1 All sets are measurable

The first part of these notes is an introduction to ordinal numbers, and by extension to set theory in general. Our main goals are as follows:

1. Introduce ordinal numbers by examples that arise quite naturally in mathematics (especially in analysis), define ordinal numbers and prove some of their properties.

2. Construct ω_1, the first uncountable ordinal as well as the smallest uncountable set.

3. Discuss why all *reasonable* sets *should* be measurable.

The second part is a crash course in mathematical logic, starting with naive set theory, propositional logic and end with the statements of Gödel's incompleteness theorems.

1.1 Ordinals and well-orders

1.1.1 Borel hierarchy

The Borel algebra is the smallest σ-algebra containing all open subsets of $\mathbb{R}$; there exists a smallest one because we can intersect all such σ-algebras. That's not a terribly concrete description. Let us try to build Borel sets in a "bottom-up" manner.

$\Gamma_0 := \{A \subseteq \mathbb{R} \mid A \text{ is either open or closed}\}$

and inductively,

$$\Gamma_{n+1} := \{\bigcup_{i=0}^{\infty} A_i \mid A_i \in \Gamma_n\} \cup \{\bigcap_{i=0}^{\infty} A_i \mid A_i \in \Gamma_n\}$$

By induction one can show that Γ_n is closed under complements. Is $\Gamma_\omega := \bigcup_{n=0}^{\infty} \Gamma_n$ the Borel algebra? Not quite, since if $B = \bigcup_{n=0}^{\infty} A_n$ where $A_n \in \Gamma_n$, there is no reason why B must belong to any of the Γ_n. Indeed it may not, so we have to start over again.

$$\Gamma_{\omega+1} := \{\bigcup_{i=0}^{\infty} A_i \mid A_i \in \Gamma_\omega\} \cup \{\bigcap_{i=0}^{\infty} A_i \mid A_i \in \Gamma_\omega\}$$

and then

$\Gamma_{\omega+2}, \Gamma_{\omega+3}, \ldots \Gamma_{\omega+\omega}, \Gamma_{\omega+\omega+\omega}, \ldots \Gamma_{\omega\cdot\omega}, \Gamma_{\omega\cdot\omega\cdot\omega}, \ldots \Gamma_{\omega^\omega}, \Gamma_{\omega^{\omega^\omega}}, \ldots$

and so on and so forth. When shall we stop? Turns out it is necessary and sufficient to iterate through all the "countable indices". More precisely, the transfinite sequence $(\Gamma_\alpha)_{\alpha<\omega_1}$ is strictly increasing, where ω_1 is the first uncountable ordinal, and $\bigcup_{\alpha<\omega_1} \Gamma_\alpha$ is exactly the Borel algebra; towards the end the meaning of these words will become clearer.

The above is analogous to how we can construct the subgroup of G generated by a set S in two ways: either intersect all subgroups containing S, or consider all elements of form $g_1 g_2^{-1} g_3 g_4^{-1} \cdots g_{2n-1} g_{2n}^{-1}$, where $g_1, \ldots, g_{2n} \in S$; the latter is basically the same as defining a sequence $S_0 = S$, $S_{n+1} = \{gh^{-1} : g, h \in S_n\}$, and consider the union $\bigcup_n S_n$. The difference between group and σ-algebra is that group operations are binary, so we only need to iterate countably many times; the operations on σ-algebra, i.e., countable union and intersection, are infinitary, and it turns out we need to iterate ω_1 times.

The transfinite viewpoint is useful not just for describing the Borel sets, but also for proving some of their properties. A common method to prove a property about Borel sets is to show that open sets have that property, and that the collection of all sets with the property is a σ-algebra. One can instead prove by transfinite induction on the countable indices: show that if it holds for all sets in previous stages then it holds for the next stage Γ_α. This idea can be used to prove, e.g., the Monotone class theorem and the π-λ theorem.

1.1.2 Pre-measure

One standard way to construct the Lebesgue measure is to first prove that the pre-measure on left-closed, right-open intervals is indeed a pre-measure, i.e., countably additive, and then use Carathéodory's extension theorem. So one needs to prove that if $[a,b) = \bigcup_{i=0}^{\infty} I_i$ where $I_i = [a_i, b_i)$ and $I_i \cap I_j = \varnothing$ for $i \neq j$, then $b - a = \sum_{i=0}^{\infty}(b_i - a_i)$.

This can be proven using compactness as in, e.g., Folland, but let's prove it using transfinite induction instead. We need to understand all the possible ways to partition $[a,b)$ into disjoint left-closed right-open intervals; we ignore the easy case where there are finitely many intervals. The left end-point a must be the left end-point of a unique interval, say $[a, a_1)$. Now a_1 belongs to a unique interval, and in fact a_1 must be its left end-point, since otherwise it would overlap with $[a, a_1)$. Say this interval is $[a_1, a_2)$; then a_2 is the left end-point of a unique interval $[a_2, a_3)$, etc. Let $a_\omega = \lim_n a_n$. If $a_\omega = b$, then the intervals $[a_i, a_{i+1})$ (where $a_0 = a$) are all the intervals in the partition, so what we want boils down to $b - a = \sum_{i=0}^{\infty}(a_{i+1} - a_i)$, which clearly holds. If $a_\omega < b$, then a_ω must be the left end-point of a unique interval, which we call $[a_\omega, a_{\omega+1})$. This process potentially can continue forever: $a_{\omega+2}, a_{\omega+3}, \ldots a_{\omega+\omega}, a_{\omega+\omega+\omega}, \ldots a_{\omega\cdot\omega}, a_{\omega\cdot\omega\cdot\omega}, \ldots a_{\omega^\omega}, a_{\omega^{\omega^\omega}}, \ldots$

But it turns out that, unlike the example of σ-algebra, this cannot continue all the way up to the first uncountable ordinal ω_1: it has to stop at some countable stage α. This is due to the second-countability of $\mathbb{R}$. Now the idea is to prove the following statement by transfinite induction: if $[a,b)$ is the disjoint union of $[a_i, a_{i+1}),\ i < \alpha$ where α is some countable ordinal, then $b - a = \sum_{i<\alpha}(a_{i+1} - a_i)$; recall that the sum of a countable set of positive numbers is always well-defined. The induction at "successor stages", e.g., when α is something like $\omega + 1$ or $\omega + 2$, is easy. At "limit stages" such as $\omega + \omega$ or ω^ω, we use that fact that α is the limit of an ω-sequence of smaller ordinals. For example, $\omega + \omega = \lim_{n<\omega}(\omega + n)$ and $\omega^\omega = \lim_{n<\omega} \omega^n$. This reduces the limit case to the case $\alpha = \omega$, which was handled at the beginning, and thus finishes the inductive proof.

1.1.3 Cantor-Bendixson derivatives

This was how Cantor originally discovered ordinals. First he proved the following uniqueness theorem for trigonometric series, building on work of Riemann and others:

Theorem (Cantor). *Let $(c_n)_{n\in\mathbb{Z}} \subseteq \mathbb{C}$. If for all $x \in [0, 2\pi]$, $\displaystyle\sum_{n=-\infty}^{\infty} c_n e^{inx}$ converges to zero, then $c_n = 0$ for all n.*

Note that the situation is different from that in Fourier analysis, where the uniqueness of coefficients is automatic; there exist everywhere converging trigonometric series that are not Fourier series, such as $\sum_{n\geq 2}\frac{\sin nx}{\log n}$.

Then Cantor generalized the theorem slightly.

Theorem (Cantor)**.** *The uniqueness theorem is still true if we only assume convergence for $x \in [0, 2\pi] \setminus A$, where A is finite.*

After that Cantor generalized his theorem again. Recall that for $X \subseteq \mathbb{R}$, X' is the set of limit points of X; a point a is a limit point of X if any neighborhood of a contains some point of X other than a itself (and thus contains infinitely many points of X); equivalently a is in the closure of X and is not an isolated point of X. For example, if $X = \{-1, -\frac{1}{2}, -\frac{1}{3}, -\frac{1}{4}, \dots\}$ then $X' = \{0\}$ and $(X')' = \varnothing$. A set X is called *perfect* if $X' = X$. Examples of nonempty perfect sets include any closed interval, and the middle third Cantor set.

Note that a subset $A \subseteq [0, 2\pi]$ is finite iff $A' = \varnothing$ (by Bolzano-Weierstrass theorem). Let $A^{(0)} = A$ and inductively $A^{(n+1)} = (A^{(n)})'$.

Theorem (Cantor)**.** *The uniqueness theorem is still true if we only assume convergence for $x \in [0, 2\pi] \setminus A$, where $A^{(n)} = \varnothing$ for some n.*

A' is always closed, and if A is closed then $A' \subseteq A$. Therefore $A^{(1)}, A^{(2)}, A^{(3)}, \dots$ is a decreasing sequence of closed sets. It is natural to ask whether the sequence stabilizes, i.e., all the $A^{(n)}$ are the same from some point on. It may not, so we define $A^{(\omega)} = \bigcap_{n=1}^{\infty} A^{(n)}$. This may or may not be empty, and if it's nonempty we can keep going: $A^{(\omega+1)} = (A^{(\omega)})'$, then $A^{(\omega+2)}, A^{(\omega+3)}, \dots A^{(\omega+\omega)}$ and so on and so forth.

The following theorem was published by Lebesgue, but presumably Cantor already knew it, because after publishing the previous theorem he turned to the study of ordinal numbers.

Theorem (Cantor?-Lebesgue)**.** *The uniqueness theorem is still true if we only assume convergence for $x \in [0, 2\pi] \setminus A$, where $A^{(\alpha)} = \varnothing$ for some α.*

Similar to the previous example, one can show that the sequence $A^{(\alpha)}$ stabilizes at some countable step using second-countability of $\mathbb{R}$. Therefore if $A^{(\alpha)}$ is eventually empty then it was already empty for some countable α. Also, one can show that if $A^{(\alpha)} = \varnothing$ for some α then A must be countable. Thus the above theorem is actually subsumed by a later result, which says in fact any countable set is a "set of uniqueness". For more on sets of uniqueness, see [1] or [2].

1.1.4 Well-orders

The indices appearing in the above examples are called ordinal numbers, or ordinals. They show up whenever we want to "count past infinity". The first several ordinal numbers are listed below:

$$0, 1, 2, 3, \dots \omega, \omega+1, \omega+2, \dots \omega\cdot 2, \omega\cdot 3, \dots \omega^2, \omega^3, \dots \omega^\omega, \omega^{\omega^\omega}, \omega^{\omega^{\omega^\omega}}, \dots$$

By law the least ordinal is 0, and then $1, 2, 3, \dots$ The next one is ω, the first infinite ordinal, which means there are infinitely many ordinals before it. The one after ω is called $\omega + 1$, and then $\omega+2, \omega+3, \dots$ The next one is $\omega+\omega$, which is also denoted $\omega \cdot 2$ since that's cool. Of course we then have $\omega \cdot 3, \omega \cdot 4, \dots$ The next one is $\omega \cdot \omega$, a.k.a. ω^2. Needless to say we have $\omega^3, \omega^4, \dots, \omega^\omega, \omega^{\omega^\omega}, \omega^{\omega^{\omega^\omega}}, \dots$ The next one is denoted $\omega^{\omega^{\omega^{\cdot^{\cdot^{\cdot}}}}}$, or also ε_0, possibly because for every $\varepsilon > 0$ there is an ω in the tower whose font size is smaller than ε.

All the ordinals we have seen are countable (except ω_1), in the sense that there are only countably many ordinals before each of them; do not confuse the ω^ω here with the set of all functions from ω to ω, which is uncountable; ω^ω is countable because it is the countable limit of the countable ordinals $\omega, \omega^2, \omega^3, \dots$

Does there exist uncountable ordinal? To answer that, of course, we need to first rigorously define ordinals. So far they are just funny-looking indices, without any structure, but it should be at least intuitively clear that ordinals are linearly ordered. Therefore each ordinal α also has a linear order structure, namely we can identify α with the set of ordinals smaller than α, equipped with the natural linear order. Here comes the fun part: we can take this idea of identification seriously and (more or less) define an ordinal to be the set of smaller ordinals. Thus the empty set $\varnothing$ is the least ordinal, which we denote 0; the set $\{0\}$ whose only element is the empty set is the next ordinal, which we denote 1. In general, $n = \{0, 1, 2, \dots n-1\}$. The first infinite ordinal is $\omega = \{0, 1, 2, \dots\}$, and then we have $\omega + 1 = \{0, 1, 2, \dots \omega\}$, etc.

This definition of ordinals, in particular natural numbers, may feel bizarre at the first glance, but it's really *the* canonical way to define natural numbers. Think about it: the empty set is the only set that has zero element (and empty set is unique by the so-called Axiom of Extensionality), so it's just more than natural to define 0 to be $\varnothing$. Now we would like to define 1 to be a one-element set, and $\{0\}$ is a pretty natural choice, and so on and so forth.

An disadvantage of defining an ordinal to be the set of smaller ordinals is that, to make it fully rigorous, we run into some technicalities and axiomatic issues. Rather than the set-theoretic nature of ordinals, it's often the linear order on them, an example of *well-order*, that is important. So let's try to define well-order instead. We have seen that it always makes sense to talk about the "next" ordnal after a bunch of ordinals. For example, the next ordinal after 100 is 101, and the next ordinal after $0, 1, 2, \dots$ is ω. This leads us to the following definition.

"Definition". A well-order is a linear order such that after a bunch of elements there is always the next one.

Obviously no standard text ever uses this definition, but personally I think it is the closest to the intuition of ordinals, and it's not super difficult to make rigorous. For that we need some terminology. Suppose $(X, <)$ is a (strict) linear order, $a, b, c \in X$ and $A \subseteq X$.

- A is an *initial segment* if whenever $a \in A$ and $b \leq a$, then $b \in A$. It is *proper* if $A \neq X$. The empty set $\varnothing$ counts as an initial segment, and is proper unless X itself is empty.
- b is the *immediate successor* of a if $a < b$ and there is no c such that $a < c < b$.
- c is an *upper bound* of A if $c \geq a$ for all $a \in A$; if such a c exists then A is called *bounded*; it might seem more reasonable to call it "bounded above", but in our context everything is bounded below. Recall that c is called the *least upper bound* (or lub) if $c \leq d$ for any upper bound d, and that lub is unique if it exist.

- c is a *strict upper bound* of A if $c > a$ for all $a \in A$, in which case A is *strictly bounded.* We call c the *least strict upper bound* (or lsub) if $c \leq d$ for any strict upper bound d. Similar to lub, we have uniqueness of lsub in case it exists.

Note that if $A \subseteq X$ has a greatest element a, then the lsub of A is the same as the immediate successor of a; if A does not have greatest element, then lsub is the same as lub. So lsub is really just a shorthand for "the immediate successor of the greatest element if there is one, otherwise the lub". Also, if $A = \varnothing$, then any $a \in X$ is a strict upper bound of A, so in this case "A has lsub" means X has a smallest element.

Definition 1. A *well-order* is a linear order $(X, <)$ that satisfies any of the following equivalent definitions:

(i) Any proper initial segment has a lsub.

(ii) Any strictly bounded subset has a lsub.

(iii) Any nonempty subset has a least element.

(iv) There does not exist an infinite descending chain $a_0 > a_1 > a_2 > \cdots$.

(v) Any proper initial segment is of the form $X_a := \{x \mid x < a\}$ for some $a \in X$; this is called the initial segment below a/determined by a.

Personally my favorite definition is (i), since it literally says "after a bunch of elements there is always the next one", but (iii) is the standard definition seen in textbooks, as well as the most useful one. We will only need (iii) and occasionally the implication (iii)⇒(iv)&(v) in what follows. (iv) is perhaps the easiest one to visualize, but it is not used as the standard definition because the implication (iv)⇒(iii) requires Axiom of Choice. Note that the empty order counts as a well-order, since it trivially satisfies all the conditions, for example it does not have nonempty subset.

An ordinal can be thought of as an isomorphism class (or "order type") of well-orders, although we cannot literally say this due to Russell's paradox. Once we formalize "an ordinal is the set of smaller ordinals", we will see that an ordinal is a well-order, and moreover a canonical representative of its isomrphism class. The technical definition of ordinals is not so relevant to our main goal, the construction of ω_1, so we postpone it and mainly talk about well-orders.

Proof of equivalence. (iii)⇒(ii): If A is strictly bounded, then the set of strict upper bounds is nonempty, so by (iii) there must be a least one.

(ii)⇒(i): A proper initial segment is strictly bounded, since anything in the complement is a strict upper bound.

(i)⇒(iii): Let A be a nonempty subset, and B be the (possibly empty) set of strict lower bounds, i.e., $B := \{b \in X \mid \forall a \in A \; b < a\}$. Note that B is an initial segment, and it's proper since A is nonempty, so B has a lsub c. We claim that c is the least element of A. For any $a \in A$, we cannot have $a < c$, otherwise since c is the lsub of B, a is not a strict upper bound of B, contradicting the definition of B. On the other hand since $c \notin B$, there exists $a_0 \in A$ such that $a_0 \leq c$. Thus $a_0 = c$ and it is the least element of A.

(iii)⇒(iv): If there were an infinite descending chain $a_0 > a_1 > a_2 > \cdots$, then $\{a_i\}_i$ would be a nonempty subset without least element.

(iv)⇒(iii): If there were a nonempty subset A without least element, then we could build an infinite descending chain as follows: pick any $a_0 \in A$; it is not the least one in A so pick any $a_1 < a_0$, and then a_2, a_3, etc. We used Axiom of Choice here, or more precisely Axiom of Dependent Choice (because each choice depends on the previous ones).

(v) is a rephrasing of (i). □

We shall use $\alpha, \beta, \gamma, \ldots$ to denote well-orders, although they are usually reserved for ordinals. Here are some examples and non-examples of well-orders.

- The "shortest" well-order is the empty order, denoted by 0. Finite linear orders are well-orders.
- A sub-order of a well-order is also a well-order.
- We can "add" α and β by putting β on the right of α. Fomrally, we take the disjoint union of the underlying sets of α and β, and define that everything in α is smaller than everything in β. In particular we can add one element to the right of α, denoted $\alpha + 1$. More generally we can form $\sum_{i \in I} \alpha_i$, where I is a well-order and α_i is a well-order for each $i \in I$. Note that addition is not commutative: $1 + \omega \simeq \omega$ while $\omega + 1 \not\simeq \omega$, although one can show that it is associative.
- One can also define multiplication and exponentiation of well-orders.
- $\mathbb{Z}, \mathbb{Q}, \mathbb{R}, \mathbb{N} + \mathbb{Z}$ are not well-order. By (iv) a linear order is a well-order iff it doesn't contain a reverse copy of $\mathbb{N}$ (under Axiom of Choice).

If $\alpha \simeq \beta + 1$ for some β (i.e., α has a greatest element), it is called a *successor*. A well-order that is neither the empty order 0 nor a successor is called a *limit*; some texts regard 0 as a limit.

We also use 0, successor and limit for elements in a well-order: if α is a well-order, we denote the least element of α by 0; for $x \in \alpha$, we use $x + 1$ to denote the immediate successor of x in α, and call it a successor element; an element that is neither 0 nor a successor is a limit element.

As is well-known, Axiom of Choice is equivalent to Zorn's lemma. In fact it is also equivalent to the Well-ordering principle that any set can be well-ordered. Zorn's lemma⇒Well-ordering principle is relatively easy: to well-order X, consider the collection of well-orderings of subsets of X under the partial order "$(J, <_J)$ is an end-extension of $(I, <_I)$". Therefore one way to construct an uncountable well-order is simply to well-order your favorite uncountable set, such as $\mathbb{R}$. There are two downsides: (i) this is not very concrete, it's difficult to imagine what a well-order on $\mathbb{R}$ looks like; (ii) this relies on AC. Our main goal is to construct the "canonical" uncountable well-order, known as ω_1, in a slightly more concrete manner and without using AC (however some properties of ω_1 do rely on choice). It turns out that ω_1 is also, roughly speaking, the smallest uncountable set. Before proceeding to the construction, let us see some more examples of well-orders arising in nature.

1.1.5 Zorn's lemma

One place where uncountable well-order shows up is in the proof of Zorn's lemma from Axiom of Choice. Suppose $(P, \leq)$ is a partial order in which any chain has an upper bound. We show that there is a maximal element. Otherwise we could build a strictly increasing sequence $p_0, p_1, \ldots p_\omega, p_{\omega+1}, \ldots$

One day this sequence will become uncountable, and before long it will become even more uncountable. Eventually its size will exceed that of P, which is absurd since the elements in the sequence are distinct.

1.1.6 Long line

This is a famous counterexample in topology. Starting with the half-open interval $[0, 1)$, we can extend it by placing a copy of itself on the right. The result is topologically the same as $[0, 1)$. We repeatedly extend the half-open interval in this way. After extending it for ω or even α many times, where α is any countable ordinal, the result is still the same as $[0, 1)$ (basically because $[0, \infty)$ and $[0, 1)$ are homeomorphic), so we can continue this all the way up to ω_1, and stop when we have extended the interval for exactly uncountably many times. Formally we consider the order topology on $\omega_1 \times [0, 1)$ under lexicographical order. The result is called the closed long ray. If we remove the left end point we get the open long ray, and gluing the endpoints of two closed long rays gives the long line. These are examples of non-second-countable manifolds (with boundary).

