MY NOTES
ON TRANSITIVITY

James Holland

math.rutgers.edu/~jch258

jch258@math.rutgers.edu

Start Date: 2019-10-03
End Date: 2019-12-11

Contents

The Axioms of ZFC

1. (Extensionality, Ext) two sets are equal whenever they have the same members:
$$\forall x\ \forall y\ (x = y \leftrightarrow \forall v\ (v \in x \leftrightarrow v \in y)).$$
2. (Empty set) there is a set $\emptyset$ with no members: $\exists z \forall x (x \notin z)$.
3. (Comprehension, Comp) for each x, and for each $\mathtt{FOL}(\in)$-formula $\varphi(v, \vec{w})$, $\{v \in x : \varphi(v, \vec{w})\}$ exists:
$$\forall w_0 \cdots \forall w_n\ \forall x\ \exists z\ \forall v\ (v \in z \leftrightarrow v \in x \wedge \varphi(v, \vec{w})).$$
4. (Pairing, Pair) for any two sets x and y, the pair $\{x, y\}$ exists: $\forall x\ \forall y\ \exists z \forall v\ (v \in z \leftrightarrow (v = x \vee v = y))$.
5. (Union, Union) for any family of sets F, there is a set containing the elements of all of those sets:
$$\forall F\ \exists U\ \forall v\ (v \in U \leftrightarrow \exists x (x \in F \wedge v \in x)).$$
6. (Foundation, Found) for each x, there is a $\in$-minimal element of x, meaning a member $y \in x$ with no $z \in y$ being in x:
$$\forall x \exists y (y \in x \wedge \forall z (z \in y \to z \notin x)).$$
7. (Infinity, Inf) an infinite set exists: $\exists N (\emptyset \in N \wedge \forall x (x \in N \to x \cup \{x\} \in N))$.
8. (Replacement, Rep) the image of a function over a set is a set: for each $\mathtt{FOL}(\in)$-formula φ,
$$\forall w_0 \cdots \forall w_n \forall D \left(\forall x (x \in D \to \exists! y\ \varphi(x, y, \vec{w})) \to \exists R (y \in R \leftrightarrow \exists x (x \in D \wedge \varphi(x, y, \vec{w})))\right).$$
9. (Powerset, P) for each x, $\mathcal{P}(x)$ exists: $\forall x\ \exists P\ \forall v\ (v \in P \leftrightarrow \forall y\ (y \in v \to y \in x))$.
10. (Choice, AC) for any family of non-empty family of non-empty, disjoint sets F, there is a set C which has chosen one element from each $z \in F$:
$$\forall F\ (\emptyset \notin F \wedge \forall x, y \in F\ (x \cap y = \emptyset) \to \exists C\ \forall x \in F\ \exists! y\ (y \in x \cap C).$$

Variant Axioms and Axiom Systems

i. (Σ_n-Comprehension, Σ_n-Comp) for each x, and for each Σ_n-formula $\varphi(v, \vec{w})$, $\{v \in x : \varphi(v, \vec{w})\}$ exists.

ii. (Collection, Coll) there is a range for a relation with over a given domain: for each $\mathtt{FOL}(\in)$-formula φ,
$$\forall w_0 \cdots \forall w_n \forall D \left(\forall x \in D\ \exists y\ \varphi(x, y, \vec{w})) \to \exists R\ \forall x \in D\ \exists y \in R\ \varphi(x, y, \vec{w})\right).$$

iii. (Σ_n-Collection, Σ_n-Coll) Coll (just) holds for Σ_n-formulas.

iv. (Zorn's Lemma, AC_Z) For every (non-empty) poset $\langle A, \preccurlyeq\rangle$, if every chain is bounded in A, then A has a $\preccurlyeq$-maximal element.

v. (AC_P) If F is a non-empty set of non-empty sets, then $\prod_{x \in F} x$ is non-empty.

vi. (AC_C) Every set is bijective with an ordinal.

vii. (AC_W) Every set has a well-order.

viii. (Dependent choice, DC) for $R \subseteq X \times X$, if $\forall x \in X\ \exists y \in X\ (x\ R\ y)$ then there is a sequence $\langle x_n : n \in \omega\rangle$ such that $x_n\ R\ x_{n+1}$ for all $n \in \omega$.

ix. For every x, y, the cartesian product $x \times y$ exists.

With these axioms, we have the following theories:

- BST consists of (1)–(6) plus (ix).
- ZF^- consists of (1)–(8) plus (ii).
- $\mathsf{ZFC}^- = \mathsf{ZF}^- + \mathsf{AC}$ consists of (1)–(8) plus (10) and (ii).
- $\mathsf{ZF} = \mathsf{ZF}^- + \mathsf{P}$ consists of (1)–(9).
- $\mathsf{ZFC} = \mathsf{ZF} + \mathsf{AC}$ consists of (1)–(10).

Chapter I. Transitivity

In general, a relation R is called "transitive" iff for all relevant x, y, and z, if $x \; R \; y$ and $y \; R \; z$, then $x \; R \; z$. Classic examples of this include equality, and the ordering on the reals $<$, among many others. In the context of set theory, a collection is called transitive iff the membership relation $\in$ is a transitive relation on it: X is transitive iff $z \in y \in X$ implies $z \in X$. On its own, this property seems unmotivated or perhaps useless, but it plays one of the most fundamental roles in set theory. To hint at an important connection, consider the actual universe of sets, denoted here by V. Note that this collection is transitive since any set $y \in \mathrm{V}$ has all of its members in V too—in this case, just by virtue of existing. So transitive collections are the first candidates for models of set theory: they are an attempt to approximate V.

Transitive sets can also approximate V in truth. In particular, certain statements are called "absolute" between certain structures when all those structures agree on them. So we will be interested in absoluteness between transitive structures, as this tells us information about V.

§ 0 A. Philosophy

We begin with the philosophically basic notion of a collection: we take it as immediate that things exist, and that we can consider collections of things as abstract objects. It is in this sense that we mean that these collections "exist", and hence we can take collections of collections, and so forth. We in the real world can then reason about these collections and their properties. The simplest examples of this kind of reasoning comes from Venn diagrams, like the one pictured below.

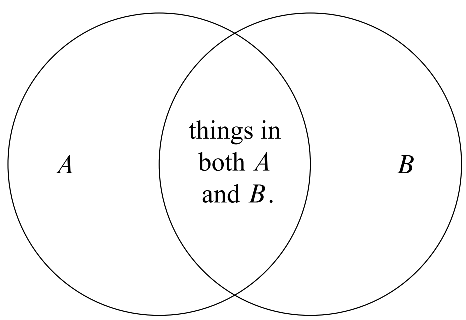

0 A • 1. Figure: Example of a Venn-diagram

The first concept we then define is the collection of all sets, the actual set theoretic universe. More precisely, we begin with the sets that are *hereditarily* sets, meaning for each x, every member of x is a set, and all of their members are too, and so on.

0 A • 2. Definition

The *universe of sets* is the structure $\mathbf{V} = \langle \mathrm{V}, \in \rangle$, where V consists of all (and only) sets, and $\in$ denotes membership.

What exactly should this universe look like? Intuitively, we start with a set with no elements: the empty set, $\emptyset$. Then, we can take the set of just this object $\{\emptyset\}$. Now we have two objects, and we can take collections of these: $\emptyset$, $\{\emptyset\}$, $\{\{\emptyset\}\}$, and $\{\emptyset, \{\emptyset\}\}$. And we can continue this iterative formation of sets. This iterative conception is at the heart of modern set theory, and I hope to further motivate why it is true through the chapter. But first, we must acknowledge how we will do this: indeed, the question of our base level axioms come into question.

We will go through the chapter introducing principles or *axioms* which are generally seen as statements true of $\mathbf{V}$ beyond any doubt. Now we are interested not just in what is true of $\mathbf{V}$, but also what we can *prove* about $\mathbf{V}$ from these axioms.

In particular, it is not immediately obvious whether certain statements are true or false. And indeed, if we are to argue that we *cannot* prove nor disprove them, we need to have agreed upon, intuitively true axioms about V. It is, of course, an open question whether our list of axioms exhausts all intuitively true statements about V. But given the power of the axioms we present, it is difficult to find simple principles that are intuitively obvious but independent of the other axioms.

How we state these axioms is important if we are to have a precise notion of proof—or the lack thereof. Usually, mathematical statements are written in a codified version of natural language, where notation replaces the normal words of English, Russian, or whatever other language. This is no different for us, but we rely on notation even more to ensure that we can carry out everything in a formal system using just basic reasoning about finite objects (namely formulas).

Note that there is an important distinction in logic between the reasoning we use in the real world and the reasoning a certain subject allows. For example, we in the real world have the ability to conclude $a = c$ from $a = b$ and $b = c$. However if we consider only the sentential/propositional connectives there—'and' and 'implies'—we cannot make the same conclusions. From the perspective of propositional logic, "$a = b$ and $a = c$" is no different from $A \wedge B$ where A and B are two completely unrelated propositions: the logic no longer considers the meaning of equality, only the meaning of these sentential connectives. To distinguish the two logics, the reasoning we use in the real world is called the *meta-theory* whereas the reasoning a certain subject (like propositional logic) allows is called the *logic system*[i].

The reasoning of a logic system is entirely formal, following from strings of symbols, but with the proper setup, it can characterize a portion of the meta-theory, like the simple example of propositional logic. The more complicated example of first-order logic is where we will state our axioms. This is both because it has the expressive power needed to present the axioms, and because there are a great deal of important results related to it, as we will see in the first section. To give a more cultural reason, first-order logic is not the only logic system one can use to study mathematics, but most other logic systems can be reformulated in terms of set theory with arguments that take place in first order logic. In fact, second order logic is sometimes called "set theory in sheep's clothing". Generally, first-order logic is the framework in which the results of set theory are given, and results *about* set theory are generally about it in this framework.

Now a priori, there's no guarantee that the world behaves in accordance with the axioms of ZFC (the standard axioms of set theory). The axioms are taken to be intuitively obvious, but in fact, we would need to reject them as part of the meta-theory if it turned out that this system were inconsistent. Furthermore, constructions allowed by ZFC like $\mathbb{R}$ and $\mathbb{N}$ can be called into question if we reject certain axioms like the existence of $\mathbb{N}$. How then do we regard such statements as "$|\mathbb{R}| > |\mathbb{N}|$"? Is this a meta-theoretic fact, or is this better regarded as a formula of first-order logic following from certain axioms? There are a few ways to address these concerns. Two major positions are presented here.

One stance is a purely formalist one. This view will neglect to say anything substantial about the meta-theory, taking only the most basic algorithmic reasoning needed for the study of logic for granted. The formalist approach then doesn't connect the reality of the meta-theory with results of axioms like ZFC, and it in some sense ignores whether the theories we study are important at all. No commitments are made for whether the natural numbers $\mathbb{N}$ exist or whether a statement like "$|\mathbb{R}| > |\mathbb{N}|$" has any meaning in the meta-theory. But the formalist will deny that $|\mathbb{R}| > |\mathbb{N}|$ has any semantic value. Instead, the formalist will view the statements about $\mathbb{N}$ or $\mathbb{R}$, for example, as merely symbols algorithmically changed from other symbols collectively called ZFC[ii]. So the results of theories in the logic system are seen purely as symbolic manipulation with no connection to the meta-theory. At best, a formalist will say the symbols in the logic system can be translated into arguments in the meta-theory where they should have been given in the first place. At worst, a formalist will say the symbols are devoid of content.

Another stance is a platonist or realist one. This view will hold that the results of axioms like ZFC in the logic system do characterize a fragment of the meta-theory—in particular, V. Not only is there a standard meaning of the statement "$|\mathbb{R}| > |\mathbb{N}|$", but there is an actual fact of the matter, and we can learn such facts through study of theories in such logic

[i] Elsewhere in the literature, you might see other words like the *object language*, *proof system*, or perhaps just *logic* to refer to logic system or how it's written.

[ii] or whatever other foundation they are studied in.

systems. By and large, a platonist stance is held by mathematicians that want to claim that their conclusions are actually true and not merely derived from playing with symbols. Indeed most of mathematics is not done through symbolic algorithms like truth tables but instead through intuitions and clever constructions. That said, a platonist stance isn't strictly necessary, since often meta-theoretic arguments can be reformulated as symbolic ones and vice-versa. In this way the two stances are not incompatible.

This work will take more of a platonist stance. More precisely, ZFC is held as a collection of true statements about **V**, and this is used to reason about ZFC as presented symbolically.

Section 1. *Logic and Model Theory*

I will assume some familiarity with the basics of first order logic, particularly the meaning of $\vdash$ and $\vDash$, as well as the associated concepts of formulas, sentences, theories, and models or structures. The reader intending to skip this section should just be aware of two things: for a signature σ, the language of first order logic is written $\mathtt{FOL}(\sigma)$; and if we have a formula with parameters, we say it is a $\mathtt{FOLp}(\sigma)$ or FOLp-formula. Rather than spend an inordinate amount of time giving the details of first-order logic, the reader is referred to any introductory logic text, like [?Endertonlogic]. So instead an overview is given as a reminder.

There are two parts to introduce first-order logic as with almost any logic system. Firstly there is a syntactic component ruling what can be said. Secondly there is a semantic component that gives meaning to these formulas. This separation is similar to the separation between the spelling of words in English, and the meaning of them. There are a number of steps in this introduction. Continuing the natural language analogy, we need to

1. determine the alphabet we're using;
2. determine how to spell words with this alphabet;
3. determine how to "reason" with these words;
4. determine the meaning of these words; and
5. connect spelling with meaning.

§ 1 A. The alphabet and its formulas

To start, the alphabet of first order logic is better regarded as a collection of alphabets that are all variations on a simpler alphabet. In particular, they all share the so called *logical symbols* given below that allow us to make basic formulas that are statements of equality and inequality: $x \neq y$, $v_3 = v_{10}$. From these basic statements—so called *atomic formulas*—we can build up larger formulas using simple rules. For things already determined to be formulas, we can connect them using formula connectives, or quantify them over some variable.

So for φ and ψ already formulas, $(\varphi \wedge \psi)$, $\neg\varphi$, $\exists x\ \varphi$, and so forth are all formulas too. This allows us to build formulas like $\exists x\ \exists y\ (\neg x = y)$ and $(x = x \wedge \neg x = x)$. We cannot, however, make ordinary mathematical statements like $x = y + z$ or $\exists z(z \cdot x^{-1} \leq x + f(z, y))$ yet. To make such statements we need a bigger alphabet. In particular, we have the concept of a signature or vocabulary to expand the logical symbols with non-logical symbols.

1 A • 1. Definition

σ is called a *signature* if and only if σ is a collection of symbols that are divided into constant symbols, relation symbols, and function symbols with the corresponding number of arguments.
The first-order language associated with σ is denoted $\mathtt{FOL}(\sigma)$.

For example, those familiar with some algebra will know that rings and fields generally use a signature of just function symbols: $\{+, \cdot, 0, 1\}$. This expands the signature usually used with groups: $\{\cdot, 1\}$. Partial orders and graphs will use only relation symbols for the order and the edges. Most importantly for us, set theory uses the signature with only one element $\{\in\}$[iii].

The rules for forming formulas change very little from when there were just logical symbols. Essentially, one just needs to respect the number of arguments for the relation and function symbols. So if f is a function symbol with two arguments, you can't write $f(x, y, z)$ or $f(t)$. The same applies to relation symbols. Building terms t_1, ..., t_n by composing function symbols and variables, we can let relations holding between terms—strings of the form $R(t_1, \cdots, t_n)$ or $t_1 = t_2$—be the basic building blocks of formulas. Then we can build the rest of the formulas in the same way as before.

Now we remark that often formulas are written in short-hand, meaning we don't include so many parantheses, and introduce symbols which are defined in the original signature. For example, $x \subseteq y$ can be defined by

$$x \subseteq y \quad \text{iff} \quad \forall z\ (z \in x \to z \in y).$$

Such defined notions affect nothing since they can be replaced by their defining formulas. In general, we're satisfied giving instructions for how to construct a formula as opposed to giving it explicitly. The same principle also holds for proofs. For an explicit example of this, the quantifier '$\exists!$' is generally used to mean "there exists a unique". We use "$\exists! x\ \varphi(x)$" merely as shorthand for $\exists x\ (\varphi(x) \wedge \forall y(\varphi(y) \to x = y))$.

§1 B. The proofs of formulas

With the notion of formula comes the notion of proof: a means of manipulating formulas. The concept of proof should be fairly familiar at this point. Note that in setting up the proof system, we should be trying to emulate valid reasoning in the meta-theory, though there is no association of meaning with formulas yet. A priori, there's no reason we couldn't allow ourselves to conclude $\varphi \wedge \psi$ from $\varphi \vee \psi$—"both" from "at least one". So there is some careful setup required in what precisely is allowed—so called *logical axioms*. The following is an informal definition, omitting what precisely a logical axiom is[iv].

1 B • 1. Definition

Let T be a collection of formulas, and φ a formula. T *proves* φ, or $T \vdash \varphi$, iff there is a sequence of formulas where every member

1. is a given assumption, i.e. a member of T; or
2. is a logical axiom, e.g. $x = x$ or $(\neg\neg\psi) \leftrightarrow \psi$; or
3. follows from previous ones by given rules of inference, e.g. ψ follows from φ and $\varphi \to \psi$.

For example, one can prove $\forall x \forall y(x + y = y + x)$ from the axioms of peano arithmetic, PA, which are then regarded as given assumptions in the above. A collection of formulas is generally called a *theory*. Note that the statement $T \vdash \varphi$ for "there is a proof of φ from the formulas T" is a meta-theoretic one about the logic system. That said, one interesting result about set theory is that it can talk about such notions through formulas inside the logic system.

And as with formulas, it's rare to give proofs as just a sequence of formulas, because they are hard to read and comprehend. Even when annotated, it's hard to see at a glance that the formulas obey the definition. For example, consider the following tedious proof of the obvious fact that $\varphi \to \varphi$ for any formula φ.

1. $(\varphi \to (\varphi \to \varphi)) \to$ $((\varphi \to ((\varphi \to \varphi) \to \varphi)) \to (\varphi \to \varphi))$	(from axiom scheme $(\varphi \to \psi) \to ((\varphi \to (\psi \to \chi)) \to (\varphi \to \chi))$ where ψ is $(\varphi \to \varphi)$ and χ is φ)
2. $\varphi \to (\varphi \to \varphi)$	(from axiom scheme $\varphi \to (\psi \to \varphi)$ where ψ is φ)
3. $(\varphi \to ((\varphi \to \varphi) \to \varphi)) \to (\varphi \to \varphi)$	(1, 2 and Modus Ponens)
4. $\varphi \to ((\varphi \to \varphi) \to \varphi)$	(from axiom scheme $\varphi \to (\psi \to \varphi)$ where ψ is $\varphi \to \varphi$)
5. $\varphi \to \varphi$	(3, 4 and Modus Ponens)

[iii] Arguably set theory uses many more symbols, e.g. '$\subseteq$', '$\emptyset$', and so forth. But these can be better regarded as short-hand for statements which use only '$\in$' and '$=$'.

[iv] Many texts make do with a list of around 15 axioms, axiom schemes, and rules of inference. So it should be clear why the exact details are omitted here.

So often proofs are given as instructions for creating a proof rather than just a sequence of formulas. This perspective is useful when arguing in the meta-theory about proofs in the logic system.

This concludes the syntactic portion of first order logic, and now we will look towards interpreting these formulas, since thus far formulas are regarded merely as a bunch of markings on paper formed in a certain way.

§ 1 C. The semantics of formulas

Now we will move on to the semantics of first order logic, looking at how to interpret these formulas and reason from them in the meta-theory. In some sense the goal is to answer "what makes a formula true?". Answering this requires first fixing a context we ask the question in, and then we build up a notion of truth in just the same way we've built up formulas. The explanations given here relate somewhat back to the real world insofar as they assume that structures, relations, functions, and so forth exist. So we might as well assume that we're working in V, where we know that these things exist.

Firstly, we have the notion of a structure. This is in some sense where we evaluate truth. For example, when we ask whether the group operation $\cdot$ is commutative, we answer relative to some particular group. The question can be asked of any group, but the answer depends on the group we evaluate in. In the same way, we can ask questions in a fixed signature, but the answer depends on the structure.

1 C • 1. Definition

Let σ be a signature. A $\mathtt{FOL}(\sigma)$-*structure* or *model* is a pair $\mathsf{M} = \langle M, \varsigma\rangle$ where M is the *universe* of M, and

1. For every n-place relation symbol R in σ, there is one $R^{\mathsf{M}} \in \varsigma$ with R^{M} a relation on tuples of M; and
2. For every n-place function symbol f in σ, there is one $f^{\mathsf{M}} \in \varsigma$ with f^{M} a function from tuples of M to M.

Intuitively, ς tells us how the model interprets the symbols of the signature σ, and the members of ς are the *interpretations* of the members of σ. For example, the signature $\sigma = \{\preccurlyeq\}$ has models which are really just any set equipped with a binary relation. For example $\langle \mathbb{N}, <\rangle$ is a $\{\preccurlyeq\}$-model, and so is any graph $\langle G, E\rangle$ where E is the edge relation of the graph. Under this definition, for any signature σ, any σ-model is also an $\emptyset$-model where there are no non-logical symbols, and the only statements are about equality[v]. In fact for any σ-model is also a δ-model for any $\delta \subseteq \sigma$.

By following the construction of any given formula, this association of a symbol in σ with the interpretation in ς presents how to tell whether any given formula is true or false in a given structure in the natural way we read formulas. Note that there will always be a fact of the matter in any given structure of whether a formula is true or false in it, even if it isn't possible to determine practically. Explicitly, we have the following definition.

1 C • 2. Definition

Let σ be a signature with R in σ a relation symbol. Let φ and ψ be $\mathtt{FOL}(\sigma)$-sentences; and let M a $\mathtt{FOL}(\sigma)$-model with various $m_i \in M$. Write

$$\begin{array}{lll}
\mathsf{M} \vDash R(m_1, \cdots, m_n) & \text{if and only if} & R^{\mathsf{M}}(m_1, \cdots, m_n) \text{ holds,}\\
\mathsf{M} \vDash m_1 = m_2 & \text{if and only if} & m_1 = m_2,\\
\mathsf{M} \vDash \varphi \wedge \psi & \text{if and only if} & \mathsf{M} \vDash \varphi \text{ and } \mathsf{M} \vDash \psi,\\
\mathsf{M} \vDash \neg\varphi & \text{if and only if} & \mathsf{M} \nvDash \varphi,\\
\mathsf{M} \vDash \forall x \varphi(x) & \text{if and only if} & \mathsf{M} \vDash \varphi(m) \text{ for every } m \in M,\\
\mathsf{M} \vDash \exists x \varphi(x) & \text{if and only if} & \mathsf{M} \vDash \varphi(m) \text{ for some } m \in M.
\end{array}$$

Implicit in this is the ability to interpret terms in the signature, and this is done exactly as one would expect. For example, the interpretation of $f(m_1, g(m_2))$ is just $f^{\mathsf{M}}(m_1, g^{\mathsf{M}}(m_2))$. Free variables are left uninterpreted, so this is why we only deal with sentences. Also note that we are mixing formal symbols and non-formal ones, leaving the parameters implicit when needed. It's important to realize that parameters can only be used when we've fixed a particular model. Parameters—like $m_1, m_2 \in M$ in the above—are not symbols in the language, and so cannot be

[v] We can still say meaningful things in this language, but mostly this is about the number of things: $\exists x \forall y (x = y)$ will require that there is only one element, for example. Some systems also drop the need for equality, in which case there are no formulas without relation symbols.

referenced in general. In some sense, parameters are used here merely to build up a notion of truth. With this concept firmly in place, notation will be slightly abused in the following ways.

1 C • 3. Definition

Let σ be a signature, and let φ and ψ be formulas, and T a theory all in $\mathtt{FOL}(\sigma)$. Let $\mathbf{M}$ be a $\mathtt{FOL}(\sigma)$-structure. Write

$$\begin{array}{lll} \mathbf{M} \vDash T & \text{if and only if} & \mathbf{M} \vDash \theta \text{ for every } \theta \in T. \\ \varphi \vDash \psi & \text{if and only if} & \text{every } \sigma\text{-model } \mathbf{M} \text{ with } \mathbf{M} \vDash \varphi \text{ also has } \mathbf{M} \vDash \psi. \\ T \vDash \psi & \text{if and only if} & \text{every } \sigma\text{-model } \mathbf{M} \text{ with } \mathbf{M} \vDash T \text{ also has } \mathbf{M} \vDash \psi. \end{array}$$

For example, $(\varphi \wedge \psi) \vDash \varphi$, since any model $\mathbf{M} \vDash (\varphi \wedge \psi)$ has $\mathbf{M} \vDash \varphi$ by Definition 1 C • 2.

These definitions comprise all the semantics of first order logic, and they all take place all in the meta-theory, meaning that $\varphi \vDash \psi$ if there is a meta-theoretic argument about models of φ. Alternatively, it might be the case that all models of φ also model ψ merely by chance with no intelligible reason behind it. So far this situation hasn't been ruled out. It is up to the next subsection to dispel this possibility.

§ 1 D. Connecting syntax and semantics

We now have the basic setup for working in mathematics. On the one hand, we can symbolically manipulate our way to various formulas, and on the other, we can argue in the meta-theory about whether certain structures satisfy a given formula. The central question, however, is whether there is any connection between the two, that is, whether '$T \vdash \varphi$' and '$T \vDash \varphi$' have any relationship.

Clearly, we should have set up our proof system to be sound, that is to say that if $T \vdash \varphi$ then $T \vDash \varphi$. This way we aren't making any "mistakes" in our symbolic manipulations. Proving that any given proof system is in fact sound can be done fairly easily through meta-theoretic arguments about structures. Mostly this amounts to checking that each logical axiom and rule of inference holds in every model.

Quite a striking result in the study of first order logic is the completeness theorem which says that the converse also holds with our notion of proof.

1 D • 1. Theorem (Completeness)

Let σ be a signature, and let T be a theory, and φ a formula in $\mathtt{FOL}(\sigma)$. Therefore $T \vDash \varphi$ implies $T \vdash \varphi$.

Proof .:.

Suppose $T \vDash \varphi$, but $T \nvdash \varphi$. This means $T \cup \{\neg\varphi\}$ is consistent (assuming the proof system is good), meaning that it doesn't prove a contradiction $\varphi \wedge \neg\varphi$. Note that $T \cup \{\neg\varphi\}$ cannot have a model, however, as this model would satisfy T and $\neg\varphi$, contradicting that $T \vDash \varphi$. To get our contradiction, we will construct a model of $T \cup \{\neg\varphi\}$ out of syntax.

Call a $\mathtt{FOL}(\sigma)$-theory T *complete* iff for every $\mathtt{FOL}(\sigma)$-sentence φ, either φ is in T, or $\neg\varphi$ is in T. By well-ordering the $\mathtt{FOL}(\sigma)$-sentences, we can successively decide whether to put a given sentence in an expansion T_0 or not according to whether the resulting expansion of T would be consistent (i.e. put it in if it is, if it's not, then leave it out). Hence we can expand $T \cup \{\neg\varphi\}$ to a theory T_0 which is consistent and complete: just the result of this process.

Now by well-ordering T_0, for each existental statement φ, being $\exists x\ \psi(x)$ in T_0, associate a unique constant c_φ, and add in the statement $\psi(c_\varphi)$ to the new theory T_1 in the expanded signature σ_1. Also expand to make sure T_1 is still consistent and complete now in $\mathtt{FOL}(\sigma_1)$. Repeating this process infinitely many times to take the closure under this propety, we end up with a complete, consistent (assuming the proof system is good) theory T_ω in an expanded signature σ_ω such that if $\exists x\ \psi(x)$ is in T_ω, then $\psi(c)$ is in T_ω for some constant symbol c in $\mathrm{cnst}_{\sigma_\omega}$.

Now we construct a model of T_ω, which is then still a model of T (by forgetting about the constants of σ_ω, we end up with a $\mathtt{FOL}(\sigma)$-model rather than a $\mathtt{FOL}(\sigma_\omega)$-model). Firstly, for c a constant symbol of σ_ω, consider the equivalence class $[c]$ consisting of all the other constants d such that $T_\omega \vdash d = c$. This is an equivalence class as T_ω is complete (assuming we've set up the proof system correctly). Now consider the structure M with universe M being the set of these equivalence classes, and with function interpretations given by

$$f^{\mathsf{M}}([d_1], \cdots, [d_n]) = [d_0] \quad \text{iff} \quad T_\omega \vdash f(d_1, \cdots, d_n) = d_0,$$

and similarly for relations (again, assuming the proof system is good, this is well-defined). The resulting structure then satisfies $\mathsf{M} \vDash T_\omega$, and so we have a model of $T \cup \{\neg\varphi\}$, and so $T \nvDash \varphi$. ⊣

This identifies the "accidental truth" of being true by chance in all models with the "justified truth" of proof. This also allows us to make conclusions from valid arguments in the meta-theory about models, and conclude that there are syntactic proofs of these results. Most important for our purposes is the fact that if $T \nvdash \varphi$, then $T \nvDash \varphi$. In particular, if T is *consistent*—meaning $T \nvdash (\varphi \wedge \neg\varphi)$—then there is a model of T. This connection between finite sequences of formulas and the existence of structures is somewhat surprising considering that structures can be very large. Furthering this relation between the finite and the infinite is the compactness theorem.

Given that proofs are finite, the compactness theorem for proofs can yield important results when paired with Completeness (1 D • 1).

1 D • 2. Theorem (Compactness)

(ZFC) Let T be a theory. Therefore T has a model if and only if each finite $\Delta \subseteq T$ has a model.

Proof .:.

If T has a model, then clearly every finite subset does too. But if T doesn't, then for any formula φ, $T \vDash (\varphi \wedge \neg\varphi)$, because no model $\mathsf{M} \vDash (\varphi \wedge \neg\varphi)$. By Completeness (1 D • 1), $T \vdash (\varphi \wedge \neg\varphi)$. Since proofs are finite, there is some finite subset $\Delta \subseteq T$ which contains all the formulas of T used in proving $(\varphi \wedge \neg\varphi)$. This finite subset then also has $\Delta \vdash (\varphi \wedge \neg\varphi)$, and so by soundness, $\Delta \vDash (\varphi \wedge \neg\varphi)$. Hence this finite subset of T can't have a model.⊣

These two theorems are very useful for their ability to generate models. As noted above, consistent theories have models which say that they're true. This is the kind of black magic that allows us to form models of peano arithmetic which aren't just the standard natural numbers. Adding to this black magic is the Löwenheim-Skolem theorem, which is the final theorem we need in the background of first order logic, and it again allows us to conclude the existence of models with extremely nice properties. The proof of this is basically a more careful version of Completeness (1 D • 1), but we are not yet equipped to prove it without knowing some more set theory. In particular, we require knowledge of infinite cardinals.

We end this section with a bit of notation that will prove useful. In particular, "FOLp" or "$\mathtt{FOLp}(\sigma)$" is used to denote "first-order logic *with parameters*". Really this is only used in the context of formulas: a formula is FOLp iff it is of the form $\varphi(\vec{v}, \vec{p})$ for some variables $\vec{v}$, and some parameters $\vec{p}$. So this is always made in the context of some (arbitrary) model. For example, the identity element in a group is FOL-definable, meaning definable without parameters. Given an arbitrary element g of the group, g^{-1} is FOLp-definable: it is the y such that $g \cdot y = y \cdot g = 1$, i.e. $\forall z((g \cdot y) \cdot z = z \cdot (g \cdot y) = z)$: we used g as a parameter here.

Section 2. Basic Set Theoretic Concepts

Recall the following definition from Subsection 0 A.

2 • 1. Definition

The *universe of sets* is the structure $\mathbf{V} = \langle \mathrm{V}, \in \rangle$, where V consists of all sets, and $\in$ denotes membership.

The axioms of set theory are various true statements (assumed) about $\mathbf{V}$, although there will be other structures that also satisfy them. Before stating the axioms precisely, I want to give some examples. If we consider the natural numbers, $\mathbb{N}$ is typically used to denote the collection of all natural numbers. So, for example, $1 \in \mathbb{N}$, and $4 \in \mathbb{N}$. We can also consider smaller sets. If we can list out all the elements of a set, we may denote the set by enclosing the members in braces: the set of 1, 4, and 8 is $\{1, 4, 8\}$.

Note that in general, $x \neq \{x\}$. To see why this is true, consider a more physical analogy: if we take x to be a marble, $\{x\}$ is a bag with one marble in it, whereas x is just the marble itself. The concept is flexible enough to allow us to collect together many things at once. If we don't wish to list out explicitly all the elements of a set, we may instead write a description in the following form: $\{x : \varphi(x)\}$ is taken to be the set of all x that satisfy φ. For example, $\{x : x \text{ is a person}\}$ is the set of all people. Similarly, $\{x \in \mathbb{N} : x^2 = 1\}$ is the set of all natural numbers that square to 1. As sets are determined by their members (two sets are the same iff they have the same elements) this set is just $\{1\}$. This is a *subset* of $\mathbb{N}$ in that all of its elements are in $\mathbb{N}$: it contains fewer members. We write $x \subseteq y$ to denote that x is a a subset of y, translated as $\forall z (z \in x \rightarrow z \in y)$ in first-order logic.

There are other ways of forming sets. For example, if x is a set, we can consider the *powerset*, the set of all collections formed from elements of x. Formally, $\mathcal{P}(x) = \{t : \forall y (y \in t \rightarrow y \in x)\}$. Additionally, we have operations on sets, like union and intersection. For example, regarding lines as sets of points, the intersection of two (non-parallel) lines is always the set containing exactly one point. In particular, $\{\langle x, y \rangle \in \mathbb{R}^2 : y = 2x + 3\}$ is a line, as is $\{\langle x, y \rangle \in \mathbb{R}^2 : y = -x\}$ is another, and their intersection is

$$\{\langle x, y \rangle \in \mathbb{R}^2 : y = 2x + 3\} \cap \{\langle x, y \rangle \in \mathbb{R}^2 : y = -x\} = \{\langle -1, 1 \rangle\}.$$

Now that we have some basic intuition set up, let's consider the following true statements about $\mathbf{V}$, which are axioms of ZFC.

2 • 2. Definition (Axioms)

(*Extensionality*) two sets are equal whenever they have the same members:

$$\forall x \ \forall y \ (x = y \leftrightarrow \forall v \ (v \in x \leftrightarrow v \in y)).$$

(*Empty set*) there is a set $\emptyset$ with no members: $\exists z \forall x (x \notin z)$.

(*Comprehension*) for each x, and for each $\mathtt{FOLp}(\in)$-formula $\varphi(v)$, $\{v \in x : \varphi(v)\}$ exists: for φ a $\mathtt{FOL}(\in)$-formula,

$$\forall w_0 \cdots \forall w_n \ \forall x \ \exists z \ \forall v \ (v \in z \leftrightarrow v \in x \wedge \varphi(v, \vec{w})).$$

Extensionality is perhaps the most definition-like axiom, contained in the idea of a set.

2 • 3. Corollary

Suppose $\{x\} = \{y\}$. Therefore $x = y$.

The empty set will provide the base for our universe in the following sense.

2 • 4. Result

For every set A, $\emptyset \subseteq A$. Moreover, $A \subseteq \emptyset$ implies $A = \emptyset$.

Proof .:.

$\emptyset \subseteq A$ since every element of $\emptyset$ (of which there are none) is an element of A. Now suppose $A \subseteq \emptyset$. Thus $\forall x (x \in A \rightarrow x \in \emptyset)$. For each x, $x \in \emptyset$ is false, and thus $x \notin A$, Hence $\forall x (x \notin A)$, and therefore A and $\emptyset$ have the same elements: no elements. By extensionality, $A = \emptyset$. ⊣

Comprehension[vi] really is a *scheme*, meaning that for each FOL-formula, we get a different axiom. It is an attempt to formalize the idea of $\{x : \varphi(x)\}$. It's important to realize, however, that the full generality is inconsistent we can only consider the subset $\{x \in A : \varphi(x)\}$ for some set A. The idea is that we can't take arbitrary collections and call them sets, as seen in the following theorem.

2 • 5. Theorem (Russell's Paradox)

There is no set $\{x : x = x\}$. Equivalently, $\neg\exists s \forall x (x \in s)$.

Proof .:.

If there were such a set, call it V. Now consider by comprehension the subset $a = \{x \in V : x \notin x\}$. By hypothesis, $a \in V$. Now we can question whether $a \in a$ or not. If $a \in a$, then a meets the definition: $a \notin a$, a contradiction. Hence $a \notin a$. But this means that a doesn't meet the definition of a, meaning $a \in a$, again a contradiction. So either way we have a contradiction, and so the hypothesis that V exists is false. $\dashv$

So comprehension at least says that we can consider (definable) subsets. In some sense, the issue is that the collection of all x is too big to be a set: V is not a set. So comprehension says that if we have a set, then all the subsets are small enough to be sets too.

§ 2 A. A word on classes versus sets

We often want to talk about collections that aren't sets. Russell's Paradox (2 • 5) gives one such example: the collection of all sets, V. There are other, less *ad hoc* collections we will want to consider later, but this raises the question of how do we talk about these things?

2 A • 1. Definition

Let $\mathbf{A}$ be a model of set theory. A *class* of $\mathbf{A}$ is a collection $\mathrm{C} \subseteq A$ which is FOLp($\in$)-definable. A class is a *proper class* iff it is not a set, meaning not in A.

As a bit of notation, classes will generally be written upright: like 'V', 'L', 'Ord', 'HOD', instead of 'V', 'L', 'Ord', 'HOD'. Another thing to realize is that I will write "subset" even for a class, so even though V is a proper class, I will write that V is a subset of V rather than the more cumbersome "sub-collection" or "sub-class".

Note that these collections are not necessarily a part of the set theoretic universe: every set is a class, but not vice-versa. Comprehension tells us that the intersection of a set with a class is a set[vii]. There are more complicated understandings of classes that allow more general collections of the universe. But at that point, we get into the realm of "class theory" rather than set theory. And before learning class theory, one needs to start with a good understanding of set theory. In essence, the typical model of class theory will be one that satisfies ZFC for sets in an expanded language with constant symbols for *at least* the definable collections of sets.

But under our definition, classes really are just short-hand for formulas. So often results for sets generalize to results for classes just by virtue of classes being definable. That said, it's still important to realize that classes are not always sets, and certain theorems do not always generalize to classes. The basic problem is that classes can't be collected together to get another class: for C a class, {C} isn't necessarily a class, because it's not a collection of sets.

§ 2 B. Ordered pairs

So far, the set theory presented is relatively uninteresting, because the axioms do not allow us to form sets with new elements: we may only take subsets. Moreover, even if we have these sets, it's not completely clear what the benefit of them is. To motivate things a little more, sets are seen as a foundation of mathematics, both practically, and philosophically. Often, when one needs to make things mathematically precise, it is done using sets[viii]. So to begin with,

[vi] also called *separation*

[vii] So in some sense, every part of a class is a set, although the totality might not be.

[viii] there are other theories some people put forth as a foundation of mathematics, but their proponents often either defer the serious paradoxical issues for set theory to deal with, or fail to start from philosophically basic notions.

we will first show that we can formalize an ordered pair $\langle x, y\rangle$, in that we have a construction where $\langle a, b\rangle = \langle a', b'\rangle$ if and only if $a = a'$ and $b = b'$. This will allow us to talk about sequences, functions, relations, and so forth. To do this, we need some additional axioms that reflect what's true of **V**.

2 B • 1. Definition (Axiom)

(*Pairing*) for any two sets x and y, the pair $\{x, y\}$ exists: $\forall x\ \forall y\ \exists z\ \forall v\ (v \in z \leftrightarrow (v = x \vee v = y))$.

2 B • 2. Definition

For x, y sets, the *ordered pair* of x and y, $\langle x, y\rangle$ is the set $\{\{x\}, \{x, y\}\}$.

As a side note, if $x = y$, then $\langle x, y\rangle$ collapses down to $\{\{x\}\}$, since $\{x, y\} = \{x, x\} = \{x\}$ because the two have the same members. Now let's prove the single point of having an ordered pair: that the entries are uniquely determined by the ordered pair.

2 B • 3. Result

Let x, x', y, y' be sets. Therefore $\langle x, y\rangle = \langle x', y'\rangle$ iff $x = x'$ and $y = y'$.

Proof .:.

Clearly if $x = x'$ and $y = y'$, then $\langle x, y\rangle = \langle x', y'\rangle$. So suppose $\langle x, y\rangle = \langle x', y'\rangle$, meaning that these sets have the same members. The members of these sets are $\{x\}$ and $\{x, y\}$, and $\{x'\}$ and $\{x', y'\}$.

If $x \neq y$ and $x' \neq y'$, then the two-element sets must be equal, and the one-element sets must be equal, implying $x = x'$ and $\{x, y\} = \{x', y'\}$. Since we already know $x = x'$, we must have $y = y'$. If $x = y$, then $\langle x, y\rangle = \{\{x\}\}$. Hence both elements of $\langle x', y'\rangle$ are equal to this: $\{x\} = \{x, y\} = \{x'\} = \{x', y'\}$, implying that $x' = y' = x = y$. The same idea holds if $x' = y'$. $\dashv$

We can also refer to the left and right coordinate of an ordered pair in this way: given an ordered pair z, the left-coordinate is just the x satisfying $\exists y(\langle x, y\rangle = z)$. In fact, using another axiom, we can restrict the search for such a y to an element of the union of z.

2 B • 4. Definition (Axiom)

(*Union*) for any family of sets F, there is a set containing the elements of all of those sets:

$$\forall F\ \exists U\ \forall v\ (v \in U \leftrightarrow \exists x(x \in F \wedge v \in x)).$$

We denote the union by $\bigcup F$, in this case. For just two sets, write $x \cup y = \{a : a \in x \vee a \in y\}$ rather than the more clumsy $\bigcup\{x, y\}$, which exists by union and pairing. For a concrete example of a union, consider $x = \{1, 2\}$, and $y = \{2, 4, 10\}$. Therefore $x \cup y = \{1, 2, 4, 10\}$. A related concept, which we could already form through comprehension, is the intersection of two sets: $x \cap y = \{a : a \in x \wedge a \in y\}$. More generally, for a non-empty family, F, the intersection $\bigcap F = \{a : \forall x \in F(a \in x)\}$, which can be written as a subset of each particular $x \in F$. Similarly, we can take complements: $x \setminus y = \{a \in x : a \notin y\}$. Using the same x and y example from before, $x \setminus y = \{1\}$ while $x \cap y = \{2\}$. Note that we have the following trivial facts about intersection, union, and so forth, mostly which just follow from properties of sentential connectives:

- $x \cap x = x$, $x \cup x = x$, $x \cup \emptyset = x$;
- $x \cap \emptyset = \emptyset$, $x \setminus x = \emptyset$;
- $x \setminus (x \cap y) = x \setminus y$;
- $x \cap y \subseteq x$, and if $a \subseteq x$ and $a \subseteq y$, then $a \subseteq x \cap y$;
- $x \cap (y \cap z) = (x \cap y) \cap z$, and similarly for union;
- $(x \cap y) \cup z = (x \cup z) \cap (y \cup z)$, and $(x \cup y) \cap z = (x \cap z) \cup (y \cap z)$;
- if $x \subseteq a$ and $y \subseteq a$, then $x \cup y \subseteq a$;
- if $x \subseteq y$ and $y \subseteq a$, then $x \subseteq a$; and
- $x \subseteq y$ iff $y \cup x = y$ iff $x \cap y = x$ iff $x \setminus y = \emptyset$.
- $x \subseteq y$ implies $\bigcup x \subseteq \bigcup y$.

These also have a related definition, since sets having completely different elements is very useful.

2 B • 5. Definition

Two sets x and y are *disjoint* iff $x \cap y = \emptyset$. A family of sets F consists of *disjoint sets* or *pairwise disjoint sets* iff $x \cap y = \emptyset$ for all $x, y \in F$.

Now ordered pairs on their own are fine, but we still need to be able to do more with them to do any basic set theory. Obviously using pairing, we can form $\{\langle x, y\rangle, \langle x', y'\rangle\}$. We can also form $\{\langle x, y\rangle, \langle x', y'\rangle, \langle x'', y''\rangle\}$ using another application of pairing and union:

$$\{\langle x, y\rangle, \langle x', y'\rangle, \langle x'', y''\rangle\} = \{\langle x, y\rangle, \langle x', y'\rangle\} \cup \{\langle x'', y''\rangle, \langle x'', y''\rangle\}.$$

We have two potential routes to form arbitrary sets of pairs—excluding finite applications of pairing and union—powerset (with comprehension), and replacement. First we introduce replacement.

2 B • 6. Definition

A $\texttt{FOLp}(\in)$-formula $\varphi(x, y)$ *defines a function over* D iff for every $x \in D$ there is a unique y with $\varphi(x, y)$. Symbolically, $\forall x\ (x \in D \to \exists! y\ \varphi(x, y))$.

Replacement then says in effect that if we can definably transform elements of a set, then the set of the transformations exist.

2 B • 7. Definition (Axiom)

(*Replacement*) the image of a function over a set is a set: for each $\texttt{FOL}(\in)$-formula φ,

$$\forall w_0 \cdots \forall w_n \forall D\ \big(\varphi(x, y, \vec{w}) \text{ defines a function over } D \to \exists R(y \in R \leftrightarrow \exists x(x \in D \land \varphi(x, y, \vec{w})))\big).$$

It should be clear that D is the intended domain of the function defined by φ, and R is the range of φ restricted to D. So replacement is saying that R exists: if I can define a function from a set, then the range is a set. So if we consider the function mapping xs and ys to $\langle x, y\rangle$, we get a cartesian product as the range.

2 B • 8. Definition

The *cartesian product* $A \times B$ of A and B is the set of all pairs from A and B: $\{\langle a, b\rangle : a \in A \land b \in B\}$.

2 B • 9. Result

Let A and B be arbitrary. Therefore $A \times B$ exists[ix].

Proof .:.

For each $a \in A$ consider the formula $\varphi(b, p, a)$ which is just that $p = \langle a, b\rangle$. This is of course shortened, but the defining notions can be replaced here. Regardless, it's clear that this defines a function over B, where b maps to $\langle a, b\rangle$ for our fixed $a \in A$. So replacement says that there is some

$$R_a = \{p : \exists b \in B\ \varphi(b, p, a)\} = \{\langle a, b\rangle : b \in B\}.$$

This is an individual slice of the cartesian product. So consider the function $\psi(r, a)$ which states $r = R_a$, i.e. $r = \{\langle a, b\rangle : b \in B\}$. (We can do this by taking the even longer formula $\psi(r, a)$ to be $\forall x\ (x \in r \leftrightarrow \exists b \in B\ \varphi(b, x, a))$.) This defines a function over A, and so another application of replacement yields the set

$$P = \{\{\langle a, b\rangle : b \in B\} : a \in A\}.$$

Hence $\bigcup\{\{\langle a, b\rangle : b \in B\} a \in A\} = \{\langle a, b\rangle : a \in A \land b \in B\} = A \times B$. $\dashv$

§ 2 C. Relations

The cartesian product is the basis for most of basic set theory, since it allows us to consider relations and fuctions, and thus define sequences, and notions of size. Really, if sets are supposed to be devoid of all structure beyond membership, this idea allows us to put structure back into play, and thus work with more complicated ideas all within set theory.

[ix] Note that we've also shown that the cartesian product of classes exists as well. In particular, for A and B classes, we have the $\texttt{FOLp}$-formula defining A × B by $x \in \mathrm{A} \times \mathrm{B}$ iff $\exists y\ \exists z\ (x = \langle y, z\rangle \land y \in \mathrm{A} \land z \in \mathrm{B})$.

2 C • 1. Definition

A *relation* is a subset $R \subseteq A \times B$ for some A, B. For any relation R, $\mathrm{dom}(R) = \{x : \exists y\ (\langle x, y\rangle \in R)\}$, and similarly, $\mathrm{ran}(R) = \{y : \exists x\ (\langle x, y\rangle \in R)\}$.

The existence of the domain and range of R can be shown by the union axiom: $x, y \in \bigcup\langle x, y\rangle = \{x, y\} \cup \{x\} = \{x, y\}$ so that $\langle x, y\rangle \in R$ implies $x, y \in \bigcup\bigcup R$. Hence we can take the appropriate subset to define the domain and range. Alternatively, we can use replacement. But resorting to the more basic axioms can be insightful.

Note that then if R is a relation, every subset of R is a relation too. Moreover, the union of relations are relations. Really a relation is just a set R where $z \in R$ implies $z = \langle x, y\rangle$ for some x and y. So the relation doesn't need to be over the same set or have some intuitive reason behind relating elements. Note that for R a relation, we will often write $x\ R\ y$ for $\langle x, y\rangle \in R$. Note that we can have the relation defined on three sets (or more) just by having $\langle x, y\rangle \in R$ always having y an ordered pair of some form. We will make this more formal or official later, so for now we focus on binary relations. Again, we get some immediate facts: for R and S relations,

- $\mathrm{dom}(R \cup S) = \mathrm{dom}(R) \cup \mathrm{dom}(S)$;
- $\mathrm{ran}(R \cup S) = \mathrm{ran}(R) \cup \mathrm{ran}(S)$;
- $\mathrm{dom}(R \cap S) \subseteq \mathrm{dom}(R) \cap \mathrm{dom}(S)$; and
- $\mathrm{ran}(R \cap S) \subseteq \mathrm{ran}(R) \cap \mathrm{ran}(S)$.

Given any relation, we can form the inverse, where we swap all the entries of the ordered pairs:

2 C • 2. Definition

For R a relation, define $R^{-1} = \{\langle y, x\rangle : \langle x, y\rangle \in R\}$ to be the *inverse* or *converse* of R.

The existence of R^{-1} can be shown through a variety of methods, notably replacement. Note that this behaves exactly as one would expect:

2 C • 3. Result

Let R be a relation. Therefore R^{-1} is a relation, and $(R^{-1})^{-1} = R$. Moreover, for S a relation, $(R \cap S)^{-1} = R^{-1} \cap S^{-1}$.

Proof .:.

Clearly R^{-1}, as a set of ordered pairs, is a relation. Moreover, $(R^{-1})^{-1} = \{\langle x, y\rangle : \langle y, x\rangle \in R^{-1}\} = \{\langle x, y\rangle.\langle x, y\rangle \in R\} = R$. To see that the inverse of an intersection is the intersection of the inverses, let $\langle y, x\rangle \in (R \cap S)^{-1}$. Therefore $\langle x, y\rangle \in R \cap S$ and so $\langle y, x\rangle \in R^{-1}$ and $\langle y, x\rangle \in S^{-1}$. Similarly, if $\langle y, x\rangle \in R^{-1} \cap S^{-1}$, then $\langle x, y\rangle$ must be in both R and in S, so that $\langle y, x\rangle \in (R \cap S)^{-1}$. So the two sets have the same elements, and so must be equal. ⊣

One of the most important kinds of relations is a partial order, notable mostly for the notion of transitivity.

2 C • 4. Definition

Let R be a relation. Write $x\ R\ y$ for $\langle x, y\rangle \in R$. We say R is a realtion *over* X iff $\mathrm{dom}(R) \cup \mathrm{ran}(R) = X$.

- R is *transitive* iff $\forall x\ \forall y\ \forall z\ (x\ R\ y \wedge y\ R\ z \rightarrow x\ R\ z)$.
- R is *symmetric* iff $\forall x\ \forall y\ (x\ R\ y \leftrightarrow y\ R\ x)$.
- R is *antisymmetric* iff $\forall x\ \forall y\ (x\ R\ y \wedge y\ R\ x \rightarrow x = y)$.
- R is *total* iff $\forall x\ \forall y\ (x, y \in \mathrm{dom}(R) \cup \mathrm{ran}(R) \rightarrow (x\ R\ y \vee x = y \vee y\ R\ x))$.
- R is *reflexive* iff $\forall x\ (x \in \mathrm{dom}(R) \cup \mathrm{ran}(R) \rightarrow \langle x, x\rangle \in R)$.
- R is a *partial order* iff it is transitive, and antisymmetric.
- R is *linear* iff it is transitive, antisymmetric, and total.

A relation R is called a *strict order* if $\langle x, x\rangle \notin R$ for all x.

We now get some very easy results about various relations that the reader should check to confirm their intuitions.

- The identity relation $\mathrm{id}_A = \{\langle x, x\rangle : x \in A\}$ is symmetric and antisymmetric.
- R is symmetric iff $R = R^{-1}$.

- If R is a linear order then $R \cap (A \times A)$ is a linear order for any set A.
- If R is antisymmetric, and $\operatorname{dom}(R) \cup \operatorname{ran}(R)$ has more than one element, then $R^{-1} \neq R$.
- If R and S are reflexive, then $R \cup S$ is reflexive.
- If R and S are antisymmetric, then $R \cup S$ is antisymmetric.
- If R is antisymmetric, and $S \subseteq R$, then S is antisymmetric.
- If R is transitive, then R^{-1} is transitive.
- If R is a partial order, then $R \cup \mathrm{id}_{\operatorname{dom}(R) \cup \operatorname{ran}(R)}$ is a reflexive partial order.
- If R is a partial order, then $R \setminus \mathrm{id}_{\operatorname{dom}(R) \cup \operatorname{ran}(R)}$ is a strict partial order.

The relations which are of fundamental importance to set theory are well-founded relations, and equivalence relations.

2 C • 5. Definition

A relation R is *well-founded* iff for every subset $X \subseteq \operatorname{dom}(R) \cup \operatorname{ran}(R)$, there is an R-minimal element of X, meaning an $x \in X$ with no $y \in X$ with $y \; R \; x$.

When we investigate well-founded linear orders. It turns out that they are *canonical* in the sense that they are all initial segments of each other (up to isomorphism). We will investigate well-founded relations later on. For now, consider some terminology regarding equivalence relations.

2 C • 6. Definition

A relation R is an *equivalence relation* iff R is reflexive, symmetric, and transitive.
An *equivalence class* of R is a $X \subseteq \operatorname{dom}(R) \cup \operatorname{ran}(R)$ such that $x \; R \; y$ for every $x, y \in X$.
For $x \in \operatorname{dom}(R) = \operatorname{ran}(R)$, write $[x]_R$, the *equivalence class of* x, for $\{y \in \operatorname{dom}(R) : x \; R \; y\}$.
For X an equivalence class of R, a *representative* of X is an $x \in \operatorname{dom}(R)$ such that $X = [x]_R$.
For X a set, a partition is a set P such that $\forall x (x \in X \to \exists! Y (Y \in P \land x \in Y))$ and $\forall Y \in P (Y \subseteq X)$.

