

The Foundations of Mathematics

©2005,2006,2007 Kenneth Kunen

Kenneth Kunen

October 29, 2007

Contents

0	Introduction	3
0.1	Prerequisites	3
0.2	Logical Notation	3
0.3	Why Read This Book?	5
0.4	The Foundations of Mathematics	5
I	Set Theory	10
I.1	Plan	10
I.2	The Axioms	10
I.3	Two Remarks on Presentation.	14
I.4	Set theory is the theory of everything	14
I.5	Counting	15
I.6	Extensionality, Comprehension, Pairing, Union	17
I.7	Relations, Functions, Discrete Mathematics	24
	I.7.1 Basics	24
	I.7.2 Foundational Remarks	29
	I.7.3 Well-orderings	31
I.8	Ordinals	33
I.9	Induction and Recursion on the Ordinals	42
I.10	Power Sets	46
I.11	Cardinals	48
I.12	The Axiom of Choice (AC)	56
I.13	Cardinal Arithmetic	61
I.14	The Axiom of Foundation	66
I.15	Real Numbers and Symbolic Entities	73
II	Model Theory and Proof Theory	78
II.1	Plan	78
II.2	Historical Introduction to Proof Theory	78
II.3	NON-Historical Introduction to Model Theory	80
II.4	Polish Notation	81
II.5	First-Order Logic Syntax	85

II.6 Abbreviations	91
II.7 First-Order Logic Semantics	92
II.8 Further Semantic Notions	98
II.9 Tautologies	105
II.10 Formal Proofs	106
II.11 Some Strategies for Constructing Proofs	110
II.12 The Completeness Theorem	116
II.13 More Model Theory	127
II.14 Equational Varieties and Horn Theories	132
II.15 Extensions by Definitions	135
II.16 Elementary Submodels	138
II.17 Other Proof Theories	142
III Recursion Theory	143
Bibliography	144

Chapter 0

Introduction

0.1 Prerequisites

It is assumed that the reader knows basic undergraduate mathematics. Specifically:

You should feel comfortable thinking about abstract mathematical structures such as groups and fields. You should also know the basics of calculus, including some of the theory behind the basics, such as the meaning of limit and the fact that the set \mathbb{R} of real numbers is uncountable, while the set \mathbb{Q} of rational numbers is countable.

You should also know the basics of logic, as is used in elementary mathematics. This includes truth tables for boolean expressions, and the use of predicate logic in mathematics as an abbreviation for more verbose English statements.

0.2 Logical Notation

Ordinary mathematical exposition uses an informal mixture of English words and logical notation. There is nothing “deep” about such notation; it is just a convenient abbreviation which sometimes increases clarity (and sometimes doesn’t). In Chapter II, we shall study logical notation in a formal way, but even before we get there, we shall use logical notation frequently, so we comment on it here.

For example, when talking about the real numbers, we might say

$$\forall x[x^2 > 4 \rightarrow [x > 2 \vee x < -2]] \text{ ,}$$

or we might say in English, that for all x , if $x^2 > 4$ then either $x > 2$ or $x < -2$.

Our logical notation uses the propositional connectives $\vee, \wedge, \neg, \rightarrow, \leftrightarrow$ to abbreviate, respectively, the English “or”, “and”, “not”, “implies”, and “iff” (if and only if). It also uses the *quantifiers*, $\forall x$ and $\exists x$ to abbreviate the English “for all x ” and “there exists x ”.

Note that when using a quantifier, one must always have in mind some intended *domain of discourse*, or *universe* over which the variables are ranging. Thus, in the

above example, whether we use the symbolic “ $\forall x$ ” or we say in English, “for all x ”, it is understood that we mean for all real numbers x . It also presumes that the various functions (e.g. $x \mapsto x^2$) and relations (e.g. $<$) mentioned have some understood meaning on this intended domain, and that the various objects mentioned (4 and ± 2) are in the domain.

“ $\exists! y$ ” is shorthand for “there is a *unique* y ”. For example, again using the real numbers as our universe, it is true that

$$\forall x[x > 0 \rightarrow \exists! y[y^2 = x \wedge y > 0]] \quad ; \quad (*)$$

that is, every positive number has a unique positive square root. If instead we used the rational numbers as our universe, then statement (*) would be false.

The “ $\exists!$ ” could be avoided, since $\exists! y \varphi(y)$ is equivalent to the longer expression $\exists y[\varphi(y) \wedge \forall z[\varphi(z) \rightarrow z = y]]$, but since uniqueness statements are so common in mathematics, it is useful to have some shorthand for them.

Statement (*) is a *sentence*, meaning that it has no free variables. Thus, if the universe is given, then (*) must be either true or false. The fragment $\exists! y[y^2 = x \wedge y > 0]$ is a *formula*, and makes an assertion about the free variable x ; in a given universe, it may be true of some values of x and false of others; for example, in \mathbb{R} , it is true of 3 and false of -3 .

Mathematical exposition will often omit quantifiers, and leave it to the reader to fill them in. For example, when we say that the commutative law, $x \cdot y = y \cdot x$, holds in \mathbb{R} , we are really asserting the sentence $\forall x, y[x \cdot y = y \cdot x]$. When we say “the equation $ax + b = 0$ can always be solved in \mathbb{R} (assuming $a \neq 0$)”, we are really asserting that

$$\forall a, b[a \neq 0 \rightarrow \exists x[a \cdot x + b = 0]] \quad .$$

We know to use a $\forall a, b$ but an $\exists x$ because “ a, b ” come from the front of the alphabet and “ x ” from near the end. Since this book emphasizes logic, we shall try to be more explicit about the use of quantifiers.

We state here for reference the usual truth tables for $\vee, \wedge, \neg, \rightarrow, \leftrightarrow$:

Table 1: Truth Tables

φ	ψ	$\varphi \vee \psi$	$\varphi \wedge \psi$	$\varphi \rightarrow \psi$	$\varphi \leftrightarrow \psi$
T	T	T	T	T	T
T	F	T	F	F	F
F	T	T	F	T	F
F	F	F	F	T	T

φ	$\neg \varphi$
T	F
F	T

Note that in mathematics, $\varphi \rightarrow \psi$ is always equivalent to $\neg \varphi \vee \psi$. For example, $7 < 8 \rightarrow 1 + 1 = 2$ and $8 < 7 \rightarrow 1 + 1 = 2$ are both true; despite the English rendering

of “implies”, there is no “causal connection” between $7 < 8$ and the value of $1 + 1$. Also, note that “or” in mathematics is always inclusive; that is $\varphi \vee \psi$ is true if one or both of φ, ψ are true, unlike the informal English in “Stop or I’ll shoot!”.

0.3 Why Read This Book?

This book describes some basic ideas in set theory, model theory, proof theory, and recursion theory; these are all parts of what is called mathematical logic. There are three reasons one might want to read about this:

1. As an introduction to logic.
2. For its applications in topology, analysis, algebra, AI, databases.
3. Because the foundations of mathematics is relevant to philosophy.

1. If you plan to become a logician, then you will need this material to understand more advanced work in the subject.

2. Set theory is useful in any area of math dealing with uncountable sets; model theory is closely related to algebra. Questions about decidability come up frequently in math and computer science. Also, areas in computer science such as artificial intelligence and databases often use notions from model theory and proof theory.

3. The title of this book is “Foundations of Mathematics”, and there are a number of philosophical questions about this subject. Whether or not you are interested in the philosophy, it is a good way to tie together the various topics, so we’ll begin with that.

0.4 The Foundations of Mathematics

The foundations of mathematics involves the *axiomatic method*. This means that in mathematics, one writes down axioms and proves theorems from the axioms. The justification for the axioms (why they are interesting, or true in some sense, or worth studying) is part of the motivation, or physics, or philosophy, not part of the mathematics. The mathematics itself consists of logical deductions from the axioms.

Here are three examples of the axiomatic method. The first two should be known from high school or college mathematics.

Example 1: Geometry. The *use* of geometry (in measurement, construction, etc.) is prehistoric, and probably evolved independently in various cultures. The axiomatic development was first (as far as we know) developed by the ancient Greeks from 500 to 300 BC, and was described in detail by Euclid around 300 BC. In his *Elements* [12], he listed axioms and derived theorems from the axioms. We shall not list all the axioms of geometry, because they are complicated and not related to the subject of this book. One such axiom (see Book I, Postulate 1) is that any two distinct points determine a unique

line. Of course, Euclid said this in Greek, not in English, but we could also say it using logical notation, as in Section 0.2:

$$\forall x, y [[\text{Point}(x) \wedge \text{Point}(y) \wedge x \neq y] \rightarrow \exists! z [\text{Line}(z) \wedge \text{LiesOn}(x, z) \wedge \text{LiesOn}(y, z)]] \quad .$$

The intended domain of discourse, or universe, could be all geometric objects.

Example 2: Group Theory. The group *idea*, as applied to permutations and algebraic equations, dates from around 1800 (Ruffini 1799, Abel 1824, Galois 1832). The axiomatic treatment is usually attributed to Cayley (1854) (see [4], Vol 8). We shall list all the group axioms because they are simple and will provide a useful example for us as we go on. A *group* is a *model* $(G; \cdot)$ for the axioms $GP = \{\gamma_1, \gamma_2\}$:

$$\begin{aligned} \gamma_1. & \forall xyz [x \cdot (y \cdot z) = (x \cdot y) \cdot z] \\ \gamma_2. & \exists u [\forall x [x \cdot u = u \cdot x = x] \wedge \forall x \exists y [x \cdot y = y \cdot x = u]] \end{aligned}$$

Here, we're saying that G is a set and \cdot is a function from $G \times G$ into G such that γ_1 and γ_2 hold in G (with “ $\forall x$ ” meaning “for all $x \in G$ ”, so G is our universe, as discussed in Section 0.2). Axiom γ_1 is the associative law. Axiom γ_2 says that there is an identity element u , and that for every x , there is an inverse y , such that $xy = yx = u$. A more formal discussion of models and axioms will occur in Chapter II.

From the axioms, one proves theorems. For example, the group axioms imply the cancellation rule. We say: $GP \vdash \forall xyz [x \cdot y = x \cdot z \rightarrow y = z]$. This turnstile symbol “ \vdash ” is read “proves”.

This formal presentation is definitely *not* a direct quote from Cayley, who stated his axioms in English. Rather, it is influenced by the mathematical logic and set theory of the 1900s. Prior to that, axioms were stated in a natural language (e.g., Greek, English, etc.), and proofs were just given in “ordinary reasoning”; exactly what a proof *is* was not formally analyzed. This is still the case now in most of mathematics. Logical symbols are frequently used as abbreviations of English words, but most math books assume that you can recognize a correct proof when you see it, without formal analysis. However, the Foundations of Mathematics should give a precise definition of what a mathematical statement is and what a mathematical proof is, as we do in Chapter II, which covers model theory and proof theory.

This formal analysis makes a clear distinction between *syntax* and *semantics*. GP is viewed as a set of two *sentences* in *predicate logic*; this is a formal language with precise rules of formation (just like computer languages such as C or java or \TeX or html). A *formal proof* is then a finite sequence of sentences in this formal language obeying some precisely defined rules of inference – for example, the *Modus Ponens* rule (see Section II.10) says that from $\varphi \rightarrow \psi$ and φ you can infer ψ . So, the sentences of predicate logic and the formal proofs are syntactic objects. Once we have given a precise definition, it will not be hard to show (see Exercise II.11.11) that there really is a formal proof of cancellation from the axioms GP

Semantics involves meaning, or *structures*, such as groups. The syntax and semantics are related by the Completeness Theorem (see Theorem II.12.1), which says that $GP \vdash \varphi$ iff φ is true in all groups.