The long line is not compact, but it has the property that the closure of any countable subset is compact; this property is termed ω-bounded. There is a structure theorem (Bagpipe theorem) for ω-bounded surfaces, extending the classification of compact surfaces.

1.1.7 Free abelian groups

A subgroup H of a free abelian group G (of arbitrary rank) is also free abelian; this fact is useful, e.g., when one is working with singular homology and would like to know that the cycle groups Z_n and boundary groups B_n are free abelian. More generally, a submodule of a free module over a PID (principal ideal domain) is also free. The case when G has infinite rank is usually proved using Zorn's lemma. There is another proof using a well-ordered basis, which is arguably more intuitive since it is a straightforward generalization of one proof for finite rank case. The presentation below is a slight rephrasing of the proof in [3].

Let us first prove the finite rank case. Say G has rank n and a basis $\{e_0, \dots, e_{n-1}\}$. For $0 \le k \le n$, let G_k be the subgroup generated by $e_0, \dots, e_{k-1}$ (so $G_0 = 0$ and $G_n = G$) and $H_k = G_k \cap H$. If we express all elements in terms of the basis and view them as column vectors, then H_k is the subgroup of H consisting of column vectors that are zero at the k-th entry and below. We claim that if H_{k+1} is strictly bigger than H_k, then it is generated over H_k by a single vector. Indeed the set of k-th entries of all vectors in H_{k+1} is a subgroup of $\mathbb{Z}$, hence cyclic, so there is a vector v_k whose k-th entry generates that group; v_k would then generate H_{k+1} over H_k; if $H_{k+1} = H_k$ we let v_k be the zero vector. By induction $v_0, \dots, v_k$ generate H_{k+1}, and therefore $v_0, \dots, v_{n-1}$ generate H. If v_k is nonzero then its last nonzero entry is at k, which implies that they form a free set of generators (if we remove the zero vectors).

The proof for general case is almost verbatim. Consider a basis of G and well-order it as $\{e_0, e_1, e_2, \dots e_\omega, e_{\omega+1}, \dots\}$. Elements of G are now infinite column vectors, but with only finitely many nonzero entries. Let G_α be the group generated by $\{e_\beta : \beta < \alpha\}$; note that $\bigcup_{n<\omega} G_n = G_\omega$, and more generally $\bigcup_{\beta<\alpha} G_\beta = G_\alpha$ for limit ordinal α. Let $H_\alpha = G_\alpha \cap H$. The same argument shows that a single vector v_α generates $H_{\alpha+1}$ over H_α. By transfinite induction, these v_α form a free generating set of H.

1.2 Construction of ω_1

We return to our goal of constructing ω_1. Part (iii) of the following lemma will be particularly useful for us.

Lemma 2. *(i) If α is a well-order and $f : \alpha \to \alpha$ is an order-embedding, namely $x < y$ implies $f(x) < f(y)$, then $f(x) \geq x$ for all $x \in \alpha$.*

(ii) A well-order cannot be isomorphic to a proper initial segment.

(iii) Different initial segments of a well-order are non-isomorphic.

(iv) The only automorphism of a well-order is identity.

(v) If two well-orders α and β are isomorphic then the isomorphism is unique.

Proof. (i) If $f(x) < x$ then applying f we get $f(f(x)) < f(x)$; applying f again we get $f(f(f(x))) < f(f(x))$, etc. Thus there is a strictly decreasing sequence $x > f(x) > f(f(x)) > \cdots$, a contradiction.

Alternatively one can argue as follows. If $A = \{x \in \alpha \mid f(x) < x\}$ were nonempty, it must has a least element. It cannot be the least element 0 of α since nothing is smaller than 0. It cannot be a successor because if $f(x) \geq x$ then $f(x+1) \geq x+1$. It cannot be a limit because if $f(y) \geq y$ for all $y < x$, then $f(x) > y$ for all $y < x$, so $f(x) \geq x$. Thus A must be empty.

(ii) Recall from Definition 1 (v) that any proper initial segment is of form $\alpha_x = \{y \in \alpha \mid y < x\}$, for some $x \in \alpha$. An isomorphism $f : \alpha \to \alpha_x$ can be viewed as an order embedding $f : \alpha \to \alpha$ whose image is α_x, in particular $f(x) < x$, contradicting part (i).

(iii) First note that two initial segments I and J can always be compared, i.e., either $I \subseteq J$ or $J \subseteq I$. This is true in any linear order, not just well-order.

Say $I \subsetneq J$, then I can be regarded as a proper initial segment of the well-order J, so by (ii) they are non-isomorphic.

(iv) If f is an automorphism then so is f^{-1}. By (i) we have $f(x) \geq x$ and $f^{-1}(x) \geq x$ for all x, so $f(x) = x$.

(v) follows from (iv), since it is true in any category that if $A \simeq B$, then the automorphisms of A are in bijection with isomorphisms from A to B. □

The single most important property of well-orders is that given two well-orders there is a unique way to align them, and hence we can compare their lengths.

Theorem 3. *Let α and β be any well-orders. Exactly one of the following happens.*

(i) α is isomorphic to a proper initial segment of β.

(ii) β is isomorphic to a proper initial segment of α.

(iii) α is isomorphic to β.

Proof. The idea is to align the 0-th element of α with the 0-th element of β, and then the 1-st, the 2-nd,...the ω-th,...until either α or β runs out of elements. At each stage there is only one possible thing to do: the next element in α is matched with the next element in β. If for example, α runs out of element first, then we are in case (i).

Formally, let $\mathcal{F}$ be the set of all order-isomorphisms $f : I \to J$ where I, J are initial segments of α, β respectively. $\mathcal{F}$ is ordered under extension of maps. The maximal element would be the desired isomorphism. This is a bit similar to Zorn's lemma, except that $\mathcal{F}$ is actually linear: we shall show that for any $f, g \in \mathcal{F}$, one of them extends the other. First notice that either $\mathrm{dom}(f) \subseteq \mathrm{dom}(g)$ or $\mathrm{dom}(g) \subseteq \mathrm{dom}(f)$ since they are initial segments.

Say $\mathrm{dom}(f) \subseteq \mathrm{dom}(g)$. Suppose for contradiction that $\{x \in \mathrm{dom}(f) \mid f(x) \neq g(x)\}$ is nonempty, so there is a least element a. But then

$$\begin{aligned} f(a) &= \text{the lsub of } \{f(x) \mid x \in \alpha, x < a\} \\ &= \text{the lsub of } \{g(x) \mid x \in \alpha, x < a\} \\ &= g(a), \end{aligned}$$

a contradiction; the first equality holds because if $f(a)$ were bigger than the lsub of $\{f(x) \mid x \in \alpha, x < a\}$, then $\mathrm{ran}(f)$ would fail to be an initial segment; the third equality holds for the same reason (in plain words, if f and g agree up to some point, they have to agree at the next point). Thus g extends f, as desired.

Therefore $\mathcal{F}$ contains a maximal map h, defined by $h(a) = f(a)$ for any $f \in \mathcal{F}$ s.t. $a \in \mathrm{dom}(f)$. Either $\mathrm{dom}(h) = \alpha$ or $\mathrm{ran}(h) = \beta$, otherwise we would be able to extend h one step further.

We have thus proved that one of (i)-(iii) happens. That only one can happen follows from Lemma 2, e.g., if both (i) and (iii) happen, then β would be isomorphic to a proper initial segment of itself. From Lemma 2 we also know that, e.g., in the scenario of (i), α is isomorphic to a unique initial segment of β, and the isomorphism is unique. □

There is a quick way to define ω_1 using AC and the above theorem, as follows. Temporarily call a well-order $(X, <)$ ω_1-*like* if the set X is uncountable, but any proper initial segment $X_a = \{x \in X \mid x < a\}, \ a \in X$ is countable. Such a well-order exists: use Zorn's lemma to well-order your favorite uncountable set X to get an uncountable well-order $(X, <)$; if $(X, <)$ is ω_1-like then we are done; otherwise, let $a \in X$ be smallest such that X_a is uncountable, and $(X_a, <)$ would be ω_1-like. Also, the above theorem easily implies that there is up to isomorphism a unique ω_1-like well-order $(X, <)$. We claim that X is the smallest uncountable set, namely for any uncountable set Y, there is an injective map $f : X \to Y$. To see that, use Zorn's lemma to well-order Y as $(Y, <)$, and by the theorem we can compare the lengths of X and Y. It cannot be the case that Y is isomorphic to a proper initial segment of X, by the definition of being ω_1-like, and therefore either they are isomorphic or X is isomorphic to some Y_b; either way we get the desired injective map. For details see [3].

Thus we have constructed the smallest uncountable set, and as can be seen from the proof, this $(X, <)$ is also the shortest uncountable well-order. As mentioned previously, the unsatisfactory thing about this argument is that it relies heavily on AC. We are going to construct ω_1 in a more concrete way, and prove its properties with as little choice as possible.

By Theorem 3, if Ω is a set of pairwise non-isomorphic well-orders, then Ω itself is naturally ordered by defining $\alpha < \beta$ if α is shorter than β, namely α is isomorphic to a proper initial segment of β. Transitivity is clear, and the theorem shows the trichotomy law. Moreover:

Proposition 4. $(\Omega, <)$ *is a well-order.*

Proof. First observe the following: if X is a linear order such that $X_a = \{x \in X \mid x < a\}$ is a well-order for any $a \in X$, then X is a well-order. Consider any nonempty subset $A \subseteq X$, and pick an arbitrary $a \in A$. If a is not the least element of A, then $X_a \cap A$ is nonempty so has a least element, which must be the least element of A.

Thus it suffices to show that Ω_α is a well-order for each $\alpha \in \Omega$. We show that by embedding Ω_α into α. Ω_α is the subset of Ω consisting of well-orders shorter than α. If β is such a well-order, then by definition β is isomorphic to a unique proper initial segment of α, which must be of form α_x. Define $f(\beta)$ to be that x; it is not too difficult to see that f is indeed an order embedding. □

We are now ready to construct ω_1, the first uncountable well-order. It should be longer than any countable well-order and nothing else. Recalling the guiding idea that an ordinal is just the set of smaller ordinals, we are tempted to make the following definition.

"Definition". ω_1 is the set of countable well-orders, well-ordered by comparing their lengths.

The first issue with this definition is that there are many isomorphic countable well-orders. To get a genuine linear order, a naive attempt is to consider "the set of isomorphism classes" of countable well-orders, or perhaps choose a representative from each class. That leads to the second issue: each isomorphism class is truly a *proper class*, namely it is too large to be a set, and we cannot collect a bunch of proper classes into a set. Below is one way to circumvent these issues.

Definition 5. ω_1 is the set of all isomorphism classes of well-orders on natural numbers, well-ordered by comparing their lengths.

To clarify this definition and why it resolves the issues, we need to discuss Russell's paradox. In set theory everything is a set; elements of a set are again sets. The existence of certain sets are guaranteed by axioms, and the existence of others needs to be proved. In the early days of set theory there was just a single axiom: the collection of all things satisfying any prescribed property is a set. In particular we can prescribe no property at all, so the collection of all sets is a set. This was found to be problematic:

Theorem (Russell's paradox)**.** *No set can contain all sets.*

Proof. If V is a set that contains all sets, then in particular $V \in V$. Strange as it may seem, this is not yet a contradiction, unless we impose that a set cannot contain itself. But we can derive a contradiction as follows. Call a set x strange if $x \in x$. Consider the subset $W \subseteq V$ consisting of all sets that are *not* strange, namely,

for any set x, $x \in W$ iff $x \notin x$.

Now we run into problem if we ask whether W is strange: since the above is true for any x, in particular it should apply to W, so we get the contradiction

$W \in W$ iff $W \notin W$. □

Thus the collection of all sets is too large to be a set, and is an example of a proper class. For basically the same reason, no set can contain all groups, or even all trivial groups. Set theoretically, a *group* is a pair $(G, \cdot)$ where G is a set and $\cdot$ is a binary function on G satisfying certain properties. A *binary function* is a subset of the Cartesian product $G \times G \times G$ satisfying certain properties; *Cartesian product* is defined using ordered pairs, and an *ordered pair* can be defined in various ways. The point is that if two groups have different underlying sets then they are different. For each set x we can create a trivial group whose underlying set is $\{x\}$. Thus there are as many trivial groups as there are sets, so the collection of all trivial groups is too large to be a set.

An ordinal is morally just an isomorphism class of well-orders. But by the same argument as above, each isomorphism class is too large to be a set. So we cannot define ω_1 to be the set of all isomorphism classes of countable well-orders. Instead we restrict our attention to those countable well-orders whose underlying set is actually the set of natural numbers $\mathbb{N}$. The collection of all such well-orders *is* a set. Each of them is a pair $(\mathbb{N}, <)$ where $<$ is a subset of $\mathbb{N} \times \mathbb{N}$. Below are some examples.

$0, 1, 2, \ldots$ represents the class ω

$1, 2, 3, \ldots 0$ represents the class $\omega + 1$

$1, 3, 5, \ldots 0, 2, 4, \ldots$ represents the class $\omega \cdot 2$

$0, 2, 4, \ldots 1, 3, 5, \ldots$ also represents the class $\omega \cdot 2$

The reason of considering isomorphism classes instead of choosing a set of representatives is, again, we want to show that the basic theory of well-orders doesn't depend on AC. Two isomorphism classes are compared in the expected way: consider *any* representative from each class and compare those two; this doesn't require choice, just like you don't need choice to define the group operation on a quotient group. For notational brevity we usually don't distinguish between an order and an isomorphism class. Technically, under our definition ω_1 does not contain the finite well-orders. To be more uniform one may artificially add them into ω_1, or consider well-orders on subsets of natural numbers instead.

The following lemma will be useful in the proof of the main theorem.

Lemma 6. *$\alpha \leq \beta$ iff there exists an order embedding $f : \alpha \to \beta$.*

By $\alpha \leq \beta$ we mean either case (i) or (iii) in Theorem 3 happens, namely α is isomorphic to an initial segment of β which is not necessarily proper. In particular, a sub-order of α cannot be longer than (but could be as long as) α.

Proof. The forward direction is clear. Suppose there is an embedding $f : \alpha \to \beta$ yet β is isomorphic to a proper initial segment of α. Composing the two maps gives an embedding of α into a proper initial segment of itself, contradicting Lemma 2. □

Theorem 7. *(i) Any countable well-order is strictly shorter than ω_1, so ω_1 is uncountable.*

(ii) $A \subseteq \omega_1$ is bounded iff it is countable.

(iii) ω_1 is the smallest uncountable set.

Proof. (i) This actually follows from the if direction of (ii), but that part requires choice. We give a proof that doesn't use choice.

Let α be a countable well-order. Define a map $f : \alpha \to \omega_1$ as follows: α_x, the initial segment below some $x \in \alpha$, is also a countable well-order (possibly finite, hence the remark above, but it's not a huge issue). Therefore α_x is isomorphic to a unique isomorphism class of well-orders on $\mathbb{N}$, and we let $f(x)$ be that class.

Clearly f is an embedding. In fact it can be checked that f is an isomorphism between α and an initial segment of ω_1, so $\alpha \leq \omega_1$. Alternatively we can use Lemma 6 as a shortcut: because there exists an embedding $f : \alpha \to \omega_1$, we must have $\alpha \leq \omega_1$. Since this works for any countable α, we also have $\alpha + 1 \leq \omega_1$, so $\alpha \simeq \omega_1$ is impossible by Lemma 2.

(ii) Suppose $A \subseteq \omega_1$ has α as an upper bound, so $A \subseteq (\omega_1)_\alpha$. Recall from the proof of Proposition 4 that $(\omega_1)_\alpha$ embeds into α, and therefore A embeds into α. Since α is countable, so is A. This means ω_1 is ω_1-like, as it should be.

A countable set is bounded because given countable well-orders $\alpha_0, \alpha_1, \alpha_2, \ldots$ we can form $\alpha = \sum_{i=0}^{\infty} \alpha_i$, which is a countable well-order. Each α_i embeds into α, so $\alpha_i \leq \alpha$ by Lemma 6. Strictly speaking, we are given countably many isomorphism classes, and need to *choose* a representative from each class, so some amount of choice is needed here.

(iii) Recall how we compare size of sets in set theory:

A and B have the same size, denoted $|A| = |B|$, if there exists a bijective map $f : A \to B$.

A is smaller than or equal to B, denoted $|A| \leq |B|$, if there exists an injective map $f : A \to B$.

A is strictly smaller than B, denoted $|A| < |B|$, if $|A| \leq |B|$ but not $|A| = |B|$.

Intuitively, if $f : A \to B$ is injective then A has the same size as the image of f, which is a subset of B, and a subset should not be larger. For finite sets, $|A| < |B|$ under this definition iff A has smaller number of elements than B; this is essentially Pigeonhole principle. It was Cantor who proposed that we might as well extend the definition to infinite sets. It is a theorem (Schröder-Bernstein) that if $|A| \leq |B|$ and $|B| \leq |A|$ then $|A| = |B|$.

Now suppose A is a set and $f : A \to \omega_1$ is injective. Applying Lemma 6 to the inclusion of $f(A)$ into ω_1, we see that $f(A)$ is isomorphic to an initial segment A' of ω_1. If A' is proper then it is countable by first part of (ii); otherwise $A' = \omega_1$, so $|A| = |\omega_1|$.

We have thus proved that if $|A| \leq |\omega_1|$ then either $|A| = |\mathbb{N}|$ or $|A| = |\omega_1|$. It turns out AC is equivalent to the statement "for any sets A, B, either $|A| \leq |B|$ or $|B| \leq |A|$". So under AC, if A is uncountable then $|\omega_1| \leq |A|$. Without AC there might exist uncountable sets that are incomparable with ω_1. For example it might happen that $|\mathbb{R}| \not\leq |\omega_1|$ and $|\omega_1| \not\leq |\mathbb{R}|$. □

1.3 Beyond ω_1

Recall that in set theory we count from zero, so one may wonder what is ω_0. It is simply an unpopular name for ω. For all purposes ω can be identified with the set of natural numbers $\mathbb{N}$, so we might also say that ω_1 is the set of isomorphism classes of well-orders on ω_0

Unsurprisingly, whenever there are subscripts we would like to keep going. So let us define ω_2 to be the collection of all isomorphism classes of well-orders on ω_1; it is the first well-order whose size is greater than ω_1, and the second smallest uncountable set under AC. Needless to say we then have $\omega_3, \omega_4, \ldots$ Next we take the limit of this sequence. One can either form the infinite sum $\sum_n \omega_n$ (pictorially, arrange $\omega, \omega_1, \omega_2, \ldots$ in a line), or take the direct limit of $\omega \to \omega_1 \to \omega_2 \to \cdots$, where each map is the isomorphism to the unique initial segment (pictorially, extend ω to ω_1, then to ω_2, etc.). It's not too hard to show that either way we get the same thing, which is denoted ω_ω. Below is an initial segment of this process.

$$\omega, \omega_1, \omega_2, \ldots \omega_\omega, \omega_{\omega+1}, \omega_{\omega+2}, \ldots \omega_{\omega\cdot 2}, \ldots \omega_{\omega^2}, \ldots \omega_{\omega^\omega}, \ldots \omega_{\varepsilon_0}, \ldots \omega_{\omega_1}, \ldots \omega_{\omega_\omega}, \ldots$$

Each of them represents a "milestone" in the endless stream of ordinals where there is a change in size. These special ordinals are called *cardinals*. The α-th cardinal is denoted either ω_α or $\aleph_\alpha$ ($\aleph$ is the first Hebrew letter Aleph); people often use ω_α to emphasize its length, and $\aleph_\alpha$ to emphasize its size. The limit of $\omega_\omega, \omega_{\omega_\omega}, \ldots$ is denoted $\omega_{\omega_{\ddots}}$ If we call it θ, then it has the hilarious property that $\omega_\theta = \theta$, namely the θ-th cardinal is θ itself.

Now (at last) we give the set theoretic definition of ordinals. Recall that the idea is to identify an ordinal with the set of smaller ordinals, and this leads to

$0 = \varnothing$

$1 = \{0\} = \{\varnothing\}$

$2 = \{0, 1\} = \{\varnothing, \{\varnothing\}\}$

$3 = \{0, 1, 2\} = \{\varnothing, \{\varnothing\}, \{\varnothing, \{\varnothing\}\}\}$

$\omega = \{0, 1, 2, \ldots\}$

$\omega + 1 = \{0, 1, 2, \ldots \omega\}$

and so on. Then ω_1 is the set of all countable ordinals. Note that in general, $\alpha + 1 = \alpha \cup \{\alpha\}$. Informally, ordinals are all the sets generated from the operation $\alpha \mapsto \alpha \cup \{\alpha\}$ and taking union, starting from the empty set.

How to make this into a rigorous definition? Note that under this definition, "smaller than" coincides miraculously with the set membership $\in$, and for each ordinal α, $(\alpha, \in)$ is a well-order. Then "the set of smaller ordinals" translates to "a transitive set of ordinals"; a set x is *transitive* if whenever $y \in x$ and $z \in y$, we have $z \in x$; equivalently $y \in x$ implies $y \subseteq x$. The idea is that we don't want sets such as $\{1, 3, 6\}$ or $\{0, 2, 4, \ldots\}$ to count as ordinals. Now one way to rigorously define ordinals is:

(i) An ordinal is a transitive set of ordinals.