For example, id_X is an equivalence relation over X with $[x]_= = \{x\}$ for all $x \in X$. But an equivalence relation is more general than equality. But in essence, an equivalence relation still acts like it in the following sense.

2 C • 7. Result

For R an equivalence relation and $x, y \in \operatorname{dom}(R)$, $x \; R \; y$ iff $[x]_R = [y]_R$.

Proof .:.

If $[x]_R = [y]_R$, then by reflexivity, $y \in [y]_R = [x]_R$ implies $x \; R \; y$. So suppose $x \; R \; y$. If $a \in [x]_R$ then $x \; R \; a$. By symmetry, $a \; R \; x$. Since $x \; R \; y$, symmetry yields that $a \; R \; y$ and symmetry again yields $y \; R \; a$, i.e. $a \in [y]_R$. Thus $[x]_R \subseteq [y]_R$. The same argument shows $[y]_R \subseteq [x]_R$. Therefore $x \; R \; y$ implies $[x]_R = [y]_R$. ⊣

2 C • 8. Corollary

For R an equivalence relation, $[x]_R = [y]_R$ or $[x]_R \cap [y]_R = \emptyset$ for all $x, y \in \operatorname{dom}(R)$.

Proof .:.

Suppose $a \in [x]_R \cap [y]_R$. By transitivity and symmetry, $x \; R \; a \land a \; R \; y$ implies $x \; R \; y$ so that $[x]_R = [y]_R$. ⊣

Hence, the set of equivalence classes partitions the domain of R.

2 C • 9. Corollary

For R an equivalence relation, $\{[x]_R : x \in \operatorname{dom}(R)\}$ is a partition of $\operatorname{dom}(R)$.

Conversely, partitions give rise to equivalence classes, and thus equivalence relations and partitions can be seen as the same thing.

2 C • 10. Result

Let X be a set and let P be a partition of X. Therefore the relation $R = \{\langle a, b \rangle \in X \times X : \exists Y \in P (a \in Y \land b \in Y)\}$ is an equivalence relation over $\operatorname{dom}(R) = X$.

Proof .:.

Symmetry is immediate by the commutativity of $\wedge$. Since each $x \in X$ has some $Y \in P$ with $x \in Y$, reflexivity is true of R, and this shows $\operatorname{dom}(R) = X$. So it suffices to show transitivity. Suppose $x, y \in Y \in P$ and $y, z \in Y' \in P$. As a partition, there is only one $Y'' \in P$ with $y \in Y''$ so that $Y = Y'$ and thus $x, y, z \in Y \in P$, which yields $x\ R\ z$. ⊣

The main point of equivalence classes is just that they give a new notion of equality by considering the equivalence classes instead of the equivalence relation so directly. This allows us to say things like "x and y are the same modulo R". Similarly, it allows us to define other relations so long as they respect the equivalence relation. In doing this, note that often the equivalence class $[x]_R$ will have multiple elements: $[x]_R = [y]_R$ although $x \neq y$. So if we are to make a definition about $[x]_R$ that makes reference to x, we need to ensure that this gives the same thing if we were to choose y instead as our representative.

2 C • 11. Result

Let $\approx$ be an equivalence relation on X, and let $R \subseteq X \times X$. Suppose $x\ R\ y$ iff $x'\ R\ y'$ for $x \approx x'$ and $y \approx y'$. Therefore the relation $R_{/\approx}$ over $X_{/\approx} = \{[x]_\approx : x \in X\}$ defined by

$$[x]_\approx\ R_{/\approx}\ [y]_\approx \quad \text{iff} \quad x\ R\ y$$

is well-defined, meaning independent of the choice of representatives.

Proof .:.

Suppose $[x]_\approx = [x']_\approx$ and $[y]_\approx = [y']_\approx$. Therefore $x\ R\ y$ iff $x'\ R\ y'$, meaning $[x]_\approx\ R_{/\approx}\ [y]_\approx$ iff $x'\ R\ y'$. ⊣

This is the idea from algebra that allow us to "mod out" by an equivalence relation, like via the orbits induced by other groups or ideals of a ring, generating a new group or ring. There are many applications, which we will see later.

§ 2 D. Functions

2 D • 1. Definition

A *function* is a relation $f \subseteq A \times B$ such that for each $x \in A$ there is exactly one $y \in B$ with $\langle x, y\rangle \in f$. We write $f : A \to B$, and $y = f(x)$ in this case.

- A function $f : A \to B$ is *injective* iff $f(x) \neq f(y)$ for all $x \neq y$ in A.
- A function $f : A \to B$ is *surjective* iff $\operatorname{ran}(f) = B$.
- A function is bijective iff it is injective and surjective.

We also call such functions injections, surjections, or bijections. Note that a function being surjective depends on how we regard it: obviously $f : \operatorname{dom}(f) \to \operatorname{ran}(f)$ is surjective. Note also that in this text, instead of $\operatorname{ran}(f)$ for the "range of f", we will write $\operatorname{im}(f)$ for the "image of f". This is merely a personal preference to distinguish relations from functions. Clearly if f is a function $f : A \to B$ and $B \subseteq C$, then we can also regard $f : A \to C$.

Because the objects we deal with in set theory are sets—in particular, sets that are *hereditarily* sets, meaning all their members are also sets, and the same holds for them too—we need to make the distinction between the "pointwise image" of a function as opposed to the "value" of a function. To motivate the example, consider the set $A = \{a, b, \{a\}\}$ and a function f with domain A. In general, there is a difference between $f(a)$ and $f(\{a\}) \neq \{f(a)\}$. But sometimes we do want to consider the set of values of a function, like $\{f(a)\}$. Similarly, sometimes we want to take a function, but restrict our attention to a smaller subset of its domain. To denote the difference, we have the following definition.

2 D • 2. Definition

Let $f : A \to B$ be a function over sets A, B. Let $X \subseteq A$. Write the *pointwise image* of f under X as $f"X = \{f(x) : x \in X\}$.
Write the *restriction* of f to X as $f \restriction X = \{\langle a, b\rangle \in f : a \in X\} = f \cap (X \times B)$.

So in the example above, $f"\{a\} = \{f(a)\}$ while $f(\{a\}) \neq f"\{a\}$. Note that $\operatorname{im}(f) = f"\operatorname{dom}(f)$. Since restriction allows us to chain our domain, $\operatorname{dom}(f \restriction X) = X$; we can also write $f"X = \operatorname{im}(f \restriction X)$. We also have the following operations on functions: composition and inverses (which might not be functions).

2 D • 3. Definition

Let $f : A \to B$ be a function over sets A, B. Let f^{-1} be the relation $\{\langle b, a\rangle : \langle a, b\rangle \in f\} \subseteq B \times A$. For $g : B \to C$, the composition $g \circ f$ is defined by $\{\langle a, c\rangle : \exists b \in B(f(a) = b \wedge g(b) = c)\}$.

It should be clear that $g \circ f$ is also a function, now from $\mathrm{dom}(f)$ to $\mathrm{im}(g)$.

Functions are fundamental to mathematics, as they are a means of transformations. More than functions, really, the importance is placed on the properties of functions. Most graphs of most functions will be set-theoretic haze: just a bunch of points with no discernible relationship between the points beyond satisfying the definition of being a function. So most applications will care about functions that preserve certain relationships. These are typically called homomorphisms, embeddings, and so on. We have already defined one such property: preserving inequality, or injectivity. But the key thing for now is to recognize that functions can be interpreted in purely set theoretic terms.

Let me take a moment to talk further about bijections, injections, and surjections. When letting their sheep out to graze, one technique that shepards used to make sure all sheep were accounted for was to pick up a pebble every time a sheep left. Then a pebble was dropped for every sheep that returned. So if there were any left over pebbles, there were sheep left out. Stated in terms of functions, there was a function f : sheep $\to$ pebbles which was *injective*—two different sheep get two different pebbles—and *surjective*—every pebble corresponds to a sheep—and hence bijective. Going back to the example, this means we have the same number of pebbles as sheep, and we have confimed this *without counting*. So bijections really form a notion of size between two sets: we merely rename the elements via the bijection. For a very simple example, consider $\{a, b, c\}$ and $\{\alpha, \beta, \gamma\}$. Renaming a 'α', b 'β', and c 'γ', we get $\{a, b, c\}$ should have the same number of elements as $\{\alpha, \beta, \gamma\}$, which it clearly does, and we did this without directly counting both and then seeing that the two numbers line up. In some sense, counting just adds a third set of numbers, and then a bijections to the numbers as a means of counting each set.

So to remove the middle-man of numbers, which we have not yet introduced in set theoretic terms yet, we have the following definition.

2 D • 4. Definition

Let A and B be sets. Write $A =_{\text{size}} B$ iff there is a bijection $f : A \to B$.

Ideally, we'd like to say the *cardinality* of A and B are the same. But without further technology in the form of ordinals, we have no means of saying this. Instead, we will say that the cardinality of a set A is the *class* of $\{B : A =_{\text{size}} B\}$. We also have a notion of order on these equivalence classes in the following sense.

2 D • 5. Definition

Let A and B be sets. Write $A \leq_{\text{size}} B$ iff there is an injection $f : A \to B$.

For example, $A \subseteq B$ has $A \leq_{\text{size}} B$. Note that this is in essence the only way to have a size less than or equal to a set in the following sense.

2 D • 6. Result

$A \leq_{\text{size}} B$ iff there is some $C =_{\text{size}} B$ with $A \subseteq C$.

Proof .:.

To see this, note that if $A \leq_{\text{size}} B$, then the injection $f : A \to B$ witnessing this has $f''A \subseteq B$. So take $C = (B \setminus f''A) \cup A$, where clearly $A \subseteq C$. Ostensibly, $C =_{\text{size}} B$ since it seems we can consider the function $F : C \to B$ defined by

$$F(b) = \begin{cases} b & \text{if } b \in B \setminus f''A \\ f(a) & \text{if } a \in A. \end{cases}$$

The only issue with this is that $A \cap (B \setminus f''A)$ might not be empty, which would make the above ill-defined. But assuming $A \cap B = \emptyset$, then F is a bijection. To remove the assumption $A \cap B = \emptyset$, consider instead $C = ((B \setminus f''A) \times \{\emptyset\}) \cup (A \times \{\{\emptyset\}\})$. with $F(\langle b, \emptyset\rangle) = b$ and $F(\langle a, \{\emptyset\}\rangle) = f(a)$. This yields the appropriate bijection. ⊣

We will see later that $A \leq_{\text{size}} B$ and $B \leq_{\text{size}} A$ implies $A =_{\text{size}} B$, as suggested by the notation. But the long proof of this isn't instrumental to us for now. What's important is the notion of bijection giving a notion of size.

We have the following easy properties of size and bijections. Note that "$f : A \to B$" is not just a statement that $f \subseteq A \times B$, but that f is a function with f defined on all of A (so $\operatorname{dom}(f) = A$) and $\operatorname{im}(f) \subseteq B$.

- If $f : A \to B$ and $g : B \to C$ are injective, then $g \circ f : A \to C$ is injective.
- If $f : A \to B$ is surjective, and $g : B \to C$ is surjective, then $g \circ f : A \to C$ is surjective.
- If $f : A \to B$ and $g : B \to C$ are bijections, then $g \circ f : A \to C$ is a bijection;
- equivalently, if $A =_{\text{size}} B$ and $B =_{\text{size}} C$ then $A =_{\text{size}} C$.
- If $f : A \to B$ is a bijection, then $f^{-1} : B \to A$ is a bijection;
- equivalently, $A =_{\text{size}} B$ iff $B =_{\text{size}} A$ for all A and B.
- If $f : A \to B$ is injective, then $f : A \to \operatorname{im} f$ is a bijection;
- equivalently, $X =_{\text{size}} f"X$ for $f : A \to B$ injective with $X \subseteq A$.

All of this has been done without the notion of counting, but the benefit of being able to count is that it opens up a new theory of "numbers". So we will return to the notion of size or *cardinality* later, after we have introduced the ordinals. But now we should have a basic intuition for functions and size.

§ 2 E. Transitive sets

Let's take a moment to look at so-called "transitive" sets. In some sense, this is a misnomer, since it is not the set that is transitive, but the membership relation.

2 E • 1. Definition

A set x is *transitive* iff membership into x, meaning $\{\langle a, b\rangle : a \in b \wedge b \in x\}$, is transitive.

So x being transitive is the same as saying $a \in b \in x$ implies $a \in x$. Equivalently, $b \in x$ implies $b \subseteq x$[x]. In some sense, this means that transitive xs not only contain various a with $a \in b \in x$, but that we go *all the way down* to the basis of the universe: $\emptyset$. To prove this, we need an additional axiom. In another sense, x being transitive means that the structure $\langle x, \in\rangle$ is a *submodel* of $\mathbf{V}$: they both interpret $\in$ in the same way. As a result of this, we get some nice model-theoretic results. Below is just one example of this showing that transitive sets have nice *absoluteness* properties that we will consider later.

2 E • 2. Result

Let X be transitive. Let $a, b \in X$. Therefore $\mathbf{X} = \langle X, \in\rangle \vDash a \subseteq b$ iff $\mathbf{V} \vDash a \subseteq b$.

Proof .:.

To say that $a \subseteq b$ is just short-hand for $\forall y\ (y \in a \to y \in b)$. Since $\mathbf{X}$ and $\mathbf{V}$ interpret $\in$ the same way, if $y \in X$, $\mathbf{X} \vDash$ "$y \in a \to y \in b$" iff $\mathbf{V} \vDash$ "$y \in a \to y \in b$". Since $X \subseteq \mathrm{V}$, y ranges over more sets in V than in X: if $\mathbf{V} \vDash$ "$a \subseteq b$", then $\mathbf{X} \vDash$ "$a \subseteq b$". The other direction, if $\mathbf{V} \vDash$ "$a \nsubseteq b$", then there must be some element $y \in a$ with $y \notin b$. But since X is transitive with $a, b \in X$, $y \in a \in X$ implies $y \in X$. Hence $\mathbf{V} \vDash$ "$y \in a \wedge y \notin b$" implies $\mathbf{X} \vDash$ "$y \in a \wedge y \notin b$", because they interpret $\in$ in the same way. But then $\mathbf{X} \vDash$ "$a \nsubseteq b$". ⊣

Finding examples of transitive sets and examples of non-transitive sets is easy. In particular,

1. $\emptyset$ is transitive. $\{\emptyset\}$ is transitive.
2. If x is transitive, then $x \cup \{x\}$ is transitive (any element $b \in x \cup \{x\}$ is still a subset since $b \subseteq x \subseteq x \cup \{x\}$).
3. Writing $0 = \emptyset$, $1 = \{0\}$, and $2 = \{0, 1\}$, then from the above, 0, 1, 2, and $\{0, 1, 2\}$ are transitive, but $\{1\}$, $\{0, 2\}$, and $\{2\}$ are not.

[x] Of course, we cannot have a set where $\forall b\ (b \subseteq x \to b \in x)$ by the same reasoning as in Russell's Paradox (2 • 5).

4. If x is transitive and $y \subseteq x$, then $x \cup \{y\}$ is transitive.

Now we introduce the axiom of foundation. To motivate the axiom, it's difficult to think of a set which could be an element of itself. Considering a more physical picture, you can't place a box (completely) inside itself—the concept wouldn't make any sense. Indeed, Russell's Paradox (2 • 5) partly goes through because we consider that the collection of everything that exists is an element of itself. This would suggest we should assume $\forall x\ (x \notin x)$ as an axiom. This would rule out some direct approaches, but we could still code the counter-intuitive situations through other loops: $x \in y$ and $y \in x$, for example.

The axiom of foundation rules out loops of arbitrary length, and has a great number of consequences. Intuitively, the idea can be motivated as above, but it can also be motivated though the iterative conception of what a collection is: namely, collections are built up of smaller things that have come before in a certain sense. This will turn out to be equivalent to the axiom. Explicitly, foundation merely states that membership is well-founded.

2 E • 3. Definition (Axiom)

(*Foundation*) for each x, there is a $\in$-minimal element of x: $\forall x \exists y (y \in x \wedge \forall z (z \in y \to z \notin x))$.

2 E • 4. Corollary

Assume the axiom of foundation. Therefore:

1. We never have $x \in x$.
2. In fact, there are no finite loops $x_0 \in x_1 \in \cdots \in x_n \in x_0$.
3. If $x \neq \emptyset$ is transitive, $\emptyset \in x$ is the $\in$-minimal element of x.
4. x is transitive iff $x \cup \{x\}$ is transitive.

Proof .:.

1. Suppose $x \in x$. By foundation, there is a $\in$-minimal element of $\{x\}$, which must be x. So any $y \in x$ has $y \notin \{x\}$ by minimality. But $x \in x$ has $x \in \{x\}$, so we have a contradiction.
2. Consider the set $\{x_0, \cdots, x_n\}$, which exists by finite applications of union and pairing. This has no $\in$-minimal element, since any x_i has $x_{i-1} \in x_i$ for $i > 0$ or else $x_n \in x_i$ for $i = 0$.
3. If x is transitive, then every element $y \in x$ is a subset of x. Hence if $y \neq \emptyset$ is $\in$-minimal, then there is some $z \in y \in x$, which yields $z \in x$ and $z \in y$, contradicting the minimality of y. Hence any $\in$-minimal element must be $\emptyset$.
4. We know that x being transitive implies $x \cup \{x\}$ is transitive. For the other direction, if $x \cup \{x\}$ is transitive, then any $a \in b \in x \cup \{x\}$ must have either $a \in x$ or $a = x$. But a cannot equal x without us having a finite loop: either $x \in b \in x$ or $x \in b = x$. Hence $a \in b \in x \cup \{x\}$ requires $a \in x$. This clearly implies that x is transitive since $a \in b \in x \subseteq x \cup \{x\}$ implies $a \in x$. ⊣

Important for later is the idea that any set is contained in a transitive set, which should seem rather clear: just continually add in the elements missing. To formalize this, however, we need some more ideas in general: the natural numbers. In general, we need ideas which will take the form of ordinals. In particular, we need a better idea of how to talk about *rank*. If $\emptyset$ is the base of the universe, then $\{\emptyset\}$ is just above it, and so has a rank one higher. Similarly, collections built from these like $\{\emptyset, \{\emptyset\}\}$ and $\{\{\emptyset\}\}$ are a rank higher than that.

§ 2 F. Formula abbreviations

We will often make abbreviations to our formulas to change their domain of discourse. For example, instead of writing $\forall x (x \in A \to \varphi(x))$, we will write $\forall x \in A\ \varphi(x)$. Similarly, instead of $\exists x (x \in A \wedge \varphi(x))$, we will write $\exists x \in A\ \varphi(x)$. These are standard translations of the more natural language ways of phrasing the formulas: "for all x in A, $\varphi(x)$ is true" and "there is an x in A where $\varphi(x)$ is true". We may also do this with other properties. For example, $\forall x < a\ \varphi(x)$ stands for $\forall x (x < a \to \varphi(x))$. Mostly this just serves to simplify formulas and make them easier to read, which we have already done with other abbreviations like $\subseteq$, $\cup$, and so forth.

Section 3. *Well-orders and Ordinals*

We will primarily be working with well-orders. Ordinals themselves are the "canonical" well-orders in that they are well-ordered by membership. They will also be special transitive sets, giving some credence to the axiom of foundation, since these canonical examples of transitive sets are well-ordered.

3 • 1. Definition

A relation R is a *well-order* iff R is linear, and well-founded.

We will see later that all well-orders on their domain and range are isomorphic to ordinals with the membership relation. First we must figure out what the ordinals are, and what properties they have.

§ 3 A. Introducing ordinals

3 A • 1. Definition

A set α is an *ordinal* iff α is transitive, and $\in$ well-orders α.

Note by foundation that $\in$ is well-founded on any set α. In the absence of the axiom of foundation, the requirement that $\in$ be well-founded isn't redundant. For the remainder of this section, we will not assume the axiom of foundation to show that the ordinals behave the same regardless. The well-founded property of membership on ordinals is used extensively in the arguments below. In essence, the results say that collection of ordinals themselves is linearly ordered by $\in$, rather than just each individual ordinal.

3 A • 2. Result

Let α, β be ordinals. Therefore,

1. Any $\gamma \in \alpha$ is an ordinal.
2. $\alpha \in \beta \vee \alpha = \beta$ is equivalent to $\alpha \subseteq \beta$.
3. $\alpha \in \beta, \beta \in \alpha$, or $\alpha = \beta$.
4. $\alpha \cup \beta$ is an ordinal.

Proof .:.

- For $\delta \in \alpha$, suppose $y \in x \in \delta$. We know $y, x \in \alpha$. Since α is linearly ordered by $\in$, it follows that either $\delta \in y$ or $y \in \delta$. Clearly $\delta \in y$ is impossible by well-foundedness. Hence $y \in \delta$ verifies that δ is transitive. Anti-symmetry follows from antisymmetry on α: $\gamma \subseteq \alpha$. Similarly, totality follows from the totality on α.

- Clearly if $\alpha \in \beta$ or $\alpha = \beta$ then $\alpha \subseteq \beta$ by transitivity. So suppose $\alpha \subseteq \beta$ for α an ordinal, but that the conclusion fails: $\alpha \neq \beta$ and $\alpha \notin \beta$. Without loss of generality, take β as the least failure in the sense that for each $\alpha' \in \beta, \alpha \subseteq \alpha'$ implies $\alpha \in \alpha'$ or $\alpha = \alpha'$ (to do this, take any ordinal β_0 witnessing the failure, and then consider the subset $\{\beta \in \beta_0 : \beta \text{ has it fail}\}$ and thus take a minimal element β by well-foundedness of $\in$ on ordinals).

  Consider $\beta \setminus \alpha$ as a subset of β. Since β is well-ordered by $\in$, there is a least element $\alpha' \in \beta \setminus \alpha$. Now suppose $\gamma \in \alpha$. Clearly $\gamma \in \alpha'$ by totality of $\in$ on β. Hence $\alpha \subseteq \alpha'$. By minimality of $\beta, \alpha \in \alpha'$ or $\alpha = \alpha'$. Therefore $\alpha \in \beta$, a contradiction.

- Let α be fixed. Let β be an ordinal with $\alpha \notin \beta, \alpha \neq \beta$, and $\beta \notin \alpha$. Without loss of generality, take β as the least failure in the sense that for each $\alpha' \in \beta, \alpha \in \alpha', \alpha = \alpha'$ or $\alpha' \in \alpha$ (to do this, just take any ordinal β_0 witnessing the failure, and then consider the subset $\{\beta \in \beta_0 : \beta \text{ has it fail}\}$ and thus take a minimal element β by well-foundedness).

  Clearly if $\alpha \subseteq \alpha' \in \beta$ for any $\alpha' \in \beta$, then (2) yields that $\alpha \in \beta$. So then $\alpha' \in \alpha$ for every $\alpha' \in \beta$. But then $\beta \subseteq \alpha$ so that $\beta \in \alpha$ or $\beta = \alpha$ by (2), again, a contradiction.

- That $\alpha \cup \beta$ is transitive is immediate: any $y \in \alpha \cup \beta$ has $y \in \alpha$ or $y \in \beta$. So if $x \in y$, then $x \in \alpha$ or $x \in \beta$ and hence $x \in \alpha \cup \beta$. Well-foundedness follows from the property holding on α and on β: for any subset X, $\alpha \cap X$ has a minimal element α_X and $\beta \cap X$ has a minimal element β_X, and one of these must be minimal for $\alpha \cup \beta$. Antisymmetry is trivial. Totality follows from (1) and (3). ⊣

Some easy examples of ordinals can be gotten from Subsection 2 E. In particular, $\emptyset$ is an ordinal, and we have the following result.

3 A • 3. Result

Let α be an ordinal. Therefore $\alpha \cup \{\alpha\}$ is an ordinal.

Proof .:.

We know by Corollary 2 E • 4 that $\alpha \cup \{\alpha\}$ (or rather the memberhsip relation on it) is transitive. So all that suffices to be shown is antisymmetry, and totallity of $\in$. Since antisymmetry is vacuously true for well-founded relations, as in Corollary 2 E • 4, we only need to show totality. But this follows from Result 3 A • 2: all elements of $\alpha \cup \{\alpha\}$ are ordinals, and so can be related by $\in$.

In particular, for $\emptyset$, we have $\{\emptyset\}$, $\{\emptyset, \{\emptyset\}\}$, and so on as ordinals. To make the notation a bit nicer, we will use the extremely suggestive notation below.

3 A • 4. Definition

For α an ordinal, write $\alpha + 1$ for $\alpha \cup \{\alpha\}$. Write $\beta < \alpha$ for $\beta \in \alpha$. Write $\emptyset = 0$.

Hence 0, $1 = 0 + 1$, $2 = 1 + 1$, $3 = 2 + 1$ are all ordinals. Note further that then every ordinal $\alpha = \text{pred}_<(\alpha)$ so that, for example, $5 = \{0, 1, 2, 3, 4\}$ (which has five elements). Note that the use of $+1$ is appropriate here as a kind of successor operation.

3 A • 5. Corollary

Let α be an ordinal. Therefore there is no ordinal β between α and $\alpha + 1$.

Proof .:.

Obviously, $\alpha \in \beta \in \alpha + 1$ requires $\beta = \alpha$ or $\beta \in \alpha$. Since $\alpha \notin \alpha$ by well-foundedness, we must have $\beta \in \alpha$, contradicting antisymmetry and that $\alpha + 1$ is an ordinal. ⊣

So far we are able to construct $n = 1 + \cdots + 1$ (n additions of 1) for each natural number n. But (provably) we can't show that the set of all of these ordinals exists from the axioms thus far. To do this, we must introduce the axiom of infinity: that there exists an infinite set of these.

3 A • 6. Definition (Axiom)

(*Infinity*) The set of natural numbers (or a set containing them) exists: $\exists N(\emptyset \in N \wedge \forall x \in N(x \cup \{x\} \in N))$.

The definition isn't able to properly say that the set of natural numbers exists without the notion of an ordinal. So we have to note the following result to then define the set of natural numbers. Clearly the result follows from foundation, but to get better acquainted with ordinals, we don't resort to this fact.

3 A • 7. Theorem

For any non-empty set X of ordinals,
- $\sup X := \bigcup X$ is an ordinal, and $\bigcup X \geq \alpha$ for each $\alpha \in X$;
- $\inf X := \bigcap X$ is an ordinal, and $\bigcap X \leq \alpha$ for each $\alpha \in X$;
- $\inf X \in X$ so that $\inf X = \min X$ is the minimum element of X.

Proof .:.

- It's clear that $\sup X \geq \alpha$ for each $\alpha \in X$. To see that $\sup X$ is an ordinal, transitivity follows from the transitivity of each ordinal in X: $x \in y \in \sup X$ has $y \in \alpha$ for some α and hence $x \in \alpha \subseteq \sup X$ implies

$x \in \sup X$. Antisymmetry is trivial, and totality follows easily from Result 3 A • 2.

- It should be clear that then $\inf X \leq \alpha$ for each $\alpha \in X$. To see that $\inf X$ is an ordinal, if $y \in x \in \alpha$ for each $\alpha \in X$, then $y \in \alpha$ for each $\alpha \in X$ so that $\inf X$ is transitive. Antisymmetry is again trivial, and totality is again easy to see as $\inf X$ is still a set of ordinals by Result 3 A • 2.

- Since every $\alpha \in X$ has $\alpha \leq \sup X$, it's easy to see that $\alpha < \sup(X) + 1$ so that $X \subseteq \sup(X) + 1$. As $\in$ is well-founded on $\sup(X) + 1$, it follows that X has a minimal element $\min X$, which is an ordinal. As $\bigcap X \leq \min X$, by (2), it suffices to show $\min X \subseteq \bigcap X$. But this is clear: every element $\alpha \in X$ has $\min X \subseteq \alpha$ so that $\min X \subseteq \bigcap X$. Hence $\min X = \bigcap X \in X$. ⊣

Thus far, we've only seen ordinals where $\sup X = \max X \in X$ or else $X = \emptyset$. But this won't always be true in general. In fact, there is a whole class of ordinals where this is false. Such ordinals are called limit ordinals, and in fact all ordinals can be broken down in this way. As a hint of what to come, the set of natural numbers will be a limit ordinal, and in fact the least such.

3 A • 8. Definition

Let $\alpha \neq 0$ be an ordinal. α is a *successor* ordinal iff $\alpha = \beta + 1$ for some ordinal β. α is a *limit* ordinal iff $\alpha = \sup \alpha$.

This classifies all ordinals.

3 A • 9. Theorem

Let α be an ordinal. Therefore $\alpha = 0$, or $\alpha = \sup(\alpha) + 1$, or $\alpha = \sup \alpha$.

Proof .:.

Let $\alpha \neq 0$. If $\alpha = \sup \alpha$, then for each $\beta < \alpha$, there is an $\gamma < \alpha$ with $\beta < \gamma$. In particular, by Corollary 3 A • 5, $\beta + 1 < \alpha$. So it's easy to see that $\alpha = \sup \alpha$ is equivalent to $\forall \beta < \alpha\ (\beta + 1 < \alpha)$. So if $\alpha \neq \sup \alpha$, there is some $\beta < \alpha$ with $\beta + 1 \not< \alpha$. Thus $\beta < \alpha \leq \beta + 1$ so by Corollary 3 A • 5, $\alpha = \beta + 1$. But then for every $\gamma < \alpha$, $\gamma \leq \beta$, implying $\beta = \sup \alpha$ and thus $\alpha = \sup(\alpha) + 1$. ⊣

Let's now collect the major properties of ordinals that we know so far.

3 A • 10. Theorem

For all ordinals α, β,

1. α is a set of ordinals;
2. $\alpha = 0$, α is a successor ordinal, or α is a limit ordinal;
3. 0 is the least ordinal;
4. the ordinals are well-ordered by $\in$;
5. $\alpha \cup \beta = \max(\alpha, \beta)$;
6. $\alpha \cap \beta = \min(\alpha, \beta)$;
7. $\alpha \leq \beta$ iff $\alpha \subseteq \beta$ (although not all sets $x \subseteq \beta$ are ordinals);
8. $\inf \alpha \leq \sup \alpha \leq \alpha < \alpha + 1$, and for $\alpha > 0$, $\inf \alpha < \sup \alpha$.

Proof .:.

1. follows from Result 3 A • 2 (1).
2. follows from Theorem 3 A • 9.
3. follows from Definition 3 A • 4: $\gamma < 0$ implies $\gamma \in \emptyset$, which is always false.
4. has linearity follow from transitivity and Result 3 A • 2 (2). To show well-foundedness, let X be a non-empty set (or class) of ordinals. Taking $\alpha \in X$ yields that $X \cap (\alpha + 1) \subseteq \alpha + 1$ which then has a least element $\beta \in X \cap (\alpha + 1)$. Any least element $\gamma \in X$ must have $\gamma \leq \alpha$ and thus $\gamma \in X \cap (\alpha + 1)$ so that β is the least element of X.
5. follows from Result 3 A • 2 (4) and (2).

6. follows from Theorem 3 A • 7.
7. follows from Result 3 A • 2 (4).
8. follows from Theorem 3 A • 7.

Now formally, we've defined a well-order to be a certain kind of set, which would make (4) false: the collection of all ordinals doesn't constitute a set. But it's easy to see what is meant by $\in$ well-ordering the ordinals (just the defining conditions without the additional requirement that the relation—$\in$ here—be a set).

3 A • 11. Result (Burali–Forti Paradox)

$\neg\exists s \forall x (x \text{ is an ordinal} \to x \in s)$. Informally, the collection Ord of all ordinals is not a set. In particular, there is no largest ordinal.

Proof .:.

There is no largest ordinal, since the largest ordinal α has $\alpha + 1 > \alpha$ by the reasoning above: $\alpha + 1 = \alpha$ implies $\alpha \in \alpha$, contradicting well-foundedness (even a set $\{\alpha\}$ has no least element, since a least element β requires $\forall z \in \beta\ (z \notin \alpha)$, which isn't true for $\beta = \alpha$).

To show that Ord can't be a set, by Theorem 3 A • 10, $\in$ well-orders Ord. Since each $\alpha \in \text{Ord}$ is transitive, it follows that $\alpha \subseteq \text{Ord}$ and hence Ord is transitive. Therefore Ord is an ordinal. But then Ord is the largest ordinal, contradicting the idea above. ⊣

Let's return to the idea of natural numbers. Notice that by our classification, every natural number is a successor ordinal, and in particular is of the form $0 + 1 + \cdots + 1$ for some (natural) number of $+$s.

3 A • 12. Definition

Write ω for the least limit ordinal, the set of natural numbers.

To see why ω should be the set of natural numbers, note that the supremum of the natural numbers must be a limit ordinal: n is a natural number implies $n + 1$ is too, so if $n < \sup \mathbb{N}$ then $n + 1 < \sup \mathbb{N}$, meaning $\sup \mathbb{N}$ is a limit ordinal. Moreover, $\sup \mathbb{N}$ must be the least limit ordinal, since every $n < \sup \mathbb{N}$ is a natural number, which means it's either a successor or 0. So this implies $\omega = \mathbb{N}$, but we haven't shown that ω actually exists, yet.

3 A • 13. Result

ω, the set of natural numbers and the least limit ordinal, exists.

Proof .:.

Let N be as in the axiom of infinity. Take the subset $N' = \sup\{\alpha \in N : \alpha \text{ is an ordinal}\}$ so that N' is an ordinal. We need to show that $\omega \leq N'$. If N' has a limit ordinal below it, then clearly ω is least by definition. So if N' has no limit ordinals below it, we want to show that $N' = \omega$.

Let $\alpha \in N'$ be the least such that $\alpha \in \omega \setminus N'$. As ω is the least limit ordinal, α must be a successor or 0. If $\alpha = 0$, then $0 \in N'$ by hypothesis that $1 \in N$ so $0 < 1 \leq N'$. If $\alpha = \beta + 1$, then $\beta \in \omega$. By the minimality of β, $\beta \in N'$ so that by the hypothesis on N, $\beta + 1 = \alpha \in N$ and hence $\alpha + 1 \in N$ so that $\alpha < N'$, a contradiction. Therefore $\omega \subseteq N'$. But then as ordinals, $\omega < N'$ or $\omega = N'$. Since N' has no limit ordinals below it, $N' = \omega$.⊣

It would seem that the reasoning alone gives the existence of ω, but really the idea only only characterizes ω. We still need the existence of N as in the axiom of infinity to ensure the existence of ω.

With the natural numbers at our fingertips, we can show that ω satisfies all the usual properties that we want, namely the axioms of peano arithmetic, PA. To do this, we need a notion of addition and multiplication of ordinals. To do this, we need a better way of defining operations on ω.

As a side note, we have a characterization of ω in meta-theoretic terms (able to be reached from 0 by finite applications of adding 1). What we've done now is show that in **V**, this coincides with the characterization of ω as the least limit ordinal. This formal characterization, however, isn't necessarily the set of natural numbers. Consider the following from model theory: in the language $\texttt{FOL}(\in, c)$ where c is a constant symbol, the theory of set theory adjoined with "$\omega > c > 1 + \cdots + 1$" ($n$ times) for each real-world natural number $n \in \mathrm{V}$ yields a theory T_n that is consistent assuming that set theory is consistent (just interpret c as $n + 1$ in V). Hence every finite subset of the theory $T = \{\varphi : \varphi \text{ in } T_n \text{ for some } n \text{ a natural number}\}$ is consistent so that T itself has a model by Compactness (1 D • 2). But in this model $\mathsf{M} \vDash T$, we have $\mathsf{M} \vDash$ "$\omega > c^{\mathsf{M}} > 1 + \cdots + 1$" ($n$ times) for each real-world natural number $n < \omega$. So ω^{M} can't be the same as ω in the real-world V. All of this is to say that we must be careful about using our *intuitive, meta-theoretic* characterization of ω to *formally* prove things about it from set theory. To ensure that we can prove all of the intuitive properties of ω *formally*, we resort to the principle of induction.

§ 3 B. Finitary recursion and induction

Recall the defining property of ω: if $0 \in \omega$, and $n \in \omega$ then $n + 1 \in \omega$ (and this is all there is in ω). In particular, this yields the following result, called the principle of induction.

3 B • 1. Theorem (Induction on ω)

Let $\varphi(x)$ be a $\texttt{FOLp}(\in)$-formula. Suppose $\varphi(0)$ and $\varphi(n) \to \varphi(n + 1)$. Therefore $\forall n \in \omega\ \varphi(n)$.

Proof .:.

Consider the set $X = \{n \in \omega : \neg\varphi(n)\}$. This has a least element $x \in X$. Note that $x \neq 0$ by the hypothesis. Since ω is the least limit ordinal, $x = \sup(x) + 1$ is a successor. Therefore by minimality, $\varphi(\sup(x))$ holds and so $\varphi(\sup(x) + 1)$ holds, contradicting that $\sup(x) + 1 = x \in X$. ⊣

Really, this is just a consequence of ω being well-ordered. But this reflects the properties of arithmetic that ω should have. The key thing here is that by specifying what happens at 0, and what happens at successor stages, we can define something on all of ω. This idea is referred to as *recursion*.

The formal statement of recursion is long and clunky. So to better understand it, we give some examples. Firstly, we would normally define addition by n by $f_n(x) = x + 1 + \cdots + 1$ where we add 1 ns. The issue with this is that this definition is informal and meta-theoretic, in some sense. It's not clear how we would define this function purely in terms of set theory without resorting to "n-times". Surely for each x this makes sense, but the map sending $n \mapsto f_n$ isn't so obviously well defined (consider non-standard models with different ωs). To get around this, for each $x < \omega$ consider the map defined by $f_x(0) = x$ and $f_x(n + 1) = f_x(n) + 1$. Using induction, this defines f_x on all $n < \omega$. Moreover, intuitively, this f_x satisfies $f_x(n) = x + n$.

Once we have f_x for each $x < \omega$, we can consider the map sending $\langle x, n\rangle$ to $f_x(n)$. This map, call it '$+$', sends $\langle x, n\rangle$ to $x + n$ in the usual sense.

To define this whole process more formally, what we're doing is specifying what happens at the start, and then what happens at successor stages. So we are given functions f and g, and we define the function h starting with $f(0)$, and finding the next values based on g and the previous value:

$$h(0) = f(0)$$
$$h(n + 1) = g(n, h(n)).$$

So to calculate $h(2)$, we start with $h(0) = f(0)$, and then calculate $h(1) = g(1, f(0))$, and then calculate $h(2) = g(2, g(1, f(0)))$. In principle, we could then keep going to define $h(4)$, $h(5)$, and so on, meaning $h(n)$ will be some particular number for each n. This means the function h is determined by these conditions in the sense that it is the unique function satisfying them. Formally, we have the following theorem. The proof of this theorem is very technical, and long, and not terribly illuminating, mostly just making precise and formal the intuitive idea of "starting and 0 and defining what happens next determines it on all of the natural numbers". It is included for those interested in the precise details, but for those uninterested, it can be skipped.

3 B • 2. Theorem (Recursion on ω)

Let f with $0 \in \operatorname{dom}(f)$ be a function. Let g be a function from ordered pairs with the first entries being natural numbers: $\omega = \operatorname{dom}(\operatorname{dom}(g))$. Therefore, there is a unique function h where $\operatorname{dom}(h) = \omega$ and

$$h(0) = f(0)$$
$$h(n+1) = g(n, h(n)).$$

Proof .:.

To show existence, we proceed by induction to show that for each $n \in \omega$, $\varphi(n, h)$ defines a unique function h_n, which is supposed to represent $h \restriction n$. Once we do this, we pull together all of the h_ns to define h.

Consider the formula $\psi(n, h)$ given formally below:

$$\operatorname{dom}(h) = n < \omega \wedge \forall k < n \left(\begin{matrix} \big(k = 0 \wedge \langle 0, f(0)\rangle \in h\big) \\ \vee \exists! v\ \exists m\ \big(k = m + 1 \wedge \langle m, v\rangle \in h \wedge \langle k, g(m, v)\rangle \in h\big) \end{matrix} \right).$$

Informally, $\psi(n, h)$ says

h is a function with domain n and obeys the recusive definition up to n.

One may easily check the following facts:

1. if $\psi(n, h)$, then h is a function;
2. if $\psi(n, h)$ and $m \leq n$, then $\psi(m, h \restriction m)$; and
3. if $\psi(n, h)$ for $n = n^* + 1$, then $\psi(n + 1, h \cup \{\langle n, g(n^*, h(n^*))\rangle\})$.

We want to now show that for each $n < \omega$, there is exactly one h with $\psi(n, h)$. This will allow us to use replacement to collect all of these approximations to the h of the theorem together.

Claim 1

$\forall n < \omega\ \exists! h\ \psi(n, h)$.

Proof .:.

There are two parts to this: the existence of h, and the uniqueness of h. Existence holds by induction: since $h_0 = \emptyset$ exists trivially, and h_{n+1} satisfying $\psi(n + 1, h_{n+1})$ exists by (3) above. So induction shows that for each $n < \omega$, there exists such an h where $\psi(n, h)$.

To show there is at most one h with $\psi(n, h)$, let $n + 1 < \omega$ is the least where this fails (it vacuously holds for $n = 0$). Thus we have two functions $h_0 \neq h_1$ where $\psi(n+1, h_0)$ and $\psi(n+1, h_1)$. Note by (2) above, $\psi(n, h_0 \restriction n)$ and $\psi(n, h_1 \restriction n)$ hold. So by the minimality of $n+1$, $h_0 \restriction n = h_1 \restriction n$. So the only place the two functions can differ is at n: $h_0(n) \neq h_1(n)$. But in satisfying ψ, we must have that for $k = n = m + 1$, $\langle k, g(m, h(m))\rangle \in h_0, h_1$, i.e. $h_0(n) = g(n, h_0(m)) = g(n, h_1(m)) = h_1(n)$, a contradiction. ⊣

Thus by replacement, we have the set $\{h_n : n \in \omega\}$ where $\psi(n, h_n)$ for each $n < \omega$. Therefore $\bigcup_{n \in \omega} h_n = h$ is a function with domain ω, and for each $n < \omega$, h satisfies $\psi(n, h \restriction n)$. Thus $h(0) = (h \restriction 1)(0) = f(0)$ and $h(n + 1) = (h \restriction n + 2)(n + 1) = g(n, (h \restriction n + 2)(n)) = g(n, h(n))$, showing that h shows the existence of such a function as in the theorem statement.

Now for uniqueness, suppose $h' \neq h$ also satisfied the hypothesis. Therefore for each $n < \omega$, $\psi(n, h' \restriction n)$ holds so that uniqueness of the parts yields $h' \restriction n = h \restriction n$ for each $n < \omega$. Hence $h'(n) = (h' \restriction n+1)(n) = (h \restriction n+1)(n) = h(n)$ for each $n < \omega$. Thus $h' = h$. ⊣

The above theorem isn't actually given in its fullest generality: we are allowed more variables. As long as the order we proceed in is well-founded, we are guaranteed the result by the same idea as above. To give a reason for this, we must consider the *transfinite* versions of these.

§ 3 C. Transfinite recursion and induction

The existence of limit ordinals is incredibly powerful, as it allows us to form larger and larger ordinals beyond just ω. To go further, we need a better way of defining or constructing these ordinals. To do this, we use the notion of transfinite recursion and induction. Intuitively, $\omega+1$, $\omega+2$, $\omega+3$, and so on have all been defined. If we wish to define $\omega+\omega$, we could do this as the least limit ordinal after ω, but this clumsy characterization isn't sustainable to define $\alpha+\beta$ for general ordinals $\alpha+\beta$. To do this, we use the characterization of ordinals into 0, successors, and limits. If we specify the definition at 0, at successors, and at limits, we will have defined it everywhere. The idea of transfinite recursion makes this explicit.

Again, first we have the fundamental property that allows us to do this: transfinite induction. The idea was already noticed in Theorem 3 A • 10 (4). But to make it explicit, we have the following theorem.

3 C • 1. Theorem (Transfinite Induction)

Let $\varphi(x,\vec{w})$ be a $\texttt{FOL}(\in)$-formula with $\vec{v}$ parameters. Suppose $\varphi(\alpha,\vec{v})$ holds whenever $\forall\beta<\alpha\ \varphi(\beta,\vec{v})$. Therefore for every ordinal α, $\varphi(\alpha,\vec{v}))$.

Proof .:.

Otherwise, take α the least such that $\neg\varphi(\alpha,\vec{v}))$. Thus for every $\beta<\alpha$, $\varphi(\beta,\vec{v})$. Hence by hypothesis $\varphi(\alpha,\vec{v})$, a contradiction. ⊣

This also applies to the natural numbers, but stated this way allows us to incorporate limit ordinals. If we had simply left the same sort of statement as in Induction on ω (3 B • 1), we wouldn't necessarily have the result for ω, much less all ordinals α. In particular, consider the property of being 0 or a successor ordinal. Clearly this holds for 0 and if it holds for α, it holds for $\alpha+1$. But this never allows one to reason their way to the limit ordinals: only successors of successors and so on.

To make the notion of transfinite recursion formal, we need three functions specifying what happens at stage 0, what happens at successor stages, and what happens at limit stages. This idea of breaking down into cases proceeds in precisely the same way as in Recursion on ω (3 B • 2). But there is a slightly easier way to state it formally. Rather than breaking down into more and more cases with more and more classifications, the main idea of recursion is just that we can calculate the next value from the previous ones. So the value at ω should be determined by the values on all $n<\omega$. Stated formally, this yields the much more compact version below.

3 C • 2. Theorem (Transfinite Recursion)

Let α be an ordinal. Let f be a function, writing $f(x)=\emptyset$ for $x\notin\mathrm{dom}(f)$. Therefore there is a unique function g with domain α such that for all $\beta<\alpha$, $g(\beta)=f(g\restriction\beta)$.

Proof .:.

Assuming existence, uniqueness follows easily by induction on α. For α the least such where this fails, there are then functions $g\neq g'$ where $g(\beta)=f(g\restriction\beta)$ and $g'(\beta)=f(g'\restriction\beta)$ for all $\beta<\alpha$. But by minimality of α, $g'\restriction\beta=g\restriction\beta$ so that $g(\beta)=f(g'\restriction\beta)=g(\beta)$, meaning $g(\beta)=g'(\beta)$ for all $\beta<\alpha$, and thus $g=g'$.

To show existence, proceed as in Recursion on ω (3 B • 2). In particular, consider the formula $\psi(\beta,g)$ which says that g is a function with domain β and $\forall\gamma<\beta\ (g(\gamma)=f(g\restriction\gamma))$. By induction on β, we can show $\exists!g\ \psi(\beta,g)$. To see this, let β be least where this fails. Hence for each $\gamma<\beta$, $\exists!g\ \psi(\gamma,g)$. By replacement we get a set $\{g_\gamma:\gamma<\beta\}$. One can easily see that $\psi(\gamma,g)$ implies $\psi(\delta,g\restriction\delta)$ for any $\delta<\gamma<\beta$. Hence the union $g=\bigcup\{g_\gamma:\gamma<\beta\}$ is a function with domain β, and one can easily check that for each $\gamma<\beta$, $g(\gamma)=f(g\restriction\gamma)$. Uniqueness follows from the uniqueness of each $g_\gamma=g\restriction\gamma$ as in Recursion on ω (3 B • 2). ⊣

The idea above actually extends to Ord in the sense we can get define an output on every ordinal α. Although we won't get a g such that $\mathrm{Ord}\subseteq\mathrm{dom}(g)$ (since g needs to be a set), we can still define what the output will be at any given α by considering the resulting function with domain $\alpha+1$. Uniqueness ensures that this output doesn't vary with the change in domain. So it makes sense to say that this defines a function on all of Ord, even though only the approximations to this function exist. Formally, we might say $\varphi(\alpha,y)$ holds iff $\exists g$ (g is a function with $\mathrm{dom}(g)=\alpha+1\wedge\forall\beta<\alpha+1(g(\beta)=f(g\restriction\beta))\wedge y=g(\alpha)$). The reasoning above tells us that $\forall\alpha\in\mathrm{Ord}\ \exists!y\varphi(\alpha,y)$. So this is the sense in which we have defined a function on all of Ord.

§ 3 D. A word on sequences and functions

Although much of this section has been stated in terms of functions, it's perhaps most intuitive to think of functions from ordinals as sequences: for each entry in a sequence, there is a subsequent entry, and there should always be a least point in the sequence where something happens. In most other branches of math, the only sequences that appear are those of length ω, or else finite.

3 D • 1. Definition

A *sequence* is a function f with $\mathrm{dom}(f)$ as an ordinal (or $\mathrm{dom}(f) = \mathrm{Ord}$, in which case f is a class). The *length* of a sequence is its domain.

This notion of a sequence is incredibly important if we want to define functions with more than just finitely many inputs. Thus far, if we wanted a function from tuples in A, B, and C to D, we'd need to consider $f : A \times B \times C \to D$. The introduction of sequences allows us to consider tuples instead as sequences: $\langle a, b, c\rangle$ can be identified with the function $f : 3 \to A \cup B \cup C$ where $f(0) = a$ and $f(1) = b$ and $f(2) = c$, identifying each entry with *where* it is in the tuple. And we can generalize this, allowing us to talk about infinite products.

3 D • 2. Definition

Let I be a set, and suppose $\{A_i : i \in I\}$ is a family of sets. Therefore the *cartesian product* $\prod_{i\in I} A_i$ is the set of functions $f : I \to \bigcup_{i\in I} A_i$ such that $f(i) \in A_i$ for each $i \in I$.

In particular, for α an ordinal, we write $A^\alpha = \prod_{\beta<\alpha} A$, generalizing $A^n = A \times \cdots \times A$ (n times) for $n < \omega$. Note that the finite product of non-empty sets is non-empty. That infinite products of non-empty sets are non-empty is equivalent to an axiom yet to be introduced. We will have no need of it for now, but it should be noted.

Really, the inherent notion of a sequence just comes from any well-order. So we should investigate further what well-orders exist. As it turns out, the ordinals will exhaust all the well-orders in V.

§ 3 E. The model theory of well-orders

We have defined what it means for a structure $\mathsf{A} = \langle A, R\rangle$ to be a well-order: R well-orders $A = \mathrm{dom}(R) \cup \mathrm{ran}(R)$. This property, however, is not expressible in first-order logic alone. To see this, we use compactness and the existence of ω.

3 E • 1. Result

Let σ be a signature with a binary relation symbol R. Let T be a $\mathtt{FOL}(\sigma)$-theory such that T contains the axioms of partial orders. Therefore, if T has an infinite, well-ordered model, then T has an ill-founded (i.e. not well-founded) model. Hence being a well-order isn't $\mathtt{FOL}$-expressible.

Proof .:.

Let $\mathsf{A} \vDash T$ be an infinite well-order. Using Recursion on ω (3 B • 2), let a_0 be some, fixed R^{A}-minimal element of A. By well-foundedness, for $n+1$, let a_{n+1} be the R^{A}-least element $a \in A$ such that

$$\forall i \leq n \ (\mathsf{A} \vDash \text{“}a_i \ R^{\mathsf{A}} \ a\text{”}).$$

Now in the expanded signature $\mathcal{L} = \{R\} \cup \{c_n : n \in \omega\}$ with new constant symbols c_n for each $n < \omega$, consider the theory $T' = T \cup \{c_{n+1} \ R \ c_n : n < \omega\}$. Note that any model of T' is ill-founded. Since any model of T' is also a model of T, it suffices to show that T' is consistent. To do this, we use Compactness (1 D • 2).

For any finite subset $\Delta \subseteq T'$, there is a largest $n < \omega$ where c_n occurs in a formula of Δ. Taking N to be this, we can interpret Δ in the expansion A' of A where c_n is interpretted as a_{N-n}: $\mathsf{A} \vDash \text{“}a_{N-(n+1)} \ R \ a_{N-n}\text{”}$ so clearly $\mathsf{A}' \vDash \text{“}c_{n+1} \ R \ c_n\text{”}$. Since $\mathsf{A}' \vDash T$, it follows that $\mathsf{A}' \vDash \Delta$, and thus Δ is consistent. As Δ was arbitrary, T' is consistent. By Completeness (1 D • 1), there is a model $\mathsf{B} \vDash T'$ which then models T, but is ill-founded. ⊣

So the property of being a well-order is a property of the the set theoretic universe. Depending on the (non-**V**) model of set theory, certain sets may or may not be well-founded, because the models don't have the set witnessing the ill-foundedness. This is a weakness of first-order logic, but it is no challenge to the legitimacy of the concept. Really, this idea just expresses the inadequacy of first-order formulas to properly characterize these notions. This is a common part of logic, as even group theory is subject to the limitation: the property of being a cyclic group isn't first-order expressible, for example. This is merely something we must live with.

Clearly, however, being a well-order is preserved under isomorphisms. In fact, our goal here will be to show that the ordinals are the *canonical* well-orders in the sense that every well-order is isomorphic to a particular ordinal (under membership). To do this, we proceed in a similar way as when we introduced ordinals. Before this, we introduce some definitions that should be familiar from model theory.

3 E • 2. Definition

Let $\mathsf{A} = \langle A, <_A \rangle$ and $\mathsf{B} = \langle B, <_B \rangle$ be structures where $<_A$ and $<_B$ are relations.

- A function $f : A \to B$ is a *homomorphism* iff $a <_A a' \to f(a) <_B f(a')$ for every $a, a' \in A$.
- A function $f : A \to B$ is an *embedding* iff $a <_A a \leftrightarrow f(a) <_B f(a')$ for every $a, a' \in A$ and f is injective.
- A function $f : A \to B$ is an *isomorphism* iff f is an embedding, and f is surjective.

If A is a linear order, an *initial segment* of A is a substructure with universe $\mathrm{pred}_{<_A}(a_0) = \{a \in A : a <_A a_0\}$ for some $a_0 \in A$.

So for each ordinal α, $\mathrm{pred}_{\in}(\alpha) = \alpha$. Now we consider the following result about well-orders. Note that for $X \subseteq A$ and $<_A \subseteq A \times A$, we continue to write $\langle X, <_A \rangle$ for the sake of readability when really we mean $\langle X, <_A \cap (X \times X) \rangle$. Note that if A is a well-order, then its initial segments are well-orders too.

