aleph_\alpha) = \text{cf}(\alpha) = \aleph_\alpha.$$

Since $\text{cf}(\alpha) \leq \alpha$ and $\aleph_\alpha \geq \alpha$, we conclude that $\alpha = \aleph_\alpha$. Hence, every regular limit cardinal larger than \aleph_0 has to be a fixed point of the aleph function.

Theorem 3.16 allows us to construct fixed points of \aleph quite easily; for example, the first fixed point is the limit of the sequence

$$\kappa_0 = 0, \kappa_{n+1} = \aleph_{\kappa_n} \quad \longrightarrow \quad 0, \aleph_0, \aleph_{\aleph_0}, \aleph_{\aleph_{\aleph_0}}, \aleph_{\aleph_{\aleph_{\aleph_0}}}, \dots$$

Let us denote it κ_0 and notice that it is singular since $\text{cf}(\kappa_0) = \omega$. Similarly, any fixed point constructed using this method will have cofinality ω . Theorem 3.16 allows these fixed points to be arbitrarily large, so we can summarize this as:

Lemma 4.44. *There are arbitrarily large singular cardinals \aleph_α such that $\aleph_\alpha = \alpha$.*

The question is: can a fixed point of \aleph be a regular cardinal? People started using the term *inaccessible* to describe these cardinals, as they cannot be reached from below neither by taking successor cardinals nor by forming suprema of smaller cardinals. Today, we call such numbers *weakly inaccessible*.

Definition 4.45 (Weakly inaccessible cardinal). A cardinal κ is *weakly inaccessible* if it is an uncountable regular limit cardinal.

Fact 4.46. *ZFC cannot prove the existence of weakly inaccessible cardinals.*

Remark. Weakly inaccessible cardinals were first introduced by F. Hausdorff in 1908. Only much later was it shown that the existence of such cardinals cannot be proved in ZFC (provided ZFC is consistent). Therefore, it is consistent with ZFC to assume that every uncountable limit cardinal number is singular.

We will return to the concept of cardinals whose existence cannot be proved in Section 5.

4.4 Cardinal Arithmetic

Cardinal arithmetic is the study of powers, infinite sums, and infinite products of cardinal numbers. In this section, we will always assume AC, so every set is equivalent to a unique cardinal number. Let us first recall that if κ and λ are non-zero cardinals and at least one of them is infinite, then

$$\kappa + \lambda = \kappa \cdot \lambda = \max\{\kappa, \lambda\}.$$

Definition 4.47. For a set a and class B , we define the class of all mappings from a to B as

$${}^a B := \{f \mid f : a \rightarrow B\}.$$

Exercise 4.48. If a and B are sets, then ${}^a B$ is a set.

Definition 4.49 (Cardinal power, AC). For cardinals κ and λ , we define the cardinal number $\kappa^\lambda := |{}^\lambda \kappa|$ as the size of the set of all functions from λ to κ .

So if A and B are sets, then $|{}^A B| = |B|^{|A|}$.

Observation 4.50. In particular, 2^λ is the cardinality of the set of all characteristic functions of subsets of λ , so $2^\lambda = |\mathcal{P}(\lambda)|$.

This also shows why we need AC to define the cardinal power: if 2^ω is defined, then $\mathcal{P}(\omega) \approx \mathbb{R}$ can be well-ordered.

Using this notation, we can formulate Cantor's theorem in ZFC as

$$(\forall \kappa \in \text{Cn}) : 2^\kappa > \kappa, \quad \text{or equivalently as} \quad (\forall \kappa \in \text{Cn}) : 2^\kappa \geq \kappa^+.$$

Exercise 4.51. Show from the definition of cardinal power that

- (a) if $0 < \kappa \leq \mu$ and $\lambda \leq \nu$ then $\kappa^\lambda \leq \mu^\nu$,
- (b) $\kappa^{\mu+\nu} = \kappa^\mu \cdot \kappa^\nu$,
- (c) $(\kappa^\mu)^\nu = \kappa^{\mu \cdot \nu}$.

Hence cardinal power is monotone and follows the usual properties of exponents.

One can show by induction that the cardinal power of natural numbers n and m is the same as the corresponding ordinal power. It is also not difficult to determine the value κ^λ when one of the arguments is a natural number:

Theorem 4.52. For any cardinals κ, λ and $n \in \omega$ it holds that

- (a) $0^0 = 1, \quad \lambda > 0 \implies 0^\lambda = 0,$
- (b) $\kappa^0 = 1, \quad 1^\lambda = 1,$
- (c) if $\kappa \geq \omega$ and $n > 0$, then $\kappa^n = \kappa, \quad \dots$ by induction from $\kappa \cdot \kappa = \kappa$
- (d) if $\lambda \geq \omega$ and $2 \leq \kappa \leq \lambda$, then $\kappa^\lambda = 2^\lambda. \quad \dots$ in particular $\kappa^\kappa = 2^\kappa$

Proof. (d) holds because ${}^\lambda \kappa \subseteq \mathcal{P}(\lambda \times \kappa)$ and from our assumption $\lambda \cdot \kappa = \lambda$, so

$$2^\lambda \leq \kappa^\lambda \leq |\mathcal{P}(\lambda \times \kappa)| = 2^{\lambda \cdot \kappa} = 2^\lambda. \quad \square$$

Example. $\omega^\omega = 2^\omega$, so the Cartesian product of countably many countably infinite sets is uncountable.

4.4.1 Cardinality of the Continuum

It is well-known that $2^\omega = |\mathcal{P}(\omega)| = |\mathbb{R}|$. Since \mathbb{R} is sometimes referred to as “the continuum,” people have started calling the cardinality of these sets the *cardinality of the continuum*, and it is commonly denoted by $\mathfrak{c} := 2^\omega = \beth_1$. Other sets with this cardinality are for example the irrational numbers (as $\mathbb{Q} \approx \omega$) and transcendental numbers (as Cantor famously proved that there are only countably many algebraic numbers).

Exercise 4.53. The set of all open subsets of \mathbb{R} has cardinality \mathfrak{c} . Note that the set of all subsets of \mathbb{R} has cardinality $2^\mathfrak{c}$.

Hint. Each open set is a countable disjoint unions of open intervals.

Exercise 4.54. The set of all continuous functions $f : \mathbb{R} \rightarrow \mathbb{R}$ has cardinality \mathfrak{c} . Note that the set of all real functions has (by definition) size $\mathfrak{c}^\mathfrak{c} = 2^\mathfrak{c}$.

Hint. Realize that if a continuous function f is known on the rational numbers \mathbb{Q} , then it is known everywhere. There is only one way to continuously fill in the gaps between the rational points.

Corollary 4.55. *There exists a real function g that intersects the graph of every continuous function $f : \mathbb{R} \rightarrow \mathbb{R}$.*

Proof. The previous exercise implies that there is a way to bijectively assign to each real number r a continuous function $f_r : \mathbb{R} \rightarrow \mathbb{R}$. We define $g(r) := f_r(r)$. \square

Exercise 4.56. The set of all well-orderings of \mathbb{R} has cardinality $2^\mathfrak{c}$. Note that the set of all binary relations on \mathbb{R} has size $2^{|\mathbb{R} \times \mathbb{R}|} = 2^{\mathfrak{c} \cdot \mathfrak{c}} = 2^\mathfrak{c}$.

Hint. Consider all ordinal types of well-orders of \mathbb{R} , and for each order-type consider the possible ways to order \mathbb{R} using this type.

4.4.2 Infinite Sums and Products

Definition 4.57 (AC). The (infinite) *sum* and *product* of a collection of cardinal numbers $\langle \kappa_i \mid i \in I \rangle$ are the cardinals

$$\sum_{i \in I} \kappa_i := \left| \bigcup_{i \in I} (\{i\} \times \kappa_i) \right|, \quad \prod_{i \in I} \kappa_i := \left| \prod_{i \in I} \kappa_i \right|.$$

Remark. We can define the sum without AC since $\bigcup (\{i\} \times \kappa_i)$ is well-ordered lexicographically, but it is required for the product: it might not be well-orderable, or $\prod \kappa_i$ might be the empty set (Exercise 1.10).

Exercise 4.58. The lexicographic ordering of $\prod_{i \in \omega} 2$ is not a well-ordering.

It is easy to see that the sum and product of cardinals will not change when we permute the indices in I . If λ_i are cardinals such that $\kappa_i \leq \lambda_i$ for all $i \in I$, and if $J \subseteq I$, then

$$\sum_{i \in I} \kappa_i \leq \sum_{i \in I} \lambda_i, \quad \sum_{i \in J} \kappa_i \leq \sum_{i \in I} \kappa_i.$$

Furthermore, if $\langle A_i \mid i \in I \rangle$ is a collection of sets, then

$$\left| \bigcup_{i \in I} A_i \right| \leq \sum_{i \in I} |A_i|, \quad \left| \prod_{i \in I} A_i \right| = \prod_{i \in I} |A_i|.$$

If, in addition, the sets are disjoint, then $\left| \bigcup_{i \in I} A_i \right| = \sum_{i \in I} |A_i|$.

Observation 4.59. *If $\kappa_i = \kappa$ for all $i \in I$, then $\sum \kappa_i = |I| \cdot \kappa$ and $\prod \kappa_i = \kappa^{|I|}$.*

Definition 4.60. If X is a set and λ is a cardinal, define

- (a) $[X]^\lambda := \{x \subseteq X \mid |x| = \lambda\}$, ... subsets of size λ
- (b) $[X]^{<\lambda} = \{x \subseteq X \mid |x| < \lambda\}$ subsets of size $< \lambda$.

If $n \geq k$ are natural numbers, then $[n]^k = \binom{n}{k}$, and $[X]^{<\omega}$ is the set of all finite subsets of X . Furthermore, if $\lambda > |X|$, then $[X]^{<\lambda} = \mathcal{P}(X)$ and $[X]^\lambda = \emptyset$.

Lemma 4.61 (AC). *If X is an infinite set of size $|X| = \kappa$ and $\lambda \in \text{Cn}$, then*

- (a) $|[X]^\lambda| = \kappa^\lambda$, if $\lambda \leq \kappa$, ... if $\lambda > \kappa$, then $|[X]^\lambda| = 0$
- (b) $|[X]^{<\lambda}| = \sum_{\mu < \lambda} \kappa^\mu$, if $\lambda \leq \kappa^+$, ... if $\lambda \geq \kappa^+$ then $|[X]^{<\lambda}| = 2^\kappa$

Proof. (a), WLOG assume $X = \kappa$, so we want to show $[\kappa]^\lambda \approx \kappa^\lambda$. For any set $x \in [\kappa]^\lambda$ there is a map $f \in {}^\lambda \kappa$ such that $x = \text{Rng}(f)$. Using AC, we choose for each $x \in [\kappa]^\lambda$ such a mapping f_x . We have constructed an injection $[\kappa]^\lambda \rightarrow {}^\lambda \kappa$, showing $[\kappa]^\lambda \preceq \kappa^\lambda$. To show the other inequality, realize that every mapping $f : \lambda \rightarrow \kappa$ is an element of $[\lambda \times \kappa]^\lambda$, so $\kappa^\lambda \preceq [\lambda \times \kappa]^\lambda$. But since $\lambda \times \kappa \approx \kappa$, we have $[\lambda \times \kappa]^\lambda \approx [\kappa]^\lambda$.

(b) follows from (a) since the sets $[X]^\mu$ are disjoint for different values of μ . □

Corollary 4.62. *If X is an infinite set, then $|[X]^{<\omega}| = |X|$.*

Proof. This is a direct consequence of (b) and basic cardinal properties:

$$|[X]^{<\omega}| = \sum_{n < \omega} |X|^n = \sum_{n < \omega} |X| = \omega \cdot |X| = \max\{\omega, |X|\} = |X|. \quad \square$$

Definition 4.63 (Weak power). For cardinals κ and λ , define the cardinals

$$\kappa^{<\lambda} := \sum_{\mu < \lambda} \kappa^\mu, \quad \kappa^{\leq \lambda} := \sum_{\mu < \lambda^+} \kappa^\mu.$$

Observation 4.64. *If κ is an infinite cardinal then $\kappa^{\leq \kappa} = \kappa^\kappa = 2^\kappa$.*

Proof. The previous lemma implies that $\kappa^{\leq \kappa} = 2^\kappa$ and $\kappa^\kappa = 2^\kappa$. □

Lemma 4.65 (Infinite sum). *Let $\langle \kappa_i \mid i \in I \rangle$ be a collection of nonzero cardinals. If the set I or at least one of the cardinals κ_i is infinite, then*

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

Proof. Let $\kappa = \sup \kappa_i$. Then $\sum \kappa_i \approx \bigcup (\{i\} \times \kappa_i) \subseteq I \times \kappa \approx |I| \cdot \kappa = \max\{|I|, \kappa\}$. We have shown “ \leq ”. For $\sum_{i \in I} \kappa_i \geq \max\{|I|, \kappa\}$, we show that

- $\sum_{i \in I} \kappa_i \geq \sum_{i \in I} 1 = |I|$, and
- $\sum_{i \in I} \kappa_i \geq \kappa$ since $\sum_{i \in I} \kappa_i \geq \kappa_j$ for all j , so $\sum_{i \in I} \kappa_i$ is an upper bound, and κ is the least upper bound. \square

Remark. Note that “ \geq ” also follows from Lemma 4.27, but we would need AC.

Corollary 4.66. *If $\langle X_i \mid i \in I \rangle$ is a collection of sets such that $|X_i| = \kappa_i$, and $\sup\{\kappa_i \mid i \in I\} \geq |I|$ is an infinite cardinal, then*

$$\left| \bigcup_{i \in I} X_i \right| = \sup_{i \in I} \kappa_i = \sum_{i \in I} \kappa_i.$$

Proof. WLOG all X_i are nonempty. Let $\kappa = \sup \kappa_i$. The previous lemma claims that $\kappa = \sum \kappa_i$. We know that $\left| \bigcup X_i \right| \leq \sum \kappa_i$, and we also know that $\kappa_j \leq \left| \bigcup X_i \right|$ for all j , so $\left| \bigcup X_i \right|$ is an upper bound. Since κ is the least upper bound, we have $\kappa \leq \left| \bigcup X_i \right|$. \square

Observation 4.67. *If κ is an infinite cardinal with cofinality $\text{cf}(\kappa)$, then there are cardinals $\langle \kappa_i \mid i < \text{cf}(\kappa) \rangle$ such that $\kappa_i < \kappa$ for all i and $\kappa = \sum \kappa_i$*

Proof. Let $X = \langle \alpha_i \mid i \in \text{cf}(\kappa) \rangle$ be a cofinal subset of κ . Then $\bigcup \alpha_i = \kappa$, and by Corollary 4.66 $\kappa = \sup |\alpha_i|$. We let $\kappa_i = |\alpha_i|$; since $\alpha_i < \kappa$ and κ is a cardinal, we have $\kappa_i < \kappa$. \square

Corollary 4.68. *A cardinal κ is singular \iff there are cardinals $\lambda < \kappa$ and $\langle \kappa_i \mid i < \lambda \rangle$ such that $\kappa_i < \kappa$ for all i and $\kappa = \sum \kappa_i$.*

Proof. (\implies) By the previous observation. (\impliedby) If such cardinals exist, then $\kappa = \sup\{\kappa_i \mid i < \lambda\}$, meaning that κ has a cofinal subset of size $\lambda < \kappa$. \square

Remark. This is a cardinal arithmetic counterpart of Theorem 4.38.

Theorem 4.69 (Tarski). *Let κ be an infinite cardinal with cofinality $\text{cf}(\kappa)$. If $\langle \kappa_i \mid i < \text{cf}(\kappa) \rangle$ is such that $1 \leq \kappa_i < \kappa$ for all i and $\sup \kappa_i = \kappa$, then*

$$\prod_{i < \text{cf}(\kappa)} \kappa_i = \kappa^{\text{cf}(\kappa)}.$$

Proof. The inequality “ \leq ” is obvious. To prove the other inequality, notice that we can WLOG assume that all $\kappa_i \geq 2$: if the number of i such that $\kappa_i \geq 2$ were less than $\text{cf}(\kappa)$, then we could remove all the $\kappa_j = 1$, and $\sup \kappa_i$ would not change; thus, we would obtain a strictly shorter cofinal sequence.

If $\text{cf}(\kappa) = \kappa$, then $\kappa^{\text{cf}(\kappa)} = 2^\kappa$, and the inequality holds. Assume $\text{cf}(\kappa) < \kappa$. Notice that we can further assume that the sequence $\langle \kappa_i \mid i < \text{cf}(\kappa) \rangle$ is strictly increasing; indeed, we can achieve this by rearranging and thinning the original sequence. Let $\lambda = \text{cf}(\kappa)$ and note that $\kappa = \sum_{i < \lambda} \kappa_i$. We obtain

$$\kappa^\lambda = \left(\sum_{i < \lambda} \kappa_i \right)^\lambda = \sum_{f \in {}^\lambda \lambda} \prod_{\alpha < \lambda} \kappa_{f(\alpha)} \leq \sum_{f \in {}^\lambda \lambda} \prod_{i < \lambda} \kappa_i = 2^\lambda \prod_{i < \lambda} \kappa_i = \prod_{i < \lambda} \kappa_i.$$

Intuitively, the second equality holds from a simple distributive argument, similar to when we want to multiply out an expression such as $(k_0 + k_1 + \dots + k_n)^r$. Formally, $(\sum_{i < \lambda} \kappa_i)^\lambda$ is equivalent to the set of all functions $g : \lambda \rightarrow \lambda \times \kappa$. And $\sum_{f \in {}^\lambda \lambda} \prod_{j < \lambda} \kappa_{f(j)}$ is equivalent to considering the functions $f : \lambda \rightarrow \lambda$, and for each f , considering all the functions $h_f : \lambda \rightarrow \kappa$ such that $h_f(\alpha) \in \kappa_{f(\alpha)}$. Clearly, every pair f, h_f defines a unique function $\lambda \rightarrow \lambda \times \kappa$. Conversely, every function $g : \lambda \rightarrow \lambda \times \kappa$ corresponds to a unique function $g' : \lambda \rightarrow \kappa$ (since $\lambda \times \kappa \approx \kappa$), and for every such g' , there clearly exists a unique pair f, h_f .

The inequality holds since for each $f \in {}^\lambda \lambda$ there exists a strictly increasing $g \in {}^\lambda \lambda$ such that $f(\alpha) \leq g(\alpha)$ for every $\alpha < \lambda$. This follows from the fact that $\lambda = \text{cf}(\kappa)$ is a regular cardinal.

The next equality holds because $|{}^\lambda \lambda| = \lambda^\lambda = 2^\lambda$. The final equality holds since $2^\lambda \leq \prod_{i < \lambda} \kappa_i$ as $\kappa_i \geq 2$ for \square

4.4.3 König's Inequality

König's Inequality talks about the relation of $\sum \kappa_i$ and $\prod \kappa_i$, and generalizes Cantor's theorem. We will need to assume AC to ensure that $\prod \kappa_i$ is well-defined.

Lemma 4.70 (AC). *If $\langle \kappa_i \mid i \in I \rangle$ is a collection of cardinals $\kappa_i \geq 2$, then*

$$\sum \kappa_i \leq \prod \kappa_i.$$

Proof. If $|I| = 0$, then $\sum = 0$ and $\prod = 1$ ($\times \kappa_i$ is the set of some mappings from the empty set, and there is always the empty mapping). If $|I| = 1$, then $\sum = \prod$. If $|I| = 2$, consider cardinals κ_0 and κ_1 . We want to find bijective mappings of κ_0 and κ_1 to disjoint sets $A_0, A_1 \subseteq \kappa_0 \times \kappa_1$. Imagine $\kappa_0 \times \kappa_1$ as a grid with horizontal axis κ_0 and vertical axis κ_1 . We map κ_0 to the first row, and κ_1 to the first column — but since they overlap at $(0,0)$, we map one of the zeroes (for example $0 \in \kappa_1$) to $(1,1)$.

If $|I| \geq 3$, map $\kappa_i \setminus \{0\}$ to the “ i -th edge of the box”, that is the “vector” $(0, \dots, 0, *, 0, \dots, 0)$, where $*$ is a wildcard at index i . And map the zero $0 \in \kappa_i$ to the “point” $(1, \dots, 1, 0, 1, \dots, 1)$, where the 0 is at index i . The formal definition of the mappings $\kappa_i \rightarrow \times \kappa_i$ is left to the reader. \square

Theorem 4.71 (König's inequality, AC). *If $I \neq \emptyset$ and $\kappa_i < \lambda_i$ for all $i \in I$, then*

$$\sum \kappa_i < \prod \lambda_i.$$

König's inequality is a generalization of Cantor's theorem. If $\kappa_i = 1$ and $\lambda_i = 2$ for all i , then König's inequality implies $|I| < 2^{|I|}$.

Proof. Notice that if $\lambda_j = 1$, then $\kappa_j = 0$; so if we remove j from I , the sum and product will not change. Hence, WLOG all $\lambda_i \geq 2$. By the previous lemma:

$$\sum \kappa_i \leq \sum \lambda_i \leq \prod \lambda_i.$$

We need to show strictness. Similar to how one can prove Cantor's theorem, we use the diagonal method. For contradiction, let $\sum \kappa_i = \prod \lambda_i$. Then there is a bijection between $\sum \kappa_i$ and $\prod \lambda_i$, and it induces a collection $\langle X_i \mid i \in I \rangle$ of

disjoint sets X_i such that $|X_i| = \kappa_i$ and $\bigcup X_i = \times \lambda_i$. Note that the elements of the cartesian product of a collection of sets $\times \lambda_i$ are functions $f : I \rightarrow \bigcup \lambda_i$. Let

$$Y_i = \{f(i) \mid f \in X_i\} \subseteq \lambda_i.$$

We have $|Y_i| \leq |X_i| = \kappa_i < \lambda_i$, so $\lambda_i \setminus Y_i$ is nonempty. Let $g \in \times \lambda_i$ be defined as $g(i) = \min(\lambda_i \setminus Y_i)$ and notice that this g is not in any X_i , a contradiction. \square

Corollary 4.72 (AC). *If $\kappa \geq 2$ and $\lambda \geq \omega$ are cardinals, then*

- (a) $\text{cf}(2^\lambda) > \lambda$,
- (b) $\text{cf}(\kappa^\lambda) > \lambda$,
- (c) $\lambda^{\text{cf}(\lambda)} > \lambda$.

Proof. (a) follows from (b).

(b) Let $\langle \kappa_i \mid i < \lambda \rangle$ be a sequence of cardinals $\kappa_i < \kappa^\lambda$. Then

$$\sup_{i \in \lambda} \kappa_i \leq \sum_{i \in \lambda} \kappa_i < \prod_{i \in \lambda} \kappa_i = (\kappa^\lambda)^\lambda = \kappa^{\lambda \cdot \lambda} = \kappa^\lambda.$$

Where we used Lemma 4.65 and König's inequality for $\lambda_i = \kappa^\lambda$ for all i .