This *is* a rigorous definition by transfinite induction on set membership. We start by declaring $0 = \varnothing$ to be an ordinal; then $1 = \{\varnothing\}$ is an ordinal since it's a transitive set of odinals, and similarly

$2 = \{0, 1\}$ is an ordinal; on the other hand, $\{1\}$ is not an ordinal because it's not transitive, and $\{0, 1, \{1\}\}$ is not an ordinal since it has an element $\{1\}$ that is not an ordinal, etc.

Obviously this is not a super convenient definition to work with. Some equivalent definitions are:

(ii) An ordinal is a transitive set X such that $(X, \in)$ is a well-order.

(iii) An ordinal is a transitive set X such that $(X, \in)$ is a linear order. The equivalence of this definition with others requires the Axiom of Regularity, which roughly says there cannot be an infinite sequence of sets $a_0 \ni a_1 \ni a_2 \ni \cdots$ that are descending in membership.

(iv) An ordinal is a transitive set of transitive sets.

The above definition of ordinals, proposed by von Neumann, has some nice consequences such as $\alpha < \beta$ iff $\alpha \in \beta$, and the lub of a collection of ordinals is simply their union. Moreover, every isomorphism class of well-orders contains exactly one ordinal.

Another thing named after von Neumann is the von Neumann hierarchy, defined as follows:

$V_0 = \varnothing$

$V_1 = \mathcal{P}(V_0) = \{\varnothing\}$, the power set of V_0

$V_2 = \mathcal{P}(V_1) = \{\varnothing, \{\varnothing\}\}$

$V_3 = \mathcal{P}(V_2) = \{\varnothing, \{\varnothing\}, \{\{\varnothing\}\}, \{\varnothing, \{\varnothing\}\}\}$

$\vdots$

$V_\omega = \bigcup_{n<\omega} V_n$

$V_{\omega+1} = \mathcal{P}(V_\omega)$

and so on and so forth. Under the Axiom of Replacement, we can keep doing this for all ordinals α. The collection of all ordinals is denoted Ord; of course it is not a set but a proper class. We then define the von Neumann universe $V = \bigcup_{\alpha \in \mathrm{Ord}} V_\alpha$ (which again is a proper class). Under Axiom of Regularity, V contains all sets. Recall that Axiom of Regularity roughly says there is no infinite descending chain of membership $a_0 \ni a_1 \ni a_2 \ni \cdots$ In particular $a \in a$ or $a \in b \in a$ cannot happen. If such weird things do happen then they must happen outside of V.

V_ω is the "finite world"; it can also be regarded as "the set of natural numbers", because the two structures $(V_\omega, \in)$ and $(\mathbb{N}, +, \cdot)$ are "the same" in a precise sense, namely they are *bi-interpretable*. $V_{\omega+10}$ is probably enough to handle most, if not all of ordinary mathematics. Some people would argue that actually a small fragment of $V_{\omega+1}$ already suffices. This is one reason they dislike ZFC as a foundation of mathematics: it gives us far too much. Still there are good reasons to care about ZFC. For example, the methods of Gödel and Cohen discussed below do not work in set theories that are too weak. Perhaps another reason is that the von Neumann hierarchy is there, so why not?

1.4 Continuum Hypothesis

Recall that the set of real numbers $\mathbb{R}$ is uncountable by Cantor's diagonal argument. Here is a sketch: $\mathbb{R}$ is in bijection with the interval $(0, 1)$, which can more or less be identified with $2^{\mathbb{N}}$, the set of 0-1 sequences, by considering binary expansion of a real number. It suffices to show that no

$f : \mathbb{N} \to 2^{\mathbb{N}}$ is surjective, i.e., given countably many 0-1 sequences one can always find a sequence different from all of them. Arrange these sequences into an "$\mathbb{N}$ by $\mathbb{N}$ matrix", and consider the diagonal sequence; it agrees with the n-th sequence at the n-th digit. Now flip the diagonal sequence, and we get a sequence that disagrees with the n-th sequence at the n-th digit, hence different from all of them. $2^{\mathbb{N}}$ can also be identified with the power set $\mathcal{P}(\mathbb{N})$. The above argument generalizes verbatim to show there is no surjection $f : X \to \mathcal{P}(X)$ for any set X (in case you wonder if we used the order on $\mathbb{N}$ anywhere, we didn't; the diagonal sequence isn't constructed recursively but all at once).

It can be proved that under AC, every set X is in bijection with ω_α for some α. Because ω_1 is the smallest uncountable set, $\mathbb{R}$ is at least ω_1. Could it be larger, like ω_2, ω_3, ω_{100}, or $\omega_{\omega+1}$? Equivalently, could there be a set whose size is strictly between $\mathbb{N}$ and $\mathbb{R}$? This question was settled completely by work of Gödel and Cohen.

Theorem (Gödel 1940, Cohen 1963)**.** *It depends.*

Gödel showed that $\mathbb{R}$ could be ω_1, by finding a *model* of set theory that satisfies $|\mathbb{R}| = \omega_1$. Cohen discovered the method of *forcing*, a very flexible way to create new models, which can make the size of $\mathbb{R}$ pretty much anything we want, subject to a simple constraint known as König's theorem: if the size of $\mathbb{R}$ is ω_α, then the ordinal α is either a successor or a limit that doesn't have a countable unbounded subset. For example, it could be $\omega_2, \omega_3, \omega_{100}, \omega_{\omega+1}, \omega_{\omega_1}$, etc., but not ω_ω or ω_{ω^2}.

A model is a concrete mathematical structure, in contrast to axioms and proofs, which are essentially just strings of symbols. For example, below are the familiar group axioms.

$\forall x \forall y \forall z (x \cdot y) \cdot z = x \cdot (y \cdot z)$

$\forall x \; e \cdot x = x \cdot e = x$

$\forall x \exists y \; x \cdot y = y \cdot x = e$

In contrast, a group is an actual set equipped with a binary operation that satisfies these axioms. Some of them satisfy the additional axiom $\forall x \forall y \; x \cdot y = y \cdot x$, and some of them do not. Therefore, we say that the abelian axiom is *independent* from the basic group axioms.

Of course, anyone who has taken abstract algebra knows there are abelian and non-abelian groups. Less obvious is an example from plane geometry. Is the parallel postulate true? Turns out it depends on which model of plane geometry we are looking at. There exist models where the parallel postulate holds, such as the ordinary Cartesian plane $\mathbb{R}^2$, and also models where it fails, such as the Beltrami-Klein model, Poincaré disk model or upper half plane model. We didn't specify what exactly a "model of plane geometry" is. In logic one interprets that as a structure consisting of a set of things called points, another set of things called lines, along with several relations between them that satisfies Hilbert's axioms for plane geometry (as opposed to, say, a Riemannian manifold). For more on this see [4].

Completely analogously, whether the Continuum Hypothesis holds depends on the particular model of set theory. A model of set theory is a pair (M, E), E a subset of $M \times M$, that satisfies for example the ZFC axioms; usually E is just set membership $\in$ restricted to M. The existence of various different models of set theory might at first seem more surprising than the existence of various models of group axioms or plane geometry, since set theory is supposed to work as a definite foundation, but it's really the same phenomenon. What *could* be confusing is that we are discussing models of set theory within set theory. As will be explained in the last section, the second "set theory" in the previous sentence is the metatheory, and the first "set theory" is the target theory, and it's important to clearly distinguish them.

We describe Gödel and Cohen's constructions very informally. Gödel starts with a model M of ZFC and throws some sets away, only keeping those "absolutely necessary" sets, resulting in a model L called the constructible universe. For example, all ordinals remain in L, and if A, B are in L we definitely want to take their intersections, union, Cartesian product, etc., so those are also kept. This process could discard some real numbers, and hence L's version of $\mathbb{R}$ could be smaller than M's version. L's version of ω_1 could also be smaller, since although L contains all the countable ordinals, it may not contain all the bijections needed to recognize the countablility of them. Because L is defined in a "constructible" way, Gödel was able to analyze the structure of L closely and show that $|\mathbb{R}| = \omega_1$ (and much more) holds in L.

Conversely, Cohen violates $|\mathbb{R}| = \omega_1$ by starting with a model M and throwing in some real numbers, along with all sets they "generate" over M. Firstly, for there to be any new real number out there, M has to be small enough to start with; indeed it is usually assumed to be countable. Also one needs to add things in a controlled manner, so that the resulting structure is still a model of set theory. It turns out one can achieve that by throwing in some real numbers that are "random enough", or generic over M; for example, as mentioned above M is usually countable, and we don't want to add a real number that "codes" the countability of M. Cohen's forcing method adds G, a large set of new reals into M, resulting in a model $M[G]$, while also ensuring that $M[G]$'s version of ω_1 is the same as M's version, so $M[G]$ would be a model of $|\mathbb{R}| \neq \omega_1$.

Gödel's method of shrinking a model and Cohen's method of enlarging a model are both extremely important tools for logicians. They have been used to establish tons of independence results in set theory, some of which are related to other fields of mathematics. It is a major open question whether these are "essentially the only methods to create new models".

Also, although we now know that CH is independent from usual set theory axioms, it remains an interesting question whether it *should* be true, or more broadly if there exists a "one true model of set theory", where CH along with all other questions has a definite answer.

1.5 Inaccessible cardinal

Another thing related to independence is the so-called inaccessible cardinal. We know that ω_1 is much longer than ω, and ω_{n+1} is much longer than ω_n. But ω_ω is a bit different: there is a countable sequence $\omega, \omega_1, \omega_2, \ldots$ that is unbounded below ω_ω.

This is not the case for ω_1, since any countable subset of ω_1 is bounded according to Theorem 7 (but remember that this needs a bit of choice). This is also why the Borel hierarchy stops at stage ω_1: if $A_n \in \bigcup_{\alpha<\omega_1} \Gamma_\alpha$ for each $n < \omega$, then there is some countable ordinal α such that Γ_α already contains every A_n, and therefore $\bigcup_n A_n \in \Gamma_{\alpha+1}$. This property deserves a name.

Definition 8. A cardinal ω_α is *regular* if any unbounded subset has size ω_α. Otherwise it is *singular.*

It can be shown that equivalently, ω_α is regular iff there is no partition $\omega_\alpha = \bigcup_{i \in I} X_i$ of ω_α into a "smaller number of smaller pieces", meaning I as well as each X_i are all strictly smaller than ω_α. Since ω is not a finite union of finite sets, and ω_1 is not a countable union of countable sets, these two are regular. Similarly it can be proved that a successor cardinal, namely $\omega_{\alpha+1}$ is always regular. On the other hand, $\omega_\omega, \omega_{\omega^2}, \omega_{\varepsilon_0}, \omega_{\omega_1}, \omega_{\omega_{\ddots}}$ are singular. Does there exist any regular limit cardinal (ω_α where α is a limit ordinal)? Such a cardinal is called *inaccessible*; it would have to be much larger than all the cardinals above; it has to satisfy $\omega_\alpha = \alpha$ and much more. Intuitively Ord, the collection of all ordinals, should be inaccessible since there is no way "something" is comparable to "everything". Even ignoring the obvious issue that Ord is not a set, this argument still leads to some fierce philosophical debates. Mathematically here is all that we know:

Theorem (Corollary of Gödel's second incompleteness theorem)**.** *Either inaccessible cardinal is inconsistent (leads to contradiction), or it is consistent but the consistency is unprovable.*

Very roughly speaking, Gödel's second incompleteness theorem says mathematics cannot prove its own consistency. Inaccessible cardinal is so powerful that it actually implies the consistency of mathematics, or more precisely the consistency of ZFC axioms. Inaccessible cardinal is related to Grothendieck universe used in category theory, which is a set so huge and rich that all the usual mathematics (including ZFC set theory) can be carried out inside the set; in other words it is a model of mathematics, and therefore establishes its consistency.

Here is the takeaway: one either *believes* that the statement IC = "there exists an inaccessible cardinal" is consistent, or tries to derive a contradiction from it. If IC is indeed consistent then it is independent from usual mathematics. The situation is strikingly different from CH, which is unconditionally independent from usual mathematics.

So why believe in the consistency of inaccessible cardinal? One reason is that nobody has ever found a contradiction; of course various people claimed to have found contradiction, but none of their work is peer-reviewed. One might hope for a more positive reason to believe in inaccessible cardinal. That leads us to the topic in our title.

1.6 All sets are measurable

Some typical non-Lebesgue measurable sets are the Vitali sets, sets that appear in Banach-Tarski paradox, and Bernstein sets. A Vitali set is constructed by declaring $x, y \in \mathbb{R}$ to be equivalent if $x - y \in \mathbb{Q}$, and choose a number from each equivalence class; in other words it is a set of representative for $\mathbb{Q}$-cosets of $\mathbb{R}$. Similarly, to double a ball one considers the rotation group of sphere and certain coset representatives. A Bernstein set $X \subseteq \mathbb{R}$ is such that both X and $\mathbb{R} \setminus X$ intersect any uncountable closed sets (and therefore both have inner measure zero). One can build a Bernstein set by first well-ordering all uncountable closed sets in a row, and then proceeding by transfinite induction.

It is apparent that in all three examples, we are using Axiom of Choice heavily. Some people would even complain that the above paragraph doesn't count as "construction" of non-measurable sets, because AC is highly non-constructive: it merely tells us there exists a set of representatives or there exists a well-ordering of all uncountable closed sets, without providing a specific one.

Naturally we ask: can we construct a non-measurable set without using AC? By "without using AC" we really mean "without using too much AC", rather than discarding AC completely. Here are some weaker variants of AC:

- AC_ω (Axiom of Countable Choice) is AC restricted to countable family of sets, i.e., given countably many sets one can choose an element from each of them.
- $\mathsf{AC}_\omega^{\mathbb{R}}$ is AC_ω restricted to subsets of reals, i.e., given countably many subsets of $\mathbb{R}$ one can choose an element from each of them. This is enough to show the regularity of ω_1 (and thus the Borel hierarchy terminates at ω_1).
- DC (Axiom of Dependent Choice) is a principle strictly stronger than AC_ω, although it's still much weaker than full AC. It says if $(\mathbb{P}, <)$ is a partial order without maximal element, then there exists an infinite increasing chain $p_0 < p_1 < p_2 < \cdots$; intuitively, one can make a countable sequence of choices where the set to choose from at each step depends on the choices made thus far.

Without $\mathrm{AC}_\omega^{\mathbb{R}}$, it might happen that $\mathbb{R}$ is a countable union of countable sets, in which case the Lebesgue measure is not σ-additive. The stronger DC is sometimes needed in analysis, for example when one invokes Hahn-Banach theorem for non-separable spaces, or when one chooses a subsequence satisfying certain properties, and then a subsequence of that subsequence…as in the proof of Arzela-Ascoli. It is also used in the proof of "every PID is a UFD". It turns out AC_ω actually suffices for Arzela-Ascoli, but it seems open whether DC is strictly necessary for "every PID is a UFD".

The takeaway is that some amount of choice is needed to do usual mathematics. Many schools of constructive mathematics accept AC_ω or even DC. So let us modify our question: can one construct a non-measurable set AC is weakened to $\mathsf{AC}_\omega^{\mathbb{R}}$?

Turns out this depends on whether one believes in inaccessible cardinal.

Theorem (Solovay 1970, Shelah 1984)**.** *The following are equivalent:*

(i) One cannot prove the existence of non-measurable set if AC *is weakened to* $\mathsf{AC}_\omega^{\mathbb{R}}$*.*

(ii) Inaccessible cardinal is consistent.

The theorem is often stated with DC instead of $\mathsf{AC}_\omega^{\mathbb{R}}$, because Shelah's papers are notoriously difficult to read, and the simplified proof that appeared later used DC.

Here is what they actually did: assuming the existence of an inaccessible cardinal, Solovay used Cohen's forcing method to produce a model in which all sets are measurable, and also DC holds. Conversely, starting from "all sets are measurable" and $\mathsf{AC}_\omega^{\mathbb{R}}$, Shelah showed that ω_1 is inaccessible from L's point of view, where L is Gödel's constructible universe, so L is a model of "there exists an inaccessible cardinal".

Inaccessible cardinals are used in algebraic geometry in the form of Grothendieck universes. They are also used in type theory, even in the popular proof assistant Lean, and the fact that Lean hasn't crashed so far might be taken as evidence for consistency of inaccessible cardinals, and thus for the consistency of "all sets are measurable". But personally I find the other direction of the equiconsistency result more appealing: nobody has ever managed to build a non-measurable set without using choice, which is evidence for consistency of inaccessible cardinal. Of course, maybe inaccessible cardinal actually is inconsistent, the choice-free non-measurable set is just around the corner, and all set theorists will lose their jobs tomorrow. In this regard, doing set theory is kind of an extreme sport.

Here is something to know for those who doubt the consistency of ZFC itself. In [5] it is argued that a version of Shelah's proof goes through in the system $\mathsf{BZ}+\mathsf{DC}$ (BZ is bounded Zermelo set theory, a fairly weak fragment of ZF) to produce a model of ZFC, assuming all sets are measurable. Therefore, whoever wants to disprove ZFC needs to construct a non-measurable set in $\mathsf{BZ}+\mathsf{DC}$.

Inaccessible cardinal is at the lower end of the *large cardinal* hierarchy. Large cardinals are to logicians as Riemann hypothesis and its generalizations are to number theorists, the only difference being that they are provably unprovable. Inner model theory, a central topic in set theory, is concerned with "explaining" why large cardinals should be consistent, which often leads to results similar to Solovay-Shelah theorem, establishing correspondence between large cardinals, which are high in the clouds, and properties of reals, which are more down to earth. When a large cardinal notion is shown to correspond to a certain property of reals, it is usually viewed as strong evidence for the consistency of the large cardinal. Current state of art of inner model theory has reached somewhere above Woodin cardinals, but is still significantly short of supercompact cardinal. As an example, if there are infinitely many Woodin cardinals then all projective sets are determined, and conversely if all projective sets are determined, one can find an inner model (in place of L in Shelah's result) which has many Woodin cardinals.

We briefly explain what projective sets and determinacy are. They belong to the field known as *descriptive set theory*, which started with Lebesgue's paper *Sur les fonctions représentables analytiquement*. In this paper Lebesgue initiated a systematic study of functions that are analytically representable: the smallest class of functions containing all polynomials and closed under pointwise convergence. These are exactly what we now call Borel functions, because by Stone-Weierstrass approximation theorem it contains all continuous functions, and then it's not hard to produce all characteristic functions of Borel sets, and then all Borel functions. In the paper Lebesgue proved a "Borel implicit function theorem", the most basic case of which is the following: if $F(x,y)$ is a two-variable Borel function and $f(x)$ is defined by

$f(x) =$ the unique y such that $F(x,y)=0$

then $f(x)$ is Borel. In other words, if the graph of a function f is a Borel subset of the Cartesian plane, then f is Borel. In his proof Lebesgue used the simple lemma that projection of a Borel subset of the plane onto the x-axis is Borel. Ten years later, the young Suslin discovered that the lemma was false: the projection of a Borel set may not be Borel. He termed these sets analytic, and together with his teacher Luzin proved various properties of analytic sets:

- Images and preimages of analytic sets under Borel functions are analytic. Analytic sets are closed under countable unions and intersections, but not under complements; complement of an analytic set is called coanalytic.
- If both A and A^c are analytic then A is in fact Borel; from this it's easy to show that Lebesgue's implicit function theorem was actually true, despite the incorrect "lemma".
- Analytic sets have *perfect set property*, i.e., an analytic set is either countable or contains a nonempty perfect set, a set whose Cantor-Bendixson derivative is itself, and thus analytic sets cannot serve as counterexamples to the Continuum Hypothesis.
- Analytic sets are Lebesgue measurable.

- Analytic sets are *Baire measurable*. A set F is nowhere dense if its closure $\overline{F}$ has empty interior. A set is meager if it is a countably union of nowhere dense sets. Meager sets are analogous to sets of measure zero in several ways. A set is Baire measurable if it differs from some open set by a meager set. Note that every closed set F is Baire measurable since $F \setminus \mathrm{Int}(F)$ is closed and has empty interior, from which it's easy to show that Borel sets are Baire measurable.

Unfortunately Suslin died young. His work was continued by Luzin, among others. Analytic sets are closed under continuous images, but continuous images of coanalytic sets produce a new class of sets. Iterating the operations of taking continuous images and complements, we get the projective sets. It was soon noted that the problem of whether one could extend the above results about analytic sets to projective sets seemed extremely difficult. Luzin was so frustrated that he famously said the following [7][8]:

One does not know and one will never know whether projective sets are measurable.

Today we know these problems are difficult for the reason that they are entangled with large cardinals, and are thus independent.

Another early result in descriptive set theory was Kondo-Addison uniformization theorem, which says if $A \subseteq \mathbb{R} \times \mathbb{R}$ is coanalytic, then there exists a function f whose domain is the projection of A onto the x-axis, and whose graph is a coanalytic set contained in A. This can be interpreted as a kind of "constructive choice principle". It follows that every Borel surjection has a coanalytic right inverse (in general there is no Borel right inverse). Attempts to extend this result to projective sets were futile, until the axiom of *Projective Determinacy* was proposed in the 60s. This is an axiom defined in terms of infinite games and imply all the nice properties of analytic/coanalytic sets extend to projective sets. By work of various people, Projective Determinacy was shown to be equiconsistent with (roughly) infinitely many Woodin cardinals. Determinacy has since become an integral part of set theory.