3 E • 3. Lemma

Let $\mathsf{A} = \langle A, <_A \rangle$ be a well-order. Let $a \in A$. Write $\mathrm{pred}_{<_A}(a)$ for $\{x \in A : x <_A a\}$. Therefore $\langle \mathrm{pred}_{<_A}(a), <_A \rangle$ is a well-order.

This can be seen just by noting that all of the properties are inhereted from the well-order on A: transitivity, antisymmetry, and totality all hold since we're taking all variables in $\mathrm{pred}_{<_A}(a)$, and well-foundedness also clearly holds, since we're taking a subset of $\mathrm{pred}_{<_A}(a)$. In fact, for any subset $X \subseteq A$, $\langle X, <_A \rangle$ is well-founded if $\langle A, <_A \rangle$ is.

3 E • 4. Lemma

Let $\mathsf{A} = \langle A, <_A \rangle$ be a well-order. Therefore, $\mathsf{A} \ncong \langle \mathrm{pred}_{<_A}(a), <_A \rangle$ for any $a \in A$.

Proof .:.

Let $f : A \to A_{<a}$ be an isomorphism where $\mathsf{A}_{<a} = \langle \mathrm{pred}_{<_A}(a), <_A \rangle$. Consider the subset $X = \{x \in A : f(x) \neq x\}$. Note that X is non-empty, since $a \in X$, for example: $a \notin A_{<a}$ cannot be in the image of f. Consider the $<_A$-least element $x \in X$ so that $f``\mathrm{pred}_{<_A}(x) = \mathrm{pred}_{<_A}(x)$, but $f(x) \neq x$. By injectivity, it follows that $f(x) \not<_A x$ and thus $f(x) >_A x$ by totality. Since f is an isomorphism, there must be some $x' \in A$ where $x = f(x')$. But then $\mathsf{A}_{<a} \vDash$ "$f(x) >_A f(x')$" requires $\mathsf{A} \vDash$ "$x >_A x'$" as an embedding. But by minimality, this implies $f(x') = x' \neq x$, a contradiction. ⊣

Using this, we get the following, which will allow us to show that any two well-orders can be compared in the sense that they are either isomorphic to each other, or to an initial segment In particular, when we restrict an isomorphism to an initial segment, we get an isomorphism between initial segments.

3 E • 5. Lemma

Let $\mathsf{A} = \langle A, <_A \rangle$ and $\mathsf{B} = \langle B, <_B \rangle$ be well-orders. Let $f : A \to B$ be an isomorphism. Therefore, for any $a \in A$, $f \restriction \mathrm{pred}_{<_A}(a)$ is an isomorphism between $\langle \mathrm{pred}_{<A}(a), <_A \rangle$ and $\langle \mathrm{pred}_{<B}(b), <_B \rangle$ for some $b \in B$.

Proof .:.

Write $\mathsf{A}_{<a}$ for $\langle \mathrm{pred}_{<_A}(a), <_A \rangle$ and similarly for $b \in B$. Let $a \in A$ be $<_A$-least such that the result fails. Let b be the least element of $B \setminus f``A_{<a}$. We will show that $f``A_{<a} = B_{<b}$, and thus that $f \restriction A_{<a}$ is an isomorphism

between $\mathbf{A}_{<a}$ and $\mathbf{B}_{<b}$.

By minimality, $B_{<b} \subseteq f``A_{<a}$, so suppose the reverse doesn't happen: there is some $a_0 \in A_{<a}$ with $f(a_0) >_B b$. As an isomorphism, there is some $a' \in A$ with $b = f(a')$ so that $\mathbf{B} \vDash$ "$f(a_0) >_B f(a')$". As an embedding, this means $\mathbf{A} \vDash$ "$a_0 >_A a'$" so that $a' \in A_{<a}$, contradicting that $b \notin f``A_{<a}$. ⊣

3 E • 6. Lemma

Let $\mathbf{A} = \langle A, <_A \rangle$ and $\mathbf{B} = \langle B, <_B \rangle$ be two well-orders. Suppose $\mathbf{A} \not\cong \mathbf{B}$, and $\mathbf{B}$ is not isomorphic to an initial segment of $\mathbf{A}$. Therefore there is a unique $b_0 \in B$ with $\mathbf{A} \cong \langle \mathrm{pred}_{<_B}(b_0), <_B \rangle$.

Proof .:.

Write $\mathbf{A}_{<a}$ for $\langle \mathrm{pred}_{<_A}(a), <_A \rangle$ and similarly for $b \in B$. Uniqueness clearly holds by Lemma 3 E • 4: $\mathbf{A} \cong \mathbf{B}_{<b_0}$ and $\mathbf{A} \cong \mathbf{B}_{<b_1}$ implies $\mathbf{B}_{<b_1} \cong \mathbf{B}_{<b_0}$. So if $b_1 \neq b_0$, then $b_0 <_B b_1$ or $b_1 <_B b_0$, and we contradict Lemma 3 E • 4 in either case.

Now suppose existence fails. Without loss of generality, let $\mathbf{A}$ be minimal in the following sense: for every $a \in A$, there is a unique $b \in B$ such that $\mathbf{A}_{<a} \cong \mathbf{B}_{<b}$. (Otherwise just choose the least $a \in A$ where this fails, and consider the structure $\mathbf{A}_{<a}$ instead. This new structure still has $\mathbf{B}$ not isomorphic to an initial segment, nor isomorphic to it as a whole.) So let $f = \{\langle a, b \rangle : \mathbf{A}_{<a} \cong \mathbf{B}_{<b}\}$ be the function such that $\mathbf{A}_{<a} \cong \mathbf{B}_{<f(a)}$. Note that f must be injective since if $x <_A y$, then $\mathbf{A}_{<x}$ is an initial segment of $\mathbf{A}_{<y}$: thus $\mathbf{B}_{<f(x)} \cong \mathbf{A}_{<x} \not\cong \mathbf{A}_{<y} \cong \mathbf{B}_{<f(y)}$ by Lemma 3 E • 4.

Claim 1

f is an embedding. Given that f is already injective, we mean $x <_A y \rightarrow f(x) <_B f(y)$ for all $x, y \in A$.

Proof .:.

Otherwise, $f(x) \geq_B f(y)$ so that $\mathbf{A}_{<y} \cong \mathbf{B}_{<f(y)}$ is an initial segment of $\mathbf{B}_{<f(x)} \cong \mathbf{A}_{<x}$. Composing the isomorphisms, we get that $A_{<y}$ is isomorphic to an initial segment of $A_{<x}$, contradicting Lemma 3 E • 4. Explicitly, take $f_y : A_{<y} \rightarrow B_{<f(y)}$ and $f_x : B_{<f(x)} \rightarrow A_{<x}$ to be isomorphisms. By Lemma 3 E • 5, $f_x \restriction B_{<f(y)}$ is an isomorphism with an initial segment $A_{<a} \subseteq A_{<x}$ so that $f_x \circ f_y : A_{<y} \rightarrow A_{<a}$ is an isomorphism. ⊣

So all that suffices is to show that f is surjective onto some initial segment. f is an isomorphism between A and $\mathrm{im}\, f$. Taking b_0 the least element of $B \setminus \mathrm{im}\, f$, we get that $B_{<b_0} \subseteq \mathrm{im}\, f$ by minimality of b_0. To show that $\mathrm{im}\, f \subseteq B_{<b_0}$, suppose $f(a_0) >_B b_0$ so that there is an isomorphism $g : B_{<f(a_0)} \rightarrow A_{<a_0}$. Thus $g \restriction B_{<b_0}$ is an isomorphism between $\mathbf{B}_{<b_0}$ and $\mathbf{A}_{<a}$ for some $a \in B$ by Lemma 3 E • 5. But then $b_0 = f(a)$ contradicts that $b_0 \notin \mathrm{im}\, f$. Hence $\mathrm{im}\, f \subseteq B_{<b_0}$, and so we have equality, and thus f is an isomorphism. ⊣

Stated more loosely, for any two well-orders, either they are isomorphic, or one is isomorphic to an initial segment of the other. As a corollary of this, the ordinals exhaust all of the well-orderings in $\mathbf{V}$.

3 E • 7. Corollary

For every well-order $\mathbf{A}$, there is a unique ordinal α such that $\mathbf{A} \cong \langle \alpha, \in \rangle$.

Of course, Ord is well-ordered by $\in$, but Ord $\notin$ V by Burali–Forti Paradox (3 A • 11), so this isn't an issue: every quantifier ranges over sets: we're only considering structures in V while $\langle \mathrm{Ord}, \in \rangle \notin$ V.

3 E • 8. Definition

Let $\mathbf{A}$ be a well-order. The *order type* of $\mathbf{A}$ is the unique ordinal α with $\mathbf{A} \cong \langle \alpha, \in \rangle$.

More than just getting a unique order type, we also get that the isomorphism is unique:

3 E • 9. Result

Let $\mathbf{A}$ be a well-order, and $f : A \rightarrow \alpha$ and $g : A \rightarrow \alpha$ isomorphisms. Therefore, $f = g$.

Proof .:.

Assume not, and let $a \in A$ be $<_A$-minimal such that $f(a) \neq g(a)$. For the sake of definiteness, assume $f(a) < g(a)$. Since g is an isomorphism, there is some $b \in A$ where $g(b) = f(a) < g(a)$. In other words, $\mathsf{V} \models$ "$g(b) < g(a)$" so that as an embedding, $\mathsf{A} \models$ "$b <_A a$" so by minimality of a, $f(a) = g(b) = f(b)$, contradicting that f is injective. $\dashv$

There are, of course, other questions one can ask of well-orders in the context of model theory, like when two ordinals are elementarily equivalent under membership, for example. But for now, we will only make use of the fact that well-orders are isomorphic to ordinals.

Section 4. *Other Well-founded Relations*

Recall Axiom (2 E • 3), the axiom of foundation. To further motivate why this axiom should be true, we will show the following result, which holds even in the absence of foundation. In essence, the result says that all well-founded models of set theory in **V** are isomorphic to transitive sets. So the axiom of foundation in some sense takes the converse to be true: all transitive sets are well-founded.

4 • 1. Theorem (The Mostowski Collapse)

Let $\mathbf{A} = \langle A, <_A \rangle$ (in V) be well-founded such that $\mathbf{A}$ satisfies the axiom of extensionality. Therefore $\mathbf{A} \cong \langle T, \in \rangle$ for a unique transitive set T.

Although we can prove the theorem outright at this point, to get a better perspective on what is going on with the proof, we will introduce a useful idea: rank. Although all well-orders are isomorphic to ordinals, well-founded, extensional structures are not in general. But they can still make use of ordinals according to *chains*, which are then well-ordered. Really, this just means indexing the *levels* of the structure like with a tree.

The most fundamental idea behind rank functions is given by Transfinite Recursion (3 C • 2), and so often we want the process to stop at some ordinal. The following lemma, a consequence of the axiom of replacement, will be useful in doing this. Note that the lemma further reinforces the idea that some collections are simply "too big" to be sets. In essence, we will use this to say that there can't be Ord-many levels of an ewfset.

4 • 2. Lemma

Let A be a set. Therefore there is no surjection $f : A \to \mathrm{Ord}$.

Proof .:.

Otherwise, the formula $\varphi(x, y, f)$ given by $\langle x, y \rangle \in f$ defines a function on A. By replacement, $f"A = \mathrm{Ord}$ exists, contradicting Burali–Forti Paradox (3 A • 11). ⊣

The general idea of a rank function is given below.

4 • 3. Lemma

Let $\mathbf{A} = \langle A, <_A \rangle$ be well-founded. Therefore there is a unique function $f : A \to \mathrm{Ord}$ such that $f(a)$ is 0 if a is $<_A$-minimal, and otherwise $f(a) = \sup\{f(b) + 1 : b <_A a\}$.

Proof .:.

Uniqueness is immediate: for f, g two such functions and a $<_A$-minimal where $f(a) \neq g(a)$, we have that $f(a) = \sup\{f(b) + 1 : b <_A a\}$. By minimality of a, this supremum is $\sup\{g(b) + 1 : b <_A a\} = g(a)$, which means $g(a)$ is $f(a)$, a contradiction.

We construct such an f by transfinite recursion. Firstly, as $\mathbf{A}$ is well-founded, define by transfinite recursion

$$\begin{aligned} X_0 &= \emptyset \\ X_{\alpha+1} &= \left\{a \in A : a \text{ is } <_A \text{-minimal in } A \setminus \bigcup\nolimits_{\beta \leq \alpha} X_\beta \right\} \\ X_\gamma &= \emptyset, \text{ for } \gamma \text{ a limit.} \end{aligned}$$

If $X_{\alpha+1}$ is ever empty, then we stop, and so $X_\alpha = A$. Then we define $f : A \to \mathrm{Ord}$ by taking $f(x)$ to be the least (and only) α such that $x \in X_{\alpha+1}$. By Lemma 4 • 2, this process stops at some $\alpha \in \mathrm{Ord}$ so that $f \in \mathrm{V}$.

Note that $x, y \in X_\alpha$ implies x and y are $<_A$-incomparable: $x \not<_A y$ and $y \not<_A x$ (otherwise, they wouldn't be minimal). Hence $f(x) = f(y)$ implies x and y are $<_A$-incomparable.

Moreover, the contrapositive then tells us that if $x <_A y$, then $f(x) \neq f(y)$, and in fact $f(x) < f(y)$, as otherwise $f(y) < f(x)$ implies $x <_A y$ is not actually $<_A$-minimal in $A \setminus \bigcup_{\beta < f(y)} X_\beta$, because $x \in A \setminus$

$\bigcup_{\beta<f(x)} X_\beta \subseteq A \setminus \bigcup_{\beta<f(y)} X_\beta$. Therefore, $f(a) \geq \sup\{f(b)+1 : b <_A a\}$ for all $a \in A$.

Now if $f(a) > \beta = \sup\{f(b)+1 : b <_A a\}$, then by the definition of the X_αs, a wasn't minimal in $A \setminus \bigcup_{\gamma<\beta+1} X_\gamma$, meaning that there is some $b \in A \setminus \bigcup_{\gamma<\beta+1} X_\gamma$ with $b <_A a$. Taking a $<_A$-minimal such b yields that $f(b) = \beta+1$, contradicting the definition of β. ⊣

The point of having a rank function is to proceed by induction on the levels.. Indeed, the proof above just defines the function f by induction on the levels of A. So if we can prove something for the elements inductively by level, then we can prove it for the whole ewfset. So we have the following definition. By uniqueness, we are justified in using "the" rank function, and defining the following as aspects of the structure alone, independent of any choice of rank function.

4 • 4. Definition

For well-founded $\mathsf{A} = \langle A, <_A\rangle$, the *rank function* on A is the function $\mathrm{rank} : A \to \mathrm{Ord}$ such that

- $\mathrm{rank}(a) = 0$ if a is $<_A$-minimal; and
- $\mathrm{rank}(a) = \sup\{\mathrm{rank}(b)+1 : b <_A a\}$ for a not $<_A$-minimal.

A structure $\mathsf{A} = \langle A, <_A\rangle$ is *extensional* iff it satisfies the axiom of extensionality:

$$\{z \in A : z\ R\ x\} = \{z \in A : z\ R\ y\} \quad \text{implies} \quad x = y.$$

For A an extensional, well-founded structure, we can use the rank function to define the following.

- the *levels* of A are the sets $\mathrm{lvl}_\alpha(\mathsf{A}) = \{a \in A : \mathrm{rank}(a) = \alpha\}$ for all $\alpha \in \mathrm{Ord}$.
- the *height* or *length* of A is $\mathrm{ht}(\mathsf{A}) = \sup\{\mathrm{rank}(a)+1 : a \in A\} = \mathrm{im}\,\mathrm{rank}$.

We include the "+1" in the definition of height (and rank) to ensure that every element has a smaller rank than the height (or rank of the element we're considering). So the empty relation has height 0, and the ewfset with one element has height 1 while the single element has rank 0. Note that for A a set, Lemma 4 • 2 implies that the height of A is an ordinal, and not just Ord itself. Note some other immediate facts.

4 • 5. Result

Let $\mathsf{A} = \langle A, <_A\rangle$ be well-founded with rank function, rank. Therefore, the following hold.

1. If $a <_A b$, then $\mathrm{rank}(a) < \mathrm{rank}(b)$.
2. If $a, b \in A$ are *comparable*—i.e. $a <_A b$ or $b <_A a$—then $\mathrm{rank}(a) < \mathrm{rank}(b)$ iff $a <_A b$.

Proof .:.

1. Clearly $a <_A b$ implies $\mathrm{rank}(b) > \sup\{\mathrm{rank}(x) : x <_A b\} \geq \mathrm{rank}(a)$ by definition of rank.
2. If a and b are comparable, then either $a <_A b$ (in which case $\mathrm{rank}(a) < \mathrm{rank}(b)$ implies $a <_A b$ by (1)), or $b <_A a$ (in which case $\mathrm{rank}(a) < \mathrm{rank}(b)$ implies $b <_A a$ vacuously by (1)).

Note that we cannot ensure in general that $\mathrm{rank}(a) < \mathrm{rank}(b)$ implies $a <_A b$, since, for example, taking $<_A = \{\langle 0,1\rangle, \langle 2,3\rangle\}$ yields a well-founded relation with $\mathrm{rank}(2) = 0$, $\mathrm{rank}(1) = 1$, but $2 \not<_A 1$. But this concept of rank is what allows us to *collapse* a well-founded, extensional set to a transitive set. We cannot do with with the above example, because it does not satisfy extensionality. It is extensionality that ensures we can *uniquely* describe elements by talking about their predecessors.

Proof of The Mostowski Collapse (4 • 1) .:.

As A satisfies extensionality, there is only one $<_A$-minimal element, $a_\emptyset$. This is because any other $a \neq a_\emptyset$ must then have $\mathrm{pred}_{<_A}(a) \neq \mathrm{pred}_{<_A}(a_\emptyset) = \emptyset$. Hence there is some element of $\mathrm{pred}_{<_A}(a)$, which means a isn't minimal.

Proceed by recursion on the levels of A to define an isomorphism. Since there is only one $<_A$-minimal element $a_\emptyset$, define $f_0(a_\emptyset) = \emptyset$. At limit stage γ define $f_\gamma = \bigcup_{\alpha<\gamma} f_\alpha$. At successor stage $\alpha+1$, consider $\mathrm{lvl}_{\alpha+1}(\mathsf{A})$.

Define $f_{\alpha+1}$ by
$$f_{\alpha+1}(x) = \begin{cases} f_\alpha(x) & \text{if } x \in \mathrm{dom}(f_\alpha) \\ \{f_\alpha(y) : y <_A x\} & \text{if } x \in \mathrm{lvl}_{\alpha+1}(\mathbf{A}). \end{cases}$$
This process stops at $\mathrm{ht}(\mathbf{A})$. Note that this process is well-defined: inductively, $\mathrm{dom}(f_\alpha) = \bigcup_{\beta \le \alpha} \mathrm{lvl}_\beta(\mathbf{A})$, and if $y <_A x \in \mathrm{lvl}_{\alpha+1}(\mathbf{A})$, then $\mathrm{rank}(y) < \mathrm{rank}(x) = \alpha$ so that y is in the domain of f_α. Taking $f = \bigcup_{\alpha < \mathrm{ht}(\mathbf{A})} f_\alpha$, it follows that $f(x) = \{f(y) : y <_A x\}$ for all $x \in A$.

Note that $T = \mathrm{im}\, f$ is transitive: if $x \in f(a) \in T$, then $x = f(b)$ for some $b <_A a$, and thus $x = f(b) \in T$. So it suffices to show that f is an isomorphism between $\mathbf{A}$ and $\langle T, \in \rangle$.

Surjectivity of $f : A \to T$ is immediate. For injectivity, let $a \in A$ be $<_A$-minimal where $f(a) = f(b)$ for some b. Let $f(x) \in f(b)$ for some $x <_A b$ so that $f(x) \in f(a)$ and thus $f(x) = f(y)$ for some $y <_A a$. By minimality of a, $y = x$ and therefore $x <_A a$. The same idea shows that if $x <_A a$ then $x <_A b$, and thus $a = b$ by extensionality.

Now if $a <_A b$ then $f(a) \in \{f(x) : x <_A b\} = f(b)$. Similarly, suppose $f(a) \in f(b)$. Thus $f(a) = f(x)$ for some $x <_A b$. By injectivity, $a = x$ and thus $a <_A b$.

To see that T is unique, suppose $g : A \to D$ is an isomorphism with D transitive. Let $a \in A$ be of least rank such that $f(a) \neq g(a)$. Note that by extensionality and the inductive hypothesis, $f(a) = \{f(x) : x <_A a\} = \{g(x) : x <_A a\} = g(a)$, a contradiction. $\dashv$

So again, The Mostowski Collapse (4 • 1) should highlight the importance of transitive sets, as they allow us to consider any sort of well-founded, extensional[xi] ewfset. This also motivates the axiom of foundation, which says that membership is well-founded. We will not accept foundation as an axiom just yet, though.

4 • 6. Definition

Let $\mathbf{A} = \langle A, <_A \rangle$ be well-founded and extensional. The *mostowski collapsing map* of $\mathbf{A}$ is an isomorphism $\pi : A \to T \subseteq \mathrm{V}$ defined by recursion on rank: for every $a \in A$, $\pi(a) = \{\pi(b) : b <_A a\}$. The *transitive collapse* of $\mathbf{A}$ is then $\langle \mathrm{im}\, \pi, \in \rangle$.

The proof of The Mostowski Collapse (4 • 1) shows that π is well-defined, unique, and is in fact an isomorphism.

Note that there is a slightly more general version of The Mostowski Collapse (4 • 1): we don't require that $\mathbf{A} \in \mathrm{V}$, but instead that at least $\mathrm{pred}_{<_A}(a) \in \mathrm{V}$ for each $a \in A$. For example, $\mathbf{V}$ satisfies this, as $\mathrm{pred}_\in(x) = x \in \mathrm{V}$ for each $x \in \mathrm{V}$. The proof remains the same, as we never needed $\mathrm{ht}(\mathbf{A})$ to be an ordinal: it could be Ord itself, as with V. The point of this generalization is just in case we have a well-founded, partially ordered structure that is not a set. Then we can collapse it down to a transitive class (not necessarily a set) under membership. For now, we will have no use of this generality, but it will be incredibly important later, as we will collapse down various collections into "inner models".

To be slightly more precise than the previous paragraph, for A and R classes, if $\mathrm{pred}_\mathrm{R}(x)$ is a set for each $x \in \mathrm{A}$, then we can define the mostowski collapse as in Definition 4 • 6 as a class, and so yield the image T as a transitive class, which is still isomorphic under membership to A under R.

§ 4 A. Powerset and the cumulative hierarchy

As a consequence of the axiom of foundation, we have the following *iterative* characterization of V in the sense that all collections are formed from things that already exist. In this sense, starting with $\mathrm{V}_1 = \{\emptyset\}$, we can take the set of collections of elements in V_1, which is $\mathrm{V}_2 = \{\emptyset, \{\emptyset\}\}$. Then we can take the set of all collections of elements in this: $\mathrm{V}_3 = \{\emptyset, \{\emptyset\}, \{\{\emptyset\}\}, \{\emptyset, \{\emptyset\}\}\}$, and so on. More precisely, by Lemma 4 • 3, there is a rank function on $\mathbf{V}$. But what exactly is this rank function? By uniqueness, we just need to give an example of one. A first stab at this would be at stage α to define the $\alpha + 1$st level by $\{y : y \subseteq \mathrm{lvl}_\alpha(\mathbf{V})\}$. This seems finte, but it's not particularly useful, as it's unclear

[xi] in the sense of satisfying the axiom of extensionality

that this results in a set. So in doing defining the rank function, we will introduce another axiom, saying that these levels exist: we can continue to define V_α for all α.

4 A • 1. Definition (Axiom)

(*Powerset*) The poweset $\wp(x) = \{y : y \subseteq x\}$ exists: $\forall x\ \exists P\ \forall y\ (y \in P \leftrightarrow y \subseteq x)$.

Note that although comprehension allows us to say that all sorts of subsets of x exist, without the powerset axiom, we cannot in general form the set of all of these at once. But once we know we can collect these together, we get some immediate properties.

- $x \in \wp(x)$, $\emptyset \in \wp(x)$;
- if x is transitive, $x \subseteq \wp(x)$;
- if $x \subseteq y$, then $\wp(x) \subseteq \wp(y)$;
- $\wp(x) \cap \wp(y) = \wp(x \cap y)$;
- $\wp(x) \cup \wp(y) \subseteq \wp(x \cup y)$.

Now consider the following collection. Regardless of whether foundation holds, we can still define it in V.

4 A • 2. Definition

Define the *cumulative hierarchy* to be the collection $\mathrm{WF} = \bigcup_{\alpha \in \mathrm{Ord}} V_\alpha$ given by transfinite recursion:

$$V_0 = \emptyset, \qquad V_{\alpha+1} = \wp(V_\alpha), \quad \text{and} \quad V_\gamma = \textstyle\bigcup_{\alpha<\gamma} V_\gamma, \text{ for } \gamma \text{ a limit ordinal.}$$

Note that WF is not a set, since we will have a surjection from WF onto the ordinals. To see this, consider the following easy to show facts.

4 A • 3. Result

For every $\alpha \in \mathrm{Ord}$, and x,
1. If x is transitive, $\wp(x)$ is transitive.
2. WF is transitive, in particular $V_\alpha \subseteq V_\beta$ for $\alpha < \beta$.
3. if $x \in V_\alpha$, then $\{x\} \in V_{\alpha+1}$.
4. V_α is closed under (finite) unions, intersections, and complements.
5. $\alpha \in V_{\alpha+1}$ for each $\alpha \in \mathrm{Ord}$, hence $\mathrm{Ord} \subseteq \mathrm{WF}$.
6. For each $x \in \mathrm{WF}$, the least α with $x \in V_\alpha$ is a successor ordinal.
7. $\langle \mathrm{WF}, \in \rangle$ is well-founded with rank function $\mathrm{rank}(x)$ as the least α with $x \in V_{\alpha+1}$.
8. $x \in \mathrm{WF}$ iff $x \subseteq \mathrm{WF}$.

Proof .:.

1. Suppose x is transitive, and let $z \in y \in \wp(x)$, i.e. $z \in y \subseteq x$. Therefore $z \in x$ so by transitivity, $z \subseteq x$, and thus $z \in \wp(x)$.

2. Proceed by induction on α to show that V_α is transitive. For $\alpha = 0$, this is immediate. For $\alpha + 1$, use (1) and the inductive hypothesis. For γ a limit, if $y \in x \in \bigcup_{\alpha<\gamma} V_\alpha$, then $y \in x \in V_\alpha$ for some $\alpha < \gamma$, in which case $y \in V_\alpha$ by the inductive hypothesis, and thus $y \in V_\gamma$. Hence every V_α is transitive, and for the same reason as with the limit ordinal, $\mathrm{WF} = \bigcup_{\alpha \in \mathrm{Ord}} V_\alpha$ is transitive too.

3. This is clear, as $\{x\} \subseteq V_\alpha$, and thus $\{x\} \in \wp(V_\alpha) = V_{\alpha+1}$.

4. For any two subsets $x, y \subseteq V_\beta$ for some $\beta < \alpha$, $x \cap y$, $x \cup y$, and $x \setminus y$ are all still subsets of V_β, and hence are elements of $V_{\beta+1} \subseteq V_\alpha$.

5. For $\alpha = 0$, clearly $V_1 = \wp(\emptyset) = \{\emptyset\}$ has $0 \in V_1$. For the successor $\alpha + 1$, $\alpha \in V_{\alpha+1}$ so that $\{\alpha\} \in V_{\alpha+2}$ and as transitive sets, using (4), $\alpha + 1 = \alpha \cup \{\alpha\} \in V_{\alpha+2}$. For limit γ, $\beta \in V_{\beta+1}$ for all $\beta < \gamma$. As $\beta < \gamma$ implies $\beta + 1 < \gamma$, this means $\beta \in V_\gamma$ for every $\beta < \gamma$. Therefore $\gamma \subseteq V_\gamma$ and so $\gamma \in V_{\gamma+1}$.

6. Let $x \in \mathrm{WF}$ be in V_α for α least. If α is a limit ordinal, then clearly $x \in \bigcup_{\beta<\alpha} V_\beta$ implies $x \in V_\beta$ for some

$\beta < \alpha$, a contradiction. Also, $x \notin V_0 = \emptyset$, hence α must be a successor.

7. Write f for this function. Note that if $x \in y$ then clearly $f(x) \leq f(y)$, as $x \in y \in V_{f(y)+1}$ implies $x \in V_{f(y)+1}$. To see that $f(x) \neq f(y)$, $y \in V_{f(y)+1}$ implies $y \subseteq V_{f(y)}$ and hence $x \in V_{f(y)}$, implying that $f(y) \geq f(x) + 1 > f(x)$. So $x \in y$ implies $f(x) < f(y)$.

   Now suppose $X \subseteq$ WF. If there is no $\in$-minimal element, then $f"X$ has no $\in$-minimal element, contradicting the well-foundedness of the ordinals.

   To see that this function f is really a rank function, we need to show that $f(x) = \beta = \sup\{f(y)+1 : y \in x\}$. So clearly, the above argument gives that $f(x) \geq \beta$. And clearly, $y \in V_{f(y)+1}$ for each $y \in x$ implies $x \subseteq \bigcup_{y \in x} V_{f(y)+1} = V_\beta$, and hence $x \in V_{\beta+1}$ shows that $f(x) \leq \beta$. Therefore $f(x) = \beta$, and f is a rank function.

8. If $x \in$ WF, then $x \subseteq$ WF by transitivity. For the other direction, if $x \subseteq$ WF, then $x \subseteq V_\alpha$ for $\alpha = \sup\{\text{rank}(y) : y \in x\}$. Hence $x \in V_{\alpha+1} \subseteq$ WF. ⊣

We can prove more about the class WF, in particular, that it consists of well-founded transitive sets. To do this, with the added technology of the natural numbers, we have the following definition. Note that we can still make this definition in the absence of foundation.

4 A • 4. Definition

Let x be a set. Define $\text{trcl}(x)$, the *transitive closure* of x to be $\bigcup_{n\in\omega} \bigcup^n x$, where $\bigcup^n$ is defined by recursion on ω: $\bigcup^0 x = x$, $\bigcup^{n+1} x = \bigcup(\bigcup^n x)$.

Hence every set is contained in its transitive closure. Of course, the transitive closure $\text{trcl}(x)$ is indeed transitive, since $y \in \text{trcl}(x)$ implies $y \in \bigcup^n x$ and hence $y \subseteq \bigcup^{n+1} x \subseteq \text{trcl}(x)$. The key reason that this should be a motivation for the axiom of foundation, is that we only every need to "go down" ω many times. Foundation will tell us that we only need to go down $< \omega$-many times, although the number of times may be arbitrarily high. Let's first prove some quick results about the transitive closure.

4 A • 5. Result

For every x,

1. $\text{trcl}(x)$ is transitive, and is the $\subseteq$-minimal transitive set containing x: if $x \subseteq T$ where T is transitive, then $\text{trcl}(x) \subseteq T$.
2. If x is transitive, then $\text{trcl}(x) = x$.
3. If $x \in y$, then $\text{trcl}(x) \subseteq \text{trcl}(y)$ (assuming foundation, $\text{trcl}(x) \subsetneq \text{trcl}(y)$).
4. $\text{trcl}(x) = x \cup \bigcup_{y\in x} \text{trcl}(y)$.

Proof .:.

1. If T is transitive and $x \subseteq T$, then clearly $\bigcup^0 x \subseteq T$. And inductively, $\bigcup^n x \subseteq T$ implies $\bigcup^{n+1} x \subseteq T$ by transitivity. Hence $\bigcup_{n\in\omega} \bigcup^n x \subseteq T$, meaning $\text{trcl}(x) \subseteq T$.
2. If x is transitive, then $\text{trcl}(x) \subseteq x$ by (1), and since clearly $x \subseteq \text{trcl}(x)$, we have equality.
3. If $x \in y \subseteq \text{trcl}(y)$ so that $x \subseteq \text{trcl}(y)$. Using (1), we again have that $\text{trcl}(x) \subseteq \text{trcl}(y)$. Assuming foundation, $x \notin \text{trcl}(x)$ (otherwise we would have a finite loop), but $x \in \text{trcl}(y)$.
4. Note that $T = x \cup \bigcup_{y\in x} \text{trcl}(y)$ is transitive, since $y \in x$ implies $y \subseteq \text{trcl}(y) = T$, and if $y \in T \setminus x$, then clearly y is in a transitive set, and hence is a subset of T. Therefore, by (1), $\text{trcl}(x) \subseteq T$. But $T \subseteq \text{trcl}(x)$, since $x \subseteq \text{trcl}(x)$, and (3) implies $\text{trcl}(y) \subseteq \text{trcl}(x)$ for each $y \in x$. ⊣

With this, we have the following, demonstrating why the notation "WF" is used.

4 A • 6. Theorem

Let x be a transitive set. Therefore $x \in$ WF iff $\langle x, \in\rangle$ is well-founded.

Proof .:.

Suppose $x \in \mathrm{WF}$. Therefore $\langle x, \in \rangle$ is well-founded, just by the fact that WF is well-founded and transitive: $x \subseteq \mathrm{WF}$. So suppose $\langle x, \in \rangle$ is well-founded.

Note that $x \subseteq \mathrm{WF}$. To see this, otherwise $A = x \setminus \mathrm{WF}$ has a $\in$-minimal element $a \in A$. Thus $a \subseteq \mathrm{WF}$ so that $a \in \mathrm{WF}$ by Result 4 A • 3. But then $x \subseteq \mathrm{WF}$ yields $x \in \mathrm{WF}$ by the same reasoning. ⊣

So if all of this was motivation, let us give the actual result.

4 A • 7. Theorem

The axiom of foundation implies $\mathrm{V} = \mathrm{WF}$.

Proof .:.

Assuming the axiom of foundation, for each $x \in \mathrm{V}$, $\mathrm{trcl}(x)$ is a transitive set where $\langle \mathrm{trcl}(x), \in \rangle$ is well-founded (just by virtue of V being well-founded). Hence $\mathrm{trcl}(x) \in \mathrm{WF}$. But then $x \subseteq \mathrm{WF}$, and so $x \in \mathrm{WF}$. Hence every element of the universe is an element of WF, and so the two are equal. ⊣

With all of that out of the way, we will now finally accept the axiom of foundation as a part of the axioms of set theory. The rank function on $\mathrm{WF} = \mathrm{V}$ is incredibly useful, as it allows us to proof properties of V through induction on rank. The cumulative hierarchy also gives a nice, stratified picture of the un0iverse, as seen below.

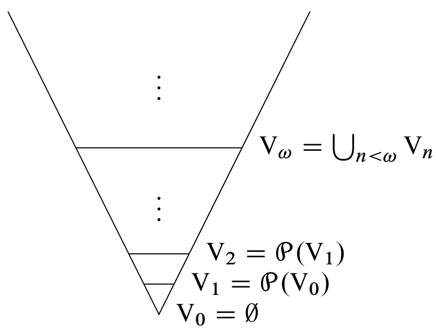

4 A • 8. Figure: The set theoretic universe

The well-foundedness of the universe also gives that any model embeded in V is then well-founded as well. This is just because any infinite decreasing sequence $\mathsf{A} \vDash$ "$a_{n+1} \in a_n$" for $\{a_n : n \in \omega\} \subseteq A$ implies $\mathsf{V} \vDash$ "$f(a_{n+1}) \in f(a_n)$" for each $n \in \omega$, where $f : A \to \mathrm{V}$ is an embedding. Now this relies on a separate, stronger characterization of well-foundedness than first-order logic alone is able to give. So we present the following meta-theoretic result.

4 A • 9. Result

Let $\mathsf{A} = \langle A, \in^{\mathsf{A}} \rangle$ be a structure. Consider the following propositions:

1. A is well-founded.
2. There are no infinite $\in^{\mathsf{A}}$-decreasing sequences of elements of A.
3. A satisfies the axiom of foundation.

Therefore (1) implies (2) and (3), but (3) doesn't imply (2) and thus doesn't imply (1) either.

Proof .:.

To see that (1) implies (2), not that any infinite $\in^{\mathsf{A}}$-decreasing sequence of elements of A is a function from some ordinal α to A. Restricting to ω yields the sequence $\langle a_n : n \in \omega \rangle$ still $\in^{\mathsf{A}}$-decreasing, which gives the set $\{a_n : n \in \omega\} \in \mathrm{V}$ with no $\in^{\mathsf{A}}$-minimal element. Hence A isn't well-founded.

To see that (1) implies (3), if A doesn't satisfy the axiom of foundation, then for some $a \in A$, $\mathsf{A} \vDash$ "$\forall x \in a\ \exists y \in a\ (y \in x)$". Hence the set $\{x \in A : \mathsf{A} \vDash \text{“}x \in a\text{”}\} \in \mathrm{V}$ has no $\in^{\mathsf{A}}$-minimal predecessor. Therefore A isn't well-founded.

To see that (3) doesn't imply (1) nor (2), we use compactness to give a model where (3) holds, but (2)—and thus (1)—fails. In particular, consider ordinal ω in **V**: $\mathbf{N} = \langle \omega, \in \rangle$. As we saw before, the theory of this model can be "misinterpreted" to give an ill-founded model. Clearly **N** satisfies the axiom of foundation, because **N** is well-founded. Therefore, we can consider the theory of **N**:

$$\mathrm{Th}(\mathbf{N}) = \{\varphi \text{ a } \mathtt{FOL}(\in)\text{-sentence} : \mathbf{N} \vDash \varphi\}$$

Now consider the additional constant symbols $\{c_n : n \in \omega\}$. Intuitively, each c_n should count "backwards". Formalizing this, let T be the theory $\mathrm{Th}(\mathbf{N}) \cup \{\text{"}c_{n+1} \in c_n\text{"} : n \in \omega\}$. Note that T has a model by compactness: for each finite subset $\Delta \subseteq T$, there is some largest $N \in \omega$ where c_N occurs in Δ (because Δ is finite). Therefore the model $\mathbf{N}'$ interpretting c_0 as $N \in \omega$, and c_1 as $N-1$ and so on—meaning $c_n^{\mathbf{N}} = N - n$ for all $n \leq N$—has $\mathbf{N}' \vDash \text{"}c_{n+1} \in c_n\text{"}$, and $\mathbf{N}' \vDash \mathrm{Th}(\mathbf{N})$, because we haven't changed any of the structure, just given names to some elements. So $\mathbf{N}'$ is a model of Δ. Hence every finite subset of T has a model, and so T has a model **A**. Thus **A** satisfies the axiom of foundation in $\mathrm{Th}(\mathbf{N}) \subseteq T$, but **A** also has the infinite $\in^{\mathbf{A}}$-decreasing sequence $\langle c_n^{\mathbf{A}} : n \in \omega \rangle$ in V. Therefore (3) holds, but neither (1) nor (2) holds for **A**. ⊣

We will later see that (2) is actually equivalent to (1), but this requires the axiom of choice in the form of König's theorem on trees.

Let us now think about the ranks of sets, and how they can be computed. Recall that the rank of a set x is the least $\alpha \in \mathrm{Ord}$ such that $x \in \mathrm{V}_{\alpha+1}$. The reason for the "+1" is that V_α for α a limit is never the least such that a set appears in it: $\mathrm{V}_\alpha = \bigcup_{\beta<\alpha} \mathrm{V}_\beta$. So defining it in this way allows us to say that there is always a set of rank α for $\alpha \in \mathrm{Ord}$. Another, easier to remember definition is that the rank of x is the least α with $x \subseteq \mathrm{V}_\alpha$.

4 A • 10. Result

For every set x and y,
- if $y \subseteq x$ then $\mathrm{rank}(y) \leq \mathrm{rank}(x)$.
- $\mathrm{rank}(\mathrm{trcl}(x)) = \mathrm{rank}(x)$;
- $\mathrm{rank}(\{x\}) = \mathrm{rank}(x) + 1$;
- $\mathrm{rank}(\mathcal{P}(x)) = \mathrm{rank}(x) + 1$; and
- $\mathrm{rank}(x \cup y) = \max(\mathrm{rank}(x), \mathrm{rank}(y))$.
- $\mathrm{rank}(x) = x$ for $x \in \mathrm{Ord}$.

Proof .:.
- The least α such that $x \subseteq \mathrm{V}_\alpha$ thus also has $y \subseteq \mathrm{V}_\alpha$.
- For $x \subseteq \mathrm{V}_\alpha$ with α least, $\bigcup x \subseteq \mathrm{V}_\alpha$ by transitivity. Hence inductively, $\mathrm{trcl}(x) = \bigcup_{n\in\omega} \bigcup^n x \subseteq \mathrm{V}_\alpha$ and thus $\mathrm{trcl}(x) \in \mathcal{P}(\mathrm{V}_\alpha) = \mathrm{V}_{\alpha+1}$. This establishes that $\mathrm{rank}(\mathrm{trcl}(x)) \leq \mathrm{rank}(x)$. Since $x \subseteq \mathrm{trcl}(x)$ (1) implies the other inequality and thus the two are equal.
- By Result 4 A • 3, rank is a rank function, and thus $\mathrm{rank}(\{x\}) = \mathrm{rank}(x) + 1$.
- Since $\{x\} \subseteq \mathcal{P}(x)$, it follows by (1) and (3) that $\mathrm{rank}(\mathcal{P}(x)) \geq \mathrm{rank}(x) + 1$. For the other direction, note that $y \subseteq x \subseteq \mathrm{V}_\alpha$ implies $y \in \mathrm{V}_{\alpha+1}$ and thus $\mathcal{P}(x) \subseteq \mathrm{V}_{\alpha+1}$ so that $\mathrm{rank}(\mathcal{P}(x)) \leq \mathrm{rank}(x) + 1$. Hence the two are equal.
- Let $\mathrm{rank}(x) < \mathrm{rank}(y) = \alpha$. Therefore $x, y \subseteq \mathrm{V}_\alpha$ and thus $x \cup y \subseteq \mathrm{V}_\alpha$, implying $\mathrm{rank}(x \cup y) \leq \alpha$. Since $y \subseteq x \cup y$, (1) implies the other inequality.
- Proceed by induction on α. For $\alpha = 0$, this is clear. For $\alpha + 1$,

$$\mathrm{rank}(\alpha \cup \{\alpha\}) = \max(\mathrm{rank}(\alpha), \mathrm{rank}(\{\alpha\})) = \mathrm{rank}(\alpha) + 1 = \alpha + 1$$

by (3) and the inductive hypothesis. For limit α, as a rank function, $\mathrm{rank}(\alpha) = \sup_{\beta<\alpha}(\mathrm{rank}(\beta) + 1) = \sup_{\beta<\alpha}(\beta + 1) = \alpha$. ⊣

At this point, calculating ranks might seem completely worthless, but they help to understand just how the universe is built up, and at what stages certain sets come into play. For now, we don't have much use for it, but later on, the levels of the cumulative hierarchy (and other hierarchies) will play a big role in understanding their larger structure through the use of reflection properties—properties of the larger structure holding in smaller parts. For example, just by calculating ranks, one can see that for limit α, V_α is closed under the powerset operation as well as taking unions, pairs, cartesian products, and so on. In this sense, which we will make precise later, the levels of the cumulative hierarchy model a great portion of set theory.

Section 5. Ordinals and Cardinality

It is nearly impossible to have a discussion about set theory that doesn't eventually devolve into a discussion about cardinals. What are cardinals? What is cardinality? These are things that need to be addressed, but to address them, we need a better understanding of ordinals.

§ 5 A. Ordinal Arithmetic

Recall that we can add 1 to ordinals: $\alpha + 1 = \alpha \cup \{\alpha\}$. Using Transfinite Recursion (3 C • 2), we can also define addition between ordinals in general. The motivating picture is that $\alpha + \beta$ is just the order of α placed before the order of β. In particular, we could define $\alpha + \beta$ to be the unique ordinal corresponding to this well-order via Corollary 3 E • 7. But instead, we have the following definition.

5 A • 1. Definition

Define $+ : \mathrm{Ord} \times \mathrm{Ord} \to \mathrm{Ord}$ as follows: for each $\alpha \in \mathrm{Ord}$,

- $\alpha + 0 = \alpha$;
- $\alpha + (\beta + 1) = (\alpha + \beta) + 1$;
- $\alpha + \gamma = \sup_{\beta < \gamma} \alpha + \beta$ for γ a limit.

Define $\cdot : \mathrm{Ord} \times \mathrm{Ord} \to \mathrm{Ord}$ as follows: for each $\alpha \in \mathrm{Ord}$,

- $\alpha \cdot 0 = \alpha$;
- $\alpha \cdot (\beta + 1) = (\alpha \cdot \beta) + \alpha$;
- $\alpha \cdot \gamma = \sup_{\beta < \gamma} \alpha \cdot \beta$ for γ a limit.

Really, we've defined the class $+_\alpha : \mathrm{Ord} \to \mathrm{Ord}$ for each $\alpha \in \mathrm{Ord}$ by Transfinite Recursion (3 C • 2), and so have $+$ defined as the class $\{\langle\langle\alpha, \beta\rangle, \gamma\rangle : +_\alpha(\beta) = \gamma\}$. So the more formally-minded can be put at ease by knowing that these classes are well-defined[xii]

Note that these definitions (restricted to ω) then coincide with the definitions of addition and multiplication on the natural numbers. In particular, given these definitions, that 0 isn't a successor, and Induction on ω (3 B • 1), all of the axioms of peano arithmetic are satisfied by ω under these interpretations. Formally, this means the following.

5 A • 2. Theorem

ZFC $\vdash$ PA$^\omega$, where PA$^\omega$ is the set of axioms of peano arithmetic with all quantifiers restricted to ω, and $+, \cdot, 0, 1$ replaced by the defining `FOL`($\in$)-formulas. In particular, ZFC $\vdash$ Con(PA) by Completeness (1 D • 1).

We haven't quite made precise what all of this means (which we will get to in the next couple sections), but the idea is just that ZFC will show things like the commutativity of $+$ and $\cdot$ on ω. But unlike normal addition, we don't have the same sort of cancellation laws *for general ordinals*, and in fact, commutativity does not hold in general.

5 A • 3. Lemma

For each $\alpha, \beta, \gamma \in \mathrm{Ord}$, if $\beta < \gamma$, then $\alpha + \beta < \alpha + \gamma$. However, it's possible that $\beta + \alpha = \gamma + \alpha$.

Proof .:.

The example to the second sentence can be given easily: take $\alpha = \omega$ with $\beta = 1 < 2 = \gamma$. As a limit ordinal, $\beta + \alpha = \sup_{n \in \omega} 1 + n = \omega = \sup_{n \in \omega} 2 + n = \omega$.

To show the first, proceed by induction on γ. For $\gamma = 0$, this is immediate. For $\gamma + 1$, $\beta < \gamma + 1$ implies $\alpha + \beta \leq \alpha + \gamma$ by the inductive hypothesis. By definition, this is strictly less than $(\alpha + \gamma) + 1 = \alpha + (\gamma + 1)$. For γ a limit, $\beta < \gamma$ implies $\alpha + \beta < (\alpha + \beta) + 1 = \alpha + (\beta + 1) \leq \alpha + \gamma$. $\dashv$

[xii] And of course, the same idea applies to multiplication.

5 A • 4. Corollary

Let $\alpha, \beta, \gamma \in \mathrm{Ord}$. Therefore $\alpha + \beta = \alpha + \gamma$ iff $\beta = \gamma$.

Proof .:.

One direction is immediate. If $\beta \neq \gamma$, say $\beta < \gamma$, then $\alpha + \beta < \alpha + \gamma$, and so the two are unequal. ⊣

We can characterize $\alpha + \beta$ as a copy of α followed by a copy of β. This is formalized by a long definition, but the idea is to produce a copy of α disjoint from a copy of β by considering $\alpha \times \{0\}$ and $\beta \times \{1\}$ instead. We use the map $x \mapsto \langle x, 0\rangle$ to define the isomorphic order on $\alpha \times \{0\}$, and similarly for $\beta \times \{1\}$. This means $\langle x, 0\rangle < \langle y, 0\rangle$ iff $x < y$, and similarly with a 1 in place of a 0, so we are justified in calling them "copies". We put these two orders together just by saying every element of the copy of β is above all elements of the copy of α. So this is how we formalize this "copy of α followed by a copy of β". The characterization is then easy, although very formal.

5 A • 5. Theorem

For each $\alpha, \beta \in \mathrm{Ord}$, $\alpha + \beta$ is the order type of $<_{A_{\alpha,\beta}}$ on $A_{\alpha,\beta} = (\alpha \times \{0\}) \cup (\beta \times \{1\})$ given by

$$\langle \gamma, n\rangle <_A \langle \delta, m\rangle \quad \text{iff} \quad (n = 0 \wedge m = 1) \vee (n = m \wedge \gamma < \delta).$$

Proof .:.

Note that for each $\gamma < \beta$; $A_{\alpha,\gamma} = A_{\alpha,\beta} \setminus \{\langle \delta, 1\rangle : \gamma \leq \delta < \beta\}$. Similarly, the uniform definition of the ordering yields that $A_{\alpha,\gamma}$ is an $<_{A_{\alpha,\beta}}$-initial segment of $A_{\alpha,\beta}$, and in fact the order $<_{A_{\alpha,\gamma}}$ is equal to $<_{A_{\alpha,\beta+1}}$ $\cap (A_{\alpha,\beta} \times A_{\alpha,\beta})$.

Proceed by induction on β. For $\beta = 0$, this is immediate: $A_{\alpha,\beta} = \alpha \times \{0\}$ and $<_{A_{\alpha,\beta}}$ is the same as the order on $\alpha = \alpha + 0 = \alpha + \beta$.

For $\beta + 1$, Note that $\langle \beta, 1\rangle$ is $<_{A_{\alpha,\beta+1}}$-maximal. If we consider $A_{\alpha,\beta+1} \setminus \{\langle \beta, 1\rangle\}$, we get $A_{\alpha,\beta}$. By the inductive hypothesis, and the idea above, it follows that

$$\left\langle \mathrm{pred}_{<_{A_{\alpha,\beta+1}}}(\langle \beta, 1\rangle), <_{A_{\alpha,\beta+1}} \right\rangle = \langle A_{\alpha,\beta}, <_{A_{\alpha,\beta}}\rangle \cong \alpha + \beta.$$

Hence adding on a single element at the end yields $\alpha + (\beta + 1) = (\alpha + \beta) + 1 \cong \langle A_{\alpha,\beta+1}, <_{A_{\alpha,\beta+1}}\rangle$.

For limit β, it should be obvious that $A_{\alpha,\beta} = \bigcup_{\gamma<\beta} A_{\alpha,\gamma}$ and $<_{A_{\alpha,\beta}} = \bigcup_{\gamma<\beta} <_{A_{\alpha,\gamma}}$. The inductive hypothesis tells us that the order type of $\mathsf{A}_{\alpha,\beta} = \langle A_{\alpha,\beta}, <_{A_{\alpha,\beta}}\rangle$, say τ, is at least $\sup_{\gamma<\beta} \alpha + \gamma = \alpha + \beta$. Moreover, if $\tau > \alpha + \beta$, then there must be some initial segment of τ and thus of $\mathsf{A}_{\alpha,\beta}$ which has order type $\alpha + \beta$, which contradicts that each initial segment has order type $\alpha + \gamma$ for $\gamma < \beta$.

We have similar sorts of properties for multiplication. The characterization for ordinal multiplication in that $\alpha \cdot \beta$ is the order type of β copies of α: $\omega \cdot 4 = \omega + \omega + \omega + \omega$, for example.

5 A • 6. Lemma

For each $\alpha, \beta, \gamma \in \mathrm{Ord}$, if $\alpha > 0$, then $\beta < \gamma$, then $\alpha \cdot \beta < \alpha \cdot \gamma$. However, it's possible that $\beta \cdot \alpha = \gamma \cdot \alpha$.

Proof .:.

Again, the example to the second sentence can be given easily: $2 \cdot \omega = \sup_{n\in\omega} 2 \cdot n = \omega = 1 \cdot \omega$ although $1 < 2$.

Proceed by induction on γ. For $\gamma = 0$, the statement is vacuously true. For $\gamma + 1$, by the inductive hypothesis, $\alpha \cdot \beta \leq \alpha \cdot \gamma = \alpha \cdot \gamma + 0 < \alpha \cdot \gamma + \alpha = \alpha \cdot (\gamma + 1)$. For γ a limit, we easily have $\beta + 1 < \gamma$ and hence $\alpha \cdot \beta < \alpha \cdot \beta + \alpha = \alpha \cdot (\beta + 1) \leq \alpha \cdot \gamma$. ⊣

5 A • 7. Corollary

For each $\alpha, \beta, \gamma \in \mathrm{Ord}$, if $\alpha > 0$, then $\alpha \cdot \beta = \alpha \cdot \gamma$ iff $\beta = \gamma$.

Proof .:.

One direction is immediate. Now if $\beta \neq \gamma$, say $\beta < \gamma$, then $\alpha \cdot \beta < \alpha \cdot \gamma$, and so the two are unequal. $\dashv$

Implicit in the restriction that $\alpha > 0$ in Lemma 5 A • 6 is that this doesn't work for $\alpha = 0$. This is, of course, true, since $\alpha \cdot \beta = 0$ for all β when $\alpha = 0$. Although Definition 5 A • 1 only states $\beta \cdot 0 = 0$ for all β, we can inductively show $0 \cdot \beta = 0$ easily: $\beta = 0$ is immediate, and since $0 \cdot (\beta + 1) = 0 \cdot \beta + 0 = 0 + 0 = 0$, it holds at successors, and so trivially at limits.

The characterization of $\alpha \cdot \beta$ as β copies of α, like addition before, relies on a very formal construction to make these "copies" precise. We do this as before by tagging each copy of α: the γth copy of α is $\alpha \times \{\gamma\}$. Hence we're ordering $\alpha \times \beta$. We ensure the γ_0th copy of α is completely before the γ_1th copy of α whenever $\gamma_0 < \gamma_1 < \beta$ by a complicated definition. But once one understands the construction, the idea is easy.