(c) Due to Observation 4.67 there are cardinals $\langle \kappa_i \mid i \in \text{cf}(\lambda) \rangle$ such that $\kappa_i < \lambda$ and $\sup \kappa_i = \lambda$. We use König's inequality for $\lambda_i = \lambda$ for all i :

$$\lambda = \sup_{i \in \text{cf}(\lambda)} \kappa_i \leq \sum_{i \in \text{cf}(\lambda)} \kappa_i < \prod_{i \in \text{cf}(\lambda)} \lambda = \lambda^{\text{cf}(\lambda)}. \quad \square$$

4.4.4 Continuum Hypothesis

The *continuum hypothesis* CH is the statement of ZFC

$$2^{\aleph_0} = \aleph_1$$

and was first proposed by Cantor in 1878. Later, in 1908, Hausdorff proposed the *generalized continuum hypothesis* GCH, which states

$$(\forall \alpha \in \text{On}) : 2^{\aleph_\alpha} = \aleph_{\alpha+1}.$$

We have already encountered the continuum hypothesis at the end of Section 3.4.3, where we mentioned that Gödel (1940) showed that GCH cannot be disproved from ZFC, and Cohen (1963) showed that CH cannot be proved in ZFC. Thus CH and GCH are independent in ZFC. Moreover, Kruse and Rubin (1960) showed that GCH implies AC over ZF.

Let us summarize what our findings up to this point imply about the *continuum function* $\aleph_\alpha \mapsto 2^{\aleph_\alpha}$. We have derived the following inequalities:

Proposition 4.73. *For any ordinals α and β the following holds:*

- (i) $\alpha \leq \beta \implies 2^{\aleph_\alpha} \leq 2^{\aleph_\beta}$, ... from monotonicity
- (ii) $2^{\aleph_\alpha} > \aleph_\alpha$, ... Cantor's theorem

(iii) $\text{cf}(2^{\aleph_\alpha}) > \aleph_\alpha$.

... part (a) of Corollary 4.72

Note that (iii) implies (ii) since $2^{\aleph_\alpha} \geq \text{cf}(2^{\aleph_\alpha})$.

These results impose certain restrictions on the values of the continuum function. For instance, since $\text{cf}(2^{\aleph_0}) > \aleph_0$, we know that 2^{\aleph_0} cannot be \aleph_ω . However, beyond these basic bounds, ZFC tells us remarkably little about what these values actually are, particularly when \aleph_α is a regular cardinal. Easton (1970) showed that the above inequalities are all the axioms of ZFC imply about the values of 2^{\aleph_α} for regular \aleph_α . More formally:

Fact 4.74 (Easton). *If $G : \text{On} \rightarrow \text{On}$ is a function that satisfies*

(i) $\alpha \leq \beta \implies \aleph_{G(\alpha)} \leq \aleph_{G(\beta)}$, ... simulating (i)

(ii) $\text{cf}(\aleph_{G(\alpha)}) > \aleph_\alpha$, ... simulating (iii)

then it is consistent with ZFC to assume that $2^{\aleph_\alpha} = \aleph_{G(\alpha)}$ holds for all regular cardinals \aleph_α .

Remark. This also means that $\mathfrak{c} = 2^{\aleph_0}$ could be arbitrarily large; it could even be an inaccessible cardinal.

Woodin (1981) showed that GCH can fail completely for singular cardinals as well. In particular:

Fact 4.75 (Woodin). *It is consistent with ZFC to assume*

$$(\forall \alpha) : 2^{\aleph_\alpha} = \aleph_{\alpha+2},$$

provided that the existence of a supercompact cardinal is consistent with ZFC.

Finally, let us ponder whether GCH could be true for some initial segment of infinite cardinals. It can be shown from Easton's result that any reasonably defined regular cardinal \aleph_α , for example

$$\aleph_0, \aleph_1, \aleph_{100}, \aleph_{\omega+1}, \aleph_{\omega_1+1},$$

can be the first on which GCH breaks. This means that $2^{\aleph_\alpha} > \aleph_{\alpha+1}$, but $2^{\aleph_\beta} = \aleph_{\beta+1}$ for all $\beta < \alpha$.

The question whether GCH can be violated for the first time by a singular cardinal has remained open for a long time. Silver (1974) showed that this cannot happen for singular cardinals with uncountable cofinality, and Magidor (1977) demonstrated that Silver's result probably cannot be extended to cardinals with countable cofinality. More specifically:

Fact 4.76 (Silver). *If \aleph_α is a singular cardinal with $\text{cf}(\aleph_\alpha) > \omega$, and $2^{\aleph_\beta} = \aleph_{\beta+1}$ holds for all $\beta < \alpha$, then also $2^{\aleph_\alpha} = \aleph_{\alpha+1}$.*

Fact 4.77 (Magidor). *The statement*

$$(\forall n < \omega) : 2^{\aleph_n} = \aleph_{n+1} \quad \wedge \quad 2^{\aleph_\omega} = \aleph_{\omega+2}$$

is consistent with ZFC, provided it is consistent with ZFC to assume the existence of a certain pair of large cardinals.

Therefore, if the existence of these cardinals is consistent, then it is impossible to prove Silvers theorem for \aleph_ω , which is a singular cardinal with cofinality ω , and Silvers theorem thus cannot be generalized for all singular cardinals.

Remark. The function $\aleph_\alpha^{\aleph_\beta}$ has also been studied extensively, and its behavior is much more complex than the behavior of 2^{\aleph_α} . We will not study its values here, but [1] contains a range of results and their simplification if we assume GCH.

5 Large Cardinals

There are many problems in set theory that lead to the question of whether there exist cardinal numbers with certain properties. If the existence of such numbers cannot be proved in ZFC, then we call them *large cardinals*. There is an entire hierarchy of large cardinals, and the weakly inaccessible cardinals we discovered at the end of Section 4.3 are the smallest of them all. Let us recall that a cardinal κ is *weakly inaccessible* if it is an uncountable regular limit cardinal.

The typical reason why it is impossible to prove the existence of cardinals with certain properties is that they would be large enough to imply the consistency of ZFC, contradicting Gödel's second incompleteness theorem. In this section, we will attempt to provide an explanation of this interesting concept.

Definition 5.1. A cardinal κ is *strongly limit* if all $\lambda < \kappa$ satisfy $2^\lambda < \kappa$.

Definition 5.2 (Inaccessible cardinal). A cardinal κ is *strongly inaccessible* (or just *inaccessible*) if it is an uncountable regular strongly limit cardinal.

Clearly, each strongly limit cardinal is limit, and each weakly inaccessible cardinal is strongly inaccessible. Intuitively, weakly inaccessible cardinals cannot be reached from below by taking successor cardinals or by forming suprema of smaller cardinals, and strongly limit cardinals cannot be reached even by using the power set operation.

Observation 5.3. GCH implies that all weakly inaccessible cardinals are strongly inaccessible.

Proof. Let κ be weakly inaccessible and $\lambda < \kappa$ be an infinite cardinal. Then $2^\lambda = \lambda^+$ and $\lambda^+ < \kappa$ because κ is limit. \square

5.1 Cumulative Hierarchy of Sets V_α

The *von Neumann universe* is a proper class built in stages, which are referred to as the *cumulative hierarchy of sets*. We start with the empty set, and at each stage, we collect every possible subset. We define the hierarchy recursively as follows:

$$V_0 := \emptyset, \quad V_{\alpha+1} := \mathcal{P}(V_\alpha), \quad V_\lambda := \bigcup_{\alpha < \lambda} V_\alpha, \quad \mathbf{WF} := \bigcup_{\alpha \in \text{On}} V_\alpha,$$

where λ is a limit ordinal. The class \mathbf{WF} is called the *von Neumann universe*.

Lemma 5.4. For all ordinals α it holds that

- (a) V_α is a transitive set, ... that is, $x \in y \in V_\alpha \implies x \in V_\alpha$
 (b) $V_\beta \subseteq V_\alpha$ for all $\beta < \alpha$.

Proof. By transfinite induction. Note that V_α is transitive if and only if $A \in V_\alpha$ implies $A \subseteq V_\alpha$ for all A . Both claims are clearly true for $\alpha = 0$ and for limit α (assuming the claim holds for all $\beta < \alpha$). Finally, if V_α is transitive and $A \in V_\alpha$, so $A \subseteq V_\alpha$ and $A \in V_{\alpha+1}$. Hence $V_\alpha \subseteq V_{\alpha+1}$. If $A \in V_\alpha$, then $A \subseteq V_\alpha \subseteq V_{\alpha+1}$, so $A \subseteq V_{\alpha+1}$. Thus $V_{\alpha+1}$ is transitive. \square

Corollary 5.5. **WF** is a transitive class since it is a union of transitive sets.

Fact 5.6. **WF** is the \subseteq -largest transitive class on which the relation \in is well-founded. This is why **WF** is sometimes called the well-founded core.

Since **On** is a transitive class, this implies that $\text{On} \subseteq \mathbf{WF}$, and it can be shown that for all ordinals α it holds that

$$\alpha = V_\alpha \cap \text{On}.$$

Let us recall that the *axiom of foundation* is the following statement:

$$(\forall A \neq \emptyset)(\exists x \in A)(x \cap A = \emptyset).$$

Hence, for all $y \in A$, we have $y \not\in x$; in other words, x is a \in -minimal element of the set A . The axiom of foundation is thus the statement that the relation \in is well-founded on the universal class **V** (each set has a minimal element). Since **V** is transitive, we must have $\mathbf{WF} = \mathbf{V}$. It is then easy to see that:

Theorem 5.7. *The following statements are equivalent:*

- (1) *the axiom of foundation,*
- (2) $\mathbf{V} = \mathbf{WF}$,
- (3) *for every set x , there is some ordinal α such that $x \in V_\alpha$.*

This tidy behavior prompted von Neumann to propose the axiom of foundation in 1925. It is not strictly required to prove anything important, 99.9% of mathematics works without it, but it ensures that the universal class V behaves nicely and it makes certain constructions possible. We will now give a quick overview of these constructions; for more details, see [1].

Well-founded induction and recursion Foundation allows us to formulate an induction and recursion principle that is similar to transfinite induction but works for all sets. The proofs are similar to transfinite induction / recursion but require some extra steps, as not all sets are transitive (unlike ordinals).

Theorem 5.8 (Well-founded induction). *If A is a class such that for every set x it holds that*

$$x \subseteq A \implies x \in A,$$

then $A = \mathbf{V}$. Similar to transfinite induction, we can equivalently formulate this via a property $\varphi(x)$ as follows: if for all sets x we have

$$\text{If } \varphi(y) \text{ holds for all } y \in x, \text{ then } \varphi(x).$$

Then $\varphi(x)$ holds for all sets x .

Proof sketch. Suppose that $\mathbf{V} \setminus A \neq \emptyset$. If it is a set, then from the axiom of foundation, it has a \in -minimal element x . Then $x \notin A$, but none of $y \in \mathbf{V} \setminus A$ are in x , so all $y \in x$ are in A . Because $x \subseteq A$, we should have $x \in A$, a contradiction. If $\mathbf{V} \setminus A$ is a proper class, then it is still possible to show that it has a minimal element, but it is a bit more involved. \square

Theorem 5.9 (About construction by well-founded recursion). *If $G : \mathbf{V} \rightarrow \mathbf{V}$ is a class map, then there exists a unique class map $F : \mathbf{V} \rightarrow \mathbf{V}$ satisfying*

$$(\forall x) : F(x) = G(F \upharpoonright x).$$

So we define $F(x)$ using the elements of x and their images.

Rank of a set We can use Foundation to assign each set an ordinal *rank* that essentially measures its complexity. Well-founded recursion implies that there is a unique map $\text{rank} : \mathbf{V} \rightarrow \text{On}$ such that

$$\text{rank}(x) = \sup\{\text{rank}(y) + 1 \mid y \in x\}$$

holds for every set x .

Exercise 5.10. Show that $\text{rank}(x) = \min\{\alpha \mid x \subseteq V_\alpha\}$.

It is easy to see that $V_\alpha = \{x \mid \text{rank}(x) < \alpha\}$ and

$$\text{rank}(x) = \alpha \iff x \in V_{\alpha+1} \setminus V_\alpha,$$

which corresponds with the idea that the cumulative hierarchy divides the universe into distinct complexity levels.

Scott's trick We have actually already encountered the idea of rank when we defined the cardinalities of sets. When working in ZF, $|x|$ is not defined for sets x that cannot be well-ordered. For such sets, we can let

$$|x| := \{y \mid y \approx x \wedge (\forall z)(z \approx x \Rightarrow \text{rank}(y) \leq \text{rank}(z))\},$$

the class of all equinumerous sets with minimal rank. It is easy to see that this definition of cardinality satisfies $x \approx y \iff |x| = |y|$. Notice that $|x|$ is a set since $|x| \subseteq V_{\alpha+1}$, where α is the rank of any $y \in |x|$.

Theorem 5.11 (About representing equivalence classes). *If \sim is an equivalence relation on a class A , then there exists a mapping $\tau : A \rightarrow \mathcal{P}(A)$ such that*

$$x \sim y \iff \tau(x) = \tau(y)$$

holds for all $x, y \in A$.

Proof. For $x \in A$ we define

$$\tau(x) := \{y \in A \mid y \sim x \wedge (\forall z \in A)(z \sim x \Rightarrow \text{rank}(y) \leq \text{rank}(z))\}.$$

This is a set for the same reason as $|x|$ above, and the desired property of $\tau(x)$ follows from the symmetry and transitivity of \sim . \square

Types of ordered sets Ordinals serve as canonical types of well-ordered sets, and the above theorem will allow us to do something similar for all ordered sets. Let \mathcal{O} be the class of all ordered sets (a, \leq_r) such that \leq_r is a partial order on the set a . Isomorphisms of ordered sets define an equivalence relation \simeq on \mathcal{O} . According to the previous theorem, there is a mapping τ that satisfies

$$(a, \leq_r) \simeq (b, \leq_s) \iff \tau((a, \leq_r)) = \tau((b, \leq_s)).$$

We say that $\tau((a, \leq_r))$ is the *type* of the ordered set (a, \leq_r) . If we know the type of an ordered set, then we know all of the properties associated with its order.

It is clear that a similar method can be employed to define types of other mathematical structures. The reason we cannot simply let $\tau(x)$ be the equivalence class to which x belongs is that this is probably a proper class. For example, the class of all graphs isomorphic to your favorite graph is a proper class unless you restrict the definition of graphs so that the vertices can only be elements of some set X . But then you cannot have graphs of arbitrary size, as the maximum number of vertices is limited by $|X|$.

5.2 Beth Numbers \beth_α

As for the sizes of the sets V_α , clearly all V_n for $n \in \omega$ are finite. This, unfortunately, does not directly imply that V_ω is countable, as in ZF a countable union of finite sets might be uncountable. Fortunately, the sets V_n have an inherent structure. In particular, they are transitive. The reader will easily verify that the mapping $f : V_\omega \rightarrow \omega$ defined recursively as $f(\emptyset) := 0$ and for $x > 0$

$$f(x) := \sum_{y \in x} 2^{f(y)}$$

is a bijection between V_ω and ω .

Furthermore, notice that $|V_{\alpha+1}| = 2^{|V_\alpha|}$. The sizes of the sets V_α for infinite ordinals α are known as *beth numbers*, conventionally written as $\beth_0, \beth_1, \beth_2, \dots$, where \beth (“beth”) is the second letter of the Hebrew alphabet.

Definition 5.12. The *beth numbers* \beth_α are defined recursively as

- (i) $\beth_0 := \aleph_0$,
- (ii) $\beth_{\alpha+1} := 2^{\beth_\alpha}$,
- (iii) $\beth_\lambda := \sup\{\beth_\alpha \mid \alpha < \lambda\}$ for limit ordinals λ .

Observation 5.13. *It is easy to verify that:*

- (a) $|V_{\omega+\alpha}| = \beth_\alpha$ for all ordinals α ,
- (b) $\beth_\alpha \geq \aleph_\alpha$ for all ordinals α ,
- (c) \beth_ω is the first strongly limit cardinal,
- (d) CH is equivalent to $\beth_1 = \aleph_1$,
- (e) GCH is equivalent to $(\forall \alpha \in \text{On}) : \beth_\alpha = \aleph_\alpha$.

Proposition 5.14. *The class of all α such that $\beth_\alpha = \aleph_\alpha$ is a proper class.*

Proof. The beth function is continuous (from definition) and increasing (Cantor's theorem), so it is a normal function. By Theorem 3.16 (iii), the class of ordinals α for which $\beth_\alpha = \alpha$ is a proper class, and by Proposition 3.18, whenever $\beth_\alpha = \alpha$, then also $\aleph_\alpha = \alpha$, so $\beth_\alpha = \aleph_\alpha$. \square

Observation 5.15. $\kappa \in \text{Cn}$ is strongly limit $\iff \kappa = \beth_\lambda$ for a limit ordinal λ .

Proof. Clearly \beth_λ is a strongly limit cardinal whenever λ is a limit ordinal. Let κ be a strongly limit cardinal. Since the beth numbers are unbounded, there exists a least beth number $\beth_\alpha \geq \kappa$. If $\alpha = \beta + 1$ were isolated, then $\beth_\beta < \kappa$, but $2^{\beth_\beta} = \beth_\alpha \geq \kappa$, so α is limit. κ is an upper bound of the set $\{\beth_\delta \mid \delta < \alpha\}$, but \beth_α is its supremum, so $\beth_\alpha \leq \kappa$. \square

5.3 Constructible Universe L_α

At the end of Section 3.4.3, we mentioned that Gödel introduced the *constructible universe* L . Its construction is similar to that of the cumulative hierarchy, but we restrict set formation to only definable subsets. Informally, denote by $\text{Def}(X)$ the set of subsets of X definable by first-order formulas over (X, \in) with parameters from X . The hierarchy of constructible sets is defined recursively as follows:

$$L_0 := \emptyset, \quad L_{\alpha+1} := \text{Def}(L_\alpha), \quad L_\lambda := \bigcup_{\alpha < \lambda} L_\alpha, \quad L := \bigcup_{\alpha \in \text{On}} L_\alpha,$$

where λ is a limit ordinal. The class L is called the *constructible universe*.

To understand $\text{Def}(X)$, recall that the full power set $\mathcal{P}(X)$ contains *every* conceivable collection of elements from X . In contrast, $\text{Def}(X)$ contains only those subsets that we can explicitly describe using the sets we already have. Formally, a subset $Y \subseteq X$ belongs to $\text{Def}(X)$ if there is a first-order formula ϕ (possibly with parameters from X) such that $Y = \{u \in X \mid (X, \in) \models \phi(u)\}$. When we say the formula is evaluated “over (X, \in) ,” we mean that all quantifiers \forall and \exists in ϕ are restricted to range strictly over the elements of X , rather than the entire universe V (so $\forall a$ is interpreted as $\forall a \in X$).

Comparing V_α and L_α It is clear from the definition that $L_\alpha \subseteq V_\alpha$ for all α . Furthermore, for every finite n , we have

$$L_n = V_n, \quad |L_n| = |V_n| < \omega,$$

from this also

$$L_\omega = V_\omega, \quad |L_\omega| = |V_\omega| = \omega.$$

Then

$$|V_{\omega+1}| = |\mathcal{P}(V_\omega)| = 2^\omega,$$

but it can be shown that

$$|L_{\omega+1}| = \omega.$$

Fact 5.16. *For all ordinals α it holds that*

- (a) L_α is a transitive set,
- (b) $L_\beta \subseteq L_\alpha$ for all $\beta < \alpha$,
- (c) $\alpha = L_\alpha \cap \text{On}$,
- (d) $|L_\alpha| = |\alpha|$ for all $\alpha \geq \omega$.

Corollary 5.17. L is a transitive class and $\text{On} \subseteq L$.

Gödel showed that L is what we call a *transitive inner model* of ZF: Given any model (M, \in^M) of ZF, the class L^M within this model is transitive, and (L^M, \in^L) , where \in^L is the restriction of \in^M to L^M , is also a model of ZF. Gödel showed that this inner model additionally satisfies AC and GCH. Hence, if ZF is consistent (has a model M), then ZFC and ZFC + GCH are also consistent (they have a model, namely L^M). This does not contradict Gödel's second incompleteness theorem because ZF cannot prove that it has a model.

Gödel essentially showed that assuming $\mathbf{V} = L$ in ZF implies AC and GCH. The statement $\mathbf{V} = L$ is quite important, and it has a name:

Axiom 5.18 (Axiom of Constructibility). $\mathbf{V} = L$.

Recall that Cohen constructed a model of ZFC in which GCH fails, so $\mathbf{V} = L$ cannot be true in this model. Hence, the axiom of constructibility cannot be proved from ZFC. Similarly:

Proposition 5.19. *The axiom of constructibility cannot be disproved from ZFC.*

Proof sketch. We show that if ZFC is consistent (has a model M), then the axiom of constructibility holds in the inner model L^M .

The hierarchy L_α is defined using ordinals, and since $\text{On} \subseteq L$, the inner model L^M has the same ordinals as M . Hence, for every L_α in M , there is L'_α in L^M . One only needs to show that $L_\alpha = L'_\alpha$ for all α . Clearly, $L_0 = L'_0$ since M and L^M share the same empty set. $L_{\alpha+1} = L'_{\alpha+1}$ holds because the transitivity of L^M implies that $\text{Def}(L_\alpha)$ in M means the same as $\text{Def}(L'_\alpha)$ in L^M . For limit α , the sets L_α and L'_α are simply the unions of all previous stages, so $L_\alpha = L'_\alpha$. Therefore, the constructible universe L' in L^M is the entire universe $V' = L^M$. \square

5.4 Significance of the Axiom of Infinity

Gödel's proofs of the consistency of ZF + AC and ZFC + GCH are what we call *relative consistency proofs*. If we accept AC and GCH as axioms, nothing will break. As we will see, the same cannot be said for the axioms of infinity.

We have already seen a shadow of inaccessible cardinals — the cardinal \aleph_0 . It is not uncountable, but it is a regular and strongly limit cardinal. This is the key to understanding large cardinals. Let us denote by Inf the axiom of infinity, by $\text{ZFC}^- := \text{ZFC} - \text{Inf}$ the theory we obtain from ZFC by removing the axiom of infinity, and by $\text{ZFC}_{\text{fin}} := \text{ZFC}^- + \neg \text{Inf}$ the theory of finite sets; that is, ZFC with the axiom of infinity replaced by its negation (there are no limit ordinals). One can prove that all sets in ZFC_{fin} are finite. For a theory T , we denote by $\text{Con}(T)$ the statement that T is consistent. See Section 3.4.3 for a more formal definition.

We should also note that ZF_{fin} can prove the axiom of choice (for finite sets), but we include it to makes things simpler.

Lemma 5.20. V_ω is a model of both ZFC^- and ZFC_{fin} .

Proof. We need to show that the axioms of ZFC_{fin} hold in V_ω . V_ω is the set of all hereditary finite sets — sets whose elements are finite sets, whose elements are finite sets, and so on, all the way down to the empty set. Hence $\neg\text{Inf}$ holds, and the axioms of ZFC^- hold because “finite sets behave nicely”:

V_ω is transitive (extensionality) and the relation \in is well-founded on V_ω (foundation). The union of finite sets is finite (pairing, union); we include all subsets of every set (schema of specification); the power set of a finite hereditary set is finite hereditary (power set); and the image of a finite set under a function is finite (replacement). Since finite sets can always be well-ordered and such a well-ordering is a hereditary finite set, the axiom of choice is also true. \square

Remark. It is important to understand that when we are proving that a sentence φ holds in some $M \subseteq \mathbf{V}$, we interpret the quantifiers $\forall x$ and $\exists x$ as $\forall x \in M$ and $\exists x \in M$. So, for example, the extensionality of sets in \mathbf{V} does not necessarily imply the extensionality of sets in M .