2 Proofs and models

The main goal of the second part is to *state* Gödel's incompleteness theorems, and to give some idea of their significance and what the proof is like. Even just to state the theorems, we unevitably have to delve into the details of first order logic– formulas, proofs, structures, satisfaction, etc. This part can be safely ignored unless you really want to torture yourself with all the logical subtleties.

2.1 Naive set theory

Before the 19th century, there was no coherent definition of numbers, functions, convergence, etc., and things worked *fine*. Most theorems people proved were mostly correct. But the need of mathematical rigor steadily increased. A notable example was Fourier analysis, which forced people to face functions with infinitely many discontinuities, and rethink about the foundation of analysis. By the end of the 19th century, infinitesimal was phased out in favor of ε-δ argument (and models of hyperbolic geometry were discovered around the same time). But to really put analysis on a firm ground, one needs a clear notion of real numbers. People invented various equivalent ways to define real numbers out of rational numbers, like decimal expansions, Cauchy sequences (due to Cantor), and Dedekind cuts.

Rational numbers are in turn defined from natural numbers. But what is a natural number? For set theorists a natural number is an element of the smallest limit ordinal ω; recall that an ordinal is a transitive set of transitive sets, and is a limit if it is neither 0 nor a successor, where successor means something of form $\alpha \cup \{\alpha\}$. This may seem ad hoc and more complicated than one might wish, but at least we finally have a fully rigorous definition of natural numbers (and thus real numbers) in terms of sets—except now the question becomes, what is a set?

It should be clear that, if we want any mathematical rigor at all, we cannot keep tracing back like this forever. We have to agree on some *primitive notions*, which are things taken for granted, that will serve as cornerstones for everything else. Set theorists choose *set* and *set membership* as their primitive notions. Set theory is versatile enough to express most of other mathematics: numbers are sets, groups are sets, manifolds are sets... Category theorists would prefer objects and morphisms since they emphasize relations and structures. Interestingly, "structural set theory" is often bi-interpretable with "material (conventional) set theory", so one may say they are different flavors of the same thing.

We shall go over the *naive* ZFC set theory axioms, in the sense that we state the axioms in English rather than formal language.

- Axiom of Separation. If X is a set and $P(x)$ is a property, there is a set Y consisting of exactly the elements x in X that satisfy the property $P(x)$.

  The point of starting with a set X is that, as shown by Russell, "the collection of all sets satisfying a prescribed property is a set" leads to contradiction. The strategy is to weaken that by restricting to members of some fixed set X. The word "property" may seem a little vague, which it is, but for now let's not worry about that.

  Separation along is not strong enough to ensure the existence of various sets we need to do mathematics: imagine that the only set in the world is the empty set, then Axiom of Separation holds, because any subset of the empty set is empty. Most of the rest of ZFC axioms are there to ensure enough sets exist.

- Axiom of Pairing. For any sets a and b, there exists a set whose only elements are a and b. This set is denoted $\{a, b\}$ and called the unordered pair containing a and b. It is allowed that $a = b$, and in that case the axiom says the singleton $\{a\}$ exists.

  Now we can define what a function is, say with domain A and codomain B. The idea is that a function is a graph that passes the vertical line test. A graph should be a subset of $A \times B$, the Cartesian product, which should be the set of ordered pairs. But what is an ordered pair? To define (a, b) we need to create an asymmetry between the sets a and b. There are various ways, and the most popular one is Kuratowski's definition $(a, b) = \{\{a\}, \{a, b\}\}$; this works smoothly even with $a = b$, and even if one does not impose the Axiom of Regularity to rule out weird things such as $a = \{a\}$.

  We can also define ordered n-tuple. There are two straightforward definitions: an n-tuple is a function whose domain is $n = \{0, \cdots, n-1\}$, or, inductively, an n-tuple is an ordered pair (a, b) where a is an $n-1$-tuple. It takes some effort to see that the second definition is really a single definition with a parameter n, instead of infinitely many definitions, one for each n.

- Axiom of Union. For any set X, there exists a set which is the union of all members of X, denoted $\bigcup X$.

  What we usually write as $\bigcup_{n=0}^{\infty} A_n$ is $\bigcup\{A_0, A_1, A_2, \dots\}$ in this notation. $A \cup B$ is defined to be $\bigcup\{A, B\}$, where $\{A, B\}$ exists by Axiom of Pairing.

- Axiom of Power Set. For any set X, there exists a set denoted $\mathcal{P}(X)$, whose elements are all the subsets of X.

  With paring, union and separation, we can now prove the basic (but by no means obvious!) fact that the cartesian product $A \times B$ of two sets exist. The ordered pair (a, b) is defined to be $\{\{a\}, \{a, b\}\}$, which is an element of $\mathcal{P}(\mathcal{P}(A \cup B))$, so we just consider the subset of $\mathcal{P}(\mathcal{P}(A \cup B))$ consisting of all elements of the form $\{\{a\}, \{a, b\}\}$. As a corollary, the collection of all functions from A to B is also a set.

- Axiom of Infinity. There exists a limit ordinal, namely an ordinal that is neither zero nor a successor; this implies the existence of a smallest limit ordinal, which we denote ω. A more common way to state infinity is that there exists an inductive set, which is a set X that contains $\varnothing$, and whenever $a \in X$, $a \cup \{a\} \in X$ (be aware that an inductive set may not be transitive, let along an ordinal). Then ω can be defined as the smallest inductive set and proved to exist.

  With the axioms so far, one can already construct usual arithmetical operations on ω, real numbers and many other standard objects in analysis and algebra.

- Axiom of Extensionality. If two sets have exactly the same elements, then they are in fact the same. That is, a set is determined by its elements.

  This axiom is often overlooked, yet very important for our intuition about sets to actually work. It ensures that, e.g., there is only one empty set.

- Axiom of Choice. If $\{X_i\}_{i \in I}$ is a family of nonempty sets, then there exists a function f with domain I, such that $f(i) \in X_i$ for every $i \in I$.

  A family $\{X_i\}_{i \in I}$ is nothing but a function with domain I, where X_i is the value of the function at i. The axiom rougly says we can make infinitely many choices at once. This is one of the most notorious axioms in set theory. Another formulation of the axiom is that the infinite product $\prod_{i \in I} X_i$ is nonempty.

  Axiom of choice differs from the previous set existence axioms in an important way: the choice function it provides is far from unique. That is why AC is often regarded as highly non-constructive.

- Axiom of Replacement. Suppose X is a set, and $P(x, y)$ is a property such that for any set x in X, there exists a unique y such that $P(x, y)$ holds; then $\{y : P(x, y) \text{ for some } x \in X\}$ is a set; think of this as "replacing" each $x \in X$ by its corresponding y.

  Replacement is rarely mentioned outside set theory but extremely important, and you may already be using it subconsciously. Because of its importance, we will illustrate the use of replacement by several examples. First let's reprove that the cartesian product $A \times B$ of two sets exists; recall that previously we gave a somewhat ad hoc proof that relied on the power set axiom and the particular definition of Kuratowski ordered pair $(a, b) = \{\{a\}, \{a, b\}\}$.

Alternatively we can argue as follows. Fix an $a \in A$; every $b \in B$ corresponds to the ordered pair (a, b), so by replacement $\{a\} \times B$ is a set. Now consider the correspondence $a \mapsto \{a\} \times B$ and apply replacement again; we get that $\{\{a\} \times B : a \in A\}$ is a set. Applying the union axiom to this set gives us $A \times B$. Note that this does not rely on a particular definition of ordered pair.

For the next example, let's use replacement to prove that a countable union of countable sets $\bigcup_n A_n$ is countable (say they are all infinite and disjoint). The obvious thing to do is to choose a bijection $f_n : \omega \to A_n$ for each n, assemble them together to get a bijection $g : \omega \times \omega \to \bigcup_n A_n$, and compose that with a bijection between $\omega \times \omega$ and ω. There are two subtle things here. Firstly, as the word "choose" suggests, this proof doesn't work without the Axiom of Choice. Secondly, to apply choice we need a family of sets to start with. Denoting by $\mathcal{F}_n$ the set of bijections between A_n and ω (which is nonempty by assumption), we want to apply choice to the family $\{\mathcal{F}_n\}_{n\in\omega}$. But why does this family exist? There is an argument using power set, similar to the first proof of the existence of cartesian product, but it is more natural to apply replacement to the correspondence $n \mapsto \mathcal{F}_n$.

Our final example illustrates how replacement allows us to "do things forever". Suppose we have an abelian group A and want to form an exact sequence $\cdots \to F_2 \to F_1 \to F_0 \to A \to 0$, aka a free resolution. First we want a surjective homomorphism $F_0 \to A$ from some free abelian group onto A. One natural option is to let F_0 be the ablian group with basis A, and send every basis element to itself. Then we take the kernel K_0 of this map and look for a surjective homomorphism $F_1 \to K_0$; again the natural choice is to let F_1 be the free group on the set K_0, etc. Now we are tempted to "continue this forever" to get the free resolution, but how to justify that using our axioms? We can consider the correspondence $n \mapsto \overrightarrow{F_n}$, where $\overrightarrow{F_n}$ is the "n-step canonical free resolution of A"; then we piece them together to get the complete free resolution. This is a toy example of the method of "transfinite recursion" that is used all the time in set theory.

Replacement also ensures the existence of the von Neumann ordinal $\omega + \omega = \{0, 1, 2, \dots \omega, \omega + 1, \omega + 2, \dots\}$. Without replacement, $\omega + \omega$ may not exist as an ordinal, although one can define a linear ordering on ω that "looks like" $\omega + \omega$, such as $0, 2, 4, \dots 1, 3, 5, \dots$ Similarly, replacement guarantees the existence of large sets like $\omega_1, \aleph_{\omega+\omega}$ and $V_{\omega+\omega}$, and thus is often viewed as the axiom that grants ZFC the most power; it could sometimes be more controversial than Axiom of Choice in discussions about axiomatic foundations. This is also one reason we mostly dealt with well-orders instead of von Neumann ordinals in the previous sections: the existence of most von Neumann ordinals like $\omega + \omega$, ω^ω, ε_0 and ω_1 depends on replacement, but problems such as the Continuum Hypothesis are still meaningful without ω_1, i.e., we can still ask if there is a bijection between $\mathcal{P}(\omega)$ and the set Ω of isomorphism classes of well-orders on ω.

- Axiom of Regularity. It roughly says there is no infinite descending chain $a_0 \ni a_1 \ni a_2 \ni \cdots$ The precise statement is that every nonempty set X has an element a that contains no element of X. Under Axiom of Choice (or its weak version DC, the Axiom of Dependent Choice) this is equivalent to non-existence of infinite descending chain, but without choice it is stronger. Like replacement, regularity is seldom mentioned in ordinary mathematics, but it could make life easier. For example, if $a = \{a\}$ then a, (a, a), $((a, a), a)$ and the trivial group whose only element is a are all the same set.

> Regularity is often regarded as having no impact on “ordinary mathematics”, but that’s not entirely correct, see [6].

A reasonably complete list of set theory axioms was first proposed by Zermelo, whose initial goal was actually not to provide a foundation for mathematics, but to give a convincing proof of the Well-ordering Principle from the Axiom of Choice. Replacement was not on his list, whose importance was pointed out later by Fraenkel, hence the name Zermelo-Fraenkel set theory.

Probably the only thing that is still unsatisfactory is the vagueness of the word “property”, which appears in both separation and replacement. A famous example of what could go wrong with naive set theory is Berry paradox. Let’s consider the set of all natural numbers that can be defined in under sixty English letters; for example, “the fifth Fermat prime” is a nineteen-letter phrase that defines the number 65537, and “the largest Fermat prime” is not currently known to define a number because we don’t know if there is any Fermat prime larger than 65537. In any case, the set of all such numbers is finite because the English alphabet is finite, so there are only finitely many sentences under sixty letters, and only a handful of them define a number. In particular, there exist numbers not definable in under sixty letters. Now let n be the smallest such number. Then n is defined by “the smallest natural number not definable in under sixty letters”, a definition with fifty-five letters! Other similar paradoxes also confused pioneers in set theory like Cantor and Hilbert.

So we need to make the notion of “property” more precise. People gradually came to the agreement that a “property” should mean a formula in *first order logic*, and that we should replace naive set theory with formal set theory built upon first order logic. Once we do that, self-referential paradoxes such as Berry paradox are no longer paradoxes; rather they become meaningful mathematical theorems.

First order logic is of course not the only option; there are alternatives like type theory, which is what underlies the popular proof assistant Lean. The key point is one needs to have *some* formal framework, so that checking whether a given mathematical definition is valid or whether a given mathematical proof is correct is, at least in theory, algorithmic. First order logic, more specifically the Hilbert style deductive system, provides a relatively simple framework.

2.2 Propositional logic

First order logic is itself an extension of propositional logic, so let’s first discuss propositional logic as a warm-up. The primitive symbols in propositional logic are the following:

- a fixed infinite list of symbols $p_0, p_1, p_2, \ldots$ called propositional atoms
- the negation symbol $\neg$ and the implication symbol $\rightarrow$
- the left parenthesis (and the right parenthesis)

We can use these symbols to form various strings such as

$\neg\neg p_1 \qquad)p_{100}\neg \rightarrow p_2 \qquad (p_0 \rightarrow p_1) \rightarrow (\neg p_1 \rightarrow \neg p_0)$

A formula is a "grammatical" string. More precisely, we use the following rules to determine if a given string counts as a formula:

1. Each propositional atom is a formula.
2. If φ is a formula then so is $(\neg\varphi)$, the concatenation of (, $\neg$, φ and).
3. If φ and ψ are formulas then so is $(\varphi \to \psi)$, the concatenation of (, φ, $\to$, ψ and).
4. These are all the formulas.

Note that according to these rules, technically $(p_0 \to p_1) \to (\neg p_1 \to \neg p_0)$ is not a formula; the correct way to write it is $((p_0 \to p_1) \to ((\neg p_1) \to (\neg p_0)))$. However, this is an acceptable shorthand in semi-formal discussion, as will be explained later. The purpose of parentheses is to avoid ambiguity, as will be shown in the Unique Readability Theorem, although there is also a way of doing without parentheses.

Now if you are a pedantic person like me, you may ask what is the definition of a "string", or what is "concatenation". The answer is they are exactly that. Just use your *common sense*! Yes, eventually mathematics as well as logic is built upon common sense. To avoid infinite regress we need various primitive notions, which are described rather than defined, and it has to be taken for granted that everybody understands the description in the same way. Below when we prove certain basic properties about formulas, we will even use such things as mathematical induction, and we need to assume that to be part of common sense.

Item 4 "these are all the formulas" is admittedly a bit vague. Some texts use the following definition: the set of formulas is the smallest set that contains the propositional atoms and is closed under the rules in items 2 and 3. But then we need to assume the notion of "set" to be part of common sense, which some people disagree with. Alternatively, we can actually describe a recursive algorithm that determines if a given string of symbols is a formula. First we need several facts about formulas. We categorize the following as a *metatheorem*, which is roughly speaking a result about formulas, in contrast to the theorems that will come later, which *are* formulas.

Metatheorem. *(i) Every formula has the same number of left parentheses and right parentheses.*

(ii) If φ is a formula and is not a propositional item, then to the left of any symbol in φ, except the first one, there are strictly more left parentheses than right parentheses.

Proof. (i) We prove this by induction on formation of formulas. This means if we have a certain property about formulas, and if we can prove:

(a) propositional atoms have the property

(b) if φ has the property then so does $(\neg\varphi)$

(c) if both φ and ψ have the property then so does $(\varphi \to \psi)$

then all formulas have that property. The validity of induction on formation of formulas is part of common sense. Some texts use "induction on length of formula" instead, which is in my opinion basically the same.

A propositional atom certainly has the same number of left and right parentheses, because it has none of them. If the theorem is true for φ, then it is also true for $(\neg\varphi)$, which has one more (and one more). Similarly for (c).

(ii) This is again proved by induction on formation of formulas. “To the left” means strictly to the left. For example, in the formula $((p_0 \to p_1) \to ((\neg p_1) \to (\neg p_0)))$ there are two (and one) to the left of the red parenthesis.

Say we want to prove it for $(\neg\varphi)$ assuming it is true for φ. The case when we cut the formula either at $\neg$ or at the first symbol of φ is clear. So choose any symbol in φ except the first one; suppose within φ there are a left parentheses and b right parentheses to the left of the symbol; by assumption we have $a \geq b$ (in fact $a > b$ unless φ is atomic, but that's not needed). Then in $(\neg\varphi)$ there are $a+1$ left parentheses and b right patentheses to the left of the symbol, and $a+1 > b$ as desired.

The proof for $(\varphi \to \psi)$ is similar, but there are more cases to consider. □

Similarly, we have the following fundamental result.

Metatheorem (Unique Readability). *For a formula θ, exactly one of the following holds:*

(a) It is a propositional item.

(b) It is equal to $(\neg\varphi)$ for a unique φ.

(c) It is equal to $(\varphi \to \psi)$ for unique φ and ψ.

Proof sketch. First we prove by formation of formulas that at least one of the three cases happen. Then we need to show that at most one happens. For example, (b) and (c) cannot both happen, because (b) implies the second symbol of θ is $\neg$, so the first symbol of the φ in (c) is $\neg$, but it is not hard to show by induction that a formula never starts with $\neg$.

Finally we need to prove that, e.g., if $(\varphi_1 \to \psi_1)$ and $(\varphi_2 \to \psi_2)$ are the same formula, then φ_1 equals φ_2 and ψ_1 equals ψ_2. This is because using (ii) of the previous metatheorem, it can be shown that, unless φ_1 or ψ_1 is atomic, the middle $\to$ in $(\varphi_1 \to \psi_1)$ is the unique symbol in the formula such that to the left of the symbol, the number of (is exactly one plus the number of). □

Every logical system has its own version of unique readability. Why is $6 \div 2(1+2)$ equal to 1 according to some calculators and 9 according to others? Because they interpret this ambiguous expression differently. We typically use parentheses to clear up ambiguity. Do they actually work? If we put parenthese everywhere such as $(6 \div (2 \cdot (1+2)))$, is there a unique way to evaluate the expression? The positive answer is provided by the unique readability theorem.

From the unique readability theorem, we can extract the following algorithm that determines if a given string of symbols is a formula.

- if the string has only one symbol, then:
  - if the symbol is one of the propositional atoms, it is a formula
  - otherwise, it is not a formula

- if the string has more than one symbol, then:
  - if the string does not start with (and end with), it is not a formula
  - otherwise, delete the two parentheses
    - if the resulting string starts with $\neg$, delete it and check if the remaining string is a formula
    - if the resulting string doesn't start with $\neg$, then search for the first $\rightarrow$ such that there are equal numbers of (and) to the left; if there is no such $\rightarrow$ then it is not a formula; if there is one, which must be unique in case the original string is a formula, delete it and check if the remaining two strings are both formulas

Note that this algorithm is guaranteed to stop and return an answer, because at each step the length of the formula to be checked decreases. We can in fact *define* a formula to be a string for which the above algorithm returns a yes, but that feels a little weird to me.

Although in theory parentheses ensures unique readability, in practice too many parentheses make the formula impossible for human to read. The solution is to agree on some shorthands. For example, we agree that $\neg$ binds stronger than anything else, and also drop the outmost pair of parentheses. Then $((p_0 \rightarrow p_1) \rightarrow ((\neg p_1) \rightarrow (\neg p_0)))$ becomes $(p_0 \rightarrow p_1) \rightarrow (\neg p_1 \rightarrow \neg p_0)$.

We also introduce three new logical connectives $\vee$ (or), $\wedge$ (and), $\leftrightarrow$ (iff), and agree that $\varphi \vee \psi$ is the shorthand for $\neg\varphi \rightarrow \psi$, $\varphi \wedge \psi$ is the shorthand for $\neg(\neg\varphi \vee \neg\psi)$, and $\varphi \leftrightarrow \psi$ is the shorthand for $(\varphi \rightarrow \psi) \wedge (\psi \rightarrow \varphi)$. Under standard convention, $\neg$ binds the strongest, and the next is $\wedge$; further ambiguity is usually handled by parentheses. For example, $\neg p_0 \rightarrow p_3 \wedge p_1$ should be interpreted as $(\neg p_0) \rightarrow (p_3 \wedge p_1)$, and rewritten in primitive symbols it is $((\neg p_0) \rightarrow (\neg((\neg(\neg p_3)) \rightarrow (\neg p_1))))$. We don't count $\vee$, $\wedge$, $\leftrightarrow$ as primitive symbols, as that would make proof by induction on formulas slightly more complicated.

There is a way to avoid ambiguity that is theoretically more elegant than parentheses, known as *Polish prefix notation.* Here our only primitive symbols are the propositional atoms, negation symbol and implication symbol, without parentheses. The rules of formula formation are that if φ and ψ are formulas, then so are $\neg\varphi$ and $\rightarrow \varphi\psi$. One can again prove a unique readability theorem. Roughly speaking, one proves by induction the unique readability together with the following fact: any proper initial segment of a formula is *not* a formula.