5 A • 8. Theorem

Let $\alpha, \beta \in \mathrm{Ord}$. Therefore $\alpha \cdot \beta$ is order type of $<_{\alpha\times\beta}$ on $\alpha \times \beta$ defined by

$$\langle \alpha_0, \beta_0 \rangle <_{\alpha\times\beta} \langle \alpha_1, \beta_1 \rangle \quad \text{iff} \quad (\beta_0 < \beta_1) \vee (\beta_0 = \beta_1 \wedge \alpha_0 < \alpha_1).$$

Proof .:.

Note that for each $\gamma < \beta$, $\alpha \times \gamma = (\alpha \times \beta) \setminus (\alpha \times (\beta \setminus \gamma))$. Similarly, the uniform definition of the ordering yields that $\alpha \times \gamma$ and its order form an $<_{\alpha\times\beta}$-initial segment of $\alpha \times \beta$. In fact, $<_{\alpha\times\gamma} = <_{\alpha\times\beta} \cap ((\alpha \times \gamma) \times (\alpha \times \gamma))$.

Proceed by induction on β. For $\beta = 0$, this is immediate, as both $\alpha \cdot \beta$ and $\alpha \times \beta$ are $\emptyset$. For $\beta + 1$, by the inductive hypothesis, $\alpha \times (\beta + 1)$ is just the order on $\alpha \times \beta$ followed by the normal order on $\alpha \times \{\beta\}$, which is isomorphic to $\alpha \cdot \beta$ followed by α (using Theorem 5 A • 5 for the formal details). Hence this is just $\alpha \cdot \beta + \alpha = \alpha \cdot (\beta + 1)$. The limit case follows similarly as before. $\dashv$

We can continue to define further ordinal operations. In particular, ordinal exponentiation. This will be the last one we develop, as it is hardly every used, but it does give a good picture of the ordinals and how we can describe them.

5 A • 9. Definition

Define ordinal exponentiation as follows: for each $\alpha \in \mathrm{Ord}$,

- $\alpha^0 = 1$;
- $\alpha^{\beta+1} = (\alpha^\beta) \cdot \alpha$;
- $\alpha^\gamma = \sup_{\beta<\gamma} a^\beta$ for γ a limit.

There is another characterization of ordinal exponentiation in terms of functions with finite support, but it is almost never used in practice, and is instead left to the exercises. But the point is that ordinal exponentiation allows us to express more and more ordinals. In particular, we have the following picture of ordinals, beginning with the natural numbers, and buiding from there using our operations.

$$\underset{0}{\circ} \quad \underset{1}{\circ} \quad \underset{2}{\circ} \;\cdots\; \underset{\omega}{\circ} \quad \underset{\omega+1}{\circ} \;\cdots\; \underset{\omega+\omega}{\circ} \;\cdots\; \underset{\omega+\omega+\omega}{\circ} \;\cdots\; \underset{\omega\cdot\omega}{\circ} \;\cdots\; \underset{\omega\cdot\omega\cdot\omega}{\circ} \;\cdots\; \underset{\omega^2}{\circ} \;\cdots\; \underset{\omega^\omega}{\circ} \;\cdots$$

And this picture, of course, never ends: we can continue to add and multiply ordinals to get larger and larger ordinals like ω^{ω^ω}, $\omega^{\omega^{\omega^\omega}}$, and so on. Actually, taking the supremum of these exponentials—ω raised to ω n-times for $n < \omega$—yields a truly gargantuan ordinal called ε_0 that satisfies $\omega^{\varepsilon_0} = \varepsilon_0$.

Now I would like to raise the question, which ordinals are important? Obviously, this isn't something inherent to the ordinals themselves but instead how we view them. But the question is still one that warrants an answer, given that the ordinals are the canonical well-orders. Are there any other ordinals that are "canonical" in a sense? The answer turns out to be yes. We will take two approaches to answer this question: one the easier route working in $\mathbf{V}$, and a harder route where we deprive ourselves of an important axiom to show that certain things exist or hold in general.

§ 5 B. Cardinals with choice

This picture of the ordinals is useful as it provides a clear idea of "counting" in set-theoretic terms: we proceed lining up the elements of a given set with ordinals just as a child (or adult) might count something by lining it up with their fingers, associating each finger with a number.

The ordinals play the role of the fingers when counting. The issue is that it doesn't follow from the other axioms that every set can be counted in this way. To motivate the axiom of choice, which we need to demonstrate this, consider the following argument.

For some ordinal $\alpha < \beta$, consider the set $X_\alpha \neq \emptyset$. Since each X_α is non-empty, consider some $x_\alpha \in X_\alpha$. Thus $\{x_\alpha : \alpha < \beta\}$ exists. This is equivalent to the axiom of choice. Although we can ensure each X_α has an element, our finite notion of proof can't ensure give these x_αs all at once if there are infinitely many X_αs.

To further motivate the idea, consider the following definition, extending a previous one.

5 B • 1. Definition

Let A and B be sets. Write $A \leq_{\text{size}} B$ iff there is an injection $f : A \to B$. Write $A \geq_{\text{size}} B$ iff there is a surjection $f : A \to B$. Write $A =_{\text{size}} B$ iff there is a bijection $f : A \to B$.

It should be clear that $A \leq_{\text{size}} B$ reflects the notion that A has fewer (or as many) elements than B, because any such injective $f : A \to B$ is really just a bijection $f : A \to \text{im}(f)$ where $\text{im}(f) \subseteq B$. Given that A and $\text{im}(f)$ have the same size, and $\text{im}(f) \subseteq B$, it makes sense to say that A is no bigger than B.

Similarly, it should be clear that $A \geq_{\text{size}} B$ reflects the notion that A has more (or as many) elements than B, since a surjection *covers* all of B with the transformed elements of A (and many elements of A might be forced to go to the same element of B just to fit inside).

It should also be intuitive that $A \leq_{\text{size}} B$ and $A \geq_{\text{size}} B$ implies $A =_{\text{size}} B$. Proving this with what we know thus far, however, is quite difficult, being impossible. So consider the following axiom that allows us to show that this is true.

5 B • 2. Definition (Axiom)

(*Choice*) for any family of non-empty family of disjoint sets F, there is a set C which has chosen one element from each $z \in F$:

$$\forall F\ (\emptyset \notin F \wedge \forall x, y \in F\ (x \cap y = \emptyset) \to \exists C\ \forall x \in F\ !y\ (y \in x \cap C).$$

We call such a set C a *choice set*[xiii]. Really this axiom is just due to the fact that all of our proofs and formulas are finite. In the real world, each $x \in F$ is non-empty, so there is an element $a_x \in x$. So then we can consider $C = \{a_x : x \in F\}$ as a perfectly good set by replacement. The issue is that the finite nature of proofs and formulas cannot incorporate all of this in a finite number of formulas: it requires a potentially infinite number of existential instantiations. But once we have this axiom, we can show that $A \leq_{\text{size}} B$ and $A \geq_{\text{size}} B$ implies $A =_{\text{size}} B$. With this, we have the following.

5 B • 3. Result

Moreover, for $f : B \to A$ a surjection, there is an injection $f' : A \to B$ such that $f(f'(a)) = a$ for all $a \in A$.

Proof .:.

For each $a \in A$, consider the set $f^{-1}(a) = \{b \in B : f(b) = a\}$. As f is surjective, $f^{-1}(a)$ is non-empty for each $a \in A$. Hence the family, which exists by replacement, $\{f^{-1}(a) : a \in A\} = F$ is a family of non-empty sets. Let C then be as in the axiom of choice: for each $f^{-1}(a)$, there is exactly one $b \in C \cap f^{-1}(a)$. Now consider the function

$$f' = \{\langle a, b\rangle \in A \times B : a \in A \wedge b \in C \cap f^{-1}(a)\}.$$

This is an injection $a \neq a' \in A$ requires $f^{-1}(a) \cap f^{-1}(a') = \emptyset$ (any common element b would need to have $a = f(b) = a'$). Moreover, f' is defined on all of A, since $C \cap f^{-1}(a)$ is has an element for each a; and

[xiii] You will occasionally see "choice function" as well, since this defines the function taking $x \in F \setminus \{\emptyset\}$ to the unique element of $x \cap C$.

$f(f'(a)) = a$ because $f'(a) \in f^{-1}(a)$ so that $f(f'(a)) = a$. ⊣

This implies the otherwise intuitive fact below.

5 B • 4. Corollary

For all sets A and B, $A \leq_{\text{size}} B$ iff $B \geq_{\text{size}} A$.

Proof .:.

If $B \geq_{\text{size}} A$, then Result 5 B • 3 tells us that $A \leq_{\text{size}} B$. Cleraly if $A \leq_{\text{size}} B$, as witnessed by the injection $f : A \to B$, then for any fixed $a_0 \in A$, we get a surjection $g : B \to A$ defined by

$$g(b) = \begin{cases} f^{-1}(b) & \text{if } b \in \text{im } f \\ a_0 & \text{otherwise.} \end{cases}$$

This is a surjection, because $A = \text{dom } f = \text{im } f^{-1} \subseteq \text{im } g \subseteq A$ implies $\text{im } g = A$. ⊣

One of the important consequences of choice is that it allows us to count.

5 B • 5. Theorem

For each set A, there is an ordinal α such that $A =_{\text{size}} \alpha$.

Proof .:.

We construct a bijection by transfinite recursion, using the axiom of choice just once. In particular, we define a sequence of approximations to a bijection $f : \alpha \to A$ where $f \restriction_\beta = f_\beta : \beta \to A$ such that $f_\beta \subseteq f_\gamma$ for $\beta < \gamma$. The way to understand the process is just that we use the axiom of choice to choose the "next" element from A. Starting from $\emptyset$ and taking unions at limit stages, this defines the whole process.

Formally, we consider $\mathcal{P}(A) \setminus \{\emptyset\}$. Because this isn't necessarily a family of disjoint sets, consider $P'(A) = \{x \times \{x\} : x \in \mathcal{P}(A) \setminus \{\emptyset\}\}$, tagging each element with names for each subset it appears in. Thus each subset $X \subseteq A$ can be identified with $X \times \{X\} = \{\langle y, X\rangle : y \in X\}$. This $P'(A)$ is a family of non-empty, disjoint sets, and thus there is a set C as in the axiom of choice. Note that this defines a choice function $\mathcal{C} : \mathcal{P}(A) \setminus \{\emptyset\} \to A$ by taking $\mathcal{C}(X)$ to be the unique y where $\langle y, X\rangle \in (X \times \{X\}) \cap C$. Using this C, we can define our sequence of f_αs. In particular, $f_0 = \emptyset$ is an injection, and for γ a limit, define $f_\gamma = \bigcup_{\beta<\gamma} f_\beta$. For the successor case, suppose $f_\alpha : \alpha \to A$ has been defined. If $A = \text{im } f_\alpha$, we let $f = f_\alpha$ and are done. Otherwise, let $f_{\alpha+1}(\alpha) = \mathcal{C}(A \setminus \text{im } f_\alpha)$.

Note that this process has to stop at some point, because otherwise there is a surjection $g : A \to \text{Ord}$ defined by taking $g(a)$ to be the least ordinal $\alpha \in \text{Ord}$ where $a \in \text{im } f_\alpha$, or else $g(a) = 0$. Replacement implies $\text{im } g = \text{Ord}$ is a set, contradicting Burali–Forti Paradox (3 A • 11).

So once $\text{im } f_\alpha = A$, define $f = f_\alpha$. Consider the following easy to see facts about f:

- when f_α is defined, $\text{dom}(f_\alpha) = \alpha$.
- In particular, $\text{dom}(f)$ is an ordinal, and f_α is defined iff $\alpha \leq \text{dom}(f)$.
- $f_\alpha \subseteq f_\beta$ for all $\alpha < \beta \leq \text{dom}(f)$.
- Hence $f_\alpha = f \restriction \alpha$ and so $\text{im } f_\alpha = f"\alpha$.

By construction $f(\alpha) \in A \setminus f"\alpha$. In particular, f is injective, since for $\alpha < \beta$, $f(\beta) \in A \setminus f"\beta$, yet $f(\alpha) \in f"\beta$. Since f is injective by construction, f is thus a bijection between an ordinal and A. ⊣

Note that this ordinal is not necessarily unique. For example, $A = \omega + 1$ has the same size as ω, because we can send $\omega \mapsto 0$ and for $n \in \omega$, we can send $n \mapsto n + 1$. This is clearly surjective onto ω, and it's injective too. So really, just reordering the elements allows us to see that the two have the same size regardless of order[xiv]. The notion of couting given by the ordinals is incredibly important, and leads to the next idea of size: cardinality, being the smallest ordinal of the same size.

[xiv] Clearly ω and $\omega + 1$ are not *isomorphic* as orders, but disregarding order, they have the same size.

5 B • 6. Definition

Let A be a set. Define the *cardinality* of A, written $|A|$, to be the least ordinal α such that $A =_{\text{size}} \alpha$. An ordinal κ is a *cardinal* iff $\kappa = |\kappa|$.

Hence $A =_{\text{size}} B$ is equivalent to $|A| = |B|$. So in particular, $\omega + 1$ is not a cardinal. We have a number of other examples of cardinals: the finite numbers and ω, for instance. To show this, note the following easy to see facts about cardinality.

5 B • 7. Result

For sets A and B, writing $A <_{\text{size}} B$ for $A \leq_{\text{size}} B$ while $A \neq_{\text{size}} B$,

1. $A \leq_{\text{size}} B$ iff $|A| \leq |B|$.
2. $A \leq_{\text{size}} B$ iff $B \geq_{\text{size}} A$ (from Corollary 5 B • 4).
3. $A \leq_{\text{size}} B$ and $A \geq_{\text{size}} B$ implies $A =_{\text{size}} B$.
4. $A =_{\text{size}} B$, $A <_{\text{size}} B$, or $B <_{\text{size}} A$.
5. For $\alpha \leq \beta \in \text{Ord}$, $|\alpha| \leq |\beta|$.

Proof .:.

1. Let $f : A \to B$ be injective. Let $c_A : A \to |A|$ and $c_B : B \to |B|$ be bijections. Define the function $g : |A| \to |B|$ by taking $g(\alpha)$ to be the least β such that $\beta \notin (c_B \circ f \circ c_A^{-1})$"$\alpha$. Note that g is therefore order preserving and hence is an embedding from $|A|$ to $|B|$. If g is bijective, then it is an isomorphism and hence $|A| =_{\text{size}} |B|$, giving that $|A| = |B|$. Otherwise, by Lemma 3 E • 6, $|A|$ is then isomorphic to an intial segment of $|B|$, and as a cardinal, $|A|$ must be this initial segment, meaning $|A| < |B|$.

   For the other direction, if $|A| \leq |B|$, then bijections $c_A : A \to |A|$ and $c_B : B \to |B|$ yield the injection $c_B^{-1} \circ c_A : A \to B$.

3. This is immediate from (1) and (2): $A \leq_{\text{size}} B$ implies $|A| \leq |B|$. $A \geq_{\text{size}} B$ is equivalent to $A \geq_{\text{size}} B$ which is just saying $|A| \geq |B|$, and therefore $|A| = |B|$. Using bijections $c_A : A \to |A|$ and $c_B : B \to |B| = |A|$ yields the bijection $c_B^{-1} \circ c_A : A \to B$ telling us that $A =_{\text{size}} B$.

4. This follows from the same relation happening for ordinals.

5. Clearly $\alpha \leq_{\text{size}} \beta$ since the identity function $\text{id} \restriction \alpha = \{\langle x, x\rangle \in \alpha \times \alpha : x \in \alpha\}$ is an injection from α to β. So by (1), $|\alpha| \leq |\beta|$. ⊣

So this notion of counting gives some very nice properties regarding size, most of which should be expected, and allows us to write $|A| \geq |B|$ instead of $A \geq_{\text{size}} B$ and so forth. So we will abandon the "size" inequalities until we develop the theory of cardinals without the axiom of choice. Beyond the above results, we also get the following famous principle.

5 B • 8. Corollary (The Pigeonhole Principle)

For all sets A and B, suppose $|A| < |B|$. Therefore, if $f : B \to A$, then f is not injective. Moreover, any $f : A \to B$ is not surjective.

Proof .:.

If $f : B \to A$ is injective, then $B \leq_{\text{size}} A$ and hence $|B| \leq |A|$, contradicting that $|A| < |B|$. Similarly, if $f : A \to B$ is surjective, then $|A| \geq |B|$, again a contradiction. ⊣

Now let's get on to proving what the cardinals are. Examples of non-cardinals are abundant. For example, $\omega + \omega$ can be put in bijection with ω since we can rename the first copy of ω with even numbers, and the second copy of ω with the odd numbers. It will be a goal to show that there exist larger cardinals than ω, since even $\omega \cdot \omega$ can be shown to have cardinality ω. Firstly, we have that every natural number is a cardinal number.

5 B • 9. Result

Let $n \in \omega$. Therefore n is a cardinal.

Proof .:.

Proceed by induction on n. For $n = 0$ this is immediate: a bijection $f : 0 \to m$ will have $f \subseteq 0 \times m = \emptyset$ so that $f = \emptyset$ and thus $0 = \emptyset = \operatorname{im} f = m$.

For $n + 1$, it suffices to show that $|n + 1| > n$ by (5) of Result 5 B • 7. So suppose $f : n + 1 \to n$ is a bijection. Consider $f``n$ which then has size $n$. But $f``n = n \setminus \{f(n)\}$. Now we show that this is impossible. If $n = 0$ or $n = 1$, this is clearly impossible, because $n = 0$ has $f = \emptyset$, and $n = 1$ has $1 \setminus \{f(1)\} = \emptyset$, which requires that $f\restriction 1 : 1 \to \emptyset$ is a bijection.

For $n = n^* + 1$ where $n^* \geq 1$, there is a clear bijection between n^* and $n \setminus \{f(n)\}$, as we will show. Explicitly, take $g : n^* \to n \setminus \{f(n)\}$ where

$$g(k) = \begin{cases} k & \text{if } k < f(n) \\ k+1 & \text{if } k \geq f(n). \end{cases}$$

This is a bijection. Clearly it's injective, so it suffices to show surjectivity. To see this, any $k \in n \setminus \{f(n)\}$ has $k \neq f(n)$ and $k \leq n^*$. If $k < f(n) \leq n^*$ then we obviously have $g(k) = k \in \operatorname{im} g$. If $f(n) < k \leq n^*$, then $k > 0$ and hence there is some $k^* \in \omega$ where $k = k^* + 1$ (this is where we use the fact that ω is the least limit ordinal) and this satisfies $f(n) \leq k^*$. Hence $g(k^*) = k^* + 1 = k$. So g is surjective, meaning g is a bijection between $n \setminus \{f(n)\}$ and n^*. Since $f\restriction n : n \to n \setminus \{f(n)\}$ is a bijection, we have a bijection between n and $n^* < n$, contradicting the inductive hypothesis. Therefore no such f can exist. ⊣

We also have that ω is a cardinal.

5 B • 10. Result

The supremum of cardinals is a cardinal. In particular $\omega = \sup_{n\in\omega} n$ is a cardinal.

Proof .:.

Let X be a set of cardinals with $\chi = \sup X$. Clearly if X has a maximal element, then χ is this, and so $\chi \in X$ is a cardinal. So suppose X has no maximal element. If $|\chi| < \chi$, then there is some cardinal $\kappa \in X$ with $|\chi| \leq \kappa$. But since there is some larger cardinal $\lambda \in X$ with then $\chi \geq \lambda > \kappa$, it follows by (5) from Result 5 B • 7 that $|\chi| \geq \lambda > \kappa \geq |\chi|$, a contradiction. Therefore $|\chi| \geq \chi$. We always have by definition of cardinality that $|\chi| \leq \chi$, and so $|\chi| = \chi$. ⊣

Hence we have a dichotomy between ω and the smaller sets, which has already been talked about. Formally, we have the following.

5 B • 11. Definition

A set A is finite iff $|A| < \omega$. A set is infinite iff $|A| \geq \omega$.

So this gives us limit cardinals like ω. But what comes after ω? Certainly there are no infinite cardinals that are successor ordinals.

5 B • 12. Result

et κ be an infinite cardinal. Therefore κ is not a successor ordinal.

Proof .:.

Let $\alpha + 1$ be a successor ordinal. Write $|\alpha| = \lambda$. Consider a bijection $b : \alpha \to \lambda$. Now consider the bijection defined by

$$f(\xi) = \begin{cases} b(\xi+1) & \text{if } \xi \in \omega \\ b(0) & \text{if } \xi = \alpha \\ b(\xi) & \text{otherwise.} \end{cases}$$

This has $f : \alpha + 1 \to \lambda \leq \alpha$ as a bijection, meaning $|\alpha + 1| \neq \alpha + 1$. ⊣

But are there any cardinals larger than ω? The answer to this question is an emphatic yes. In fact, there are just as many cardinals as there are ordinals. And consistently, there are just as many ordinals as there are sets. To generate these cardinals, consider the following theorem, often considered the result that gave birth to the field of set theory, and inspired Russell's Paradox (2 • 5).

5 B • 13. Theorem (Cantor's Theorem)

Let X be a set. Therefore $|X| < |\mathcal{P}(X)|$.

Proof .:.

Let $f : X \to \mathcal{P}(X)$. Consider the set $A = \{x \in X : x \notin f(x)\}$. This is a definable subset of X, and so clearly $A \in \mathcal{P}(X)$. If f were surjective, then $A = f(a)$ for some $a \in A$. So we can ask whether $a \in A$ or not. If $a \in A$, then it meets the definition: $a \notin f(a) = A$, which is a contradiction. Hence $a \notin A$. But this means a doesn't meet the definition: $a \in f(a) = A$. Again, we have a contradiction, and so $A \in \mathcal{P}(X) \setminus \mathrm{im}(f)$. ⊣

Therefore $|\omega| < |\mathcal{P}(\omega)|$, and thus there are larger cardinals than ω. In fact, this theorem gives that there is *no* largest cardinal, since any cardinal κ has $\mathcal{P}(\kappa) > \kappa$. With this information under our belt, consider the following definition.

5 B • 14. Definition

Define by transfinite recursion the infinite cardinals.

$$\begin{aligned} \aleph_0 &= \omega \\ \aleph_{\alpha+1} &= \text{the least cardinal greater than } \aleph_\alpha \\ \aleph_\gamma &= \sup_{\beta<\gamma} \aleph_\beta, \text{ for } \gamma \text{ a limit.} \end{aligned}$$

Although the two are the same as sets, when referring to $\aleph_\alpha$ as an ordinal rather than a cardinal, write ω_α.

So this allows us to consider truly large sets: $\aleph_2$, $\aleph_{\omega_1}$, $\aleph_{\omega_\omega}$, and so on.

5 B • 15. Corollary

The sequence of $n < \omega$ and $\aleph_\alpha$s exhausts all of the cardinals and cardinalities.

Proof .:.

Proceed by induction on α where $\alpha = |\alpha|$. Clearly if $\alpha < \omega$ then we're done. Otherwise, consider $X = \{\beta < \alpha : |\beta| = \beta\}$. This is a set of cardinals, and its supremum λ is then a cardinal by Result 5 B • 10. Note that inductively each $\beta \in X$ has $\beta = \aleph_\gamma$ for some γ. In particular, for $\delta = \sup\{\gamma + 1 : \aleph_\gamma \in X\}$, we have that $X = \{\aleph_\gamma : \gamma < \delta\}$ and thus $\sup X = \aleph_\delta$. Because $\alpha \geq \sup X$, either $\alpha = \sup X = \aleph_\delta$, or $\alpha > \sup X$, and is thus the least cardinal greater than $\aleph_\delta$, meaning $\alpha = \aleph_{\delta+1}$. ⊣

But the definition of the alephs raises the question that allowed us to even consider larger cardinals: what is $|\mathcal{P}(\omega)|$? Where on the long line of alephs is this? Note that the above tells us that $|\mathcal{P}(\omega)| \geq \aleph_1$, but it's not clear whether this equality holds or not. The statement that $|\mathcal{P}(\omega)| = \aleph_1$ is often referred to as *the continuum hypothesis* or CH. Many set theorists have—often very complicated—reasons for thinking that CH is false and instead that $|\mathcal{P}(\omega)| = \aleph_2$[xv]. We will return to this question after investigating what cardinality looks like in a world without choice.

§ 5 C. Cardinality without choice

In the world of choice, the equivalence relation of $=_{\text{size}}$ has canonical representatives in the form of ordinals called *cardinals*, and so every set can be compared in size. In particular, $\leq_{\text{size}}$, $\geq_{\text{size}}$ and $=_{\text{size}}$ are all just different parts of a single linear order: modulo $=_{\text{size}}$, $\leq_{\text{size}}$ (the existence of an injection) is one direction and $\geq_{\text{size}}$ (the existence of a surjection) is the reverse direction.

[xv]Most of this will not be covered in this book, but for those interested, a search for PFA will lead one in the right direction. Be warned, however, that the proof that ZFC $\vdash$ PFA $\to$ $|\mathcal{P}(\omega)| = \aleph_2$ is incredibly long, dealing with complicated set theoretic postulates independent of the other axioms, and full of a technical method called "forcing".

In general, $=_{\text{size}}$ is still an equivalence relation, and $\leq_{\text{size}}$ is still a partial order modulo $=_{\text{size}}$, but it's not necessarily the case that it's linear, nor that $A \leq_{\text{size}} B$ is equivalent to $B \geq_{\text{size}} A$. How, then, do we define cardinality? How do we choose canonical representatives for the equivalences classes of $=_{\text{size}}$? The issue is that we can't, and so in a choiceless context, we don't even try to define representatives of the $=_{\text{size}}$ equivalence classes in general. We can still do this for ordinals, yielding the same notion of what a cardinal is, but this is only because the ordinals have a canonical order on them. Without choice, there isn't always a well-order on sets.

5 C • 1. Lemma

The axiom of choice (AC) is equivalent to the statement that every set has a well-order.

Proof .:.

Let X be an arbitrary set, and suppose AC holds. By Theorem 5 B • 5, there is a bijection $f : X \to \alpha$ for some $\alpha \in \text{Ord}$. Hence the order $W = \{\langle x, y\rangle \in X \times X : f(x) < f(y)\}$ induced by f makes f an isomorphism between $\langle X, W\rangle$ and $\langle \alpha, \in\rangle$, meaning W is a well-order.

Now suppose every set can be well-ordered. Let F be an arbitrary set of disjoint, non-empty sets. Consider $X = \bigcup F$. This has a well-order W. Hence each $x \in F$ has a W-least element, called $a_x \in x$. Moreover, $a_x \notin y$ for each $y \in F \setminus \{x\}$ since the elements of F are pairwise disjoint. Therefore, the set $C = \{a_x : x \in F\}$ works as a choice set for F. ⊣

5 C • 2. Corollary

AC holds for families of sets of ordinals. Hence all parts of Result 5 B • 7 holds for $A, B \in \text{Ord}$.

Ostensibly, as with choice, for arbitrary X we can take the $\in$-least element of $\{\alpha \in \text{Ord} : \alpha =_{\text{size}} X\}$ and thus arrive at a cardinality for X as before. The issue is that it's not clear this set is non-empty, and in fact, if choice fails then this *will* be empty for some X as Lemma 5 C • 1 shows.

5 C • 3. Definition

Let X be a set. The *choiceless-cardinality* of X is the equivalence class $[X]_{\text{size}} = \{A : A =_{\text{size}} X\}$.
A *cardinal* is still an ordinal κ such that κ is $<$-minimal in $\text{Ord} \cap [\kappa]_{\text{size}}$.

Note that $[X]_{\text{size}}$ will be a class rather than a set. With these concepts, we still have the following results about $\leq_{\text{size}}$. Namely, that $\leq_{\text{size}}$ is antisymmetric modulo $=_{\text{size}}$. A similar result was shown with choice: Result 5 B • 7 (3), where $A \leq_{\text{size}} B$ and $A \geq_{\text{size}} B$ implies $A =_{\text{size}} B$. But this was done by comparing cardinality rather than defining a bijection outright.

5 C • 4. Theorem (Cantor–Bernstein)

Let A and B be sets. Suppose $A \leq_{\text{size}} B$ and $B \leq_{\text{size}} A$. Therefore $A =_{\text{size}} B$.

Proof .:.

Let $\mathbb{A} : A \to B$ and $\mathbb{B} : B \to A$ be injections, witnessing the hypothesis. We will categorize the elements of B in the following way. Call elements $b \in B \setminus \text{im}(\mathbb{A})$ *starting points*.

For each starting point $b_0 \in B$, we can then identify the path it takes by going to A via $\mathbb{B}$, then then back to B via $\mathbb{A}$. Write $(\mathbb{A} \circ \mathbb{B})^n$ for $(\mathbb{A} \circ \mathbb{B}) \circ (\mathbb{A} \circ \mathbb{B}) \circ \cdots \circ (\mathbb{A} \circ \mathbb{B})$, meaning $\mathbb{A} \circ \mathbb{B}$ composed n-times for each $n \in \mathbb{N}$. So a point $b \in B$ is *on the path of* b_0 iff $b = (\mathbb{A} \circ \mathbb{B})^n(b_0)$ for some $n \in \mathbb{N}$, possibly 0. Now we define $f : A \to B$ via replacement by

$$f(a) = \begin{cases} \mathbb{B}^{-1}(a) & \text{if } \mathbb{A}(a) \text{ is on the path of a starting point,} \\ \mathbb{A}(a) & \text{otherwise.} \end{cases}$$

This makes sense as $\mathbb{B}$ is injective: $\mathbb{B}^{-1}$ is a function. To see this, if $\langle a, b\rangle, \langle a, b'\rangle \in \mathbb{B}^{-1}$ for $b \neq b'$, then $\mathbb{B}(b) = \mathbb{B}(b') = a$ contradicts injectivity. Hence f is a function defined on all of A, and clearly $\text{im}\, f \subseteq B$.

So it suffices to show that f is injective, and surjective.

Claim 1

f is injective.

Proof .:.

Suppose $f(a) = f(a')$ for $a \neq a'$. If both $\mathbb{A}(a)$ and $\mathbb{A}(a')$ are on the path of a starting point, then $f(a) = \mathbb{B}^{-1}(a) = f(a') = \mathbb{B}^{-1}(a')$. This contradicts that $\mathbb{B}$ is a function: $\langle a, b\rangle, \langle a', b\rangle \in \mathbb{B}^{-1}$ implies $\mathbb{B}(b)$ is both a and a'. So this case can't happen. Similarly, if neither $\mathbb{A}(a)$ nor $\mathbb{A}(a')$ is on the path of a starting point, then $f(a) = \mathbb{A}(a) = f(a') = \mathbb{A}(a')$ contradicts the injectivity of $\mathbb{A}$.

So suppose for the sake of definiteness that $\mathbb{A}(a)$ is on the path of a starting point, but $\mathbb{A}(a')$ isn't. Note that $f(a)$ is then on the path of a starting point, because $\mathbb{A} \circ \mathbb{B}(f(a)) = \mathbb{A} \circ \mathbb{B}(\mathbb{B}^{-1}(a)) = \mathbb{A}(a)$ on the path of a starting point. $\mathbb{A}(a)$ of course is not itself a starting point, since it's in the image of $\mathbb{A}$, but $f(a)$ might be. Anyway, $f(a)$ being on a path means that $f(a) = f(a') = \mathbb{A}(a')$ is too, a contradiction. ⊣

All that remains to be shown is that f is surjective. To do this, let $b \in B$. If b is on the path of a starting point, $a = \mathbb{B}(b)$ yields $f(a) = b$. If b is not on the path of a starting point, then certainly b itself is not a starting point, meaning $b \in \operatorname{im}(\mathbb{A})$. So taking such an a with $\mathbb{A}(a) = b$ yields that $\mathbb{A}$ isn't on the path of a starting point, and thus $f(a) = \mathbb{A}(a) = b$. Thus f is surjective, and so a bijection. ⊣

As detailed above, it's tempting for each X to define $\{\alpha \in \mathrm{Ord} : \alpha =_{\text{size}} X\}$, and then take $|X|$ to be the $<$-least element of this class. Although we can't do this because the class might be empty, we still at least have the following result, showing that the ordinals can still overtake any set in the $\leq_{\text{size}}$-ordering.

5 C • 5. Theorem (Hartogg's Number)

Let X be a set. Therefore there is a cardinal κ such that $\kappa \not\leq_{\text{size}} X$.

Proof .:.

Consider the approximations to a well-order of X. In particular, consider the set

$$\mathcal{W} = \{W \in \mathcal{P}(X \times X) : W \text{ is a well-order of } \operatorname{dom}(W) \cup \operatorname{ran}(W).\}$$

Now by Corollary 3 E • 7, we can consider the set of the corresponding order types.

$$\mathcal{O} = \{\alpha \in \mathrm{Ord} : \exists W \in \mathcal{W}(\langle \alpha, \in\rangle \cong \langle \operatorname{dom}(W) \cup \operatorname{ran}(W), W\rangle)\}.$$

Note that $\mathcal{O}$ must be an ordinal, since it is transitive: $\beta < \alpha \in \mathcal{O}$ has that the well-order $W \in \mathcal{W}$ with order type α can be restricted to an initial segment with order type β and thus $\beta \in \mathcal{O}$. So it suffices to show that $\mathcal{O} \not\leq_{\text{size}} X$.

Suppose $f : \mathcal{O} \to X$ is an injection. Therefore the order $W = \{\langle f(\alpha), f(\beta)\rangle \in X \times X : \alpha < \beta\}$ yields a well-order of a subset of X that is isomorphic to $\mathcal{O}$. In particular, $W \in \mathcal{W}$ and $\mathcal{O} \in \mathcal{O}$, contradicting that the ordinals are well-founded. ⊣

If choice holds, the cardinal described above is just any cardinal greater than $|X|$. But without choice, it's not clear that $X =_{\text{size}} \alpha$ for any $\alpha < \kappa$, as this would guarantee by Cantor–Bernstein (5 C • 4) that X can be well-ordered, and thus any family $F \subseteq \mathcal{P}(X)$ would have a choice set just by selecting the least-elements in the non-empty sets. So if every X has a cardinality, then we always get choice sets, and thus the axiom of choice holds.

Hartogg's Number (5 C • 5) is especially useful in confirming the other properties of the cardinals that we know from Subsection 5 B. There, ω_1 was shown to exist from a well-order of $\mathcal{P}(\omega)$. But without choice, it's possible for $\mathcal{P}(\omega)$ to have no well-order. How then do we show that there are larger cardinalities? We use Hartogg's Number (5 C • 5). Note that we still have the usual properties of $\leq_{\text{size}}$ due to choice holding on the ordinals by Corollary 5 C • 2.

5 C • 6. Corollary

For each cardinal $\kappa \in \mathrm{Ord}$, there is a cardinal $\lambda > \kappa$.

Hence without choice we can still define $\aleph_1, \aleph_2, \cdots, \aleph_\omega$, and so on. It's just that not every set needs to be in bijection with one of these.

§ 5 D. cofinality and cardinal arithmetic

We now return to the world of choice, although often it is unnecessary for this subsection. There will be times when it is needed, but mostly this is just in requiring that functions from κ to λ can be well-ordered for κ and λ ordinals.

As introduced before, there are operations defined on ordinal numbers: addition, multiplication, and exponentiation, for example. We have similar operations on cardinals, although they do not obey the same rules. It will happen that everything becomes either dramatically simpler, or else impossible to know. We begin with some notable properties of cardinals. We begin with addition.

5 D • 1. Definition

Let X and Y be sets. Write $X \sqcup Y$ for $(X \times \{0\}) \cup (Y \times \{1\})$ the disjoint union.
Let κ and λ be cardinals. Define $\kappa + \lambda$ to be the cardinality of $\kappa \sqcup \lambda$. Define $\kappa \cdot \lambda$ to be the cardinality of $\kappa \times \lambda$.

Note that the cardinality of the ordinal addition $\kappa + \lambda$ is the cardinal addition $\kappa + \lambda$. To make this more apparent what is meant, $|\omega + \omega_1| = \aleph_0 + \aleph_1$ for example. Unlike with ordinal addition, where $\omega + 1 \neq \omega$, both of these cardinal operations simplify to just being the maximum of the two cardinals. First we have some immediate properties about these operations, just following from the existence of easy to find injections or surjections. Below, for κ, λ, and θ cardinals:

- Cardinal addition and multiplication are commutative.
- $\kappa < \lambda$ implies $\theta \cdot \kappa \leq \theta \cdot \lambda$ and $\theta + \kappa \leq \theta + \lambda$ (possibly with equality, as we shall see).
- $\kappa + 0 = \kappa$, and $\kappa \cdot 0 = 0$.
- $\kappa + 1 = \kappa$, and $\kappa \cdot 1 = \kappa$.
- $\kappa + \lambda \leq \kappa \cdot \lambda$ when $\lambda \neq 0$.
- $\alpha \leq \aleph_\alpha$ (possibly with equality, as we shall see).

Trivially, however, these facts won't be important to know, since we will get that $\kappa + \lambda = \kappa \cdot \lambda = \max(\kappa, \lambda)$. To show this, we first consider the case where $\kappa = \lambda$.

5 D • 2. Lemma

Let κ be an infinite cardinal. Therefore $\kappa \cdot \kappa = \kappa + \kappa = \kappa$.

Proof .:.

Clearly $\kappa \leq \kappa + \kappa \leq \kappa \cdot \kappa$ so it suffices to show $\kappa \cdot \kappa \leq \kappa$. We consider the following ordering on Ord $\times$ Ord. Write $\langle \alpha_0, \beta_0 \rangle \prec \langle \alpha_1, \beta_1 \rangle$ iff

1. $\max(\alpha_0, \beta_0) < \max(\alpha_1, \beta_1)$; or else
2. $\alpha_0 < \alpha_1$; or else
3. $\beta_0 < \beta_1$.

In essence, we have a lexicographic order where things in the square $\gamma \times \gamma$ always precede things in the square $\delta \times \delta$ for $\gamma < \delta$. As a result, this means we follow the edges of increasingly bigger squares.

Claim 1

$\prec$ is a (class) well-order of Ord $\times$ Ord.

Proof .:.

That $\prec$ is a linear order should be easy to see from the definition: transitivity follows from from progressing through the cases each time, and the other requirements follow from $<$ being a linear order on Ord. To show that $\prec$ is well-founded, let X be a set of pairs of ordinals. Consider the set $Y = \{\max(\alpha, \beta) : \langle \alpha, \beta \rangle \in X\}$. This has a $\in$-least element $\alpha_0 \in Y$ so consider the class $\{\langle \alpha, \beta \rangle \in X : \max(\alpha, \beta) = \alpha_0\}$. Now we similarly choose the $\in$-least first entry in this set, and of those entries with the same max and same first entry, we conside the $\in$-least second entry. This gives a $\prec$-least element of X just by definition of $\prec$. $\dashv$

Proceed by induction on γ to show $\aleph_\gamma \cdot \aleph_\gamma = \aleph_\gamma$ by showing $\prec_\gamma = \prec \cap ((\aleph_\gamma \times \aleph_\gamma) \times (\aleph_\gamma \times \aleph_\gamma))$ has order type $\aleph_\gamma$. Because we prioritize smaller squares, for each $\langle \alpha, \beta \rangle \in$ Ord $\times$ Ord, $\mathrm{pred}_\prec(\langle \alpha, \beta \rangle)$ is a set, and in particular, it has order type at most (using ordinal multiplication) $\max(\alpha,\beta)\cdot\max(\alpha,\beta)$. So for $\alpha,\beta<\aleph_\gamma$, the inductive hypothesis tells us that this $\mathrm{pred}_{\prec}(\langle\alpha,\beta\rangle)$ has cardinality $|\max(\alpha,\beta)|\cdot|\max(\alpha,\beta)| = |\max(\alpha,\beta)| < \aleph_\gamma$, and thus the order type of this initial segment is an ordinal strictly less than $\aleph_\gamma$. Thus every initial segment of $\prec_\gamma$ has order type strictly less than $\aleph_\gamma$, and therefore the order type of $\prec_\gamma$ is at most $\aleph_\gamma$. Since clearly the order type is at least $\aleph_\gamma$ (consider $\{\langle\alpha,0\rangle : \alpha<\aleph_\gamma\}$, still well-ordered by $\prec$ and isomorphic to $\langle\aleph_\gamma,<\rangle$, and use Lemma 3 E • 6), we have equality and thus $|\aleph_\gamma\times\aleph_\gamma| = \aleph_\gamma$. ⊣

We can then conclude that $\kappa\cdot\lambda = \kappa+\lambda = \max(\kappa,\lambda)$ for infinite cardinals κ and λ.

5 D • 3. Corollary

Let $\kappa<\lambda$ be cardinals with λ infinite. Therefore $\kappa\cdot\lambda=\kappa+\lambda=\lambda$.

Proof .:.

$\lambda = 1\cdot\lambda\le\kappa\cdot\lambda\le\lambda\cdot\lambda=\lambda$ by Lemma 5 D • 2, and similarly for addition. ⊣

Now we will discuss some aspects of cardinal arithmetic that are more complicated in the sense that it's impossible to write down precisely which $\aleph_\alpha$ the answer is. But there are still interesting results we can give.

5 D • 4. Definition

Let A and B be sets. Define ${}^AB = \{f\in\mathcal{P}(A\times B) : f \text{ is a function from } A \text{ to } B\}$.
For κ and λ cardinals, define $\kappa^\lambda = |{}^\lambda\kappa|$.
We often write $\kappa^{<\lambda}$ for $\sup_{\theta<\lambda}\kappa^\theta$.

Now obviously we get the following facts about exponentiation: for all cardinals $\kappa,\lambda,\theta\in\mathrm{Ord}$, and all sets A,

- ${}^{\emptyset}A=\{\emptyset\}$, ${}^1A = 1\times A$.
- $\kappa<\lambda$ implies $\theta^\kappa\le\theta^\lambda$;
- $\kappa<\lambda$ implies $\kappa^\theta\le\lambda^\theta$;
- $\kappa^0=1$, $\kappa^1=\kappa$.
- $\kappa^2=\kappa\cdot\kappa$, and so successively, $\kappa^n=\kappa$ for each $n\in\omega$, if κ is infinite. Therefore $\kappa^{<\omega}=\kappa$.

The notation of exponentiation makes sense for this operation because if $f : A\sqcup B\to C$, meaning $f\in{}^{A\sqcup B}C$, then we can view f according to how it acts on A and how it acts on B. In particular, every function in ${}^{A\sqcup B}C$ can be viewed as a pair of functions in ${}^AC\times{}^BC$, and vice versa (because we're taking the disjoint union). Moreover, the idea of evaluation just gives that ${}^A({}^BC)$ is effectively the same as ${}^{A\times B}C$ in that every function $f : A\to{}^BC$ can be uniquely identified with the map $g : A\times B\to C$ where $g(a,b)=f(a)(b)$. Hence we get the following facts about cardinal exponentiation.

5 D • 5. Result

Let κ, λ, and θ be cardinals. Therefore,
1. $\theta^{\kappa+\lambda}=\theta^\kappa\cdot\theta^\lambda$.
2. $(\theta^\kappa)^\lambda=\theta^{\kappa\cdot\lambda}$.

Moreover, these concepts also allow us to view the powerset as exponentiation.

5 D • 6. Result

Let X be a set. Therefore $|\mathcal{P}(X)| = |{}^X2| = 2^{|X|}$.

Proof .:.

Each subset corresponds to its characteristic function: $A\subseteq X$ yields $\chi_A : X\to 2$ where $\chi_A(x)=1$ iff $x\in A$ and otherwise $\chi_A(x)=0$. Hence $\chi_A^{-1}(1)=A$ for all $A\subseteq X$. In particular, if $\chi_A=\chi_B$ then they both have the same preimage of 1 and so $A=B$. Similarly, every function $f : X\to 2$ yields a unique subset of X just by the preimage of 1: $A_f=\{x\in X : f(x)=1\}=f^{-1}(1)$ and so $\chi_{A_f}=f$. Hence the map $F:\mathcal{P}(X)\to{}^X2$ where $A\mapsto\chi_A$ is a bijection. Therefore $|\mathcal{P}(X)|=|{}^X2|$ which is just $2^{|X|}$ by definition. ⊣

In particular, $|\mathcal{P}(\aleph_0)| = 2^{\aleph_0} > \aleph_0$ by Cantor's Theorem (5 B • 13), which more generally gives the following.

5 D • 7. Corollary

Let κ be a cardinal. Therefore $2^\kappa > \kappa$.

Note that $\kappa^\kappa = 2^\kappa$ for infinite κ, since

$$2^\kappa \leq \kappa^\kappa \leq (2^\kappa)^\kappa = 2^{\kappa\cdot\kappa} = 2^\kappa.$$

Another proof that $2^\kappa > \kappa$ follows from a very useful theorem. First, note that we can generalize exponentiation to other products, and we generalize multiplication to other sums.

5 D • 8. Definition

Let I be a set, and let $\{\kappa_i : i \in I\}$ be a set of ordinals. The cardinal sum $\sum_{i\in I} \kappa_i$ is the cardinality of the union $\bigcup_{i\in I} \kappa_i \times \{i\}$.
The cardinal product $\prod_{i\in I} \kappa_i$ is the cardinality of the cartesian product $\prod_{i\in I} \kappa_i$.

Obviously we have $\sum_{i\in I} \kappa_i \leq \prod_{i\in I} \kappa_i$ just by looking at the map sending $\langle \alpha, i\rangle \in \bigcup_{i\in I} \kappa_i \times \{i\}$ to the function in the cartesian product $\prod_{i\in I} \kappa_i$ where $i \mapsto \alpha$ and $j \mapsto 0$ for every $j \in I$ with $j \neq i$. We also have the following easy to confirm properties.

- $I \subseteq J$ with $\{\kappa_j : j \in J\}$ a set of cardinals implies $\sum_{i\in I} \kappa_i \leq \sum_{j\in J} \kappa_j$; and
- if in addition, $\emptyset \notin J$, $\prod_{i\in I} \kappa_i \leq \prod_{j\in J} \kappa_j$.
- $\kappa_i \leq \theta_i$ implies $\sum_{i\in I} \kappa_i \leq \sum_{i\in I} \theta_i$, and similarly for products.
- $\sum_{i\in I} 1 = |I|$ and $\prod_{i\in I} 2 = 2^{|I|}$; and more generally,
- $\sum_{i\in I} \kappa = |I| \cdot \kappa$ and $\prod_{i\in I} \kappa = \kappa^{|I|}$.

Mostly we will look at sums as given by partitions: if we can cover a set, then the cardinality is given by how many pieces we need, and how big the pieces are.

5 D • 9. Result

Let X be a set, and $P \subseteq \mathcal{P}(X)$ a partition of X such that $|P|$ is infinite. Therefore $|X| = \sum_{Y\in P} |Y| = |P| \cdot \sup_{Y\in P} |Y|$.

Proof .:.

Since X can be written as the disjoint union $X = \bigcup_{Y\in P} Y$, it's clear that $\{|Y| : Y \in P\}$ is a set of cardinals, and $\bigcup_{Y\in P} |Y| \times \{Y\}$ is in bijection with X, just by sending $\langle \alpha, Y\rangle$ to the $f_Y(\alpha)$ where $f_Y : |Y| \to Y$ is a bijection. As a result, $|X| = \sum_{Y\in P} |Y|$.

This is the same as $|P| \cdot \sup_{Y\in P} |Y|$, since $|P| = \sum_{Y\in P} 1 \leq \sum_{Y\in P} |Y| \leq \sum_{Y\in P} |P| = |P| \cdot |P| = |P|$. $\dashv$

Infinite sums in general work like this.

5 D • 10. Corollary

Let I be a set, and $\{\kappa_i : i \in I\}$ a set of cardinals. Therefore $\sum_{i\in I} \kappa_i = |I| \cdot \sup_{i\in I} \kappa_i$.

A less trivial theorem is the following[xvi], giving an alternative proof of Cantor's Theorem (5 B • 13).

5 D • 11. Theorem (König's Theorem)

Let I be a set (used only as an index), and let $\{\kappa_i : i \in I\}$ and $\{\theta_i : i \in I\}$ be two sets of cardinals. Suppose $\kappa_i < \theta_i$ for all $i \in I$. Therefore $\sum_{i\in I} \kappa_i < \prod_{i\in I} \theta_i$.

Proof .:.

Without loss of generality, instead consider the situation where we have pairwise disjoint families

$$\{\mathrm{K}_i : i \in I\}, \qquad |\mathrm{K}_i| = \kappa_i$$

[xvi] named after Kőnig Gyula who often published under the pseudonym "Julius König".

$$\{\Theta_i : i \in I\}, \qquad |\Theta_i| = \theta_i.$$

For example, $\mathrm{K}_i = \kappa_i \times \{i\}$ and $\Theta_i = \theta_i \times \{i\}$ works. Let $\mathrm{K} = \bigcup_{i \in I} \mathrm{K}_i$, and Θ equal the cartesian product $\prod_{i \in I} \Theta_i$. It's clear that there's an injection from K to Θ, just because $\kappa_i < \theta_i$: send $\langle \alpha, i \rangle \in \mathrm{K}$ to the map $f \in \Theta$ where $f(i) = \langle \alpha, i \rangle$ and $f(j) = \langle \alpha + 1, j \rangle$ for $i \neq j \in I$.

Now suppose we had a surjection $F : \mathrm{K} \to \Theta$. We will diagonalize out of this using evaluation maps: for $x \in I$ and $f \in \Theta$, $\iota_x(f) = f(x)$. Let $F_i = (\iota_i \circ F) \upharpoonright \mathrm{K}_i$, a function from K_i to $\bigcup_{i \in I} \Theta_i$.

Since $|\Theta_i| > |\mathrm{K}_i|$, as a function from K_i, F_i can never cover Θ_i. So let $g(i) \in \Theta_i \setminus \mathrm{im}(F_i)$ for each $i \in I$. The resulting function g cannot be in the image of F. To see this, if we let $k \in \mathrm{K}$ be such that $F(k) = g$, then we know $k \in \mathrm{K}_i$ for exactly one $i \in I$. Hence $k \in \mathrm{dom}(F_i)$ and so $g(i) \in \Theta_i \setminus \{F_i(k)\}$ by construction. Yet $F_i(k) = (\iota_i \circ F)(k) = F(k)(i) = g(i)$, a contradiction. ⊣

5 D • 12. Corollary

Let κ be a cardinal. Therefore $2^\kappa > \kappa$.

Proof .:.

Since $2 > 1$, $\kappa = \sum_{i \in \kappa} 1 < \prod_{i \in \kappa} 2 = 2^\kappa$ by König's Theorem (5 D • 11). ⊣

This raises the question, how much more can we know about 2^κ, and λ^κ more generally? Because cardinal exponentiation grows in both arguments, we at least know that $2^\kappa = \kappa^\kappa$ since $2^\kappa \leq \kappa^\kappa \leq (2^\kappa)^\kappa = 2^{\kappa \cdot \kappa} = 2^\kappa$ by Result 5 D • 5 and Lemma 5 D • 2. For now, the question we will address is just when does 2^λ cross from being below κ to being above? To give the best possible answer we can give, we need a new concept.

5 D • 13. Definition

A *poset* is a structure $\mathbf{A} = \langle A, <_A \rangle$ where $<_A$ is a partial order over A.
Let $\mathbf{A} = \langle A, <_A \rangle$ be a poset. A subset $X \subseteq A$ is *cofinal* in $\mathbf{A}$ iff $\forall a \in A\, \exists x \in X\ (a = x \vee a <_A x)$.
For α an ordinal, the *cofinality* of α, written $\mathrm{cof}(\alpha)$, is the least order type of a cofinal subset of α.

If α is an ordinal, we say that $X \subseteq \alpha$ has order type β for the more formal statement that $\langle X, \in \rangle$ has order type $\langle \beta, \in \rangle$. For linear orders, being cofinal is the same as being unbounded. So $\mathrm{cof}(\alpha)$ is also the least order type of an unbounded subset of α. One may expect that if $\alpha > \beta > \mathrm{cof}(\alpha)$, then there is a cofinal subset of α with order type β. But this may not be true, paradoxically. The main reason is that after using $\mathrm{cof}(\alpha)$ many elements of β, we might run out of room to place the other elements of β while preserving the order. First, note the following easy examples: for α an ordinal,

- $\mathrm{cof}(\alpha + 1) = 1$, as witnessed by $\{\alpha\} \subseteq \alpha + 1$.
- $\mathrm{cof}(\aleph_\omega) = \omega$, as witnessed by $\{\aleph_n : n < \omega\}$.
- $\mathrm{cof}(\aleph_\alpha) = \mathrm{cof}(\alpha)$ for α a limit, by the same reason.
- $\mathrm{cof}(\aleph_0) = \omega$.
- $\mathrm{cof}(\alpha) \leq \alpha$ for each ordinal α.

An arguably easier way to characterize cofinality is with functions.

5 D • 14. Definition

An ordinal λ is cofinal in $\alpha \in \mathrm{Ord}$ iff there is an increasing function $f : \lambda \to \alpha$ such that $\mathrm{im}\, f$ is cofinal in α.

This is an alternative way to characterize it in the following sense.

5 D • 15. Result

Let α be an ordinal. There is a subset of α of order-type β iff β is cofinal in α.

Proof .:.

Obviously if β is cofinal in α, then there is a subset of α of order-type β: $\mathrm{im}\, f$ where $f : \beta \to \alpha$ is increasing. So suppose $X \subseteq \alpha$ has order-type β. Thus there is a function $f : \beta \to X$ which is an isomorphism and thus order preserving, and increasing in particular. It follows that f witnesses that β is cofinal in α. ⊣

Let's investigate what kinds of ordinals can be cofinalities. Note that being unbounded in an ordinal isn't unique: for α an ordinal, obviously both $\mathrm{cof}(\alpha)$ and α itself have unbounded sequences in α. For a less trivial example, $\omega+\omega+\omega$ has $\{\omega+\omega+n : n \in \omega\}$ as a subset with order type ω, $\{\omega+\alpha+m : \alpha \leq \omega \wedge m \in \omega\}$ as a subset with order type $\omega+\omega$, and both are unbounded in $\omega+\omega+\omega$.

Nevertheless, we do get a kind of uniqueness in the following sense.

5 D • 16. Lemma

Let β be cofinal in α. Therefore $\mathrm{cof}(\beta) = \mathrm{cof}(\alpha)$

Proof .:.