After the Completeness Theorem, model theory and proof theory diverge. Proof theory studies more deeply the structure of formal proofs, whereas model theory emphasizes primarily the semantics – that is, the mathematical structure of the models. For example, let G be an infinite group. Then G has a subgroup $H \subseteq G$ which is countably infinite. Also, given any cardinal number $\kappa \geq |G|$, there is a group $K \supseteq G$ of size κ . Proving these statements is an easy algebra exercises *if* you know some set theory, which you will after reading Chapter I.

These statements are part of model theory, not group theory, because they are special cases of the Löwenheim–Skolem–Tarski Theorem (see Theorems II.16.4 and II.16.5), which applies to models of arbitrary theories. You can also get H, K to satisfy all the first-order properties true in G . For example if G is non-abelian, then H, K will be also. Likewise for other properties, such as “abelian” or “3-divisible” ($\forall x \exists y (yyy = x)$). The proof, along with the definition of “first-order”, is part of model theory (Chapter II), but the proof uses facts about cardinal numbers from set theory, which brings us to the third example:

Example 3: Set Theory. For infinite sets, the basic work was done by Cantor in the 1880s and 1890s, although the idea of sets — especially finite ones, occurred much earlier. This is our first topic, so you will soon see a lot about uncountable cardinal numbers. Cantor just worked naively, not axiomatically, although he was aware that naive reasoning could lead to contradictions. The first axiomatic approach was due to Zermelo (1908), and was improved later by Fraenkel and von Neumann, leading to the current system *ZFC* (see Section I.2), which is now considered to be the “standard” axioms for set theory.

A philosophical remark: In model theory, *every* list of sentences in formal logic forms the axioms for some (maybe uninteresting) axiomatic theory, but informally, there are two different uses to the word “axioms”: as “*statements of faith*” and as “*definitional axioms*”. The first use is closest to the dictionary definition of an axiom as a “truism” or a “statement that needs no proof because its truth is obvious”. The second use is common in algebra, where one speaks of the “axioms” for groups, rings, fields, etc.

Consider our three examples:

Example 1 (Classical Greek view): these are *statements of faith* — that is, they are obviously true facts about real physical space, from which one may then derive other true but non-obvious facts, so that by studying Euclidean geometry, one is studying the structure of the real world. The intended universe is fixed – it could be thought of as all geometric objects in the physical universe. Of course, Plato pointed out that “perfect” lines, triangles, etc. only exist in some abstract idealization of the universe, but no one doubted that the results of Euclidean geometry could be safely applied to solve real-world problems.

Example 2 (Everyone’s view): these are *definitional* axioms. The axioms do not capture any deep “universal truth”; they only serve to define a useful class of structure. Groups occur naturally in many areas of mathematics, so one might as well encapsulate their properties and prove theorems about them. Group theory is the study of groups in general, not one specific group, and the intended domain of discourse is the particular group under discussion.

This view of Example 2 has never changed since the subject was first studied, but our view of geometry has evolved. First of all, as Einstein pointed out, the Euclidean axioms are false in real physical space, and will yield incorrect results when applied to real-world problems. Furthermore, most modern uses of geometry are not axiomatic. We define 3-dimensional space as \mathbb{R}^3 , and we discuss various metrics (notions of distance) on it, including the Euclidean metric, which approximately (but not exactly) corresponds to reality. Thus, in the modern view, geometry is the study of geometries, not one specific geometry, and the Euclidean axioms have been downgraded to mere definitional axioms — one way of describing a specific (flat) geometry.

Example 3 (Classical (mid 1900s) view): these are statements of faith. *ZFC* is the theory of everything (see Section I.4). Modern mathematics might seem to be a mess of various axiom systems: groups, rings, fields, geometries, vector spaces, etc., etc. This is all subsumed within set theory, as we’ll see in Chapter I. So, we postulate once and for all these *ZFC* axioms. Then, from these axioms, there are no further assumptions; we just make definitions and prove theorems. Working in *ZFC*, we say that a group is a *set* G together with a product on it satisfying γ_1, γ_2 . The product operation is really a function of two variables defined on G , but a function is also a special kind of *set* — namely, a set of ordered pairs. If you want to study geometry, you would want to know that a metric space is a *set* X , together with some distance function d on it satisfying some well-known properties. The distances, $d(x, y)$, are real numbers. The real numbers form the specific *set* \mathbb{R} , constructed within *ZFC* by a set-theoretic procedure which we shall describe later (see Definition I.15.4).

We study set theory first because it is the foundation of everything. Also, the discussion will produce some technical results on infinite cardinalities which are useful in a number of the more abstract areas of mathematics. In particular, these results are needed for the model theory in Chapter II; they are also important in analysis and topology and algebra, as you will see from various exercises in this book. In Chapter I, we shall state the axioms precisely, but the proofs will be informal, as they are in most math texts. When we get to Chapter II, we shall look at formal proofs from various axiom systems, and *GP* and *ZFC* will be interesting specific examples.

The *ZFC* axioms are listed in Section I.2. The list is rather long, but by the end of Chapter I, you should understand the meaning of each axiom and why it is important. Chapter I will also make some brief remarks on the interrelationships between the axioms; further details on this are covered in texts in set theory, such as [18, 20]. These interrelationships are not so simple, since *ZFC* does not settle everything of interest. Most notably, *ZFC* doesn’t determine the truth of the Continuum Hypotheses, *CH*.

This is the assertion that every uncountable subset of \mathbb{R} has the same size as \mathbb{R} .

Example 3 (Modern view): these are definitional axioms. Set theory is the study of models of *ZFC*. There are, for example, models in which $2^{\aleph_0} = \aleph_5$; this means that there are exactly four infinite cardinalities, called $\aleph_1, \aleph_2, \aleph_3, \aleph_4$, strictly between countable and the size of \mathbb{R} . By the end of Chapter I, you will understand exactly what *CH* and \aleph_n mean, but the models will only be hinted at.

Chapter III covers *recursion theory*, or the theory of algorithms and computability. Since most people have used a computer, the informal notion of *algorithm* is well-known to the general public. The following sets are clearly decidable, in that you can write a program which tests for them in your favorite programming language (assuming this language is something reasonable, like C or java or python):

1. The set of primes.
2. The set of axioms of *ZFC*.
3. The set of valid C programs.

That is, if you are not concerned with efficiency, you can easily write a program which inputs a number or symbolic expression and tells you whether or not it's a member of one of these sets. For (1), you input an integer $x > 1$ and check to see if it is divisible by any of the integers y with $x > y > 1$. For (2), you input a finite symbolic expression and see if it is among the axiom types listed in Section I.2. Task (3) is somewhat harder, and you would have to refer to the C manual for the precise definition of the language, but a C compiler accomplishes task (3), among many other things.

Deeper results involve proving that certain sets which are *not* decidable, such as the following:

4. The set of C programs which halt (say, with all values of their input).
5. $\{\varphi : ZFC \vdash \varphi\}$.

That is, there is *no* program which reads a sentences φ in the language of set theory and tells you whether or not $ZFC \vdash \varphi$. Informally, “mathematical truth is not decidable”. Certainly, results of this form are relevant to the foundations of mathematics. Chapter III will also be an introduction to understanding the meaning of some more advanced results along this line, which are not proved in this book. Such results are relevant to many areas of mathematics. For example, $\{\varphi : GP \vdash \varphi\}$ is not decidable, whereas $\{\varphi : AGP \vdash \varphi\}$ is decidable, where *AGP* is the axioms for abelian groups. The proofs involve a lot of group theory. Likewise, the solvability of diophantine equations (algebraic equations over \mathbb{Z}) is undecidable; this proof involves a lot of number theory. Also, in topology, simple connectivity is undecidable. That is, there's no algorithm which inputs a polyhedron (presented, for example, as a finite simplicial complex) and tells you whether or not it's simply connected. This proof involves some elementary facts about the fundamental group in topology, plus the knowledge that the word problem for groups is undecidable.

This book only touches on the basics of recursion theory, but we shall give a precise definition of “decidable” and explain its relevance to set theory and model theory.

Chapter I

Set Theory

I.1 Plan

We shall discuss the axioms, explain their meaning in English, and show that from these axioms, you can derive all of mathematics. Of course, this chapter does not contain all of mathematics. Rather, it shows how you can develop, from the axioms of set theory, basic concepts, such as the concept of number and function and cardinality. Once this is done, the rest of mathematics proceeds as it does in standard mathematics texts.

In addition to basic concepts, we describe how to compute with infinite cardinalities, such as $\aleph_0, \aleph_1, \aleph_2, \dots$.

I.2 The Axioms

For reference, we list the axioms right off, although they will not all make sense until the end of this chapter. We work in predicate logic with binary relations $=$ and \in .

Informally, our universe is the class of all *hereditary sets* x ; that is, x is a set, all elements of x are sets, all elements of elements of x are sets, and so forth. In this (Zermelo-Fraenkel style) formulation of the axioms, proper classes (such as our domain of discourse) do not exist. Further comments on the intended domain of discourse will be made in Sections I.6 and I.14.

Formally, of course, we are just exhibiting a list of sentences in predicate logic. Axioms stated with free variables are understood to be universally quantified.

Axiom 0. Set Existence.

$$\exists x(x = x)$$

Axiom 1. Extensionality.

$$\forall z(z \in x \leftrightarrow z \in y) \rightarrow x = y$$

Axiom 2. Foundation.

$$\exists y(y \in x) \rightarrow \exists y(y \in x \wedge \neg \exists z(z \in x \wedge z \in y))$$

Axiom 3. Comprehension Scheme. For each formula, φ , without y free,

$$\exists y \forall x(x \in y \leftrightarrow x \in z \wedge \varphi(x))$$

Axiom 4. Pairing.

$$\exists z(x \in z \wedge y \in z)$$

Axiom 5. Union.

$$\exists A \forall Y \forall x(x \in Y \wedge Y \in \mathcal{F} \rightarrow x \in A)$$

Axiom 6. Replacement Scheme. For each formula, φ , without B free,

$$\forall x \in A \exists! y \varphi(x, y) \rightarrow \exists B \forall x \in A \exists y \in B \varphi(x, y)$$

The rest of the axioms are a little easier to state using some defined notions. On the basis of Axioms 1,3,4,5, define \subseteq (subset), \emptyset (or 0; empty set), S (ordinal successor function), \cap (intersection), and $\text{SING}(x)$ (x is a singleton) by:

$$\begin{aligned} x \subseteq y &\iff \forall z(z \in x \rightarrow z \in y) \\ x = \emptyset &\iff \forall z(z \notin x) \\ y = S(x) &\iff \forall z(z \in y \leftrightarrow z \in x \vee z = x) \\ w = x \cap y &\iff \forall z(z \in w \leftrightarrow z \in x \wedge z \in y) \\ \text{SING}(x) &\iff \exists y \in x \forall z \in x(z = y) \end{aligned}$$

Axiom 7. Infinity.

$$\exists x(\emptyset \in x \wedge \forall y \in x(S(y) \in x))$$

Axiom 8. Power Set.

$$\exists y \forall z(z \subseteq x \rightarrow z \in y)$$

Axiom 9. Choice.

$$\emptyset \notin F \wedge \forall x \in F \forall y \in F(x \neq y \rightarrow x \cap y = \emptyset) \rightarrow \exists C \forall x \in F(\text{SING}(C \cap x))$$

- ✦ $ZFC =$ Axioms 1–9. $ZF =$ Axioms 1–8.
- ✦ ZC and Z are ZFC and ZF , respectively, with Axiom 6 (Replacement) deleted.
- ✦ Z^- , ZF^- , ZC^- , ZFC^- are Z , ZF , ZC , ZFC , respectively, with Axiom 2 (Foundation) deleted.