Because of how elegant Polish notation is, especially in first order logic, we adopt the following convention: officially we use Polish notation, but for readability we write formulas in the old way. For example, when we write $\neg p_0 \rightarrow p_3 \wedge p_1$ we really mean $\rightarrow \neg p_0 \rightarrow \neg\neg p_3 \neg p_1$.

We haven't said a single word about what formulas are for. They are used to model mathematical reasoning. For example, work of Frey, Serre and Ribet in the 80s showed that the Taniyama-Shimura conjecture implies Fermat's last theorem. The conjecture was proved true by work of several people in the 90s, most notably Wiles, and as a corollary Fermat's last theorem is true. If we call Taniyama-Shimura conjecture p_0 and Fermat's last theorem p_1, the reasoning here is:

from $p_0 \rightarrow p_1$ and p_0, conclude p_1

This is the celebrated *modus ponens.* In Hilbert style propositional logic, this is the only inference rule, and everything can be proven from the axioms together with modus ponens. There are many other proof systems; for example, natural deduction is a system that has a bunch of inference rules but no axioms, which allows for more flexibility and more transparent proofs. However, for results such as Gödel's incompleteness theorems, the efficiency of the proof system doesn't matter, so we choose to work with Hilbert style proof which is slightly easier to defined.

Below are the *axiom schema* in Hilbert style propositional logic.

Pr1 $\varphi \to (\psi \to \varphi)$

Pr2 $(\varphi \to (\psi \to \theta)) \to ((\varphi \to \psi) \to (\varphi \to \theta))$

Pr3 $(\neg\varphi \to \psi) \to ((\neg\varphi \to \neg\psi) \to \varphi)$

What "schema" means is that, for example, every formula of the form $\varphi \to (\psi \to \varphi)$ is an axiom. Examples include

$p_1 \to (p_2 \to p_1)$

$(\neg p_1 \leftrightarrow p_3) \to (p_{100} \to (\neg p_1 \leftrightarrow p_3))$

$(p_1 \to (p_2 \to p_1)) \to (((\neg p_1 \leftrightarrow p_3) \to (p_{100} \to (\neg p_1 \leftrightarrow p_3))) \to (p_1 \to (p_2 \to p_1)))$

and so on and so forth; similar for the other two schema. Thus each axiom scheme gives rise to infinitely many axioms, which can make some people uneasy. But it should at least seem plausible that we can write down an algorithm that checks if a given formula is an instance of Pr1, so everything is algorithmic, and for most purposes algorithmic is as good as finite.

A few words on the axioms. They are chosen based on economy rather than transparency. Pr1 may feel a little weird if you read → as "implies"; I personally like to read it as "if φ is true, then it is true under any circumstance ψ". Or we can rewrite $\varphi \to \psi$ as $\neg\varphi \vee \psi$ (where we are using the intuitive though non-trivial fact that φ is equivalent to $\neg\neg\varphi$), and then Pr1 becomes $\neg\varphi \vee (\neg\psi \vee \varphi)$, whose validity follows from that of $\neg\varphi \vee \varphi$. Pr2 is somewhat like modus ponens: if φ implies both $\psi \to \theta$ and ψ, then it implies θ. Pr3 is law of contradiction: if from $\neg\varphi$ we can prove both ψ and $\neg\psi$, then φ must in fact be true.

A *theory* is just a set of formulas, like $\{p_0, \neg p_1 \to p_3\}$ or $\{p_n : n \text{ is even}\}$. The set is allowed to be infinite. If you want to avoid talking about arbitrary infinite sets, you can focus on theories that are recursive (there is an algorithm that checks if a given formula is in the set), or at least recursively enumerable (there is an algorithm that lists the formulas in the set, possibly with repetition).

Given a theory Γ, a Γ-*proof* consists of a sequence of formulas $\varphi_1, \varphi_2, \ldots, \varphi_n$ such that for each $1 \le i \le n$, one of the following happens:

1. φ_i is in Γ.
2. φ_i is an axiom, i.e., an instance of Pr1, Pr2 or Pr3.
3. (modus ponens) There exist $1 \le j, k < i$ such that φ_k is equal to $\varphi_j \to \varphi_i$

If $\varphi_1, \varphi_2, \ldots, \varphi_n$ is a Γ-proof, we say φ_n is a *theorem* of Γ, or *provable* from Γ. As an example, below is a (surprisingly complicated) proof of $p_0 \to p_0$ from the empty theory.

1	$p_0 \to ((p_0 \to p_0) \to p_0)$	Pr1
2	$(p_0 \to ((p_0 \to p_0) \to p_0)) \to ((p_0 \to (p_0 \to p_0)) \to (p_0 \to p_0))$	Pr2
3	$(p_0 \to (p_0 \to p_0)) \to (p_0 \to p_0)$	1, 2 and MP
4	$p_0 \to (p_0 \to p_0)$	Pr1
5	$p_0 \to p_0$	3, 4 and MP

This is what a 99% rigorous mathematical proof looks like (to get to 100% we need to write everything in our official Polish notation). Of course there is nothing special about the propositional atom p_0; we can replace it everywhere by any other atom, or indeed any other formula φ, and the proof would still be valid. Thus we actually proved the *theorem scheme* $\varphi \to \varphi$.

At least when Γ is finite (or more generally recursive), deciding whether a sequence of formulas is a valid Γ-proof is completely algorithmic. Is there an algorithm deciding whether a given formula φ is a theorem of Γ? This is not immediately clear, since a priori we need to exhaust all the infinitely many proofs to see if φ is a theorem. But the answer turns out to be yes, and uses the method of truth table. For simplicity let's assume Γ to be empty. Then φ is a theorem iff its truth value is 1 under any assignment to the atoms that appear in φ; since only finitely many atoms appear, this is at least in theory a finite procedure. Let's call a formula a *tautology* if its truth value is 1 under any assignment. That every theorem is a tautology can be proved by induction on length of proof. The converse, namely every tautology is a theorem, is known as the *completeness theorem* for propositional logic, and is a baby version of Gödel's completeness theorem for first order logic.

For brevity we won't get into the details of truth table and the completeness theorem for propositional logic, but here's a remark if you have seen the proof of completeness: when we extend the consistent theory to a maximal consistent theory, some kind of infinitary (non-algorithmic) reasoning is necessary. The precise result in reverse mathematics is as follows: the completeness of propositional logic is equivalent to weak Kőnig's lemma over the base system RCA_0.

2.3 First order logic

Hopefully you already got a taste of what formal logic is like, so we will try to be brief in this section on first order logic. Before writing formulas in first order logic, we need to fix a *language*, which consists of three parts: constant symbols, function symbols, and relation symbols. As a first example, the language of set theory has just one binary relation symbol $\in$. When doing set theory, the primitive symbols are:

- a fixed infinite list of *variables* $\mathtt{v}_0, \mathtt{v}_1, \mathtt{v}_2, \ldots$
- logical connectives $\neg$ and $\to$
- quantifier $\forall$ (for all)
- equality symbol $=$
- the non-logical symbol $\in$, which is a binary relation symbol

Note the distinguished font for variables. When we need to talk about unspecified variables, we use English letters such as x, y, z or A, B, C, with or without subscripts. For example, we will say things like "every formula of the form $x = x$ is a logical axiom" or "suppose $x_1, x_2, \ldots, x_m$ are all the free variables of the formula φ".

The formula formation rules are as follows:

1. For any variables x and y, $= xy$ and $\in xy$ are formulas.

   These are called *atomic formulas*. For example, $= \mathtt{v}_0\mathtt{v}_2$ and $\in \mathtt{v}_{100}\mathtt{v}_{100}$ are both formulas. Although officially we are using Polish notation, for readability we actually write $\mathtt{v}_0 = \mathtt{v}_2$ and $\mathtt{v}_{100} \in \mathtt{v}_{100}$ most of the time.

2. If φ and ψ are formulas, so are $\neg\varphi$ and $\to \varphi\psi$.

3. If φ is a formula and x is any variable, then $\forall x\varphi$ is a formula.

4. These are all the formulas.

We also employ various standard shorthands: $\exists x\varphi$ for $\neg\forall x\neg\varphi$, $x \neq y$ for $\neg x = y$, etc.

The flexibility of first order logic lies in the fact that, by changing the language we get different frameworks that are convenient for the study of different subjects. The language of set theory is $\{\in\}$. The language of group theory is $\{\cdot, i, e\}$, where $\cdot$ is a binary function symbol, i is a unary function symbol called inverse, and e is a constant symbol. The language of ring theory is $\{+, \cdot, 0, 1\}$, where 0 and 1 are constant symbols, while $+$ and $\cdot$ are binary function symbols. The language of arithmetic is the language of ring plus a binary relation symbol $<$. There is also the language of plane geometry, which will be discussed later. The logical symbols, namely variables, logical connectives, quantifier and equality symbol are common to all subjects.

The formula formation rules for a general language is more complicated than the case of set theory, which has neither constant nor function symbols. First, we assume every function or relation symbol comes with a positive integer called its *arity*; some texts allows zero-arity function or relation, but it doesn't make that much a difference. We first define when a string of symbols is a *term*:

1. Each variable and constant symbol is a term.

2. If f is a function symbol of arity n and $t_1, \ldots, t_n$ are terms, then $ft_1 \cdots t_n$ is a term. For clarify we often write $f(t_1, \ldots, t_n)$ instead, but officially neither the parentheses nor the comma are primitive symbols.

3. These are all the terms.

For example, $(\mathtt{v}_0 \cdot ((1 + 1) + 1) + \mathtt{v}_1 \cdot \mathtt{v}_2) \cdot 0$ is a term in the language of ring theory. You can try to write it in Polish notation if you want. Be aware that we *cannot* simplify this term to 0, despite that it's called "the language of ring theory". We can only do that if we manage to prove the formula $(\mathtt{v}_0 \cdot ((1 + 1) + 1) + \mathtt{v}_1 \cdot \mathtt{v}_2) \cdot 0 = 0$ from some theory, say the standard ring theory axioms. Also we cannot simplify $((1 + 1) + 1)$ to 3 because 3 is not even a primitive symbol, unless we agree that 3 is a shorthand for $((1 + 1) + 1)$.

Although $\mathtt{v}_0 + \mathtt{v}_1$ looks much more comfortable than $+\mathtt{v}_0\mathtt{v}_1$, the Polish notation is in fact quite natural, considering that we may need to deal with function or relation symbols of arity at least 3. Such symbols don't arise in practice very often, but are theoretically useful.

Now comes the formula formation rules:

1. If t_1 and t_2 are terms, $= t_1t_2$ is a formula.
2. If E is a relation symbol of arity n and $t_1, \ldots, t_n$ are terms, $Et_1 \cdots t_n$ is a formula.
3. If φ and ψ are formulas, so are $\neg\varphi$ and $\rightarrow \varphi\psi$.
4. If φ is a formula and x is any variable, then $\forall x\varphi$ is a formula.
5. These are all the formulas.

So the difference with set theory is that, in set theory the only terms are variables, so the only atomic formulas are $x = y$ and $x \in y$, but in general there are more possibilities due to the existence of function and constant symbols.

As with propositional logic, we have the unique readability theorem.

Metatheorem (Unique readability)**.** *1. If t is a term, it is either a variable, or a constant symbol, or equal to $ft_1 \ldots t_n$ for a unique n-ary function symbol f and unique terms $t_1, \ldots, t_n$.*

2. If θ is a formula, it is either $t_1 = t_2$ for unique terms t_1, t_2, or $Et_1 \ldots t_n$ for unique n-ary relation symbols E and unique terms $t_1, \ldots, t_n$, or $\neg\varphi$ for unique φ, or...(you know what I mean)

A further complication of first order logic compared to propositional logic is the notion of *free* and *bounded* variables. Roughly speaking, a variable x is bounded if it is in the scope of a quantifier, and otherwise free. For example, consider the formula $\exists \mathtt{v}_1\ \mathtt{v}_1 + \mathtt{v}_1 = \mathtt{v}_0$ in the language of ring theory; the variable $\mathtt{v}_0$ is free and $\mathtt{v}_1$ is bounded; this formula reads "$\mathtt{v}_0$ is divisible by two". It is possible for a variable to be both free and bounded, as in $\mathtt{v}_0 \neq 1 \wedge \forall \mathtt{v}_0\ (\mathtt{v}_0 \cdot \mathtt{v}_0) \cdot \mathtt{v}_0 = 1$, so being free or bounded is actually a property of an *occurence* of variable. In the above formula, the first red $\mathtt{v}_0$ is free while the three green $\mathtt{v}_0$'s are bounded. This formula turns out to be equivalent to $\mathtt{v}_0 \neq 1 \wedge \forall \mathtt{v}_1\ (\mathtt{v}_1 \cdot \mathtt{v}_1) \cdot \mathtt{v}_1 = 1$, that is, we can always rename bounded variables. Below is the precise recursive definition of free and bounded occurences:

- For an atomic formula $t_1 = t_2$ or $Et_1 \cdots t_n$, every occurence of every variable is free.
- The free and bounded occurences of variables in $\neg\varphi$ are the same as that of φ. Similar for $\varphi \rightarrow \psi$.
- All occurences of x in $\forall x\varphi$ are bounded. For variables other than x it's the same as that of φ.

Checking whether a certain occurence is free or bounded is algorithmic. A formula without free occurences of variables is called a *sentence*. Intuitively, a formula with free variables represents a property of an element, like $\exists \mathtt{v}_1\ \mathtt{v}_1 + \mathtt{v}_1 = \mathtt{v}_0$, while a sentence describes a property of a structure, such as $\forall \mathtt{v}_0 \forall \mathtt{v}_1 (\mathtt{v}_0 + \mathtt{v}_1 = \mathtt{v}_1 + \mathtt{v}_0)$, which expresses commutativity of the operation $+$.

Fixing a language, we use the convention that a *theory* is simply a set of sentences in that language, although some texts allow arbitrary formulas. Below are some examples of theories.

Group theory Gr: the language is $\{\cdot, i, e\}$, and the axioms are

Gr1 $(x \cdot y) \cdot z = x \cdot (y \cdot z)$

Gr2 $x \cdot e = x \qquad e \cdot x = x$

Gr3 $x \cdot i(x) = e \qquad i(x) \cdot x = e$

Remarks:

1. We just said a theory should be a set of sentences, but these are obviously not sentences. When we present a list of formulas and call them a theory, we are actually refering to the *universal closures* of the formulas. The universal closure of a formula φ whose free variables are $x_1, \ldots, x_n$ is $\forall x_1 \cdots \forall x_n \varphi$. So the "real" group theory is

   Gr1 $\forall x \forall y \forall z \ (x \cdot y) \cdot z = x \cdot (y \cdot z)$

   Gr2 $\forall x \ x \cdot e = x \qquad \forall x \ e \cdot x = x$

   Gr3 $\forall x \ x \cdot i(x) = e \qquad \forall x \ i(x) \cdot x = e$

2. Keep in mind that technically x, y, z are not variables; they are what we use to refer to variables (or variables for variables if you wish). The actual variables are $\mathtt{v}_0, \mathtt{v}_1, \mathtt{v}_2, \ldots$ So how to interpret these formulas? Does $\forall x \ x \cdot e = x$ give rise to infinitely many axioms, one for each choice of variable? That's one way to interpret it. Or you can also think of x as a fixed variable like $\mathtt{v}_0$, because from $\forall \mathtt{v}_0 \ \mathtt{v}_0 \cdot e = \mathtt{v}_0$ it is possible to prove $\forall \mathtt{v}_1 \ \mathtt{v}_1 \cdot e = \mathtt{v}_1$, $\forall \mathtt{v}_2 \ \mathtt{v}_2 \cdot e = \mathtt{v}_2$, etc. From some point on we will forget about the ugly $\mathtt{v}_0, \mathtt{v}_1, \mathtt{v}_2$ and stick to x, y, z, but technically there is a distinction.

3. Of course we usually write x^{-1} instead of $i(x)$. In practice we also use shorthands such as x^3 for $(x \cdot x) \cdot x$. But we do need to make sure our shorthands are "legal", namely they can be rewritten using the primitive symbols. Consider the following 'shorthand"

   $x \neq e \to \forall n \ x^n \neq e$

   that attempts to say "the group is torsion free". How to unravel this into a formula in primitive symbols? Well, you can't. The property of being torsion-free can be expressed using *infinitely* many axioms:

   $x \neq e \to \ x \cdot x \neq e$

   $x \neq e \to \ (x \cdot x) \cdot x \neq e$

   $x \neq e \to \ ((x \cdot x) \cdot x) \cdot x \neq e$

   $\vdots$

   and it can be proved that no single formula is equivalent to all of the above, so it's impossible to rewrite $x \neq e \to \forall n \ x^n \neq e$ as a single formula.

Strict linear order LO: the language has a binary relation symbol $<$, and the axioms are

1. $\neg(x < x)$
2. $x < y \wedge y < z \to x < z$
3. $x < y \vee x = y \vee y < x$

Graph theory G: the language has a binary relation symbol E, and the only axiom is $xEy \to yEx$, which corresponds to the fact that a graph is basically a symmetric binary relation. More precisely, this is the axiom for graphs without multiedge but possibly with loop. To exclude loops we may add the axiom $\neg(xEx)$.

Note that the language of graph theory is really the same as that of strict linear order, as well as set theory, because it doesn't matter how we decide to write or read a binary relation symbols. All that matter are the axioms. We can change all the $\in$ in ZFC axioms to $<$ and read it as "smaller than" from now on, and the math would not change in any way. As a quote often attributed to Hilbert goes, "one must be able to say at all times—instead of points, straight lines, and planes—tables, chairs, and beer mugs".

Peano Arithmetic PA: the language is $\{+, \cdot, 0, 1, <\}$. We will omit the algebraic axioms such as commutativity, associativity and distributivity, and focus on the most important axiom(s), the *induction scheme*:

$\varphi(0) \wedge \forall x(\varphi(x) \to \varphi(x+1)) \to \forall x\varphi(x)$

Here $\varphi(x)$ can be any formula, so the induction scheme is actually an algorithm that generates infinitely many axioms. It's not necessary for there to be free occurences of x in $\varphi(x)$, but that's usually the case. $\varphi(0)$ is the formula obtained from $\varphi(x)$ by *substituting* x with the constant symbol 0, a process that will be defined later, and similarly $\varphi(x+1)$ is obtained from $\varphi(x)$ by substituting x with the term $x+1$.

Recall that we should really take the universal closure of a formula before putting it into a theory. For example, if $\varphi(x)$ is the formula $\exists z\exists w(w < y \wedge x = z \cdot y + w)$, then the corresponding instance of the induction scheme is $\forall y(\varphi(0) \wedge \forall x(\varphi(x) \to \varphi(x+1)) \to \forall x\varphi(x))$, which is used to prove that given a positive number y, any number x can be expressed as $z \cdot y + w$ where $w < y$.

Set theory ZFC: the language consists of a binary relation symbol $\in$; and the axioms are formalizations of the naive set theory axioms. For example, the Axiom of Power Set is

$\forall X\exists Y\forall S(S \in Y \leftrightarrow S \subseteq X)$

where $S \subseteq X$ is a shorthand for $\forall x(x \in S \to x \in X)$. The Axiom of Separation says given any set X and any property we can form a subset Y, and a property is basically a formula, so we actually have the Axiom Scheme of Separation:

$\forall X\exists Y\forall z(z \in Y \leftrightarrow z \in X \wedge \varphi)$

which gives rise to infinitely many axioms. But we need to be a little careful here: the property φ should not mention Y. More precisely, we should require that Y does not have free occurence in φ, otherwise we can get some absurd "axioms" such as $\forall X\exists Y\forall z(z \in Y \leftrightarrow z \in X \wedge z \notin Y)$. The same remark applies to replacement. A side note: Zermelo himself was not a big fan of first order logic, and pursued second order logic instead; it is now widely agreed that first order logic is more appropriate as a foundation, while second order logic is more or less set theory.

Before introducing the proof system, we still need to define substitution. If we have a formula φ, a variable x and a term t, we get a new formula $\varphi[t/x]$ by replacing every free occurence in x by t. It takes some time to define this rigorously by induction first on terms, and then on formulas. It takes even more time to show, e.g., the result $\varphi[t/x]$ is indeed a formula. This would have been a bit easier if we defined a formula not as a string but as a tree.

We call a substitution *appropriate* if every occurence of every variable in t is still free in $\varphi[t/x]$. Here is an example of an inappropriate substitution: if we substitute the free variable x in the formula $\exists y(y \neq x)$ by the term y, we get the obviously problematic formula $\exists y(y \neq y)$. Note that substituting x with itself is always appropriate.

Now we can state the *logical axioms* in first order logic, for a fixed language. There are three groups of axioms: the propositional axioms, the axioms about quantifiers, and the axioms about equality.

Pr1 $\varphi \to (\psi \to \varphi)$

Pr2 $(\varphi \to (\psi \to \theta)) \to ((\varphi \to \psi) \to (\varphi \to \theta))$

Pr3 $(\neg\varphi \to \psi) \to ((\neg\varphi \to \neg\psi) \to \varphi)$

Q1 $\forall x\varphi \to \varphi[t/x]$, where t is a term and the substitution is required to be appropriate.