Enumerate $X = \{x_\xi : \xi < \beta\}$, and let $Y = \{y_\xi : \xi < \mathrm{cof}(\alpha)\}$ be cofinal by definition of $\mathrm{cof}(\alpha)$. For each $y \in Y \subseteq \alpha$, as X is cofinal in α, there is some $x \in X$ with $y < x$. So for $y_\xi \in Y$, let $x'_\xi \in X$ be the least element of X such that $y_\xi < x'_\xi$. Hence $\{x'_\xi : \xi < \mathrm{cof}(\alpha)\}$ is a subset of X that is cofinal with order type $\mathrm{cof}(\alpha)$. Since $\langle X, <\rangle \cong \langle \beta, <\rangle$, taking the relevant transformation of the x'_ξs yields then that $\mathrm{cof}(\beta) \leq \mathrm{cof}(\alpha)$.

But any cofinal subset of β of order type $\mathrm{cof}(\beta)$ yields a cofinal subset of X of order type $\mathrm{cof}(\beta)$, and thus a cofinal subset of α of order type $\mathrm{cof}(\beta)$. So by minimality of α, $\mathrm{cof}(\alpha) \leq \mathrm{cof}(\beta)$. Therefore $\mathrm{cof}(\beta) = \mathrm{cof}(\alpha)$. ⊣

As a result, cofinalities are their own cofinality.

5 D • 17. Corollary

Let α be an ordinal. Therefore $\mathrm{cof}(\mathrm{cof}(\alpha)) = \mathrm{cof}(\alpha)$.

More than just this, it turns out that they will be cardinals.

5 D • 18. Theorem

Let α be an ordinal. Therefore $\mathrm{cof}(\alpha)$ is a cardinal.

Proof .:.

Let $\kappa = |\mathrm{cof}(\alpha)|$ with $b : \kappa \to \mathrm{cof}(\alpha)$ a bijection. For each $\xi < \kappa$, define $f(\xi)$ to be the least element of $\mathrm{cof}(\alpha)$ larger than $\max(\sup_{\gamma<\xi} b(\gamma), \sup_{\gamma<\xi} f(\gamma))$.

This is well defined, since the max will always be less than $\mathrm{cof}(\alpha)$. To see this, otherwise, If either supremum (take f for definiteness) has $\sup_{\gamma<\xi} f(\gamma) = \mathrm{cof}(\alpha)$, then $\{f(\gamma) : \gamma < \xi\}$ is a cofinal subset of $\mathrm{cof}(\alpha)$ with order type ξ so that $\mathrm{cof}(\alpha) = \mathrm{cof}(\xi) \leq \xi < \kappa \leq \mathrm{cof}(\alpha)$, which is a contradiction. Therefore $f(\xi)$ is always defined for $\xi < \kappa$.

By definition, f is increasing. Moreover, $\mathrm{im}\, f$ is cofinal in $\mathrm{cof}(\alpha)$, since b is a bijection: each $\zeta < \mathrm{cof}(\alpha)$ has $b(\gamma) = \zeta$ for some $\gamma < \kappa$ so that $f(\gamma+1) > b(\gamma) = \zeta$. Because $\mathrm{im}\, f$ has order type κ, $\mathrm{cof}(\alpha) = \mathrm{cof}(\kappa) \leq \kappa \leq \mathrm{cof}(\alpha)$. Hence $\kappa = \mathrm{cof}(\alpha)$ is a cardinal. ⊣

Hence being a cofinality is a property of cardinals. We also introduce some notation.

5 D • 19. Definition

Let κ be a cardinal. κ is *regular* iff $\mathrm{cof}(\kappa) = \kappa$. κ is *singular* iff $\mathrm{cof}(\kappa) < \kappa$.
For κ a cardinal, κ^+ is the least cardinal greater than κ.

Note that regular cardinals appear all over the place, as do singular cardinals. In particular, all successor cardinals are regular.

5 D • 20. Result

Let κ be a cardinal. Therefore κ^+ is regular.

Proof .:.

Let $X \subseteq \kappa^+$ be cofinal with order type $\alpha = \mathrm{cof}(\kappa^+) < \kappa^+$. By Theorem 5 D • 18, $\alpha \leq \kappa$. Note that X actually forms a partition of κ^+ by looking at the spaces between elements of X. For now, write

$$[\beta, \alpha) = \{\xi \leq \alpha : \beta \leq \xi < \alpha\} = \alpha \setminus \beta text.$$

Define $\alpha \approx \beta$ iff $X \cap [\beta, \alpha) = \emptyset$ and $X \cap [\alpha, \beta) = \emptyset$, meaning $\alpha \approx \beta$ iff (for $\alpha < \beta$) $\alpha \notin X$ and there are no elements of X strictly between α and β. Note that this is an equivalence relation: it's clearly symmetric and reflexive. $\approx$ is transitive since if $\alpha \approx \beta \approx \gamma$ with $\alpha < \gamma$, then one of the following holds:

- $\beta < \alpha < \gamma$, in which case $[\alpha, \gamma) \cap X \subseteq [\beta, \gamma) \cap X = \emptyset$;
- $\alpha < \gamma < \beta$, in which case $[\alpha, \gamma) \cap X \subseteq [\alpha, \beta) \cap X = \emptyset$;
- $\alpha < \beta < \gamma$, in which case $[\alpha, \gamma) = [\alpha, \beta) \cup [\alpha, \gamma)$ so that the intersection with X is $\emptyset \cup \emptyset = \emptyset$.

Note that each equivalence class of $\approx$ has size at most κ, since a class C is bounded by an element of X as it's cofinal: $C \subseteq \sup C + 1 < \sup X = \kappa^+$.

Since the number of equivalence classes is $|\kappa^+_{/\approx}| = |X| = |\alpha| = \alpha$, it follows that as the partition covers κ^+,

$$\kappa^+ \leq \sum_{C \in \kappa^+_{/\approx}} |C| \leq \sum_{C \in \kappa^+_{/\approx}} \kappa = |\kappa^+_{/\approx}| \cdot \kappa = \alpha \cdot \kappa = \max(\alpha, \kappa) < \kappa^+,$$

a contradiction. Thus $\alpha = \mathrm{cof}(\kappa^+) \geq \kappa^+$, and so we have equality. ⊣

Where does all of this talk of regularity get us? Recall that we started this rabbit hole with a question: for which λ is $2^\lambda > \kappa$? It turns out that the answer to this question is unknowable in the sense that different models of set theory will give different answers. But, we do know at least the following, which is often also referred to as "König's theorem".

5 D • 21. Theorem

Let κ be a cardinal. Therefore $\kappa < \kappa^{\mathrm{cof}(\kappa)}$. Moreover, $\kappa < \mathrm{cof}(2^\kappa)$.

Proof .:.

Let $X = \{x_\alpha : \alpha < \mathrm{cof}(\kappa)\}$ be an increasing enumeration of a cofinal subset of κ. By König's Theorem (5 D • 11) and Corollary 5 D • 10, we get that

$$\kappa = \sup_{\alpha < \mathrm{cof}(\kappa)} x_\alpha = \sum_{\alpha < \mathrm{cof}(\kappa)} x_\alpha < \prod_{\alpha < \mathrm{cof}(\kappa)} \kappa = \kappa^{\mathrm{cof}(\kappa)}.$$

Moreover, if we instead choose a κ-length increasing enumeration $Y = \{y_\alpha : \alpha < \kappa\} \subseteq 2^\kappa$, we get that

$$\sup_{\alpha < \kappa} y_\alpha \leq \sum_{\alpha < \kappa} y_\alpha < \prod_{\alpha < \kappa} 2^\kappa = (2^\kappa)^\kappa = 2^{\kappa \cdot \kappa} = 2^\kappa.$$

Hence Y isn't cofinal in 2^κ, and therefore $\mathrm{cof}(2^\kappa) > \kappa$. ⊣

The concept of cofinality also is the source of many other results about regular cardinals, especially successor cardinals.

5 D • 22. Lemma

Let α be an ordinal and $X \subseteq \alpha$. If $|X| < \mathrm{cof}(\alpha)$ then $\sup X < \alpha$.

Proof .:.

If $\sup X = \alpha$ then for β the order type of X, noting that then $\beta \leq |X|$, we have by Lemma 5 D • 16 that $\mathrm{cof}(\alpha) = \mathrm{cof}(\beta) \leq |X| < \mathrm{cof}(\alpha)$, a contradiction. ⊣

5 D • 23. Corollary

Let κ be a regular cardinal. Suppose $2^{<\kappa} = \kappa$. Therefore $\kappa^{<\kappa} = \kappa$.

Proof .:.

$\kappa^{<\kappa}$ is the cardinality of $\bigcup_{\beta\in\kappa} {}^{\beta}\kappa$. Because functions from β to κ are bounded in κ by Lemma 5 D • 22, each such function is a function from β to some $\alpha < \kappa$. In particular, $\bigcup_{\beta\in\kappa} {}^{\beta}\kappa = \bigcup_{\alpha,\beta\in\kappa} {}^{\beta}\alpha$. Thus

$$\kappa^{<\kappa} \leq \kappa \cdot \sup_{\alpha,\beta} |\alpha|^{|\beta|} \leq \kappa \cdot \sup_{\alpha,\beta\in\kappa} 2^{|\alpha|\cdot|\beta|} = \kappa \cdot 2^{<\kappa} = \kappa \dashv$$

The hypothesis that $2^{<\kappa} = \kappa$ is an odd, and seems very strong. This is really a statement similar to the continuum hypothesis.

§ 5 E. The continuum hypothesis

We know that $2^{\aleph_0}$ is some cardinal, and thus is $\aleph_\alpha$ for some α. Cantor's Theorem (5 B • 13) tells us that $2^{\aleph_0} \geq \aleph_1$. Theorem 5 D • 21 tells us that $\mathrm{cof}(2^{\aleph_0}) \geq \aleph_1$ as well. But this is really all we can know.

5 E • 1. Definition

CH is the statement that $2^{\aleph_0} = \aleph_1$.

Without choice, it's not clear that $\mathcal{P}(\omega)$ has an ordinal cardinality. So there are a number of formulations of CH that are equivalent under choice, and we must be careful which we choose if we are in a choiceless context.

Now we introduce a term that is so essential to much of mathematics, it's a wonder we have gotten so far without its introduction.

5 E • 2. Definition

Let X be a set. X is *countable* iff $|X| \leq \aleph_0$.

5 E • 3. Result

CH is equivalent to the statement CH$'$: for every $X \subseteq \mathcal{P}(\omega)$, either X is countable, or $X =_{\text{size}} \mathcal{P}(\omega)$.

Proof .:.

If CH is true, every $X \subseteq \mathcal{P}(\omega)$ has $|X| \leq \aleph_1$ and therefore $|X| < \aleph_1$ or $|X| \leq \aleph_0$. If CH fails, then $2^{\aleph_0} > \aleph_1$. So the bijection $b : \mathcal{P}(\omega) \to 2^{\aleph_0}$ yields a preimage $b^{-1}``\aleph_1$ of size $\aleph_1$ that is a subset of $\mathcal{P}(\omega)$. Hence CH$'$ fails. $\dashv$

CH$'$ is in essence an equivalent formulation of CH, but it is often more appropriate of a formulation, because we can ask if it holds in restricted contexts. Really, CH is a statement about well-orders while CH$'$ is a statement more about subsets of $\mathcal{P}(\omega)$. As such, we can ask which families $X \subseteq \mathcal{P}(\mathcal{P}(\omega))$ have an analogous version of CH$'$ hold of them. We will see later that CH$'$ holds of closed subsets of $\mathbb{R}$, for instance: every closed subset is either countable or of size $2^{\aleph_0}$.

First, we note that $\mathbb{R}$ has size $2^{\aleph_0}$ so that the analogous version of CH$'$ makes sense for subsets of $\mathcal{P}(\mathbb{R})$ rather than just $\mathcal{P}(\mathcal{P}(\omega))$. To be formal, this requires a specific construction of the real numbers, which is not done here. Instead, we rely on a more informal knowledge of real numbers as decimal expansions.

5 E • 4. Theorem

$|\mathbb{R}| = 2^{\aleph_0}$.

Proof .:.

It suffices to show that $|\mathbb{R}| = \aleph_0^{\aleph_0}$ since this is equal to $2^{\aleph_0}$. Each real $r \in \mathbb{R}$ can be identified with a decimal expansion: $r = r_0.r_1r_2r_3\cdots$, meaning an $\omega+1$-length sequence in ω, where $r_0 \in \omega$ and $r_n \in 10$ for $n > 0$. The number of such sequences is $\aleph_0 \cdot 10^{\aleph_0}$, and so there are that many real numbers. But

$$2^{\aleph_0} \leq 10^{\aleph_0} \leq \aleph_0^{\aleph_0} \leq (2^{\aleph_0})^{\aleph_0} = 2^{\aleph_0},$$

and thus $|\mathbb{R}| = \aleph_0 \cdot 2^{\aleph_0} = 2^{\aleph_0}$. $\dashv$

Really this just says that $\mathbb{R} =_{\text{size}} \wp(\omega)$ which is clearly $2^{\aleph_0}$ by the above argument. In a choiceless context, we still get that $\mathbb{R} =_{\text{size}} \wp(\omega) =_{\text{size}} {}^{\omega}2$, but it's not clear that this has an ordinal cardinality: that it can be well-ordered.

One has to have a little care about the decimal expansion in the proof of Theorem 5 E • 4 to ensure that it is unique, for example, $1.000\cdots = 0.999\cdots$[xvii]. But this can be done just by specifying that each decimal expansion should end in an infinite sequence of 0s if it has one that ends in 9s.

With this section, we have introduced all of the axioms of what is commonly referred to as set theory[xviii]. The whole collection of axioms (as well as their actual first-order formulas) are written at the beginning of the document.

[xvii] To see this, note that $r_0.r_1r_2\cdots$ is formally just $\sum_{n\in\omega} r_n \cdot 10^{-n}$ and $0.99999\cdots$ is then equal to

$$\lim_{N\to\omega} \sum_{n=1}^{N} \frac{9}{10^n} = \lim_{N\to\omega} \frac{10^N - 1}{10^N} = \lim_{N\to\omega} 1 - \frac{1}{10^N} = 1.$$

[xviii] There are two notions of "set theory": one is just "set theory" in the sense of "the axioms of sets"; and the other is the field of study with the same name. Often I will use "set theory" as a more informal way of writing ZFC.

Section 6. Another Look at Model Theory

Some have described the field of set theory as being more about the model theory of set theory. Regardless of opinion about this, it does note of a relationship between the two. With the ideas of cardinality at our disposal, we may investigate further some properties of first-order logic. Then we will look more precisely at how these theorems interact with ideas surrounding set theory.

§ 6 A. Further into first-order logic and model theory

The first result we will consider is the idea of a model *generated* by a set, and formulas. There are two or three versions of this theorem. The first two versions are certainly useful for logic, and have the most applications outside of logic, especially algebra, in detailing what is first-order expressible. The third version is the most useful for our purposes, and implies the other two. First we introduce a definition.

6 A • 1. Definition

Let A, and B be $\mathtt{FOL}(\sigma)$-models.
A is a submodel of B, written $\mathsf{A} \subseteq \mathsf{B}$, iff the interpretations of A are the same in B, but restricted to being functions and relations over A.
A is an elementary submodel of B, written $\mathsf{A} \preccurlyeq \mathsf{B}$, iff $A \subseteq B$, and for all $\mathtt{FOLp}(\sigma)$-formulas with parameters in $A \cap B = A$, we have $\mathsf{A} \vDash \varphi$ iff $\mathsf{B} \vDash \varphi$.

It should be clear that being an elementary submodel implies being a submodel just by looking at the atomic $\mathtt{FOLp}(\sigma)$-formulas. But being a submodel does not entail being elementary. For example, the order of the real numbers on the unit interval $\langle (0,1), \leq \rangle$ is the same as for the closed unit interval $\langle [0,1], \leq \rangle$ so that they are submodels: $\langle (0,1), \leq \rangle \subseteq \langle [0,1], \leq \rangle$. But $\langle [0,1], < \rangle \vDash$ "$\exists x \forall y (y \leq x)$" while $\langle (0,1), \leq \rangle \nvDash$ "$\exists x \forall y (y \leq x)$": $\langle [0,1], \leq \rangle$ has a maximal element whereas $\langle (0,1), \leq \rangle$ does not. In essence, being an elementary submodel is the strongest amount of agreement two models can have on first-order formulas. So note the following properties of elementary submodels: for all $\mathtt{FOL}(\sigma)$-models A, B, and C;

- $\mathsf{A} \preccurlyeq \mathsf{A}$.
- $\mathsf{A} \preccurlyeq \mathsf{B} \preccurlyeq \mathsf{A}$ iff $\mathsf{A} = \mathsf{B}$ (since $A \subseteq B \subseteq A$, and they interpret the signature the same way).
- $\mathsf{A} \preccurlyeq \mathsf{B} \preccurlyeq \mathsf{C}$ implies $\mathsf{A} \preccurlyeq \mathsf{C}$.
- $\mathsf{A} \preccurlyeq \mathsf{C}$ and $\mathsf{B} \preccurlyeq \mathsf{C}$ implies $\mathsf{A} \preccurlyeq \mathsf{B} \leftrightarrow A \subseteq B$.

The next theorem, one of the versions of the Löwenheim–Skolem theorem, then tells us that we can *generate* elementary submodels using arbitrary subsets of the original model we start with.

6 A • 2. Theorem (Taking a Skolem Hull)

Let A be an infinite $\mathtt{FOL}(\sigma)$-model, and $X \subseteq A$. Therefore there is a model $\mathsf{Hull}^{\mathsf{A}}(X)$ called the *skolem hull* of X, such that

1. $X \subseteq \mathrm{Hull}^{\mathsf{A}}(X) \subseteq A$;
2. $|\mathrm{Hull}^{\mathsf{A}}(X)| \leq |X| \cdot |\sigma| \cdot \aleph_0$;
3. $\mathsf{Hull}^{\mathsf{A}}(X) \preccurlyeq \mathsf{A}$.

To prove this result, we essentially do a careful proof of Completeness (1 D • 1), building up a model from X by closing under the functions of σ and whatever witnesses existential statements need from A. So the following combinatorial result will be useful in showing that we do not add too many elements in building up the skolem hull.

6 A • 3. Lemma

Let X be a set. Let f be a function with $X \subseteq \operatorname{dom} f$. Therefore the *closure* of X under f—meaning the $\subseteq$-least set Y with $X \subseteq Y$ and $f''Y \subseteq Y$—has size at most $|X| \cdot \aleph_0$.

Proof .:.

Write $X_0 = X$, and define $X_{n+1} = X_n \cup f"X_n$. Let $Y = \bigcup_{n\in\omega} X_n$. Note that for each $x \in Y$, $f(x) \in X_{n+1}$ where $x \in X_n$. Hence $f(x) \in Y$. Thus Y is closed under f. Moreover, for each $n \in \omega$, $|X_{n+1}| \le |X_n| + |X_n| = 2|X_n|$ because $|f"X_n| \le |X_n|$. Therefore, inductively, $|X_n| \le \aleph_0 \cdot |X|$ for each $n \in \omega$. Therefore the union Y has $|X| \le |Y| \le \aleph_0 \cdot \aleph_0 \cdot |X| = \aleph_0 \cdot |X|$. Regardless of whether Y is the $\subseteq$-least set containing X, any $Z \subseteq Y$ which is the real closure of X has $|Z| \le |X| \cdot \aleph_0$. ⊣

As a result, we can close under entire sets of functions as well, and still we can bound the size of the resulting set.

6 A • 4. Corollary

Let X be a set. Let σ be a set of functions with $X \subseteq \operatorname{dom} f$ for each $f \in \sigma$. Therefore the *closure* of X under σ—meaning the $\subseteq$-least set Y with $X \subseteq y$ and $f"Y \subseteq Y$ for each $f \in \sigma$—has size at most $|X| \cdot |\sigma| \cdot \aleph_0$.

Proof .:.

As before, write $X_0 = X$, and define

$$X_{n+1} = X_n \cup \bigcup_{f\in\sigma} (\text{the closure of } X_n \text{ under } f).$$

Thus by Lemma 6 A • 3, $|X_{n+1}| \le |X_n| + |X_n| \cdot |\sigma| \cdot \aleph_0 = |X_n| \cdot |\sigma| \cdot \aleph_0$ For each $n \in \omega$. So inductively, it follows that $|X_n| \le |X| \cdot |\sigma|^n \cdot \aleph_0 = |X| \cdot |\sigma| \cdot \aleph_0$. Taking the union $Y = \bigcup_{n\in\omega} X_n$ yields that Y is closed under each $f \in \sigma$ as in Lemma 6 A • 3, and moreover, $|Y| \le |X| \cdot |\sigma| \cdot \aleph_0^2$. ⊣

Therefore, when we build up the skolem hull, we aren't adding too many elements to X. Note that in the following proof of Taking a Skolem Hull (6 A • 2), indirectly confirm that we have an elementary submodel by the idea of skolem functions: functions which map existential statements to elements that witness them. This allows us to see that the agreement between $\mathbf{A}$ and $\mathbf{Hull}^{\mathbf{A}}(X)$ includes existential statements. The propositional connectives are practically free, and so by induction on formulas, this implies the hull is an elementary submodel.

Proof of Taking a Skolem Hull (6 A • 2) .:.

For each existential $\mathtt{FOL}(\sigma)$-formula $\psi(\vec{x})$ being $\exists v \varphi(v, \vec{x})$, add the function symbol f_ψ (with arity being the length of $\vec{x}$) to the signature. Thus we now consider the signature

$$\sigma' = \sigma \cup \{f_\psi : \psi \text{ is an existential } \mathtt{FOL}(\sigma)\text{-formula}\}.$$

We interpret the functions f_ψ in the model $\mathbf{A}$ by the axiom of choice: for $\psi(\vec{x})$ being $\exists v\ \varphi(v, \vec{x})$, if $\mathbf{A} \vDash$ "$\exists v\ \varphi(v, \vec{x})$", choose $f_\psi^{\mathbf{A}}(\vec{x}) \in W$ such that $\mathbf{A} \vDash \varphi(f_\psi^{\mathbf{A}}(\vec{x}), \vec{x})$. Obviously, if $\mathbf{A} \nvDash \exists v\ \varphi(v, \vec{x})$, then we can set $f_\psi^{\mathbf{A}}(\vec{x})$ to be any particular, fixed element of W that we want (this is only done to ensure that f'_ψ is indeed a function defined over all of W). Hence we can consider the $\mathtt{FOL}(\sigma')$-model $\mathbf{A}'$ with these new interpretations, noting that we have only added interpretations: $X \subseteq W' = W$, for instance.

With this, by Corollary 6 A • 4, we can consider the closure of X under the functions of σ', yielding $\mathrm{Hull}^{\mathbf{A}}(X)$. This clearly has $X \subseteq \mathrm{Hull}^{\mathbf{A}}(X) \subseteq W$, meaning (1) holds. Moreover, by Corollary 6 A • 4, $|\mathrm{Hull}^{\mathbf{A}}(X)| \le |X| \cdot |\sigma| \cdot |\aleph_0|$, meaning (2) holds.

Now we take the model $\mathbf{Hull}^{\mathbf{A}}(X)$ to have the same function and relation interpretations as $\mathbf{A}$, but restricted to $\mathrm{Hull}^{\mathbf{A}}(X)$. To show (3), suppose $w_0, \cdots, w_n \in \mathrm{Hull}^{\mathbf{A}}(X)$. We proceed by induction on the $\mathtt{FOL}(\sigma)$-formula $\varphi(\vec{x})$ to show that $\mathbf{Hull}^{\mathbf{A}}(X) \vDash$ "$\varphi(\vec{w})$" iff $\mathbf{A} \vDash$ "$\varphi(\vec{w})$".

- For $\varphi(\vec{x})$ atomic, the result is immediate by definition.
- For $\varphi(\vec{x})$ being $\neg\psi(\vec{x})$, the inductive hypothesis clearly give the result.
- For $\varphi(\vec{x})$ being $\psi(\vec{x}) \wedge \chi(\vec{x})$, $\mathbf{A} \vDash$ "$\psi(\vec{w}) \wedge \chi(\vec{w})$" iff it models each individually. By the inductive hypothesis, this is equivalent to $\mathbf{Hull}^{\mathbf{A}}(X)$ modeling each individually, meaning $\mathbf{Hull}^{\mathbf{A}}(X) \vDash$ "$\varphi(\vec{w})$".
- For $\varphi(\vec{x})$ being $\exists v\ \psi(v, \vec{x})$, $\mathbf{A} \vDash$ "$\exists v\ \psi(v, \vec{w})$" iff $\mathbf{A} \vDash$ "$\psi(f_\varphi^{\mathbf{A}}(\vec{w}), \vec{w})$". Since $\mathbf{Hull}^{\mathbf{A}}(X)$ is closed under these skolem functions, by the inductive hypothesis, this is equivalent to $\mathbf{Hull}^{\mathbf{A}}(X) \vDash$ "$\psi(f_\varphi^{\mathbf{A}}(\vec{w}), \vec{w})$", iff $\mathbf{Hull}^{\mathbf{A}}(X) \vDash$ "$\exists v\ \psi(v, \vec{w})$".

Hence by induction on $\mathtt{FOL}(\sigma)$-formulas, it follows that $\mathbf{Hull}^{\mathbf{A}}(X) \preccurlyeq \mathbf{A}$, and thus (1)–(3) hold. ⊣

This method of taking skolem hulls, which are effectively the models first-order generated by the set X we look at, is incredibly powerful, and yields the next result, which has some counter intuitive consequences.

6 A • 5. Theorem (Löwenheim–Skolem)

Let T be a $\mathtt{FOL}(\sigma)$-theory with an infinite model. Therefore, for every cardinal $\kappa \geq |\sigma| \cdot \aleph_0$, there is a model $\mathbf{M} \vDash T$ with $|M| = \kappa$.

Proof .:.

Let $\mathbf{A} \vDash T$ be an infinite model. For the downward version, suppose$|\sigma| \cdot \aleph_0 \leq \kappa \leq |A|$, and let $X \subseteq A$ be a subset of size κ. Therefore, by Taking a Skolem Hull (6 A • 2), there is a hull $\mathbf{Hull}^{\mathbf{A}}(X) \preccurlyeq A$ with $X \subseteq \mathrm{Hull}^{\mathbf{A}}(X)$ meaning $\kappa \leq |\mathrm{Hull}^{\mathbf{A}}(X)|$. Moreover, $|\mathrm{Hull}^{\mathbf{A}}(X)| \leq \kappa \cdot |\sigma| \cdot \aleph_0 = \kappa$ so that the hull has size κ. By elementarity, each sentence $\varphi \in T$ has $\mathbf{A} \vDash \varphi$ iff $\mathbf{Hull}^{\mathbf{A}}(X) \vDash \varphi$. Therefore $\mathbf{Hull}^{\mathbf{A}}(X) \vDash T$, and so the hull works.

For the upward version, we use compactness: let $\kappa > |A|$. Consider the expanded signature $\sigma' = \sigma \cup \kappa$, where each ordinal $< \kappa$ is a constant symbol (devoid of its meaning as an ordinal). Now consider the theory $T' = T \cup \{\text{“}\alpha \neq \beta\text{”} : \alpha \neq \beta < \kappa\}$. Any $\mathtt{FOL}(\sigma')$-model $\mathbf{B} \vDash T'$ has $\mathbf{B} \vDash T$ and must have at least κ many elements: the interpretations of the ordinal symbols which are all different. T' has a model by Compactness (1 D • 2): every finite subset $\Delta \subseteq T$ is modeled by an expansion $\mathbf{A}'$ of $\mathbf{A}$ by interpretting the *finitely many* ordinal-symbols in Δ as just different elements of A. Since A has infinitely many elements, we can always do this. Therefore $\mathbf{A}' \vDash \Delta$, and so T' has a model, which is then of size $\lambda \geq \kappa \geq |\sigma| \cdot \aleph_0$. Using the downward statement, it follows that we have a model $\mathbf{M} \vDash T'$ of size exactly κ. ⊣

Some immediate consequences of this are that if ZFC is consistent, then there is a *countable* model in addition to a model of size $\aleph_1$, and models of every cardinality. One might be very confused about this, since supposedly ω and thus $\mathcal{P}(\omega) >_{\text{size}} \aleph_0$ should be in the model, but this isn't necessarily true: the model will contain fewer subsets than in the real world, as not every subset of ω can be described by formulas, and Löwenheim–Skolem (6 A • 5) really only deals with what is minimally required of the formulas of $\mathtt{FOL}$. In essence, the model won't realize its $\mathcal{P}(\omega)$ is small, because it doesn't conain the necessary bijection between its interpretation of $\mathcal{P}(\omega)$ and $\aleph_0$.

Now often in model theory, one deals with chains of elementary submodels. So it's nice to have the following theorem.

6 A • 6. Theorem (Tarski–Vaught Theorem)

Let $\mathbf{A}_\alpha$ be a $\mathtt{FOL}(\sigma)$-model for each $\alpha < \gamma \in \mathrm{Ord}$ with γ a limit ordinal. Suppose $\mathbf{A}_\alpha \preccurlyeq \mathbf{A}_\beta$ for all $\alpha < \beta < \gamma$. Therefore there is a model $\bigcup_{\alpha<\gamma} \mathbf{A}_\alpha$ where $\mathbf{A}_\alpha \preccurlyeq \bigcup_{\beta<\gamma} \mathbf{A}_\beta$ for all $\alpha < \gamma$.

Proof .:.

The “direct limit” $\bigcup_{\alpha\in\gamma} \mathbf{A}_\alpha$ is just given by the union of the corresponding models: the universe the union of the universes, the relations are the unions of the relations, and the functions are the unions of the functions. The constants are necessarily the constants as interpreted by $\mathbf{A}_0$.

Using this, any $\mathtt{FOLp}(\sigma)$-formula φ with parameters in $\mathbf{A}_\alpha$ that is atomic clearly has $\mathbf{A}_\alpha \vDash \varphi$ iff $\bigcup_{\beta<\gamma} \mathbf{A}_\beta \vDash \varphi$. Similarly, the propositional connectives follow easily. For the existential case, it should be clear that $\mathbf{A}_\alpha \vDash$ “$\exists x\ \varphi(x)$” implies $\mathbf{A}_\alpha \vDash$ “$\varphi(a)$” for some $a \in A_\alpha$ and thus by the inductive hypothesis, $\bigcup_{\beta<\gamma} \mathbf{A}_\beta \vDash$ “$\varphi(a)$” and thus $\bigcup_{\beta<\gamma} \mathbf{A}_\beta \vDash$ “$\exists x\ \varphi(x)$”. For the reverse direction, since $\bigcup_{\beta<\gamma} \mathbf{A}_\beta \vDash$ “$\varphi(a)$” for some $a \in \bigcup_{\beta<\gamma} A_\beta$, we have that $a \in A_\beta$ for some $\beta < \gamma$. Therefore by the inductive hypothesis, $\mathbf{A}_\beta \vDash$ “$\varphi(a)$” and thus $\mathbf{A}_\beta \vDash$ “$\exists x\ \varphi(x)$”. But by elementarity, it follows that $\mathbf{A}_\alpha \vDash$ “$\exists x\ \varphi(x)$”. ⊣

Often we don't want to consider an elementary submodel directly, but instead a model which *maps* to an elementary submodel by way of an embedding.

6 A • 7. Definition

Let A and B be $\mathtt{FOL}(\sigma)$-models. For $f : A \to B$ an injective map, the structure $f``\mathsf{A}$ is the structure with universe $f``A$ and with interpretations of σ given by f applied to the interpretations in A.
$f : A \to B$ is an *embedding* (A is *embedded* in B) iff $f``\mathsf{A} \subseteq \mathsf{B}$.
$f : A \to B$ is an *elementary embedding* (A is *elementarily embedded* in B) iff $f``\mathsf{A} \preccurlyeq \mathsf{B}$.

An alternative characterization of being an elementary embedding would be that for every $\mathtt{FOL}(\sigma)$-formula $\varphi(\vec{x})$ and $\vec{a}$ members of A, $\mathsf{A} \vDash$ "$\varphi(\vec{a})$" iff $\mathsf{B} \vDash$ "$\varphi(f(\vec{a}))$". This characterization is arguably a better way of thinking about it. Similarly, f is an embedding iff $\mathsf{A} \vDash$ "$R(\vec{a})$" iff $\mathsf{B} \vDash$ "$R(f(\vec{a}))$" for every relation R and $\vec{a}$ in A, and similarly for functions: $\mathsf{A} \vDash$ "$F(\vec{a}) = a_0$" iff $\mathsf{B} \vDash$ "$F(f(\vec{a})) = f(a_0)$".

But with this added concept, we also can generalize the union model of Tarski–Vaught Theorem (6 A • 6) to the direct limit proper. The direct limit is essentially a least upper bound with respect to embedability. The figure below represents the general idea: any M with embeddings following the diagram is "larger" than the direct limit in that the direct limit embeds in M.

$$\begin{tikzcd}
\mathsf{A} \arrow[dr, "f_{\mathsf{A},\mathsf{C}}"'] \arrow[drr, dashed, "f_{\mathsf{A},\infty}"] \arrow[drrr, dotted, bend left, "f_{\mathsf{A},\mathsf{M}}"] & & & \\
& \mathsf{C} \arrow[r, dashed, "f_{\mathsf{C},\infty}"] & \lim_{\mathcal{F}} \mathcal{A} \arrow[r, dashed] & \mathsf{M} \\
\mathsf{B} \arrow[ur, "f_{\mathsf{B},\mathsf{C}}"] \arrow[urr, dashed, "f_{\mathsf{B},\infty}"'] \arrow[urrr, dotted, bend right, "f_{\mathsf{B},\mathsf{M}}"'] & & &
\end{tikzcd}$$

6 A • 8. Figure: The direct limit embeddings

And here the diagram commutes.

6 A • 9. Definition

Let $\mathcal{A}$ be a set of $\mathtt{FOL}(\sigma)$-models and $\mathcal{F}$ be a set of (elementary) embeddings between models of $\mathcal{A}$. $\langle \mathcal{A}, \mathcal{F} \rangle$ is called a *directed system of (elementary) embeddings* iff

- for each $\mathsf{A}, \mathsf{B} \in \mathcal{A}$, there is at most one $f : A \to B$ in $\mathcal{F}$, denoted $f_{\mathsf{A},\mathsf{B}}$ with $f_{\mathsf{A},\mathsf{A}} = \mathrm{id} \restriction A$;
- for each $\mathsf{A}, \mathsf{B} \in \mathcal{A}$ there is some $\mathsf{C} \in \mathcal{A}$ with $f_{\mathsf{A},\mathsf{C}}, f_{\mathsf{B},\mathsf{C}} \in \mathcal{F}$; and
- if $f_{\mathsf{A},\mathsf{B}}, f_{\mathsf{B},\mathsf{C}} \in \mathcal{F}$, then there is an (elementary) embedding $f_{\mathsf{A},\mathsf{C}} \in \mathcal{F}$ with $f_{\mathsf{B},\mathsf{C}} \circ f_{\mathsf{A},\mathsf{B}} = f_{\mathsf{A},\mathsf{C}}$.

For $\langle \mathcal{A}, \mathcal{F} \rangle$ a directed system of embeddings, the *direct limit* is the $\mathtt{FOL}(\sigma)$-model $\mathbf{lim}_{\mathcal{F}}\, \mathcal{A}$ such that

1. there is an embedding $f_{\mathsf{A},\infty} : A \to \lim_{\mathcal{F}} \mathcal{A}$ such that $f_{\mathsf{B},\infty} \circ f_{\mathsf{A},\mathsf{B}} = f_{\mathsf{A},\infty}$ whenever $f_{\mathsf{A},\mathsf{B}}$ exists; and
2. for every model M satisfying (1) in place of $\mathbf{lim}_{\mathcal{F}}\, \mathcal{A}$, there is an embedding $f : \lim_{\mathcal{F}} \mathcal{A} \to M$ such that $f \circ f_{\mathsf{A},\infty} = f_{\mathsf{A},\mathsf{M}}$ for all $\mathsf{A} \in \mathcal{A}$.

The general idea behind directed systems of embeddings is that we can continually embedd things in a "larger" model, and we can do so in a way where the embeddings work well together. The idea behind the direct limit is that there should be an upper bound to this: one where everything in $\mathcal{A}$ embedds into it (in a way that works nicely with), and it's the "least" such model. So the direct limit should be thought of as a least upper bound on a poset[xix] $\langle \mathcal{A}, R \rangle$ where $\mathsf{A}\ R\ \mathsf{B}$ iff $f_{\mathsf{A},\mathsf{B}} \in \mathcal{F}$, although the least upper bound isn't necessarily in $\mathcal{A}$.

Of course, we should confirm that every directed system of embeddings has a direct limit, which mostly just amounts to checking that a certain construction works. The general idea behind the construction is that we just take the disjoint copies of all the models, and then take the union as in Tarski–Vaught Theorem (6 A • 6). From there, we mod out by "eventual equivalence" when transformed by elements of $\mathcal{F}$.

[xix] The directed set doesn't necessarily form a poset, since the existence of embeddings need not be antisymmetric, but it will at least be reflexive and transitive.

6 A • 10. Result

Let $\langle \mathcal{A}, \mathcal{F} \rangle$ be a directed system of embeddings between $\texttt{FOL}(\sigma)$-models. Therefore the direct limit $\mathbf{lim}_{\mathcal{F}} \mathcal{A} = \mathsf{D}$ exists, and is isomorphic to the disjoint union of $\mathcal{A}$ modulo eventual equivalence through the embeddings of $\mathcal{F}$.

Proof .:.

Without loss of generality, each $A \cap B = \emptyset$ for $\mathsf{A}, \mathsf{B} \in \mathcal{A}$ just by replacing A with the isomorphic model A' with universe $A \times \{A\}$, tagging each element of the universe with the model it comes from. This ensures that the union $\bigcup \mathcal{A}$ is a disjoint union. This union is defined as in Tarski–Vaught Theorem (6 A • 6): the universe is the disjoint union $\bigcup_{\mathsf{A} \in \mathcal{A}} A$, and the relations and functions are the disjoint unions of the corresponding relations and functions. This union model does not interpret the constant symbols of σ. To form a $\texttt{FOL}(\sigma)$-model D, consider the relation on $\bigcup_{\mathsf{A} \in \mathcal{A}} A$ defined by (for $x \in \mathsf{A}$ and $y \in \mathsf{B}$) $x \approx y$ iff there is some C where $f_{\mathsf{A},\mathsf{C}}(x) = f_{\mathsf{B},\mathsf{C}}(y)$, meaning that x and y are eventually equal in the embeddings.

Claim 1

$\approx$, eventual equivalence through the embeddings of $\mathcal{F}$, is an equivalence relation on $\bigcup_{\mathsf{A} \in \mathcal{A}} A$.

Proof .:.

Clearly $\approx$ is reflexive (take $f_{\mathsf{A},\mathsf{A}} = \mathrm{id}$ to witness this) and symmetric (by symmetry of $=$). $\approx$ is transitive, since if $f_{\mathsf{A},\mathsf{B}}(x) = f_{\mathsf{A}',\mathsf{B}}(y)$ and $f_{\mathsf{A}',\mathsf{C}}(y) \approx f_{\mathsf{A}'',\mathsf{C}}(z)$, then for some M with $f_{\mathsf{B},\mathsf{M}}, f_{\mathsf{C},\mathsf{M}} \in \mathcal{F}$, it follows by injectivity and the embeddings working well together that

$$\begin{aligned} f_{\mathsf{A},\mathsf{M}}(x) &= f_{\mathsf{B},\mathsf{M}} \circ f_{\mathsf{A},\mathsf{B}}(x) \\ &= f_{\mathsf{B},\mathsf{M}} \circ f_{\mathsf{A}',\mathsf{B}}(y) \\ &= f_{\mathsf{A}',\mathsf{M}}(y) \\ &= f_{\mathsf{C},\mathsf{M}} \circ f_{\mathsf{A}',\mathsf{C}}(y) \\ &= f_{\mathsf{C},\mathsf{M}} \circ f_{\mathsf{A}'',\mathsf{C}}(z) = f_{\mathsf{A}'',\mathsf{M}}(z). \end{aligned}$$

Hence $\approx$ is transitive, and so an equivalence relation. ⊣

Now consider the model D with universe $(\bigcup_{\mathsf{A} \in \mathcal{A}} A)_{/\approx}$—the equivalence classes of $\approx$—and corresponding relation and function interpretations as per Result 2 C • 11: $R^{\mathsf{D}}([x_0]_{\approx}, \cdots, [x_n]_{\approx})$ iff for $x_i \in A_i$ and $\mathsf{B} \in \mathcal{A}$ such that $f_{\mathsf{A}_i,\mathsf{B}} \in \mathcal{F}$ for each $i \leq n$, $R^{\mathsf{B}}(f_{\mathsf{A}_i,\mathsf{B}}(x_0) \cdots, f_{\mathsf{A}_n,\mathsf{B}}(x_n))$. As the $f \in \mathcal{F}$ are embeddings and so respect R, this will be well-defined. We do the same process for the functions of σ. Note that the constant symbols work out nicely after modding out by $\approx$: each constant symbol c of σ is interpreted as $[c^{\mathsf{A}}]_{\approx}$ for any $\mathsf{A} \in \mathcal{A}$. As embeddings, the constant symbols are mapped to the corresponding constant symbols, and thus eventual equivalence always holds between c^{A} and c^{B}. This completes the construction of D. Now we must show that D is the direct limit.

Firstly, note that each $\mathsf{A} \in \mathcal{A}$ has an embedding $f_{\mathsf{A},\mathsf{D}} : A \to D$ defined by $a \mapsto [a]_{/}$. This is an embedding, because for $a_0, \cdots, a_n \in A$, $R^{\mathsf{D}}([a_0]_{\approx}, \cdots, [a_n]_{\approx})$ by definition is equivalen to $R^{\mathsf{B}}(f_{\mathsf{A},\mathsf{B}}(a_0), \cdots, f_{\mathsf{A},\mathsf{B}}(a_n))$ for some B with $f_{\mathsf{A},\mathsf{B}} \in \mathcal{A}$. In particular, for $\mathsf{B} = \mathsf{A}$, this is just $R^{\mathsf{A}}(a_0, \cdots, a_n)$. The same idea applies for functions and constants to show that $f_{\mathsf{A},\mathsf{D}}$ is an embedding. Moreover, this embedding plays nicely with the $f \in \mathcal{F}$, since eventual equivalence yields $f_{\mathsf{A},\mathsf{B}}(a) \approx a$ so that $f_{\mathsf{B},\mathsf{D}} \circ f_{\mathsf{A},\mathsf{B}}(a) = [f_{\mathsf{A},\mathsf{B}}(a)]_{\approx} = [a]_{\approx} = f_{\mathsf{A},\mathsf{D}}(a)$. So (1) holds of Definition 6 A • 9.

To see that D is the least such model—that (2) holds of Definition 6 A • 9—suppose M has the same property. Let $f_{\mathsf{D},\mathsf{M}}$ be defined by, for $x \in A$ and $\mathsf{A} \in \mathcal{A}$, $f_{\mathsf{D},\mathsf{M}}([x]_{\approx}) = f_{\mathsf{A},\mathsf{M}}(x)$. This is well defined since if $x \approx f_{\mathsf{B},\mathsf{A}}(x)$ then

$$f_{\mathsf{D},\mathsf{M}}(f_{\mathsf{B},\mathsf{A}}(x)) = f_{\mathsf{B},\mathsf{M}} \circ f_{\mathsf{A},\mathsf{B}}(x) = f_{\mathsf{A},\mathsf{M}}(x) = f_{\mathsf{D},\mathsf{M}}(x)$$

for any $\mathsf{B} \in \mathcal{A}$ with $f_{\mathsf{A},\mathsf{B}} \in \mathcal{F}$. In fact, $f_{\mathsf{D},\mathsf{M}}$ will be injective since $[x]_{\approx} \neq [y]_{\approx}$ implies the transformations of x and y by $f \in \mathcal{F}$ are always different so that applying the embedding $f_{\mathsf{A},\mathsf{M}}$ where A contains transformations of both x and y, the transformations are still different in M. The reverse direction holds in the same way. To see that $f_{\mathsf{D},\mathsf{M}}$ respects the relations, functions, and constant symbols of σ, suppose c is a constant symbol of σ. $f_{\mathsf{D},\mathsf{M}}(c^{\mathsf{D}}) = f_{\mathsf{A},\mathsf{M}}(c^{\mathsf{A}}) = c^{\mathsf{M}}$ for any $\mathsf{A} \in \mathcal{A}$ as $f_{\mathsf{A},\mathsf{M}}$ is an embedding and $c^{\mathsf{D}} = [c^{\mathsf{A}}]_{\approx}$. For the relation R, if

$R^{\mathsf{D}}([x_0]_{\approx}, \cdots, [x_n]_{\approx})$, then R^{A} holds of the transformations of the $\vec{x}$ where A contains all of these transformed eleemnts. But then applying $f_{\mathsf{A},\mathsf{M}}$ yields that the relation holds of the $f_{\mathsf{A},\mathsf{M}}$ transformations of the $\vec{x}$ so that the relation holds of the $f_{\mathsf{D},\mathsf{M}}$ transformations. The reverse direction is the same, and the argument for functions proceeds similarly. Therefore $f_{\mathsf{D},\mathsf{M}}$ is an embedding, and so D is the direct limit. ⊣

Thus Tarski–Vaught Theorem (6 A • 6) can be reformulated as saying that if we have a chain of elementary embeddings, then each is elementarily embedded in the direct limit. So to generalize this, we have the following result, whose proof is precisely the same as Tarski–Vaught Theorem (6 A • 6), although translated through the elementary embeddings of $\mathcal{F}$ instead of the elements themselves.

6 A • 11. Corollary

Let $\langle \mathcal{A}, \mathcal{F} \rangle$ be a directed system of elementary embeddings. Therefore A is elementarily embedded in $\mathbf{lim}_{\mathcal{F}} \mathcal{A}$ for each $\mathsf{A} \in \mathcal{A}$.

The point of all of this talk about elementary embeddings will become clear in the next chapter. But it is an important idea if we want to learn about V, as the first-order truths of V are then reflected in any model it elementarily embeds into. So-called large cardinals often state the existence of elementary embeddings from V into another model, and so commonly uses the techniques of this subsection.

§ 6 B. Logic within set theory

One might worry that, since the above ideas depend on set theory, although the meta-theoretic ZFC can prove the above results about first-order logic the formal[xx] ZFC can't. Readers worried about this can put their minds at ease. But although objects in ZFC are hereditarily sets, we can still code non-set things like formulas using sets.

Rather than give a tedious account of the syntax of first-order logic, and an even more tedious account of how to formalize this, we merely give an impression on how these things are formalized in ZFC.

6 B • 1. Definition

The *logical symbols of formal first-order logic* is the set ω, consisting of the codes for logical symbols

$$`\wedge' = 0, \quad `\neg' = 1, \quad `\exists' = 2, \quad `(' = 3, \quad `)' = 4, \quad `,' = 5$$

and variables $`v_n' = n + 6$ for $n < \omega$.
For A a set, a *variable assignment* for A is a function $f : \{`v_n' : n \in \omega\} \to A$, i.e. a function $f : \omega \setminus 6 \to A$.
For σ a set of relations, functions, and constants, a σ*-formula* is a $\varphi \in (\omega \sqcup \sigma)^{<\omega}$ obeying the usual syntax rules.
For σ a set of relations, functions, and constants, a σ*-proof* is a finite sequence of σ-formulas that obeys the usual syntax rules for proofs.

Once we have the syntax of first-order logic in ZFC, we can start to address the satisfaction relation. This is done by induction on formulas. Firstly, we have a couple definitions that allows us to more precisely see why we can do this in ZFC: the relations are well-founded.

6 B • 2. Definition

Suppose $\preccurlyeq_n$ is a linear order on B_n for each $n \in \omega$. Then the *length-prioritized lexicographic ordering* $\preccurlyeq_{\text{lex}} \subseteq \bigcup_{N \in \omega} \prod_{n=0}^{N} B_n$ is the order defined by, for $f : n \to \bigcup_{k \in \omega} B_k$ and $g : m \to \bigcup_{k \in \omega} B_k$ where $n, m \in \omega$,

$$f \preccurlyeq_{\text{lex}} g \leftrightarrow f = g \vee |f| < |g| \vee (|f| = |g| \wedge \text{the least } k \in \omega \text{ with } f(k) \neq g(k) \text{ has } f(k) <_k g(k)).$$

This is really just the dictionary order on ${}^{<\omega}B$ where each component potentially has its own ordering. This is best understood when each $<_n$ is the same ordering on $B = B_n$. In particular, working with triplets, $\prec \, = \, <_0 \, = \, <_1 \, = \, <_2$, $\preccurlyeq_{\text{lex}}$ just orders ${}^3B$ as follows: $\langle a, b, c \rangle <_{\text{lex}} \langle a', b', c' \rangle$ iff

- $a \prec a'$; or
- $a = a'$ and $b \prec b'$; or
- $a = a'$ and $b = b'$ and $c \prec c'$.

[xx]"Formal" here means relating to formulas, i.e. "syntactic", rather than the opposite of casual.

Sequences of different lengths are compared in the same way, but the longer one comes after in the order. To save space, we would also write the above as the more intelligible conditions:

- $a \prec a'$; or else
- $b \prec b'$; or else
- $c \prec c'$.

Definition 6 B • 2 just generalizes this to larger product sequences with more relations. What's important for us is when this is a well-ordering.

6 B • 3. Lemma

Suppose each $\leqslant_n \subseteq B_n \times B_n$ is a well-order of B_n. Therefore $\leqslant_{\text{lex}}$ is a well-order of $\bigcup_{N\in\omega} \prod_{n=0}^{N} B_n$.

Proof .:.

It should be clear that $\leqslant_{\text{lex}}$ is a linear order of $B^{<\omega}$: transitivity follows since each $<_n$ is transitive. Totality clearly holds since any two distinct sequences differ somewhere, and since each $<_n$ is total, wherever they differ is ordered. Clearly anti-symmetry holds by anti-symmetry of each $<_n$. So $\leqslant_{\text{lex}}$ is clearly linear, and all that suffices is to show well-foundedness.

Let $\langle f_n : n \in \omega\rangle$ be $<_{\text{lex}}$-decreasing. Therefore $\langle \text{dom}(f_n) : n \in \omega\rangle$ is non-increasing. So without loss of generality, we can assume each $\text{dom}(f_n) < k$ for some $k \in \omega$. For each $m < k$ consider $\langle f_n(m) : n < \omega \wedge m \in \text{dom}(f_n)\rangle$. If each of these is finite or eventually stabilizes, then eventually f_{n+1} is an intial segment of f_n. If this were the case, then the only way for $\langle f_n : n \in \omega\rangle$ to be $<_{\text{lex}}$-decreasing is for their lengths to be decreasing, contradicting the well-foundedness of ω. Thus for some m, $\langle f_n(m) : n < \omega \wedge \text{dom}(f_n)\rangle$ is infinite and doesn't stabilize. Take the least $m \in \omega$ for which this happens. Therefore, eventually, $f_{n+1}(m) <_m f_n(m)$, contradicting the well-foundedness of $<_m$. ⊣

The point of having $<_{\text{lex}}$ prioritize length is to ensure that inductive hypotheses hold for subformulas: for ψ a subformula of φ, $\psi \leqslant_{\text{lex}} \varphi$. Hence, we can proceed by induction on formulas.

6 B • 4. Corollary

For any signature σ well-ordered by $<_\sigma$, the σ-formulas are well-ordered by $<_{\text{lex}}$.

Proof .:.

Strictly speaking, the order is on $\omega \cup \sigma$ where were merely place all elements of ω before σ. In other words, for α the order-type of $\langle \sigma, <_\sigma\rangle$, the order on $\omega \sqcup \sigma$ is given by $\omega + \alpha$. Hence $<_{\text{lex}}$ well-orders $(\omega \sqcup \sigma)^{<\omega}$, which contain all of the σ-formulas. ⊣

For set structures, V can define the satisfaction relation by induction on formulas (the property of being a subformula is well-founded, as formulas are certain finite sequences of an alphabet). In fact, we can define this relation *uniformly*.

6 B • 5. Definition

For A a set, and σ a signature, an *interpretation* of σ in A is a map ς with $\sigma = \text{dom}(\varsigma)$ where for R an n-placed relation, $\varsigma(R) = R^{\mathsf{A}} \subseteq A^n$, and similarly for functions and constants.

6 B • 6. Theorem

Let σ be a signature, A be a set, v a variable assignment for A, ς an interpretation of σ in A, and x a $\sigma \sqcup A$-formula coding the real-world formula $\psi(\vec{y})$. Therefore, there is a $\texttt{FOL}(\in)$-formula "models$(\sigma, A, \varsigma, v, x)$" such that

$$\langle A, \varsigma\rangle \vDash \text{“}\psi(v(\vec{y}))\text{”} \quad \text{iff} \quad \mathsf{V} \vDash \text{“models}(\sigma, A, \varsigma, v, x)\text{”}.$$

Proof .:.

In particular, models$(\sigma, A, \varsigma, v, x)$ iff

- x is a $\sigma \sqcup A$-formula;
- ς is a set of relations, functions, and constants over A;
- v is a *variable assignment*, a function mapping variables to elements of A;

- there is a function $f_v : (A \cup \sigma)^{<\omega} \to \{0, 1\}$ such that for any $z \in (A \cup \sigma)^{<\omega}$,
  - z is of the form "$R(\vec{y})$" where $R \in \sigma$ is a relation symbol, and $v(\vec{y}) \in \varsigma(R) \leftrightarrow f_v(z) = 1$,
  - z is of the form "$y_0 = y_1$" and $v(y_0) = v(y_1) \leftrightarrow f_v(z) = 1$,
  - z is of the form "$\psi \wedge \theta$" and $f_v(z) = 1$ iff both $f_v(\text{"}\psi\text{"}) = 1$ and $f_v(\text{"}\theta\text{"}) = 1$.
  - z is of the form "$\neg\psi$" and $f(\text{"}\psi\text{"}) = 1$ iff $f(z) = 0$, and
  - z is of the form "$\forall t\ \psi(\vec{y}, t)$" and $f(z) = 1$ iff any function h_w obeying these rules for all $z' \leqslant_{\text{lex}} z$—where w is any variable assignment w for A with $v \setminus \{\langle a, v(a)\rangle\} \subseteq w$—has $h_w(\text{"}\psi(a, t)\text{"}) = 1$; and
- $f_v(x') = 1$ where x' has every free variable y in x replaced by $v(y)$. ⊣

The above is really only a partial proof, since it only applies to signatures with no function symbols. But this isn't an issue with the result, it just makes the defining formula even longer to have an auxiliary function interpreting terms through the variable assignment.