Most of elementary mathematics takes place within ZC^- (approximately, Zermelo's theory). The Replacement Axiom allows you to build sets of size \aleph_ω and bigger. It also lets you represent well-orderings by von Neumann ordinals, which is notationally useful, although not strictly necessary.

Foundation says that \in is well-founded – that is, every non-empty set x has an \in -minimal element y . This rules out, e.g., sets a, b such that $a \in b \in a$. Foundation is never needed in the development of mathematics.

Logical formulas with defined notions are viewed as abbreviations (or macros) for formulas in $\in, =$ only. In the case of defined predicates, such as \subseteq , the macro is expanded by replacing the predicate by its definition (changing the names of variables as necessary), so that the Power Set Axiom abbreviates:

$$\forall x \exists y \forall z ((\forall v (v \in z \rightarrow v \in x)) \rightarrow z \in y) .$$

In the case of defined functions, one must introduce additional quantifiers; the Axiom of Infinity above abbreviates

$$\exists x \left(\exists u (\forall v (v \notin u) \wedge u \in x) \wedge \forall y \in x \exists u (\forall z (z \in u \leftrightarrow z \in y \vee z = y) \wedge u \in x) \right) .$$

Here, we have replaced the “ $S(y) \in x$ ” by “ $\exists u (\psi(y, u) \wedge u \in x)$ ”, where ψ says that u satisfying the property of being equal to $S(y)$.

We follow the usual convention in modern algebra and logic that basic facts about $=$ are logical facts, and need not be stated when axiomatizing a theory. So, for example, the converse to Extensionality, $x = y \rightarrow \forall z (z \in x \leftrightarrow z \in y)$, is true by logic – equivalently, true of all binary relations, not just \in . Likewise, when we wrote down the axioms for groups in Section 0.4, we just listed the axioms γ_1, γ_2 which are specific to the product function. We did not list statements such as $\forall xyz (x = y \rightarrow x \cdot z = y \cdot z)$; this is a fact about $=$ which is true of any binary function.

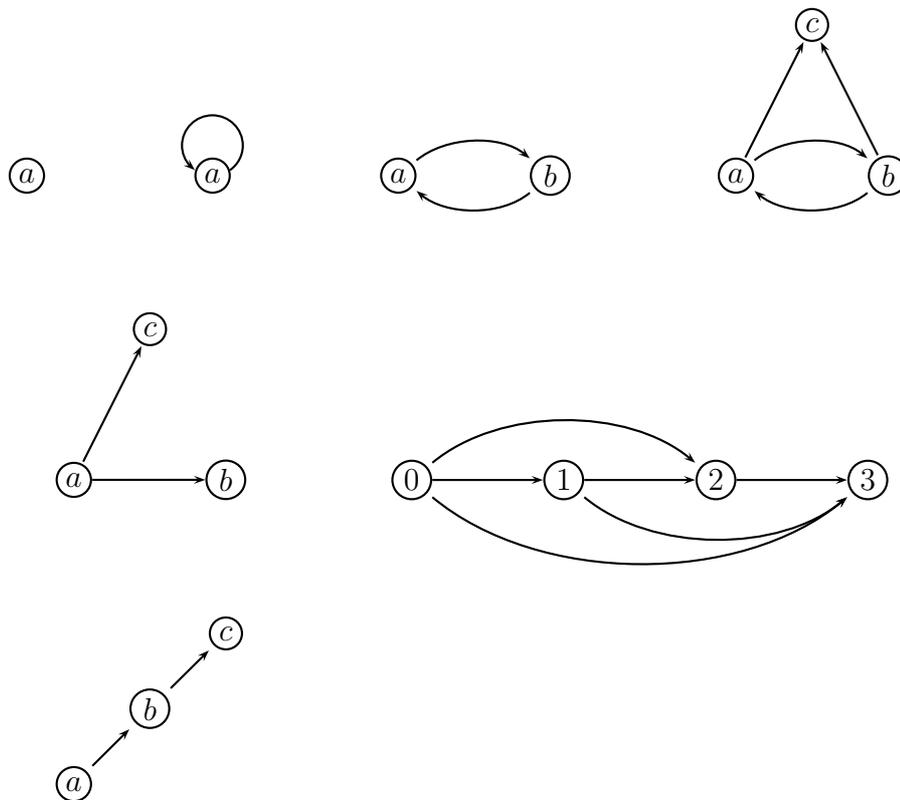
In most treatments of formal logic (see Chapter II; especially Remark II.8.16), the statement that the universe is non-empty (i.e., $\exists x (x = x)$) is also taken to be a logical fact, but we have listed this explicitly as Axiom 0 to avoid possible confusion, since many mathematical definitions do allow empty structures (e.g., the empty topological space).

It is possible to ignore the “intended” interpretation of the axioms and just view them as making assertions about a binary relation on some non-empty domain. This point of view is useful in seeing whether one axiom implies another. For example, #2 of the following exercise shows that Axiom 2 does not follow from Axioms 1,4,5.

Exercise I.2.1 Which of Axioms 1,2,4,5 are true of the binary relation E on the domain D , in the following examples?

1. $D = \{a\}; E = \emptyset$.
2. $D = \{a\}; E = \{(a, a)\}$.

3. $D = \{a, b\}; E = \{(a, b), (b, a)\}$.
4. $D = \{a, b, c\}; E = \{(a, b), (b, a), (a, c), (b, c)\}$.
5. $D = \{a, b, c\}; E = \{(a, b), (a, c)\}$.
6. $D = \{0, 1, 2, 3\}; E = \{(0, 1), (0, 2), (0, 3), (1, 2), (1, 3), (2, 3)\}$.
7. $D = \{a, b, c\}; E = \{(a, b), (b, c)\}$.



Hint. It is often useful to picture E as a digraph (directed graph). The *children* of a node y are the nodes x such that there is an arrow from x to y . For example, the children of node c in #4 are a and b . Then some of the axioms may be checked visually. Extensionality says that you never have two distinct nodes, x, y , with exactly the same children (as one has with b and c in #5). Pairing says that given any nodes x, y (possibly the same), there is a node z with arrows from x and y into z . One can see at sight that this is true in #2 and false in the rest of the examples. \square

Throughout this chapter, we shall often find it useful to think of membership as a digraph, where the children of a set are the members of the set. The simple finite models presented in Exercise I.2.1 are primarily curiosities, but general method of models (now using infinite ones) is behind all independence proofs in set theory; for example, there

are models of all of ZFC in which the Continuum Hypothesis (CH) is true, and other models in which CH is false; see [18, 20].

I.3 Two Remarks on Presentation.

Remark I.3.1 In discussing any axiomatic development — of set theory, of geometry, or whatever — be careful to distinguish between the:

- Formal discussion: definitions, theorems, proofs.
- Informal discussion: motivation, pictures, philosophy.

The informal discussion helps you understand the theorems and proofs, but is not strictly necessary, and is not part of the mathematics. In most mathematics texts, including this one, the informal discussion and the proving of theorems are interleaved. In this text, the formal discussion starts in the middle of Section I.6.

Remark I.3.2 Since we're discussing *foundations*, the presentation of set theory will be a bit different than the presentation in most texts. In a beginning book on group theory or calculus, it's usually assumed that you know nothing at all about the subject, so you start from the basics and work your way up through the middle-level results, and then to the most advanced material at the end. However, since set theory is fundamental to all of mathematics, you already know all the middle-level material; for example, you know that \mathbb{R} is infinite and $\{7, 8, 9\}$ is a finite set of size 3. The focus will thus be on the really basic material and the advanced material. The basic material involves discussing the meaning of the axioms, and explaining, on the basis of the axioms, what exactly *are* 3 and \mathbb{R} . The advanced material includes properties of uncountable sets; for example, the fact that \mathbb{R} , the plane $\mathbb{R} \times \mathbb{R}$, and countably infinite dimensional space $\mathbb{R}^{\mathbb{N}}$ all have the same size. When doing the basics, we shall use examples from the middle-level material for motivation. For example, one can illustrate properties of functions by using the real-valued functions you learn about in calculus. These illustrations should be considered part of the informal discussion of Remark I.3.1. Once you see how to derive elementary calculus from the axioms, these illustrations could then be re-interpreted as part of the formal discussion.

I.4 Set theory is the theory of everything

First, as part of the motivation, we begin by explaining why set theory is the foundation of mathematics. Presumably, you know that set theory is important. You may not know that set theory is *all*-important. That is

- *All* abstract mathematical concepts are set-theoretic.
- *All* concrete mathematical objects are specific sets.

Abstract concepts all reduce to set theory. For example, consider the notion of function. Informally, a function gives you a specific way of corresponding y 's to x 's (e.g., the function $y = x^2 + 1$), but to make a precise definition of “function”, we identify a function with its graph. That is, f is a function iff f is a set, all of whose elements are ordered pairs, such that $\forall x, y, z[(x, y) \in f \wedge (x, z) \in f \rightarrow y = z]$. This is a precise definition if you know what an ordered pair is.

Informally, (x, y) is a “thing” uniquely associated with x and y . Formally, we define $(x, y) = \{\{x\}, \{x, y\}\}$. In the axiomatic approach, you have to verify that the axioms let you construct such a set, and that $ZFC \vdash$ “unique”, that is (see Exercise I.6.14):

$$ZFC \vdash \forall x, y, x', y'[(x, y) = (x', y') \rightarrow x = x' \wedge y = y'] .$$

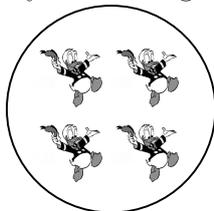
Then, a group is a special kind of pair (G, \cdot) where \cdot is a function from $G \times G$ into G . $G \times G$ is the set of all ordered pairs from G . So, we explain everything in terms of sets, pairs, and functions, but it all reduces to sets. We shall see this in more detail later, as we develop the axioms.

An example of concrete object is a specific number, like: $0, 1, 2, -2/3, \pi, e^{2i}$. Informally, 2 denotes the concept of twoness – containing a thing and another thing and that's all. We could define $two(x) \leftrightarrow \exists y, z[x = \{y, z\} \wedge y \neq z]$, but we want the official object 2 to be a specific set, not a logical formula. Formally, we pick a specific set with two distinct things in it and call that 2. Following von Neumann, let $2 = \{0, 1\}$. Of course, you need to know that 0 is a specific set with zero elements – i.e. $0 = \emptyset$ and $1 = \{0\} = \{\emptyset\}$ (a specific set with one element), which makes $2 = \{\emptyset, \{\emptyset\}\}$. You can now verify $two(2)$.