Q2 $\forall x(\varphi \to \psi) \to (\varphi \to \forall x\psi)$, where it is required that x doesn't have free occurence in φ.

Eq1 $x = x$

Eq2 $x = y \to y = x$

Eq3 $x = y \wedge y = z \to x = z$

Eq4 $x_1 = y_1 \wedge \cdots \wedge x_n = y_n \to fx_1 \cdots x_n = fy_1 \cdots y_n$, for each n-ary function symbol f.

Eq5 $x_1 = y_1 \wedge \cdots \wedge x_n = y_n \wedge Ex_1 \cdots x_n \to Ey_1 \cdots y_n$, for each n-ary relational symbol E.

A few remarks. The three propositional schema are basically the same as before, though of course we now plug in first order formulas instead of propositional formulas. Since subtituting a variable x with itself is appropriate, and $\varphi[x/x]$ is the same as φ, a special case of Q1 is $\forall x\varphi \to \varphi$. One way to understand Q2 is that $\forall x(\varphi \to \psi) \to (\forall x\varphi \to \forall x\psi)$ always holds, and if x doesn't appear freely in φ then $\forall x\varphi$ is the same as φ. Each of Eq1 to Eq5 is a scheme; for example, Eq1 gives rise to infinitely many logical axioms $\mathsf{v}_0 = \mathsf{v}_0$, $\mathsf{v}_1 = \mathsf{v}_1$, $\mathsf{v}_2 = \mathsf{v}_2$, etc.

Fix a language and a theory, namely a set of sentences. The theory is allowed to be infinite, but all the theories that we actually use are recursive, e.g., there are algorithms that check whether a given formula is a PA or ZFC axiom. We can finally define what a proof is; it consists of a sequence of formulas $\varphi_1, \varphi_2, \ldots, \varphi_n$ such that for each $1 \leq i \leq n$, one of the following happens:

1. φ_i is a sentence in the theory.

2. φ_i is an instance of one of the ten logical axiom schema.
3. (modus ponens) There exist $1 \leq j, k < i$ such that φ_k is equal to $\varphi_j \to \varphi_i$.
4. (universal generalization) There exists $j < i$ such that φ_i is $\forall x \varphi_j$ for some variable x.

A formula is said to be a theorem of the theory, or provable from the theory, if it appears as the last formula in some proof, or equivalently if it appears in some proof (it's clear from definition that an initial segment of a proof is itself a proof).

Universal generalization is most often used in the following way. Say we are doing set theory and want to show all sets have a certain property φ, so our goal is to prove $\forall x \varphi$. Suppose we already know that $\forall x \psi$ for another property ψ. *Fix an arbitrary set* x, and after some hard work we manage to use the fact that x has property ψ to deduce that it has property φ. Since the x we chose was arbitrary, we may conclude that in fact $\forall x \varphi$. Universal generalization is used at the last step, and Q1 is used in the form $\forall x \psi \to \psi$ when we fix an x.

Below is a proof of the uniqueness of empty set using Axiom of Extensionality. It's not entirely formal since it uses many shorthands, and we omit the proof of many logical tautologies (theorems of the empty theory). We want to show $\forall x \forall y (\forall z (z \notin x) \wedge \forall z (z \notin y) \to x = y)$. We prove this by fixing arbitrary x, y s.t. $\forall z (z \notin x) \wedge \forall z (z \notin y)$, deduce that $x = y$, and then generalize that. To show $x = y$ we need to apply extensionality, and for that we fix an arbitrary z.

1 $z \notin x \to (z \in x \to z \in y)$	tautology $\varphi \to \varphi \vee \psi$
2 $\forall z (z \notin x \to (z \in x \to z \in y))$	1 and UG
3 $\forall z (z \notin x \to (z \in x \to z \in y)) \to$ $(\forall z (z \notin x) \to \forall z (z \in x \to z \in y))$	tautology $\forall x (\varphi \to \psi) \to (\forall x \varphi \to \forall x \psi)$
4 $\forall z (z \notin x) \to \forall z (z \in x \to z \in y)$	2,3 and MP
5 $\forall z (z \notin y) \to \forall z (z \in y \to z \in x)$	similar to 1-4
6 $\forall z (z \notin x) \wedge \forall z (z \notin y) \to \forall z (z \in x \leftrightarrow z \in y)$	4,5 plus some logic
7 $\forall z (z \in x \leftrightarrow z \in y) \to x = y$	extensionality
8 $\forall z (z \notin x) \wedge \forall z (z \notin y) \to x = y$	6,7 plus some logic
9 $\forall x \forall y (\forall z (z \notin x) \wedge \forall z (z \notin y) \to x = y)$	8 and UG twice

In principle this is what we should write when doing math on paper or computer everyday, but in practice nobody ever does this, unless they are forced to do so by the logic professor, or if they are working with formal proof assistants such as Lean (but even Lean is much more intuitive than this). The proofs we read in most textbooks are written in either "highly abbreviated formal language", or simply in English. But it's good to know that if necessary, one can always translate an English proof into a formal proof, i.e., strings of symbols that follow strict rules, and that the validity of a proof can be checked by an algorithm.

Besides making things more rigorous, a major consequence of formalizing mathematics in first order logic is that we can now study mathematical reasoning itself, and prove results about formulas, proofs, theories, etc. A somewhat silly example we already saw is "every formula has the same number of left parentheses as right parentheses" (assuming one doesn't use Polish notation, and instead includes parentheses as primitive symbols). A less trivial example is unique readability. The deduction theorem of first order logic is another useful metatheorem, which says if T is a theory, φ is a sentence, ψ is a formula provable from $T \cup \{\varphi\}$, then $\varphi \to \psi$ is provable from T. Note that these are fundamentally different from, say, the uniqueness of empty set, which is a ZFC theorem. A ZFC theorem is a formula that appears in some proof; examples include uniqueness of empty set, existence of ω_1, Fermat's last theorem, etc. In contrast, the theorems about parentheses or unique readability are metatheorems, and they are proved using common sense. Here is an example I personally find helpful to understand the distinction: Fermat's last theorem is a ZFC theorem. On the other hand, "Fermat's last theorem can be proved in ZFC" is a (silly) metatheorem; it is proved by exhibiting the ZFC proof. By the way, "Fermat's last theorem can be proved in PA" is a metastatement that is open as of 2024, but widely believed to be true. There is a whole field called *reverse mathematics* devoted to finding the "weakest system" needed to prove a given statement.

Two of the most famous metatheorems are Gödel's first and second incompleteness theorems. A theory is *consistent* if it doesn't prove contradiction, namely a formula of the form $\theta \wedge \neg\theta$. It follows easily from the law of contradiction that if a theory proves a contradiction, then it proves every formula, which isn't very interesting. So when we are working in a theory like PA or ZFC, we would like the theory to be consistent in the first place.

- Gödel's first incompleteness: Consider a theory in some language that is strong enough to do some basic number theory, such as PA or ZFC (but not, e.g., group theory or graph theory). If the theory is consistent, namely it doesn't prove a formula of form $\theta \wedge \neg\theta$, then there is a formula φ such that the theory proves neither φ nor $\neg\varphi$.
- Gödel's second incompleteness: Same assumption as above, but for definiteness let's focus on ZFC. The second incompleteness theorem provides a concrete formula $\mathrm{Con}(\ulcorner\mathsf{ZFC}\urcorner)$ that is not provable in ZFC. Of course there are many things that are not provable in ZFC; what's special about $\mathrm{Con}(\ulcorner\mathsf{ZFC}\urcorner)$ is that it seems to be saying "ZFC is consistent".

These are often summarized as "any formal system has its limits", and "no formal system can prove its own consistency". Towards the end we will explain the incompleteness theorems in more detail.

Since we are talking about axioms and proofs, it's only appropriate to mention synthetic geometry. Euclid's *Elements* was the birthplace of axiomatic reasoning, and for over two milleniums the gold standard of mathematical rigor. Below are Euclid's five postulates of plane geometry.

1. A straight line segment can be drawn joining any two points.
2. Any straight line segment can be extended indefinitely in a straight line.
3. Given any straight line segment, a circle can be drawn having the segment as radius and one endpoint as center.
4. All right angles are congruent.

5. If two lines are drawn which intersect a third in such a way that the sum of the inner angles on one side is less than two right angles, then the two lines inevitably must intersect each other on that side if extended far enough.

By the 19th century people had started to realize these axioms were not perfect. Some of them are too vague, and there are other hidden assumptions used in Euclid's proofs. It was Hilbert who worked out what is effectively a first order theory for plane geometry, making Euclid the gold standard of mathematical rigor again, which was significant since first order logic hadn't been invented at that time. It is more natural to present Hilbert's axioms in *two-sorted* first order logic, which roughly means we have two lists of variables $A_0, A_1, A_2, \dots$ and $l_0, l_1, l_2, \dots$ We refer to A_is as points and the l_is as lines. The language of Hilbert's geometry consists of four relation symbols, which stand for incidence, between-ness, segment congruence and angle congruence respectively.

- The relation symbol I for incidence has point arity 1 and line arity 1, namely in formula formation it can only appear in the form $I(A, l)$, where A is a point variable and l is a line variable. Its intended meaning is "A lies on l".
- The relation symbol for between-ness has point arity 3, namely it is about three points A, B, C. Unable to come up with a nice symbol, we will simply use the English phrase "B is between A and C"; implicitly the three points should lie on the same line.
- The relation symbol for segment congruence has point arity 4, and is written $AB \cong CD$, although more accurately it should be something like $E_{sc}(A, B, C, D)$.
- The relation symbol for angle congruence has point arity 6, and is written $\angle ABC \cong \angle DEF$. Again, technically we should write something like $E_{ac}(A, B, C, D, E, F)$.

Below are some examples of Hilbert's axioms for plane geometry, HG. Althought they are stated in English for clarity, it's straightforward to rewrite them as first order formulas in the above language.

- For any different points A, B there is a unique line such that $I(A, l)$ and $I(B, l)$. This corresponds to Euclid's first postulate.
- For any differnt points A, B there exists a point C such that B is between A and C. In other words the segment AB can be extended, so this is Euclid's second postulate.
- Suppose $\triangle ABC$ is a triangle, which is short for "A, B, C are pairwise different points and they don't lie on the same line". If the line l intersects the segment AB, namely if there is a point between A, B that lies on l, then l must also intersect either BC or AC. This is called Pasch's axiom; it ensures we are doing plane geometry instead of, say, 3D geometry.
- For any line l and any point A not on l, there exists a unique line l' through A which does not intersect l. This is Euclid's fifth.
- If $\triangle ABC$ and $\triangle A'B'C'$ are two triangles, $AB \cong A'B'$, $A'C' \cong A'C'$ and $\angle BAC \cong \angle B'A'C'$, then $BC \cong B'C'$. This is the SAS criterion for congruent triangles.

- It can be proved from the between-ness axioms that a point on the line l divides it into two rays, which gives two directions. For any points A, B, any point P on a line l and any direction on l at P, there exists a unique point Q on l such that $AB \cong PQ$ and the direction from P to Q coincides with the given direction. Roughly speaking, this axiom says we can "move a segment around", which is essential to ruler and compass constructions. There is a similar but more complicated axiom which says we can move an angle around.

There are many books that develop plane geometry using HG. For example, Hartshorne[4] discusses in detail how the proof of each theorem in *Elements* can be modified to be a proof in HG, and much more.

At this point we can also formulate the famous question about Euclid's fifth, or Parallel Axiom: is it provable from the rest of HG? Numerous mathematicians experimented with the negation of Parallel Axiom. Initially they hoped to derive a contradiction, thus proving the Parallel Axiom by law of contradiction. Later, Gauss, Lobachevsky and Bolyai independently concluded that there was no contradiction, and founded the industry known as hyperbolic geometry. But only because we haven't found a contradiction doesn't mean there is no contradiction: maybe it is still lurking somewhere! One had to wait until Beltrami and Klein who provided *models* for hyperbolic geometry, thus firmly establishing its consistency. Building a model is in some sense the only way to prove the consistency of a theory. It is easiest to discuss model theory after we formalize first order logic within some set theory such as ZFC, and this is what we will do in the next section.

2.4 Second formalization, model theory

In this section we are going to *redo everything of the previous section* inside ZFC. We will define formulas and proofs as (essentially) certain finite sequences of natural numbers, and call them *coded formulas* and *coded proofs* (the terminology is non-standard). This is to distinguish them from formulas and proofs of the previous section, which are strings of symbols that we (in principle) write on paper and computer to do mathematics. Similarly, theories such as Gr, PA or HG will be defined as sets of natural numbers, and to distinguish this "second formalization" from the original theories, we denote them as $\ulcorner\mathsf{Gr}\urcorner$, $\ulcorner\mathsf{PA}\urcorner$ and $\ulcorner\mathsf{HG}\urcorner$, and call them *coded theories.* In particular, ZFC itself has a second formalization $\ulcorner\mathsf{ZFC}\urcorner$.

There are two reasons to do this. First, this is essential for the incompleteness theoerms. The first incompleteness theorem says, e.g., there is a sentence φ in the language of set theory that is neither provable nor disprovable from ZFC. Very crudely speaking, the sentence is "I am not provable". There is obviously some self-reference going on here, and that gives some idea of why the second formalization might be useful.

The second reason, as will be explained shortly, is to enable us to do *model theory*, the study of general first order structures. If all we want is a formalization of first order logic, then PA or indeed in much weaker theories like Robinson's Q already suffice, essentially because a finite sequence of natural numbers can be coded into a single natural number, say using prime factorization, and the previous section was mostly concerned with finite strings. This is why the incompleteness theorems apply to PA and Q as well. So why choose ZFC specifically? Because model theory is the study of general structures such as groups, rings, Cartesian plane or Poincaré disk, so we formalize model theory in ZFC, for the same reason why most of mathematics, including group theory, ring theory or plane geometry, is usually formalized in a system like ZFC.

First let's get this straight: group theorists almost never study groups by proving theorems in Gr, because there isn't much one can prove. Gr can prove some basic facts like $\forall x \forall y (x \cdot y = e \to y \cdot x = e)$, which says a left inverse must also be a right inverse, but most basic notions in group theory such as finite groups, simple groups and free groups cannot be expressed using first order formulas, let alone group representation. As another example, consider plane geometry. There is quite a bit one can prove in HG, but still it's not always the most convenient framework in which to study plane geometry. For instance, it's not even clear how to formalize the notion of area in HG, or how to define π. Also, from time to time we would like to introduce coordinates, or even use analysis to solve geometry problems, which are, at least a priori, not available in HG.

So instead of working in the theory Gr, most of the time group theorists just study groups in the usual way, namely using Sylow's theorems, Jordan-Hölder theorem, representation theory, homological algebra, etc. Some of these involve infinite sets, but they can all be formalized in ZFC. Now this doesn't necessarily mean group theorists consciously work in ZFC—most of them are probably not familiar with the specifics of first order logic or ZFC axioms. Some of them might even claim that they don't work in ZFC but rather just reason with common sense. This can lead to some fierce philosophical debate, though, since people have different opinions about potential infinity vs actual infinity, and about the kind of reasoning applicable to infinite sets. Below are selected opinions on this topic.

- Ultra-platonism: The von Neumann universe of sets is something that actually exists outside our physical universe. ZFC is an imperfect approximation to truth in this ideal mathematical universe. Statements such as continuum hypothesis or projective determinacy have definite truth values in this universe.
- Platonism: The set of all natural numbers exists. Its power set may or may not exist, but in any case it's perfectly fine to talk about particular sets of natural numbers, such as the set of all primes, the set of all Fermat primes, the set of all indices of Turing machines that halt, etc.
- Finitism: There is no such thing as the collection of all natural numbers as a whole, but we can interpret statements such as twin-prime conjecture or unsolvability of halting problem in terms of potential infinity, so they are meaningful. PA is a valid system to reason about natural numbers.
- Strict-finitism: Only computable things exist. "The set of all primes" is no more than a way to talk about the algorithm that lists all the primes. A strict-finitist rejects PA, but may be ok with PRA (primitive recursive arithmetic), a very weak fragment of PA.
- Ultra-finitism: Only functions computable in polynomial time exist. Exponentiation is not a total function.

There are many other opinions not listed here, such as various schools of constructivism, because I don't really know how they work.

Back to group theory. Does it make sense to talk about a presentation of a group with infinitely many relations? If so, must the relations be given by an algorithm? Does it make sense to consider the set of all 2 by 2 real invertible matrices and put a topology on it? The answers to these questions vary depending on one's philosophical view towards infinity.

Formalism provides a compromise approach to these questions—just ignore them. We don't try to answer whether infinite groups exist "in real life", whatever that means, or how we can reason about them. Instead, we fix a theory such as ZFC and define the abstract notion of a group, which is a nonempty set together with a binary operation that satisfies certain properties, and then go on to prove ZFC theorems about groups. How this relates to "groups in real life" is another question that we simply ignore. Just shut up and prove! As we have seen, writing ZFC proofs boils down to simple manipulations of finite strings of symbols, and with a bit of care, checking whether a given string is a valid formula/proof can be done in polynomial time. So once we manage to prove, e.g., the classification of finite simple groups in ZFC, possibly with the help of infinitary methods like representation theory or homological algebra, even the ultra-finitists have to agree that this is indeed a proof of the sentence "every finite simple group is of one of the forms..." in ZFC, although they may dismiss this as word game with no implication for groups in real life.

That explains why one would want to study groups in ZFC rather than Gr, and the same goes with ring theory, field theory, plane geometry, etc. Model theory is the study of general first order sturctures—a nonempty set equipped with some functions and relations with finite arity, which include all the above examples. Before discussing models in ZFC, we need to set up first order logic within ZFC, aka the second formalization. Model theory is not directly relevant for understanding the incompleteness theorems, but it is at the heart of the *completeness theorem*, another fundamental result of Gödel (**Caution:** the word "completeness" in completeness theorem has a totally different meaning than in incompleteness theorems). Besides, model theory is an important field of logic that not only is fundamental in modern set theory, but also has applications to other subjects such as real and complex algebraic geometry, and is definitely worth an introduction.

If the above philosophical discussion seems confusing, I highly recommend [9, Chapter III], especially section III.2, where the necessity of developing formal logic twice is discussed.

Now we turn to the real math. Let's first recall how basic mathematical objects are defined in set theory. We start with the smallest limit ordinal ω, which is guaranteed to exist by Axiom of Infinity. Elements of ω are called natural numbers; recall that due to how we define ordinals, each natural number is also the set of smaller natural numbers. An ordered pair is a set of form $(a, b) = \{\{a\}, \{a, b\}\}$. A function f is a set of ordered pairs such that whenever (a, b) and (a, c) both belong to f, we have $b = c$. Note that from the set f we can read its domain and range, for example $\mathrm{dom}(f) = \{a \in \bigcup\bigcup f : \exists b\ (a, b) \in f\}$, but we cannot read the codomain.

A remark on "definition". In the previous section we only discussed how to write proofs in ZFC, such as the uniqueness of empty set. We didn't say how to define new things. So what does it mean exactly to define a function f as a set of ordered pairs that satisfies such and such? It means from now on, we use the phrase "f is a function" to abbreviate the following formula with free variable f:

$$\forall x(x \in f \to \exists a \exists b\ x = (a, b)) \land \forall a \forall b \forall c((a, b) \in f \land (a, c) \in f \to b = c)$$

where $x = (a, b)$ and $(a, b) \in f$ are abbreviations of some other formulas that you can try to figure out, based on our discussion about ordered pair. Different people might come up with slightly different formulas, but that's not a problem since any two reasonable choices should be provably equivalent in ZFC.

We use $f(a) = b$ to mean $(a, b) \in f$. The empty set $\varnothing$ is a function. If A is any set, then the set of all functions from $\varnothing$ to A is $\{\varnothing\}$; if A is nonempty, then there is no function from A to $\varnothing$. A set is finite if there is a bijection between the set and some natural number. Unless otherwise specified, by sequence we always mean a function f with $\mathrm{dom}(f) = n$ for some natural number $n \in \omega$, and n is called the length of the sequence; note that the empty set is a sequence of length 0. We can define such things as addition and multiplication on ω, or more generally we can prove the *recursion theorem*; a basic form of recursion theorem is that for any natural number C and any function $g : \omega \times \omega \to \omega$, there exists a unique function $f : \omega \to \omega$ satisfying $f(0) = C$ and $f(n + 1) = g(n, f(n))$ for all $n \in \omega$. The uniqueness of f is proved by mathematical induction, which holds because of the definition of ω as the smallest limit ordinal. The existence is proved using finite approximation: first prove by induction that for each natural number N, there exists a (unique) f_N such that $\mathrm{dom}(f_N) = N + 1 = \{0, 1, \dots, N\}$ and f_N satisfies the desired property for all $n < N$. Then f can be defined by $f(n) = f_n(n)$. The idea here is very similar to how we compare two well-orders in Theorem 3. Concatenation of two or arbitrary number of fintite strings can also be defined this way.