Doing this then allows us to confirm by the same sort of proofs before that Completeness (1 D • 1), Compactness (1 D • 2), and so forth hold. But the important thing about Theorem 6 B • 6 is that ZFC has the ability to understand when something is true in a given set model. We will often use Theorem 6 B • 6 without stating so, because the idea of a formula being true of a (set) structure is so widely used. Of course, we may not have access to classes since they aren't objects in the universe.

§ 6 C. Common applications to set theory

For now, our main application will be with respect to Taking a Skolem Hull (6 A • 2) and elementarity. The great thing about taking skolem hulls of transitive sets is that we end up with well-founded sets, and thus can collapse them.

6 C • 1. Result

Let $\mathbf{A} = \langle A, R\rangle$ be well-founded. Therefore any $\mathbf{B} = \langle B, R'\rangle$ embedded in $\mathbf{A}$ is also well-founded.

Proof .:.

Let $f : B \to A$ be an embedding and let $X \subseteq B$ be arbitrary. Since $\mathbf{A}$ is well-founded, $f"X \subseteq A$ has an R-minimal element $a \in f"X$. Thus for every $y \in f"X$, $\neg y\ R\ a$. As an embedding, $\neg(f^{-1}(y)\ R'\ f^{-1}(a))$ for each $y \in f"X$, meaning $\neg(x\ R'\ f^{-1}(a))$ for each $x \in X$. Therefore $f^{-1}(a)$ is R'-minimal. Thus $\mathbf{B}$ is also well-founded. ⊣

6 C • 2. Corollary

Let T be a transitive set and $X \subseteq T$. Therefore $\mathbf{Hull}^{\langle T, \in\rangle}(X)$ is well-founded, and is isomorphic to the transitive collapse $\mathbf{cHull}^{\langle T, \in\rangle}(X)$, which is then elementarily embedded in $\langle T, \in\rangle$. Moreover, if X is transitive, X is left uncollapsed: the collapsing map $\pi : \text{Hull}^{\langle T, \in\rangle}(X) \to \text{cHull}^{\langle T, \in\rangle}(X)$ has $\pi \restriction X = \text{id} \restriction X$.

Proof .:.

Write T' for $\text{cHull}^{\langle T, \in\rangle}(X)$. By Taking a Skolem Hull (6 A • 2), $\mathbf{Hull}^{\langle T, \in\rangle}(X) \preccurlyeq \langle T, \in\rangle$ so that the hull is well-founded. By elementarity, the hull satisfies the axiom of extensionality. By The Mostowski Collapse (4 • 1), the hull is isomorphic to the transitive $\langle T', \in\rangle$ by the map inductively defined by $\pi(x) = \{\pi(a) : \mathbf{Hull}^{\langle T, \in\rangle}(X) \vDash \text{"}a \in x\text{"}\}$. Note that as a substructure of $\mathbf{V}$, for $a, x \in H$, $\mathbf{Hull}^{\langle T, \in\rangle}(X) \vDash \text{"}a \in x\text{"}$ iff $a \in x$. Moreover, $\mathbf{Hull}^{\langle T, \in\rangle}(X) \vDash \text{"}a \in x\text{"}$ implies $a \in H$ just by virtue of the semantics. Therefore $\mathbf{Hull}^{\langle T, \in\rangle}(X) \vDash \text{"}a \in x\text{"}$ iff $a \in x \cap H$ and thus $\pi(x)$ is equal to $\{\pi(a) : a \in x \cap H\}$. In particular, if X is transitive, the inductive hypothesis tells us that $\pi(x)$ for $x \in X$ is equal to $\{\pi(a) : a \in x \cap H\} = \{a : a \in x \cap H\} = x \cap H$. Since X is transitive, $x \subseteq X \subseteq H$ so that $x \cap H = x$. Therefore $\pi(x) = x$ and so $\pi \restriction X = \text{id} \restriction X$ by induction on rank. ⊣

Note that the use of "collapse" especially makes sense here, because every $\pi(x) \in \text{cHull}^{\mathbf{T}}(X)$ has $\text{rank}(\pi(x)) \leq$

rank(x). Of course, strict inequality requires that $x \not\subseteq \mathrm{Hull}^{\mathsf{T}}(X)$. Using Tarski–Vaught Theorem (6 A • 6) and direct limits in general, we can build up skolem hulls to have less and less collapsed while still being relatively small.

In particular, if we take the hull that includes all of an ordinal, we get a model that contains all of the ordinals below it. Using the elementary chains, this allows us to conclude the following, showing we can get ordinals in our uncollapsed model before collapsing.

6 C • 3. Corollary

Let T be a transitive set with $\kappa \in T$ an uncountable, regular cardinal and $X \subseteq T$ of size $< \kappa$. Therefore, there is an elementary $\mathsf{H} \preccurlyeq \langle T, \in \rangle$ with $H \cap \mathrm{Ord}$ an ordinal, $|H| < \kappa$, and $X \subseteq H$.

Proof .:.

Take the skolem hull $\mathsf{H}_0 = \mathsf{Hull}^{\langle T, \in \rangle}(X)$. This may not have $H_0 \cap \mathrm{Ord}$ as an ordinal although it will satisfy that $\mathsf{H}_0 \preccurlyeq \langle T, \in \rangle$ and $|H_0| \leq \aleph_0 \cdot 1 \cdot |X| < \kappa$. For H_n already defined, if $H_n \cap \mathrm{Ord}$ is an ordinal, then stop the process, and take $\mathsf{H} = \mathsf{H}_n$. Otherwise let $\mathsf{H}_{n+1} = \mathsf{Hull}^{\langle T, \in \rangle}(H_n \cup \sup(H_n \cap \mathrm{Ord}))$. As a regular cardinal, $\sup(H_n \cap \mathrm{Ord}) < \kappa$ because inductively $|H_n| < \kappa$, which also tells us that $|H_{n+1}| < \kappa$. Define H_ω to be the direct limit of the H_ns for $n < \omega$ as in Tarski–Vaught Theorem (6 A • 6).

Note that $\mathsf{H}_\omega \preccurlyeq \langle T, \in \rangle$ with $X \subseteq H_\omega$ and $|H_\omega| \leq \aleph_0 \cdot \sup_{n \in \omega} |H_n|$. As each $|H_n| < \kappa$ and κ has cofinality $\kappa > \omega$, it follows that $\sup_{n \in \omega} |H_n| < \kappa$ and thus $|H_\omega| < \kappa$. To see that $H_\omega \cap \mathrm{Ord}$ is an ordinal, it suffices to show that $H_\omega \cap \mathrm{Ord}$ is transitive. For $\beta \in H_\omega \cap \mathrm{Ord}$, it follows that $\beta \in H_n \cap \mathrm{Ord}$ for some $n < \omega$. Thus $\beta \subseteq H_{n+1}$ and so $\beta \subseteq H_\omega \cap \mathrm{Ord}$. ⊣

Theorems and ideas like this will play a big role in what we can do with small models of fragments of set theory as well as inner models (which haven't been defined yet). To proceed further in this direction, we will need to consider the `FOLp` agreement between $\mathbf{V}$ and other transitive sets in the next section.

Section 7. Absoluteness

Absoluteness in some sense refers to how correct our definitions of concepts are. Formally, a definition is absolute between two models if the two agree on what the definition applies to. For example, $x \in y$ is absolute between any two transitive models containing x and y: $\mathbf{A} \vDash$ "$x \in y$" iff $\mathbf{B} \vDash$ "$x \in y$" since they both interpret membership the same way.

7 • 1. Definition

Let $\varphi(\vec{x})$ be a FOL-formula. Let $\mathbf{A}$ and $\mathbf{B}$ be models. We say that φ is *absolute* between $\mathbf{A}$ and $\mathbf{B}$ iff $\mathbf{A} \vDash \varphi(\vec{a})$ iff $\mathbf{B} \vDash \varphi(\vec{a})$ whenever $\vec{a}$ are parameters belonging to both A and B.

7 • 2. Corollary

If $\mathbf{A}$ and $\mathbf{B}$ model some theory T, and if $T \vdash$ "$\forall \vec{x}\ (\varphi \leftrightarrow \psi)$", then φ is absolute between $\mathbf{A}$ and $\mathbf{B}$ iff ψ is.

This general definition isn't much to work with. It does, however, tell us that many of our set-theoretic conceptions are not absolute between models of set theory. For example, the argument in Result 4 A • 9 shows that well-foundedness isn't absolute. A similar argument shows that even a set being infinite isn't even absolute between models of set theory: consider any particular infinite set $A \in \mathrm{V}$ and the signature $\sigma \cup A \cup \{`A'\}$ where we have a constant symbolf for A and every element of A. Then consider the theory $T = \mathsf{ZFC} + \{$"$a \in A \wedge A$ is finite" $: a \in A\}$. We can always interpret the symbol 'A' as some finite subset of A, and in particular, for any finite subtheory of T, the set of all constants appearing in the subtheory. This shows that each finite subset of T is satisfiable, and thus that T has a model where A is finite.

So these ideas mean first-order logic on its own doesn't tell us much about the deeper structure of $\mathbf{V}$. They only tell us is that asking whether something is absolute in general, without any further restrictions, is not a good question to ask. So for the most part, we will restrict our view to models which are transitive. And in doing so, we also can refine this notion a bit. Note that we are assuming, as structures, that the models are non-empty.

7 • 3. Definition

For $A \subseteq \mathrm{V}$ transitive, write $\mathbf{A}$ for $\langle A, \in \rangle$, and call $\mathbf{A}$ a *transitive model*. Let $\mathbf{A} \subseteq \mathbf{B}$ be two transitive models, and let $\varphi(\vec{x})$ be a FOL-formula.

- φ is *downward-absolute between* $\mathbf{A}$ *and* $\mathbf{B}$ iff $\mathbf{B} \vDash \varphi(\vec{a})$ implies $\mathbf{A} \vDash \varphi(\vec{a})$ whenever $\vec{a}$ are in $A \cap B = A$.
- φ is *upward-absolute between* $\mathbf{A}$ *and* $\mathbf{B}$ iff $\mathbf{A} \vDash \varphi(\vec{a})$ implies $\mathbf{B} \vDash \varphi(\vec{a})$ whenever $\vec{a}$ are in A.
- φ is *downward-absolute* iff φ is downward-absolute between all transitive models and submodels.
- φ is *upward-absolute* iff φ is upward-absolute between all transitive models and submodels.
- φ is *absolute* iff φ is absolute between all transitive models.

Equivalently, φ is absolute iff φ is absolute between $\mathbf{V}$ and its transitive submodels. Because membership for transitive models is always the same, we get that "$x \in y$" is absolute. "$x = \emptyset$" is absolute: for $\mathbf{A}$ transitive with $x \in A$, x is non-empty iff there is a $y \in x \subseteq A$. By absoluteness of "$y \in x$" and transitivity, this is equivalent to $\mathbf{A} \vDash$ "$\exists y\ (y \in x)$", which just says $\mathbf{A} \vDash$ "$x \neq \emptyset$". Hence "$x \neq \emptyset$" (and thus "$x = \emptyset$") is absolute.

The above idea should indicate that absoluteness can often be proven in a kind of inductive way, beginning with simple formulas like "$x \in y$" and working with increasingly more complex formulas. We can prove a great number of absoluteness results by studying this kind of complexity, which is mostly just due to the number of quantifiers. But because we're working with transitive models, *bounded* quantifiers do not increase complexity: a bounded quantifier $Qx \in X$ ranges over the same elements (namely the elements of X) in $\mathbf{V}$ as in the transitive model $\mathbf{A}$, because both properly understand what it means to be a member of X. And this is really the best understanding of transitivity: properly understanding membership. The following hierarchy of formulas is called the *Lévy hierarchy* after Azriel Lévy (לוי עזריאל).

7 • 4. Definition

Let φ be a FOL-formula. A *bounded quantifier* is a quantifier of the form "$\exists x \in X$" or "$\forall x \in X$" for some x and X, being short-hand for "$\exists x(x \in X \wedge \cdots)$" and "$\forall x(x \in X \to \cdots)$" respectively.

- φ is Σ_0 (and Π_0) iff all quantifiers occurring in φ are bounded.
- φ is Σ_{n+1} iff φ is of the form $\exists x\ \psi$ where ψ is Π_n.
- φ is Π_n iff φ is of the form $\neg\psi$ where ψ is Σ_n.

Note that the X in "$\exists x \in X$" is a variable rather than a parameter. We also get a variant hierarchy where we allow parameters. In particular, something is $\Sigma_n(A)$ iff it satisfies the same definition, but allows parameters in A. If we allow parameters, we get more formulas. For example, we could bound quantifiers by elements of A, allowing more $\Sigma_0(A)$-formulas than the standard Σ_0-formulas.[xxi]

This is the first hierarchy of formulas we will encounter, although we will encounter many more later revolving around $\mathbb{R}$ and $\mathbb{N}$. The fact that transitive sets understand bounded quantifiers (when they contain the parameters), tells us that Σ_0-formulas are absolute. Note that we can formalize this by first noting that truth in transitive classes can be known by $\mathbf{V}$, just by bounding quantifiers.

7 • 5. Definition

Let C be a (FOLp-definable) transitive class and φ a FOLp-formula with parameters in C. The formula φ^{C} is the formula where each quantifier "$\exists x$" and "$\forall x$" is replaced by "$\exists x \in \mathrm{C}$" and "$\forall x \in \mathrm{C}$", respectively.

Alternatively, we can define φ^{C} by induction on φ:

"$(x = y)^{\mathrm{C}}$"	is	"$x = y$";
"$(x \in y)^{\mathrm{C}}$"	is	"$x \in y$";
"$(\neg\varphi)^{\mathrm{C}}$"	is	"$\neg\varphi^{\mathrm{C}}$";
"$(\varphi \wedge \psi)^{\mathrm{C}}$"	is	"$(\varphi^{\mathrm{C}} \wedge \psi^{\mathrm{C}})$"; and
"$(\exists x\ \varphi)^{\mathrm{C}}$"	is	"$\exists x \in \mathrm{C}\ \varphi$".

More explicitly, since membership in C is really a formula, for C is defined by ψ, then "$(\exists x\ \varphi)^{\mathrm{C}}$" is "$\exists x\ (\psi(x) \wedge \varphi^{\mathrm{C}})$". The same idea applies to C a set with $\psi(x)$ just being $x \in C$. Note that if we already have a bounded quantifier, the restriction of "$\exists x \in X$" to C then gives "$\exists x \in X \cap \mathrm{C}$", and similarly "$\forall x \in X$" maps to "$\forall x \in X \cap \mathrm{C}$". Therefore, we can recast truth about φ in $\mathbf{C}$ as truth of φ^{C} in $\mathbf{V}$. Note that we already knew how to do with with sets by Theorem 6 B • 6, but not classes in general.

7 • 6. Lemma

Let C be a transitive class and φ a FOLp-formula with parameters in C. Therefore $\langle \mathrm{C}, \in \rangle = \mathbf{C} \vDash \varphi$ iff $\mathbf{V} \vDash \varphi^{\mathrm{C}}$.

Proof .:.

Proceed by induction on formula complexity. As a transitive class, $\mathbf{C} \vDash$ "$x \in y$" iff $\mathbf{V} \vDash$ "$x \in y$" so that the result holds if φ is atomic. The sentential connectives follow easily from the inductive hypothesis. So suppose φ is of the form $\exists x\ \psi$.

- If $\mathbf{C} \vDash$ "$\exists x\ \psi$", then $\mathbf{C} \vDash$ "$\psi(c)$" for some $c \in \mathrm{C}$ and thus inductively—since $\psi^{\mathrm{C}}(c)$ is $(\psi(c))^{\mathrm{C}}$—$\mathbf{V} \vDash$ "$\psi^{\mathrm{C}}(c)$". Thus $\mathbf{V} \vDash$ "$\exists x \in \mathrm{C}\ \psi^{\mathrm{C}}$" which is just to say that $\mathbf{V} \vDash$ "φ^{C}".
- Conversely, if $\mathbf{V} \vDash$ "$\exists x \in \mathrm{C}\ \psi^{\mathrm{C}}$", then for some $c \in \mathrm{C}$, $\mathbf{V} \vDash$ "$\psi^{\mathrm{C}}(c)$", which inductively says $\mathbf{C} \vDash$ "$\psi(c)$" and thus $\mathbf{C} \vDash \varphi$. $\dashv$

So just by rewriting the definition, we get the following, alternative characterization of absoluteness.

[xxi] Note that, a priori, not every formula can be placed in the Lévy hierarchy as we've stated it here. Other sources will do away with this issue by allowing *blocks* of quantifiers rather than single ones. This is avoided here both to show the importance of the background assumptions, and to show how to get around the issue. In particular, "$\exists x \exists y(x \in y)$" can't be placed in the hierarchy. It is only by assuming some additional set theory that this formula is *equivalent* to one in the Lévy hierarchy.

More precisely, every formula is equivalent to one in prenex normal form: all quantifiers appear at the beginning of the formula. Assuming some basic set theory, each block of quantifiers of the form $Qx_0\ Qx_1 \cdots Qx_n\varphi$ can instead be written as $Qx(x = \langle x_0, \cdots, x_n\rangle \wedge \varphi')$ where φ' replaces each x_i with the defined notion of being the ith entry in x, something which can be said using only bounded quantifiers. This allows us to show each formula is equivalent to one in the Lévy hierarchy. In this sense, we say that φ is Π_n (or Σ_n) iff φ is equivalent to a formula which is Π_n (or Σ_n). In doing so, however, we need to specify the theory they are equivalent under.

7 • 7. Corollary

Let φ be a FOL-formula. Therefore φ is absolute iff $\mathbf{V} \vDash$ "$\varphi \leftrightarrow \varphi^{\mathrm{C}}$" for each transitive class C.

§ 7 A. Easy absoluteness results

Important examples of absolute formulas include all of the Σ_0-formulas of the Lévy hierarchy.

7 A • 1. Result

Let φ be a Σ_0-formula. Therefore φ is absolute.

Proof .:.

Proceed by structural induction on φ. Since all quantifiers in φ are bounded, and $\forall x \in X\ \psi$ is equivalent to $\neg\exists x \in X\ \neg\psi$, we only need to consider the sentential operations and the bounded quantifier $\exists x \in X\ \psi$. The sentential operations are immediate by induction. So it suffices to consider bounded quantification: suppose φ is $\exists x \in X\ \psi$. Let **A** be an arbitrary, transitive model with $X \in A$ so that $X \subseteq A$. Inductively, for all $a \in A$, $\psi^A(a)$ holds iff $\psi(a)$ holds.

- Since $X \subseteq A$, if there is an $a \in X$ such that $\psi(a)$ holds then there is an $a \in A$ such that $a \in X$ and $\psi^A(a)$ holds, i.e. φ^A holds.
- Conversely, if φ^A holds, then there is some $a \in X \cap A = X$ such that $\psi^A(a)$ holds. Inductively, this means $\psi(a)$ holds and thus $\exists x \in X\ \psi$. ⊣

More generally, this says that if φ is absolute between transitive **A** and **B**, then "$\exists x \in X\ \varphi$" is absolute between them as well when X is in both. And of course, boolean combinations[xxii] of absolute formulas are absolute as well. This is stated as follows with the same proof as Result 7 A • 1.

7 A • 2. Result

Let **A** be a transitive model. Therefore the set of FOLp-formulas absolute between **A** and **V** is closed under bounded quantification, conjunctions, and negations.

7 A • 3. Corollary

The following axioms are absolute, because they are true in all non-empty, transitive models:
- the axiom of extensionality,
- the axiom of the empty set, and
- the axiom of foundation.

Proof .:.

Since all of these are true in **V**, the only way for these to fail to be absolute is if they are false in some transitive model. So we will show this does't happen. All of these can be shown through careful analysis of the forms of the axioms.

- Extensionality says that for every x and y, "$x = y \leftrightarrow \forall v \in x\ (x \in y) \wedge \forall v \in y\ (v \in x)$" holds. The property of extensionality holding at x, y is Σ_0 and thus absolute. So if it holds for all $x, y \in$ V, then it holds for all $x, y \in$ C for any class C. Therefore extensionality holds in C, and is thus absolute.
- By (3) of Corollary 2 E • 4, the universe has the empty set in it. By the argument just after Definition 7 • 3, "$x = \emptyset$" is absolute and thus the axiom of the empty set is satisfied and thus absolute: "$\exists x\ (x = \emptyset)$" is absolute.
- The axiom of foundation says that for every $x \neq \emptyset, \exists y \in x\ \forall z \in y (z \notin x)$. Since "$x \neq \emptyset$" is absolute and the other part is Σ_0, foundation holding for x is absolute, meaning for every $x \in$ C, **C** believes foundation holds for x, because it holds in **V**. Therefore C must satisfy foundation. ⊣

Result 7 A • 1 gives a great number of absoluteness results. For the most part, we will not give the completely formal definitions that show these are Σ_0. Instead, like with much of first-order logic, we will resort to giving impressions

[xxii] meaning formulas built up from the sentential connectives starting from some given formulas

and instructions that allow one to carefully check that they are. The following are absolute all because they are defined by Σ_0-formulas.

- x being an (un-ordered) pair: everything in x is either some $y \in x$ or $z \in x$.
- x being an ordered pair.
- x being the first-coordinate of an ordered pair y: there is a $z \in y$ such that for every $w \in z$, $w = x$.
- x being a relation.
- x being the domain of R: for every $z \in x$, there is a pair $\langle z, y\rangle \in R$, and vice versa.
- x being the range of R.
- x being a function: for every y in the domain of x, there is a unique z in the range of x with $\langle y, z\rangle \in x$.
- x being the output of a function f with input y, i.e. $x = f(y)$.
- x being an injective function.
- x being a surjective function.
- x being a subset of y: every $z \in x$ is in y.
- x being transitive: every $z \in x$ is a subset of x.
- x being an ordinal: x is transitive, and $\in$ linearly orders x.

And many, many more concepts are absolute by Result 7 A • 1. As a result of the above absoluteness examples, we have some nice consequences about what it means for transitive sets to model the axioms of set theory. Most of the axioms of set theory state the closure of the universe under certain sets. Pair, for example, says that for every x and y, $\{x, y\}$ exists. Now while the property of *being* an (un-ordered) pair is absolute, this doesn't tell us that the *existence* of un-ordered pairs is absolute. Similarly, being a subset is absolute, but being the powerset isn't, because a model might contain fewer subsets than another: $4 = \{0, 1, 2, 3\}$ thinks $\mathcal{P}(2)$ exists and is $3 = \{0, 1, \{0, 1\}\}$, because every subset of 2 that 4 contains is in 3: $\{1\} \notin 4$ although $\{1\} \subseteq 2$.

But we can have a better picture of what these sorts of defined sets will look like, because their defining formulas are absolute. Really, the following is just another way to state absoluteness.

7 A • 4. Result

Let φ be a FOLp-formula absolute between a transitive model M and V. Therefore $\{x : \varphi(x)\}^{\mathsf{M}} = \{x : \varphi(x)\} \cap M$.

Proof .:.

$\mathsf{M} \vDash$ "$\varphi(x)$" iff $x \in M$ and $\varphi^M(x)$ holds. By absoluteness, this is equivalent to $x \in M$ and $\varphi(x)$. ⊣

As a result, $\mathcal{P}^{\mathsf{M}}(X) = \mathcal{P}(X) \cap M$, $\mathrm{Ord}^{\mathsf{M}} = \mathrm{Ord} \cap M$, and so on. This idea also gives an understanding of when transitive models satisfy (some of the) axioms of set theory.

7 A • 5. Corollary

Let A be a transitive model. Therefore,

- $\mathsf{A} \vDash$ Union iff $x \in A$ implies $\bigcup x \in A$.
- $\mathsf{A} \vDash$ Pair iff $x, y \in A$ implies $\{x, y\} \in A$.
- $\mathsf{A} \vDash$ Comp iff $\{x \in y : \varphi^A(x)\} \in A$ for each $y \in A$ and FOLp-formula φ.
- $\mathsf{A} \vDash$ P iff $x \in A$ implies $\mathcal{P}(x) \cap A \in A$.
- $\mathsf{A} \vDash$ Inf if $\omega \in A$.

Proof .:.

For Union, Pair, and Comp, by absoluteness, the only way A can interpret "$x = \bigcup y$" and "$x = \{y, z\}$" is the same way V does. The only way A can interpret "$x = \{y \in z : \varphi(y)\}$" is the way V interprets "$x = \{y \in z : \varphi^A(y)\}$". Hence A being closed under unions, pairing, or comprehension as A interprets it (i.e. satisfying Union, Pair, or Comp) is the same as being closed under unions, pairing, or comprehension as V interprets it.

For P, note that being a subset is absolute. Hence
$$\mathcal{P}^{\mathsf{A}}(x) = \{y \in A : \mathbf{A} \vDash \text{“}y \subseteq x\text{”}\} = \{y \in A : y \subseteq x\} = \mathcal{P}(x) \cap A.$$
Hence being closed under powersets (i.e. satisfying P) is the same as being closed under powersets intersected with the universe.

For `Inf`, clearly if $\omega \in A$, then by the absoluteness results above, **A** satisfies `Inf`. The reverse may not hold, since
$$\omega \cup \{\{1\} + n : n < \omega\} \cup \big\{\omega \cup \{\{1\} + n : n < \omega\}\big\},$$
where $x + 1 = x \cup \{x\}$ and $x + n = ((x + 1) + \cdots) + 1$, is a transitive set that models the axiom of infinity, but ω is merely a subset of the universe (and a subset of a set in the universe), not a set inside it. ⊣

Of course, not everything turns out to be absolute, but we can get partial absoluteness for some formulas, as we've used in Corollary 7 A • 3. For example, we have the following easy consequences of Result 7 A • 1.

7 A • 6. Result

Π_1-formulas are downward absolute. Σ_1-formulas are upward absolute.

Proof .:.

Let φ be $\forall x\ \theta$ where θ is Σ_0. If $\mathbf{V} \vDash \forall x\ \theta$, then, in particular, $\theta(a)$ holds for every $a \in$ C. By absoluteness, $\theta(a) \leftrightarrow \theta^{\mathrm{C}}(a)$ and thus $\mathbf{C} \vDash \forall x\ \theta$.

Let ψ be $\exists x\ \theta$ where θ is Σ_0. Thus $\neg\psi$, being $\forall x\ \neg\theta$ is downward absolute. Taking the contrapositive means that ψ is upward absolute: $\mathbf{V} \vDash$ “$\neg\psi \to \neg\psi^{\mathrm{C}}$” implies $\mathbf{V} \vDash$ “$\psi^{\mathrm{C}} \to \psi$”. ⊣

Again, more generally, when φ is absolute between $\mathbf{A} \subseteq \mathbf{B}$, then $\exists x\ \varphi$ is upward absolute between them and $\forall x\ \varphi$ is downward absolute between them. We cannot ask for stronger than this mere partial absoluteness. For example, the existence of ω—meaning the axiom of infinity—is Σ_1:
$$\text{“}\exists N(\emptyset \in N \wedge \forall x \in N(x \cup \{x\} \in N))\text{”},$$
but the transitive set $\{\emptyset\}$ doesn't have such an N although **V** does. So upward absoluteness is all we can say about Σ_1-formulas in general. Similarly, in $\{\emptyset\}$, we have the Π_1-sentence “$\forall x\ (x = \emptyset)$” as true although it's false for **V**. So downward absoluteness is all we can say about Π_1-formulas in general.

§ 7 B. The Lévy hierarchy and absoluteness with some set theory

As Corollary 7 A • 3 shows, relatively few things will be absolute, especially if they require more axioms of set theory to even state properly. For example, while *being* the union of two sets is absolute, the *existence* of such a set isn't absolute. For example, $\{0, 1, \{1\}\}$ is transitive, but $2 = \{1\} \cup 1$ isn't in the set. So often it will be useful to restrict our attention to transitive models of some fragment of ZFC.

7 B • 1. Definition

Let T be a theory, and φ a `FOL`-formula.
We say that φ is Σ_n^T (or Π_n^T) iff $T \vdash$ “$\varphi \leftrightarrow \psi$” for some Σ_n (or Π_n) formula ψ.
We say that φ is Δ_n^T iff φ is both Σ_n^T and Π_n^T.

As a result, $\Sigma_n^\emptyset$ just consists of all formulas logically equivalent to Σ_n-formulas, and similarly for $\Pi_n^\emptyset$. As a result, the placement of a formula φ isn't unique: if φ is a Σ_n^T-*sentence*, then “$\forall x \varphi$”—just adding on a dummy quantifier—is a Π_{n+1}^T-sentence that is logically equivalent to φ.

Note also that if φ is Σ_n^T and $T \subseteq T'$, then φ is $\Sigma_n^{T'}$ as well. Hence absoluteness results for $T \subseteq$ ZFC extend to absoluteness results for ZFC.

7 B • 2. Corollary

$\Delta_1^\emptyset$-formulas are absolute.

Proof .:.

If φ is a $\Delta_1^{\emptyset}$-formula, then φ is $\Pi_1^{\emptyset}$ and thus downward absolute by Result 7 A • 6; and φ is $\Sigma_1^{\emptyset}$ and thus upward absolute by Result 7 A • 6. ⊣

7 B • 3. Corollary

Well-foundedness is absolute between transitive models of Lemma 4 • 3, e.g. of $\mathsf{ZF} - \mathsf{P}$.

Proof .:.

Well-foundedness is downward absolute because a relation R is ill-founded iff the following Σ_1-formula holds: $\exists x\ \forall y \in x\ \exists z \in x\ (z\ R\ y)$. This means well-foundedness is Π_1, and thus downward absolute.

Upward absoluteness holds as it is Σ_1^T for T such a theory as in the statement of the corollary: it states the existence of a pair: a function and an ordinal which constitute a rank function. By Lemma 4 • 3: if $R \subseteq A \times A$ is well-founded in a model $\mathbf{C}$ of this, then $\mathbf{C}$ believes that there is a rank function $f : A \to \mathrm{Ord}^{\mathbf{C}}$. Since the following are Σ_0 and so absolute between transitive models:

- being a function, and being $f(x)$;
- being an ordinal—which implies $\mathrm{Ord}^{\mathbf{C}} = \mathrm{Ord} \cap \mathrm{C}$;
- being 0;
- being R-minimal—which is Σ_0 as seen by "$\forall y \in A\ (\neg y\ R\ x)$";
- being $x + 1$; and
- being the supremum of a set of ordinals,

it follows that $f : A \to \mathrm{Ord} \cap C$ is still a rank function in $\mathbf{V}$. Hence there can be no infinite R-decreasing sequence in A without the ranks decreasing and so violating the well-foundedness of the ordinals. Therefore, well-foundedness is upward absolute between such models, and hence absolute between such models. ⊣

We also get that functions and sets defined by transfinite recursion using absolute notions will be absolute.

7 B • 4. Theorem

Suppose "$F(x) = y$" is absolute between transitive models of $\mathsf{ZF} - \mathsf{P}$. Let $G : \mathrm{Ord} \to \mathrm{V}$ be defined by transfinite recursion: $G(\beta) = F(G \restriction \beta)$. Therefore "$G(x) = y$" is absolute between transitive models of $\mathsf{ZF} - \mathsf{P}$.

Proof .:.

Applying transfinite recursion in $\mathbf{M}$, $G^{\mathbf{M}}$ is such that for every ordinal $\alpha \in \mathrm{Ord} \cap M$, $G^{\mathbf{M}}(\alpha) = F^{\mathbf{M}}(G^{\mathbf{M}} \restriction \alpha)$. By transfinite *induction* on α, the inductive hypothesis that $G^{\mathbf{M}} \restriction \alpha = G \restriction \alpha$ and the absoluteness of F implies $G^{\mathbf{M}}(\alpha) = F(G \restriction \alpha) = G(\alpha)$, as desired. Hence $G^{\mathbf{M}} = G \cap M$. ⊣

As a result, the rank of a set is absolute between transitive models of $\mathsf{ZF} - \mathsf{P}$ being defined by transfinite recursion.

7 B • 5. Corollary

"$\mathrm{rank}(x) = \alpha$" is absolute between transitive models of $\mathsf{ZF} - \mathsf{P}$. Hence $\mathrm{V}_\alpha^{\mathbf{M}} = \mathrm{V}_\alpha \cap M$ for transitive models $\mathbf{M} \vDash \mathsf{ZF} - \mathsf{P}$.

Proof .:.

$\mathrm{rank}(x) = \alpha$ is absolute between such models as a consequence of Lemma 4 • 3 where we take membership to be the well-founded relation. As noted, the proof works just as well for classes as for sets, since it just relies on transfinite induction and recursion.

To show $\mathrm{V}_\alpha^{\mathbf{M}} = \mathrm{V}_\alpha \cap M$, just note that we can define $x \in \mathrm{V}_\alpha$ iff $\mathrm{rank}(x) < \alpha$, which is absolute. ⊣

Some of the most common axioms satisfied are those of basic set theory.

7 B • 6. Definition

Basic set theory (BST) consists of the following axioms:

1. extensionality, empty set, foundation;
2. comprehension, pairing, union; and
3. the existence of cartesian products: $\forall x\ \forall y\ \exists z\ \forall w\ (w \in z \leftrightarrow \exists a \in x\ \exists b \in y\ (w = \langle a, b \rangle))$.

We know from Corollary 7 A • 3 that (1) is already absolute and satisfied by all transitive sets. So the addition of (2) adds more absoluteness between models we care about. (3) is not an explicit axiom of ZFC, but it does follow from both powerset and replacement. Since we will work in contexts in which either might be missing, we use the weaker result that cartesian products exist.

The reason for this is that under BST, a greater number of things are equivalent, and the Lévy hierarchy will be closed under various operations. What we mean by this is that for T a theory, a formula φ is Σ_n^T iff φ is equivalent over T to a Σ_n formula. So we mean that Σ_n^{BST} is a much larger class than $\Sigma_n^{\emptyset}$. For the most part, we will just need slight weakenings of ZFC, but working in more generality will help later. In particular, Σ_1^{BST} is closed under existential quantification: $\exists y\ \exists x\ \psi$ for ψ being Σ_0 is equivalent to $\exists z\ \exists y \in z\ \exists x \in z\ \psi$ by pairing in BST. Similarly, Π_1^{BST} is closed under universal quantification in addition to $\vee$ and $\wedge$.

7 B • 7. Result

For each $n < \omega$, Σ_n^{BST} is closed under existential quantification, disjunction, and conjunction.
Π_n^{BST} is closed under universal quantification, disjunction, and conjunction.
Δ_n^{BST} is closed under bounded quantification, disjunction, conjunction, and negation.

To get more than this, we need more set theory. For example, in full ZFC, each Σ_n^{ZFC} is closed under existential quantification and both bounded quantifiers (and similarly for Π_n^{ZFC} and universal quantification). Showing this, however, requires some complicated tricks better suited for the end of the chapter.

But these calculations give some partial absoluteness about cardinality and cofinality.

7 B • 8. Result

Being a cardinal is Π_1^{BST} and therefore downward absolute between models of BST.
κ being singular is Σ_1^{BST} and therefore upward absolute between models of BST. Hence being regular is Π_1^{BST} and so downward absolute between models of BST.

Proof .:.

κ being a cardinal is equivalent to being an ordinal (which is Σ_0) and

$$\underbrace{\text{“}\forall x\ \underbrace{\forall f\ \underbrace{(x \in \kappa \wedge f \text{ is a function from } x \text{ to } \kappa \rightarrow f \text{ is not bijective})}_{\Sigma_0}}_{\Pi_1}\text{”.}}_{\Pi_1^{\mathsf{BST}}}$$

The calculation above shows that this is equivalent over BST to a Π_1-formula.

κ being singular is equivalent to there being an increasing function whose image is cofinal in κ and whose domain is an ordinal less than κ:

$$\underbrace{\text{“}\exists x\ \exists f\ \underbrace{(x < \kappa \wedge f \text{ is a function from } x \text{ to } \kappa \wedge f \text{ is increasing} \wedge \operatorname{im} f \text{ is unbounded in } \kappa)}_{\Sigma_0}\text{”}}_{\Sigma_1^{\mathsf{BST}}} \dashv$$

It should be noted that any $T \supseteq \mathsf{BST}$ also has these absoluteness results because T will prove the equivalences that BST does. But we can do more of these kinds of calculations to get that ω is absolute when ω is in the model.

7 B • 9. Result

Finiteness is absolute between transitive models of BST.

Proof .:.

1. $x = y \cup \{y\} = y + 1$ is an absolute relation.
2. x being a limit ordinal is equivalent to x being an ordinal (absolute between transitive models) and $\forall y \in x\ (y \cup \{y\} \in x)$, which is absolute by (1).
3. x being the least ordinal of a set of ordinals X is absolute. To see this, the least ordinal of X can be defined by a Σ_0-formula: x is the least ordinal of X iff x is an ordinal and $x \in X$ and $\forall y \in X(y \text{is an ordinal} \rightarrow x = y \vee x \in y)$, which is Σ_0.
4. x being ω is the same as being the least ordinal in the class of limit ordinals (if there are any).
5. x being $n < \omega$ is just defined by iteratively considering (1).

x being finite is just to say that there is some $n \in \omega$ with a bijection $f : x \rightarrow n$. This form is Σ_1^{BST} and thus upward absolute between transitive models of BST. For downward absoluteness, suppose $\mathbf{A} \vDash \mathsf{BST}$ with $x \in A$. If x is really finite, then there is some $n < \omega$ where we can then write out that $x = \{x_0, \cdots, x_{n-1}\}$ and so define f by $f = \{\langle x_0, 0\rangle, \cdots, \langle x_{n-1}, n-1\rangle\}$, just using a single Σ_0-formula. Since $n \in A$ by pairing and union, it follows by cartesian products and comprehension that $f \in A$ and thus $f : x \rightarrow n$ is a bijection showing x is finite in $\mathbf{A}$. ⊣

BST is mostly brought up because almost every model we would like to consider will be a model of it. With stronger theories, like ZF, we get more absoluteness, and learn more about **V**.

§ 7 C. Toy models for set theory

So far we've investigated the absoluteness between models of various fragments of set theory, but we haven't given many concrete examples of what these models look like. Firstly, the levels of the cumulative hierarchy serve as a nice introduction to models of (fragments of) set theory. Just by their form, we immediately get some axioms holding in them.

7 C • 1. Result

Let $\alpha \in \mathrm{Ord}$. Therefore $\mathbf{V}_\alpha = \langle \mathrm{V}_\alpha, \in\rangle \vDash \mathsf{BST} - \mathsf{Pair}$ − "the existence of cartesian products".

Proof .:.

Extensionality, empty set, and foundation all hold by Corollary 7 A • 3, because V_α is transitive. It suffices by Corollary 7 A • 5 to show that V_α is closed under pairing, unions, and subsets. But just by the rank argument given in Result 4 A • 10, $x, y \in \mathrm{V}_\alpha$ implies $x \cup y \in \mathrm{V}_\alpha$, and $y \subseteq x \in \mathrm{V}_\alpha$ implies $y \in \mathrm{V}_\alpha$. In particular, $\{z \in x : \varphi^{\mathrm{V}_\alpha}(z)\} \in \mathrm{V}_\alpha$. ⊣

The issue with pairing and the existence of cartesian products is that they increase rank. Thus for V_α to be closed under these, α should be a limit ordinal. If this is the case, then $\mathbf{V}_\alpha$ models much more than just BST. In particular, $\mathbf{V}_\alpha$ models almost all of ZFC.

7 C • 2. Result

Let α be a limit ordinal. Therefore $\mathbf{V}_\alpha = \langle \mathrm{V}_\alpha, \in\rangle \vDash \mathsf{ZFC} - \mathsf{Rep} - \mathsf{Inf}$.

Proof .:.

Recall that $\mathrm{V}_\beta \subseteq \mathrm{V}_\gamma$ for $\beta < \gamma$. Hence as a limit ordinal, $\mathrm{V}_{\beta+n} \subseteq \mathrm{V}_\alpha$ for any $\beta < \alpha$ and $n < \omega$. Let $x \in \mathrm{V}_{\alpha_x+1}$ and $y \in \mathrm{V}_{\alpha_y+1}$ be arbitrary with $\alpha_x, \alpha_y < \alpha$.

- For Pair, $x, y \in \mathrm{V}_{\max(\alpha_x,\alpha_y)+1}$ and thus $\{x, y\} \in \mathrm{V}_{\max(\alpha_x,\alpha_y)+2}$. As a limit ordinal, $\alpha_x, \alpha_y < \alpha$ implies $\max(\alpha_x, \alpha_y) + 2 < \alpha$ and thus $\{x, y\} \in \mathrm{V}_\alpha$. So by Corollary 7 A • 5, $\mathbf{V}_\alpha \vDash \mathsf{Pair}$.
  $\mathbf{V}_\alpha$ satisfies Comp trivially, since any subset $y \subseteq x \in \mathrm{V}_{\alpha_x+1}$ (like $y = \{z \in x : \varphi^{\mathrm{V}_\alpha}(z)\}$) has $y \subseteq x \subseteq \mathrm{V}_{\alpha_x}$ and thus $y \in \mathrm{V}_{\alpha_x+1}$ by definition of cumulative hierarchy. In fact, $\mathcal{P}(x) \in \mathrm{V}_{\alpha_x+2} \subseteq \mathrm{V}_\alpha$. So since $\mathcal{P}^{\mathrm{V}_\alpha}(x) = \mathcal{P}(x) \cap \mathrm{V}_\alpha = \mathcal{P}(x)$, it follows that $\mathbf{V}_\alpha \vDash \mathsf{P}$ too.

- For Union, if $x \in V_{\alpha_x+1}$, then $\text{rank}(x) > \sup\{\text{rank}(y)+1 : y \in x\}$ Hence $y \subseteq V_{\alpha_x+1}$ for each $y \in x$. Therefore $\bigcup x \subseteq V_{\alpha_x+1}$ and so $\bigcup x \in V_{\alpha_x+2}$. By Corollary 7 A • 5, $\mathbf{V}_\alpha \vDash$ Union. As a result, the existence of cartesian products holds, since $x \times y \in \wp(\wp(\wp(x \cup y))) \in V_\alpha$ for each $x, y \in V_\alpha$, and being the cartesian product is absolute.
- For AC, for any non-empty family of non-empty, disjoint sets $F \in V_\alpha$, there is a set C in V that has chosen one element from each set in F. Note that $C \in \wp(\bigcup F) \in V_\alpha$ so that $C \in V_\alpha$. ⊣

So by Corollary 7 A • 5, since $\text{rank}(\omega) = \omega$, $\mathbf{V}_\alpha \vDash$ Inf whenever $\alpha > \omega$ is a limit. This should be taken to be evidence of the consistency of ZFC; the two independent hurdles to this being the axiom of infinity, and replacement.

7 C • 3. Corollary

Let $\alpha > \omega$ be a limit ordinal. Therefore $\mathbf{V}_\alpha \vDash \text{ZFC} - \text{Rep}$. Moreover, $\mathbf{V}_\omega \vDash \text{ZFC} - \text{Inf}$, and in fact $\mathbf{V}_\omega \vDash \neg\text{Inf}$.

Proof .:.

$\omega \in V_\alpha$ implies $\mathbf{V}_\alpha \vDash \text{ZFC} - \text{Rep}$ by Result 7 C • 2 and Corollary 7 A • 5. As for $\mathbf{V}_\omega$, note that every element of V_ω is finite: by induction, $V_0 = \emptyset$, and $|V_{n+1}| = 2^{|V_n|+1}$ is finite as well. Hence $V_\omega = \bigcup_{n<\omega} V_n$ has no infinite set in it and so $\mathbf{V}_\omega \vDash \neg\text{Inf}$, as any set following such a definition would be infinite (in $\mathbf{V}$).

To see that $\mathbf{V}_\omega \vDash$ Rep, suppose φ is a `FOLp`-formula that defines a function over $D \in V_\omega$, which is to say

$$\mathbf{V}_\omega \vDash \text{“}\forall x \in D\ \exists! y\ \varphi(x, y)\text{”}.$$

We now wish to show that the image of φ is in V_ω. Note that in $\mathbf{V}$, there is then a function $f : D \to V_\omega$. As D is finite, there is some finite subset $R \subseteq V_\omega$ with $f : D \to R$. As each $r \in R$ has rank $n_r < \omega$ and R is finite, it follows that the rank of R is $\max\{n_r + 1 : r \in R\} < \omega$ and thus $R \in V_\omega$. Hence $\mathbf{V}_\omega \vDash$ Rep. ⊣

Of course, by Gödel's incompleteness theorem—assuming that ZFC is consistent—we can't construct from ZFC alone a model of ZFC, as this would imply ZFC $\vdash$ Con(ZFC). But the issues with replacement and the axiom of infinity can be dealt with at the cost of the powerset axiom.

Note that in the proof of Corollary 7 C • 3, the reason why replacement holds in $\mathbf{V}_\omega$ is due to rank being bounded: the domain is small enough, and so the outputs are bounded. If we could ensure that our toy model was "regular" in a similar sense as with V_ω, we can ensure replacement holds. To make this idea precise, we have the following definition.

7 C • 4. Definition

Let $\kappa \geq \aleph_0$ be a cardinal. Let H_κ be the set of *hereditarily* $<\kappa$*-sized sets* defined by $x \in H_\kappa$ iff $|x| < \kappa$ and every $y \in \text{trcl}(x)$ has $|y| < \kappa$.

To give a concrete example, $V_\omega = H_{\aleph_0}$ by similar reasoning as in Corollary 7 C • 3. There's an alternative characterization of H_κ for regular κ. For the most part, we will not be interested in H_κ for singular κ, since it will not model as much set theory as with regular cardinals.

7 C • 5. Result

For $\kappa \geq \aleph_0$ a regular cardinal, $x \in H_\kappa$ iff $|\text{trcl}(x)| < \kappa$.

Proof .:.

If $|\text{trcl}(x)| < \kappa$, then clearly each $y \in \text{trcl}(x)$ has $|y| < \kappa$ because the transitive closure is transitive: $y \subseteq \text{trcl}(x)$. So suppose $|x| < \kappa$ and every $y \in \text{trcl}(x)$ has $|y| < \kappa$. Note that $\text{trcl}(x) = \bigcup_{n<} \bigcup^n x$. Clearly $|\bigcup^0 x| = |x| < \kappa$. Inductively, $|\bigcup^n x| < \kappa$ so that $\bigcup^{n+1} x = \bigcup(\bigcup^n x)$ is the union of $<\kappa$-many sets of size $<\kappa$. As a regular cardinal, it follows that this has size $<\kappa$. Hence $\text{trcl}(x)$, being the union of countably many sets of size $<\kappa$, has size $<\kappa$. ⊣

Clearly $H_\kappa \subseteq H_\lambda$ for $\kappa < \lambda$, and these are all transitive. Now just by it's definition, it's not clear that H_κ is a set. But by dealing just with regular cardinals—for each singular cardinal λ, $H_\lambda \subseteq H_{\lambda^+}$—we can show that each H_κ is a set. The proof of this is non-trivial, and we be the first real use of The Mostowski Collapse (4 • 1).

7 C • 6. Result

For $\kappa \geq \aleph_0$ a regular cardinal, $\mathrm{H}_\kappa \subseteq \mathrm{V}_\kappa$ is a set.

Proof .:.

Let $x \in \mathrm{H}_\kappa$ be arbitrary. Write $T = \operatorname{trcl}(x \cup \{x\})$. Proceed by induction on rank to show every $y \in T$ has rank less than κ. For $y = \emptyset$, this is obvious. For y of rank $\alpha + 1$ with $\alpha < \kappa$, because κ is a limit ordinal, $\alpha + 1 < \kappa$ so y has rank $< \kappa$.

For y of rank γ a limit, $\gamma = \sup\{\operatorname{rank}(z) + 1 : z \in y\}$. Note that $|y| < \kappa$, and each $\operatorname{rank}(z) + 1 < \kappa$ for $z \in y$ by induction. In other words, we have a function from $|y| < \kappa$ to κ. This is then bounded in κ, because κ is regular. Hence γ, being at most this bound, is less than κ.

Thus each $y \in T$ has rank $< \kappa$, and in particular, $x \in T$ has rank $< \kappa$. Therefore $\mathrm{H}_\kappa \subseteq \mathrm{V}_\kappa$. ⊣

Let's now investigate how much set theory H_κ will satisfy. We of couse have the basics.

7 C • 7. Lemma

Let $\kappa \geq \aleph_0$ be a regular cardinal. Therefore $\mathbf{H}_\kappa = \langle \mathrm{H}_\kappa, \in \rangle \vDash \mathsf{BST} + \mathsf{Rep}$.

Proof .:.

- For Pair, note that $\operatorname{trcl}(\{x, y\}) = \operatorname{trcl}(x) \cup \operatorname{trcl}(y) \cup \{x, y\}$ by (4) of Result 4 A • 5. So if $x, y \in \mathrm{H}_\kappa$, then $|\operatorname{trcl}(\{x, y\})| < \kappa + \kappa + \kappa = \kappa$ and therefore $\{x, y\} \in \mathrm{H}_\kappa$.
- For Union, $\operatorname{trcl}(\bigcup x) \subseteq \operatorname{trcl}(x)$ so if $x \in \mathrm{H}_\kappa$, then $\operatorname{trcl}(\bigcup x)$ has size $\leq |\operatorname{trcl}(x)| < \kappa$ and thus $\bigcup x \in \mathrm{H}_\kappa$.
- For Comp, since $y \subseteq x$ implies $\operatorname{trcl}(y) \subseteq \operatorname{trcl}(x)$, it follows that $x \in \mathrm{H}_\kappa$ impies $y \in \mathrm{H}_\kappa$. Hence $\mathcal{P}(x) \subseteq \mathrm{H}_\kappa$. In particular, all definable subsets x are in H_κ, and thus $\mathbf{H}_\kappa \vDash \mathsf{Comp}$.
- For replacement, we argue as with $\mathbf{V}_\omega$. Suppose φ defines a function from $D \in \mathrm{H}_\kappa$. We wish to show that the image of φ under D, being R, is in H_κ. As each $r \in R$ has $|\operatorname{trcl}(r)| < \kappa$ and $|R| \leq |D| < \kappa$, it follows that $\operatorname{trcl}(R) = R \cup \bigcup_{r \in R} \operatorname{trcl}(r)$ has size $< \kappa$, being the union of $< \kappa$-many sets each of size $< \kappa$. Therefore $R \in \mathrm{H}_\kappa$ so that $\mathbf{H}_\kappa \vDash \mathsf{Rep}$.
- The existence of cartesian products follows from replacement. ⊣

More than just basic set theory, we get all of the axioms, except perhaps for powerset. The issue with powerset is Cantor's Theorem (5 B • 13): the powerset will have a higher cardinality. For example, $\mathrm{H}_{\aleph_1}$, the hereditarily countable sets, will contain ω, but $\mathcal{P}(\omega) \cap \mathrm{H}_{\aleph_1} = \mathcal{P}(\omega)$ will not be a set in $\mathrm{H}_{\aleph_1}$ because it will be too large: $|\mathcal{P}(\omega)| \geq \aleph_1$.

Now as we've seen, $\mathbf{H}_{\aleph_0} = \mathbf{V}_\omega \vDash \mathsf{ZFC} - \mathsf{Inf}$. With uncountable κ, however, we gain Inf at the expense of P.

7 C • 8. Theorem

Let $\kappa > \aleph_0$ be a regular cardinal. Therefore $\mathbf{H}_\kappa \vDash \mathsf{ZFC} - \mathsf{P}$.

Proof .:.

As $\operatorname{trcl}(\omega) = \omega < \kappa$, $\omega \in \mathrm{H}_\kappa$ so that $\mathbf{H}_\kappa \vDash \mathsf{Inf}$. For AC, for any non-empty family $F \in \mathrm{H}_\kappa$ of non-empty, disjoint sets, a choice set $C \in \mathrm{V}$ has $C \in \mathcal{P}(\bigcup F) \subseteq \mathrm{H}_\kappa$ so that $C \in \mathrm{H}_\kappa$. The rest follow from Lemma 7 C • 7. ⊣

As a result, if a regular cardinal $\kappa > \aleph_0$ has $\mathrm{H}_\kappa = \mathrm{V}_\kappa$, then $\mathbf{V}_\kappa \vDash \mathsf{ZFC}$. So the existence of such κ cannot be proven to exist just within ZFC. Such axioms stating the existence of such κ are effectively stronger axioms of infinity, since $\mathsf{ZFC} - \mathsf{Inf}$ cannot prove the existence of a model of $\mathsf{ZFC} - \mathsf{Inf}$ although ZFC can. The analogy being that $\mathsf{ZFC} + \mathsf{LC}$ (LC standing for "large cardinals") can prove the existence of a model of ZFC although ZFC can.

7 C • 9. Definition

A cardinal κ is *weakly inaccessible* iff κ is regular and a limit cardinal: $\lambda < \kappa$ implies $\lambda^+ < \kappa$.
A cardinal κ is *strongly inaccessible* or just *inaccessible* iff κ is regular and a *strong limit* cardinal: $\lambda < \kappa$ implies $2^\lambda < \kappa$.

Note that being weakly or strongly inaccessible is downward absolute between models of ZF: being regular is downward absolute, and being a limit cardinal is equivalent to $\{\alpha < \kappa : |\alpha| = \alpha\}$ being unbounded in κ, which is clearly downward absolute. Similarly, being a strong limit is downward absolute between models of ZFC.

7 C • 10. Corollary

Let κ be strongly inaccessible. Therefore $\mathrm{V}_\kappa = \mathrm{H}_\kappa$ and $\mathbf{V}_\kappa \vDash \mathsf{ZFC}$.

Proof .:.