The set of all natural numbers is denoted by $\mathbb{N} = \omega = \{0, 1, 2, 3, \dots\}$. Of course, $3 = \{0, 1, 2\}$. Once you have ω , it is straightforward to define the sets $\mathbb{Z}, \mathbb{Q}, \mathbb{R}, \mathbb{C}$. The construction of $\mathbb{Z}, \mathbb{Q}, \mathbb{R}, \mathbb{C}$ will be outlined briefly in Section I.15, with the details left to analysis texts. You probably already know that \mathbb{Q} is countable, while \mathbb{R} and \mathbb{C} are not.

I.5 Counting

Basic to understanding set theory is learning how to count. To count a set \mathcal{D} of ducks,



you first have to get your ducks in a row, and then you pair them off against the natural numbers until you run out of ducks:

I.6 Extensionality, Comprehension, Pairing, Union

We begin by discussing these four axioms and deriving some elementary results from them.

Axiom 1. Extensionality:

$$\forall x, y [\forall z (z \in x \leftrightarrow z \in y) \rightarrow x = y] .$$

As it says in Section I.2, all free variables are implicitly quantified universally. This means that although the axiom is listed as just $\forall z (z \in x \leftrightarrow z \in y) \rightarrow x = y$, the intent is to assert this statement for all x, y .

Informal discussion: This says that a set is determined by its members, so that if x, y are two sets with exactly the same members, then x, y are the same set. Extensionality also says something about our intended domain of discourse, or universe, which is usually called V (see Figure I.1, page 18). Everything in our universe must be a set, since if we allowed objects x, y which aren't sets, such as a duck (D) and a badger (B), then they would have no members, so that we would have

$$\forall z [z \in B \leftrightarrow z \in D \leftrightarrow z \in \emptyset \leftrightarrow \text{FALSE}] ,$$

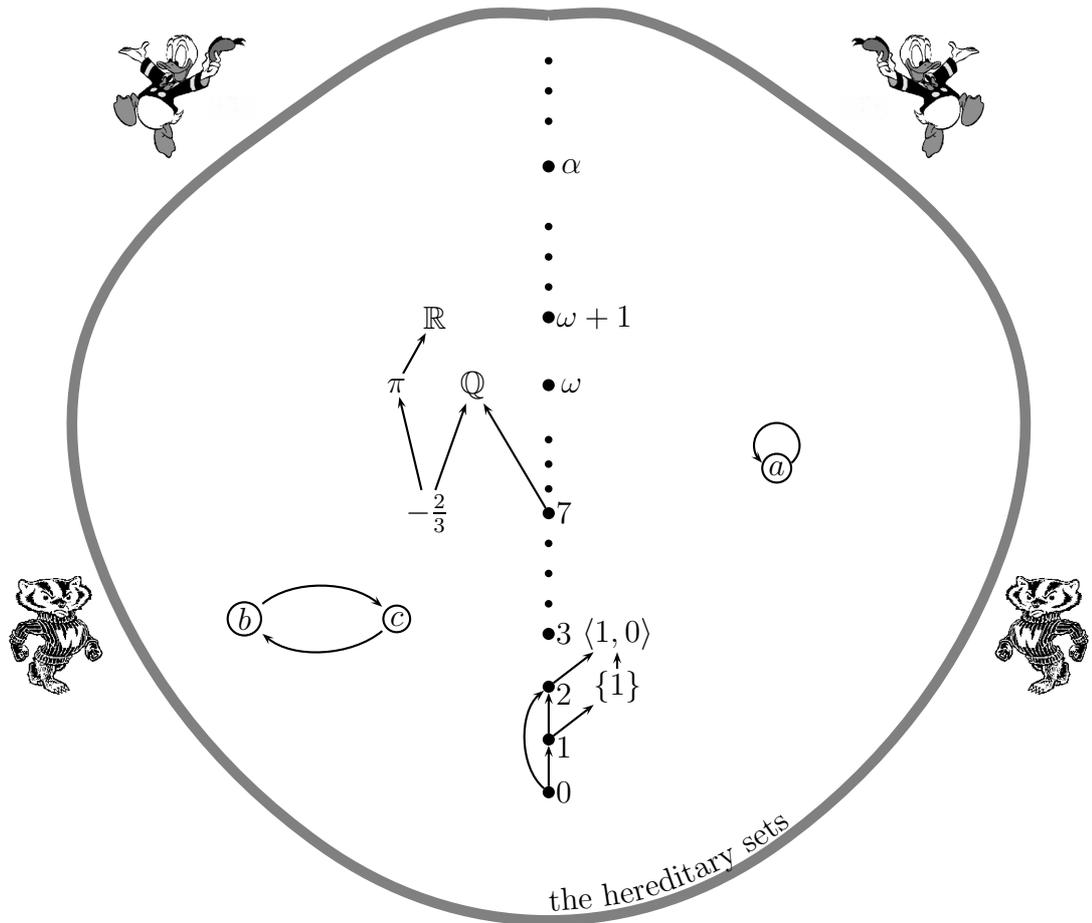
whereas B, D, \emptyset are all different objects. So, physical objects, such as B, D , are not part of our universe.

Now, informally, one often thinks of sets or collections of physical objects, such as $\{B, D\}$, or a set of ducks (see Section I.5), or the set of animals in a zoo. However, these sets are also not in our mathematical universe. Recall (see Section 0.2) that in writing logical expressions, it is understood that the variables range only over our universe, so that a statement such as " $\forall z \dots$ " is an abbreviation for "for all z in our universe \dots ". So, if we allowed $\{B\}$ and $\{D\}$ into our universe, then $\forall z (z \in \{B\} \leftrightarrow z \in \{D\})$ would be true (since B, D are not in our universe), whereas $\{B\} \neq \{D\}$.

More generally, if x, y are (sets) in our universe, then all their elements are also in our universe, so that the hypothesis " $\forall z (z \in x \leftrightarrow z \in y)$ " really means that x, y are sets with exactly the same members, so that Extensionality is justified in concluding that $x = y$. So, if x is in our universe, then x must not only be a set, but all elements of x , all elements of elements of x , etc. must be sets. We say that x is *hereditarily* a set (if we think of the members of x as the children of x , as in Exercise I.2.1, then we are saying that x and all its descendents are sets). Examples of such hereditary sets are the numbers 0, 1, 2, 3 discussed earlier.

So, the quantified variables in the axioms of set theory are intended to range over the stuff inside the blob in Figure I.1 — the hereditary sets. The arrows denote membership (\in). Not all arrows are shown. Extensionality doesn't exclude sets a, b, c such that $a = \{a\}$ and $c \in b \in c$. These sets are excluded by the Foundation Axiom (see Section I.14), which implies that the universe is neatly arrayed in levels, with all arrows sloping up. Figure I.1 gives a picture of the universe under ZF^- , which is set theory without the Foundation Axiom.

Figure I.1: The Set-Theoretic Universe in ZF^-



Formal discussion: We have our first theorem. There is at most one empty set:

Definition I.6.1 $\text{emp}(x)$ iff $\forall z(z \notin x)$.

Then the Axiom of Extensionality implies:

Theorem I.6.2 $\text{emp}(x) \wedge \text{emp}(y) \rightarrow x = y$.

Now, to prove that there *is* an empty set, you can't do it just by Extensionality, since Extensionality is consistent with $\neg[\exists x \text{emp}(x)]$:

Exercise I.6.3 Look through the models in Exercise I.2.1, and find the ones satisfying Extensionality plus $\neg[\exists x \text{emp}(x)]$.

The usual proof that there is an empty set uses the Comprehension Axiom. As a first approximation, we may postulate:

Garbage I.6.4 *The Naive Comprehension Axiom (NCA) is the assertion: For every property $Q(x)$ of sets, the set $S = \{x : Q(x)\}$ exists.*

In particular, $Q(x)$ can be something like $x \neq x$, which is always false, so that S will be empty.

Unfortunately, there are two problems with NCA:

1. It's vague.
2. It's inconsistent.

Regarding Problem 2: Cantor knew that his set theory was inconsistent, and that you could get into trouble with sets which are too big. His contradiction (see Section I.11) was a bit technical, and uses the notion of cardinality, which we haven't discussed yet. However, Russell (1901) pointed out that one could rephrase Cantor's paradox to get a really simple contradiction, directly from NCA alone:

Paradox I.6.5 (Russell) *Applying NCA, define $R = \{x : x \notin x\}$. Then $R \in R \leftrightarrow R \notin R$, a contradiction.*

Cantor's advice was to avoid inconsistent sets (see [10]). This avoidance was incorporated into Zermelo's statement of the Comprehension Axiom as it is in Section I.2. Namely, once you *have* a set z , you can form $\{x \in z : Q(x)\}$. You can still form $R = R_z = \{x \in z : x \notin x\}$, but this only implies that *if* $R \in z$, then $R \in R \leftrightarrow R \notin R$, which means that $R \notin z$ — that is,

Theorem I.6.6 *There is no universal set: $\forall z \exists R [R \notin z]$.*

So, although we talk informally about the universe, V , it's not really an object of study in our universe of set theory.

Regarding Problem 1: What is a property? Say we have constructed ω , the set of natural numbers. In mathematics, we should not expect to form $\{n \in \omega : n \text{ is stupid}\}$. On a slightly more sophisticated level, so-called “paradoxes” arise from allowing ill-formed definitions. A well-known one is:

Let n be the least number which cannot be defined using forty words of less. But I've just defined n in forty words of less.

Here, $Q(x)$ says that x can be defined in 40 English words or less, and we try to form $E = \{n \in \omega : Q(n)\}$, which is finite, since there are only finitely many possible definitions. Then the least $n \notin E$ is contradictory.

To avoid Problem 1, we say that a property is something defined by a *logical formula*, as described in Section 0.2 — that is, an expression made up using \in , $=$, propositional

(boolean) connectives (and, or, not, etc.), and variables and quantifiers \forall, \exists . So, our principle now becomes: For each logical formula φ , we assert:

$$\forall z[\exists y\forall x(x \in y \leftrightarrow x \in z \wedge \varphi(x))] .$$

Note that our proof of Theorem I.6.6 was correct, with $\varphi(x)$ the formula $x \notin x$. With a different φ , we get an empty set:

Definition I.6.7 \emptyset denotes the (unique) y such that $\text{emp}(y)$ (i.e., $\forall x[x \notin y]$).

Justification. To prove that $\exists y[\text{emp}(y)]$, start with any set z (there is one by Axiom 0) and apply Comprehension with $\varphi(x)$ a statement which is always false (for example, $x \neq x$) to get a y such that $\forall x(x \in y \leftrightarrow \text{FALSE})$ — i.e., $\forall x(x \notin y)$. By Theorem I.6.2, there is at most one empty set, so $\exists!y \text{emp}(y)$, so we can name this unique object \emptyset . \square

As usual in mathematics, before giving a name to an object satisfying some property (e.g., $\sqrt{2}$ is the unique $y > 0$ such that $y^2 = 2$), we must prove that that property really is held by a unique object.

In applying Comprehension (along with other axioms), it is often a bit awkward to refer back to the statement of the axiom, as we did in justifying Definition I.6.7. It will be simpler to introduce some notation:

Notation I.6.8 For any formula $\varphi(x)$:

- \Rightarrow If there is a set A such that $\forall x[x \in A \leftrightarrow \varphi(x)]$, then A is unique by Extensionality, and we denote this set by $\{x : \varphi(x)\}$, and we say that $\{x : \varphi(x)\}$ exists.
- \Rightarrow If there is no such set, then we say that $\{x : \varphi(x)\}$ doesn't exist, or forms a proper class.
- \Rightarrow $\{x \in z : \varphi(x)\}$ abbreviates $\{x : x \in z \wedge \varphi(x)\}$.