Theorem 5.21. $\text{ZFC} \vdash \text{Con}(\text{ZFC}^-)$, and thus, by Gödel’s second incompleteness theorem, we have the two following results:

- (a) ZFC^- cannot prove the axiom of infinity,
- (b) the consistency of ZFC^- does not imply the consistency of ZFC .

Proof. Because ZFC can define V_ω and prove that (V_ω, \in) is a model of ZFC^- , it follows that ZFC proves the consistency of ZFC^- . If ZFC^- proved Inf , it would thus be able to prove its own consistency. And if $\text{Con}(\text{ZFC}^-)$ implied $\text{Con}(\text{ZFC})$, ZFC would be able to prove its own consistency, as it proves $\text{Con}(\text{ZFC}^-)$. \square

Therefore, in accepting the axiom of infinity, we have to take a leap of faith. This is in direct contrast to accepting the axiom of choice — we cannot break the consistency of ZF by adding AC , but we might be breaking the consistency of ZF^- by adding Inf ; no one can know for sure. In a certain sense, the axiom of infinity significantly increases the power of the theory, allowing it to prove the consistency of its weaker version.

A little bit of history Even though ZF is a formalist, axiomatic theory, its axioms should be obviously true statements. But what proof do we have (in the real world) of the existence of infinite sets? Philosophers differentiate between two views or concepts of infinity. No matter how many natural numbers one writes down, there will always be only finitely many of them. However, it will always be possible to add a new one; hence, natural numbers are *potentially infinite*. Similarly, one can always extend a line segment, so line segments also represent a potential infinity. If someone somehow created *all* the natural numbers (it would not be possible to write down any more; all of them would be there), or created an *entire* line (impossible to extend any further), then one would create what is known as an *actual infinity*.

Philosophers have pondered for centuries the question of whether an actual infinity can exist, but the consensus has generally been negative. The best proof of actual infinity we have to this day comes from the Bohemian philosopher and

theologian Bernard Bolzano, and his 1851 book “The Paradoxes of the Infinite.” The idea of the proof is as follows:

Consider the set of all “truths.” Let S_1 be the proposition: “There are truths.” This is true: if there were no truths, reasoning would be impossible. Let S_2 be the proposition: “The proposition S_1 is true.” Let S_3 be the proposition: “The proposition S_2 is true.” And so on. By induction, all of the propositions S_n are true, and they represent distinct truths since they talk about different objects. Therefore, there are infinitely many truths. God, in His perfection, must see all truths; hence, He possesses the actual infinity of all truths.

Mathematicians have historically worked almost exclusively with potential infinity. The birth of set theory came in 1874 when Cantor published an article in which he proved that the real numbers are uncountable, but the algebraic numbers (roots of polynomials with integer coefficients) are countable. He discussed the infinities represented by this very carefully, not yet addressing them as actual infinity. We know from a letter he sent to Dedekind that Cantor later found Bolzano’s book and used Bolzano’s arguments to defend the actual infinity represented by ω and \aleph_0 , which he explicitly introduced in his later papers.

Not everyone in the mathematical world agreed with Cantor. Namely, the German mathematician Leopold Kronecker strongly disagreed with Cantor’s work and prevented Cantor from publishing in the oldest and most prestigious mathematics journal at the time. However, Kronecker was far from the only one standing against Cantor. Many famous French mathematicians, including Borel, Lebesgue, and Poincaré, were skeptical of some aspects of Cantor’s set theory for a long time, before finally accepting it. Cantor’s work was finally widely accepted at the beginning of the 20th century, when it was backed by Hilbert who famously said “No one shall expel us from the paradise that Cantor has created.”

My main source for this history section was a lecture given by prof. Petr Vopěnka, available here: https://www.youtube.com/watch?v=_b_rPG3bu0Y.

5.5 Large Cardinal Axioms

Similar to how assuming the existence of a limit ordinal allows us to prove the consistency of ZFC^- , assuming the existence of an inaccessible cardinal allows us to prove the consistency of ZFC . To illustrate how large cardinal proofs are done, we will prove (following [24]) that the existence of strongly inaccessible cardinals cannot be proved in ZFC .

Lemma 5.22. *The following are equivalent for any uncountable cardinal κ :*

- (1) κ is strongly inaccessible,
- (2) for every $x \in V_\kappa$, if $f : x \rightarrow \kappa$, then $\sup(\text{Rng } f) < \kappa$,
- (3) for every $\alpha < \kappa$, if $f : 2^\alpha \rightarrow \kappa$, then $\sup(\text{Rng } f) < \kappa$.

Proof. (2) \Rightarrow (3): Trivial.

(3) \Rightarrow (1): If κ is not regular, then there is some $\alpha < \kappa$ and $f : \alpha \rightarrow \kappa$ with $\sup(\text{Rng } f) = \kappa$, this f can be trivially extended to 2^α . If κ is not strongly limit, let $\alpha < \kappa$ be the first ordinal such that $\kappa \leq 2^\alpha$. Since $\kappa \preceq 2^\alpha$, there is an injection from κ into 2^α , which can be reversed; therefore, there is $f : 2^\alpha \rightarrow \kappa$ which is surjective, so $\sup(\text{Rng } f) = \kappa$.

(1) \Rightarrow (2): It is enough to check that (2) holds for $x = V_\alpha$ for all $\alpha < \kappa$, since if $x \in V_\kappa$, there is some $\alpha < \kappa$ such that $x \subseteq V_\alpha$ (κ is a limit ordinal). Thus it is enough to show that if $\alpha < \kappa$, then $|V_\alpha| < \kappa$, since κ is a regular cardinal. We prove this by transfinite induction on α . It clearly holds for $\alpha = 0$. Suppose it holds for α , then $|V_\alpha| = \lambda < \kappa$, so $|V_{\alpha+1}| = |\mathcal{P}(V_\alpha)| = 2^\lambda < \kappa$ since κ is strongly limit. Suppose α is limit and for all $\beta < \alpha$ we have $|V_\beta| = \lambda_\beta < \kappa$. Since $\alpha < \kappa$ and κ is regular, we have

$$|V_\alpha| = \left| \bigcup_{\beta < \alpha} V_\beta \right| = \sup_{\beta < \alpha} \lambda_\beta < \kappa.$$

The second equality holds thanks to Corollary 4.66. □

Lemma 5.23. *If κ is strongly inaccessible, then V_κ is a model of ZFC.*

Proof. Since V_κ is transitive, Extensionality and Foundation hold. Because κ is uncountable, $\omega \in V_\kappa$ since $\kappa \subseteq V_\kappa$; thus Infinite holds. Since κ is a limit ordinal, Power set, Union, and Pairing hold. This also implies that Choice holds: if $x \in V_\kappa$, then $x \times x \in V_\kappa$ since $x \times x \in \mathcal{P}(\mathcal{P}(\mathcal{P}(x))) \in V_\kappa$, and V_κ is transitive. Since Choice holds in the outer theory, x has a well-ordering $R \subseteq x \times x$. Repeating what we did earlier yields $R \in V_\kappa$ since $R \in \mathcal{P}(x \times x) \in V_\kappa$.

Note that Specification follows from Replacement, so it is enough to show that. Let $A \in V_\kappa$ and let φ be such that $V_\kappa \models \varphi(x,y)$ corresponds to a function with domain $\supseteq A$. We want to show that $B := \{y \mid (\exists x \in A) : \varphi(x,y)\} \in V_\kappa$. Suppose $\varphi(x,y)$ holds and define $f(x)$ as the first ordinal α such that $y \in V_{\alpha+1}$. Then f is a function from A to κ , so by (2) in the previous lemma, f must be bounded by some $\lambda < \kappa$, so $\forall y \in B$ we have $y \in V_{\lambda+1}$, so $B \in V_{\lambda+2} \subseteq V_\kappa$. □

Denote by Inf^* the statement “there exists a strongly inaccessible cardinal.” Then we can state a similar theorem about Inf^* and ZFC, as we did in the previous section about Inf and ZFC^- in Theorem 5.21.

Corollary 5.24. *ZFC cannot prove the existence of strongly inaccessible cardinals, and the consistency of ZFC does not imply the consistency of $\text{ZFC} + \text{Inf}^*$.*

Proof. Direct consequence of Gödel’s second incompleteness theorem and the previous lemma’s claim that $\text{ZFC} + \text{Inf}^* \vdash \text{Con}(\text{ZFC})$. □

Remark. Using a similar argument as in Lemma 5.23, one can show that if κ is weakly inaccessible, then L_κ is a model of ZFC. Therefore ZFC cannot prove the existence of weakly inaccessible cardinals either.

One could now imagine some property P such that if a cardinal κ satisfies P , then V_κ is a model of $\text{ZFC} + \text{Inf}^*$, and thus $\text{ZFC} + \text{Inf}^*$ cannot prove the existence of such a cardinal. This process can then be iterated with increasingly powerful cardinal properties. In this light, we can view large cardinals as generalizations of the axiom of infinity — by assuming the existence of a large cardinal, we significantly increase the strength of our theory. However, we might be adding some inconsistencies.

Many problems in set theory reduce to “if it is consistent with ZFC to assume the existence of certain large cardinals, then it is consistent to assume that something holds,” or “if a certain large cardinal exists, then something holds.” The

following hierarchy displays some large cardinals you might encounter, ordered by consistency strength:

$$\text{Inaccessible} < \text{Mahlo} < \dots < \text{Measurable} < \text{Woodin} < \text{Supercompact} < \dots$$

The “gap” represented by the dots above corresponds to a massive jump in logical complexity. The cardinals on the left (*inaccessible* and *Mahlo*) are essentially just “very tall” milestones. They are consistent with the constructible universe L discussed earlier. Intuitively, L is the “minimalist” version of set theory, containing no “random” or “chaotic” sets. If we assume the axiom of constructibility $\mathbf{V} = L$ (that every set is constructible), inaccessible and Mahlo cardinals can still exist. They fit inside the tidy, rigid structure of L .

However, once we cross the gap to *measurable* cardinals, this is no longer true. These cardinals are so powerful that their existence contradicts $\mathbf{V} = L$. They require the universe to contain objects (specifically, complex ultrafilters or embeddings) that are too “rich” to be constructible and thus are not included in L . Their existence implies that the universe must be much “wider” than L , containing extra information that L cannot see.

In the missing region (the dots), we find cardinals that bridge this gap. *Weakly compact* cardinals sit at the very top of the hierarchy compatible with L . Above them, we find *Erdős* and *Ramsey* cardinals, which are defined by partition properties (a generalization of Ramsey’s theorem). These are one of the first cardinals that truly break out of the constructible universe, paving the way for the measurable cardinals above them (every measurable cardinal is Ramsey). We will discuss these “partition” cardinals in more detail in Section 7.3.5.

The large cardinal hierarchy is highly complex, and I suggest reading [24] as an introduction. For a more thorough exposition, see [22]. For an overview of large cardinals and their properties, see THE UPPER ATTIC of Cantor’s attic.

6 Topics in Pure Set Theory

In this section, we study several central tools of higher infinite combinatorics: filters and ultrafilters, stationary sets, and infinite trees. These structures capture different but closely related ways of organizing large sets and are a core part of the study of large cardinals.

6.1 Filters and Ultrafilters

Throughout this section, X is always a nonempty set, and given a set $A \subseteq X$, we denote by \overline{A} the set $X \setminus A$. For a nonempty system of subsets $\mathcal{A} \subseteq \mathcal{P}(X)$, we denote by \mathcal{A}^* the set $\{\overline{A} \mid A \in \mathcal{A}\}$.

Definition 6.1 (Filter, ultrafilter). Let $X \neq \emptyset$. A nonempty system of subsets $\mathcal{F} \subseteq \mathcal{P}(X)$ is a *filter* on X if

- (i) $\emptyset \notin \mathcal{F}$,
- (ii) if $A \in \mathcal{F}$ and $A \subseteq B$, then $B \in \mathcal{F}$, ... \mathcal{F} is *upward closed*
- (iii) if $A, B \in \mathcal{F}$, then $A \cap B \in \mathcal{F}$.

A filter \mathcal{F} that also satisfies

$$(iv) \forall A \subseteq X \text{ either } A \in \mathcal{F} \text{ or } \overline{A} \in \mathcal{F}$$

is called an *ultrafilter* on X . Given an ultrafilter \mathcal{U} on X , we say that a subset $A \subseteq X$ is *\mathcal{U} -large* if $A \in \mathcal{U}$.

Intuition. A filter on X allows us to express which subsets of X we consider to be “large.” (i) The empty set is not large. (ii) Adding something to a large set will not make it small. (iii) An intersection of two large sets is large. (iv) Whenever we split X into two parts, one of them has to be large.

Observation 6.2. *For every filter \mathcal{F} on X it holds that*

- (a) *condition (i) is equivalent to $\mathcal{F} \neq \mathcal{P}(X)$,* *... from (ii)*
- (b) $X \in \mathcal{F}$, *... from $\mathcal{F} \neq \emptyset$ and (ii)*
- (c) $\forall A, B \in \mathcal{F} : A \cap B \neq \emptyset$, *... from (iii) and (i)*
- (d) \mathcal{F} is closed under finite intersections, *... from (iii)*
- (e) $\nexists A \subseteq X$ such that $A, \overline{A} \in \mathcal{F}$, *... from (iii) and (i)*

The last observation implies that the “or” in (iv) is exclusive.

Ideals are the dual notion to filters, allowing us to define which subsets of X are “small.” This means that we obtain the definition of an ideal by replacing the symbols \emptyset , \cap , and \subseteq in the conditions (i)–(iii) for X , \cup , and \supseteq .

Definition 6.3 (Ideal, prime ideal). Let $X \neq \emptyset$. A nonempty system of subsets $\mathcal{I} \subseteq \mathcal{P}(X)$ is an *ideal* on X if

- (i) $X \notin \mathcal{I}$,
- (ii) if $A \in \mathcal{I}$ and $B \subseteq A$, then $B \in \mathcal{I}$,
- (iii) if $A, B \in \mathcal{I}$, then $A \cup B \in \mathcal{I}$.

An ideal \mathcal{I} that also satisfies

$$(iv) \forall A \subseteq X \text{ either } A \in \mathcal{I} \text{ or } \overline{A} \in \mathcal{I}$$

is called a *prime ideal* on X .

For example, the system of all subsets of \mathbb{R} with Lebesgue measure zero is an ideal. Analogous observations, as we did for filters, apply to ideals as well.

Exercise 6.4. Show that \mathcal{F} is a filter $\iff \mathcal{F}^*$ is an ideal

Definition 6.5 (Dual filter, dual ideal). If \mathcal{F} is a filter on X , then \mathcal{F}^* is the *dual ideal* of \mathcal{F} . If \mathcal{I} is an ideal on X , then \mathcal{I}^* is the *dual filter* of \mathcal{I} .

Clearly $(\mathcal{F}^*)^* = \mathcal{F}$ and $(\mathcal{I}^*)^* = \mathcal{I}$. It is also easy to see that a filter \mathcal{F} is an ultrafilter $\iff \mathcal{F}^*$ is a prime ideal. This duality between filters and ideals allows us to study only one of these notions. Statements about one can be converted to statements about the other via De Morgan's laws:

$$\bigcup_{i \in I} A_i = \bigcap_{i \in I} \overline{A_i}, \quad \bigcap_{i \in I} A_i = \bigcup_{i \in I} \overline{A_i}.$$

Lemma 6.6. *If \mathcal{U} is an ultrafilter on X and $A = A_1 \cup \dots \cup A_k \in \mathcal{U}$, then $A_j \in \mathcal{U}$ for at least one index j . Furthermore, if the sets A_i are disjoint, then the index j is unique.*

Proof. Suppose that $A_j \notin \mathcal{U}$ for all j . Then from (iv) we have $\overline{A_j} \in \mathcal{U}$ for all j and (iii) implies that $\bigcap \overline{A_j} = \overline{A} \in \mathcal{U}$. Hence, both A and \overline{A} are in \mathcal{U} , a contradiction. Notice that if the sets A_i are disjoint, then (iii) and (i) further imply that the index j must be unique. \square

Definition 6.7 (Maximal filter). A filter \mathcal{F} on X is *maximal* if there is no filter \mathcal{F}' on X such that $\mathcal{F} \subset \mathcal{F}'$.

Theorem 6.8. *For a filter \mathcal{F} on X , the following conditions are equivalent:*

- (i) \mathcal{F} is an ultrafilter,
- (ii) \mathcal{F} is a maximal filter,
- (iii) whenever $A \cup B \in \mathcal{F}$, then $A \in \mathcal{F}$ or $B \in \mathcal{F}$.

Proof. We already know (i) \implies (iii) from Lemma 6.6.

(iii) \implies (ii) Suppose $\mathcal{F} \subset \mathcal{F}'$ for some filter \mathcal{F}' . Then there is some $A \in \mathcal{F}'$ such that $A \notin \mathcal{F}$, hence $\overline{A} \in \mathcal{F}$ since $A \cup \overline{A} = X \in \mathcal{F}$. Because $\mathcal{F} \subseteq \mathcal{F}'$, also $\overline{A} \in \mathcal{F}'$, so both A and \overline{A} are in \mathcal{F}' , a contradiction.

(ii) \implies (i) Suppose that \mathcal{F} is not an ultrafilter. Then there exists a set $A \subseteq X$ such that $A, \overline{A} \notin \mathcal{F}$. Note that for all $B \in \mathcal{F}$ we have $A \cap B \neq \emptyset$; otherwise $B \subseteq \overline{A}$, so we would have $\overline{A} \in \mathcal{F}$. Define a filter

$$\mathcal{F}' := \{B' \subseteq X \mid (\exists B \in \mathcal{F}) : (A \cap B) \subseteq B'\}.$$

It is indeed a filter: $\emptyset \notin \mathcal{F}'$ since $A \cap B \neq \emptyset$ for all B ; it is upward closed from definition; and it is intersection closed because if $B'_1, B'_2 \in \mathcal{F}'$, then $A \cap B_1 \subseteq B'_1$ and $A \cap B_2 \subseteq B'_2$ for some $B_1, B_2 \in \mathcal{F}$. Since $B_1 \cap B_2 \in \mathcal{F}$, we have $A \cap B_1 \cap B_2 \subseteq B'_1 \cap B'_2 \in \mathcal{F}'$.

Notice that \mathcal{F}' is a filter such that $\mathcal{F} \subseteq \mathcal{F}'$ and $A \in \mathcal{F}'$ but $A \notin \mathcal{F}$, so \mathcal{F} is not maximal. \square

Example. For every set X and every nonempty $Y \subseteq X$, the set

$$\mathcal{F}_Y := \{A \subseteq X \mid Y \subseteq A\}$$

is a filter on X . Filters of this form are called *principal filters*.

If X is an infinite set, then the set

$$\mathcal{I}_{fin} := \{A \subseteq X \mid A \text{ is finite}\}$$

is an ideal on X . Its dual filter

$$\mathcal{F}_{co} := \{A \subseteq X \mid \bar{A} \text{ is finite}\}$$

is the *cofinite filter* on X .

Lemma 6.9. *A principle filter \mathcal{F}_Y on X is an ultrafilter $\iff |Y| = 1$.*

Proof. “ \implies ” Suppose that $x, y \in Y$ and $x \neq y$. Then $Y = (Y \setminus \{x\}) \cup \{x\} \in \mathcal{F}_Y$, but neither $Y \setminus \{x\}$ nor $\{x\}$ are in \mathcal{F}_Y , contradicting the previous lemma.

“ \impliedby ” We know that \mathcal{F}_Y is a filter. To show that it is an ultrafilter, let $Y = \{x\}$ be a singleton and suppose that $A \subseteq X$. Either $x \in A$, in which case $A \in \mathcal{F}$, or $x \notin A$, in which case $x \in \bar{A}$, so $\bar{A} \in \mathcal{F}$. \square

Definition 6.10 (Nonprincipal ultrafilter). Ultrafilters of the form

$$\mathcal{U}_x = \{A \subseteq X \mid x \in A\}$$

for some $x \in X$ are called *principal*. Ultrafilters that are not principal are called *nonprincipal* or *free*, and they are the interesting ones.

Lemma 6.11. *For an ultrafilter \mathcal{U} on X , the following conditions are equivalent:*

- (i) \mathcal{U} is a principal ultrafilter,
- (ii) \mathcal{U} contains a finite set,
- (iii) $\bigcap \mathcal{U} = \bigcap_{A \in \mathcal{U}} A \neq \emptyset$.

In particular, it follows from (ii) that all ultrafilters on a finite set X are principal.

Proof. (i) \implies (iii) If $\mathcal{U} = \{A \subseteq X \mid x \in A\}$ for some $x \in X$, then $x \in \bigcap \mathcal{U}$.

(iii) \implies (ii) Let $A_0 = \bigcap \mathcal{U}$ and note that $\bar{A}_0 \notin \mathcal{U}$ (since $A_0 \subseteq A$ for all $A \in \mathcal{U}$), so $A_0 \in \mathcal{U}$ because \mathcal{U} is an ultrafilter. We claim that $|A_0| = 1$. Suppose $|A_0| > 1$, so A_0 has a proper nonempty subset $B \subseteq A_0$. But then neither B nor \bar{B} contains A_0 as a subset, so $B, \bar{B} \notin \mathcal{U}$, a contradiction with \mathcal{U} being an ultrafilter.

(ii) \implies (i) Let $A_0 \in \mathcal{U}$ be a set of the smallest (finite) size.

If $A_0 = \{x\}$ is a singleton, then for all $B \in \mathcal{U}$ we have $x \in B$, since otherwise $B \cap A_0 = \emptyset$. Hence $\mathcal{U} \subseteq \mathcal{U}_x = \{A \subseteq X \mid x \in A\}$. To show equality, suppose that $x \in A$ but $A \notin \mathcal{U}$. Then $\bar{A} \in \mathcal{U}$ since \mathcal{U} is an ultrafilter, but $x \notin \bar{A}$.

If $|A_0| > 1$, then A_0 contains a proper nonempty subset $B \subset A_0$. Because \mathcal{U} is an ultrafilter and $B \notin \mathcal{U}$ (since A_0 has the smallest size in \mathcal{U}), we have $\bar{B} \in \mathcal{U}$. Hence $\bar{B} \cap A_0 \in \mathcal{U}$, but $|\bar{B} \cap A_0| < |A_0|$, which contradicts the minimality of A_0 as well. \square

Corollary 6.12. *An ultrafilter \mathcal{U} on an infinite set X is nonprincipal \iff it contains the cofinite filter \mathcal{F}_{co} .*

Proof. If an ultrafilter is nonprincipal, then it contains no finite sets, so it must contain all their complements. On the other hand, if it $\bar{A} \in \mathcal{U}$ for every finite $A \subseteq X$, then also $A \notin \mathcal{U}$ for every finite set $A \subseteq X$, so \mathcal{U} contains no finite sets, and thus is a principal ultrafilter. \square

Theorem 6.13 (AC, The Ultrafilter Lemma, Tarski, 1930). *Every filter \mathcal{F} on a set X can be extended to an ultrafilter $\mathcal{U} \supseteq \mathcal{F}$ on X .*

Proof. We will use Zorn's lemma. Let $\mathbf{F} \subseteq \mathcal{P}(\mathcal{P}(X))$ be the set of all filters on X . Notice that if \mathcal{C} is a chain of filters in the partially ordered set (\mathbf{F}, \subseteq) , then $\bigcup \mathcal{C} = \bigcup_{\mathcal{F} \in \mathcal{C}} \mathcal{F}$ is also a filter and an upper bound of \mathcal{C} . Hence, we can apply Zorn's lemma to obtain a \subseteq -maximal filter \mathcal{F}' above \mathcal{F} . Theorem 6.8 implies that \mathcal{F}' is an ultrafilter. \square

Remark. We have used Zorn's lemma to prove this theorem, but the full power of the axiom of choice is not required, as the ultrafilter lemma is equivalent to the *Boolean prime ideal theorem* BPI, which is sometimes used as a weak form of the axiom of choice. BPI is independent in $\mathbf{ZF} + \mathbf{DC}$, and DC is independent in $\mathbf{ZF} + \mathbf{AC}_\omega + \mathbf{BPI}$. Most notably, if it is consistent to assume the existence of a strongly inaccessible cardinal, then there is a model (constructed by Solovay, see Fact 7.21) of $\mathbf{ZF} + \mathbf{DC}$ in which ω has no nonprincipal ultrafilter.