We already know that $\mathrm{H}_\kappa \subseteq \mathrm{V}_\kappa$ since κ is regular. So it suffices to show that $\mathrm{V}_\kappa \subseteq \mathrm{H}_\kappa$. Firstly, note that $|\mathrm{V}_\alpha| < \kappa$ for $\alpha < \kappa$. This is obvious for $\alpha = 0$. For $\alpha < \kappa$, inductively $|\mathrm{V}_\alpha| = \lambda < \kappa$ implies $|\mathrm{V}_{\alpha+1}| = 2^\lambda < \kappa$ as κ is strongly inaccessible. Similarly, if $\gamma < \kappa$ is a limit, $|\mathrm{V}_\gamma| = \sup_{\alpha<\gamma} |\mathrm{V}_\alpha|$. Since $\gamma < \kappa$ and inductively each $|\mathrm{V}_\alpha| < \kappa$ for $\alpha < \gamma$, this must have size $|\mathrm{V}_\gamma| < \kappa$ since κ is regular.

But for any $x \in \mathrm{V}_{\alpha+1}$ with $\alpha < \kappa$, it follows that $\mathrm{trcl}(x) \in \mathrm{V}_{\alpha+1}$ and thus $|\mathrm{trcl}(x)| \le |\mathrm{V}_{\alpha+1}| < \kappa$ and therefore $x \in \mathrm{H}_\kappa$. By Theorem 7 C • 8, all axioms of ZFC except possibly P are satisfied by $\mathbf{V}_\kappa = \mathbf{H}_\kappa$. By Result 7 C • 2, P is satisfied too, and thus all axioms of ZFC. $\dashv$

Weakly inaccessible cardinals will have their uses later: $\mathbf{L}_\kappa \vDash \mathsf{ZFC}$ for weakly inaccessible κ, for example.

§ 7 D. Reflection theorems

The levels of V are able to capture a lot of information about V itself. This idea generalizes to other classes $\mathrm{M} \subseteq \mathrm{V}$ that have a similar construction as the cumulative hierarchy.

7 D • 1. Definition

A transitive class M is *stratified* iff there is a (class) function mapping $\alpha \in \mathrm{Ord}$ to $\mathrm{M}_\alpha \in \mathrm{V}$ such that
- $\mathrm{M} = \bigcup_{\alpha\in\mathrm{Ord}} \mathrm{M}_\alpha$ and $\mathrm{M}_\gamma = \bigcup_{\alpha<\gamma} \mathrm{M}$ for limit γ;
- $\alpha \le \beta$ implies $\mathrm{M}_\alpha \subseteq \mathrm{M}_\beta$;
- $\mathrm{M}_\alpha \in \mathrm{M}$ for each $\alpha \in \mathrm{Ord}$;
- Each M_α is transitive.

Note that, for example V is stratified as witnessed by the cumulative hierarchy. Being stratified entails that **M** satisfies some weakenings of the axioms of ZFC. In particular, **M** might not satisfy comprehension. So whereas the following axioms are equivalent to the usual axioms under the theory {Comp, Ext}, they are weaker in its absence.

7 D • 2. Definition

- (wPair) $\{x, y\} \subseteq z$ for some z: $\forall x\ \forall y\ \exists z\ (x \in z \wedge y \in z)$.
- (wUnion) $\bigcup F \subseteq z$ for some z: $\forall F\ \exists U\ \forall v\ (\exists x(x \in F \wedge v \in x) \to v \in U)$.
- (wP) for each x, $\wp(x) \subseteq z$ for some z: $\forall x\ \exists P\ \forall v\ (v \subseteq x \to v \in P)$.
- (wRep) the image of a function over a set is contained a set: for each $\mathtt{FOL}(\in)$-formula φ,
$$\forall w_0 \cdots \forall w_n\ \forall D\ \left(\forall x \in D\ \exists! y\ \varphi(x, y, \vec{w})\right) \to \exists R\ \forall x \in D\ \exists y \in R\ \varphi(x, y, \vec{w})\big).$$

Writing wZFC for ZFC (and similarly wZF for ZF) replacing axioms with these weak versions, we have the following. Note that wZFC is equivalent to ZFC in the sense that wZFC $\vdash \varphi$ for each $\varphi \in$ ZFC and vice versa. But if we remove axioms like comprehension—as we do below—then the resulting theories are not equivalent.

7 D • 3. Result

If M is stratified, then $\mathbf{M} \vDash \mathsf{wZF} - \mathsf{Comp} - \mathsf{Inf}$

Proof .:.

Extensionality, empty set, and foundation all hold by virtue of M being stratified.
- For wPair, if $x, y \in \mathrm{M}$, then $x \in \mathrm{M}_\alpha$ and $y \in \mathrm{M}_\beta$ for some ordinals $\alpha, \beta \in \mathrm{Ord}$. Therefore, as $\mathrm{M}_\alpha, \mathrm{M}_\beta \subseteq \mathrm{M}_{\max(\alpha,\beta)}$, we have $x, y \in \mathrm{M}_{\max(\alpha,\beta)} \in \mathrm{M}$ witnessing the axiom.

- For wUnion, if $x \in \mathrm{M}$, then $x \in \mathrm{M}_\alpha$ for some $\alpha \in \mathrm{Ord}$ and as a transitive set with $x \subseteq \mathrm{M}_\alpha$, $\bigcup x \subseteq \mathrm{trcl}(x) \subseteq \mathrm{M}_\alpha$ and therefore $\mathrm{M}_\alpha \in \mathrm{M}$ witnesses the axiom.
- For wP, suppose $x \in \mathrm{M}$. For each $y \in \mathcal{P}(x) \cap \mathrm{M}$, let α_y be the least α with $y \in \mathrm{M}_\alpha$. In $\mathbf{V}$, we can thus consider the supremum $\beta = \sup_{y \in \mathcal{P}(x) \cap \mathrm{M}} \alpha_y$. Hence $\mathcal{P}(x) \cap \mathrm{M} \subseteq \mathrm{M}_\beta \in \mathrm{M}$ witnesses the axiom.
- For wRep, suppose φ defines a function in $\mathbf{M}$ on some $D \in \mathrm{M}$. For $x \in D$, let α_x be the least α with $y \in \mathrm{M}_\alpha$, where y is the output of x. By replacement in $\mathbf{V}$, the supremum $\beta = \sup_{x \in D} \alpha_x < \mathrm{Ord}$. But then M_β contains the pointwise output of D. Therefore $\mathrm{M}_\beta \in \mathrm{M}$ witnesses the axiom. ⊣

For any stratified M, we get that the levels of M reflect the truth of M itself. To show this, we need some restricted versions of Tarski–Vaught Theorem (6 A • 6).

7 D • 4. Lemma

Let φ be a FOLp-formula and $M_0 \subseteq M$ be non-empty, transitive classes. Therefore the following are equivalent:

1. φ and all of its subformulas are absolute between $\mathbf{M}_0$ and $\mathbf{M}$.
2. for each subformula of φ (possibly including φ itself) of the form "$\exists y\ \psi(\vec{x}, y)$", for all $\vec{x} \in M^{<\omega}$,
$$\exists y \in M\ \psi^M(\vec{x}, y) \to \exists y \in M_0\ \psi^M(\vec{x}, y).$$

Proof .:.

The same proof for Tarski–Vaught Theorem (6 A • 6) applies to show that (1) implies (2).

So assume (2) holds. Proceeding by induction on subformulas, given a subformula ψ of φ, we can assume each proper subformula of ψ is absolute between $\mathbf{M}_0$ and $\mathbf{M}$. If ψ is atomic or of the form "$\chi \wedge \theta$" or "$\neg\theta$" then clearly ψ is absolute between $\mathbf{M}_0$ and $\mathbf{M}$.

So consider the subformula "$\exists y\ \psi$". Note that then ψ^M iff ψ^{M_0} by the inductive hypothesis. Therefore, $\exists y \in M_0\ \psi^{M_0}$ implies $\exists y \in M_0\ \psi^M$ and therefore $\exists y \in M\ \psi^M$. By (2), the reverse implications hold. So we know for any $\vec{m}$ in M,
$$\mathbf{M} \vDash \text{“}\exists y\ \psi(\vec{m}, y)\text{”} \quad \text{iff} \quad \exists y \in M\ \psi^M(\vec{m}, y) \quad \text{iff} \quad \exists y \in M_0\ \psi^{M_0}(\vec{m}, y)\ ⊣$$

This allows us to perform an induction on formulas so that when we close under the property of (2), we get absoluteness.

7 D • 5. Theorem (The Reflection Principle)

Let M be stratified, and let φ be a FOLp-formula. Therefore, there are arbitrarily large $\alpha \in \mathrm{Ord}$ where φ is absolute between $\mathbf{M}$ and $\mathbf{M}_\alpha$.

Proof .:.

Proceed by induction on φ to show the variant result that φ and all of its subformulas are absolute between $\mathbf{M}$ and $\mathbf{M}_\alpha$ for arbitrarily large $\alpha \in \mathrm{Ord}$. For φ being "$x = y$" or "$x \in y$", this is obvious, as they are absolute between all transitive models, which $\mathbf{M}$ and $\mathbf{M}_\alpha$ are.

Let $\beta \in \mathrm{Ord}$ be arbitrary. For each subformula of φ of the form "$\exists y\ \psi$", we will show there is an $\alpha > \beta$ where
$$\mathbf{M} \vDash \text{“}\forall \vec{x}\ \left(\exists y\ \psi(\vec{x}, y) \to \exists y \in M_\alpha\ \psi(\vec{x}, y)\right)\text{”}$$
and thus by (2) of Lemma 7 D • 4, conclude that φ and its subformulas are absolute between $\mathbf{M}$ and $\mathbf{M}_\alpha$.

For each subformula "$\exists y\ \psi$" and $\vec{x} \in M$, let $F_\psi(\vec{x})$ be the least ordinal $\alpha \in \mathrm{Ord}$ such that $\exists y \in \mathrm{M}_\alpha\ \psi^{\mathrm{M}}(\vec{x}, y)$ (if there is no such ordinal, set $\alpha = 0$). Such an ordinal $\neq 0$ will exist if $\mathbf{M} \vDash$ "$\exists y\ \psi(\vec{x}, y)$", since M is stratified. So F_ψ points to a level where there is a witness to ψ.

For $\alpha \in \mathrm{Ord}$, consider
$$G(\alpha) = \sup\{F_\psi(\vec{x}) : \vec{x} \in \mathrm{M}_\alpha^{<\omega} \wedge \text{“}\exists y\ \psi\text{”is one of the existential subformulas of } \varphi\}.$$
This means that in $\mathbf{M}$, every input in M_γ has its witness to ψ^{M} somewhere in $\mathrm{M}_{G(\gamma)}$. So now we just continually apply G and then union up to get a model closed under this.

Take $\alpha_0 > \beta$ arbitrary. Let $\alpha_{n+1} = \max(G(\alpha_n), \alpha_n + 1)$ and set $\alpha = \sup_{n<\omega} \alpha_n$. Clearly α is a limit ordinal and hence $\mathrm{M}_\alpha = \bigcup_{n<\omega} \mathrm{M}_{\alpha_n}$. But then any $\vec{x} \in \mathrm{M}_\alpha^{<\omega}$ has $\vec{x} \in \mathrm{M}_{\alpha_n}^{<\omega}$ for some $n < \omega$ and therefore if there is a y where $\psi^{\mathrm{M}}(\vec{x}, y)$, there is a y in $\mathrm{M}_{\alpha_{n+1}}$. Therefore M_α and M satisfy (2) of Lemma 7 D • 4. Hence φ and all of its subformulas are absolute between $\mathbf{M}_\alpha$ and $\mathbf{M}$, and $\alpha > \beta$. ⊣

An alternative proof of The Reflection Principle (7 D • 5) can be given by more combinatorial means[xxiii], but this is not done here, since the relevant concepts will not be introduced until Chapter II. Note that The Reflection Principle (7 D • 5) is equivalent to the result holding for finitely many formulas φ, as we can just take the single formula which is the conjunction of the finitely many. In particular, we have the following.

7 D • 6. Corollary

ZFC is not finitely axiomatizable: there is no finite set of FOL(∈)-formulas T such that $T \vdash \varphi$ iff ZFC $\vdash \varphi$.

Proof .:.

For each model $\mathbf{M} \vDash$ ZFC, the hierarchy $\mathrm{V}_\alpha^{\mathbf{M}}$ witnesses that M is stratefied in $\mathbf{M}$. If there were such a finite collection, the conjunction of these finitely many formulas is a formula φ. By The Reflection Principle (7 D • 5) in $\mathbf{M}$, since $\mathbf{M} \vDash \varphi$, there is some $\mathbf{V}_\alpha^{\mathbf{M}} \vDash \varphi$ and therefore $\mathbf{V}_\alpha^{\mathbf{M}} \vDash$ ZFC. Consider the (according to $\mathbf{M}$) least $\alpha \in \mathrm{Ord}^{\mathbf{M}}$ where $\mathbf{V}_\alpha^{\mathbf{M}} \vDash \varphi$. By the same argument above, by the absoluteness of rank and thus the V_αs between transitive submodels of $\mathbf{M}$, there is some $\beta \in \mathrm{V}_\alpha^{\mathbf{M}}$ where $\mathbf{V}_\beta^{\mathbf{V}_\alpha} = \mathbf{V}_\beta^{\mathbf{M}} \vDash \varphi$, contradicting the minimality of α in $\mathbf{M}$. ⊣

The above corollary highlights an important idea regarding the relativity of transitivity. In principle, everything we've done thus far has been in an arbitrary model of ZFC, and so the notions of "transitive", "well-founded", and so forth are notions relative to this background model. In particular, for $\mathbf{M} \vDash$ ZFC,

- A model $\mathbf{N}$ is transitive in $\mathbf{M}$ iff $\mathrm{trcl}^{\mathbf{M}}(N) = N$.
- A relation R is well-founded in $\mathbf{M}$ iff there is no $m \in M$ with $\mathbf{M} \vDash$ "$\forall x \in m \; \exists y \in m \; (y \; R \; x)$".

As we've seen, these can differ between different models of set theory. But the same absoluteness results above hold; it's just that they are restricted to transitive models *of our given model* rather than the more philosophically based notion of $\mathbf{V}$.

Corollary 7 D • 6 also highlights an important distinction between a theorem and a theorem scheme. The Reflection Principle (7 D • 5) is a theorem scheme in that for each φ, we get a different theorem. The Reflection Principle (7 D • 5) is not equivalent to any single formula by the same sort of reasoning as in Corollary 7 D • 6. That said, we're still effectively working in an arbitrary model of ZFC, and so a coded version of The Reflection Principle (7 D • 5) still holds in an arbitrary model of ZFC, it's just that the coded notion of "formula" etc. in a non-standard model may not agree with the actual universe, just like well-foundedness.

In particular, if $\mathbf{M}$ has $\langle \omega^{\mathbf{M}}, \in^{\mathbf{M}} \rangle \ncong \langle \omega, \in \rangle$, then $\mathbf{M}$ will have an $n \in \omega^{\mathbf{M}}$ that $\mathbf{M}$ thinks is a coded formula, but doesn't correspond to any real-world formula, because n isn't even an actual natural number. For a more concrete example, ZFC $\nvdash$ Con(ZFC) implies there are models of ZFC where $\neg$Con(ZFC). In such a model $\mathbf{M}$, the coded proof that ZFC $\vdash$ "$\varphi \wedge \neg\varphi$" corresponds to one of these "natural numbers" of $\mathbf{M}$, and not the code of an actual proof.

We can actually get a slightly stronger reflection theorem. Note that for $\kappa = \mathrm{Ord}$—and thus writing M_κ for M—this is the same as The Reflection Principle (7 D • 5).

7 D • 7. Theorem (The Reflection Theorem)

Let M be stratified. Let $\kappa > \aleph_0$ be a regular cardinal (allowing for $\kappa = \mathrm{Ord}$). Let φ be a FOLp-formula. Therefore, there are arbitrarily large $\alpha < \kappa$ where φ is absolute between $\mathbf{M}_\kappa$ and $\mathbf{M}_\alpha$.

[xxiii] In particular, each formula is absolute between $\mathbf{M}_\alpha$ and $\mathbf{M}$ on a *club* of $\alpha \in \mathrm{Ord}$. Since the intersection of two clubs is a club, the propositional connectives are dealt with easily in the induction on formulas. The existential case "$\exists x \; \psi$" can be dealt with by considering the map sending parameters to the least β with a witness in M_β. Taking the supremum of such βs and then closing the club for ψ under this yields another club that gets the job done, similar to the approach taken above, but avoiding Lemma 7 D • 4 at the cost of giving a background on clubs.

Proof .:.

Proceed by induction on φ. Let $\beta \in \mathrm{Ord}$ be arbitrary. As before, we only need to deal with the existential subformulas of φ. For each subformula of φ of the form "$\exists y\ \psi$", we will show there is an α with $\beta < \alpha < \kappa$ where

$$\text{“}\forall \vec{x} \in \mathrm{M}_\kappa\ \left(\exists y \in \mathrm{M}_\kappa\ \psi^{\mathrm{M}_\kappa}(\vec{x}, y) \to \exists y \in M_\alpha\ \psi(\vec{x}, y)\right)\text{”}.$$

For $\vec{x} \in M$, let $F_\psi(\vec{x})$ be the least ordinal $\alpha < \kappa$ such that $\exists y \in \mathrm{M}_\alpha\ \psi^{\mathrm{M}_\kappa}(\vec{x}, y)$ (if there is no such ordinal, set $\alpha = 0$). So F_ψ points to a level where there is a witness to ψ (if there is one). For $\alpha < \kappa$, consider

$$G(\alpha) = \sup\{F_\psi(\vec{x}) : \vec{x} \in \mathrm{M}_\alpha^{<\omega} \wedge \text{“}\exists y\ \psi\text{”is one of the existential subformulas of } \varphi\}.$$

As before, we just continually apply G and then union up to get a model closed under this.

Take α_0 to be arbitrary such that $\beta < \alpha_0 < \kappa$. Let $\alpha_{n+1} = \max(G(\alpha_n), \alpha_n + 1)$ and set $\alpha = \sup_{n<\omega} \alpha_n$. As $\mathrm{cof}(\kappa) > \omega$ and inductively each $\alpha_n < \kappa$, it follows that $\alpha < \kappa$. Consider $\mathrm{M}_\alpha = \bigcup_{n<\omega} \mathrm{M}_{\alpha_n} \subseteq \mathrm{M}_\kappa$. As before, M_α and M_κ satisfy (2) of Lemma 7 D • 4. Hence φ and all of its subformulas are absolute between M_α and M_κ, and $\kappa > \alpha > \beta$. ⊣

The point of this will be to have the ability to take skolem hulls of the levels of a stratified model, and end up with smaller models of the same statements. With The Reflection Principle (7 D • 5), we can't take a skolem hull of M and expect it to be in the model of set theory, since M is a proper class, and not a set. But M_κ for $\kappa \in \mathrm{Ord}$ *is* a set, and so we can take the skolem hull.

Section 8. *The First Inner Models*

We begin with the definition of an inner model. The general picture of an inner model is just a "skinny" version of the background model V, where the class of all ordinals constitutes the "backbone" of the universe.

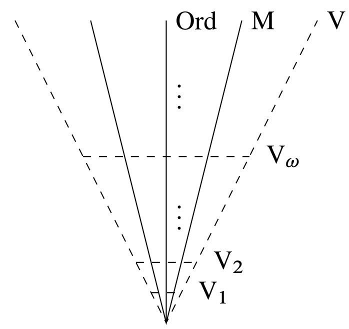

8 • 1. Figure: An inner model

8 • 2. Definition

A class $M \subseteq V$ is an *inner model* iff
- M is transitive;
- $\mathrm{Ord} \subseteq M$; and
- $\mathbf{M} = \langle M, \in \rangle \models \mathsf{ZFC}$

If we replace ZFC in the last condition with some theory T, we say M is an inner model *of* T.

So clearly V is an inner model. Moreover, as we've defined things, for any model $\mathbf{W} \models \mathsf{ZFC} - \mathsf{Found}$, $\mathrm{WF}^{\mathbf{W}} \subseteq W$ is an inner model of ZFC. Of course, in $\mathbf{V}$, both of these are just V. So these examples are not particularly illuminating for us. The goal of this section is to introduce two more inner models: one of which is very rigid, and one of which is very flexible.

Note that being an inner model is a scheme, and not a singular formula. It's saying that $\mathrm{Ord} \subseteq M$ is transitive (a single sentence) and that $\mathbf{M} \models \varphi$ for each $\varphi \in \mathsf{ZFC}$ (infinitely many sentences).

§ 8 A. The constructible universe and definability

Recall that the levels of V were defined by iteratively taking the powerset operation. Gödel's definition of the constructible universe, L, does the same, but restricts to subsets which are definable over the previous levels.

8 A • 1. Definition

Let $L = \bigcup_{\alpha<\mathrm{Ord}} L_\alpha$ where $\alpha \mapsto L_\alpha$ is defined by transfinite recursion: $L_0 = \emptyset$, $L_\gamma = \bigcup_{\alpha<\gamma} L_\alpha$ for γ a limit, and

$$L_{\alpha+1} = \{x \in \mathcal{P}(L_\alpha) : x \text{ is definable over } \langle L_\alpha, \in \rangle\}.$$

Here, x being definable over $\mathbf{L}_\alpha$ means that there is a $\mathtt{FOLp}$-formula $\varphi(y)$ where $\mathbf{L}_\alpha \models$ "$\varphi(y)$" iff $y \in x$.

The importance of L to set theory is hard to overstate. There are three main ideas why. Firstly, every transitive model of enough set theory has an interpretation of L, and this interpretation is the same across all transitive models of $\mathsf{ZF} - \mathsf{P}$. Secondly, it's the *only* model with this property, demonstrating a strong minimality condition. In fact, it's defining formula is so rigid that any transitive model elementarily equivalent to one of the L_α levels is actually one of the L_α levels. Thirdly, as a result of all of this, L is always the smallest inner model. And thus it can be seen as the transitive model "generated" by the theory of ZFC in that it is the smallest such model. This is analogous to the situation with arithmetic, where $\mathbb{N}$ is the smallest model of the peano axioms, and so can be thought of as being generated by them.

To confirm all of this, we begin with showing that $L \models \mathsf{ZFC}$. First we will show that L is stratefied, which gives a great portion of set theory. Note that $L_\alpha \in L_{\alpha+1}$ as witnessed by the formula "$x = x$". Furthermore, because we're allowing parameters, $\mathrm{L}_\alpha \subseteq \mathrm{L}_\beta$ for $\alpha \le \beta$. One might think the only thing needed to confirm that L is stratified is that each L_α is transitive. But in fact, we need to ensure that the function taking α to L_α is definable: that L is a class.

To do this, we need to understand how to formalize definability within set theory. We know from Theorem 6 B • 6 that we can do a lot of the work based on this in ZF, since we can look at the full powerset, and then restrict to those subsets which have a first-order definition as per Theorem 6 B • 6. But to help us later, it will be useful to work in ZF − P, which requires instead a reliance on the replacement axiom. So rather than rely on Theorem 6 B • 6, we will instead think of closing a given set under operations corresponding to the logical operations.

8 A • 2. Definition

Let A be a set. For $n, m, k \in \omega$, define
- $\mathrm{Exists}^n_A(D) = \{\tau \in A^n : \exists x \in A\ (\tau^\frown x \in D)\}$;
- $\mathrm{Memb}^{n,m,k}_A = \{\tau \in A^k : n, m \in \mathrm{dom}(\tau) \land \tau(n) \in \tau(m)\}$;
- $\mathrm{Equal}^{n,m,k}_A = \{\tau \in A^k : n, m \in \mathrm{dom}(\tau) \land \tau(n) = \tau(m)\}$

Define $\mathrm{FOL}(A)$ to be the closure of $\{\mathrm{Memb}^{n,m}_A, \mathrm{Equal}^{n,m}_A : n, m \in \omega\}$ under Exists^n_A, intersections, and complements in A^n for $n < \omega$.

Note that $\mathrm{Exists}^n_A(D)$ corresponds to the existential quantificatier while $\mathrm{Memb}^{n,m,k}_A$ and $\mathrm{Equal}^{n,m,k}_A$ correspond to membership and equality. Similarly, intersections correspond to conjunction. Complements correspond to negations. Hence starting with the atomic formulas and closing under existential quantification, conjunction, and relative complement, we get all of the first-order formulas, and this corresponds precisely to looking at their defined sets, closing under these operations. Hence $\mathrm{FOL}(A)$ corresponds to the `FOL`-defined subsets of $A^{<\omega}$.

Note that we can define $\mathrm{FOL}(A)$ by finitary recursion, just repeatedly applying the operations to the sets in the previous stage, starting with the first stage of $\{\mathrm{Memb}^{n,m}_A, \mathrm{Equal}^{n,m}_A : n, m \in \omega\}$. Thus without powerset, $\mathrm{FOL}(A)$ exists.

8 A • 3. Definition

Let A be a set. For $\sigma \in A^n$ where $n < \omega$, define $\mathrm{FOL}_\sigma(A)$ to be the closure of $\mathrm{FOL}(A)$ under the operations of Definition 8 A • 2 and the operation

$$\mathrm{Param}^\sigma_A(D) = \{\tau \in A^n : \sigma^\frown\tau \in D\}.$$

Define $\mathrm{FOLp}(A)$ to be $\bigcup_{\sigma \in A^{<\omega}} \mathrm{FOL}_\sigma(A)$, the set of all subsets of $A^{<\omega}$ that are `FOLp`-definable.[xxiv]

As before, in ZF − P, $\mathrm{FOLp}(A)$ exists.

8 A • 4. Corollary

The function $\alpha \mapsto \mathrm{L}_\alpha$ is definable, and hence L is a class: $x \in \mathrm{L}$ iff $\exists\alpha \in \mathrm{Ord}(x \in \mathrm{L}_\alpha)$.

Proof .:.

Using Definition 8 A • 3 we can talk about which sets are `FOLp`-definable over L_α. So we can define recursively $x = \mathrm{L}_\alpha$ iff there exists a function L with $\mathrm{dom}(L) = \alpha + 1$ and $L(\alpha) = x$ such that
- $L(0) = \emptyset$;
- for every $\beta < \alpha$, $L(\beta + 1) = \{y \in \mathrm{FOLp}(L(\beta)) : y \subseteq \mathrm{L}_\beta\}$.
- for every limit ordinal $\gamma \le \alpha$, $L(\gamma) = \bigcup_{\beta<\gamma} L(\beta)$. ⊣

Now we can show that L is stratified, and hence get a large portion of ZFC by Result 7 D • 3

8 A • 5. Lemma

For each $\alpha \in \mathrm{Ord}$, L_α is transitive, and hence L is stratified.

[xxiv] In principle, we should be a bit careful arguing about $\mathrm{FOLp}(A)$ as the `FOLp`-definable subsets of A, since $\mathrm{FOLp}(A)$ is a formally defined concept that may not mesh with the real world notions. For example, in model with a non-standard ω, it's not immediately obvious that all "finite" subsets of A are in $\mathrm{FOLp}(A)$, even though all the *actual* finite subsets are. The proof that all the subsets of size $< \omega^{\mathsf{M}}$ are in $\mathrm{FOLp}^{\mathsf{M}}(A)$ for any M is a bit tedious, going through the details of $\mathrm{FOLp}(A)$ being defined as a closure of certain operations. It is not particularly difficult, but it is not particularly enlightening, and doesn't serve to help understand the general ideas.

Proof .:.

Proceed by induction on α. For $\alpha = 0$, $L_0 = \emptyset$ is obviously transitive. For $\alpha + 1$, each $x \in L_{\alpha+1}$ is a subset of L_α. Clearly each $a \in L_\alpha$ is in $L_{\alpha+1}$ defined by the FOLp-formula "$y \in a$": $a = \{y \in L_\alpha : y \in a\} \in \mathrm{FOLp}(L_\alpha)$. Hence $x \subseteq L_\alpha \subseteq L_{\alpha+1}$. The result for limits holds clearly by the inductive hypothesis. $\dashv$

Thus by Result 7 D • 3, $\mathbf{L} \vDash \mathsf{wZF} - \mathsf{Comp} - \mathsf{Inf}$. To confirm that $\mathbf{L} \vDash \mathsf{Inf}$, we will show $\mathrm{Ord} \subseteq L$ and thus $\omega \in L$ showing Inf holds by Corollary 7 A • 5.

8 A • 6. Lemma

Ord $\subseteq$ L so that $\mathbf{L} \vDash \mathsf{Inf}$.

Proof .:.

We show by induction that $\alpha \in L_{\alpha+1} \setminus L_\alpha$. For $\alpha = 0$, $L_0 = \emptyset$. Since $L_\alpha \in L_{\alpha+1}$ for each α, it follows that $0 \in L_{0+1} = L_1 \setminus L_0$. Inductively, $\alpha \in L_{\alpha+1} \setminus L_\alpha$. Therefore $\alpha + 1 = \alpha \cup \{\alpha\} \subseteq L_{\alpha+1}$ is FOLp-definable over $\mathbf{L}_{\alpha+1}$ by "$x < \alpha \vee x = \alpha$", showing $\alpha + 1 \in L_{\alpha+2}$. Since $\alpha \notin L_{\alpha+1}$, $\alpha + 1 \notin L_{\alpha+1}$ so that $\alpha + 1 \in L_{\alpha+2} \setminus L_{\alpha+1}$. This deals with the successor case.

For the limit case, the inductive hypothesis tells us that $\alpha \subseteq L_\alpha$ and that α is the least ordinal not in L_α ($\alpha \notin L_\beta$ for any $\beta < \alpha$ as otherwise this would imply by transitivity of L_β that $\beta \in L_\beta$). Therefore α is definable over $\mathbf{L}_\alpha$ by "x is an ordinal". Hence $\alpha \in L_{\alpha+1}$. Hence each ordinal α is in $L_{\alpha+1} \subseteq L$ so that Ord $\subseteq$ L. $\dashv$

Hence we only need to show that comprehension and choice hold in $\mathbf{L}$. To do this, we use Corollary 7 A • 5.

8 A • 7. Theorem

$\mathbf{L} \vDash \mathsf{Comp}$ and therefore L is an inner model of ZF.

Proof .:.

Let φ be arbitrary. We want to show that for each $A \in L$, $A_\varphi = \{x \in A : \varphi^{L}(x)\} \in L$. To see this, note that $A \in L_\alpha$ for some $\alpha \in \mathrm{Ord}$. Note that there are arbitrarily large $\beta \in \mathrm{Ord}$ where φ is absolute between $\mathbf{L}$ and $\mathbf{L}_\alpha$ by The Reflection Principle (7 D • 5). In particular, there is some L_β where $A \in L_\beta$, and $\forall x\ (\varphi^{L_\beta}(x) \leftrightarrow \varphi^{L}(x))$. As a result, φ defines A_φ over $\mathbf{L}_\beta$ and thus $A_\varphi \in L_{\beta+1} \subseteq L$. Therefore, $\mathbf{L} \vDash \mathsf{Comp}$ so that $\mathbf{L} \vDash \mathsf{wZF}$. As wZF is equivalent to ZF, $\mathbf{L} \vDash \mathsf{ZF}$. $\dashv$

So all that remains is the axiom of choice. Note that all of the above work on L didn't use the axiom of choice. So if we were to start in a universe $\mathbf{W} \vDash \mathsf{ZF} + \neg\mathsf{AC}$, we would still have $\langle L^{\mathbf{W}}, \in^{\mathbf{W}}\rangle \vDash \mathsf{ZFC}$. The basic idea behind the proof is to well-order all of the sets in L according to the formulas that defined the sets.[xxv]

8 A • 8. Theorem

L is an inner model of ZFC.

Proof .:.

It suffices to show $\mathbf{L} \vDash \mathsf{AC}$. For each $x \in L$, write α_x for the least α where $x \in L_{\alpha_x+1}$. Well-order the (codes of the) FOL($\in$)-formulas with $\leqslant_{\text{lex}}$ as in Definition 6 B • 2. For each $x \in F$ let φ_x be the (code of the) $\leqslant_{\text{lex}}$-least formula which defines x over $\mathbf{L}_{\alpha_x}$ for some parameters $\vec{w}_x$ of L_{α_x}. Now define by recursion the order $<_{L_\alpha} \subseteq L_\alpha \times L_\alpha$ by taking $<_{L_0} = \emptyset$, and $<_{L_\gamma} = \bigcup_{\alpha<\gamma} <_\alpha$ for γ a limit, and $x <_{L_{\alpha+1}} y$ iff

1. $x, y \in L_\alpha$ and $x <_{L_\alpha} y$; or else
2. $\alpha_x < \alpha_y$; or else
3. "φ_x" $<_{\text{lex}}$ "φ_y"; or else
4. $\vec{w}_x <_{\text{lex}} \vec{w}_y$ under the $<_{L_\alpha}$-order.

It follows by induction and Lemma 6 B • 3 that each $<_{L_\alpha}$ is a well-order of L_α. In fact, $<_{L_\alpha} \in L_{\alpha+\omega}$, because the notions above are all easily definable.

[xxv] Again, formally, we would do this by ordering them by the lexicographically least sequence of operations that yield the element to be in L. We are merely thinking of this sequence of operations as a formula built up by the corresponding syntactic operations.

So let $F \in \mathrm{L}$ be a non-empty family of non-empty, disjoint sets. Note that $F \subseteq \mathrm{L}_\alpha$ for some α. Consider

$$C = \{y \in \bigcup F : \exists x \in F(y \text{ is the } <_{\mathrm{L}_\alpha}\text{-least element of } x)\}.$$

It follows that C is a choice set for F, and is in L. Thus $\mathbf{L} \vDash \mathsf{AC}$. ⊣

§ 8 B. L as a canonical inner model

As stated before, L has many "canonicity" properties. In particular, it has a strong minimality condition, being contained (up to a given height) in any transitive model of $\mathsf{ZF}-\mathsf{P}$. As a result, it's the smallest inner model, and is determined by its theory. We state these three facts as follows. Firstly, we have the absoluteness of L, leading to L being the smallest inner model.

8 B • 1. Theorem (Absoluteness of L)

For any transitive model $\mathbf{M} \vDash \mathsf{ZF}-\mathsf{P}$, writing $\mathrm{L}_{\mathrm{Ord}}$ for L;

1. For each $\alpha \in \mathrm{Ord} \cap \mathrm{M}$, $\mathrm{L}_\alpha \subseteq \mathrm{M}$.
2. $\mathrm{L}^{\mathbf{M}} = \mathrm{L}_{\mathrm{Ord}\cap\mathrm{M}}$.

In particular, $\mathrm{L}^{\mathbf{L}} = \mathrm{L}$. In fact, if $\mathrm{Ord} \subseteq \mathrm{M}$, then all of L is contained in M.

8 B • 2. Corollary (Smallest Inner Model)

$\mathrm{L} \subseteq \mathrm{M}$ for any inner model M of $\mathsf{ZF}-\mathsf{P}$.

Next, since we can write "$\mathrm{V} = \mathrm{L}$" as a $\mathtt{FOL}(\in)$-sentence, considering it as an axiom yields the following.

8 B • 3. Theorem (Condensation)

Suppose $\mathbf{M} \vDash \mathsf{ZF}-\mathsf{P}+$ "$\mathrm{V} = \mathrm{L}$" where M is transitive. Therefore $\mathrm{M} = \mathrm{L}_{\mathrm{Ord}\cap\mathrm{M}}$.

This theorem can be strengthened significantly, although we will prove stronger versions later. In particular, if $\mathbf{M} \preccurlyeq_{\Sigma_1} \mathbf{L}_\alpha$ for some $\alpha \in \mathrm{Ord}$ or $\alpha = \mathrm{Ord}$, then $\mathbf{M} \cong \mathbf{L}_\beta$ for some $\beta \leq \alpha$. Here "$\preccurlyeq_{\Sigma_1}$" refers to being an elementary substructure with respect to Σ_1-formulas.

To show the above results, we need to show the absoluteness of the construction of L. Firstly, note the following absoluteness result.

8 B • 4. Lemma

"$y = \mathrm{FOLp}(x)$" is absolute between transitive models of $\mathsf{ZF}-\mathsf{P}$.

Proof .:.

This follows since the closure of a set under these operations is given by recursion. Given that each of the operations is clearly absolute, it follows that the output of this is absolute by Theorem 7 B • 4. ⊣

Proof of Absoluteness of L *(8 B • 1)* .:.

Proceed by induction on α to show $\mathrm{L}^{\mathbf{M}}_\alpha = \mathrm{L}_\alpha$ and thus $\mathrm{L}_\alpha \subseteq \mathrm{M}$ for $\alpha \in \mathrm{Ord} \cap \mathrm{M}$. Clearly, for $\alpha = 0$, $\mathrm{L}_\alpha = \emptyset \in \mathrm{M}$. Similarly, by the absoluteness of unions and the inductive hypothesis, for limit $\gamma \in \mathrm{Ord} \cap \mathrm{M}$, $\mathrm{L}^{\mathbf{M}}_\gamma = \bigcup_{\alpha<\gamma} \mathrm{L}^{\mathbf{M}}_\alpha = \bigcup_{\alpha<\gamma} \mathrm{L}_\alpha = \mathrm{L}_\gamma$. For the successor stage $\alpha+1$, Lemma 8 B • 4 tells us that $\mathrm{L}^{\mathbf{M}}_{\alpha+1} = \mathrm{FOLp}^{\mathbf{M}}(\mathrm{L}^{\mathbf{M}}_\alpha) = \mathrm{FOLp}(\mathrm{L}_\alpha) = \mathrm{L}_{\alpha+1}$. Hence $\mathrm{L}_\alpha \in \mathrm{M}$ for each $\alpha \in \mathrm{Ord} \cap \mathrm{M}$.

We have $\mathrm{L}^{\mathbf{M}} = \bigcup_{\alpha\in\mathrm{Ord}\cap\mathrm{M}} \mathrm{L}^{\mathbf{M}}_\alpha = \bigcup_{\alpha\in\mathrm{Ord}\cap\mathrm{M}} \mathrm{L}_\alpha = \mathrm{L}_{\mathrm{Ord}\cap\mathrm{M}} \subseteq \mathrm{M}$. In particular, for M an inner model, $\mathrm{L}^{\mathbf{M}} = \mathrm{L} \subseteq \mathrm{M}$. ⊣

This shows the first two canonical properties of L: Absoluteness of L (8 B • 1) and Smallest Inner Model (8 B • 2). To show the third, Condensation (8 B • 3), we first should note that the sentence "$\mathrm{V} = \mathrm{L}$" does indeed exist, being defined through FOLp: for every x there is an ordinal α such that some function L defined on $\alpha+1$ has $x \in L(\alpha)$ and L satisfies the properties as laid out by Corollary 8 A • 4. More succinctly, "$\forall x\, \exists\alpha \in \mathrm{Ord}\, (x \in \mathrm{L}_\alpha)$". The proof of condensation is easy given these first two properties.

Proof of Condensation (8 B • 3) .:.

Since $\mathbf{M} \vDash$ "$\mathrm{V} = \mathrm{L}$", $\mathrm{M} = \mathrm{V}^{\mathbf{M}} = \mathrm{L}^{\mathbf{M}} = \mathrm{L}_{\mathrm{M} \cap \mathrm{Ord}}$ by Absoluteness of L (8 B • 1). ⊣

Being such a minimal model allows us to say more about absoluteness.

8 B • 5. Theorem

Suppose φ is upward absolute between inner models of ZF − P. Suppose $\mathrm{L} \vDash \varphi$. Therefore φ is absolute between inner models of ZF − P.

Proof .:.

Suppose $\mathbf{M} \vDash \mathsf{ZF} - \mathsf{P}$. Therefore $\mathrm{L}^{\mathbf{M}} = \mathrm{L}_{\mathrm{Ord} \cap \mathrm{M}} = \mathrm{L} \subseteq \mathrm{M}$ has $\mathbf{L} \vDash \varphi$. By upward absoluteness, $\mathbf{M} \vDash \varphi$. ⊣

If absoluteness is generally regarded as $\varphi \leftrightarrow \varphi^{\mathrm{M}}$ being true for all appropriate $\mathbf{M}$, the above tells us that this is equivalent to $\varphi^{\mathrm{L}} \leftrightarrow \varphi^{\mathrm{M}}$ for all appropriate $\mathbf{M} \vDash \mathsf{ZF} - \mathsf{P}$.

§ 8 C. Applications and properties of L

The importance of Condensation (8 B • 3) comes from its use with Taking a Skolem Hull (6 A • 2) in the form of Corollary 6 C • 2. In particular, since any skolem hull is elementarily equivalent to a level of L, when we collapse it, it becomes a level of L.

If we investigate further the levels of L, we get some quick examples of models of "$\mathrm{V} = \mathrm{L}$". Note that the levels of L, although defined similarly, develop differently to the levels of V. In particular, $\mathrm{V}_\alpha \neq \mathrm{L}_\alpha$ in general, even if we assume $\mathrm{V} = \mathrm{L}$. An easy example of this is that in ZFC, $\mathcal{P}(\omega) \subseteq \mathrm{V}_{\omega+1}$, meaning $|\mathrm{V}_{\omega+1}| \geq 2^{\aleph_0} > \aleph_0$. But $\mathrm{L}_{\omega+1}$ has only countably many new elements, corresponding to the defining formulas and parameters[xxvi]. Hence $|\mathrm{L}_{\omega+1}| = \aleph_0$. So the point is that subsets of ω don't appear in $\mathrm{L}_{\omega+1}$. In particular, one should not make the mistake of thinking $\mathrm{V}_\alpha^{\mathrm{L}} = \mathrm{L}_\alpha$. This is (almost always) false.

8 C • 1. Lemma

Let $\alpha \geq \omega$. Therefore $|\mathrm{L}_\alpha| = |\alpha|$.

Proof .:.

Proceed by induction on α. For $\alpha = \omega$, this is clear as L_ω is the countable union of sets, each of which is countable by induction: $\mathrm{L}_0 = \emptyset$ is clearly countable, and $\mathrm{L}_{n+1} \subseteq \mathcal{P}(\mathrm{L}_n)$ which is also finite for $n < \omega$. For $\alpha + 1$, $\mathrm{L}_{\alpha+1}$ is the closure of L_α under countably many operations and is thus $|\mathrm{L}_{\alpha+1}| \leq |\mathrm{L}_\alpha| \cdot \aleph_0$. Since clearly $\omega \subseteq \mathrm{L}_\alpha$ for $\alpha \geq \omega$ and $\mathrm{L}_\alpha \subseteq \mathrm{L}_{\alpha+1}$, it follows that the reverse inequality holds and in fact $|\mathrm{L}_{\alpha+1}| = |\mathrm{L}_\alpha| = |\alpha| = |\alpha + 1|$.

For limit γ, $|\mathrm{L}_\gamma| = \left|\bigcup_{\alpha<\gamma} \mathrm{L}_\alpha\right| \leq |\gamma| \cdot \sup_{\alpha<\gamma} |\mathrm{L}_\gamma| = |\gamma| \cdot \sup_{\alpha<\gamma} |\alpha| = |\gamma|$. ⊣

This allows us to more precisely understand what the levels of L look like.

8 C • 2. Result

Let $\kappa > \aleph_0$ be a regular cardinal. Therefore $\mathbf{L}_\kappa \vDash \mathsf{ZFC} - \mathsf{P} +$ "$\mathrm{V} = \mathrm{L}$".

Proof .:.

Once we show $\mathbf{L}_\kappa \vDash \mathsf{ZF} - \mathsf{P}$, by absoluteness, $\mathrm{L}^{\mathbf{L}_\kappa} = \mathrm{L}_\kappa = \mathrm{V}^{\mathbf{L}_\kappa}$ so that $\mathbf{L}_\kappa \vDash$ "$\mathrm{V} = \mathrm{L}$". So let $x, y \in \mathrm{L}_\kappa$ be arbitrary. Thus $x \in \mathrm{L}_\alpha$ and $y \in \mathrm{L}_\beta$ for some $\alpha, \beta < \kappa$.

- For Pair, assume without loss of generality that $\alpha < \beta$. Thus $x, y \in \mathrm{L}_\beta$ and so $\{x, y\} \in \mathrm{L}_{\beta+1}$.
- For Union, $\bigcup x \subseteq \mathrm{trcl}(x) \subseteq \mathrm{L}_\alpha$ and thus $\bigcup x \in \mathrm{L}_{\alpha+1}$ as it is easily definable.

[xxvi] More formally, it's the closure of L_ω under countably many operations, and hence adds only countably many elements.

- For Comp, use The Reflection Theorem (7 D • 7). In particular, the same proof as Theorem 8 A • 7 applies to show $\mathsf{L}_\kappa \vDash \mathsf{Comp}$: for each φ, there are arbitrarily large $\gamma < \kappa$ (e.g. $\gamma > \alpha$ where $x \in \mathrm{L}_\alpha$) where φ is absolute between L_γ and L_κ. Therefore in $\mathrm{L}_{\gamma+1}$, the set defined by comprehension, $\{z \in x : \varphi^{\mathrm{L}_\kappa}\} = \{z \in x : \varphi^{\mathrm{L}_\gamma}\} \in \mathrm{L}_{\gamma+1} \subseteq \mathrm{L}_\kappa$.
- For wRep, suppose $D \in \mathrm{L}_\kappa$, φ is a FOLp-formula, and that $\mathsf{L}_\kappa \vDash$ "$\forall x \in D\ \exists! y\ \varphi(x, y)$". We need to find an $R \in \mathrm{L}_\kappa$ such that every $x \in D$ has a $y \in R$ with $\varphi^{\mathrm{L}_\kappa}(x, y)$. Say $D \in \mathrm{L}_\alpha$ so that $|D| \le |\mathrm{L}_\alpha| = |\alpha| < \kappa$. Thus the L-ranks of the image of D should be bounded in L_κ. Explicitly, consider the function $f \in \mathrm{L}$ defined by $\varphi^{\mathrm{L}_\kappa}$ on D. Note that as $f : \mathrm{L}_\kappa \to \mathrm{L}_\kappa$, $\operatorname{im} f \subseteq \mathrm{L}_\kappa$. Let $\gamma = \sup\{\gamma_x + 1 : f(x) \in \mathrm{L}_{\gamma_x+1}\}$. Because κ is regular and $|D| \le |\mathrm{L}_\alpha| = |\alpha| < \kappa$, we have that $\gamma < \kappa$. Hence $\operatorname{im} f \subseteq \mathrm{L}_\gamma$ and thus L_γ witnesses the axiom of wRep for φ and D. Comprehension then gives Rep.
- For AC, the definition of $<_{\mathrm{L}_\kappa}$ in Theorem 8 A • 8 yields a definable well-order of all of L_κ. Hence for any non-empty family F of non-empty, disjoint sets in L_κ, $F \subseteq \mathrm{L}_\alpha$ for some $\alpha < \kappa$ so that the $<_{\mathrm{L}_\alpha}$-least (i.e. the $<_{\mathrm{L}_\kappa}$-least) element of each $x \in F$ yields a choice set just as in Theorem 8 A • 8. $\dashv$

8 C • 3. Corollary

If $\mathsf{L} \vDash$ "$\kappa > \omega$ is a cardinal", then $\mathsf{L}_\kappa \vDash \mathsf{ZFC} - \mathsf{P} +$ "$\mathrm{V} = \mathrm{L}$".

Proof .:.

If ZFC $\vdash$ "$\kappa > \aleph_0$ is regular $\to \varphi^{\mathrm{L}_\kappa}$" for each φ of ZF $-$ P $+$ "V $=$ L", then in particular, ZFC $+$ "V $=$ L" proves this. But then ZFC proves each relativization to L, i.e. if "$(\kappa > \aleph_0$ is regular$)^{\mathrm{L}}$" then $(\varphi^{\mathrm{L}_\kappa})^{\mathrm{L}}$ which is just $\varphi^{\mathrm{L}_\kappa}$.$\dashv$

One of the more important corollaries of Result 8 C • 2 and Condensation (8 B • 3) is what happens when we take skolem hulls.

8 C • 4. Corollary

Let $\kappa > \aleph_0$ be a regular cardinal. Let $X \subseteq \mathrm{L}_\kappa$. Therefore the collapsed skolem hull $\mathrm{cHull}^{\mathsf{L}_\kappa}(X) = \mathrm{L}_\alpha$ for some $\alpha < \kappa$. Moreover, if X is transitive, then $X \subseteq \mathrm{cHull}^{\mathsf{L}_\kappa}(X)$.

Proof .:.

By The Mostowski Collapse (4 • 1), the collapsed hull models "V $=$ L":

$$\mathbf{cHull}^{\mathsf{L}_\kappa}(X) \cong \mathbf{Hull}^{\mathsf{L}_\kappa}(X) \preccurlyeq \mathsf{L}_\kappa \vDash \mathsf{ZFC} - \mathsf{P} + \text{"V} = \text{L"}.$$

Hence by Condensation (8 B • 3), $\mathrm{cHull}^{\mathsf{L}_\kappa}(X) = \mathrm{L}_\alpha$ for some α. As $\mathrm{Hull}^{\mathsf{L}_\kappa}(X) \subseteq \mathrm{L}_\kappa$, α can be calculated as $\alpha = \mathrm{Ord} \cap \mathrm{cHull}^{\mathsf{L}_\kappa}(X) \le \kappa$.

Thus it suffices to show $X \subseteq \mathrm{cHull}^{\mathsf{L}_\kappa}(X)$ when X is transitive. To do this, we show that the collapsing map fixes X. Let $\pi : \mathrm{Hull}^{\mathsf{L}_\kappa}(X) \to \mathrm{cHull}^{\mathsf{L}_\kappa}(X)$ be the collapsing isomorphism, defined inductively by $\pi(x) = \{\pi(y) : y \in x\}$. We show that $\pi(x) = x$ for each $x \in X$. Suppose not. Let $x \in X$ be the $\in$-least element of X where $\pi(x) \ne x$. Thus $\pi(x) = \{\pi(y) : y \in x\}$. As $x \subseteq X$, it follows by minimality that each $y \in x$ has $\pi(y) = y$ and hence $\pi(x) = \{y : y \in x\} = x$. $\dashv$

Now so far, we've been investigating and developing this theory for seemingly no reason. But an important application of this Corollary 8 C • 4 gives the relative consistency of lots of combinatorial principles. For now, we just show that the *generalized* continuum hypothesis (GCH) holds: $2^\kappa = \kappa^+$ for infinite cardinals κ. Recall from Cantor's Theorem (5 B • 13) (and Result 5 D • 6) that we only know $2^\kappa \ge \kappa^+$. From the method of forcing (which hasn't been introduced here), 2^κ can consistently be *any* cardinal of cofinality $> \kappa$. So L thinks 2^κ is as small as it can possibly be all of the time.

The general idea behind the proof is that all of the subsets of κ appear by stage L_{κ^+}. Recall that although $\kappa \in \mathrm{L}_{\kappa+1}$, not every subset of κ in L may appear at stage $\mathrm{L}_{\kappa+2}$, unlike V where $\kappa \in \mathrm{V}_{\kappa+1}$ but $\mathcal{P}(\kappa) \in \mathrm{V}_{\kappa+2}$.

8 C • 5. Theorem

$\mathsf{L} \vDash \mathsf{GCH}$, meaning $\mathsf{L} \vDash$ "$\forall \kappa\ (|\kappa| = \kappa \ge \aleph_0 \to 2^\kappa = \kappa^+)$".

Proof .:.

Argue in a model of "V = L" to suppress so many superscripts of $\mathbf{L}$. Let $\kappa \geq \aleph_0$ be a cardinal, and let $x \in \wp(\kappa)$ be arbitrary so that $x \in \mathrm{L}_\alpha$ for some $\alpha \in \mathrm{Ord}$. Let θ be a regular cardinal larger than $\max(\kappa, \alpha)$ (for example, $\theta = \max(\kappa^+, |\alpha|^+)$ works, but we just need it to be regular and sufficiently large). Therefore $\mathbf{L}_\theta \vDash \mathsf{ZF} - \mathsf{P} + $"V = L".

Let $H = \mathrm{cHull}^{\mathbf{L}_\theta}(\{x\} \cup \kappa)$ so that $H = \mathrm{L}_\alpha$ for some $\alpha < \theta$. As $|H| \leq \aleph_0 \cdot |\kappa \cup \{x\}| = \kappa$, it follows by Lemma 8 C • 1 that $\alpha \leq \kappa^+$. Note also that $\kappa \cup \{x\}$ is transitive, so that $\kappa \cup \{x\} \subseteq H$ by Corollary 8 C • 4. In particular, $x \in \mathrm{L}_\alpha \subseteq \mathrm{L}_{\kappa^+}$. As $x \in \wp(\kappa)$ was arbitrary, $\wp(\kappa) \subseteq \mathrm{L}_{\kappa^+}$ and therefore $2^\kappa \leq |\mathrm{L}_{\kappa^+}| = \kappa^+$. By Cantor's Theorem (5 B • 13), $\kappa^+ \leq 2^\kappa$ and thus we have equality. ⊣

Note that this shows there is no hope of proving the consistency of $\neg\mathsf{GCH}$ from ZFC with our current methods: trying to define an inner model with this true in it. Any attempts to define a class C by a formula φ to show $\mathsf{ZFC} \vdash \mathsf{ZFC}^{\mathrm{C}} + \neg\mathsf{GCH}^{\mathrm{C}}$ would also need to have ZFC + "V = L" $\vdash \mathsf{ZFC}^{\mathrm{C}} + \neg\mathsf{GCH}^{\mathrm{C}}$. But as the smallest inner model, any model $\mathbf{M} \vDash \mathsf{ZFC}$ + "V = L" has by absoluteness of L,

$$\mathrm{L}^{\mathbf{M}} = \mathrm{L}^{\mathrm{C}^{\mathbf{M}}} \subseteq \mathrm{C}^{\mathbf{M}} \subseteq \mathrm{M} = \mathrm{L}^{\mathbf{M}}$$

and thus would have $\mathbf{C} = \mathbf{L}^{\mathbf{M}} \vDash \neg\mathsf{GCH}$, a contradiction.

The regularity property of GCH in $\mathbf{L}$ is also manifested in another regularity property in the levels of L, as suggested by the fact that both $\mathbf{H}^{\mathbf{L}}_\kappa$ and $\mathbf{L}_\kappa$ model $\mathsf{ZFC} - \mathsf{P}$.

8 C • 6. Result

Let $\kappa > \aleph_0$ be a regular cardinal. Therefore $\mathrm{L}_\kappa = \mathrm{H}^{\mathbf{L}}_\kappa$.