Comprehension asserts that sets of the form $\{x : x \in z \wedge \varphi(x)\}$ always exist. We have just seen that the empty set, $\emptyset = \{x : x \neq x\}$, does exist, whereas a universal set, $\{x : x = x\}$, doesn't exist. It's sometimes convenient to “think” about this collection and give it the name V , called the *universal class*, which is then a proper class. We shall say more about proper classes later, when we have some more useful examples (see Notation I.8.4). For now, just remember that the assertion “ $\{x : x = x\}$ doesn't exist” is simply another way of saying that $\neg\exists A\forall x[x \in A]$. This is just a notational convention; there is no philosophical problem here, such as “how can we talk about it if it doesn't exist?”. Likewise, there is no *logical* problem with asserting “Trolls don't exist”, and there is no problem with thinking about trolls, whether or not you believe in them.

Three further remarks on Comprehension:

1. Some elementary use of logic is needed even to *state* the axioms, since we need the notion of “formula”. However, we're not using logic yet for formal proofs. Once the axioms are stated, the proofs in this chapter will be informal, as in most of mathematics.

2. The Comprehension Axiom is really infinite scheme; we have one assertion for each logical formula φ .

3. The terminology $\varphi(x)$ just emphasizes the dependence of φ on x , but φ can have other free variables, for example, when we define intersection and set difference:

Definition I.6.9 Given z, u :

$$\Rightarrow z \cap u := \{x \in z : x \in u\}.$$

$$\Rightarrow z \setminus u := \{x \in z : x \notin u\}.$$

Here, φ is, respectively, $x \in u$ and $x \notin u$. For $z \cap u$, the actual instance of Comprehension used is: $\forall u \forall z \exists y \forall x (x \in y \leftrightarrow x \in z \wedge x \in u)$. To form $z \cup u$, which can be bigger than z and u , we need another axiom, the Union Axiom, discussed below.

Exercise I.6.10 Look through the models in Exercise I.2.1, and find one which satisfies Extensionality and Comprehension but doesn't have pairwise unions — that is, the model will contain elements z, u with no w satisfying $\forall x [x \in w \leftrightarrow x \in z \vee x \in u]$.

You have certainly seen \cup and \cap before, along with their basic properties; our emphasis is how to derive what you already know from the axioms. So, for example, you know that $z \cap u = u \cap z$; this is easily proved from the definition of \cap (which implies that $\forall x [x \in z \cap u \leftrightarrow x \in u \cap z]$) plus the Axiom of Extensionality. Likewise, $z \cap u \subseteq z$; this is easily proved from the definition of \cap and \subseteq :

Definition I.6.11 $y \subseteq z \iff \forall x (x \in y \rightarrow x \in z)$.

In Comprehension, φ can even have z free — for example, it's legitimate to form $z^* = \{x \in z : \exists u (x \in u \wedge u \in z)\}$; so once we have officially defined 2 as $\{0, 1\}$, we'll have $2^* = \{0, 1\}^* = \{0\}$.

The proviso in Section I.2 that φ cannot have y free avoids self-referential definitions such as the Liar Paradox “This statement is false”: — that is

$$\exists y \forall x (x \in y \leftrightarrow x \in z \wedge x \notin y) \quad .$$

is contradictory if z is any non-empty set.

Note that \emptyset is the only set whose existence we have actually demonstrated, and Extensionality and Comprehension alone do not let us prove the existence of any non-empty set:

Exercise I.6.12 Look through the models in Exercise I.2.1, and find one which satisfies Extensionality and Comprehension plus $\forall x \neg \exists y [y \in x]$.

We can construct non-empty sets using Pairing:

$$\forall x, y \exists z (x \in z \wedge y \in z) .$$

As stated, z could contain other elements as well, but given z , we can always form $u = \{w \in z : w = x \vee w = y\}$. So, u is the (unique by Extensionality) set which contains x, y and nothing else. This justifies the following definition of unordered and ordered pairs:

Definition I.6.13

- ☛ $\{x, y\} = \{w : w = x \vee w = y\}$.
- ☛ $\{x\} = \{x, x\}$.
- ☛ $\langle x, y \rangle = (x, y) = \{\{x\}, \{x, y\}\}$.

The key fact about ordered pairs is:

Exercise I.6.14

$$\langle x, y \rangle = \langle x', y' \rangle \rightarrow x = x' \wedge y = y' .$$

Hint. Split into cases: $x = y$ and $x \neq y$. □

There are many other definitions of ordered pair which satisfy this exercise. In most of mathematics, it does not matter which definition was used – it is only important that x and y are determined uniquely from their ordered pair; in these cases, it is conventional to use the notation (x, y) . We shall use $\langle x, y \rangle$ when it is relevant that we are using this specific definition.

We can now begin to count:

Definition I.6.15

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

Exercise I.6.16 $\langle 0, 1 \rangle = \{1, 2\}$, and $\langle 1, 0 \rangle = \{\{1\}, 2\}$.

The axioms so far let us generate infinitely many different sets: $\emptyset, \{\emptyset\}, \{\{\emptyset\}\}, \dots$ but we can't get any with more than two elements. To do that we use the Union Axiom. That will let us form $3 = 2 \cup \{2\} = \{0, 1, 2\}$. One could postulate an axiom which says that $x \cup y$ exists for all x, y , but looking ahead, the usual statement of the Union Axiom will justify infinite unions as well:

$$\forall \mathcal{F} \exists A \forall Y \forall x [x \in Y \wedge Y \in \mathcal{F} \rightarrow x \in A] .$$

That is, for any \mathcal{F} (in our universe), view \mathcal{F} as a family of sets. This axiom gives us a set A which contains all the members of members of \mathcal{F} . We can now take the union of the sets in this family:

Definition I.6.17

$$\bigcup \mathcal{F} = \bigcup_{Y \in \mathcal{F}} Y = \{x : \exists Y \in \mathcal{F}(x \in Y)\}$$

That is, the members of $\bigcup \mathcal{F}$ are the members of members of \mathcal{F} . As usual when writing $\{x : \dots\dots\}$, we must justify the existence of this set.

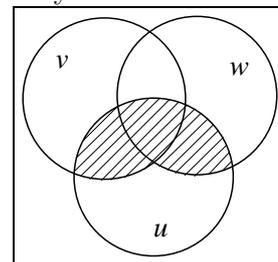
Justification. Let A be as in the Union Axiom, and apply Comprehension to form $B = \{x \in A : \exists Y \in \mathcal{F}(x \in Y)\}$. Then $x \in B \leftrightarrow \exists Y \in \mathcal{F}(x \in Y)$. \square

Definition I.6.18 $u \cup v = \bigcup\{u, v\}$, $\{x, y, z\} = \{x, y\} \cup \{z\}$, and $\{x, y, z, t\} = \{x, y\} \cup \{z, t\}$.

You already know the basic facts about these notions, and, in line with Remark I.3.2, we shall not actually write out proofs of all these facts from the axioms. As two samples:

Exercise I.6.19 Prove that $\{x, y, z\} = \{x, z, y\}$ and $u \cap (v \cup w) = (u \cap v) \cup (u \cap w)$.

These can easily be derived from the definitions, using Extensionality. For the second one, note that an informal proof using a Venn diagram can be viewed as a shorthand for a rigorous proof by cases. For example, to prove that $x \in u \cap (v \cup w) \leftrightarrow x \in (u \cap v) \cup (u \cap w)$ for all x , you can consider the eight possible cases: $x \in u, x \in v, x \in w$, $x \in u, x \in v, x \notin w$, $x \in u, x \notin v, x \in w$, etc. In each case, you verify that the left and right sides of the “ \leftrightarrow ” are either both true or both false. To summarize this longwinded proof in a picture, draw the standard Venn diagram of three sets, which breaks the plane into eight regions, and note that if you shade $u \cap (v \cup w)$ or if you shade $(u \cap v) \cup (u \cap w)$ you get the same picture. The shaded set consists of the regions for which the left and right sides of the “ \leftrightarrow ” are both true.



We can also define the intersection of a family of sets; as with pairwise intersection, this is justified directly from the Comprehension Axiom and requires no additional axiom:

Definition I.6.20 When $\mathcal{F} \neq \emptyset$,

$$\bigcap \mathcal{F} = \bigcap_{Y \in \mathcal{F}} Y = \{x : \forall Y \in \mathcal{F}(x \in Y)\}$$

Justification. Fix $E \in \mathcal{F}$, and form $\{x \in E : \forall Y \in \mathcal{F}(x \in Y)\}$. \square

Note that $\bigcup \emptyset = \emptyset$, while $\bigcap \emptyset$ would be the universal class, V , which doesn't exist.

We can now count a little further:

Definition I.6.21 The ordinal successor function, $S(x)$, is $x \cup \{x\}$. Then define:

$$\begin{aligned} 3 &= S(2) = \{0, 1, 2\} \\ 4 &= S(3) = \{0, 1, 2, 3\} \\ 5 &= S(4) = \{0, 1, 2, 3, 4\} && \text{etc. etc. etc.} \\ 6 &= S(5) = \{0, 1, 2, 3, 4, 5\} \\ 7 &= S(6) = \{0, 1, 2, 3, 4, 5, 6\} \end{aligned}$$

But, what does “etc” mean? *Informally*, we can define a natural number, or finite ordinal, to be any set obtained by applying S to 0 a finite number of times. Now, define $\mathbb{N} = \omega$ to be the set of all natural numbers. Note that each $n \in \omega$ is the set of all natural numbers $< n$. The natural numbers are what we count with in elementary school.

Note the “informally”. To understand this formally, you need to understand what a “finite number of times” means. Of course, this means “ n times, for some $n \in \omega$ ”, which is fine if you know what ω is. So, the whole thing is circular. We shall break the circularity in Section I.8 by formalizing the properties of the order relation on ω , but we need first a bit more about the theory of relations (in particular, orderings and well-orderings) and functions, which we shall cover in Section I.7. Once the ordinals are defined formally, it is not hard to show (see Exercise I.8.11) that if x is an ordinal, then $S(x)$ really is its successor — that is, the next larger ordinal.

If we keep counting past all the natural numbers, we hit the first infinite ordinal. Since each ordinal is the set of all smaller ordinals, this first infinite ordinal is ω , the set of all natural numbers. The next ordinal is $S(\omega) = \omega \cup \{\omega\} = \{0, 1, 2, \dots, \omega\}$, and $S(S(\omega)) = \{0, 1, 2, \dots, \omega, S(\omega)\}$. We then have to explain how to add and multiply ordinals. Not surprisingly, $S(S(\omega)) = \omega + 2$, so that we shall count:

$$0, 1, 2, 3, \dots, \omega, \omega + 1, \omega + 2, \dots, \omega + \omega = \omega \cdot 2, \omega \cdot 2 + 1, \dots$$

The idea of counting into the transfinite is due to Cantor. This specific representation of the ordinals is due to von Neumann.

I.7 Relations, Functions, Discrete Mathematics

You’ve undoubtedly used relations (such as orderings) and functions in mathematics, but we must explain them within our framework of axiomatic set theory. Subsection I.7.1 contains the basic facts and definitions, and Subsection I.7.3 contains the theory of well-orders, which will be needed when ordinals are discussed in Section I.8.