Now we begin the process of second formalization, starting with the notion of a language. Recall that a language consists of some function symbols, relation symbols and constant symbols, and each function/relation symbol comes with a positive integer called its arity. Thus in its most generality:

Definition. A *coded language* is a 5-tuple $\mathcal{L} = (\mathcal{F}, \mathcal{R}, \mathcal{C}, a_{\mathcal{F}}, a_{\mathcal{R}})$, where $\mathcal{F}, \mathcal{R}, \mathcal{C}$ are disjoint sets, and $a_{\mathcal{F}} : \mathcal{F} \to \mathbb{N}_+$ and $a_{\mathcal{R}} : \mathcal{R} \to \mathbb{N}_+$ are functions.

The three sets $\mathcal{F}, \mathcal{R}, \mathcal{C}$ can be empty, finite, countable or uncounable. Elements of $\mathcal{F}, \mathcal{R}, \mathcal{C}$ are called function symbols, relation symbols, and constant symbols respectively (you can add the adjective coded if you want). For $f \in \mathcal{F}$, the number $a_{\mathcal{F}}(f)$ is called the arity of f, and similarly for $a_{\mathcal{R}}(P)$. We often abuse notation and write for example $\mathcal{L} = \{+, \cdot, 0, 1\}$ (this is the coded language of ring), with the understanding that $\cdot$ and $+$ are binary function symbols, while 0 and 1 are constant symbols. And we don't specify what the sets $+, \cdot, 0, 1$ are exactly: just pick you favourite four sets, like $0, 1, 2, 3$, or $\{100\}, \{\{100\}\}, \{\{\{100\}\}\}, \{\{\{\{100\}\}\}\}$.

Some or all of $\mathcal{F}, \mathcal{R}, \mathcal{C}$ could be empty. They could also be infinite or even uncountable. But when they are countable, there is no loss of generality in assuming all the symbols are natural numbers.

Definition. For $\mathcal{L}$ a coded language, an *$\mathcal{L}$-structure* $\mathcal{M}$ consists of a nonempty set M, called the domain or underlying set, together with the following data:

(i) for each function symbol f of arity n, a function $f^{\mathcal{M}} : M^n \to M$;

(ii) for each relation symbol R of arity m, a set $R^{\mathcal{M}} \subseteq M^m$;

(iii) for each constant symbol c, an element $c^{\mathcal{M}} \in M$.

These are called *interpretations* of the corresponding symbols. To be more pedantic, an $\mathcal{L}$-structure is a tuple $\mathcal{M} = (M, \mathcal{F}^{\mathcal{M}}, \mathcal{R}^{\mathcal{M}}, \mathcal{C}^{\mathcal{M}})$, where $\mathcal{F}^{\mathcal{M}}$ is a function with domain $\mathcal{F}$ that sends f to some $f^{\mathcal{M}} : M^{a_{\mathcal{F}}(f)} \to M$, etc. We often abuse notation and call $(M, f, \dots E, \dots, c \dots)$ a structure.

Examples:

1. If $\mathcal{F}, \mathcal{R}, \mathcal{C}$ are all empty, then an $\mathcal{L}$-structure is the same as a nonempty set. If only $\mathcal{C}$ is nonempty, then an $\mathcal{L}$-structure is basically a nonempty set with marked elements.

2. If $\mathcal{L}$ consists of a single binary function symbol f, then an $\mathcal{L}$-structure is just a set equipped with a binary operation, also called a magma.

3. If $\mathcal{L}$ consists of a single binary relation symbol E, a structure $\mathcal{M} = (M, E^{\mathcal{M}})$ can be identified with a set $E^{\mathcal{M}} \subseteq M^2$, and is the same as a directed graph on M, possibly with loops, but without multiedge.

4. The following construction is often useful in model theory: start with an $\mathcal{L}$-structure $\mathcal{M}$ with underlying set M. Define a new language $\mathcal{L}_M$ by adding a constant symbol c_a for every $a \in M$. Now M is naturally an $\mathcal{L}_M$-structure since we can interpret c_a as a. This is one place where allowing uncountable language could be useful.

Next we would like to define, e.g., when a magma $(M, \cdot)$ satisfies the commutativity axiom. This should mean $x \cdot y = y \cdot x$ for all $x, y \in M$. If we try to unravel this into a formula in the language of set theory, we get something like

$\forall x \in M \forall y \in M \forall z \in M(((x, y), z) \in \cdot \leftrightarrow ((y, x), z) \in \cdot)$

Here $\cdot$ is a function from M^2 to M, so a typical element of $\cdot$ looks like $((x, y), z)$. Also, $\forall x \in M\varphi$ is a shorthand for $\forall x(x \in M \to \varphi)$. It is more cumbersome, though still doable, to write down a formula that expresses associativity. How to express infinitely many things at once, such as the property of being torsion free? For this particular example we can say something like "for each $x \in M$ nontrivial, consider the unique $f : \omega \to M$ satisfying $f(0) = x$ and $f(n + 1) = f(n) \cdot x$, then $f(n)$ is nontrivial for every n". A more systematic and elegant way is to first define coded formulas and theories, and then define by recursion when an $\mathcal{L}$-structure $\mathcal{M}$ satisfies a coded formula under an assignment.

It is important to maintain a clear distinction between formulas, proofs, theories and their coded counterparts, which are certain sets defined in ZFC. One common solution is to put corners around everything that is coded. Thus fix once and for all four sets $\ulcorner\forall\urcorner, \ulcorner\to\urcorner, \ulcorner\neg\urcorner, \ulcorner=\urcorner$, as well as countably many sets $\ulcorner \mathtt{v}_0\urcorner, \ulcorner \mathtt{v}_1\urcorner, \ulcorner \mathtt{v}_2\urcorner, \ldots$ For definiteness, we may let

$\ulcorner\forall\urcorner = 0 \qquad \ulcorner\to\urcorner = 1 \qquad \ulcorner\neg\urcorner = 2 \qquad \ulcorner=\urcorner = 3 \qquad \ulcorner \mathtt{v}_i\urcorner = 5^i$

Technically, we should redefine a language so that the function, relation and constant symbols are required to be distinct from the above logical symbols. Now for a language $\mathcal{L}$, we define when a finite sequence of symbols is a coded $\mathcal{L}$-terms as follows: $\langle c\rangle$, the length-one sequence whose entry is c, is a coded term; $\langle\ulcorner \mathtt{v}_i\urcorner\rangle$ is a coded term for each i; if f is an n-ary function symbol and the sequences $t_1, \ldots, t_n$ are coded terms, then so is the concatenated sequence $ft_1 \cdots t_n$; and these are all the coded terms. Since we are working in ZFC, we can formalize "these are all the coded terms" by defining the set of coded terms to be the smallest set with certain properties. Coded formulas are defined similarly using Polish notation: if t_1 and t_2 are coded terms, then $\ulcorner=\urcorner t_1 t_2$ is a coded formula; if E is a relation symbol with arity n and $t_1, \ldots, t_n$ are coded terms, then $Et_1 \ldots t_n$ is a coded formula; if φ and ψ are coded formulas, so are $\ulcorner\to\urcorner\varphi\psi$, $\ulcorner\neg\urcorner\varphi$ and $\ulcorner\forall\urcorner\ulcorner \mathtt{v}_i\urcorner\varphi$ for each i.

We can now define free occurence, bounded occurence, substitution, coded proof, coded theory and prove their basic properties. We can also prove the unique readability theorem for coded terms and formulas. Be aware that these are now theorems proved using ZFC axioms, instead of metatheorems proved using common sense. We are not going to prove the unique readability theorem, since of course it takes too much time to write the formal proof, but if we write the proof in English it would look identical to the previous proof. Here is a good place to introduce the following convenient thesis.

Formalization Thesis: If we have a metatheorem about formulas/proofs/theories, its formalized counterpart about coded formulas/coded proofs/coded theories is a ZFC theorem. In fact they are PA theorems, if we did the second formalization in PA.

This is only a thesis, i.e., a guideline, instead of a precise statement, because of the vagueness of the word "counterpart". We don't attempt to design a precise procedure to obtain the formalized counterpart of a given metatheorem like unique readability, but empirically there is only one natural way to formalize a metatheorem, up to minor variations, and the formalization is always a ZFC theorem. This is justified by the fact that results like unique readability were proven using common sense, so their formalization should be provable in ZFC and PA, which are much stronger than common sense.

Now comes the central definitions in model theory. Suppose we have an $\mathcal{L}$-structure $\mathcal{M}$ and an coded $\mathcal{L}$-formula φ. An *assignment suitable for* φ is a function s such that:

(i) $\mathrm{dom}(s)$ is a finite subset of the variable set $\{\ulcorner \mathtt{v}_0 \urcorner, \ulcorner \mathtt{v}_1 \urcorner, \ulcorner \mathtt{v}_2 \urcorner \dots\}$. Some texts let $\mathrm{dom}(s)$ be the whole variable set, which is actually cleaner, but I think when setting up the basics, it is a good idea to avoid infinite sets whenever possible.

(ii) $\mathrm{ran}(s) \subseteq M$. Recall that M is the underlying set of $\mathcal{M}$.

(iii) $\mathrm{dom}(s)$ at least contains all the free variables of φ.

An assignment suitable for a term t is defined in the same way, except (iii) is changed to "$\mathrm{dom}(s)$ at least contains all the variables that appear in t". We are going to define what it means for $\mathcal{M}$ to *satisfy* φ w.r.t. a suitable assignment s. As with anything about formulas, this is defined by induction on formulas, and the unique readability theorem plays a crucial role.

To handle atomic formulas, we need a separate induction that defines the value $t(s)$ of a term t under an assignment s. This is intuitively very clear: we just plug the value $s(x)$ into x for each variable x. Say we are dealing with the language of arithmetic and we have the term $t = 1 + (x_0 + x_1) \cdot x_4$ with three free variables, and s assigns a to x_0, b to x_1 and c to x_4. Then $t(s) = 1^{\mathcal{M}} +^{\mathcal{M}} (a +^{\mathcal{M}} b) \cdot^{\mathcal{M}} c$. The precise definition is as follows.

- If $t = \langle c \rangle$ for some constant symbol c, then note that any assignment is suitable, and we define $t(s) = c^{\mathcal{M}}$.
- If $t = \langle \ulcorner \mathtt{v}_i \urcorner \rangle$ for some i, then s is suitable for t iff $\ulcorner \mathtt{v}_i \urcorner \in \mathrm{dom}(s)$, in which case we define $t(s) = s(\ulcorner \mathtt{v}_i \urcorner)$.
- If $t = f t_1 \cdots t_n$ for some unique function symbol f and coded terms $t_1, \dots, t_n$ (we need unique readability here), then s is suitable for t iff it is suitable for each t_k, and we define $t(s) = f^{\mathcal{M}}(t_1(s), \dots, t_n(s))$.

To be more pedantic, we need to show there is a unique function defined on the set

$\{(t, s) : t$ is a term and s is an assignment suitable for $t\}$

that satisfies the above items. This is an example of recursion on well-founded relation, and is very similar to the recursion theorem on ω mentioned earlier. We omit the details.

Finally, satisfaction of formula under suitable assignment is defined inductively as follows:

- $\ulcorner = \urcorner t_1 t_2$ is satisfied under s iff $t_1(s) = t_2(s)$.
- $Et_1 \dots t_n$ is satisfied under s iff $(t_1(s), \dots t_n(s)) \in E^{\mathcal{M}}$.
- $\ulcorner \neg \urcorner \varphi$ is satisfied under s iff φ is not satisfied under s.
- $\ulcorner \rightarrow \urcorner \varphi\psi$ is satisfied under s iff either φ is not satisfied, or ψ is satisfied.
- If s is suitable for the coded formula $\ulcorner \forall \urcorner \ulcorner \mathtt{v}_i \urcorner \varphi$, for each $a \in M$ consider the assignment $s[a/x]$ defined as follows: if $\ulcorner \mathtt{v}_i \urcorner \notin \mathrm{dom}(s)$ then add $\ulcorner \mathtt{v}_i \urcorner$ to the domain of s and send it to a; if $\ulcorner \mathtt{v}_i \urcorner \in \mathrm{dom}(s)$ then change the value of s at $\ulcorner \mathtt{v}_i \urcorner$ to a. Note that $s[a/x]$ is suitable for φ. We define that $\ulcorner \forall \urcorner \ulcorner \mathtt{v}_i \urcorner \varphi$ is satisfied under s iff φ is satisfied under $s[a/x]$ for any $a \in M$.

Again, formally this is done by recursion on a certain well-founded relation. We usually abbreviate "$\mathcal{M}$ satisfies φ under the suitable assignment $x_1 \mapsto a_1, \dots, x_n \mapsto a_n$" by $\mathcal{M} \models \varphi(a_1, \dots, a_n)$. If φ is a sentence then the empty assignment (or any assignment) is suitable, so one can unambiguously talk about whether $\mathcal{M}$ satisfies a given sentence φ.

For readability, we write, e.g., $\ulcorner \forall x\ x = x \urcorner$ instead of the more accurate $\ulcorner \forall \urcorner \ulcorner \mathtt{v}_0 \urcorner \ulcorner = \urcorner \ulcorner \mathtt{v}_0 \urcorner \ulcorner \mathtt{v}_0 \urcorner$. A fact which is almost never stated in textbooks, but implicitly used all the time, is that for every "particular" sentence φ, the notion of satisfaction defined above is equivalent to the "hands-on" definition. For example, consider a structure $(M, \cdot)$; we have seen previously that the natural way to express commutativity is

$\forall x \in M \forall y \in M\ x \cdot y = y \cdot x$

or more accurately

$\forall x \in M \forall y \in M \forall z \in M(((x, y), z) \in \cdot \leftrightarrow ((y, x), z) \in \cdot)$

On the other hand, let φ be the coded formula $\ulcorner \forall x \forall y\ x \cdot y = y \cdot x \urcorner$. It can be proved in ZFC that the above is equivalent to $M \models \varphi$. which is not immediately clear since satisfaction is defined by a pretty complicated recursion.

If T is a coded $\mathcal{L}$-theory, namely a set of coded $\mathcal{L}$-sentences, and $\mathcal{M} \models \varphi$ for every $\varphi \in T$, then $\mathcal{M}$ is called a *model* of T. The coded theory T is *satisfiable* if it has a model.

Examples:

1. The Gr from previous section can easily be formalized as $\ulcorner \mathsf{Gr} \urcorner$, a coded theory in the language $\{\cdot, i, e\}$. There is also a version of $\ulcorner \mathsf{Gr} \urcorner$ in the language $\{\cdot, e\}$:

   $\ulcorner \forall x \forall y \forall z\ (x \cdot y) \cdot z = x \cdot (y \cdot z) \urcorner$

   $\ulcorner \forall x\ e \cdot x = x \urcorner \qquad \ulcorner \forall x\ x \cdot e = x \urcorner$

   $\ulcorner \forall x \exists y\ x \cdot y = e \wedge y \cdot x = e \urcorner$

Models of $\ulcorner \mathsf{Gr} \urcorner$ include $(\mathbb{Z}, +, 0)$ and $(\mathbb{Q}^*, \cdot, 1)$, but $(\mathbb{Z}, +, 1)$ is not a model, because if we interpret e as 1 then the second axiom is not satisfied, although $(\mathbb{Z}, +)$ is a model of the version of $\ulcorner \mathsf{Gr} \urcorner$ formulated in $\{\cdot\}$.

One can ask whether $\ulcorner \mathsf{Gr} \urcorner$ is a *complete*, that is whether every coded formula in the language $\{\cdot, e\}$ is either provable or disprovable from $\ulcorner \mathsf{Gr} \urcorner$ (of course it also makes sense to ask this about Gr, but throughout this section we work in ZFC). The answer is no, and a concrete example is the abelian axiom $\ulcorner \forall x \forall y (x \cdot y = y \cdot x) \urcorner$, which is neither provable nor disprovable; we also say that it is *independent* from $\ulcorner \mathsf{Gr} \urcorner$. This is not that obvious. Of course we know ZFC proves there are both abelian and non-abelian groups, in other words there are models of $\ulcorner \mathsf{Gr} \urcorner$ that satisfy $\ulcorner \forall x \forall y (x \cdot y = y \cdot x) \urcorner$ and models that don't. To conclude the independence of the formula $\ulcorner \forall x \forall y (x \cdot y = y \cdot x) \urcorner$, one needs the soundness theorem, which is the "easy" direction of the completeness theorem discussed below.

2. If there is a single binary relation symbol E, then a structure is the same as a directed graph with loops but no multiedge. But this is also the language for the theory $\ulcorner \mathsf{ZFC} \urcorner$, the formalization of ZFC. So a model of $\ulcorner \mathsf{ZFC} \urcorner$ is essentially just a very special kind of directed graph. We won't explicitly define what $\ulcorner \mathsf{ZFC} \urcorner$ is: either convince yourself that it is completely staightforward (though tedious) to describe the set of all formulas that are instances of the separation and replacement schema, or check [12, Chapter 3 §5]; this is the only source I know that spells out every single detail about formalization.

3. If $\mathcal{M}$ is a structure and φ is a coded formula with free variables, then $\{(a_1, \dots, a_n) \in M^n : \mathcal{M} \models \varphi(a_1, \dots, a_n)\}$ is called the set defined by the formula $\varphi(x_1, \dots, x_n)$. For example, in group theory $\ulcorner \forall y (x \cdot y = y \wedge y \cdot x = y) \urcorner$ defines the set of elements x that commute with all elements, i.e., the center. A slogan is that "model theory is the study of definable sets".

   Number theory can be viewed as the study of the structure $(\mathbb{N}, +, \cdot, 0, 1)$. The definable sets in this structure are extremely complicated. In contrast, the definable sets of both $(\mathbb{R}, +, \cdot, 0, 1)$ and $(\mathbb{C}, +, \cdot, 0, 1)$ are fairly tame—turns out any such set is actually definable by a quantifier-free formula, so a definable set is basically a Boolean combination of zero sets of polynomials, also known as a constructible set in algebraic geometry. In particular, $\mathbb{N}$ is not definable in $\mathbb{R}$ or $\mathbb{C}$. Number theory is on the "wild side" of model theory while geometry is on the "tame side". Perhaps that is why number theorists invented arithmetic geometry, to bring their research closer to the tame side.

4. A property of $\mathcal{L}$-structures is *axiomatizable* if there is a theory T such that the structures with that property are exactly the models of T. For example, being abelian is axiomatizable for groups, and being loopless is axiomatizable for directed graphs.

   There are many more non-axiomatizable statements than axiomatizable ones. Let's try to axiomatize infinite groups. For each *fixed* number n one can write down a formula φ_n saying "there are at least n elements"; for $n = 2$ it is $\ulcorner \exists x \exists y \; x \neq y \urcorner$; for $n = 3$ it is $\ulcorner \exists x \exists y \exists z (x \neq y \wedge y \neq z \wedge x \neq z) \urcorner$. Of course we are not allowed to write something like $\forall n \varphi_n$, but at least we can express infinitude by letting T to be the all the φ_n along with the basic group axioms, so infinite groups are axiomatizable. What about finite groups? Turns out this is non-axiomatizable; indeed, whenever you have a property that is axiomatizable but not finitely axiomatizable, the opposite property will be non-axiomatizable, by the *Compactness Theorem* of first order logic. Many other important properties such as being finitely generated, solvable, free, or p-group are not expressible in the formal language. This can be viewed as saying first order logic is extremely weak, but this weakness turns out to be a feature: there are many more tools available in the study of first order logic compared to other logics, because it's "just strong enough".

5. We didn't define what a 2-sorted first order structure is, but it's just a simple variation. The domain would consist of two parts, and some function or relation symbol might take elements from both parts as input or output. This way one can define models of Hilbert's axioms.

   Although synthetic geometry (axiomatic method) is often contrasted with analytic geometry (coordinate method), by work of Hilbert they are in a sense equivalent: any model of Hilbert's axioms is of form F^2, where F is a Pythagorean field, namely an totally ordered field where $\sqrt{1+a^2}$ exists for any a. On one hand, if F is a Pythagorean field it is easy to interpret points, lines, betweenness and congruence in F^2, making it into a model of Hilbert's axioms. Conversely, given a model Π of Hilbert's axioms, one can define "the field of segments" F, and show that Π is isomorphic to F^2 as structures in the language for plane geometry. So by the Completeness Theorem (a statement is provable iff it's true in all models), synthetic geometry is more or less the same as analytic geometry, which in turn is the same as the study of ordered field. For details see [4].

6. A (small) category can also be regarded as a 2-sorted structure in a suitable language (or a 1-sorted structure if we identify objects with identity morphisms). The two sorts are objects and morphisms. There are two functions that give us the source and target of a morphism, as well as a ternary relation that reads "f is the composition of g and h"; composition is not total (not any two morphisms can be composed), so we have to interpret composition as a relation, unless we revise the foundation of model theory to allow partial function, which doesn't seem popular.

   Universal properties in category theory are related to definability when we view the category as a structure. For example, "A is an initial object if for any object B there is a unique morphism from A to B" can clearly be written as a formula in the formal language of category. Same for monomorphism, projective object, product, etc. A functor between two categories is the same as a model-theoretic homomorphism.