Proof .:.

Suppose $x \in \mathrm{L}_\kappa$. It suffices to show that $\mathbf{L} \vDash$ "$|\mathrm{trcl}(x)| < \kappa$". As a limit ordinal, $x \in \mathrm{L}_\alpha$ for some $\alpha < \kappa$ and thus $\mathrm{trcl}(x) \subseteq \mathrm{L}_\alpha$ by transitivity. Without loss of generality, we can assume $\alpha \geq \omega$ so by Lemma 8 C • 1, $|\mathrm{trcl}(x)| \leq |\mathrm{L}_\alpha| = |\alpha| < \kappa$ and therefore $x \in \mathrm{H}^{\mathbf{L}}_\kappa$ so that $\mathrm{L}_\kappa \subseteq \mathrm{H}^{\mathbf{L}}_\kappa$.

So now we consider $x \in \mathrm{H}^{\mathbf{L}}_\kappa$. Note that $\mathrm{trcl}(x) \cup \{x\} \in \mathrm{L}_\alpha$ for some $\alpha \in \mathrm{Ord}$. Let $\theta > \alpha$ be a regular cardinal (which is then regular in $\mathbf{L}$). Consider the skolem hull $H = \mathrm{cHull}^{\mathbf{L}_\theta}(\mathrm{trcl}(x) \cup \{x\})$ which then has size $|H| \leq \aleph_0 \cdot |\mathrm{trcl}(x) \cup \{x\}| < \kappa$. By Condensation (8 B • 3), $H = \mathrm{L}_\beta$ for some β. In fact, since $\beta = \mathrm{Ord} \cap H$ and $|H| < \kappa$, $\beta < \kappa$. Moreover, Corollary 8 C • 4 implies that, as a transitive set, $\mathrm{trcl}(x) \cup \{x\} \subseteq H$ is left uncollapsed. Hence $\mathrm{trcl}(x) \cup \{x\} \in \mathrm{L}_\beta \subseteq \mathrm{L}_\kappa$. In particular, $x \in \mathrm{L}_\kappa$ and thus $\mathrm{H}^{\mathbf{L}}_\kappa = \mathrm{L}_\kappa$. ⊣

In combination with Corollary 7 C • 10, this shows that weakly inaccessible cardinals yield set models of ZFC.

8 C • 7. Theorem

Let κ be weakly inaccessible. Therefore $\mathbf{L}_\kappa \vDash \mathsf{ZFC}$.

Proof .:.

Any weakly inaccessible cardinal is regular so that $\mathrm{L}_\kappa = \mathrm{H}^{\mathbf{L}}_\kappa$. By GCH in $\mathbf{L}$, κ being weakly inaccessible is the same as κ being strongly inaccessible. Therefore, $\mathbf{L} \vDash$ "$\varphi^{\mathrm{H}_\kappa}$" and so $\varphi^{\mathrm{L}_\kappa}$ for each φ of ZFC. ⊣

8 C • 8. Corollary

If κ is (strongly or weakly) inaccessible, then there is a countable, transitive model of ZFC.

Proof .:.

$\mathbf{L}_\kappa$ is a model of ZFC so that $\mathrm{cHull}^{\mathrm{L}_\kappa}(\emptyset) \vDash \mathsf{ZFC}$ and is countable. ⊣

Note that The Reflection Principle (7 D • 5) tells us that every finite fragment of ZFC has a countable transitive model, but it's not provable in ZFC that there is a countable, transitive model of the entirety of ZFC, as this would imply $\mathsf{ZFC} \vdash \mathrm{Con}(\mathsf{ZFC})$, contradicting Gödel's incompleteness theorems.

As a final note about L for this section, the ordering $<_{\mathrm{L}} = \bigcup_{\alpha \in \mathrm{Ord}} <_{\mathrm{L}_\alpha}$ described in Theorem 8 A • 8 is definable, and is a well-order of all of L. So $\mathbf{L} \vDash \mathsf{AC}$ just because there of a much stronger principle holding: the existence of a *global* well-order. This is stronger than mere AC, which says there is a well-order of each individual set, but perhaps not of the entire universe, a proper class. Next, we will investigate what happens when there is a definable global well-order in general. This yields another inner model HOD. Of course, it's consistent that "HOD = L" holds, since any definable inner model contained L must be L itself by Absoluteness of L (8 B • 1): $\mathrm{L} = \mathrm{L}^{\mathrm{HOD}^{\mathsf{L}}} \subseteq \mathrm{HOD}^{\mathsf{L}} \subseteq \mathrm{L}$.

§ 8 D. Hereditarily ordinal definable sets

Whereas L is a very rigid inner model by Condensation (8 B • 3), the next inner model we will consider will be very flexible. It is so flexible, in fact, that V = HOD might be false in HOD, which is to say $\mathrm{HOD}^{\mathrm{HOD}}$ might not be HOD. Consistently, for any countable, transitive $\mathbf{W} \vDash \mathsf{ZFC}$, there is a $\mathbf{U}$ with W an inner model of U such that $\mathbf{U} \vDash$ "W = HOD". So no matter where we start, it's consistent we're starting from HOD of a larger model. We begin—as with L—with the closure under definability.

Recall that we needed the clumsy closure definition of definability from Definition 8 A • 2 because we wanted to work with models of $\mathsf{ZF} - \mathsf{P}$ to ensure the absoluteness of L, and in particular to get Corollary 8 C • 4. For HOD, we have no such interest, because HOD is so flexible, even under full ZFC. So we will use P with the ostensibly more complicated formula from Theorem 6 B • 6, defining what it means to have $\mathbf{A} \vDash \varphi$ for $\mathbf{A}$ a set.[xxvii]

8 D • 1. Definition

A set x is *ordinal definable* or OD iff there is an $\alpha \in \mathrm{Ord}$ such that for some (coded) formula $\varphi \in (\omega \sqcup \{\in\})^{<\omega}$ and parameters $\vec{\beta} \in \mathrm{Ord} \cap \mathrm{V}_\alpha$, $\mathbf{V}_\alpha \vDash$ "$\forall y\ (\varphi(\vec{\beta}, y) \leftrightarrow x = y)$".

The reason for taking a level of V is just to make the definition more concrete: using The Reflection Principle (7 D • 5), x is ordinal definable iff there is a formula with only ordinal parameters such that x is defined by this formula.[xxviii] And of course, every set which is ordinal definable over $\mathbf{V}_\alpha$ (using parameters $\vec{\beta}$) is ordinal definable over $\mathbf{V}$ using the ordinals $\vec{\beta}$ and now with α.

8 D • 2. Corollary

For $\varphi(\vec{w}, x)$ a $\mathtt{FOL}(\in)$-formula, $\mathsf{ZF} \vdash \forall \vec{\alpha} \in \mathrm{Ord}\ \forall x\ (\forall y\ (x = y \leftrightarrow \varphi(\vec{\alpha}, y)) \rightarrow x \in \mathrm{OD})$.

Note further some immediate consequences of this definition.

- $\mathrm{Ord} \subseteq \mathrm{OD}$;
- $x, y \in \mathrm{OD}$ implies $\{x, y\} \in \mathrm{OD}$;
- $x \in \mathrm{OD}$ implies $\mathcal{P}(x) \in \mathrm{OD}$.
- $\mathrm{V}_\alpha \in \mathrm{OD}$ for each $\alpha \in \mathrm{Ord}$, as it is definable from α.

Now, unfortunately, OD might not be transitive (otherwise $\mathrm{V} = \bigcup_{\alpha \in \mathrm{Ord}} \mathrm{V}_\alpha \subseteq \mathrm{OD} \subseteq \mathrm{V}$ implies that every set is ordinal definable). To counteract this issue, we bring in the concept of begin *hereditarily* ordinal definable. Then we can confirm the other axioms of ZFC, as the resulting class will be transitive, allowing us to use results from Section 7.

8 D • 3. Definition

The class HOD of *hereditarily ordinal definable sets* consists of all sets x such that $x \in \mathrm{OD}$ and $\mathrm{trcl}(x) \subseteq \mathrm{OD}$.

8 D • 4. Corollary

HOD is a transitive class.

[xxvii] Note there that for finite signatures, φ can be regarded as a sequence of natural numbers where each number corresponds to a symbol. This makes the formal definition of satisfaction and definability more complicated, but dramatically simplifies the presentation. In particular, when quantifying over "formulas" of any given model, we're quantifying over elements of ω where the sequence $\langle n_0, \cdots, n_m \rangle \in \omega^{<\omega}$ is encoded by the number $2^{n_0+1} 3^{n_0+1} \cdots p_m^{n_m+1} \in \omega$ where p_m is the mth prime number of ω starting with $p_0 = 2$.

[xxviii] Of course, as with other issues about definability, we should be slightly careful about this. If a model of ZFC misinterprets ω, then it will have formulas that are not actual (coded) formulas. The issue is, as before, trying to identify the real world formulas with the formulas of the model. All real world formulas yield formulas of the model, but there may be formulas of the model that are not actual formulas.

Note how this reflects the idea of "hereditarily" as in Definition 7 C • 4. There are two basic ideas about HOD that we care about:

1. **HOD** $\vDash$ ZFC; and
2. M = HOD iff there is a definable, global well-order of M over **M**.

First we define what a global well-order is.

8 D • 5. Definition

A *class well-ordering (definable in* **V***)* of a class C is a class $\preccurlyeq \subseteq C \times C$ defined over **V** by some FOLp($\in$)-formula φ such that every non-empty $X \in C$ has a $\preccurlyeq$-least element.
A *definable, global well-order* of **V** is a class well-order (definable in **V**) of V.

Clearly if a class C has a well-order (definable in **V**) $\preccurlyeq$, then **V** can *uniformly* get choice sets for sets in C just by considering the $\preccurlyeq$-least elements in whatever non-empty family of non-empty, disjoint sets. Note that HOD has such a well-order defined according to the defining formula and ordinal parameters used to define its members.

8 D • 6. Lemma

There is a (definable in **V**) class well-ordering of OD. Moreover, all initial segments of this well-ordering are sets.

Proof .:.

Consider the *Gödel ordering* of $<_{\mathrm{Ord}^{<\omega}}$ on finite sequences of ordinals where $\langle \alpha_1, \cdots, \alpha_n \rangle <_{\mathrm{Ord}^{<\omega}} \langle \beta_1, \cdots, \beta_m \rangle$ for $n, m \in \omega$ iff

- $\max(\vec{\alpha}) < \max(\vec{\beta})$; or else
- $n < m$; or else
- $\vec{\alpha} <_{\mathrm{lex}} \vec{\beta}$.

It should be clear that $<_{\mathrm{Ord}^{<\omega}}$ is well-founded, and seeing that it's linear isn't difficult. Hence $<_{\mathrm{Ord}^{<\omega}}$ is a well-order. It should also be clear that for any particular $\vec{\alpha}$, all $<_{\mathrm{Ord}^{<\omega}}$-predecessors of $\vec{\alpha}$ are contained in $(\max(\vec{\alpha}) + 1)^{<\omega}$, and thus the initial segments of $<_{\mathrm{Ord}^{<\omega}}$ are sets.

As a result, define the ordering $\preccurlyeq_{\mathrm{OD}}$ on OD by taking $x \preccurlyeq_{\mathrm{OD}} y$ iff

- the least α where x is OD in V_α is less than the least for y; or else
- the $<_{\mathrm{Ord}^{<\omega}}$-least set of parameters used to define x in this V_α $<_{\mathrm{Ord}^{<\omega}}$-precedes those for y; or else
- the formula φ used to define x with those parameters in V_α $<_{\mathrm{lex}}$-precedes that formula for y.

It should be clear that this yields a well-ordering of OD. ⊣

As HOD $\subseteq$ OD, we have the following.

8 D • 7. Corollary

HOD has a well-order definable over **V** and thus HOD $\vDash$ AC.

Proof .:.

Using $<_{\mathrm{OD}}$, if $F \in$ HOD is a non-empty family of non-empty, disjoint sets, then let F be defined by ψ in $\mathbf{V}_\alpha$ with ordinal parameters $\vec{\beta}$. In **V**, we can then consider the $<_{\mathrm{OD}}$-least elements of $x \in F$: $z \in C$ iff $\exists x\, \exists F'\, (x \in F' \wedge \psi^{V_\alpha}(\vec{\beta}, F') \wedge z \in x$ is $<_{\mathrm{OD}}$-least$)$. Note that then $C \in$ HOD, which shows that HOD $\vDash$ AC. ⊣

Now we can confirm HOD $\vDash$ ZFC. Note that HOD is, of course, non-empty as Ord $\subseteq$ HOD. The majority of the proof of this comes down to coming up with a definition to show that whatever set we're interested in x has $x \in$ OD and $x \subseteq$ HOD, implying that $x \in$ HOD because $\mathrm{trcl}(x \cup \{x\}) \in$ OD.

8 D • 8. Theorem

HOD $\vDash$ ZFC, and therefore HOD is an inner model.

Proof .:.

Let $x, y \in \mathrm{HOD}$ as witnessed by ψ_x and ψ_y with parameters $\vec{\alpha}_x$ and $\vec{\alpha}_y$ in V_{γ_x} and V_{γ_y} respectively: $x \in \mathrm{V}_{\gamma_x}$ is such that $\mathbf{V}_{\gamma_x} \models$ "$\forall z\ (\varphi(\vec{\alpha}_x, z) \leftrightarrow x = z)$", and similarly for y.

- As usual, extensionality, empty set, and foundation follow from HOD being a (non-empty) transitive class.
- For Pair, note that $\{x, y\}$ is OD defined by $z \in \{x, y\}$ iff $\mathbf{V}_{\gamma_x} \models$ "$\psi_x(\vec{\alpha}_x, z)$" $\vee$ $\mathbf{V}_{\gamma_y} \models$ "$\psi_y(\vec{\alpha}_y, z)$". This formula has parameters $\vec{\alpha}_x$, $\vec{\alpha}_y$, γ_x, and γ_y so that $\{x, y\} \in \mathrm{OD}$ and therefore $\{x, y\} \in \mathrm{HOD}$ as $\mathrm{trcl}(\{x, y\}) = \mathrm{trcl}(x) \cup \mathrm{trcl}(y) \cup \{x, y\} \subseteq \mathrm{HOD}$.
- For Union, the union $\bigcup x$ is defined by $z \in \bigcup x$ iff $\mathbf{V}_{\gamma_x} \models$ "$\exists y\ \exists x' (\psi_x(\vec{\alpha}_x, x') \wedge y \in x' \wedge z \in y)$". Thus $\bigcup x$ is OD. As $\mathrm{trcl}(\bigcup x) \subseteq \mathrm{trcl}(x) \subseteq \mathrm{HOD}$, it follows that $\bigcup x \in \mathrm{HOD}$.
- For Comp, let $\varphi(\vec{w}, z)$ be arbitrary. We want to show $x_\varphi = \{z \in x : \varphi^{\mathrm{HOD}}(\vec{t}, x, z)\} \in \mathrm{HOD}$ for each $\vec{t} \in \mathrm{HOD}^{<\omega}$. Note that x_φ can be defined in $\mathbf{V}$ by
$$v = x_\varphi \quad \text{iff} \quad \forall z \left(z \in v \leftrightarrow \exists \vec{w}\ \exists x' \left(\vec{w} = \vec{t} \wedge \mathbf{V}_{\gamma_x} \models \text{“}\psi(\vec{\alpha}, x')\text{”} \wedge \varphi^{\mathrm{HOD}}(\vec{w}, x', z)\right)\right),$$
where by $\vec{w} = \vec{t}$ we really mean the conjunction of "w_i is the unique element satisfying the defining formula for t_i with the corresponding parameters in the corresponding level of V" for each i. As a result, x_φ is OD and so $\mathrm{trcl}(x_\varphi) \subseteq \mathrm{trcl}(x) \subseteq \mathrm{HOD}$ implies $x_\varphi \in \mathrm{HOD}$.
- For P, note that $\mathcal{P}(x) \cap \mathrm{HOD} \subseteq \mathcal{P}(x) \in \mathrm{V}_{\gamma_x+1} \in \mathrm{OD}$. This implies $\mathrm{V}_{\gamma_x+1} \cap \mathrm{HOD} \in \mathrm{HOD}$, which is a set that contains every $z \in \mathrm{HOD}$ such that $z \subseteq x$. This implies wP: $\mathcal{P}(x) \cap \mathrm{HOD} \in \mathrm{V}_{\gamma_x+1} \cap \mathrm{HOD} \in \mathrm{HOD}$. By Comp, P holds.
- For Rep, let ψ define in **HOD** a function over $x \in \mathrm{HOD}$. Let $\alpha \in \mathrm{Ord}$ be such V_α contains the range of the function defined by ψ^{HOD} over x. Since $\mathrm{V}_\alpha \in \mathrm{OD}$, $\mathrm{V}_\alpha \cap \mathrm{HOD} \in \mathrm{HOD}$ and this witnesses wRep. Therefore by Comp, Rep holds. ⊣

Therefore, as L is the smallest inner model, we have that L is its own HOD.

8 D • 9. Corollary

$\mathrm{L} = \mathrm{HOD}^{\mathrm{L}}$ and thus $\mathbf{L} \models$ "$\mathrm{L} = \mathrm{HOD}$".

The issue with asking questions of HOD in general is due to its highly non-constructive nature: a subset of something (relatively) small like ω might need parameters in V_α for extremely large α to be defined, for example. In this sense, HOD requires knowing about all of the sets of V. And this is the general idea why $\mathrm{HOD}^{\mathbf{HOD}}$ might not be HOD. More precisely, because the V_αs are not absolute between inner models, being ordinal definable is not absolute.

So far we have shown $\mathrm{HOD} \models \mathsf{ZFC}$, and $\mathrm{V} = \mathrm{HOD}$ implies V has a definable, global well-order, because $\mathrm{V} = \mathrm{HOD}$ is clearly equivalent to $\mathrm{V} = \mathrm{OD}$, which provably has a definable, global well-order. Our final goal for this section is then to show the converse: if V has a definable, global well-order, then $\mathrm{V} = \mathrm{HOD}$.

8 D • 10. Theorem

Let φ be a $\mathtt{FOL}(\in)$-formula. Suppose $\preccurlyeq$, defined by $x \preccurlyeq y$ iff $\varphi(x, y)$, is a global well-order of V. Therefore $\mathrm{V} = \mathrm{HOD}$.

Proof .:.

By Lemma 4 • 3, there is a $\mathtt{FOL}$-definable (class) function $f : \mathrm{V} \to \mathrm{Ord}$ where $f(x)$ is the rank of x in $\langle \mathrm{V}, \preccurlyeq \rangle$. But then x is definable from the ordinal $f(x)$. In particular, if $f(x) = y$ is defined by $\psi(x, y)$, then $z = x$ iff $\psi(z, \alpha)$ for $\alpha = f(x)$. Hence $\mathrm{V} = \mathrm{OD}$ and therefore $\mathrm{trcl}(x \cup \{x\}) \subseteq \mathrm{OD}$ automatically. Thus $\mathrm{V} = \mathrm{HOD}$. ⊣

8 D • 11. Corollary

$\mathrm{V} = \mathrm{HOD}$ iff there is a $\mathtt{FOL}$p-definable, global well-order.

Proof .:.

One direction was just proven in Theorem 8 D • 10.. If $\mathrm{V} = \mathrm{HOD}$, then since V has a class well order of $\mathrm{HOD} = \mathrm{V}$ definable in $\mathbf{V}$, there is a global well-order. ⊣

So how is it possible for $\mathrm{HOD}^{\mathrm{HOD}}$ to not be HOD? Ostensibly, HOD has a definable, global well-order by virtue of the one from V and thus $\mathsf{HOD} \vDash$ "$\mathrm{V} = \mathrm{HOD}$", i.e. $\mathrm{HOD} = \mathrm{HOD}^{\mathsf{HOD}}$. The issue is that this definable well-order is not absolute, because it depends on the levels of V rather than the levels of HOD. Again, because $\mathrm{HOD}^{\mathsf{HOD}}$ has lost the information about the other sets in V, we don't know that the definition for the global well-order still yields a global well-order. We only have $\mathrm{HOD} = \mathrm{HOD}^{\mathsf{HOD}}$ if this (or some other) global well order is definable over HOD, not V.

Section 9. *Variants of the Axioms*

There are many propositions ostensibly stronger than other axioms, but turn out to be equivalent to the rest of ZFC. We detail some of these proposals, as well as some variant axioms that are actual weakenings of other axioms.

§ 9 A. Equivalents of the axiom of choice

Recall the official definition of AC.

9 A • 1. Definition (Axiom)

(AC) for any family of non-empty family of disjoint sets F, there is a set C which has chosen one element from each $z \in F$:

$$\forall F \; (\emptyset \notin F \wedge \forall x, y \in F \; (x \cap y = \emptyset) \rightarrow \exists C \; \forall x \in F \; !y \; (y \in x \cap C).$$

We will give three equivalent (over ZF) formulations of AC. Recall that a *chain* for a poset $\langle A, \preccurlyeq \rangle$ is just a $\preccurlyeq$-linearly order subset of A.

9 A • 2. Definition (Axiom)

(*Zorn's Lemma*, AC_Z) For every (non-empty) poset $\langle A, \preccurlyeq \rangle$, if every chain is bounded in A, then A has a $\preccurlyeq$-maximal element.
(AC_P) If F is a non-empty set of non-empty sets, then $\prod_{x \in F} x$ is non-empty.
(AC_C) Every set is bijective with an ordinal.
(AC_W) Every set has a well-order.

Note that by Lemma 5 C • 1, AC_W is equivalent to AC over models of ZF. We will use this extensively to show the equivalences $\mathsf{AC} \leftrightarrow \mathsf{AC}_Z \leftrightarrow \mathsf{AC}_C \leftrightarrow \mathsf{AC}_W$. First we have Zorn's lemma.

9 A • 3. Theorem

ZF ⊢ "$\mathsf{AC} \leftrightarrow \mathsf{AC}_Z$".

Proof .:.

- ($\mathsf{AC}_Z \rightarrow \mathsf{AC}$) This implication holds in BST. For F a non-empty family of non-empty, disjoint sets, by comprehension, union, and powerset, consider the set
$$C = \left\{ y \in \mathcal{P}\left(\bigcup F\right) : \forall x \in F \; (y \cap x = \emptyset \vee \exists! z \in x \; (z \in y)) \right\}$$
of approximations to a choice set for F. Note that the poset $\langle C, \subseteq \rangle$ exists by the existence of cartesian products. Note that for any chain $c \subseteq C$, $\bigcup c$ yields another chain so that $\bigcup c \in C$ and c is bounded by $\bigcup c$. Therefore if AC_Z holds, then there is a $\subseteq$-maximal element X of C. But then for any $x \in F$, there must be some $z \in x$ with $z \in X$, as otherwise (by pairing and union) $X \cup \{z\}$ would contradict maximality. Therefore X must be a choice set.
- ($\mathsf{AC}_W \rightarrow \mathsf{AC}_Z$) Let $\langle A, \preccurlyeq \rangle$ be a (non-empty) poset such that every chain is bounded in A. By AC_W, let $<_A$ be a well-order of A. Define by transfinite recursion an injective function f from ordinals into A.
  - Let $f(0)$ be the $<_A$-least element of A.
  - Let $f(\alpha + 1)$ to be the $<_A$-least element of $\{a \in A : f(\alpha) \prec_A a\}$.
  - For limit γ, note that $f"\gamma \subseteq A$ is a chain. Hence we can take $f(\gamma)$ to be the $<_A$-least element that bounds $f"\gamma$.

  If A has no $\preccurlyeq$-maximal element, then $f(\alpha)$ can be defined for all α. Note, however, that f is injective by transitivity. This contradicts replacement given that A is a set.

Arguably the easiest equivalence to prove is that $\mathsf{AC} \leftrightarrow \mathsf{AC}_P$. Note that any $f \in \prod_{x \in F} x$ is referred to as a *choice function* in that $f(x) \in x$ for each $x \in F$.

9 A • 4. Theorem

BST ⊢ "$\mathsf{AC}_P \to \mathsf{AC}$" and with P or Rep, "$\mathsf{AC} \to \mathsf{AC}_P$".

Proof .:.

Suppose AC_P holds. Let F be a non-empty set of disjoint, non-empty sets. Let $f \in \prod_{x \in F} x$ be a choice function. Therefore $\operatorname{im} f$ is a choice set so that AC holds.

Suppose AC holds. Let F be an aribtrary non-empty set of non-empty sets. Consider $F' = \{\{x\} \times x : x \in F\}$ which exists by the existence of cartesian products, and either replacement or powerset with comprehension. Note that F' is a non-empty family of now disjoint (by considering the first component), non-empty sets. Therefore a choice set C exists by AC. Now every element of C is of the form $\langle x, a\rangle$ for $x \in F$ and $a \in x$. Moreover, for each $x \in F$ there is exactly one a such that $\langle x, a\rangle \in C$. Hence $C \in \prod_{x \in F} x$ is a choice function. ⊣

The last equivalence is that $\mathsf{AC} \leftrightarrow \mathsf{AC}_C$. We have already proven that ZF ⊢ "$\mathsf{AC} \to \mathsf{AC}_C$" with Theorem 5 B • 5. So it suffices to show $\mathsf{AC}_C \to \mathsf{AC}_W$.

9 A • 5. Theorem

BST ⊢ "$\mathsf{AC}_C \to \mathsf{AC}_W$".

Proof .:.

Let X be arbitrary. By AC_C, there is a bijection $f : X \to \alpha$ for some ordinal $\alpha \in \mathrm{Ord}$. Therefore $\{\langle x, y\rangle : f(x) \in f(y)\} \subseteq X \times X$ is a well-order of X. ⊣

Therefore, over ZF, AC_C, AC_W, AC_Z, AC_P, and AC are all equivalent.

§ 9 B. Weakenings of the axiom of choice

If AC is saying that we can make choices infinitely often, there are then variations of this that weaken how many choices we can make, or when we can make them. In particular, the first way we can weaken AC is by restricting outselves to making just finitely many choices.

9 B • 1. Definition

The axiom of *finite choice* says that for any finite family of non-empty, disjoint sets, there is a choice set for the family.

This weakening is so weak, that it is provable.

9 B • 2. Theorem

ZF ⊢ finite choice.

Proof .:.

Suppose F is a finite family of non-empty, disjoint sets, meaning there is a bijection $f : |F| \to F$ where $|F| < \omega$. Proceed by induction on $|F|$ to show that there is a choice function for F. For $|F| = 1$, this is easy: $F = \{x\}$ for some $x \neq \emptyset$. As $x \neq \emptyset$, by basic properties of first-order logic, there is some $y \in x$ where then $\{y\}$ is a choice set for F.

For $|F| = n + 1$, let $x \in F$ be arbitrary. Inductively, $F \setminus \{x\}$ has a choice set C. Now take $y \in x$ and let $C' = C \cup \{y\}$ to get a choice set for F. By induction, all such F have a choice set. ⊣

So really finite choice is just a consequence of applying existential instantiation from `FOL`-proofs finitely many times. The reason AC is needed, is because proofs are finite, and we can't apply existential instantiation infinitely many times with a finite proof.

The second weakening of AC is that we can make *countably* many choices, which is then clearly stronger than finite choice, and is in fact independent of ZF: ZF doesn't prove it, nor disprove it.

9 B • 3. Definition

The axiom of *countable choice* says that for any non-empty, countable family (meaning there is an injection from it into ω) of disjoint, non-empty sets, there is a choice set for the family.

One important consequence of countable choice is Kőnig's theorem on trees, which requires some version of choice. First we introduce the concept of a tree, which is incredibly important in set theory in that it is a slight generalization of ordinals.

9 B • 4. Definition

A *tree* is a poset $\langle A, \leqslant\rangle$ such that for every $a \in A$, $\mathrm{pred}_{\leqslant}(a)$ is well-ordered by $\leqslant$.
A tree is *finitely splitting* iff there are finitely many least elements, and for every node $a \in A$ there are at most finitely many direct $\leqslant$-successors to a.
A *branch* is a $\subseteq$-maximal, $\leqslant$-linearly ordered subset of A.

In particular, if a tree has height $n < \omega$, then the tree is finite and so finite choice yields a branch with height n. But is there an infinite branch if the tree is infinite (of height ω)? AC and Kőnig's theorem in particular state that this is true.

9 B • 5. Theorem (Kőnig's Lemma on Trees)

Let $\mathsf{T} = \langle T, \leqslant_T\rangle$ be an finitely splitting tree of height ω. Therefore

$$\mathsf{ZF} + \text{“countable choice”} \vdash \text{“there is an infinite branch of } T\text{”}.$$

Proof .:.

First we will show that T must be countable.

Claim 1

There is an injection $f : T \to \omega$.

Proof .:.

It should be clear by induction that for each $n < \omega$, the nth level of T is finite. For each $n < \omega$, consider the set B_n of bijections from $|\mathrm{lvl}_n(T)| < \omega$ to $\mathrm{lvl}_n(T)$ (which is a set because it's a subset of ${}^{<\omega}\mathrm{lvl}_n(T)$). Note that the family $\{B_n : n < \omega\}$ of these bijections is then countable, nonempty, and each B_n is also non-empty and disjoint as the levels are disjoint. So by countable choice, there is a choice C set for this family.

So let $f_n \in C$ be $f_n : \mathrm{lvl}_n(T) \to |\mathrm{lvl}_n(T)|$ for each $n < \omega$. Let $f : T \to \omega$ be defined by $f(\tau) = 2^n \cdot 3^{f_n(\tau)}$, for $\tau \in \mathrm{lvl}_n(T)$. Since each element is only in one level, and for each level $n < \omega$, f_n is a bijection, it follows that f is an injection. ⊣

For each $\tau \in T$, let S_τ consist of all $\sigma \in T$ such that

- σ is a direct successor to τ (i.e. $\tau <_T \sigma$ and there are no $\rho \in T$ with $\tau <_T \rho <_T \sigma$); and
- there are infinitely many $\leqslant_T$-successors to σ (i.e. the set of $\rho \in T$ with $\sigma \leqslant_T \rho$ is infinite).

Note that S_τ might be empty for some τ. But clearly, for some $\tau_0 \in T$ being $\leqslant_T$-minimal, S_{τ_0} is non-empty (otherwise T will be the union of finitely many finite sets and thus be finite). Therefore $F = \{S_\tau : \tau \in T \wedge S_\tau \neq \emptyset\}$ is a countable, non-empty family of non-empty, disjoint sets. Thus by countable choice, there is a choice set C for F.

Now we proceed by recursion to give an infinite path. Let τ_0 be a $\leqslant_T$-least element of T with infinitely many successors. By the same reasoning as above one of the direct successors also must have infinitely many successors: $S_{\tau_0} \neq \emptyset$ (meaning $S_{\tau_0} \in F$). For τ_n already defined with $S_{\tau_n} \in F$, take τ_{n+1} to be the unique element of $C \cap S_{\tau_n}$. By the same reasoning before, we must have $S_{\tau_{n+1}} \in F$. Therefore the sequence defined by recursion, $\langle \tau_n : n \in \omega\rangle$, yields an infinite branch of T. ⊣

Let's consider an alternative way to state countable choice. ... This generalizes to the axiom of dependent choice.

9 B • 6. Definition

The axiom of *dependent choice* (DC) says that for $R \subseteq X \times X$, if $\forall x \in X\ \exists y \in X\ (x\ R\ y)$ then there is a sequence $\langle x_n : n \in \omega \rangle$ such that $x_n\ R\ x_{n+1}$ for all $n \in \omega$.

Just from this definition, it's not immediately clear that DC implies countable choice, but we can show this without much effort.

9 B • 7. Theorem

ZF $\vdash$ "DC $\rightarrow$ countable choice".

Proof .:.

Assume DC and suppose F is a non-empty, countable set of non-empty, disjoint sets, as witnessed by an injection $f : F \rightarrow \omega$. Without loss of generality, F is infinite as finite choice follows from ZF alone. Consider the relation $R \subseteq (\bigcup F) \times (\bigcup F)$ defined as follows. For each $x \in \bigcup F$ let $F_x \in F$ be the unique element of F that has x as an element. Thus $x \in F_x$ for each $x \in \bigcup F$. Define $x\ R\ y$ iff

1. $f(F_x) < f(F_y)$; and
2. there is no $X \in F$ where $f(F_x) < f(X) < f(F_y)$.

Since F is infinite, im f is unbounded in ω, meaning that for each $x \in \bigcup F$, there is some $y \in \bigcup F$ where $f(F_x) < f(F_y)$. Therefore by DC, there is a sequence $\langle x_n \in \bigcup F : n \in \omega \rangle$ where $f(F_{x_n}) < f(F_{x_{n+1}})$ for all $n \in \omega$. Without loss of generality (just by finite choice in ZF to add in finitely many entries in the sequence) we can assume $\min\{f(F_{x_n}) : n \in \omega\} = \min\{f(X) : X \in F\}$. Let $C = \{x_n : n < \omega\}$. This is a choice set for $\{F_{x_n} : n < \omega\}$.

So now we show that $\{F_{x_n} : n < \omega\} = F$. Since clearly $\{F_{x_n} : n < \omega\} \subseteq F$, suppose $X \in F \setminus \{F_{x_n} : n < \omega\}$. Note that $f(X) = n$ for some $n < \omega$ where then $n \geq \min\{f(F_{x_m}) : m < \omega\}$. Clearly if equality holds, then $X = F_{x_m}$ for some $m < \omega$ by injectivity of f. Thus, as $\{f(F_{x_m}) : m < \omega\}$ is unbounded in ω, there is some $m < \omega$ where $f(F_{x_m}) < f(X) < f(F_{x_{m+1}})$. But then $\neg(x_m\ R\ x_{m+1})$, a contradiction. Therefore there can be no such X and thus $\{F_{x_n} : n < \omega\} = F$, meaning C is a choice set for F. $\dashv$

The reverse does not hold, implying that DC is strictly stronger than countable choice. DC is used for a great number of theorems and basic results in analysis, particularly in the use of sequences. For instance DC will show the equivalence between continuity (in a general topological sense for $\mathbb{R}$) and sequential continuity. Indeed, in a very vague sense, much of analysis can be carried out in DC, or at least DC relativized to $\mathbb{R}$. Given this, DC is really the first serious weakening of AC used for mathematics.

To show the power of DC over countable choice, we have the following theorem. Countable choice only gave Kőnig's Lemma on Trees (9 B • 5): for a finitely branching tree of height ω, there is an infinite path. The main reason why we needed finitely many branches is to ensure the resulting tree was countable. DC does not need this restriction, only that there are no finite braches (meaning branches with a finite length). In fact, the following consequence of DC is equivalent to DC over ZF.

9 B • 8. Theorem

Therefore ZF $\vdash$ "DC $\leftrightarrow$ every tree of height ω has a branch".

Proof .:.

Assume DC and let $\mathbf{T} = \langle T, \leq_T \rangle$ be a tree of height ω. If there is some $\tau \in T$ with no $\sigma \in T$ where $\tau <_T \sigma$, then $\mathbf{T}$ has a finite branch, which can be seen just by considering $\mathrm{pred}_{<_T}(\tau)$. So suppose T has no finite branches. Thus for each $\tau \in T$, there is some $\sigma \in T$ with $\tau <_T \sigma$. Hence by dependent choice, there is a sequence of $\langle \tau_n \in T : n < \omega \rangle$ where $\tau_n < \tau_{n+1}$ for all $n < \omega$. Closing $\{\tau_n : n < \omega\}$ under $<_T$-predecessors then yields an infinite branch.

Assume every tree of height ω has a branch. Let $R \subseteq X \times X$ be such that $\forall x \in X\ \exists y \in X\ (x\ R\ y)$. Consider the tree of finite sequences

$$T = \{f : n + 1 \rightarrow X : n < \omega \wedge \forall m < n\ (f(m)\ R\ f(m+1))\},$$

ordered by $\leqslant_T$ where $\tau <_T \sigma$ iff $\tau \subsetneq \sigma$. Note that T has no finite branches, since each $f : n+1 \to X$ has some $y \in X$ where $f(n)\ R\ y$ so that $f' = f \cup \{\langle n+1, y\rangle\}$ extends f. Hence T has height ω and therefore has an *infinite* branch $f : \omega \to T$ where $f(n) \subsetneq f(n+1)$. Therefore, $\langle x_n : n < \omega\rangle$ defined by $x_n = f(m)(n)$ for any (and all) $m > n$ (since the domains are strictly increasing as m increases, this is just to make sure $n \in \mathrm{dom}(f(m))$) yields that DC holds for R. ⊣

Since the existence of finite branches will be a part of ZF, DC is really equivalent to the existence of branches when there are no finite branches. So the theorem above (and Kőnig's Lemma on Trees (9 B • 5)) is really talking about the existence of branches when they are forced to be infinite. AC says that every tree (of ordinal height) has branches, whereas these various weakenings say that only certain trees have branches:

- finite choice says finite trees have branches;
- countable choice says countable trees of height ω have branches;
- DC says all trees of height ω have branches;
- AC says all set trees have branches; and
- the extension of AC *global choice*—formally understood as "V = HOD" through Corollary 8 D • 11—says that trees of height Ord have branches.

And we can prove almost all of these to be equivalences. The odd one out is countable choice, which seems to be strictly stronger than countable trees of height ω having branches, and instead seems to be equivalent to when we can also take, in some sense, finite approximations to choice sets.

9 B • 9. Theorem

1. ZF ⊢ "finite choice ↔ finite trees have branches".
2. ZF ⊢ "countable choice ↔ countable trees of height ω have branches $\wedge$ every countable $F \neq \emptyset$ of non-empty, disjoint sets has a C' where $\forall x \in F\ (|C' \cap x| < \aleph_0)$".
3. ZF ⊢ "DC ↔ trees of height ω have branches".
4. ZF ⊢ "AC ↔ trees of height $<$ Ord have branches".
5. **W** ⊨ ZF has **W** ⊨ "V = HOD" iff class trees of **W** of height $\leq \mathrm{Ord}^{\mathbf{W}}$ have (**W**-definable) class branches.

Proof .:.

1. Since both are proven from ZF, it follows that they are equivalent over ZF.

3. This is Theorem 9 B • 8.

2. We've proven that countable choice implies countable trees of height ω have branches by Kőnig's Lemma on Trees (9 B • 5), and obviously an actual choice set C for F yields the second statement. So suppose countable trees of height ω have branches, and we can take choice sets modulo finite subsets. To show countable choice, let $F \neq \emptyset$ be an arbitrary countable set of non-empty, disjoint sets. Let C' be such that $C' \cap x$ is finite for each $x \in F$. Consider the following tree of refinements on C. First, let $f : \omega \to F$ be a bijection. Define
$$T = \{c \subseteq C' : |c| < \omega \wedge \forall n < |c|\ \exists! x\ (x \in c \cap f(n))\},$$
with $c \leqslant_T d$ iff $c \subseteq d$. This is a tree $\mathbf{T} = \langle T, \leqslant_T\rangle$ with height ω and since $C' \cap f(n)$ is finite for each n, $\mathbf{T}$ is finitely splitting. Hence T is countable, and thus there is a branch $C \subseteq T$. Note that this branch is itself a subset of C', and for each $n < \omega$, $\exists! x\ (x \in C \cap f(n))$. Hence C is a choce set for F.

4. Let $\mathbf{T} = \langle T, \leqslant_T\rangle$ be a tree. Suppose AC holds. Therefore, by Theorem 9 A • 3, Zorn's lemma holds. So consider the non-empty poset $\langle A, \preccurlyeq\rangle$ where A consists of chains of T and $a \preccurlyeq b$ iff $a \subseteq b$. Note that every chain of A is bounded in A since a chain $c \subseteq A$ has $\bigcup c$ as a chain of $\mathbf{T}$ and $a \in c$ implies $a \preccurlyeq \bigcup c$. Thus A has a $\preccurlyeq$-maximal element C, which is then a branch of $\mathbf{T}$.

   Now suppose that all set trees have branches. We will show that AC_C holds. Let X be an arbitrary set. Consider the tree where
$$T = \{c : \alpha \to X : \alpha \in \mathrm{Ord} \wedge c \text{ is injective}\},$$

where $c \leqslant_T c'$ iff $c \subseteq c'$. Note that $<_T$ is a well-founded relation, since any infinite $<_T$-decreasing sequence yields a corresponding decreasing sequence of ordinals according to the domains of the entries in the sequence: $c_{n+1} <_T c_n$ implies $\mathrm{dom}(c_{n+1}) < \mathrm{dom}(c_n)$. Moreover, $\leqslant_T$ linearly order predecessors so that $\mathsf{T} = \langle T, \leqslant_T \rangle$ is a tree. By Hartogg's Number (5 C • 5), T is a set, being a subset of ${}^{\kappa}X$ where $\kappa \in \mathrm{Ord}$ has $\kappa \not\leq_{\text{size}} X$. Therefore there is a branch $C \subseteq T$. Note that then $f = \bigcup C$ is an injective function from some ordinal $\alpha < \kappa$ to X. Moreover, f must be surjective, as any $x \in X \setminus \mathrm{im}\, f$ has $f <_T f' = f \cup \{\langle \alpha, x \rangle\}$, contradicting the maximality of f. Therefore $f : \alpha \to X$ is a bijection, showing AC_C holds.

5. Suppose V = HOD holds, and let $\preccurlyeq$ be the definable well-order of V. Let $\mathrm{T} \neq \emptyset$ and the tree order $\leqslant_{\mathrm{T}}$ be classes. Define by transfinite recursion a (possibly class) branch of T. Let t_0 be a minimal element of T. For t_α defined for $\alpha < \beta$, if there is no t with $t_\alpha <_T t$, then the closure of $\{t_\alpha : \alpha < \beta\}$ under $\leqslant_{\mathrm{T}}$-predecessors is a branch of T. Otherwise, take t_β to be the $\preccurlyeq$-least such t. The class $\{t \in \mathrm{T} : \exists\alpha \in \mathrm{Ord}\ (t \leqslant_{\mathrm{T}} t_\alpha)\}$ is then a branch of T.

   Suppose all trees of height Ord have (class) branches. This means that there is a FOLp-definable branch of every tree of height Ord. We will define a class well-ordering of V from this. Consider the class T of functions f from ordinals (to V) such that

   i. f is injective; and
   ii. If $x \in \mathrm{V}_\alpha \cap \mathrm{im}\, f$, then $\mathrm{V}_\beta \subseteq \mathrm{im}\, f$ for $\beta < \alpha$.

   This is ordered by $f \leqslant_{\mathrm{T}} g$ iff $f \subseteq g$. As before, $\leqslant_{\mathrm{T}}$ is well-founded, and linear on collections of predecessors. Thus T and $\leqslant_{\mathrm{T}}$ form a tree so that there is a FOLp-definable branch $\mathrm{C} \subseteq \mathrm{T}$.

   $\bigcup \mathrm{C}$ is then a class function from Ord (to V). Moreover, $\bigcup \mathrm{C}$ is injective. To see that $\bigcup \mathrm{C}$ is surjective, note $x \in \mathrm{V}$ with $\mathrm{rank}(x) = \alpha$ implies $x \in \mathrm{V}_{\alpha+1}$. Since $|\mathrm{V}_{\alpha+1}| < \mathrm{Ord}$ and f is injective, $\mathrm{im}\, f$ cannot be contained in $\mathrm{V}_{\alpha+1}$. Hence there is some $y \in \mathrm{V}_\gamma \cap \mathrm{im}\, f$ for $\gamma > \alpha + 1$. But then by (ii), (a sufficiently large initial segment of) f has $\mathrm{V}_{\alpha+1} \subseteq \mathrm{im}\, f$. Hence $x \in \mathrm{im}\, f$ so that $f : \mathrm{Ord} \to \mathrm{V}$ is a class bijection. This yields the class well-order of V by $x \preccurlyeq y$ iff $f^{-1}(x) < f^{-1}(y)$. ⊣

§ 9 C. The axiom scheme of collection

We begin with *Scott's trick*, a clever way of restricting formulas which ostensibly define proper classes to sets.

9 C • 1. Theorem (Scott's Trick)

Let $\approx \subseteq \mathrm{C} \times \mathrm{C}$ be a definable equivalence relation on a proper class C. Therefore, for each $c \in \mathrm{C}$ we can define the equivalence class of c under $\approx$ to be the *set*

$$[c]_\approx = \{d \in \mathrm{C} : c \approx d \wedge \mathrm{rank}(d) \text{ is the least such rank}\}.$$

Proof .:.

For each $c \in \mathrm{C}$, the class $\{\mathrm{rank}(d) \in \mathrm{Ord} : d \approx c\}$ has a least element, α. Hence $\{d \in \mathrm{V}_{\alpha+1} : c \approx d\}$ is a set. Moreover, for any $d \approx c$, it follows that $[d]_\approx = [c]_\approx$. ⊣

Scott's trick really just says that we can still make sense of equivalence classes for what might ordinarily be proper classes. This can be slightly generalized to other relations.

9 C • 2. Theorem

Let $\varphi(x, y)$ be a FOLp-formula. Suppose for every $x \in D$ there is some y where $\varphi(x, y)$. Therefore there is a *set* R where for every $x \in D$ there is a $y \in R$ with $\varphi(x, y)$ (and for every $y \in R$ there is a $x \in D$ with $\varphi(x, y)$).

Proof .:.

As with Scott's Trick (9 C • 1), for each $d \in D$, let $[d]$ be the set $\{y : \varphi(d, y) \wedge \mathrm{rank}(y) \text{ is least}\}$. Therefore, by replacement, $R = \bigcup_{d \in D} [d]$ yields a set witnessing the result. ⊣

Notice that the above theorem is really a strengthening of replacement. Whereas replacement requires φ to define a function, the above shows that we can weaken this requirement to just being a relation.

9 C • 3. Definition

The *axiom (scheme) of collection* (Coll) consists of formulas of the form

$$\forall w_0 \cdots \forall w_n \forall D \left(\forall x \in D\ \exists y\ \varphi(x, y, \vec{w})\right) \to \exists R\ \forall x \in D\ \exists y \in R\ \varphi(x, y, \vec{w})\big),$$

where φ is a $\mathtt{FOL}(\in)$-formula.

So just by examining the form of this, this is stronger than replacement, although Theorem 9 C • 2 shows they are equivalent over ZF. Note, however, that for Scott's trick and this idea to hold, we required powerset to ensure that $\mathrm{V}_\alpha = \{x : \mathrm{rank}(x) < \alpha\}$ is a set for each α. In fact, under ZF − P, Coll is *strictly* stronger than Rep, although the proof of this is quite complicated.

Given this, in the absence of powerset, we often will work with collection instead of replacement.

9 C • 4. Definition

ZF^- is the theory $\mathsf{ZF} - \mathsf{P} + \mathsf{Coll}$. Similarly ZFC^- is $\mathsf{ZF}^- + \mathsf{AC}$.

Note that we have encountered many toy models of (fragments of) set theory that model ZF^- rather than merely $\mathsf{ZF} - \mathsf{P}$. In particular, $\mathrm{H}_\kappa \vDash \mathsf{ZFC}^-$ for regular κ, showing $\mathrm{L}_\kappa \vDash \mathsf{ZFC}^-$ for κ regular in L.

9 C • 5. Result

Let $\kappa > \aleph_0$ be a regular cardinal. Therefore $\mathsf{H}_\kappa \vDash \mathsf{Coll}$ and thus $\mathsf{H}_\kappa \vDash \mathsf{ZFC}^-$.

Proof .:.

Suppose φ defines a relation over $D \in \mathrm{H}_\kappa$. By choice and collection in V, there is an $R \subseteq \mathrm{H}_\kappa$ such that for each $d \in D$ there is some unique $r \in R$ with $\varphi(d, r)$, and for each $r \in R$ there is some (possibly multiple) $d \in D$ with $\varphi(d, r)$. Hence there is a surjection from D onto R. Since then $|R| \le |D| < \kappa$ with $R \subseteq \mathrm{H}_\kappa$, it follows by regularity of κ that $R \in \mathrm{H}_\kappa$. $\dashv$

§ 9 D. Refinements of collection and comprehension

Because comprehension, collection, and replacement are all schemes—meaning that for each appropriate φ we get a new axiom—there are various refinements of these that restrict what kinds of formulas are allowed.

9 D • 1. Definition

For $n \in \mathbb{N}$, Σ_n-Comp refers to the axiom scheme of comprehension restricted to Σ_n-formulas. And similar definitions hold for Σ_n-Coll, and Σ_n-Rep.

The benefit of having these refinements is being able to have "enough" comprehension, or "enough" collection to play around with. Although a structure might not satisfy full collection or full comprehension, often it will at least satisfy Σ_0-collection or Σ_0-comprehension. This is particularly important in the fine structure theory of L and $\mathrm{L}[E]$ as explored in later chapters.

Similar to these refinements on the axiom schemes, we also have refinements on elementarity. For example,

9 D • 2. Definition

Let σ be a signature. Let A and B be $\mathtt{FOL}(\sigma)$-structures. Let $n < \omega$. A function $j : A \to B$ is a Σ_n-embedding iff for all Σ_n-formulas (defined analogously as in Definition 7 • 4) $\mathsf{A} \vDash$ "$\varphi(\vec{a})$" iff $\mathsf{B} \vDash$ "$\varphi(j(\vec{a}))$". If $j = \mathrm{id}$ so that $\mathsf{A} \subseteq \mathsf{B}$ is a submodel, we also write $\mathsf{A} \preccurlyeq_{\Sigma_n} \mathsf{B}$.

Thus for transitive classes of set theory, a Σ_0-embedding is just an embedding while a Σ_ω-embedding (meaning an embedding which is Σ_n for each $n < \omega$) is an embedding with full elementarity. And this is where the refinements of collection and comprehension become useful: if B satisfies some fragment of comprehension, and there is a sufficiently elementary embedding from A into B, then A will also satisfy some fragment of comprehension.

The Axioms of ZFC

1. (Extensionality, Ext) two sets are equal whenever they have the same members:
$$\forall x\ \forall y\ (x = y \leftrightarrow \forall v\ (v \in x \leftrightarrow v \in y)).$$
2. (Empty set) there is a set $\emptyset$ with no members: $\exists z \forall x (x \notin z)$.
3. (Comprehension, Comp) for each x, and for each $\mathtt{FOL}(\in)$-formula $\varphi(v, \vec{w})$, $\{v \in x : \varphi(v, \vec{w})\}$ exists:
$$\forall w_0 \cdots \forall w_n\ \forall x\ \exists z\ \forall v\ (v \in z \leftrightarrow v \in x \wedge \varphi(v, \vec{w})).$$
4. (Pairing, Pair) for any two sets x and y, the pair $\{x, y\}$ exists: $\forall x\ \forall y\ \exists z \forall v\ (v \in z \leftrightarrow (v = x \vee v = y))$.
5. (Union, Union) for any family of sets F, there is a set containing the elements of all of those sets:
$$\forall F\ \exists U\ \forall v\ (v \in U \leftrightarrow \exists x (x \in F \wedge v \in x)).$$
6. (Foundation, Found) for each x, there is a $\in$-minimal element of x, meaning a member $y \in x$ with no $z \in y$ being in x:
$$\forall x \exists y (y \in x \wedge \forall z (z \in y \to z \notin x)).$$
7. (Infinity, Inf) an infinite set exists: $\exists N (\emptyset \in N \wedge \forall x (x \in N \to x \cup \{x\} \in N))$.
8. (Replacement, Rep) the image of a function over a set is a set: for each $\mathtt{FOL}(\in)$-formula φ,
$$\forall w_0 \cdots \forall w_n \forall D \left(\forall x (x \in D \to \exists! y\ \varphi(x, y, \vec{w})) \to \exists R (y \in R \leftrightarrow \exists x (x \in D \wedge \varphi(x, y, \vec{w})))\right).$$
9. (Powerset, P) for each x, $\mathcal{P}(x)$ exists: $\forall x\ \exists P\ \forall v\ (v \in P \leftrightarrow \forall y\ (y \in v \to y \in x))$.
10. (Choice, AC) for any family of non-empty family of non-empty, disjoint sets F, there is a set C which has chosen one element from each $z \in F$:
$$\forall F\ (\emptyset \notin F \wedge \forall x, y \in F\ (x \cap y = \emptyset) \to \exists C\ \forall x \in F\ \exists! y\ (y \in x \cap C).$$

Variant Axioms and Axiom Systems

i. (Σ_n-Comprehension, Σ_n-Comp) for each x, and for each Σ_n-formula $\varphi(v, \vec{w})$, $\{v \in x : \varphi(v, \vec{w})\}$ exists.

ii. (Collection, Coll) there is a range for a relation with over a given domain: for each $\mathtt{FOL}(\in)$-formula φ,
$$\forall w_0 \cdots \forall w_n \forall D \left(\forall x \in D\ \exists y\ \varphi(x, y, \vec{w})) \to \exists R\ \forall x \in D\ \exists y \in R\ \varphi(x, y, \vec{w})\right).$$

iii. (Σ_n-Collection, Σ_n-Coll) Coll holds for Σ_n-formulas.

iv. (Zorn's Lemma, AC_Z) For every (non-empty) poset $\langle A, \preccurlyeq \rangle$, if every chain is bounded in A, then A has a $\preccurlyeq$-maximal element.

v. (AC_P) If F is a non-empty set of non-empty sets, then $\prod_{x \in F} x$ is non-empty.

vi. (AC_C) Every set is bijective with an ordinal.

vii. (AC_W) Every set has a well-order.

viii. (Dependent choice, DC) for $R \subseteq X \times X$, if $\forall x \in X\ \exists y \in X\ (x\ R\ y)$ then there is a sequence $\langle x_n : n \in \omega \rangle$ such that $x_n\ R\ x_{n+1}$ for all $n \in \omega$.

ix. For every x, y, $x \times y$ exists.

With these axioms, we have the following theories:

- BST consists of (1)–(6) plus (ix).
- ZF^- consists of (1)–(8) plus (ii).
- $\mathsf{ZFC}^- = \mathsf{ZF}^- + \mathsf{AC}$ consists of (1)–(8) plus (10) and (ii).
- $\mathsf{ZF} = \mathsf{ZF}^- + \mathsf{P}$ consists of (1)–(9).
- $\mathsf{ZFC} = \mathsf{ZF} + \mathsf{AC}$ consists of (1)–(10).

Section 10. *Exercises*

lots to put here