I.7.1 Basics

Definition I.7.1 R is a (binary) relation iff R is a set of ordered pairs — that is,

$$\forall u \in R \exists x, y [u = \langle x, y \rangle] .$$

xRy abbreviates $\langle x, y \rangle \in R$ and $x \not R y$ abbreviates $\langle x, y \rangle \notin R$.

Of course, the abbreviations xRy and $x \not R y$ are meaningful even when R isn't a relation. For reference, we collect some commonly used properties of relations:

Definition I.7.2

- R is transitive on A iff $\forall xyz \in A [xRy \wedge yRz \rightarrow xRz]$.
- R is irreflexive on A iff $\forall x \in A [x \not R x]$
- R is reflexive on A iff $\forall x \in A [xRx]$
- R satisfies trichotomy on A iff $\forall xy \in A [xRy \vee yRx \vee x = y]$.
- R is symmetric on A iff $\forall xy \in A [xRy \leftrightarrow yRx]$.
- R partially orders A strictly iff R is transitive and irreflexive on A .
- R totally orders A strictly iff R is transitive and irreflexive on A and satisfies trichotomy on A .
- R is an equivalence relation on A iff R is reflexive, symmetric, and transitive on A .

For example, $<$ on \mathbb{Q} is a (strict) total order but \leq isn't. For the work we do here, it will be more convenient to take the strict $<$ as the basic order notion and consider $x \leq y$ to abbreviate $x < y \vee x = y$.

Note that Definition I.7.2 is meaningful for any sets R, A . When we use it, R will usually be a relation, but we do not require that R contain only ordered pairs from A . For example, we can partially order $\mathbb{Q} \times \mathbb{Q}$ coordinatewise: $(x_1, x_2)R(y_1, y_2)$ iff $x_1 < y_1$ and $x_2 < y_2$. This does not satisfy trichotomy, since $(2, 3) \not R (3, 2)$ and $(3, 2) \not R (2, 3)$. However, if we restrict the order to a line or curve with positive slope, then R does satisfy trichotomy, so we can say that R totally orders $\{(x, 2x) : x \in \mathbb{Q}\}$.

Of course, these "examples" are not official yet, since we must first construct \mathbb{Q} and $\mathbb{Q} \times \mathbb{Q}$. Unofficially still, any subset of $\mathbb{Q} \times \mathbb{Q}$ is a relation, and if you project it on the x and y coordinates, you will get its domain and range:

Definition I.7.3 For any set R , define:

$$\text{dom}(R) = \{x : \exists y[\langle x, y \rangle \in R]\} \quad \text{ran}(R) = \{y : \exists x[\langle x, y \rangle \in R]\} \quad .$$

Justification. To see that these sets exist, observe that if $\langle x, y \rangle = \{\{x\}, \{x, y\}\} \in R$, then $\{x\}$ and $\{x, y\}$ are in $\bigcup R$, and then $x, y \in \bigcup\bigcup R$. Then, by Comprehension, we can form:

$$\{x \in \bigcup\bigcup R : \exists y[\langle x, y \rangle \in R]\} \quad \text{and} \quad \{y \in \bigcup\bigcup R : \exists x[\langle x, y \rangle \in R]\} \quad .$$

□

For a proof of the existence of $\text{dom}(R)$ and $\text{ran}(R)$ which does not rely on our specific definition (I.6.13) of ordered pair, see Exercise I.7.16.

Definition I.7.4 $R \upharpoonright A = \{\langle x, y \rangle \in R : x \in A\}$.

This \upharpoonright (restriction) is most often used for functions.

Definition I.7.5 R is a function iff R is a relation and for every $x \in \text{dom}(R)$, there is a unique y such that $\langle x, y \rangle \in R$. In this case, $R(x)$ denotes that unique y .

We can then make the usual definitions of “injection”, “surjection”, and “bijection”:

Definition I.7.6

- ☛ $F : A \rightarrow B$ means that F is a function, $\text{dom}(F) = A$, and $\text{ran}(F) \subseteq B$.
- ☛ $F : A \xrightarrow{\text{onto}} B$ or $F : A \rightarrow B$ means that $F : A \rightarrow B$ and $\text{ran}(F) = B$ (F is a surjection or maps A onto B).
- ☛ $F : A \xrightarrow{1-1} B$ or $F : A \hookrightarrow B$ means that $F : A \rightarrow B$ and $\forall x, x' \in A [F(x) = F(x') \rightarrow x = x']$ (F is an injection or maps A 1-1 into B).
- ☛ $F : A \xrightarrow[1-1]{\text{onto}} B$ or $F : A \xleftrightarrow{1-1} B$ means that both $F : A \xrightarrow{1-1} B$ and $F : A \xrightarrow{\text{onto}} B$. (F is a bijection from A onto B).

For example, with the sine function on the real numbers, the following are all true statements:

$$\begin{aligned} \sin & : \mathbb{R} \rightarrow \mathbb{R} \\ \sin & : \mathbb{R} \rightarrow [-1, 1] \\ \sin & : \mathbb{R} \xrightarrow{\text{onto}} [-1, 1] \\ \sin \upharpoonright [-\pi/2, \pi/2] & : [-\pi/2, \pi/2] \xrightarrow{1-1} \mathbb{R} \\ \sin \upharpoonright [-\pi/2, \pi/2] & : [-\pi/2, \pi/2] \xrightarrow[1-1]{\text{onto}} [-1, 1] \end{aligned}$$

Definition I.7.7 $F(A) = F \text{ ``} A = \text{ran}(F \upharpoonright A)$.

In most applications of this, F is a function. The $F(A)$ terminology is the more common one in mathematics; for example, we say that $\sin([0, \pi/2]) = [0, 1]$ and $\sin(\pi/2) = 1$; this never causes confusion because $\pi/2$ is a number, while $[0, \pi/2]$ is a set of numbers. However, the $F(A)$ could be ambiguous when A may be both a member of and a subset of the domain of F ; in these situations, we use $F \text{ ``} A$. For example, if $\text{dom}(F) = 3 = \{0, 1, 2\}$, we use $F(2)$ for the value of F with input 2 and $F \text{ ``} 2$ for $\{F(0), F(1)\}$.

The axioms discussed so far don't allow us to build many relations and functions. You might think that we could start from sets S, T and then define lots of relations as subsets of the *cartesian product* $S \times T = \{\langle s, t \rangle : s \in S \wedge t \in T\}$, but we need another axiom to prove that $S \times T$ exists. Actually, you can't even prove that $\{0\} \times T$ exists

with the axioms given so far, although “obviously” you should be able to write down this set as $\{\langle 0, x \rangle : x \in T\}$, and it “should” have the same size as T . Following Fraenkel (1922) (the “ F ” in *ZFC*), we justify such a collection by the Replacement Axiom:

$$\forall x \in A \exists! y \varphi(x, y) \rightarrow \exists B \forall x \in A \exists y \in B \varphi(x, y)$$

That is, suppose that for each $x \in A$, there is a unique object y such that $\varphi(x, y)$. Call this unique object y_x . Then we “should” be able to form the set $C = \{y_x : x \in A\}$. Replacement, plus Comprehension, says that indeed we can form it by letting $C = \{y \in B : \exists x \in A \varphi(x, y)\}$.

Definition I.7.8 $S \times T = \{\langle s, t \rangle : s \in S \wedge t \in T\}$.

Justification. This definition is just shorthand for the longer

$$S \times T = \{x : \exists s \in S \exists t \in T [x = \langle s, t \rangle]\} ,$$

and as usual with this $\{x : \dots\dots\}$ notation, we must prove that this set really exists. To do so, we use Replacement twice:

First, fix $s \in S$, and form $\{s\} \times T = \{\langle s, x \rangle : x \in T\}$ by applying Replacement (along with Comprehension), as described above, with $A = T$ and $\varphi(x, y)$ the formula which says that $y = \langle s, x \rangle$.

Next, form $D = \{\{s\} \times T : s \in S\}$ by applying Replacement (along with Comprehension), as described above, with $A = S$ and $\varphi(x, y)$ the formula which says that $y = \{x\} \times T$. Then $\bigcup D = \bigcup_{s \in S} \{s\} \times T$ contains exactly all pairs $\langle s, t \rangle$ with $s \in S$ and $t \in T$. \square

Replacement is used to justify the following common way of defining functions:

Lemma I.7.9 *Suppose $\forall x \in A \exists! y \varphi(x, y)$. Then there is a function f with $\text{dom}(f) = A$ such that for each $x \in A$, $f(x)$ is the unique y such that $\varphi(x, y)$.*

Proof. Fix B as in the Replacement Axiom, and let $f = \{\langle x, y \rangle \in A \times B : \varphi(x, y)\}$. \square

For example, for any set A , we have a function f such that $\text{dom}(f) = A$ and $f(x) = \{\{x\}\}$ for all $x \in A$.

Cartesian products are used frequently in mathematics – for example, once we have \mathbb{R} , we form the plane, $\mathbb{R} \times \mathbb{R}$. Also, a two-variable function from X to Y is really a function $f : X \times X \rightarrow Y$; so $f \subseteq (X \times X) \times Y$. It also lets us define inverse of a relation and the composition of functions:

Definition I.7.10 $R^{-1} = \{\langle y, x \rangle : \langle x, y \rangle \in R\}$.

Justification. This is a defined subset of $\text{ran}(R) \times \text{dom}(R)$. \square

If f is a function, then f^{-1} is not a function unless f is 1-1. The \sin^{-1} or arcsin function in trigonometry is really the function $(\sin \upharpoonright [-\pi/2, \pi/2])^{-1}$.

Definition I.7.11 $G \circ F = \{\langle x, z \rangle \in \text{dom}(F) \times \text{ran}(G) : \exists y[\langle x, y \rangle \in F \wedge \langle y, z \rangle \in G]\}$.

In the case where F, G are functions with $\text{ran}(F) \subseteq \text{dom}(G)$, we are simply saying that $(G \circ F)(x) = G(F(x))$.

If S and T are ordered sets, we can order their cartesian product $S \times T$ *lexicographically* (i.e., as in the dictionary). That is, we can view elements of $S \times T$ as two-letter words; then, to compare two words, you use their first letters, unless they are the same, in which case you use their second letters:

Definition I.7.12 If $<$ and \prec are relations, then their lexicographic product on $S \times T$ is the relation \triangleleft on $S \times T$ defined by:

$$\langle s, t \rangle \triangleleft \langle s', t' \rangle \leftrightarrow [s < s' \vee [s = s' \wedge t \prec t']] .$$

Exercise I.7.13 If $<$ and \prec are strict total orders of S, T , respectively, then their lexicographic product on $S \times T$ is a strict total order of $S \times T$.

Finally, we have the notion of isomorphism:

Definition I.7.14 F is an isomorphism from $(A; <)$ onto $(B; \triangleleft)$ iff $F : A \xrightarrow[\text{onto}]{1-1} B$ and $\forall x, y \in A [x < y \leftrightarrow F(x) \triangleleft F(y)]$. Then, $(A; <)$ and $(B; \triangleleft)$ are isomorphic (in symbols, $(A; <) \cong (B; \triangleleft)$) iff there exists an isomorphism from $(A; <)$ onto $(B; \triangleleft)$.

This definition makes sense regardless of whether $<$ and \triangleleft are orderings, but for now, we plan to use it just for order relations. It is actually a special case of the general notion of isomorphism between arbitrary algebraic structures used in model theory (see Definition II.8.18).