Corollary 6.14 (AC). *Every infinite set has a nonprincipal ultrafilter.*

Proof. Use the previous theorem to extend the cofinite filter. \square

This result can be further strengthened.

Definition 6.15 (Uniform filter). A filter \mathcal{F} on an infinite set X is *uniform* if all $A \in \mathcal{F}$ have the same cardinality as X .

Observation 6.16. *Every uniform ultrafilter is nonprincipal, as it contains no finite sets.*

Proposition 6.17 (AC). *Every infinite set has a uniform ultrafilter.*

We do this by extending a more powerful filter than the cofinite filter. For an infinite set X with $|X| = \kappa$, we call the ideal

$$\mathcal{I} := [X]^{<\kappa} = \{A \subseteq X \mid |A| < \kappa\}$$

the *Fréchet ideal* on X . The dual filter to \mathcal{I} is called the *Fréchet filter* on X . The Fréchet filter on a countably infinite set X is just the cofinite filter.

Proof of Proposition 6.17. Let X be an infinite set and let \mathcal{U} be an ultrafilter that extends the Fréchet filter \mathcal{F} on X . Suppose that $A \in \mathcal{U}$ has size $|A| < \kappa$. Then $\bar{A} \in \mathcal{F} \subseteq \mathcal{U}$, a contradiction with \mathcal{U} being an ultrafilter. \square

Definition 6.18 (κ -complete filter). A filter (ideal) \mathcal{F} on a set X is κ -*complete* if any intersection (union) of less than κ sets from \mathcal{F} is still in \mathcal{F} .

Every filter and every ideal is ω -complete. If a filter is ω_1 -complete, then it is closed under countable intersections. These filters have been extensively studied and are often referred to as σ -complete instead. It follows from De Morgan's laws that a filter is κ -complete \iff its dual ideal is κ -complete. Notice that every principal filter \mathcal{F}_Y is κ -complete for all cardinals κ because its intersection $\bigcap \mathcal{F}_Y = Y$ is nonempty.

The Fréchet ideal (and thus also the filter) on a cardinal λ is $\text{cf}(\lambda)$ -complete, but not $\text{cf}(\lambda)^+$ -complete. Since the Lebesgue measure is countably additive, the ideal of sets with Lebesgue measure zero is σ -complete.

Filters on algebras of sets So far, we have considered a filter on X to be a subset of $\mathcal{P}(X)$. However, we can define a filter on any *algebra* of subsets of X .

Definition 6.19 (Algebra of sets). An *algebra of subsets* of X is a set system $M \subseteq \mathcal{P}(X)$ such that (i) $X \in M$ (ii) if $A, B \in M$, then $A \cap B \in M$ (iii) if $A \in M$, then $\overline{A} \in M$. Note that (ii) and (iii) imply that if $A, B \in M$, then $A \cup B \in M$.

Definition 6.20 (M -filter). Given an algebra M of subsets of X , a set system $\mathcal{F} \subseteq M$ is an M -*filter*, if (i) $\emptyset \notin \mathcal{F}$, and for every $A, B \in M$ we have: (ii) if $A \in \mathcal{F}$ and $A \subseteq B$, then $B \in \mathcal{F}$, and (iii) if $A, B \in \mathcal{F}$, then $A \cap B \in \mathcal{F}$. We say that \mathcal{F} is an M -*ultrafilter* if for every $A \in M$ either $A \in \mathcal{F}$ or $\overline{A} \in \mathcal{F}$.

TODO this definition will be used for an equivalent characterization of weakly compact cardinals

6.2 Stationary Sets

Definition 6.21 (Club set). Let α be a limit ordinal. A set $C \subseteq \alpha$ is

- (a) *unbounded* (in α) if C is cofinal in α , that is $\sup(C) = \alpha$,
- (b) *closed* (in α) if whenever $\beta < \alpha$ and $\sup(C \cap \beta) = \beta$, then $\beta \in C$,
- (c) a *club set* (in α) if it is closed and unbounded in α .

Notice that nothing changes if we require β in the definition of a closed set to be a limit ordinal. This is because $\sup(\delta + 1) = \delta$, so the condition only becomes “active” for limit ordinals. It is also easy to see that this definition of a closed set is equivalent to Definition 3.5.

Definition 6.22. The set of *accumulation points* of a subset $A \subseteq \alpha$ is the set

$$\text{acc}(A) := \{\beta \mid \sup(A \cap \beta) = \beta\}.$$

TODO stationary set, stationary limit, mahlo cardinals,...

6.3 Infinite Trees

In graph theory, (possibly infinite) trees are acyclic graphs in which each pair of vertices is connected by a (finite) path. We will study *order-theoretic trees*, a more general concept, as we will soon see that graph trees correspond to order-theoretic trees of height at most ω .

Definition 6.23. An *order-theoretic tree* is a partially ordered set $(T, <_T)$ where for every $x \in T$ the set $(\leftarrow, x) = \{y \in T \mid y <_T x\}$ is well-ordered by $<_T$. A *subtree* of a tree $(T, <_T)$ is a subset $S \subseteq T$ with the induced order. We call the elements $x \in T$ of a tree its *nodes*.

Observation 6.24. *Equivalently, $(T, <_T)$ is a tree if $<_T$ it is well-founded and for every $x \in T$ the set (\leftarrow, x) is a chain.*

This, in particular, means that every chain in a tree is well-ordered since it is linear and well-founded.

Definition 6.25. A *branch* of a tree is a \subseteq -maximal chain.

Observation 6.26 (AC). *Every chain can be extended into a branch.*

Proof. Chains are partially ordered by \subseteq and the union of a chain is a chain. Hence, we can use Zorn's lemma to find a \subseteq -maximal chain C' for any given chain C such that $C \subseteq C'$. \square

Definition 6.27 (Tree terminology). Let T be a tree and $x \in T$.

- (1) the *subtree of x* is the set $[x, \rightarrow) := \{y \in T \mid x \leq y\}$.
- (2) $y \in T$ is a *predecessor* of x if $y < x$, *successor* of x if $x < y$, and *immediate successor* of x if it is a minimal successor of x (with respect to $<_T$).
- (3) x is a *leaf* if it has no immediate successors.
- (4) the *branching factor* of x is the cardinality of the set of its immediate successors. If every node has finite branching factor, then we say that the tree is *finitely branching*.
- (5) x is a *root* if it is a minimal element of T (with respect to $<_T$)
- (6) the *height* or *level* of x is the ordinal type of the well-ordered set (\leftarrow, x) , denoted as $|x|_T$ or $H_T(x)$.
- (7) the tree is divided into *levels*, starting at level zero. The α th level of T is

$$T(\alpha) := \{x \in T \mid |x|_T = \alpha\}.$$

- (8) the *height* of T is $H(T) := \sup\{|x|_T + 1 \mid x \in T\}$. Equivalently, it is the least α for which $T_\alpha = \emptyset$.
- (9) the *length* of a branch is its ordinal type with respect to $<_T$. It is easy to see that every branch has length at most $H(T)$.
- (10) a *cofinal branch* is a branch of length $H(T)$.

Example. Let us now look at some examples of trees.

- Every rooted graph-tree corresponds to an order-theoretic tree of height at most ω with a single root. However, trees of height more than ω have some "limit points" along their branches, meaning that the tree is not connected in the usual graph sense.
- Every ordinal $(\alpha, <)$ is a tree of height α with a single branch.
- Given a set A and an ordinal α , the set of all sequences of elements from A of length less than α , ordered by inclusion, is a tree. More precisely,

$$T = \bigcup_{\beta < \alpha} {}^\beta A,$$

so the nodes are functions, and the immediate successors of a node $f : \beta \rightarrow A$ are the nodes $f_a : \beta + 1 \rightarrow A$ where $f_a \upharpoonright \beta = f$ and $f_a(\beta) = a$ for each

$a \in A$. Think of the label of a node as a set of instructions on how to get here, specifying which edge to choose at every step of the way from the root. Draw a picture. Since every node has branching factor $|A|$ and the tree has height α , we call this tree the *complete A -ary tree* of height α .

- The complete 2-ary (binary) tree of height ω is called *Cantor's tree*. Since every level is finite (level n has size 2^n), it is a countable tree. But since every branch corresponds to an infinite sequence of 1s and 0s, the number of branches is $|\mathcal{P}(\omega)| = |\mathbb{R}| = \mathfrak{c}$, and we can use its branches to model the real numbers.

We will now try to answer the question of which trees are guaranteed to have a cofinal branch. Is there a tree of height ω with no cofinal branch? Yes — consider a tree with ω branches of lengths $1, 2, 3, \dots, n, \dots$. Then the tree has height ω , but none of the branches are infinite.

Exercise 6.28. Find for any infinite cardinal κ a tree of height κ with no cofinal branch and with levels of cardinality at most $\text{cf}(\kappa)$

Theorem 6.29 (König's lemma, AC_ω). *Every tree of height ω with finite levels has a cofinal branch.*

Remark. The requirement “every level is finite” is equivalent to “the tree has finite branching factor and finitely many roots.”

Proof. We say that a node x is good if its subtree $[x, \rightarrow)$ is infinite. Because the tree is infinite and has finitely many roots, there exists a good root r . Because r has finite branching factor, one of its immediate successors x_0 has to be good. Then we choose x_1 as a good immediate successor of x_0 and so on. We construct an infinite sequence of nodes $(x_n)_{n < \omega}$ such that x_{n+1} is an immediate successor of x_n , and $[x_n, \rightarrow)$ is infinite for all n . This is a branch of length ω .

Where is the catch? At each stage, x_n might have multiple good immediate successors, and we need to *choose* one of them. There are multiple ways to do this. Since we are making sequential choices that depend on each other, one might want to use the axiom of dependent choice (see Axiom 1.7). However, we do not need its full strength. Because every level is finite, it has a well-ordering. Since the tree has countably many levels, we can use AC_ω to *choose* a well-ordering for every level and then select x_{n+1} as the minimal good immediate successor of x_n with respect to this ordering.

Another way would be to realize that since every level is countable (namely finite) and the tree has height ω , Lemma 3.39 implies that if we assume AC_ω , then the tree is countable. Thus there exists a bijection between its nodes and ω . We can use this bijection to well-order the nodes and again pick x_{n+1} as the minimal good successor. One might ask: if all levels are finite, should not the countable union of these levels also be countable? Surprisingly, the answer is not necessarily yes. As discussed in Section 1.6, in ZF there might exist infinite sets that contain no countable subset. The idea to list the nodes as x_1, \dots, x_n for the first level, y_1, \dots, y_m for the second level, and so on, does not work: you would need to *choose* a linear ordering of each level, requiring AC_ω , or more specifically $\text{AC}_\omega^{\text{fin}}$, a restriction of AC_ω to selections from finite sets. \square

Exercise 6.30. Derive König’s lemma from Zorn’s lemma.

Fact 6.31. *If κ is an infinite cardinal with $\text{cf}(\kappa) = \omega$, then every tree of height κ with finite levels has a cofinal branch.*

Proof. Balcar–Stepanek 3.39 □

Fact 6.32 (Aronszajn). *There is a tree of height ω_1 with countable levels that has no cofinal branch.*

Proof. Balcar–Stepanek 3.41 □

Fact 6.33. *If $\lambda < \text{cf}(\kappa)$, then every tree of height κ with levels of cardinality strictly smaller than λ has a cofinal branch.*

Remark. The result by Aronszajn demonstrates that if the levels are not *strictly* smaller, then the cofinal branch might not exist.

Proof. Balcar–Stepanek 3.40. The proof uses the notion of stationary sets. □

7 Infinitary Combinatorics

Infinitary combinatorics explores the extension of ideas in combinatorics to infinite sets, where intuition often clashes with the rigid logic of set theory. This field uses powerful tools like the axiom of choice and transfinite induction to reveal that infinite structures can behave paradoxically, as seen in the Banach–Tarski decomposition, or remarkably tamely, as captured by the Compactness Principle.

7.1 Compactness Principles

There are many statements in mathematics that go under the name of “compactness principle.” You might know, for instance, the compactness theorem from first order logic: “if every finite part of a theory has a model, then the entire theory has a model,” or the de Bruijn–Erdős theorem: “if every finite subgraph of a graph G is r -colorable, then the entire graph G is r -colorable.”

Their common property is the ability to transfer statements about finite objects to statements about infinite objects, and the fact that the axiom of choice is required to prove them. The following combinatorial compactness principle can be viewed as a generalization of the de Bruijn–Erdős theorem, and was discovered by Richard Rado.

7.1.1 Rado’s Selection Principle

Definition 7.1. Let $\langle A_i \mid i \in I \rangle$ be a collection of finite nonempty sets.

- (a) A *partial selector* of the collection is a mapping f such that $\text{Dom } f \subseteq I$ and $f(i) \in A_i$ for all $i \in \text{Dom } f$.
- (b) A system of partial selectors S *covers finite subsets* of I if for every finite $u \subseteq I$ there exists $f \in S$ such that $u \subseteq \text{Dom}(f)$.

- (c) A *total selector* is any mapping $g \in \prod_{i \in I} A_i$. It is a *filtered extension* of a system of partial selectors S if for every finite $u \subseteq I$ there exists $f \in S$ defined on u such that $g \upharpoonright u = f \upharpoonright u$.

Intuition. Think of I as a set of vertices and think of A_x for $x \in I$ as a set of colors available for the vertex x . A partial selector is a partial coloring of I , a total selector is a coloring of the entire I , and a filtered extension of a system of partial selectors is a coloring of the entire I that agrees on every finite part of I with one of the partial selectors.

Theorem 7.2 (Rado's selection principle, AC). *If $\langle A_i \mid i \in I \rangle$ is a collection of finite nonempty sets, then every system of partial selectors covering finite subsets of I has a filtered extension.*

Proof (cf. [4]). Let S be a system of partial selectors covering finite subsets of I . We say that a partial selector h is *compatible* with S if for all finite $u \subseteq I$ there exists a partial selector $f \in S$ defined on the entire u such that

$$h \upharpoonright u = f \upharpoonright (u \cap \text{Dom}(h)).$$

This means that f and h agree on the elements of u that h is defined on. We say that h is *compatible with f on u* .

Denote by P the set of all partial selectors compatible with S . Notice that $P \neq \emptyset$ as $\emptyset \in P$. Indeed, the empty mapping is compatible with S if S covers finite subsets of I , which it does. Furthermore, if a total selector is compatible with S , then it is a filtered extension. We will use Zorn's lemma on P ordered by \subseteq to obtain a maximal selector, and then show that it is total.

To use Zorn's lemma, we need to show that every chain $C = \{h_j \mid j \in J\}$ has an upper bound. We claim that its union $H := \bigcup_{j \in J} h_j$ is an upper bound. Clearly, it is a partial selector; we need to show that it is compatible with S . Let $u \subseteq I$ be finite; we claim that $u \cap \text{Dom}(H)$ is contained in the domain of some $h \in C$, and it is easy to see that if h is compatible with f on u , then H is also compatible with f on u .

To find such h , order the elements of $u \cap \text{Dom}(H)$ as x_0, \dots, x_n (since we are assuming AC we can well-order the entire I). For every $i \leq n$, there exists $h_i \in C$ defined on x_i . We start with $h^0 := h_0$ defined on x_0 . Suppose that h^k is defined on all x_0, \dots, x_k . Since C is a chain, one of h^k, h_{k+1} has to extend the other, and so there exists h^{k+1} defined on all x_0, \dots, x_{k+1} . By induction up to n , there exists $h^n \in C$ defined on all x_0, \dots, x_n .

By Zorn's lemma, there exists a maximal element g of P . It remains to show that g is a total selector. For contradiction suppose that there exists $x \in I \setminus \text{Dom}(g)$. Since g is maximal, g cannot be extended by any pair (x, y) where $y \in A_x$. Let $g_y = g \cup \{(x, y)\}$ (we try to color the vertex x with color y). Because g is maximal, none of g_y are compatible with S . Thus for all $y \in A_x$, there exists a finite witness $u \subseteq I$ such that g_y is not compatible on u with any selector from S defined on u (negation of what it means to be compatible with S). Using AC, we *choose* one such witness u_y for every $y \in A_x$.

Since A_x is finite, $U := \bigcup_{y \in A_x} u_y$ is also finite, and $U' := U \cup \{x\}$ as well. Because g is compatible with S , and U' is finite, there exists $f \in S$ (defined on the entire U') compatible with g on U' . If we let $y := f(x)$, then g_y is compatible with f on u_y , which is a contradiction. \square

We will now look at some applications of Rado's selection principle.

7.1.2 Infinite Hall's Theorem

The marriage problem asks the following question: Every man knows certain women, and more men might know the same woman. Under what conditions is it possible for each man to marry a woman he knows if polyandry¹⁶ is forbidden?

Hall's theorem provides a surprisingly simple answer: it is possible \iff for every group of men, the total number of distinct women they know is at least the number of men in the group.

Definition 7.3 (Transversal). A *system of distinct representatives* of a collection of sets $\langle A_i \mid i \in I \rangle$, also called a *transversal*, is an injective function $f : I \rightarrow \bigcup A_i$ such that $f(i) \in A_i$ for every $i \in I$.

The following is clearly a necessary condition for a transversal to exist:

$$(\forall J \subseteq I) : \left| \bigcup_{j \in J} A_j \right| \geq |J|.$$

It is called *Hall's condition*, and Hall proved that, in the finite case, it is sufficient:

Theorem 7.4 (Hall's theorem). A *finite collection of finite sets* $\langle A_i \mid i \in I \rangle$ has a transversal \iff *Hall's condition is satisfied*.

Exercise 7.5. Use Rado's selection principle to derive the following infinite version of Hall's theorem:

Theorem 7.6 (Infinite Hall's theorem, AC). A *possibly infinite collection of finite sets* $\langle A_i \mid i \in I \rangle$ has a transversal \iff *Hall's condition is satisfied*.

Exercise 7.7. Will infinite Hall's theorem still hold even if the sets A_i are allowed to be countable? (It will not). What if we in addition assume that every countable subcollection $\langle A_i \mid i \in J \rangle$ for countable $J \subseteq I$ has a transversal?

7.1.3 De Bruijn–Erdős Theorem

An r -coloring of a graph $G = (V, E)$ is a function $\chi : V \rightarrow r$. It is a *proper coloring* if, for every edge $uv \in E$, we have $\chi(u) \neq \chi(v)$. If G has a proper r -coloring, then we say that it is *r -colorable*.

Theorem 7.8 (De Bruijn–Erdős, AC). *If every finite subgraph of a graph G is r -colorable, then the entire graph G is r -colorable.*

Proof. We use Rado's selection principle on the collection $\langle A_v \mid v \in V_G \rangle$ where $A_v = r$ for all v (every vertex has r possible colors). For every finite subset $U \subseteq V$, there exists a proper r -coloring of the induced subgraph $G[U]$. These colorings form a system of partial selectors S covering finite subsets of V . By Rado's selection principle, there exists a filtered extension $\chi : V \rightarrow r$. It is easy to see that this is a proper r -coloring of the entire graph G . If not, then there is an edge $uv \in E$ such that $\chi(u) = \chi(v)$. But since $\{u, v\}$ is a finite subset of V , there exists $f \in S$ such that χ and f agree on $\{u, v\}$. Because f is proper, we must have $\chi(u) = f(u) \neq f(v) = \chi(v)$, a contradiction. \square

¹⁶Polyandry is a form of polygamy in which a woman takes two or more husbands at the same time.

Corollary 7.9 (AC). *Every (even infinite) planar graphs is 4-colorable.*

Proof. The four color theorem allows us to 4-color every finite subgraph. \square

We will now formulate a generalized version of the de Bruijn–Erdős theorem. It is also one of the theorems on the list of “compactness principles,” and has special importance in Ramsey theory.

A *hypergraph* is a pair $H = (V, E)$ where V is a set of *vertices* and $E \subseteq [V]^{<\omega}$ is a set of *hyperedges* (finite sets of vertices). For $W \subseteq V$, we denote by H_W the *induced hypergraph* $(W, E \cap \mathcal{P}(W))$. An r -*coloring* of H is a function $\chi : V \rightarrow r$. We say that a hyperedge $e \in E$ is *monochromatic* with respect to χ if $\chi \upharpoonright e$ is constant (all $v \in e$ have the same color). An r -coloring of H is *proper* if no hyperedge $e \in E$ is monochromatic. H is r -*colorable* if it has a proper r -coloring. The *chromatic number* $\chi(H)$ of H is the least r such that H is r -colorable.

Theorem 7.10 (Compactness principle, AC). *Let $H = (V, E)$ be a hypergraph and $r \in \omega$. If for all finite $W \subseteq V$ we have $\chi(H_W) \leq r$, then $\chi(H) \leq r$.*

Proof. The same as the proof of the de Bruijn–Erdős theorem. \square

Exercise 7.11. Derive the compactness principle for countable V from König’s lemma. This implies that the countable case does not require the full power of the axiom of choice: it suffices to assume AC_ω .

7.1.4 Compactness in Ramsey Theory

Ramsey theory is a branch of combinatorics centered on the idea that “complete disorder is impossible.” It studies the conditions under which order must inevitably emerge from sufficiently large systems, no matter how chaotic they might appear.

The following version of the compactness principle is particularly useful in Ramsey theory, as it allows us to easily convert many infinite Ramsey-style theorems into their finite versions.

Proposition 7.12 (AC_ω). *Let $H = (V, E)$ where V is countable and $E \subseteq [V]^{<\omega}$. If $\chi(H) > r$, then there exists some finite $W \subseteq V$ such that $\chi(H_W) > r$.*

Remark. This is also clearly true for arbitrary V , but in that case, we need to assume full AC. Exercise 7.11 shows that if V is countable, AC_ω is enough.

We will demonstrate this on the classical finite and infinite versions of Ramsey’s theorem. We will prove them in Section 7.3; now we will only show that the infinite version implies the finite one.