We didn't define coded proofs in detail, but you can imagine what they are like: for a coded theory T, a coded T-*proof* is a sequence of coded formulas such that each one is either an element of T, or one of the logical axioms, or obtained from previous coded formulas by one of the two inference rules, modus ponens and universal generalization. A coded theory T is *consistent*, denoted $\mathrm{Con}(T)$, if it doesn't prove contradiction, or equivalently if it doesn't prove all coded formulas. $\mathrm{Con}(T)$ itself is a formula in ZFC with one free variable T. If we want, we can unravel $\mathrm{Con}(T)$ into a very long formula that involves only the logical symbols and membership symbol $\in$.

A priori there isn't any relation between consistency and satisfiability, i.e., existence of moedl. Consistency is about proofs, which are strings of symbols, and satisfiability is about models, which are concrete structures. It is thus highly non-trivial that they turn out to be equivalent.

Theorem (Gödel's completeness theorem)**.** *Con(T) if and only if T is satisfiable.*

The word "completeness" refers to the harmony between proof and model in first order logic, and is unrelated to the completeness of a coded theory T, which says every sentence can be decided from T. The if direction is called soundness theorem, and is relatively straightforward though very painful. As with anything we show by induction, that if T proves a sentence φ then any model of T must satisfy φ; actually to facilitate induction we need to consider a general formula φ whose free variables are *exactly* $x_1, \dots, x_n$, and show that if T proves φ then any model satisfies $\ulcorner\forall x_1 \cdots \forall x_n\urcorner\varphi$, the universal closure of φ. In particular, if T has a model at all then it cannot prove a contradiction, because a model never satisfies contradiction. For the only if direction, modern textbooks don't use Gödel's original proof, but a simpler one found by Henkin. The idea is to first appropriately enlarge the language and the theory, and then consider the model "generated" by the constant symbols.

Besides being conceptually appealing, the completeness theorem has many applications in logic. An important corollary is the compactness theorem. First observe that since a proof has finitely many lines, it can only invoke finitely many axioms. Thus if T proves a sentence φ, then a finite subset of T already proves it. Specializing to the case when φ is a contradiction, this tells us if any finite subset of T is consistent, then T itself is consistent. So far it is an almost trivial observation. But when we combine this with completeness theorem, we get the following:

Theorem (Compactness theorem)**.** *If any finite subset of T is satisfiable, then T itself is satisfiable.*

The name compactness comes from the fact that the version of this theorem in propositional logic is essentially saying "the product space 2^X is compact for any set X".

Compactness is the first choice when we want to show that some property is not axiomatizable or finitely axiomatizable. We have seen that for each n, there is a sentence φ_n saying "there are at least n elements" (this works in any language), so if we let $T = \{\varphi_n\}_n$ then a model of T is the same as an infinite structure in that language. Can we find a T' whose models are exactly the finite structures? No, because if such a T' exists, then $T' \cup \{\varphi_n\}_{n \le N}$ would be satisfiable for each N (any structure of size at least N is a model), so by compactness $T' \cup T$ is satisfiable, which is absurd. The same argument shows that if a theory has finite models of arbitrarily large size then it must have infinite models, so for example it's impossible to axiomatize the finite fields; nevertheless there is the interesting notion of pseudo-finite fields, which are fields that "look like" finite fields from the perspective of model theory. Another corollary is that we cannot find a single sentence θ that is equivalent to T, namely its models are exactly the infinite structures, since otherwise $\neg\theta$ would axiomatize the finite structures. Compactness also has some neat applications in algebra, such as Nullstellensatz and Ax-Grothendieck theorem. For Nullstellensatz see [10], which essentially uses compactness to show that $\ulcorner\mathsf{ACF}_p\urcorner$, the coded theory of algebraically closed fields in characteristic p (including the case $p = 0$), is model-complete, which means every embedding between two structures is an elementary embedding. Model theory texts typically derive this from the fact that $\ulcorner\mathsf{ACF}_p\urcorner$ admits quantifier elimination, which takes some more work.

2.5 Incompleteness theorems

We leave ZFC and come back to common sense. A theory (without the adjective coded!) is complete if for every sentence φ, either φ or its negation is provable. Unfortunately we use φ to refer to both an arbitrary formula and an arbitrary coded formula; hopefully the intended meaning will be clear from the context. If φ is a sentence that is neither provable nor disprovable, it is *independent* over the theory.

All the usual theories like PA and ZFC are recursive, i.e., there is an algorithm checking whether a given formula is a ZFC axiom. But the incompletness theorems are natural stated (and proved) for the more general recursively enumerable theories, meaning there is an algorithm that runs forever and lists all the axioms, possibly with repetition.

Metatheorem (Gödel's first incompleteness). *If a theory is consistent, recursively enumerable and "at least as strong as Robinson's* Q*", then it is incomplete.*

Robinson's Q is roughly speaking PA with the induction scheme deleted, but with one special instance of the induction scheme added back: every number is either zero or something plus one. "At least as strong as Robinson's Q" is to ensure basic reasoning about natural numbers goes through, which is certainly the case for PA and ZFC, so they are incomplete.

Proof sketch. We will try to convey the main ideas of the proof. For definiteness suppose the theory is ZFC. First we assign *Gödel numbers* to every formula (here without the adjective coded): each formula is a finite string of symbols, so if we assign a natural number to each symbol, we can view a formula as a finite string of numbers, which can in turn be converted into a single number, say using prime factorization. Denote by $\ulcorner\varphi\urcorner$ the Gödel number of the formula φ.

Within ZFC we carry out the second formalization, as done in the first half of the previous section. Crucially, the procedure of Gödel numbering can also be formalized in ZFC, which assigns a natural number (an element of ω) to each coded formula. We write down a formula $\mathrm{Prov}(n)$ in the language of ZFC with one free variable n that says "the coded formula with Gödel number n is provable from $\ulcorner\mathsf{ZFC}\urcorner$".

We want to find a sentence φ that is neither provable nor disprovable. A natural candidate is the sentence "I am not provable", but this is not a formal sentence. The idea is that we want a φ such that $\varphi \leftrightarrow \neg\mathrm{Prov}(\ulcorner\varphi\urcorner)$ is a sentence provable from ZFC. Assume we have found such a φ. If φ were provable, then by the Formalization Thesis, ZFC proves $\mathrm{Prov}(\ulcorner\varphi\urcorner)$, so by the above equivalence ZFC proves $\neg\varphi$, contradicting the assumption that ZFC is consistent. If φ were disprovable, then by the equivalence we know ZFC proves $\mathrm{Prov}(\ulcorner\varphi\urcorner)$. So we have a false metastatement "φ is provable" whose formalization is proved by ZFC. As strange as it may look, this is not quite a contradiction, and Gödel's orginal argument gets stuck here—he needed to add a mild extra assumption that ZFC is so-called ω-consistent for the above argument to go through. But he was fine with that, because the fact that φ is not disprovable is not needed in the proof of second incompleteness, which was the main goal of his paper. Later, Rosser found a better sentence φ than Gödel's, a formalized version of "if I am provable, then my negation has a shorter proof", for which the above argument goes through without the extra assumption of ω-consistency. Hence the first incompleteness theorem is also known as Gödel-Rosser theorem.

We still need to find a φ such that $\varphi \leftrightarrow \neg\mathrm{Prov}(\ulcorner\varphi\urcorner)$ is provable in ZFC. This follows from a slightly more general result:

Metatheorem (Diagonalization Lemma). *For any formula $A(n)$ with a free variable n, there is a sentence φ such that $\varphi \leftrightarrow A(\ulcorner\varphi\urcorner)$ is provable.*

Here comes the most clever part in the whole proof: for every formula $\varphi(x)$ with exactly one free variable, we consider its *self-substitution* $\varphi(\ulcorner\varphi\urcorner)$, namely the sentence obtained from $\varphi(x)$ by plugging its Gödel number into itself; this can again be formalized in ZFC. Then we consider the formula $\psi(n)$ saying "$A(m)$ where m is (the Gödel number of) the self-substitution of the coded formula with Gödel number n". Finally, $\varphi = \psi(\ulcorner\psi\urcorner)$ is the desired sentence, because by design,

$\psi(\ulcorner\theta\urcorner) \leftrightarrow A(\ulcorner\theta(\ulcorner\theta\urcorner)\urcorner)$

for any formula θ with one free variable, so letting θ be ψ we get

$\psi(\ulcorner\psi\urcorner) \leftrightarrow A(\ulcorner\psi(\ulcorner\psi\urcorner)\urcorner)$ □

The second incompleteness theorem follows relatively easily from the first, modulo another big application of the Formalization Thesis.

Metatheorem (Gödel's second incompleteness)**.** *If a theory is consistent, recursively enumerable and "at least as strong as Robinson's* Q*", then it cannot prove (the formalized version of) its own consistency.*

Proof sketch. Again, for definiteness we work with ZFC. We need to show that ZFC cannot prove $\mathrm{Con}(\ulcorner\mathsf{ZFC}\urcorner)$. Let φ be the Gödel sentence (not the Gödel-Rosser sentence), so we can prove $\varphi \leftrightarrow \neg\mathrm{Prov}(\ulcorner\varphi\urcorner)$ in ZFC, and

if ZFC is consistent then it cannot prove φ.

By the Formalization Thesis applied to the above metastatement, ZFC proves that

$\mathrm{Con}(\ulcorner\mathsf{ZFC}\urcorner) \to \neg\mathrm{Prov}(\ulcorner\varphi\urcorner)$

Thus if ZFC could prove $\mathrm{Con}(\ulcorner\mathsf{ZFC}\urcorner)$, it would be able to prove $\neg\mathrm{Prov}(\ulcorner\varphi\urcorner)$, or equivalently φ. But φ is not provable, a contradiction. □

Remarks:

1. We assume the theory to be consistent to start with, because an inconsistent theory proves everything and hence is uninteresting.

2. These two results are more accurately called the first and second incompleteness metatheorems. Of course, by the Formalization Thesis they have counterparts that are ZFC or indeed PA theorems. The first incompleteness theorem says, e.g., the coded theory $\ulcorner\mathsf{ZFC}\urcorner$ is incomplete, and the second incompleteness says $\ulcorner\mathsf{ZFC}\urcorner$ cannot prove $\ulcorner\mathrm{Con}(\ulcorner\mathsf{ZFC}\urcorner)\urcorner$, an ugly coded formula that you wouldn't ever want to write down in full detail.

3. The incompleteness metatheorems do not apply to theories like Gr or HG, because they cannot talk about natural numbers. This can be seen from the fact that despite they are both incomplete, they both have recursive complete extensions. Gr can be extended to the theory of torsion-free divisible abelian groups, which are more or less the same as vector spaces over $\mathbb{Q}$, and this is known to be complete by model-theoretic method. Similarly, HG can be extended to a system called Tarski geometry, which is more or less the same as the theory of real closed fields, another well known complete theory.

4. An informal corollary of the second incompleteness metatheorem is that "humans cannot prove the consistency of ZFC", because if we could prove it using any foreseeable method, by the

Formalization Thesis we should be able to prove $\mathrm{Con}(\ulcorner\mathsf{ZFC}\urcorner)$ in any strong enough system, including ZFC itself, and that would contradict second incompleteness.

Of course, all this is based on the assumption that ZFC is consistent. Could it actually be inconsistent? Again, many people have tried to derive a contradiction, but nobody has succeeded. The takeaway of second incompleteness is that if we cannot find a contradiction, then all we can do is to take the consistency of ZFC for granted.

5. There is a neat proof of first incompleteness metatheorem that avoids the delicate self-substitution argument, and instead uses the unsolvability of halting problem: there can't exist an algorithm that checks whether a given algorithm eventually stops. Suppose for contradiction that ZFC is complete. Then it would be decidable, namely there is an algorithm that checks, for a given sentence φ in the language of set theory, which one of φ and $\neg\varphi$ is proven by ZFC. We simply enumerating all the ZFC proofs one by one, which is possible because it is a recursive theory. We stop once we find either a proof of φ or $\neg\varphi$, which must happen by assumption. Now given an algorithm Γ, we convert it into a sentence φ_Γ in the language of set theory that expresses "the algorithm Γ halts", and plug it into the previous algorithm to decide whether it is true. So the completeness of ZFC leads to an algorithm for solving halting problem, a contradiction.

   Unlike Gödel's proof, this proof doesn't give a concrete sentence that is independent. This is actually quite a big deal, because the Gödel sentence is used in the proof of second incompleteness theorem. Another unsatisfactory issue with the above proof is that in the last step, one needs to assume that ZFC (or PA, or whatever theory we are dealing with) is *sound*, which means it only proves true statements, so that if ZFC proves that an algorithm halts then it actually halts. Soundness is even stronger than ω-consistency, the hypothesis used by Gödel in first incompleteness, which is in turn stronger than mere consistency. But it is possible to push the halting-problem-based proof a little further and drop the soundness hypothesis; see [11] for a friendly presentation.

   One may wonder how could ZFC or PA ever prove false statements. A useful example to keep in mind is that as a corollary to second incompleteness, if ZFC is consistent, then so is the theory $\mathsf{ZFC} + \neg\mathrm{Con}(\ulcorner\mathsf{ZFC}\urcorner)$, and this theory proves the false statement $\neg\mathrm{Con}(\ulcorner\mathsf{ZFC}\urcorner)$. The halting theorem proof does work to prove in ZFC that $\ulcorner\mathsf{PA}\urcorner$ is incomplete, because in ZFC one can prove the soundness of $\ulcorner\mathsf{PA}\urcorner$ by appealing to the standard model $\mathbb{N}$ of natural numbers.

6. To appreciate the assumption of recursively enumerability, consider in ZFC the coded theory known as *true arithmetic*, which is the set of all coded sentences φ in the language of arithmetic satisfied by the standard model $\mathbb{N}$ of PA. This theory is not recursively enumerable, as can be shown using the halting problem, similar to the above argument. And it is complete by definition, thus showing the necessity of the assumption in the first incompleteness theorem.

The first incompleteness metatheorem is mainly instructional in set theory, and are seldom directly used. It tells us there is a sentence independent from ZFC, in fact a very concrete one, the Gödel-Rosser sentence, but it's kind of artificial. There is a plethora of natural statements in set theory that were later found to be independent. We have already discussed the most famous one, the Continuum Hypothesis, which is independent from ZFC. Axiom of Choice is independent from the rest of ZFC. The existence of an inaccessible cardinal, if consistent, is also independent from

ZFC. One proves for example the independence of CH by constructing a (class) model of ZFC that satisfies CH and another model that doesn't, rather than some clever grammatical trick as in the proof of first incompleteness.

The second incompleteness metatheorem plays a more important role. PA cannot prove its own consistency, so Hilbert's original dream of proving the consistency of mathematics by finitary method is doomed. But ZFC certainly can prove Con(⌜PA⌝) because $\mathbb{N}$ is a model of ⌜PA⌝. Now to establish Con(⌜ZFC⌝) we need to go beyond ZFC. Denote by IC the formal statement "there exists an inaccessible cardinal". It turns out ZFC + IC proves the existence of a model of ⌜ZFC⌝: if κ is an inaccessible cardinal, then the von Neumann hierarchy up to κ, denoted V_κ, can be shown to be a model of ⌜ZFC⌝. So by the easy direction of the completeness theorem, ZFC + IC proves Con(⌜ZFC⌝). We say that ZFC + IC has higher *consistency strength* than ZFC. In fact it has strictly higher consistency strength, namely ZFC does not prove Con(⌜ZFC + IC⌝). To see why, note that of course ZFC (or just PA) proves Con(⌜ZFC + IC⌝) $\to$ Con(⌜ZFC⌝), just because the first theory contains the second. So if ZFC could prove Con(⌜ZFC + IC⌝), it would be able to prove Con(⌜ZFC⌝), which is exactly what the second incompleteness forbids. In fact a slightly stronger statement is true: ZFC cannot prove the implication Con(⌜ZFC⌝) $\to$ Con(⌜ZFC + IC⌝), as a corollary of second incompleteness applied to the theory ZFC + Con(⌜ZFC⌝). Set theorists have studied tons of extensions of ZFC, and empirically all natural theories seem to be linearly ordered (in fact well-ordered) by consistency strength, like a ladder to heaven. Moreover, within each "equiconsistency class" there is usually a representative of the form ZFC+some large cardinal.

Let's briefly discuss the metamathematical issues surrounding the proof of CH's independence over ZFC. As with most consistency proofs we use models. Of course we don't try to prove the existence of a model of ⌜ZFC + CH⌝ in ZFC, because it's impossible by second incompleteness, and it's not really what we want. What Gödel's inner model method proves is the metatheorem that if ZFC is consistent, then so is ZFC + CH. Essentially Gödel provides a complicated formula with one free variable x, abbreviated $x \in L$, such that whenever φ is a theorem of ZFC + CH, the *relativization* φ^L is a theorem of ZFC. Here φ^L is defined by induction on formulas: $(x = y)^L$ is $x = y$, $(x \in y)^L$ is $x \in y$, $(\varphi \to \psi)^L$ is $\varphi^L \to \psi^L$, $(\neg\varphi)^L$ is $\neg\varphi^L$, and $(\forall x \varphi)^L$ is $\forall x(x \in L \to \varphi)$. In particular ZFC + CH cannot prove a contradiction $\theta \wedge \neg\theta$, since otherwise ZFC would prove the contradiction $\theta^L \wedge \neg\theta^L$. Although the collection of all sets x satisfying $x \in L$ is a proper class instead of a set, in many ways we can think of L as a "model" of ZFC + CH, so it is called a class model. It turns out Gödel's provably cannot work to show the consistency of ZFC + ¬CH; one has to go "outside" of V instead, which is absurd if taken literally. To solve this, most expositions of Cohen's forcing method starts with a (set) model M of ⌜ZFC⌝ and enlarges it to a model $M[G]$ of ⌜ZFC + ¬CH⌝. So technically this proves Con(⌜ZFC⌝) $\to$ Con(⌜ZFC + ¬CH⌝) in ZFC. Then it takes some effort to explain why the argument actually shows the metatheorem "if ZFC is consistent then so is ZFC + ¬CH". An alternative approach, known as the Boolean-valued model method, is closer in spirit to Gödel's inner model and proves the metatheorem more directly.

Why should we believe in the consistency of ZFC or even PA in the first place? PA might be less controversial, since we can "imagine" a model $\mathbb{N}$. Not all people agree, though, and many have actually worked hard to disprove PA, the most famous attempt due to Edward Nelson, who even claimed a proof of inconsistency of PRA. However, a gap was found in Nelson's proof. There are ways to prove the consistency of PA without using a model (none of them is formalizable in PA itself, by second incompleteness), such as Gentzen's consistency proof, which shows that roughly speaking, consistency of PA follows from the well-foundedness of the ordinal ε_0, which sounds quite intuitive to most people with knowledge about ordinals (again there are exceptions, such as Vladimir Voevodsky). This started the topic known as *ordinal analysis* in reverse mathematics.

So it seems extremely difficult, if not impossible, to disprove PA. Let's be less ambitious and try to disprove ZFC. That's also quite difficult. In fact nobody has ever disproved inaccessible cardinal, or any of the standard large cardinal notions (below Kunen's inconsistency result), let alone ZFC. So is there any way to help us believe in the consistency of ZFC, similar to the consistency proofs for PA? For now there doesn't seem to be anything close to a Gentzen style argument. Some people claim that we can imagine a model in our mind: von Neumann's universe V, built up in a hierarchy of length Ord. This is a common argument in favor of ZFC, but at the same time very controversial. Personally I find the argument by the variant of Solovay-Shelah theorem more appealing: to believe in the consistency of ZFC, it suffices to believe that one cannot find a non-measurable set of reals in bounded Zermelo set theory with dependent choice.

We end with a random remark that Cantor and Gödel, two of the most important figures in logic, were both devout Christians, and both died miserably.

References

[1] Alexander S. Kechris (1997) *Set theory and uniqueness for trigonometric series*

[2] Walter Rudin (1990) *Set Theory: An Offspring of Analysis*, lecture given at UW-Milwaukee, https://www.youtube.com/watch?v=hBcWRZMP6xs

[3] Irving Kaplansky (1972) *Set Theory and Metric Spaces*, AMS Chelsea Publishing

[4] Robin Hartshorne (2000) *Geometry: Euclid and Beyond*, Springer New York, NY

[5] Alex Simpson (2013-09-10) *How strong is "all sets are Lebesgue Measurable" in weaker contexts than ZF?*, MathOverflow, https://mathoverflow.net/q/141801

[6] Brian Pinsky (2023-01-22) *Consequences of foundation/regularity in ordinary mathematics (over ZF–AF)?*, MathOverflow, https://mathoverflow.net/q/439141

[7] Nikolai Luzin (1925) *Sur les ensembles profectifs de M. Henri Lebesgue.* Comptes Rendus de l'Académie des Sciences, Paris, 180, 1572-1574

[8] Akihiro Kanamori (1995) *The Emergence of Descriptive Set Theory*

[9] Kenneth Kunen (2009) *The Foundations of Mathematics*, Studies in Logic Series Volume 19

[10] Nir Avni *Algebraic Geometry*, lecture notes taken by Akhil Mathew, http://math.uchicago.edu/ amathew/AGnotes.pdf

[11] Scott Aaronson (2011-07-21) *Rosser's Theorem via Turing machines*, blog post, https://scottaaronson.blog/?p=710

[12] Frank R. Drake (1974) *Set Theory*, Studies in Logic and the Foundations of Mathematics, Volume 76, Elsevier