The Replacement Axiom justifies the usual definition in mathematics of a quotient of a structure by an equivalence relation (see Definition I.7.2):

Definition I.7.15 Let R be an equivalence relation on a set A . For $x \in A$, let $[x] = \{y \in A : yRx\}$; $[x]$ is called the equivalence class of x . Let $A/R = \{[x] : x \in A\}$.

Here, forming $[x]$ just requires the Comprehension Axiom, but to justify forming A/R , the set of equivalence classes, we can let f be the function with domain A such that $f(x) = [x]$ (applying Lemma I.7.9), and then set $A/R = \text{ran}(f)$. In most applications, A has some additional structure on it (e.g, it is a group, or a topological space), and one defines the appropriate structure on the set A/R . This is discussed in books on group theory and topology. For a use of quotients in model theory, see Definition II.12.9.

A similar use of Replacement gives us a proof that $\text{dom}(R)$ and $\text{ran}(R)$ exist (see Definition I.7.3) which does not depend on the specific set-theoretic definition of (x, y) .

Exercise I.7.16 Say we've defined a "pair" $[(x, y)]$ in some way, and assume that we can prove $[(x, y)] = [(x', y')] \rightarrow x = x' \wedge y = y'$. Prove that $\{x : \exists y[[(x, y)] \in R]\}$ and $\{y : \exists x[[(x, y)] \in R]\}$ exist for all sets R .

Exercise I.7.17 The class of all groups, $(G; \cdot)$, is a proper class.

Hint. If it were a set, you could get V by elementary set operations. □

I.7.2 Foundational Remarks

1. Set theory is the theory of everything, but that doesn't mean that you could understand this (or any other) presentation of axiomatic set theory if you knew absolutely nothing. You don't need any knowledge about infinite sets; you could learn about these as the axioms are being developed; but you do need to have some basic understanding of finite combinatorics even to understand what statements are and are not axioms. For example, we have assumed that you can understand our explanation that an instance of the Comprehension Axiom is obtained by replacing the φ in the Comprehension Scheme in Section I.2 by a logical formula. To understand what a logical formula is (as discussed briefly in Section 0.2 and defined more precisely in Section II.5) you need to understand what "finite" means and what finite strings of symbols are.

This basic *finitistic reasoning*, which we do not analyze formally, is called the *metatheory*. In the metatheory, we explain various notions such as what a formula is and which formulas are axioms of our *formal theory*, which here is ZFC.

2. The informal notions of "relation" and "function" receive two distinct representations in the development of set theory: as sets, which are objects of the formal theory, and as abbreviations in the metatheory.

First, consider relations. We have already defined a relation to be a set of ordered pairs, so a relation is a specific kind of set, and we handle these sets within the formal theory *ZFC*.

Now, one often speaks informally of \in , $=$, and \subseteq as "relations", but these are not relations in the above sense – they are a different kind of animal. For example, the subset "relation", $S = \{p : \exists x, y[p = \langle x, y \rangle \wedge x \subseteq y]\}$ doesn't exist — i.e., it forms a proper class, in the terminology of Notation I.6.8 (S cannot exist because $\text{dom}(S)$ would be the universal class V , which doesn't exist). Rather, we view the symbol \subseteq as an abbreviation in the metatheory; that is, $x \subseteq y$ is an abbreviation for $\forall z(z \in x \rightarrow z \in y)$. Likewise, the isomorphism "relation" \cong is not a set of ordered pairs; rather, the notation $(A; <) \cong (B; \triangleleft)$ was introduced in Definition I.7.14 as an abbreviation for a more complicated statement. Of course, the membership and equality "relations", \in and $=$, are already basic symbols in the language of set theory.

Note, however, that many definitions of properties of relations, such as those in Definition I.7.2, make sense also for these "*pseudorelations*", since we can just plug the pseudorelation into the definition. For example, we can say that \in totally orders the set $3 = \{0, 1, 2\}$; the statement that it is transitive on 3 just abbreviates: $\forall xyz \in 3[x \in y \wedge y \in z \rightarrow x \in z]$. However, \subseteq is not a (strict) total order on 3 because irreflexivity fails. It even makes sense to say that \subseteq is transitive on the universe, V , as an abbreviation for the (true) statement $\forall xyz[x \subseteq y \wedge y \subseteq z \rightarrow x \subseteq z]$. Likewise, \in is *not* transitive on V because $0 \in \{0\}$ and $\{0\} \in \{\{0\}\}$ but $0 \notin \{\{0\}\}$. Likewise, it makes sense to assert:

Exercise I.7.18 \cong is an equivalence relation.

Hint. To prove transitivity, we take isomorphisms F from $(A; \triangleleft_1)$ to $(B; \triangleleft_2)$ and G from $(B; \triangleleft_2)$ to $(C; \triangleleft_3)$ and compose them to get an isomorphism $G \circ F : A \rightarrow C$. \square

Note that this discussion of what abbreviates what takes place in the metatheory. For example, in the metatheory, we unwind the statement of Exercise I.7.18 to a statement just involving sets, which is then to be proved from the axioms of *ZFC*.

A similar discussion holds for functions, which are special kinds of relations. For example, $f = \{(1, 2), (2, 2), (3, 1)\}$ is a function, with $\text{dom}(f) = \{1, 2, 3\}$ and $\text{ran}(f) = \{1, 2\}$; this f is an object formally defined within *ZFC*. Informally, $\bigcup : V \rightarrow V$ is also a function, but V doesn't exist, and likewise we can't form the set of ordered pairs $\bigcup = \{\langle \mathcal{F}, \bigcup \mathcal{F} \rangle : \mathcal{F} \in V\}$ (if we could, then $\text{dom}(\bigcup)$ would be V). Rather, we have a formula $\varphi(\mathcal{F}, Z)$ expressing the statement that Z is the union of all the elements of \mathcal{F} ; $\varphi(\mathcal{F}, Z)$ is

$$\forall x[x \in Z \leftrightarrow \exists Y \in \mathcal{F}[x \in Y]] \quad ,$$

in line with Definition I.6.17. We prove $\forall \mathcal{F} \exists! Z \varphi(\mathcal{F}, Z)$, and then we use $\bigcup \mathcal{F}$ to “denote” that Z ; formally, the “denote” means that a statement such as $\bigcup \mathcal{F} \in w$ abbreviates $\exists Z(\varphi(\mathcal{F}, Z) \wedge Z \in w)$. See Section II.15 for a more formal discussion of the status of defined notions in an axiomatic theory.

As with relations, elementary properties of function make sense when applied to such “*pseudofunctions*”. For example, we can say that “ \bigcup is not 1-1”; this just abbreviates the formula $\exists x_1, x_2, y [\varphi(x_1, y) \wedge \varphi(x_2, y) \wedge x_1 \neq x_2]$.

In the 1700s and 1800s, as real analysis was being developed, there were debates about exactly what a function is (see [22, 23]). There was no problem with specific real-valued functions defined by formulas, such as $f(x) = x^2 + 3x$; $f'(x) = 2x + 3$. However, as more and more abstract examples were developed, such as continuous functions which were nowhere differentiable, there were questions about exactly what sort of rules suffice to define a function.

By the early 1900s, with the development of set theory and logic, the “set of ordered pairs” notion of function became universally accepted. Now, a clear distinction is made between the notion of an arbitrary real-valued function and one which is definable (a model-theory notion; see Chapter II), or computable (a recursion-theory notion; see Chapter III). Still, now in the 2000s, elementary calculus texts (see [30], p. 37) often confuse the issue by defining a function to be some sort of “rule” which associates y 's to x 's. This is very misleading, since you can only write down countably many rules, but there are uncountably many real-valued functions. In analysis, one occasionally uses $\mathbb{R}^{\mathbb{R}}$, the set of all functions from \mathbb{R} to \mathbb{R} ; more frequently one uses the subset $C(\mathbb{R}, \mathbb{R}) \subset \mathbb{R}^{\mathbb{R}}$ consisting of all continuous functions. Both $C(\mathbb{R}, \mathbb{R})$ and $\mathbb{R}^{\mathbb{R}}$ are uncountable (of sizes 2^{\aleph_0} and $2^{2^{\aleph_0}}$, respectively; see Exercise I.15.8).

However, the “rule” concept survives when we talk about an operation defined on all sets, such as $\bigcup : V \rightarrow V$. Here, since V and functions on V do not really exist, the only way to make sense of such notions is to consider each explicit rule (i.e., formula) which defines one set as a function of another, as a way of introducing abbreviations in the

metatheory. An explicit example of this occurs already in Section I.2, where we defined the successor “function” by writing an explicit formula to express “ $y = S(x)$ ”; we then explained how to rewrite the Axiom of Infinity, which was originally expressed with the symbol “ S ”, into a statement using only the basic symbols “ \in ” and “ $=$ ”.

Lemma I.7.9 says that if we have any “rule” function and restrict it to a set A , then we get a “set-of-ordered-pairs” function. For example, for any set A , we can always form the set $\bigcup \upharpoonright A = \{(x, y) : x \in A \wedge y = \bigcup x\}$.

3. One should distinguish between individual sentences and schemes (or rules) in the metatheory. Each axiom of *ZFC* other than Comprehension and Replacement forms (an abbreviation of) one sentence in the language of set theory. But the Comprehension Axiom is a rule in the metatheory for producing axioms; that is whenever you replace the φ in the Comprehension Scheme in Section I.2 by a logical formula, you get an axiom of *ZFC*; so really *ZFC* is an infinite list of axioms. Likewise, the Replacement Axiom is really an infinite scheme.

A similar remark holds for theorems. Lemma I.7.9 is really a theorem scheme; formally, for each formula $\varphi(x, y)$, we can prove in *ZFC* the theorem:

$$\forall x \in A \exists! y \varphi(x, y) \rightarrow \exists f [f \text{ is a function} \wedge \text{dom}(f) = A \wedge \forall x \in A \varphi(x, f(x))] .$$

But Exercise I.7.18 is just one theorem; that is, it is (the abbreviation of) one sentence which is provable from *ZFC*.

I.7.3 Well-orderings

Definition I.7.19 $y \in X$ is R -minimal in X iff

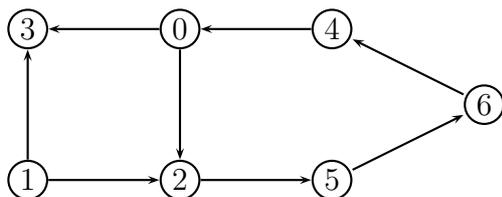
$$\neg \exists z (z \in X \wedge z R y) ,$$

and R -maximal in X iff

$$\neg \exists z (z \in X \wedge y R z) .$$

R is well-founded on A iff for all non-empty $X \subseteq A$, there is a $y \in X$ which is R -minimal in X .

This notion occurs very frequently, so we give a picture and some examples of it. In the case that A is finite, we can view R as a directed graph, represented by arrows.



If $R = \{(0, 2), (0, 3), (1, 2), (1, 3), (2, 5), (4, 0), (5, 6), (6, 4)\}$ and $X = 4 = \{0, 1, 2, 3\}$, then 0, 1 are both R -minimal in X (they have no arrows into them from elements of X), and 2, 3 are both R -maximal in X (they have no arrows from them into elements of X). It is easily seen that R is well-founded on $6 = \{0, 1, 2, 3, 4, 5\}$, but not well-founded on 7 because the cycle $C = \{0, 2, 5, 6, 4\} \subseteq 7$ has no R -minimal element.