Theorem 7.13 (Finite Ramsey’s theorem FRT). *For every size n , number of colors r , and dimension k , there exists N such that for every r -coloring χ of $[N]^k$, there exists a subset $A \subseteq N$ of size n such that $\chi \upharpoonright [A]^k$ is constant (every $a \in [A]^k$ has the same color).*

Theorem 7.14 (Infinite Ramsey’s theorem IRT). *For every number of colors r , dimension k , and for every r -coloring χ of $[\omega]^k$, there exists a subset $A \subseteq \omega$ of size ω such that $\chi \upharpoonright [A]^k$ is constant (every $a \in [A]^k$ has the same color).*

Proposition 7.15 (AC_ω). $\text{IRT} \implies \text{FRT}$.

Proof. In Ramsey theory compactness proofs, one always wants to construct a hypergraph $H = (V, E)$, the vertices of which are the objects we want to color, and the edges of which are the structures we want to find. Let us fix the number of colors r and the dimension k . We are coloring k -sets of the natural numbers, and we want to find an n -set such that all of its k -sets have the same color. Define

$$V := [\omega]^k, \quad E := \{[A]^k \mid A \in [\omega]^n\}.$$

Since $||[\omega]^k| = \omega$, assuming AC_ω is sufficient for the compactness principle to hold. The plan is to show the following:

$$\text{IRT} \implies \chi(H) > r \implies (\exists \text{ finite } W \subseteq V) \chi(H_W) > r \implies \text{FRT}.$$

Infinite Ramsey's theorem implies that for every r -coloring of $[\omega]^k$, there exists an *infinite* $A \subseteq \omega$ such that $[A]^k$ is monochromatic. Let A' be the set containing the first n elements of A . Then $[A']^k$ is a monochromatic edge of H . Notice that this implies $\chi(H) > r$, since every r -coloring of H contains a monochromatic edge.

By the compactness principle, there exists a finite $W \subseteq V$ such that $\chi(H_W) > r$: every r -coloring of H_W contains a monochromatic edge. Let $N \in \omega$ be the first number larger than all the natural numbers explicitly mentioned in W . This means that $W \subseteq [N]^k$. Therefore, for every r -coloring of $[N]^k \subseteq V$, there exists $A \subseteq N$ of size n such that $[A]^k$ is monochromatic. We have proved the finite version of Ramsey's theorem. \square

Exercise 7.16. An *arithmetic progression* of length k is an increasing sequence of natural numbers a_1, \dots, a_k such that the difference $a_{i+1} - a_i$ is the same for all i . *Van der Waerden's theorem* states that for any finite coloring of ω , there exists a color class containing arithmetic progressions of arbitrary finite length. Show that the infinite version of Van der Waerden's theorem implies its finite version: for any number of colors r and length k , there exists N such that any r -coloring of N contains a monochromatic arithmetic progression of length k .

Remark. Van der Waerden's theorem is a classical result in *partition regularity*. It was vastly strengthened by Endre Szemerédi in 1975. *Szemerédi's theorem* replaces the partition condition with a density condition, stating that any subset of the natural numbers with positive upper density contains arithmetic progressions of arbitrary length.

Intuitively, a set has positive upper density if it occupies at least some fixed percentage of the integers as one goes to infinity. Formally, a set $A \subseteq \omega$ has positive upper density if

$$\limsup_{n \rightarrow \infty} \frac{|A \cap n|}{n} > 0.$$

Since in any finite partition at least one color class must have positive upper density, Szemerédi's theorem implies Van der Waerden's theorem.

A further landmark result in this direction is the *Green–Tao theorem* (2004), which establishes that the set of prime numbers contains arithmetic progressions of arbitrary length. This result is particularly deep because the primes have upper density 0 (by the *prime number theorem*, the density decays as $1/\ln n$), and thus the result does not follow from Szemerédi's theorem.

Exercise 7.17. For finite $X \subseteq \omega$ denote by $\Sigma(X)$ the set of all $y \in \omega$ that are the sum of some $Y \subseteq X$. *Folkman's theorem* states that for any finite coloring of ω , there exist arbitrarily large finite sets $X \subseteq \omega$ such that $\Sigma(X)$ is monochromatic. Show that the infinite version of Folkman's theorem implies its finite version: for any number of colors r and size n , there exists N such that any r -coloring of N contains a set X of size n such that $\Sigma(X)$ is monochromatic.

Remark. Hindman extended Folkman's result in 1974 when he showed that for any finite coloring of ω , there exists an infinite subset $A \subseteq \omega$ such that all finite partial sums of elements from A have the same color.

Exercise 7.18. Let $S(x_1, \dots, x_n)$ be a linear homogeneous constraint

$$c_1x + c_2x_2 + \cdots + c_nx_n = 0, \quad (\forall i) c_i \in \mathbb{Z}.$$

We say that S is *regular* if for every finite coloring of ω there exist numbers $a_1, \dots, a_n \in \omega$ with the same color such that $S(a_1, \dots, a_n)$ holds. Rado proved in 1933 that S is regular \iff some nonempty subset of the coefficients c_i sums to zero. Formulate a finite version of this theorem, and deduce it using a compactness argument.

7.2 Chromatic Numbers of Infinite Graphs

Chromatic Number of the Plane

This is a famous problem presented by Hadwiger and Nelson in 1950: what is the minimum number of colors required to color the plane \mathbb{R}^2 such that no two points at distance 1 from each other have the same color? The answer is unknown and might depend on whether we accept the axiom of choice.

This problem is clearly equivalent to finding the chromatic number of the graph $G = (V, E)$, the vertices of which are the points of the plane $V = \mathbb{R}^2$, and two points a, b form an edge if the distance from a to b is 1.

Exercise 7.19. Show that $4 \leq \chi(G) \leq 7$.

Hint. To show $\chi(G) \geq 4$, find a small unit-distance graph that requires 4 colors (7 vertices are enough). To show the upper bound, find a 7-coloring of the plane. Coloring a suitable tiling of the plane might be helpful.

The lower bound was recently (2018) raised to 5 when a computer found a 1581-vertex unit-distance graph that requires five colors. But that is all we know as of writing this text: $\chi(G)$ could be 5, 6, or 7. We can increase the lower bound by finding a larger counterexample and decrease the upper bound by finding a more efficient coloring of the entire plane.

If we accept the axiom of choice, then the de Bruijn–Erdős theorem implies that if every finite unit-distance graph can be k -colored, then $\chi(G) \leq k$. This gives us another tool for lowering the upper bound, so $\chi(G)$ *might* depend on the axioms we accept.

Conditional Chromatic Number Graph

In 2003, Shelah and Soifer [30] presented a graph whose chromatic number *does* depend on the axioms of set theory, provided that a strongly inaccessible cardinal exists. The difference is in fact quite striking: $\chi(G)$ will either be 2, or it will be uncountable.

They considered a graph $G = (V, E)$, the vertices of which are the real numbers $V = \mathbb{R}$, and two numbers x, y form an edge if $|x - y| = \sqrt{2} + q$, where $q \in \mathbb{Q}$ is a rational number.

Proposition 7.20. *In ZFC it holds that $\chi(G) = 2$.*

Proof. Define for each integer $n \in \mathbb{Z}$ the set $\mathbb{Q}_n := \mathbb{Q} + n\sqrt{2}$, a copy of \mathbb{Q} shifted by $n\sqrt{2}$. Notice that the distance between $x \in \mathbb{Q}_n$ and $y \in \mathbb{Q}_{n+1}$ is $q + \sqrt{2}$ for some $q \in \mathbb{Q}$. Therefore, the graph induced by \mathbb{Q}_n and \mathbb{Q}_{n+1} is a complete bipartite graph. Furthermore, notice that every neighbor of $x \in \mathbb{Q}_n$ lies either in \mathbb{Q}_{n-1} or in \mathbb{Q}_{n+1} . This means that coloring the elements of \mathbb{Q}_n red for even n , and blue for odd n , is a 2-coloring of the subgraph induced by the set

$$S := \{q + n\sqrt{2} \mid q \in \mathbb{Q}, n \in \mathbb{Z}\}.$$

Now we will need to use a basic result from group theory; for details, see Lemma A.3 in the Appendix. It is easy to verify that S is a subgroup of \mathbb{R} . Hence, the relation

$$x \sim y \iff x - y \in S$$

is an equivalence on \mathbb{R} . Its equivalence classes are the cosets of S in \mathbb{R} , and they partition \mathbb{R} into disjoint shifted copies of S . From our previous observations, it is clear that each of these copies is a component of G , and we can color each copy with 2 colors. Where is the catch? Without the axiom of choice, it is not clear what the “even” and “odd” levels of a given copy are. We need to use the axiom of choice to *choose* a base level to start coloring from for every copy of S .

Formally, *choose* a representative from each equivalence class of \sim . For any $x \in \mathbb{R}$, denote by $f(x)$ the representative of the equivalence class to which x belongs. Now define a 2-coloring χ of \mathbb{R} as follows:

$$\chi(x) = \begin{cases} 0, & \text{if } x - f(x) = 2n\sqrt{2} + q \text{ for some } n \in \mathbb{Z} \text{ and } q \in \mathbb{Q}, \\ 1, & \text{if } x - f(x) = (2n + 1)\sqrt{2} + q \text{ for some } n \in \mathbb{Z} \text{ and } q \in \mathbb{Q}. \end{cases}$$

It is clear from the previous discussion that χ is a proper 2-coloring of G . □

In our alternative axiomatic system, we will omit AC and will instead assume the following axiom:

Axiom (LM). Every subset of \mathbb{R} is Lebesgue measurable.

Because this text does not assume prior knowledge of the Lebesgue measure, Appendix A.2 includes a short introduction to measure theory. We should note that Lebesgue measure requires AC_ω to satisfy some of its basic properties. Furthermore, AC can prove the existence of subsets of \mathbb{R} (for example, the Vitali sets) that are not Lebesgue measurable, so $\text{AC} \implies \neg\text{LM}$.

Fact 7.21 (Solovay [35], 1970). *LM is consistent with ZF + DC provided that it is consistent to assume the existence of a strongly inaccessible cardinal. Here, DC denotes the axiom of dependent choice (see Axiom 1.7).*

Remark. In 1984, Shelah [30] showed in a paper titled “Can you take Solovay’s inaccessible away?” that you cannot take Solovay’s inaccessible away.

Thus, if a strongly inaccessible cardinal exists, then ZF + DC + LM has a model. Since DC is stronger than AC_ω , this is a model of ZF + AC_ω + LM as well. The result by Solovay also implies that the weaker forms of the axiom of choice (AC_ω and DC) are not strong enough to construct a non-measurable set.

Proposition 7.22. *In ZF + AC_ω + LM, it holds that $\chi(G) \geq \aleph_1$.*

Proof. Suppose that $\chi : \mathbb{R} \rightarrow \omega$ is a coloring of \mathbb{R} with countably many colors. This coloring induces a partition $\langle A_i \mid i < \omega \rangle$ of \mathbb{R} into countably many disjoint subsets. Since the Lebesgue measure λ is countably additive, we have

$$\sum \lambda(A_i) = \lambda(\mathbb{R}) = \infty. \quad (7.1)$$

Therefore, there exists i such that $\lambda(A_i) > 0$. We will show that if $\lambda(A) > 0$, then there exist two points $x, y \in A$ that form an edge of G . This means that G has a monochromatic edge in A_i , so χ is not a proper coloring, and we need more than $|\omega| = \aleph_0$ colors.

Let $\lambda(A) > 0$. Then A cannot be too “sparse” everywhere. The density theorem implies that there exists an interval I such that

$$\frac{\lambda(I \cap A)}{\lambda(I)} \geq \frac{9}{10}.$$

Choose $q \in \mathbb{Q}$ such that $q + \sqrt{2} \in (0, \frac{\lambda(I)}{10})$. Now shift A by this small amount to get $B := A + (q + \sqrt{2})$, and notice that

$$\frac{\lambda(I \cap B)}{\lambda(I)} \geq \frac{8}{10}.$$

It is a fundamental property of the Lebesgue measure that if X and Y have finite measure, then $\lambda(X \cup Y) = \lambda(X) + \lambda(Y) - \lambda(X \cap Y)$. From this and the monotonicity of λ , we have

$$\lambda(I) \geq \lambda((I \cap A) \cup (I \cap B)) = \lambda(I \cap A) + \lambda(I \cap B) - \lambda(I \cap A \cap B), \quad (7.2)$$

and therefore

$$\frac{\lambda(I \cap A \cap B)}{\lambda(I)} \geq \frac{\lambda(I \cap A)}{\lambda(I)} + \frac{\lambda(I \cap B)}{\lambda(I)} - \frac{\lambda(I)}{\lambda(I)} \geq \frac{9}{10} + \frac{8}{10} - 1 = \frac{7}{10}.$$

Hence $I \cap A \cap B$ is nonempty. Let $x \in A \cap B$, and notice that $y := x - (q + \sqrt{2})$ lies in A . The points x, y both belong to A , and their distance is $x - y = q + \sqrt{2}$ for some $q \in \mathbb{Q}$, so they form an edge of G . \square

Remark. Exercise A.36 shows that the interval I exists even if A is not measurable, and (7.1) follows from countable subadditivity of the outer Lebesgue measure. The only place where we actually needed LM was (7.2). Even though it might seem impossible, this step could fail for non-measurable sets. Even if the outer Lebesgue measure was $\lambda^*(I \cap A) = \lambda^*(I \cap B) = \lambda(I)$, it might happen that $I \cap A \cap B = \emptyset$. Non-measurable sets can be “large” but “hollow.”

7.3 Partition Relations of Cardinals

Exercise 7.23. If $\mathbb{Q} = A_1 \cup A_2 \cup \dots \cup A_n$ is a partition of \mathbb{Q} into finitely many parts, then at least one A_i contains an order-isomorphic copy of \mathbb{Q} . That is, A_i has a countable dense subset with no maximum and no minimum.

Definition 7.24 (Partition arrow). For cardinals κ, λ, μ and a natural number $k \in \omega$, the expression

$$\kappa \longrightarrow (\lambda)_{\mu}^k$$

means that for every mapping $\chi : [\kappa]^k \rightarrow \mu$, there exists $A \subseteq \kappa$ of cardinality λ such that $\chi \upharpoonright [A]^k$ is constant. We say that χ is a *coloring* of k -sets of κ by μ colors, and that A is *homogeneous* for χ . We also say that the set $[A]^k$ is *monochromatic* with respect to the coloring χ .

We read the expression $\kappa \rightarrow (\lambda)_{\mu}^k$ as “ κ arrows¹⁷ k -sets of λ with μ colors.”

Note that the partition arrow is interesting only when $\mu > 1$.

Example. Let us look at some examples of the partition arrow relation:

- $(n-1)r+1 \rightarrow (n)_r^1$ is the pigeonhole principle with r holes, where we want n pigeons in a single hole.
- if κ is regular and $\mu < \kappa$, then $\kappa \rightarrow (\kappa)_{\mu}^1$ is the pigeonhole principle for regular cardinals (see Corollary 4.39).
- $6 \rightarrow (3)_2^2$ means that in any group of six people, there will always be a subgroup of three mutual friends or three mutual strangers.
- $(\forall n, k, r)(\exists N) N \rightarrow (n)_r^k$ is the finite version of Ramsey’s theorem for r colors and dimension k .
- $(\forall k, r) \omega \rightarrow (\omega)_r^k$ is the infinite version of Ramsey’s theorem for r colors and dimension k .

Definition 7.25 (Colorful partition arrow). For cardinals κ, λ, μ and a natural number $k \in \omega$, the expression

$$\kappa \longrightarrow (\lambda, \mu)^k$$

means that for every mapping $\chi : [\kappa]^k \rightarrow \{\text{red}, \text{blue}\}$, there either exists $A \subseteq \kappa$ of cardinality λ such that all sets in $[A]^k$ are **red**, or there exists $B \subseteq \kappa$ of cardinality μ such that all sets in $[B]^k$ are **blue**.

Example. In the language of graph theory, $(\forall k)(\forall l)(\exists N) : N \rightarrow (k, l)^2$ means that for every k and l , there exists N such that every graph on N vertices either contains a clique of size k or an independent set of size l .

Observation 7.26. *Some simple properties of the partition arrow are*

- (a) $\kappa \rightarrow (\lambda, \lambda)^k$ is equivalent to $\kappa \rightarrow (\lambda)_2^k$,
- (b) if $\kappa \rightarrow (\lambda, \mu)^k$, then $\kappa \rightarrow (\mu, \lambda)^k$,
- (c) if $\kappa_1 \rightarrow (\lambda_1)_{\mu_1}^{k_1}$ and $\kappa_1 \leq \kappa_2$, $\lambda_1 \geq \lambda_2$, $\mu_1 \geq \mu_2$, $k_1 \geq k_2$, then $\kappa_2 \rightarrow (\lambda_2)_{\mu_2}^{k_2}$.

¹⁷Česky říkáme „šipkuje.“

7.3.1 Limits of the Partition Arrow

When defining the partition arrow, we only considered finite dimensions k . The reason is the following proposition:

Proposition 7.27 (AC). *For every cardinal κ we have $\kappa \not\rightarrow (\omega)_2^\omega$.*

Proof 1. We will construct a bad coloring of $[\kappa]^\omega$. Let \sqsubseteq be a well-ordering of $[\kappa]^\omega$ given by AC. Define a coloring $\chi : [\kappa]^\omega \rightarrow \{\text{red}, \text{blue}\}$ as

$$\chi(A) := \begin{cases} \text{red}, & \text{if } (\exists B \in [A]^\omega) B \sqsubseteq A, \\ \text{blue}, & \text{otherwise.} \end{cases}$$

So the set A is red if it has a \sqsubseteq -smaller (infinite) subset, and it is blue if it is \sqsubseteq -minimal. Now suppose that some $X \in [\kappa]^\omega$ is homogeneous for χ :

(a) if $\chi(X) = \text{red}$, it means that X has a \sqsubseteq -smaller subset. Since \sqsubseteq is a well-ordering, there exists a \sqsubseteq -minimal subset $A \subseteq X$. But $\chi(A) = \text{blue}$.

(b) if $\chi(X) = \text{blue}$ let A be the \sqsubseteq -minimal subset $Y \subseteq X$ such that $X \setminus Y$ is infinite. Choose $a \in X \setminus A$, then $\chi(A \cup \{a\}) = \text{red}$. \square

Proof 2. Define an equivalence relation \sim on $[\kappa]^\omega$ as

$$X \sim Y \iff X \text{ and } Y \text{ differ only infinitely many elements.}$$

Using AC, choose a representative from each equivalence class, and denote the representative of the class to which X belongs as $f(X)$. Now define a coloring $\chi : [\kappa]^\omega \rightarrow \{\text{red}, \text{blue}\}$ as

$$\chi(A) := \begin{cases} \text{red}, & \text{if } |X \setminus f(X)| \text{ is odd,} \\ \text{blue}, & \text{if } |X \setminus f(X)| \text{ is even.} \end{cases}$$

Suppose that some $X \in [\kappa]^\omega$ is monochromatic. Pick some $x \in X$ and let $Y = X \setminus \{x\}$; then $X \sim Y$ and so $f(X) = f(Y)$. Notice that $|X \setminus f(X)|$ and $|Y \setminus f(Y)|$ have different parities; hence X and Y have different colors. \square

In both cases, we used the axiom of choice to define a “wild” coloring of $[\kappa]^\omega$. It turns out that it might be impossible to prove the above proposition without the use of the axiom of choice. It might not even be possible to prove that

$$\omega \not\rightarrow (\omega)_2^\omega,$$

holds; and a weaker choice principle like AC_ω or the axiom of dependent choice DC would not help us either.

Definition 7.28. A system of countable subsets of the natural numbers $X \subseteq [\omega]^\omega$ is said to be *Ramsey* if there exists an infinite set $A \subseteq \omega$ such that either $[A]^\omega \subseteq X$ or $[A]^\omega \cap X = \emptyset$.

Observation 7.29. $X \subseteq [\omega]^\omega$ is Ramsey \iff the coloring $\chi : [\omega]^\omega \rightarrow \{\text{red}, \text{blue}\}$ defined as

$$\chi(A) = \begin{cases} \text{red}, & \text{if } A \in X, \\ \text{blue}, & \text{if } A \notin X, \end{cases}$$

is “good:” it admits an infinite homogeneous subset.

The two “bad” colorings we constructed earlier correspond to two non-Ramsey sets. The question is: which sets are Ramsey?

Fact 7.30 (Galvin–Příkrý, 1973). *Every Borel subset of $[\omega]^\omega$ is Ramsey.*

Remark. If $X \subseteq [\omega]^\omega$ is a Borel set, then it essentially means that X can be defined using a clear, constructive rule we can actually write down, such as “contains infinitely many prime numbers.” On the other hand, using the axiom of choice is inherently non-constructive.

Fact 7.31 (Mathias, 1977). *The statement “every subset of $[\omega]^\omega$ is Ramsey” is consistent with $\text{ZF} + \text{DC}$ provided that it is consistent to assume the existence of a strongly inaccessible cardinal.*

Remark. Mathias achieved this by showing that in the model of $\text{ZF} + \text{DC} + \text{LM}$ that Solovay constructed (see Fact 7.21) is every subset of $[\omega]^\omega$ Ramsey.

Corollary 7.32. *If a strongly inaccessible cardinal exists, then $\omega \not\rightarrow (\omega)_2^\omega$ cannot be proved in $\text{ZF} + \text{DC}$.*

Remark. The question of whether $\text{ZF} +$ “every subset of $[\omega]^\omega$ is Ramsey” has a model if a strongly inaccessible cardinal does not exist is an open problem. Interestingly, this implies that if someone proves that a non-Ramsey set exists using only the tools of $\text{ZF} + \text{DC}$, then inaccessible cardinals are not consistent with ZF , and the entire large cardinal hierarchy collapses.

7.3.2 Ramsey’s Theorem

We now turn our attention to the two theorems formulated by Ramsey in 1928.

Theorem 7.33 (Infinite Ramsey’s theorem). $\omega \longrightarrow (\omega)_r^k$ for all $r, k < \omega$.

Proof. We proceed by induction on k . The base case $k = 1$ is the pigeonhole principle. Suppose $k \geq 2$ and χ is a coloring of k -sets of ω . For $t \in \omega$ define a coloring of $(k - 1)$ -sets of $\omega \setminus \{t\}$ as

$$\chi_t(A) := \chi(A \cup \{t\}).$$

Let $t_0 := 0$ and consider the coloring χ_{t_0} of $[\omega \setminus \{t_0\}]^{k-1}$. Use the induction hypothesis for $k - 1$ to obtain an infinite $A_0 \subseteq \omega \setminus \{t_0\}$ homogeneous for χ_{t_0} . Notice that every k -set of $\{t_0\} \cup A_0$ that contains t_0 has the same color with respect to χ .

We repeat this process. Let $t_1 := \min(A_0)$ and consider the coloring χ_{t_1} of $[A_0 \setminus \{t_1\}]^{k-1}$. By using the induction hypothesis, we obtain an infinite $A_1 \subseteq A_0 \setminus \{t_1\}$ homogeneous for χ_{t_1} . Note that every k -set of $\{t_0\} \cup \{t_1\} \cup A_1$ that contains either t_0 or t_1 has the same color (but they might be different).

Repeating this process indefinitely, we construct an infinite sequence t_0, t_1, t_2, \dots such that for every t_n it holds that every k -set of $\{t_i \mid i < \omega\}$ containing t_n has the same color. Since there are only finitely many colors, by the pigeonhole principle, there exists an infinite subsequence $t_{i_0}, t_{i_1}, t_{i_2}, \dots$ such that all t_{i_j} correspond to the same color. Therefore, the set $\{t_{i_j} \mid j < \omega\}$ is homogeneous for χ . \square

Notice that this proof relies on the fact that ω is well-ordered.

Fact 7.34 (Kleinbern, 1969). *In ZF, it is impossible to prove that for every infinite set X and for every finite coloring $\chi : [X]^k \rightarrow r$ there exists an infinite $A \subseteq X$ homogeneous for χ .*

We also present a proof based on the more general method of uniform ultrafilters, which

Exercise 7.35. Prove from infinite Ramsey's theorem (assuming AC) that

- (a) Every infinite ordered set has an infinite chain or an infinite antichain.
- (b) Every infinite linearly ordered set contains an infinite increasing sequence or an infinite decreasing sequence.

Theorem 7.36 (Finite Ramsey's theorem). *For all $n, r, k < \omega$ there exists $N < \omega$ such that $N \rightarrow (n)_r^k$.*

We have already encountered Ramsey's theorems when we discussed the compactness principle, and we have seen that if we assume \mathbf{AC}_ω , then it is possible to deduce the finite version from the infinite one (see Proposition 7.15). However, as one might suspect, \mathbf{AC}_ω is not *needed* to prove this theorem.

Proof sketch. There are more ways to prove the finite version of Ramsey's theorem, but after examining the proof of the infinite version presented above, it should be clear that the same method can be used to prove the finite version. In the finite case, we do not need infinitely many t_i to use pigeonhole at the end — $(n-1)r+1$ is surely enough. The only difference is that at each step t_i , we need to invoke the induction hypothesis for a sufficiently large n_i so that each A_i is large enough to keep the induction going. \square

The numbers $N(k, r, n)$ are called *Ramsey numbers*, and they are an active area of research, as their asymptotics are particularly difficult to pin down — even in the simplest case $k = 2$ and $r = 2$.

Definition 7.37. Define $R(n)$ as the least N such that $N \rightarrow (n)_2^2$.

Proposition 7.38. $2^{n/2} \leq R(n) \leq 2^{4n}$

Proof sketch. The upper bound can be deduced by simply examining the proof of finite Ramsey's theorem presented above. The lower bound is due to Erdős (1947) and his famous probabilistic argument, which goes as follows:

We are coloring $[N]^2$ randomly using 2 colors. For every $S \in [N]^n$ define a random variable X_S as 1 if S is homogeneous, and 0 if it is not. Define X_n as the sum of all these variables X_S . Then

$$\mathbb{E}[X_S] = \text{Pr}[S \text{ colored homogeneously}] = \frac{\#\text{good colorings}}{\#\text{all colorings}} = \frac{2}{2^{\binom{n}{2}}} = 2^{1-\binom{n}{2}}.$$

Thus

$$\mathbb{E}[X_n] = \sum_{|S|=n} \mathbb{E}[X_S] = \binom{N}{n} 2^{1-\binom{n}{2}}.$$

If we pick N small enough so that $\mathbb{E}[X_n] < 1$, then there must exist a coloring of N that has 0 homogeneous subsets S of size n . In that case, $R(n) > N$. By manipulating the equation $\binom{N}{n} 2^{1-\binom{n}{2}} < 1$ and using a bound on $\binom{N}{n}$, one can arrive at $2^{n/2} \leq R(n)$ (or, in fact, even something slightly better). \square

Remark. As we have seen, these bounds are fairly easy to derive, but they seem to be incredibly difficult to improve by any significant amount (although there have been some improvements).

7.3.3 An Unprovable Theorem

Recall Goodstein's theorem, and the fact that it cannot be proved in Peano arithmetic PA. In 1977, Paris and Harrington [27] found a natural extension of (finite) Ramsey's theorem that cannot be proved in PA. It goes as follows:

Definition 7.39. A set $A \subseteq \omega$ is said to be *large* if $|A| \geq \min(A)$.

Definition 7.40. For natural numbers n, r, k and N , the expression $N \xrightarrow{*} (n)_r^k$ means that for every coloring $\chi : [N]^k \rightarrow r$ there exists a large subset $A \subseteq N$ of size at least n that is homogeneous for χ .

Intuition. So $N \xrightarrow{*} (n)_r^k$ means exactly the same as $N \rightarrow (n)_r^k$, we only require the homogeneous set we find to be *large*; if we start late (its minimum is large), then the set also has to be large.

Theorem 7.41 (Paris–Harrington, AC_ω). *For all $n, r, k < \omega$ there exists $N < \omega$ such that $N \xrightarrow{*} (n)_r^k$.*

Proof. We use a compactness argument similar to when we proved Proposition 7.15. For given n, k , and r , we define a hypergraph $H = (V, E)$ where

$$V = [\omega]^k, \quad E = \{[A]^k \mid A \subseteq \omega \text{ such that } n \leq |A| < \omega \text{ and } A \text{ is large}\}.$$

Since $|[\omega]^k| = \omega$, assuming AC_ω is enough to invoke the compactness principle.

Infinite Ramsey's theorem implies that for every r -coloring of $[\omega]^k$, there exists an *infinite* $A \subseteq \omega$ such that $[A]^k$ is monochromatic. Let A' be the set containing the first $\max\{n, \min(A)\}$ elements of A . Then $[A']^k$ is a monochromatic edge of H . Therefore $\chi(H) > r$. By the compactness principle, there exists a finite $W \subseteq V$ such that $\chi(H_W) > r$: every r -coloring of H_W contains a monochromatic edge. We let $N \in \omega$ be the first number larger than all the natural numbers explicitly mentioned in W . \square

Theorem 7.42 (Paris–Harrington). *Theorem 7.41 cannot be proved in PA.*

Intuition. This implies that some kind of transfinite argument (like AC_ω) at least as strong as transfinite induction up to ε_0 is needed to prove this theorem.

Proof idea. Paris and Harrington originally proved this via the black magic of model theory. However, Ketonen and Solovay (1981) presented a purely combinatorial proof in which they showed that the numbers N from the theorem grow roughly like f_{ε_0} . Fact 3.88 implies that PA cannot prove that the numbers N exist for all n, k and r . \square

7.3.4 Transfinite Partition Relations

All of the statements in this section assume AC, as it is needed to define cardinal powers. Hence, every set can be well-ordered and is identical to a cardinal.

Imagine you are organizing a big party (for infinitely many people), where each pair of attendants must agree in advance on a language they will speak. The following exercise claims that if there are κ different languages, then there exists a party of 2^κ people in which a dialogue is impossible.

Exercise 7.43. Show that if κ is an infinite cardinal, then $2^\kappa \not\rightarrow (3)_\kappa^2$.

Hint. Define a suitable coloring of pairs of sequences of 0s and 1s of length κ . If you are stuck, try to show $\kappa \not\rightarrow (3)_\kappa^2$ first.

We know that $\omega \rightarrow (\omega)_2^2$. Does $\omega_1 \rightarrow (\omega_1)_2^2$? Sierpiński showed that it does not. In fact, not even 2^ω , which could potentially be much larger, is enough.

Proposition 7.44 (Sierpiński). $2^\omega \not\rightarrow (\omega_1)_2^2$.

To prove this proposition, we first need a lemma.

Lemma 7.45. *No subset of \mathbb{R} is ordered according to ω_1 or ω_1^* (the reverse order).*

Proof. Suppose for contradiction that there is an increasing sequence of real numbers $\langle x_\alpha \mid \alpha \in \omega_1 \rangle$ of length ω_1 . Then, for every $\alpha < \omega_1$, there exists a rational number $q \in (x_\alpha, x_{\alpha+1})$ since \mathbb{Q} is dense. Take q_α to be the smallest from this interval according to some well-ordering of \mathbb{Q} . We get an injective map $\omega_1 \rightarrow \mathbb{Q}$, a contradiction. The proof for a decreasing sequence is analogous. \square

Proof of Proposition 7.44. Let \prec be a well-ordering of \mathbb{R} and define a coloring $\chi : [\mathbb{R}]^2 \rightarrow 2$. If $x, y \in \mathbb{R}$ and $x < y$, let

$$\chi(\{x, y\}) = \begin{cases} 0, & \text{if } x \prec y, \\ 1, & \text{if } y \prec x. \end{cases}$$

Suppose that there exists a homogeneous subset $A \subseteq \mathbb{R}$ of size ω_1 . If it has color 0, then the standard order $<$ of \mathbb{R} is a well-order on A since it agrees with \prec . Because $|A| = \omega_1$, the order type of $(A, <)$ is at least ω_1 , and an initial segment of $(A, <)$ has type ω_1 , which contradicts the lemma. Similarly, if we assume that A has color 1, then $(A, <)$ is a reverse well-ordered set of type at least ω_1^* . \square

Theorem 7.46 (Generalized Sierpiński). *For every infinite cardinal κ we have*

$$2^\kappa \not\rightarrow (\kappa^+)_2^2.$$

Proof. The strategy is similar to that in Proposition 7.44. We will find a linearly ordered set of size 2^κ that has no subset ordered according to κ^+ or $(\kappa^+)^*$. Consider the set ${}^\kappa 2$ of all mappings $g : \kappa \rightarrow 2$ with the lexicographic order. Assume that $\langle g_\alpha \mid \alpha < \kappa^+ \rangle$ is an increasing sequence of length κ^+ .

We will use recursion to construct a function $h : \kappa \rightarrow 2$ and a non-decreasing sequence $\langle \alpha_\xi \mid \xi < \kappa \rangle$ such that $\alpha_\xi < \kappa^+$ and for all $\beta \geq \alpha_\xi$ we have $h \upharpoonright (\xi + 1) = g_\beta \upharpoonright (\xi + 1)$. Notice that this implies a contradiction: if we let $\alpha := \sup\{\alpha_\xi \mid \xi < \kappa\}$,

then from the regularity of κ^+ we have $\alpha < \kappa^+$, but from the construction of h we have $h = g_\beta$ for all $\beta \geq \alpha$. This is a contradiction since $g_\alpha \neq g_{\alpha+1}$.

Construction of h and the sequence: if for every $\alpha < \kappa^+$ we have $g_\alpha(0) = 0$, then we let $h(0) = 0$ and $\alpha_0 = 0$. If $g_\alpha(0) = 1$ for some α , then we let α_0 be the first ordinal with this property, and we define $h(0) = 1$. It is clear that for every $\beta \geq \alpha_0$, we have $g_\beta(0) = h(0) = 1$ since the mappings g_α are ordered lexicographically. Suppose that the values $h(\eta)$ and α_η have already been constructed for all $\eta < \xi$. Let $\beta = \sup\{\alpha_\eta \mid \eta < \xi\}$. If for every $\alpha \geq \beta$ we have $g_\alpha(\xi) = 0$, then we let $h(\xi) = 0$ and $\alpha_\xi = \beta$. If $g_\alpha(\xi) = 1$ for some $\alpha \geq \beta$, then we let α_ξ be the first ordinal with this property, and we define $h(\xi) = 1$.

The proof for a decreasing sequence is analogous: the 1s and 0s only switch roles when defining h . \square

Definition 7.47. For a cardinal κ and a natural number $n \geq 1$ define $\kappa \uparrow n$ as

$$\kappa \uparrow 1 := \kappa, \quad \kappa \uparrow (n + 1) := \kappa^{\kappa \uparrow n}.$$

Observation 7.48. If κ is infinite, then

$$\kappa \uparrow n = \kappa^{\cdot^{\cdot^{\cdot^{\kappa^{\kappa}}}}} = 2^{\cdot^{\cdot^{\cdot^{\cdot^{\kappa}}}}}$$

where the number of 2s in the tower is $n - 1$.

Proof. Theorem 4.52 implies that if κ is infinite, then $\kappa^\lambda = 2^\lambda$ for all $\lambda \geq \kappa$. \square

Fact 7.49 (Erdős–Rado, 1956). For every infinite cardinal κ and every natural number n it holds that

$$(\kappa \uparrow n)^+ \longrightarrow (\kappa^+)_\kappa^n.$$

In particular,

$$(2^\kappa)^+ \longrightarrow (\kappa^+)_2^2.$$

Observation 7.50. Sierpiński’s theorem implies that the value $(2^\kappa)^+$ in the last arrow cannot be improved.

Fact 7.51 (Erdős–Dushnik–Miller, 1941). For every infinite cardinal κ we have

$$\kappa \longrightarrow (\kappa, \omega)^2 \tag{7.3}$$

In the language of graph theory, this means that every infinite graph contains either a countably infinite independent set, or a clique with the same cardinality as the whole graph.

Remark. The arrow (7.3) no longer holds for dimension $k = 3$.

Exercise 7.52. Show that if \preceq_1 and \preceq_2 are two well-orderings of an infinite set X , then there is a subset $Y \subseteq X$ of the same cardinality as X , on which the orders agree.

Hint. Define a 2-coloring of $[X]^2$ based on whether the orders agree or not.

We have only scratched the surface of this topic. For a comprehensive collection of results regarding transfinite partition relations, see [11].

7.3.5 Partitions and Large Cardinals

Weakly compact cardinals Sierpiński’s theorem implies that if κ is a successor cardinal, then the arrow

$$\kappa \longrightarrow (\kappa, \omega)^2$$

from the Erdős–Dushnik–Miller theorem cannot be improved to

$$\kappa \longrightarrow (\kappa, \kappa)^2. \tag{7.4}$$

Ramsey’s theorem implies that (7.4) holds when $\kappa = \omega$, but does it hold for any uncountable limit cardinals?

Definition 7.53 (Weakly compact cardinal). An infinite cardinal κ is *weakly compact* if κ is uncountable and $\kappa \rightarrow (\kappa)_2^2$.

Exercise 7.54. Show that all weakly compact cardinals are strongly inaccessible. Thus ZFC cannot prove the existence of weakly compact cardinals, and it is consistent to assume that ω is the only infinite cardinal for which (7.4) holds.

Hint. To show that κ has to be strongly limit, use Sierpiński’s theorem. To show regularity, assume that there exists an increasing sequence of ordinals $\langle \alpha_i \mid i \in \lambda \rangle$ for some $\lambda < \kappa$ such that $\sup \alpha_i = \kappa$, and deduce a contradiction by defining a suitable 2-coloring of $[\kappa]^2$.

Fact 7.55. *Weakly compact cardinals are not only inaccessible, but also Mahlo, and even hyper-Mahlo.*

Definition 7.56. We say that an infinite cardinal κ has the *tree property* if every tree of height κ with levels of size strictly less than κ has a cofinal branch.

König’s lemma is equivalent to saying that ω has the tree property.

Fact 7.57. κ is weakly compact $\iff \kappa$ is strongly inaccessible and has the tree property.

Weakly compact cardinals can be similarly characterized using many other interesting properties (see for example [1] or [24]). We mention that Ramsey’s theorem generalizes to weakly compact cardinals in its full power: if κ is weakly compact, then not only $\kappa \rightarrow (\kappa)_2^2$, but in fact

$$\kappa \longrightarrow (\kappa)_\mu^k$$

for all $k \in \omega$ and $\mu < \kappa$.

Another interesting equivalent characterization of weakly compact cardinals κ is that they are Π_1^1 -inaccessible. Formally, φ is a Π_1^1 -formula if it looks like $\varphi = (\forall X) \phi(X)$, where $\forall X$ is a second-order quantifier (quantifying over subsets of the universe), and ϕ is a first order formula.

We say that κ is Π_1^1 -inaccessible if for all subsets $R \subseteq V_\kappa$ (which act as unary predicates) and all Π_1^1 sentences φ , whenever $\langle V_\kappa, \in, R \rangle \models \varphi$, then there is $\alpha < \kappa$ such that $\langle V_\alpha, \in, R \cap V_\alpha \rangle \models \varphi$. Intuitively, this means that if we try to describe κ using a property that is not “too complex” (describable using a Π_1^1 formula with a parameter R), then there exists a smaller cardinal that already satisfies this property.

It should be noted that we use full semantics for the second order quantifiers. For example, if $M \models (\forall X)(\exists y)(\forall z)\phi(X, y, z)$, then the first order variables y and z range over M , and the second order variable X ranges over $\mathcal{P}(M)$.

TODO prove that indescribability implies Mahlo (proof in Jech)

TODO prove equivalent characterizations of weakly compact cardinals (tree property, general Ramsey, ultrafilters), based on proofs in Jech and Balcar–Štěpánek.

Ineffable cardinals Weakly compact cardinals generalize Ramsey’s theorem, guaranteeing a homogeneous subset $A \subseteq \kappa$ of size κ . If we further require A to be stationary, we get *ineffable cardinals*. They were first introduced by Jensen and Kunen [21] and then generalized and thoroughly studied by Baumgartner [2] and Friedman [12].

Definition 7.58. For cardinals μ and κ such that $\text{cf}(\kappa) > \omega$, and a natural number $k \in \omega$, the expression

$$\kappa \longrightarrow (\text{Stac})_{\mu}^k$$

means that for every mapping $\chi : [\kappa]^k \rightarrow \mu$, there is a subset $A \subseteq \kappa$ stationary in κ , such that $\chi \upharpoonright [A]^k$ is constant.

Definition 7.59 (Ineffable cardinal). A cardinal κ is *ineffable* if it is uncountable, regular, and $\kappa \rightarrow (\text{Stac})_2^2$.

Observation 7.60. *Every ineffable cardinal is weakly compact.*

Fact 7.61. *Ineffable cardinals are stationary limits of subtle cardinals, which are in turn stationary limits of totally indescribable cardinals. Totally indescribably cardinals are much stronger than weakly compact cardinals which are merely Π_1^1 -indescribable.*

Ineffable cardinals strengthen Ramsey’s theorem to stationary subsets, but not in its full generality. Jensen and Kunen (1969) showed that if κ is ineffable, then

$$\kappa \longrightarrow (\text{Stac})_{\mu}^2$$

holds for all cardinals $\mu < \kappa$, but Baumgartner showed that

$$\kappa \longrightarrow (\text{Stac})_2^3$$

does not hold for all ineffable cardinals.

Definition 7.62 (Totally ineffable cardinal). A cardinal κ is *n-ineffable* if it is uncountable, regular, and $\kappa \rightarrow (\text{Stac})_2^{n+1}$. It is *totally ineffable* if it is *n-ineffable* for all $1 \leq n < \omega$.

Observation 7.63. *Every uncountable regular cardinal κ is 0-ineffable.*

Proof. Suppose we partition κ into two disjoint sets $A \cup B = \kappa$. If neither of these sets were stationary, then there would exist clubs C_A and C_B such that $C_A \cap A = \emptyset$ and $C_B \cap B = \emptyset$. Since the intersection of two clubs is a club, the set $C := C_A \cap C_B$ is a club in κ , but $C \cap \kappa = \emptyset$. \square

Note that 1-ineffable cardinals are exactly the ineffable cardinals, and that totally ineffable cardinals strengthen the stationary version of Ramsey's theorem to all dimensions.

Fact 7.64 (Baumgartner, 1975). *For $n \geq 2$, all n -ineffable cardinals are stationary limits of $(n - 1)$ -ineffable cardinals.*

Erdős and Ramsey cardinals When we discussed large cardinal axioms in Section 5.5, we mentioned that weakly compact cardinals are still consistent with the axiom of constructibility $\mathbf{V} = L$. *Erdős* and *Ramsey* cardinals are among the first large cardinals powerful enough to contradict $\mathbf{V} = L$, as they imply the existence of certain sets that are too “rich” to be constructible.

Definition 7.65. For cardinals κ, μ , an ordinal α , and $k \leq \omega$, the expression

$$\kappa \longrightarrow (\alpha)_\mu^{<k}$$

means that for every mapping $\chi : [\kappa]^{<k} \rightarrow \mu$, there exists a subset $A \subseteq \kappa$ ordered according to α such that for every $n < k$ is $f \upharpoonright [A]^n$ constant. Note that the value of f on $[A]^n$ might differ for different n .

Intuition. If A is ordered according to α , then $|A| = |\alpha|$, so this is a direct generalization of the cardinal partition arrow for ordinals, and for finding a subset that is homogeneous for all $n < k$.

Theorem 7.66. *For all $n, r, k < \omega$ there exists $N < \omega$ such that $N \rightarrow (n)_r^{<k}$.*

Proof. We will define a sequence N_1, N_2, \dots, N_{k-1} by repeated application of Ramsey's theorem and take $N := N_{k-1}$. Let

$$N_1 \rightarrow (n)_r^1, \quad N_2 \rightarrow (N_1)_r^2, \quad N_3 \rightarrow (N_2)_r^3, \quad \dots, \quad N_{k-1} \rightarrow (N_{k-2})_r^{k-1}.$$

Let χ be an r -coloring of $[N_{k-1}]^{<k}$. We use finite Ramsey's theorem for $k - 1$ and find $S_{k-1} \subseteq N_{k-1}$ of size N_{k-2} such that $[S_{k-1}]^{k-1}$ is monochromatic. Then we repeat this for $k - 2$ and find $S_{k-2} \subseteq S_{k-1}$ of size N_{k-3} such that $[S_{k-2}]^{k-2}$ is monochromatic. We construct a sequence $S_1 \subseteq S_2 \subseteq \dots \subseteq S_{k-1}$ such that $|S_1| = n$ and for each i is $[S_i]^i$ monochromatic. \square

Theorem 7.67. $\omega \rightarrow (\omega)_r^{<k}$ holds for all $k, r < \omega$.

Proof. The same as the previous theorem, we only apply the infinite version of Ramsey's theorem instead of the finite one. \square

Proposition 7.68. $\omega \not\rightarrow (\omega)_2^{<\omega}$.

Proof. We define a bad coloring $\chi : [\omega]^{<\omega} \rightarrow 2$. For $A \in [\omega]^k$ let

$$\chi(A) := \begin{cases} 0, & \text{if } k \notin A, \\ 1, & \text{if } k \in A. \end{cases}$$

Now let S be an infinite subset of ω and let $k = \min(S)$. Clearly $[S]^k$ contains both sets of color 0 (k -sets without k) and sets of color 1 (k -sets with k). \square

Earlier, we saw that $\kappa \not\rightarrow (\omega)_2^\omega$ for any cardinal κ . Paul Erdős asked if there is a cardinal κ such that $\kappa \rightarrow (\omega)_2^{<\omega}$. It turns out that the existence of such a cardinal cannot be proved, but no-one has found a way to disprove it either.

Definition 7.69 (α -Erdős cardinal). Let $\alpha \geq \omega$ be an infinite ordinal. An infinite cardinal κ is α -Erdős if $\kappa \rightarrow (\alpha)_2^{<\omega}$. The least cardinal with this property is denoted by $\kappa(\alpha)$.

Remark. The term “ α -Erdős cardinal” is sometimes reserved only for $\kappa(\alpha)$.

Clearly, if $\alpha < \beta$, then all β -Erdős cardinals are also α -Erdős, so $\kappa(\omega)$ is the smallest Erdős cardinal. Note that Proposition 7.68 implies that $\kappa(\omega) > \omega$, so all Erdős cardinals are uncountable. It can be shown (see for example [20]) that if α is limit, then

$$\kappa(\alpha) \longrightarrow (\alpha)_\mu^{<\omega}$$

holds for all $\mu < \kappa(\alpha)$.

Fact 7.70 (Erdős–Hajnal, 1958). $\kappa(\alpha)$ is a regular cardinal for all $\alpha \geq \omega$, and if $\beta > \alpha$, then $\kappa(\beta) > \kappa(\alpha)$.

Fact 7.71 (Jensen–Kunen, 1969). If α is limit, then $\kappa(\alpha)$ is not only strongly inaccessible, but subtle.

Erdős cardinals $\kappa(\alpha)$ do not have to be weakly compact nor ineffable. However, Friedman [12] showed that if an ω -Erdős cardinal exists, then there exists a smaller cardinal that is totally ineffable.

Fact 7.72 (Rowbottom, 1964 [29]). If an ω_1 -Erdős cardinal exists, then $\mathbf{V} \neq L$.

Fact 7.73 (Silver, 1970 [31]). The existence of an ω -Erdős cardinal is consistent with the axiom of constructibility $\mathbf{V} = L$. Moreover, the existence of a cardinal κ such that κ is α -Erdős for all $\alpha < \omega_1$ is also consistent with $\mathbf{V} = L$. Thus ω_1 is the exact break-off point as far as the axiom of constructibility is concerned.

Silver then remarks that it does not seem extravagant to assert that for all practical purposes, $(\exists \kappa) \kappa \rightarrow (\omega)_2^{<\omega}$ and $(\exists \kappa)(\forall \alpha < \omega_1) \kappa \rightarrow (\alpha)_2^{<\omega}$ are the strongest large cardinal axioms known to be consistent with $\mathbf{V} = L$.

Definition 7.74 (Ramsey cardinal). κ is Ramsey if $\kappa \rightarrow (\kappa)_2^{<\omega}$.

In other words, κ is Ramsey if it is κ -Erdős, so Ramsey cardinals are the fixed points of the ordinal function $\alpha \mapsto \kappa(\alpha)$. If κ is Ramsey, then

$$\kappa \longrightarrow (\kappa)_\mu^{<\omega}$$

for every $\mu < \kappa$ because κ is κ -Erdős. However, notice that Exercise 7.43 implies that

$$\kappa \longrightarrow (\kappa)_\kappa^{<\omega}$$

cannot be satisfied by any cardinal κ .

Proposition 7.68 demonstrates that ω is not a Ramsey cardinal (which is a bit ironic). Hence, all Ramsey cardinals are uncountable. If $\kappa \rightarrow (\kappa)_2^{<\omega}$, then also $\kappa \rightarrow (\kappa)_2^2$, so every Ramsey cardinal is weakly compact and at least ω_1 -Erdős. Therefore, the existence of a Ramsey cardinal implies that $\mathbf{V} \neq L$.

Ramsey cardinals are the strongest large cardinals we mention, but they are still nothing compared to the higher levels of the large cardinal hierarchy.

7.4 Banach–Tarski Paradox

We have already mentioned the Banach–Tarski paradox when we talked about the axiom of choice and its paradoxical consequences: we can decompose a solid ball into finitely many pieces and reassemble them (by applying rotations and translations) into two identical copies of the original. We say that the ball and its two copies are *equidecomposable*.

Definition 7.75. We say that two subsets $A, B \subseteq \mathbb{R}^3$ are

- (a) *congruent* $A \simeq B$ if B can be obtained from A by translations and rotations.
- (b) *equidecomposable* using n pieces $A \stackrel{n}{\simeq} B$ if there exist partitions into disjoint sets $A = A_1 \cup \dots \cup A_n$ and $B = B_1 \cup \dots \cup B_n$ such that $A_i \simeq B_i$ for all i .

Furthermore, if $A \stackrel{n}{\simeq} B'$ where $B' \subseteq B$, then we write $A \stackrel{n}{\preceq} B$ and we say that A is *sub-equidecomposable* into B .

The Banach–Tarski paradox can also be viewed as a particularly bizarre manifestation of the failure of finite additivity of the outer Lebesgue measure in ZFC (see Appendix A.2 for more details), and it can be generalized to higher dimensions as well. Interestingly, a 2-dimensional analog of this paradox is impossible, as it can be shown that two equidecomposable planar geometric figures must have the same area. However, if we also allow area-preserving affine transformations, then it is possible to turn a unit square into two planar figures with the same area as the original. This is known as the *von Neumann paradox*.

Observation 7.76. Let $A, C \subseteq \mathbb{R}^3$ be disjoint and $B, D \subseteq \mathbb{R}^3$ be disjoint. Then

- (a) if $A \stackrel{n}{\simeq} B$ and $C \stackrel{m}{\simeq} D$, then $A \cup C \stackrel{n+m}{\simeq} B \cup D$,
- (b) if $A \stackrel{n}{\preceq} B$ and $C \stackrel{m}{\preceq} D$, then $A \cup C \stackrel{n+m}{\preceq} B \cup D$.

Observation 7.77 (Transitivity). Let $A, B, C \subseteq \mathbb{R}^3$. Then

- (a) if $A \stackrel{n}{\simeq} B \stackrel{m}{\simeq} C$, then $A \stackrel{n \cdot m}{\simeq} C$,
- (b) if $A \stackrel{n}{\preceq} B \stackrel{m}{\preceq} C$, then $A \stackrel{n \cdot m}{\preceq} C$.

Proof. The reason is that the n pieces considered by the first partition intersect with the m pieces considered by the second partition. And there can be up to $n \cdot m$ intersections. Draw a picture in 2 dimensions if this is not clear. \square

TODO polish the rest of this section

Theorem 7.78. If $A \stackrel{n}{\preceq} B$ and $B \stackrel{m}{\preceq} A$, then $A \stackrel{n+m}{\simeq} B$.

Remark. This can be seen as a generalization of the Cantor–Bernstein theorem, and it is attributed to Banach (1924).

Proof. The goal is to find partitions into disjoint sets $A = A_1 \cup A_2$ and $B = B_1 \cup B_2$ such that $A_1 \overset{n}{\simeq} B_1$ and $B_2 \overset{m}{\simeq} A_2$. Denote by $\varphi : A \rightarrow B$ the injection induced by $A \overset{n}{\simeq} B$, and by $\psi : B \rightarrow A$ the injection induced by $B \overset{m}{\simeq} A$.

To obtain the desired partition, follow your favorite proof of the Cantor–Bernstein theorem. We will, for every $a \in A$, consider the sequence

$$\psi^{-1}(a), \quad \varphi^{-1}(\psi^{-1}(a)), \quad \psi^{-1}(\varphi^{-1}(\psi^{-1}(a))), \quad \dots$$

This sequence could be empty, finite, or infinite (we stop if the preimage does not exist). Denote by A_L the set

$$\{a \in A \mid \text{the sequence of preimages is of even length}\},$$

by A_O odd length sequences, and by A_∞ the infinite sequence starting points. So $A = A_L \cup A_O \cup A_\infty$. Similarly, we partition $B = B_L \cup B_O \cup B_\infty$. Now define $A_1 = A_L \cup A_\infty$, $A_2 = A_O$ and $B_1 = B_O \cup B_\infty$, $B_2 = B_L$. This works because when we add one more point to a odd sequence we get an even sequence and vice versa. \square

Exercise 7.79. Prove $S^1 \simeq^2 S^1 \setminus \{(1,0)\}$ where S^1 is the unit circle. So we can fill in a missing point in the circle for free. Similarly show that $D \simeq^{n+2} D \cup n$ line segments of length 1, where D is the unit disk (we rotate the radii). So we can generate line segments for free.

Proposition 7.80. $S \subseteq \mathbb{R}^3$ unit sphere and $D \subset S$ a countable subset, $D' := S \setminus D$. Then $S \simeq^2 D'$. So we can fill in countably many points in the sphere.

Intuition. In the unit circle version, consider that it is missing more than one point - then we can make sequences from each point and rotate them all at once. But one has to be careful if the points are spaced by some multiple of pi radians.

Proof. There exists an axis disjoint from D and a rotation α about some angle such that $\alpha[D], \alpha[\alpha[D]], \dots$, are disjoint. So every point will have its own circle missing one point. How can we pick this axis? No two points from D should rotate on each other, so for all pairs a, b from D we need to forbid all axis contained in the hyperplane orthogonal to the segment ab . D is countable, so we are forbidding countably many hyperplanes, so after we remove them, the measure of what we are left with is still the same as before. Hence we can just pick a random axis and with probability 1 it will work with any angle. Or we don't care about the axis but need to be careful about the angle. Now define partitions $A = D \cup \alpha[D] \cup \alpha[\alpha[D]] \cup \dots$ and $B = S \setminus A$. Then $A \simeq \alpha[A]$ and $B \simeq B$, so $S = A \cup B \simeq^2 \alpha[A] \cup B = (A \setminus D) \cup B = D'$. \square

Theorem 7.81 (Banach–Tarski). $\bar{S}, \bar{S}_n \subset \mathbb{R}^3$ disjoint unit balls (closed), then $\bar{S} \simeq^{10} \bar{S} \cup \bar{S}_1$. But it can be done using a smaller number of parts.

The theorem relies on the following crucial fact.

Fact 7.82. There exist rotations α by 180° , and β by 120° such that (except for $\alpha^2 = \beta^3 = e$) there is no other finite composition of α and β (for example $\alpha\beta\alpha\beta^2$) is equal to the identity. In group theory terms, α and β generate a group of rotations that is isomorphic to the free product $\mathbb{Z}_2 * \mathbb{Z}_3$ (that is $2^3 = 1$ and $3^2 = 1$, but we cannot get 1 any other way). Osofsky and Adams proved that we can take any two axes (going through the origin) with angle $\theta = 45^\circ$ between them.

The Banach–Tarski paradox does not work in \mathbb{R}^2 because no rotations with similar properties exist in \mathbb{R}^2 . But if we allow some other affine transformations (not only rotation and translation) that we can create a similar “paradox” already in \mathbb{R}^2 .

Proof of theorem. Take rotations α and β from the previous fact and the group G generated by α and β . Elements of G are the identity e and rotations of the form $(\alpha)\beta^{\epsilon_1}\alpha\beta^{\epsilon_2}\alpha\beta^{\epsilon_3}\dots$ where $\epsilon_i \in \{1,2\}$ (the rotations are applied from right to left). Think about it, these are all the compositions and (α) denotes that it might be missing. One might need to do some linear algebra to show that every element of G is a rotation about some axis (composition of two rotations is a rotation).

Denote by D the set of all intersections of the axis of the rotations from G with the ball S and note that D is countable. For every $\delta \in G$ we have $\delta[D] = D$. Indeed, if $x \in D$ then there is $\gamma \in G$ s.t. $\gamma(x) = x$ (the rotation corresponding to the axis). We define the rotation $\delta\gamma\delta^{-1}$ and notice that $\delta(x)$ is a fixed point of this rotation (hence it lies on its axis), so $\delta(x) \in D$. Similarly we consider the rotation $\delta^{-1}\gamma\delta$ with fixed point $\delta^{-1}x$ to show the other inclusion.

Denote by $D' := S \setminus D$, by definition D' has no fixed points of nontrivial elements of G , and by our previous observation if $x \in D'$ and $\delta \in G$ then $\delta(x) \in D'$. For $x \in D'$ define $S_x :=$ the orbit of x (from group theory) as $\{\gamma(x) \mid \gamma \in G\}$ that is all points where x can be mapped. Notice that these orbits form a partition of D' . Using the axiom of choice, choose one point from every orbit and collect them into a set T . Every element of D' can be uniquely expressed as $\gamma(t)$ where $\gamma \in G$ and $t \in T$ because $x \in D'$.

Define

$$A := \{\gamma(t) \mid t \in T, \gamma = e \vee \gamma = \alpha\beta^{\epsilon_1}\alpha\beta^{\epsilon_2}\alpha\beta^{\epsilon_3}\dots\}$$

and

$$B := \{\gamma(t) \mid t \in T, \gamma = \beta\alpha\beta^{\epsilon_1}\alpha\beta^{\epsilon_2}\dots\}$$

and

$$C := \{\gamma(t) \mid t \in T, \gamma = \beta^2\alpha\beta^{\epsilon_1}\alpha\beta^{\epsilon_2}\dots\}$$

Clearly $D' = A \cup B \cup C$ is a partition and $A \simeq B \simeq C \simeq A$ as $B = \beta[A]$ and $C = \beta[B]$ and $A = \beta[C]$. Also, $\alpha[B \cup C] \subset A \simeq C$. So $A \preceq^1 B$ and $B \cup C \preceq^1 C$, together $A \cup B \cup C \preceq^2 B \cup C \preceq^1 A$, so $D' = A \cup B \cup C \preceq^2 A$. Specifically, we apply the rotations $\beta^2\alpha[B \cup C] \subseteq C$ and $\beta[A] = B$. Since $\alpha[C \cup B] = A$, we apply the rotation $\alpha\beta^2\alpha$ to $B \cup C$ and $\alpha\beta$ to A to get $B \cup C \cup A \preceq^2 A$.

By the proposition about the countable subset of a sphere $D \cup D' = S \simeq^2 D' \preceq^2 A$, so $S \preceq^4 A$, similarly $S_1 \preceq^4 B$ and $S \cup S_1 \preceq^8 A \cup B$. We extend to balls \bar{S}, \bar{S}_1 and consider sets \bar{A}, \bar{B} of radii segments $0a, a \in A$ and $0b, b \in B$. Then remove origins and used the mappings for spheres $\bar{S}_0 - 0 \cup \bar{S}_1 - 0' \preceq^8 \bar{A} \cup \bar{B} \preceq \bar{S}$. There is still a lot of points in \bar{S} which are not included in $\bar{A} \cup \bar{B}$, so we can map the origins to those. We can map 0 to 0 and $0'$ into some arbitrary $y \in C$. that adds one more piece and we have $\bar{S} \cup \bar{S}_1 \preceq^9 \bar{S}$. Cantor–Bernstein together with $\bar{S} \preceq^1 \bar{S}$ give $S \cup \bar{S}_1 \simeq^{10} \bar{S}$. \square

Exercise 7.83. There exists n such that carbon atom \simeq^n the Sun.

Exercise 7.84. Similar construction does not work in \mathbb{R}^2 . Show that if α is a rotation by 180 and β by 120, then $(\alpha\beta)^6 = e$. And find an identity that is satisfied by any pair of rotations α, β in \mathbb{R}^2 .

Appendix

A.1 Group Theory

Definition A.1. A *commutative group* is a set G equipped with a binary operation $+: G \times G \rightarrow G$ that satisfies:

- (i) $(\forall x, y, z \in G) x + (y + z) = (x + y) + z$, ... associativity
- (ii) $(\forall x, y \in G) x + y = y + x$, ... commutativity
- (iii) $(\exists e \in G)(\forall x \in G) e + x = x + e = x$, ... neutral element
- (iv) $(\forall x \in G)(\exists y \in G) x + y = e$ inverse element

It is easy to show that the *neutral* and *inverse* elements are uniquely determined, and we denote them by 0 and $-x$ respectively. We denote by $x - y$ the element $x + (-y)$. It is also easy to show that $-(x + y) = -x - y$.

A nonempty subset $H \subseteq G$ is a *subgroup* of G if H , together with the restriction of $+$ to H , is a group. Equivalently, if for every $a, b \in H$ also $-a \in H$ and $a + b \in H$. Note that this implies $0 \in H$.

Example. $(\mathbb{R}, +)$, $(\mathbb{Q}, +)$, or $(\mathbb{Z}, +)$ are commutative groups, while $(\mathbb{N}, +)$ is not.

Definition A.2. If H is a subgroup of G and $x \in G$, then the set

$$x + H := \{x + h \mid h \in H\}$$

is called a *coset* of H in G .

Lemma A.3. If H is a subgroup of G , then the relation defined by

$$x \sim y \iff x - y \in H$$

is an equivalence on G . The equivalence classes of \sim are exactly the cosets of H in G . This means that G is partitioned into disjoint shifted copies of H .

Proof. We first verify that it is an equivalence. It is reflexive since $x - x = 0 \in H$. Symmetric, since if $x - y \in H$, then $y - x = -(x - y) \in H$. And transitive since if $x - y \in H$ and $y - z \in H$, then $(x - y) + (y - z) = x - z \in H$.

Denote by $[x]$ the equivalence class of x ; we will show that $[x] = x + H$. If $z \in x + H$, then $z = x + h$ for some $h \in H$ and $z - x = h$, so $x \sim z$. If $y \sim x$, then $y - x \in H$ and $y - x = h$ for some $h \in H$, so $y = x + h$ and $y \in x + H$. \square

Remark. This statement is true for non-commutative groups as well, but one needs to use the equivalence $x \sim y \iff -x + y \in H$ and must be a bit more careful. For example, $-(x + y) = -y + (-x)$, which is not the same as $-x - y$ if G is not commutative.

A.2 Measure Theory

The goal of measure theory is to generalize the notion of “length” (or volume) to a broad class of subsets of \mathbb{R} . Ideally, we seek a map $\lambda : \mathcal{P}(\mathbb{R}) \rightarrow [0, \infty]$ that assigns a length to every set, such that $\lambda([a, b]) = b - a$, and the measure of a disjoint union is the sum of the measures. It will turn out that this is impossible to do for *every* set — assuming we work in ZFC. Therefore, we will restrict our attention only to *measurable* sets.

The following text closely follows [23].

The outer Lebesgue measure We first define the *outer Lebesgue measure* λ^* for *any* set $A \subseteq \mathbb{R}$ by covering it with countably many open intervals.

Definition A.4. The *outer Lebesgue measure* of a set $A \subseteq \mathbb{R}$ is defined as

$$\lambda^*(A) := \inf \left\{ \sum_{n=1}^{\infty} (b_n - a_n) \mid A \subseteq \bigcup_{n=1}^{\infty} (a_n, b_n) \right\}$$

Note that the open intervals (a_n, b_n) can be empty (when $a_n = b_n$), and that $\lambda^*(A)$ can be ∞ .

It is easy to see that λ^* is *monotone*: if $A \subseteq B$, then $\lambda^*(A) \leq \lambda^*(B)$. We will show that λ^* also has some other good properties.

Lemma A.5. *If I is an interval, then $\lambda^*(I) = \ell(I)$, where $\ell(I)$ denotes the length of the interval I .*

Proof. We will prove it for closed intervals $I = [a, b]$; the other cases are similar.

For every real $\varepsilon > 0$, we have $I \subseteq (a - \varepsilon, b + \varepsilon)$, and thus

$$\lambda^*([a, b]) \leq b - a + 2\varepsilon.$$

The value of ε can be chosen arbitrarily small, and so $\lambda^*([a, b]) \leq b - a$.

To prove the other inequality, we need to show that whenever $\langle I_i \mid i \in \omega \rangle$ is a collection of open intervals covering the interval $[a, b]$, then $\sum \ell(I_i) \geq b - a$. The famous *Heine–Borel theorem* implies that there exists a finite subcollection of the intervals that also covers $[a, b]$. By deleting some intervals, $\sum \ell(I_i)$ can only decrease, so WLOG assume that only finitely many intervals are nonempty.

One of the intervals I_i contains the point a ; let us call it (a_1, b_1) . If $b_1 < b$, then some interval, say (a_2, b_2) , contains b_1 . We continue in a similar fashion until we arrive at an interval (a_k, b_k) containing b . We have

$$\sum \ell(I_i) \geq (b_1 - a_1) + (b_2 - b_1) + \cdots + (b_k - b_{k-1}) = b_k - a_1 > b - a. \quad \square$$

Exercise A.6. Show that if $A \subseteq \mathbb{R}$ is countable, then $\lambda^*(A) = 0$.

Remark. This shows, in particular, that \mathbb{Q} has outer measure zero. This illustrates that the result proved above, $\lambda^*(I) = \ell(I)$, has to rely on some property that differentiates real and rational numbers. The proof presented above used the Heine–Borel theorem, which relies on the fact that the real numbers are a complete metric space (unlike the rationals).

The failure of additivity Another property we would like our measure to have is *additivity*. Given a system of subsets $\langle A_i \mid i \in I \rangle$ of \mathbb{R} , we would like

$$\lambda^*\left(\bigcup_{i \in I} A_i\right) = \sum_{i \in I} \lambda^*(A_i) \quad (1)$$

to hold in as many cases as possible. At a minimum, we require *finite additivity* ((1) holds for all finite I), and *countable additivity* ((1) holds for all countable I) would be greatly appreciated.

Lemma A.7 (AC_ω). *If $\langle A_i \mid i \in \omega \rangle$ is a countable system of subsets of \mathbb{R} , then*

$$\lambda^*\left(\bigcup_{i < \omega} A_i\right) \leq \sum_{i < \omega} \lambda^*(A_i).$$

This property is called countable subadditivity of λ^ .*

Note that this lemma immediately solves the previous exercise.

Proof. For every set A_i , take a sufficiently efficient cover. More precisely, *choose* using AC_ω a cover $\langle I_j^{(i)} \mid j < \omega \rangle$ such that

$$\sum_{j < \omega} \ell(I_j^{(i)}) \leq \lambda^*(A_i) + \frac{\varepsilon}{2^{i+1}}.$$

Since we are assuming AC_ω , a countable union of countable covers is a countable cover, and we have

$$\lambda^*\left(\bigcup A_i\right) \leq \sum_{i, j < \omega} \ell(I_j^{(i)}) \leq \sum_{i < \omega} \lambda^*(A_i) + \varepsilon.$$

This is true for every $\varepsilon > 0$. □

Exercise A.8. Show that if λ^* is finitely additive, then it also must be countably additive.

Proposition A.9. *In ZFC, λ^* is not countably additive.*

Corollary A.10. *In ZFC, λ^* is not finitely additive.*

Proof of proposition. We will find countably many disjoint subsets $A_1, A_2, \dots \subseteq \mathbb{R}$ such that $\lambda^*(\bigcup A_i) \neq \sum \lambda^*(A_i)$. Define a relation \sim on \mathbb{R} as

$$x \sim y \iff x - y \in \mathbb{Q}.$$

Because \mathbb{Q} is a subgroup of \mathbb{R} , Lemma A.3 implies that \sim is an equivalence, and that its equivalence classes partition \mathbb{R} into shifted disjoint copies of \mathbb{Q} . Using the axiom of choice, select a single element contained in the interval $[0, 1]$ from each equivalence class of \sim , and collect them into a set $V \subset [0, 1]$. Such a set is called a *Vitali set*.

Since \mathbb{Q} is countable, we can enumerate all the rational numbers in the interval $[-1, 1]$ in a sequence q_0, q_1, q_2, \dots . Let $A_i := V + q_i$ be the translation of V by

q_i for each $i < \omega$. The sets A_i are clearly disjoint and contained in $[-1, 2]$, and some thought reveals that they together cover $[0, 1]$. Hence, from monotonicity

$$\lambda^*([0, 1]) = 1 \leq \lambda\left(\bigcup_{i < \omega} A_i\right) \leq \lambda^*([-1, 2]) = 3.$$

Notice that λ^* is by definition translation-invariant, so $\lambda^*(A_i) = \lambda^*(V)$ for every i . Thus, if $\lambda^*(0) = 0$, then $\sum \lambda^*(A_i) = 0$, and if $\lambda^*(V) > 0$, then $\sum \lambda^*(A_i) = \infty$. Even without knowing which of these possibilities actually holds, we can say for sure that $\sum \lambda^*(A_i)$ cannot be between 1 and 3. \square

The Lebesgue measure The second step in the construction of the Lebesgue measure is defining a suitable system \mathcal{E} of subsets of \mathbb{R} such that the outer measure λ^* restricted to \mathcal{E} becomes countably additive, and at the same time, \mathcal{E} is as rich as possible.

The *complement* of a set E is the set $E^c := \mathbb{R} \setminus E$.

Definition A.11. A set $E \subseteq \mathbb{R}$ is *measurable* if for every $A \subseteq \mathbb{R}$ we have

$$\lambda^*(A) = \lambda^*(A \cap E) + \lambda^*(A \cap E^c). \quad (2)$$

Let \mathcal{E} be the set of all measurable subsets of \mathbb{R} . The *Lebesgue measure* λ on \mathbb{R} is the restriction of λ^* to \mathcal{E} .

Observation A.12. If E is measurable, then E^c is also measurable.

To check the measurability of a set E , it suffices to verify the inequality “ \geq ” in (2), since “ \leq ” is implied by the subadditivity of λ^* .

Fact A.13. It can be shown that measurable sets can be approximated from the inside by compact sets. That is, for every $\varepsilon > 0$ and every measurable set E of finite measure, there exists a compact set $K \subseteq E$ such that $\lambda^*(E \setminus K) < \varepsilon$.

Observation A.14. If $\lambda^*(E) = 0$, then E is measurable.

Proof. Let $A \subseteq \mathbb{R}$ be arbitrary. Since $A \cap E \subseteq E$, we have $\lambda^*(A \cap E) \leq \lambda^*(E) = 0$. Similarly, $\lambda^*(A) \geq \lambda^*(A \cap E^c)$. Combining these, we get

$$\lambda^*(A) \geq \lambda^*(A \cap E^c) = \lambda^*(A \cap E^c) + \lambda^*(A \cap E). \quad \square$$

Exercise A.15. Show that the interval (a, ∞) is measurable for every $a \in \mathbb{R}$.

Lemma A.16. If A and B are measurable, then $A \cup B$ is also measurable. By induction, all finite unions of measurable sets are measurable.

Proof. Let X be the test set from the definition of measurability. We claim

$$\begin{aligned} \lambda^*(X) &\geq \lambda^*(X \cap (A \cup B)) + \lambda^*(X \cap (A \cup B)^c) \\ &= \lambda^*((X \cap A) \cup (X \cap A^c \cap B)) + \lambda^*(X \cap A^c \cap B^c). \end{aligned}$$

To get the measure of $X \cap A^c \cap B^c$, we use the measurability of B to split $X \cap A^c$:

$$\lambda^*(X \cap A^c) = \lambda^*(X \cap A^c \cap B) + \lambda^*(X \cap A^c \cap B^c).$$

Combining this with what we obtain by splitting X by A and using subadditivity:

$$\begin{aligned}\lambda^*(X) &= \lambda^*(X \cap A) + \lambda^*(X \cap A^c) \\ &= \lambda^*(X \cap A) + \lambda^*(X \cap A^c \cap B) + \lambda^*(X \cap A^c \cap B^c) \\ &\geq \lambda^*((X \cap A) \cup (X \cap A^c \cap B)) + \lambda^*(X \cap A^c \cap B^c).\end{aligned}$$

Which is what we wanted to show. \square

Corollary A.17. *If A and B are measurable, then $A \cap B$ and $A \setminus B$ are also measurable.*

Proof. $A \cap B = (A^c \cup B^c)^c$ and $A \setminus B = A \cap B^c$. \square

Exercise A.18. Show that if B_0, B_1, \dots, B_{n-1} are disjoint measurable sets, then for any test set X :

$$\lambda^*\left(X \cap \bigcup_{i < n} B_i\right) = \sum_{i < n} \lambda^*(X \cap B_i).$$

Do this by induction from the definition of a measurable set.

Proposition A.19. *The union of countably many measurable sets is measurable.*

Proof. Let $\langle E_i \mid i < \omega \rangle$ be measurable sets; we want to show that $E := \bigcup E_i$ is measurable. We will first construct a sequence of disjoint sets B_i such that $\bigcup B_i = \bigcup E_i$. We do this by letting $B_0 = E_0$, $B_1 = E_1 \setminus E_0$, $B_2 = E_2 \setminus (E_0 \cup E_1)$, and so on. Note that all B_i are measurable from what we showed earlier. Denote $F_n := \bigcup_{i < n} B_i$. This is a finite union, so F_n is also measurable. Thus, for any test set X , we have

$$\lambda^*(X) = \lambda^*(X \cap F_n) + \lambda^*(X \cap F_n^c) = \sum_{i < n} \lambda^*(X \cap B_i) + \lambda^*(X \cap F_n^c),$$

where we used the result of the previous exercise. Since $F_n \subseteq E$, we have $E^c \subseteq F_n^c$. Therefore, $\lambda^*(X \cap F_n^c) \geq \lambda^*(X \cap E^c)$ and

$$\lambda^*(X) \geq \sum_{i < n} \lambda^*(X \cap B_i) + \lambda^*(X \cap E^c)$$

holds for any n . Hence, it also holds when $n \rightarrow \infty$ and

$$\lambda^*(X) \geq \sum_{i < \omega} \lambda^*(X \cap B_i) + \lambda^*(X \cap E^c).$$

We finish this by using the countable subadditivity of the outer measure:

$$\sum_{i < \omega} \lambda^*(X \cap B_i) \geq \lambda^*\left(\bigcup_{i < \omega} (X \cap B_i)\right) = \lambda^*(X \cap E).$$

We have shown $\lambda^*(X) \geq \lambda^*(X \cap E) + \lambda^*(X \cap E^c)$ for any test set X . \square

Corollary A.20. *The intersection of countably many measurable sets is measurable.*

Proof. This follows from the fact that the intersection of a group of sets is the complement of the union of the complements of the sets. \square

Theorem A.21. *The Lebesgue measure λ is countably additive.*

Proof. Let $\langle E_i \mid i < \omega \rangle$ be a system of disjoint measurable sets. We want to show that $\lambda(\bigcup E_i) = \sum \lambda(E_i)$. We already know “ \leq ” from subadditivity; we need to show “ \geq .” Take $A := E_1 \cup E_2$ as the “testing” set in the definition of a measurable set. Since E_1 is measurable, the definition tells us that

$$\lambda(E_1 \cup E_2) = \lambda(E_1) + \lambda(E_2).$$

By induction, this can be extended to any finite collection of disjoint measurable sets. Using this and monotonicity, we have

$$\lambda\left(\bigcup_{i < \omega} E_i\right) \geq \lambda\left(\bigcup_{i < n} E_i\right) = \sum_{i < n} \lambda(E_i)$$

for every $n \in \omega$, yielding $\lambda(\bigcup_{i < \omega} E_i) \geq \sum_{i < \omega} \lambda(E_i)$. □

Exercise A.22. Show that if E_1 and E_2 are measurable and have finite measure (so $\lambda(E_1 \cup E_2) < \infty$), then

$$\lambda(E_1 \cup E_2) = \lambda(E_1) + \lambda(E_2) - \lambda(E_1 \cap E_2).$$

More generally, show that the inclusion-exclusion principle holds. If $\langle E_i \mid i \in I \rangle$ is a finite collection of measurable sets with finite measure, then

$$\lambda\left(\bigcup_{i \in I} E_i\right) = \sum_{\emptyset \neq J \subseteq I} (-1)^{|J|+1} \lambda\left(\bigcap_{j \in J} E_j\right).$$

Borel sets Set systems that are closed under complements and countable unions (like the system of all measurable sets \mathcal{E}) are of fundamental importance in measure theory, and they have a name:

Definition A.23. A σ -algebra is a nonempty set system closed under complements and countable unions. And thus, under countable intersections as well.

It is easy to see that if \mathcal{A} is a system of subsets of a set X , then the intersection \mathcal{F} of all σ -algebras containing \mathcal{A} is also a σ -algebra. We say that \mathcal{F} is the *smallest σ -algebra containing \mathcal{A}* . We also say that \mathcal{F} is the σ -algebra *generated by \mathcal{A}* .

Definition A.24. *Borel sets* are the elements of the σ -algebra \mathcal{B} generated by the system \mathcal{A} of all open intervals of \mathbb{R} . We denote

Remark. It does not matter whether “open intervals” means only finite or also infinite open intervals (a, ∞) , since (a, ∞) is a countable union of finite open intervals.

Observation A.25. *All Borel sets are measurable.*

Proof. If we show that all open intervals are measurable, then the σ -algebra \mathcal{E} contains all open intervals, and by definition $\mathcal{B} \subseteq \mathcal{E}$. Exercise A.15 shows that the interval $(a, \infty) \in \mathcal{E}$ is measurable. Because sets of measure zero are measurable, the set $\{a\}$ is measurable, and thus $\{a\} \cup (a, \infty) = [a, \infty)$ is also measurable. Its complement, $(-\infty, a)$, is thus also measurable. Now it is easy to create every finite open interval (a, b) as $(-\infty, b) \cap (a, \infty)$. □

Observation A.26. *All open and all closed sets are Borel sets.*

Proof. We will show that every open set is a union of countably many disjoint open intervals. From this, every open set is a Borel set, and closed sets are the complements of open sets, so they are Borel sets as well.

Given a nonempty open set U , define for every point $x \in U$ the largest open interval contained in U that contains x :

$$I_x := \bigcup \{(a, b) \mid x \in (a, b) \subseteq U\}.$$

Since the set U is open, I_x exists for every x . We claim that the intervals I_x and I_y are either disjoint, or identical for every $x, y \in U$. Suppose $z \in I_x \cap I_y$. Since I_x and I_y are open intervals overlapping at z , their union is also an open interval contained in U . Since I_x and I_y are the maximal intervals containing x and y , they must both contain the entire union $I_x \cup I_y$. Therefore $I_x = I_y$.

Therefore $\mathcal{I} := \{I_x \mid x \in U\}$ partitions U into a collection of disjoint intervals. Since \mathbb{Q} is dense in \mathbb{R} , we can pick a unique rational number q_I from each interval $I \in \mathcal{I}$. This defines an injection $\mathcal{I} \rightarrow \mathbb{Q}$, so \mathcal{I} must be countable. \square

Corollary A.27. *All countable subsets $A \subseteq \mathbb{R}$ are Borel sets.*

Proof. All singletons $\{x\}$ are Borel sets since they are closed, and each countable set A is a countable union of the singletons of its elements. \square

Example. The set of all rational numbers \mathbb{Q} is a Borel set since it is countable. Therefore, the set of all irrational numbers $\mathbb{R} \setminus \mathbb{Q}$ is also a Borel set.

This illustrates an important concept. To construct the irrational numbers, we took a countable union of closed sets and then took its complement. Equivalently, we could have taken a countable intersection of all the open sets $\mathbb{R} \setminus \{q\}$ for rational q . We can iterate this, taking countable unions and countable intersections of the sets we have already constructed to create new sets. This gives rise to the *Borel hierarchy*.

Definition A.28. We define the *Borel hierarchy* for ordinals α as follows:

- (i) Σ_1^0 are all open sets,
- (ii) Π_1^0 are all closed sets,
- (iii) Σ_α^0 for $\alpha > 1$ are all countable unions of Π_δ^0 sets for $\delta < \alpha$,
- (iv) Π_α^0 for $\alpha > 1$ are all countable intersections of Σ_δ^0 sets for $\delta < \alpha$.

Observation A.29. *If $\alpha < \beta$, then $\Sigma_\alpha^0 \subseteq \Sigma_\beta^0$ and $\Pi_\alpha^0 \subseteq \Pi_\beta^0$.*

Observation A.30. *A set B is in Π_α^0 if and only if its complement B^c is in Σ_α^0 .*

Example. \mathbb{Q} is a countable union of closed sets, so $\mathbb{Q} \in \Pi_2^0$. The irrational numbers are a complement of \mathbb{Q} , so they are in Σ_2^0 .

The Borel hierarchy allows us to construct increasingly complex sets; virtually every set you could define is a Borel set. Does the hierarchy ever stop?

Fact A.31. *The Borel hierarchy stabilizes at the level $\alpha = \omega_1$. It can also be shown that the number of all Borel sets is $|\mathcal{B}| = |\mathbb{R}| = \mathfrak{c}$.*

Is every measurable set a Borel set? No, far from it. Every set of outer measure zero is measurable, and it can be shown that there exist uncountable sets of measure zero; namely, the Cantor set \mathcal{C} . The number of all subsets of \mathcal{C} is $2^{\mathfrak{c}} > \mathfrak{c}$. This means that almost all subsets of \mathcal{C} are just “random noise,” and they are not Borel sets.

Fact A.32. *The measurable sets \mathcal{E} are the smallest σ -algebra containing all Borel sets and all sets of outer measure zero.*

The density theorem It can be shown that measurable sets with positive measure behave like solid objects rather than fuzzy clouds. If you pick a random point from such a set and zoom in infinitely close, in almost all cases, the set will eventually fill 100% of your view, appearing completely solid. Conversely, if you zoom in on a point outside the set, it will disappear entirely, occupying 0%.

Definition A.33 (Density). Let $E \subseteq \mathbb{R}$ be measurable. For a point $x \in \mathbb{R}$, define

$$d_E(x) := \lim_{\delta \rightarrow 0} \frac{\lambda(B(x, \delta) \cap E)}{\lambda(B(x, \delta))},$$

where $B(x, \delta)$ denotes the open ball of radius δ centered at x . If the limit exists, then we call it the *density* of E at x .

Fact A.34 (Lebesgue density theorem). *If $E \subseteq \mathbb{R}$ is measurable, then the set of those $x \in E$ for which $d_E(x)$ is undefined or smaller than 1 has measure zero. In particular, if $\lambda(E) > 0$, then E has density 1 at almost all of its points.*

Corollary A.35. *If E is measurable and $\lambda(E) > 0$, then for any $0 \leq \varepsilon < 1$, there exists an interval I such that*

$$\varepsilon \leq \frac{\lambda(I \cap E)}{\lambda(I)} \leq 1.$$

Exercise A.36. Show that this is true even if E is not measurable. Use the definition of the outer Lebesgue measure λ^* and countable subadditivity.

Lebesgue measure in higher dimensions For \mathbb{R}^n instead of \mathbb{R} , intervals in the definition of λ^* are replaced with open “boxes” (Cartesian products of open intervals). Then it can be shown that the appropriate analogs of the statements above that concern λ^* hold in \mathbb{R}^n as well. In dimension 3 (and higher), the lack of finite additivity of λ^* in ZFC manifests itself in a particularly bizarre way. In Section 7.4, we show that a 3-dimensional ball can be decomposed into finitely many disjoint subsets, which can be rearranged (by rotating and translating them) into two perfect copies of the original ball.

In the definition of a measurable set in \mathbb{R}^n , the “test” sets are $A \subseteq \mathbb{R}^n$. Borel sets in \mathbb{R}^n can be defined as the σ -algebra generated by all open subsets of \mathbb{R}^n . The density theorem holds in \mathbb{R}^n as well.

Bibliography

- [1] Petr Balcar Bohuslav a Štěpánek. *Teorie množin*. Vydání 2., opravené a rozšířené. Praha: Academia, 2001. ISBN: 80-200-0470-X.
- [2] James E. Baumgartner. “Ineffability properties of cardinals I”. In: *Infinite and Finite Sets*. Ed. by A. Hajnal, R. Rado, and V. T. Sós. Vol. 10. Colloquia Mathematica Societatis János Bolyai. Dedicated to P. Erdős on his 60th birthday. Amsterdam: North-Holland, 1975, pp. 109–130.
- [3] Stephen Budiansky. *Journey to the Edge of Reason: The Life of Kurt Gödel*. W. W. Norton, 2021. ISBN: 9781324005452.
- [4] Boris Bukh. *Walk through Combinatorics: Compactness principle*. 2013. URL: http://www.borisbukh.org/DiscreteMath12/notes_compactness.pdf (visited on 02/07/2026).
- [5] Jakub Bulín. *NAIL062 Propositional and Predicate Logic: Lecture Notes*. 2025. URL: <https://github.com/jbulin-mff-uk/nail062/raw/main/lecture/lecture-notes/lecture-notes.pdf> (visited on 11/06/2025).
- [6] Jochen Burghardt. *Matchstick representation of ordinal numbers up to ω^ω* . 2023. URL: <https://commons.wikimedia.org/wiki/File:Omega-exp-omega-normal.pdf> (visited on 11/04/2025).
- [7] Jochen Burghardt. *Order Type Examples*. 2019. URL: <https://commons.wikimedia.org/wiki/File:OrderTypeExamples.pdf> (visited on 11/04/2025).
- [8] Timothy Y Chow. “The consistency of arithmetic”. In: *The Mathematical Intelligencer* 41.1 (2019), pp. 22–30. URL: <https://arxiv.org/pdf/1807.05641>.
- [9] Walter Dean and Sean Walsh. “The prehistory of the subsystems of second-order arithmetic”. In: *The Review of Symbolic Logic* 10.2 (2017), pp. 357–396. URL: <https://arxiv.org/pdf/1612.06219>.
- [10] Herbert B Enderton. *Elements of set theory*. Gulf Professional Publishing, 1977, pp. 195–196.
- [11] Paul Erdős et al. *Combinatorial set theory: partition relations for cardinals*. Vol. 106. Elsevier, 2011.
- [12] Harvey M Friedman. “Subtle cardinals and linear orderings”. In: *Annals of Pure and Applied Logic* 107.1-3 (2001), pp. 1–34.
- [13] Jean H Gallier. “What’s so special about Kruskal’s theorem and the ordinal Γ_0 ? A survey of some results in proof theory”. In: *Annals of pure and applied logic* 53.3 (1991), pp. 199–260. URL: <https://www.cis.upenn.edu/~jean/kruskal.pdf>.
- [14] Gina Garcia Tarrach. “The Axiom of Choice and its implications in mathematics”. MA thesis. Universitat de Barcelona, 2017. URL: <https://hdl.handle.net/2445/121981>.

- [15] Kurt Gödel. “On Formally Undecidable Propositions of Principia Mathematica and Related Systems I”. In: *Kurt Gödel: Collected Works: Volume I Publications 1929-1936*. OUP Oxford, 1986, pp. 144–195. ISBN: 0-19-503964-5.
- [16] Reuben Louis Goodstein. “On the restricted ordinal theorem”. In: *The Journal of Symbolic Logic* 9.2 (1944), pp. 33–41. DOI: 10.2307/2268019.
- [17] Gro-Tsen and IkamusumeFan. *A graphical “matchstick” representation of the ordinal ω^2* . 2015. URL: https://commons.wikimedia.org/wiki/File:Ordinal_ww.svg (visited on 11/04/2025).
- [18] Joel David Hamkins. *Transfinite recursion as a fundamental principle in set theory*. Accessed: 2025-11-08. 2014. URL: <https://jdh.hamkins.org/transfinite-recursion-as-a-fundamental-principle-in-set-theory/>.
- [19] Karel Hrbáček and Tomáš Jech. *Introduction to set theory*. eng. Third edition, revised and expanded. Pure and applied mathematics. A series of monographs and textbooks ; 220. Boca Raton: Taylor & Francis, 1999. ISBN: 0-8247-7915-0.
- [20] Thomas Jech. *Set theory*. Springer, 2003.
- [21] Ronald Jensen and Kenneth Kunen. “Some Combinatorial Properties of L and V”. Handwritten notes, The Rockefeller University. Available at <https://www.math.uni-bonn.de/~raesch/jensen/> (accessed 15 February 2026). 1969.
- [22] Akihiro Kanamori. *The higher infinite: large cardinals in set theory from their beginnings*. Springer, 2003.
- [23] Ida Kantor, Jiří Matoušek, and Robert Šámal. *Mathematics++*. Vol. 75. American Mathematical Soc., 2015, pp. 6–15.
- [24] Asaf Karagila. *Lecture Notes: Large Cardinals*. 2025. URL: <https://karagila.org/files/LC-2025.pdf> (visited on 02/04/2026).
- [25] Laurie Kirby and Jeff Paris. “Accessible independence results for Peano arithmetic”. In: *Bulletin of the London Mathematical Society* 14.4 (1982), pp. 285–293. DOI: 10.1112/blms/14.4.285.
- [26] Alberto Marcone and Antonio Montalbán. “The Veblen functions for computability theorists”. In: *The Journal of symbolic logic* 76.2 (2011), pp. 575–602. URL: <https://arxiv.org/pdf/0910.5442>.
- [27] Jeff Paris and Leo Harrington. “A mathematical incompleteness in Peano arithmetic”. In: *Handbook of Mathematical Logic, edited by J. Barwise*. North-Holland, 1977, pp. 1133–1142.
- [28] Michael Rathjen. “The art of ordinal analysis”. In: *Proceedings of the International Congress of Mathematicians*. Vol. 2. European Mathematical Society. 2006, pp. 45–69.
- [29] Frederick Rowbottom. “Large Cardinals and Small Constructible Sets”. PhD thesis. University of Wisconsin–Madison, 1964.

- [30] Saharon Shelah and Alexander Soifer. “Axiom of choice and chromatic number of the plane”. In: *Journal of Combinatorial Theory, Series A* 103.2 (2003), pp. 387–391. ISSN: 0097-3165. DOI: [https://doi.org/10.1016/S0097-3165\(03\)00102-X](https://doi.org/10.1016/S0097-3165(03)00102-X). URL: <https://www.sciencedirect.com/science/article/pii/S009731650300102X>.
- [31] Jack Silver. “A large cardinal in the constructible universe”. In: *Fund. Math* 69 (1970), pp. 93–100.
- [32] Stephen George Simpson. *Subsystems of second order arithmetic*. Vol. 1. Cambridge University Press, 2009.
- [33] Will Sladek. *The termite and the tower: Goodstein sequences and provability in PA*. 2007. URL: <https://andrescaicedo.wordpress.com/wp-content/uploads/2017/09/sladekgoodstein.pdf>.
- [34] Peter Smith. *Tennenbaum’s Theorem*. (accessed 18 February 2026). 2011. URL: <https://www.logicmatters.net/resources/pdfs/TennenbaumTheorem.pdf>.
- [35] Robert M Solovay. “A model of set-theory in which every set of reals is Lebesgue measurable”. In: *Annals of Mathematics* 92.1 (1970), pp. 1–56.
- [36] Alan Mathison Turing et al. “On computable numbers, with an application to the Entscheidungsproblem”. In: *J. of Math* 58.345-363 (1936), p. 5.
- [37] Veritasium. *Math’s Fundamental Flaw*. 2021. URL: <https://www.youtube.com/watch?v=HeQX2HjkcNo> (visited on 11/06/2025).
- [38] Veritasium. *The Man Who Almost Broke Math (And Himself...) - Axiom of Choice*. 2025. URL: https://www.youtube.com/watch?v=_cr46G2K5Fo (visited on 11/08/2025).
- [39] Vsauce. *How To Count Past Infinity*. 2016. URL: <https://www.youtube.com/watch?v=SrU9YDoXE88> (visited on 11/04/2025).
- [40] Googology Wiki. *Veblen function*. 2025. URL: https://googology.fandom.com/wiki/Veblen_function (visited on 11/12/2025).
- [41] Wikipedia. *Hausdorff maximal principle — Wikipedia, The Free Encyclopedia*. 2025. URL: <https://en.wikipedia.org/w/index.php?title=Hausdorff%5C%20maximal%5C%20principle&oldid=1300391387> (visited on 10/07/2025).