Exercise I.7.20 *If A is finite, then R is well-founded on A iff R is acyclic (has no cycles) on A .*

This exercise must be taken to be informal for now, since we have not yet officially defined “finite”; the notion of “cycle” ($a_0 R a_1 R a_2 R \cdots R a_n R a_0$) also involves the notion of “finite”. This exercise fails for infinite sets. If R is any strict total order relation, it is certainly acyclic, but need not be well-founded; for example the usual $<$ on \mathbb{Q} is not well-founded because \mathbb{Q} itself has no $<$ -minimal element.

Definition I.7.19 will occur again in our discussion of Zorn’s Lemma (Section I.12) and the Axiom of Foundation (Section I.14), but for now, we concentrate on the case of well-founded total orders:

Definition I.7.21 *R well-orders A iff R totally orders A strictly and R is well-founded on A .*

Note that if R is a total order, then X can have at most one minimal element, which is then the least element of X ; so a well-order is a strict total order in which every non-empty subset has a least element. As usual, this is definition makes sense when R is a set-of-ordered-pairs relation, as well as when R is a pseudorelation such as \in .

Informally, the usual order (which is \in) on ω is a well-order. Well-foundedness just expresses the least number principle. You’ve seen this used in proofs by induction: To prove $\forall n \in \omega \varphi(n)$, you let $X = \{n \in \omega : \neg\varphi(n)\}$, assume $X \neq \emptyset$, and derive a contradiction from the least element of X (the first place where the theorem fails).

Formally, well-order is part of the definition of ordinal, so it will be true of ω by definition, but we require some more work from the axioms to prove that ω exists. So far, we’ve only defined numbers up to 7 (see Definition I.6.21).

First, some informal examples of well-ordered subsets of \mathbb{Q} and \mathbb{R} , using your knowledge of these sets. These will become formal (official) examples as soon as we’ve defined \mathbb{Q} and \mathbb{R} in *ZFC*.

Don’t confuse the “least element” in well-order with the greatest lower bound from calculus. For example, $[0, 1] \subseteq \mathbb{R}$ isn’t well-ordered; $(0, 1)$ has a greatest lower bound, $\inf(0, 1) = 0$, but no least element.

As mentioned, the “least number principle” says that \mathbb{N} is well-ordered. Hence, so is every subset of it by:

Exercise I.7.22 *If R well-orders A and $X \subseteq A$, then R well-orders X .*

Example I.7.23 Informally, you can use the natural numbers to begin counting any well-ordered set A . \emptyset is well-ordered; but if A is non-empty, then A has a least element a_0 . Next, if $A \neq \{a_0\}$, then $A \setminus \{a_0\}$ has a least element, a_1 . If $A \neq \{a_0, a_1\}$, then $A \setminus \{a_0, a_1\}$ has a least element, a_2 . Continuing this, we see that if A is infinite, it will begin with an ω -sequence, a_0, a_1, a_2, \dots . If this does not exhaust A , then $A \setminus \{a_n : n \in \omega\}$ will have a least element, which we may call a_ω , the ω^{th} element of A . We may continue this process of listing the elements of A until we have exhausted A , putting A in 1-1 correspondence with some initial segment of the ordinals. The fact that this informal process works is proved formally in Theorem I.8.19.

To display subsets of \mathbb{Q} well-ordered in types longer than ω , we can start with a copy of \mathbb{N} compressed into the interval $(0, 1)$. Let $A_0 = \{1 - 2^{-n} : n \in \omega\} \subseteq (0, 1)$: this is well-ordered (in type ω). Room has been left on top, so that we could form $A_0 \cup \{1.5, 1.75\}$, whose ω^{th} element, a_ω , is 1.5, and the next (and last) element is $a_{\omega+1} = 1.75$. Continuing this, we may add a whole block of ω new elements above A_0 inside $(1, 2)$. Let $A_k = \{k + 1 - 2^{-n} : n \in \omega\} \subseteq (k, k + 1)$, for $k \in \mathbb{N}$. Then $A_0 \cup A_1 \subseteq (0, 2)$ is well-ordered in type $\omega + \omega = \omega \cdot 2$, and $\bigcup_{k \in \omega} A_k \subseteq \mathbb{Q}$ is well-ordered in type $\omega^2 = \omega \cdot \omega$. Now, \mathbb{Q} is isomorphic to $\mathbb{Q} \cap (0, 1)$, so that we may also find a well-order of type ω^2 inside $(0, 1)$, and then add new rationals above that. Continuing this, every countable well-ordering is embeddable into \mathbb{Q} (see Exercise I.11.32).

Exercise I.7.24 If $<$ and \prec are well-orders of S, T , respectively, then their lexicographic product on $S \times T$ is a well-order of $S \times T$; see also Exercise I.7.13.

I.8 Ordinals

Now we break the circularity mentioned at the end of §I.6:

Definition I.8.1 z is a transitive set iff $\forall y \in z [y \subseteq z]$.

Definition I.8.2 z is a (von Neumann) ordinal iff z is a transitive set and z is well-ordered by \in .

Some remarks on “transitive set”: To first approximation, think of this as completely unrelated to the “transitive” appearing in the definition of total order.

Unlike properties of orderings, the notion of a transitive set doesn’t usually occur in elementary mathematics. For example, is \mathbb{R} a transitive set? If $y \in \mathbb{R}$, you think of y and \mathbb{R} as different types of objects, and you probably don’t think of y as a set at all, so you never even ask whether $y \subseteq \mathbb{R}$. But, now, since everything is a set, this is a meaningful question, although still a bit unnatural for \mathbb{R} (it will be false if we define \mathbb{R} by Definition I.15.4). However, this question does occur naturally when dealing with the natural numbers because, by our definitions of them, their elements are also natural

numbers. For example, $3 = \{0, 1, 2\}$, is a transitive set – its elements are all subsets of it (e.g., $2 = \{0, 1\} \subseteq 3$). Then, 3 is an ordinal because it is a transitive set and well-ordered by \in . $z = \{1, 2, 3\}$ is not an ordinal – although it is ordered by \in in the same *order type*, it isn't transitive – since $1 \in z$ but $1 \not\subseteq z$ (since $0 \in 1$ but $0 \notin z$). We shall see (Lemma I.8.18) that two distinct ordinals cannot be isomorphic.

One can check directly that $0, 1, 2, 3, 4, 5, 6, 7$ are indeed ordinals by our definition. It is more efficient to note that $0 = \emptyset$ is trivially an ordinal, and that the successor of an ordinal is an ordinal (Exercise I.8.11), which is easier to prove after a few preliminaries.

First, we remark on the connection between “transitive set” (Definition I.8.1) and “transitive relation” (Definition I.7.2). \in is not a transitive relation on the universe – that is, $\forall xyz[x \in y \wedge y \in z \rightarrow x \in z]$ is false (let $x = 0$, $y = \{0\}$, and $z = \{\{0\}\}$). But if you *fix* a specific z , then the statement $\forall xy[x \in y \wedge y \in z \rightarrow x \in z]$, which simply asserts that z is a transitive set, may or may not be true. You might call z a “point of transitivity” of \in .

Up to now, our informal examples have implicitly assumed that the ordinals themselves are ordered in a transfinite sequence. The order relation is exactly the membership relation, since each ordinal is the set of all smaller ordinals. We shall now make this informal concept into a theorem (Theorem I.8.5). Because of this theorem, the following notation is convenient:

Notation I.8.3 *When discussing ordinals, Greek letters (especially $\alpha, \beta, \gamma, \delta, \zeta, \eta, \xi, \mu$) “range over the ordinals”; that is, $\forall \alpha \varphi(\alpha)$ abbreviates*

$$\forall x[x \text{ is an ordinal} \rightarrow \varphi(x)] \quad .$$

Also, $\alpha < \beta$ means (by definition) $\alpha \in \beta$, and $\alpha \leq \beta$ means $\alpha \in \beta \vee \alpha = \beta$.

Informally we define the class of all ordinals to be $ON = \{x : x \text{ is an ordinal}\}$. We shall see (Theorem I.8.9) that ON is a proper class – that is it doesn't exist (see Notation I.6.8); it is a collection which is “too large” to be a set. However, the informal concept of ON yields some useful terminology:

Notation I.8.4

- ✧ $x \in ON$ abbreviates “ x is an ordinal”.
- ✧ $x \subseteq ON$ abbreviates “ $\forall y \in x[y \text{ is an ordinal}]$ ”.
- ✧ $x \cap ON$ abbreviates “ $\{y \in x : y \text{ is an ordinal}\}$ ”.

The next theorem says that ON is well-ordered by \in . Since ON doesn't really exist, we list in the theorem precisely what we are asserting:

Theorem I.8.5 *ON is well-ordered by \in . That is:*

1. \in is transitive on the ordinals: $\forall \alpha \beta \gamma[\alpha < \beta \wedge \beta < \gamma \rightarrow \alpha < \gamma]$.

2. \in is irreflexive on the ordinals: $\forall \alpha[\alpha \notin \alpha]$.
3. \in satisfies trichotomy on the ordinals: $\forall \alpha \beta[\alpha < \beta \vee \beta < \alpha \vee \alpha = \beta]$.
4. \in is well-founded on the ordinals: every non-empty set of ordinals has an \in -least member.

Before we prove the theorem, we establish three lemmas. First,

Lemma I.8.6 *ON is a transitive class. That is, if $\alpha \in ON$ and $z \in \alpha$, then $z \in ON$.*

Proof. α is a transitive set, so $z \subseteq \alpha$. Since \in well-orders α , it well-orders every subset of α (Exercise I.7.22), so \in well-orders z , so we need only show that z is a transitive set, that is, $x \in y \in z \rightarrow x \in z$. Since $z \subseteq \alpha$, we have $y \in \alpha$, so, $y \subseteq \alpha$, and hence $x \in \alpha$. But now that $x, y, z \in \alpha$, we have $x \in y \in z \rightarrow x \in z$ because the \in relation is transitive on α (this is part of the definition of “ordinal”). \square

Because of this lemma, we shall usually use Greek letters for members of α .

Lemma I.8.7 *For all ordinals α, β : $\alpha \cap \beta$ is an ordinal.*

Proof. $\alpha \cap \beta \subseteq \alpha$, so it is well-ordered by \in (see Exercise I.7.22), and $\alpha \cap \beta$ is a transitive set because the intersection of transitive sets is transitive (this is clear from the definition). \square

Once Theorem I.8.5 is proved, it will be easy to see that $\alpha \cap \beta$ is actually the smaller of α, β (see Exercise I.8.10).

On the ordinals, $<$ is \in by definition. Our third lemma says that \leq is \subseteq ; this is obvious from Theorem I.8.5, since each ordinal is the set of all smaller ordinals, but proving it directly, as a lemma to Theorem I.8.5, takes a bit more work:

Lemma I.8.8 *For all ordinals α, β : $\alpha \subseteq \beta \leftrightarrow \alpha \in \beta \vee \alpha = \beta$.*

Proof. For \leftarrow : Just use the fact that β is transitive:

For \rightarrow : Assume $\alpha \subseteq \beta$ and $\alpha \neq \beta$. We show that $\alpha \in \beta$. Let $X = \beta \setminus \alpha$. Then $X \neq \emptyset$, so let ξ be the \in -least member of X . Then $\xi \in \beta$, so we are done if we show that $\xi = \alpha$:

If $\mu \in \xi$, then $\mu \in \beta$ (since β is transitive) and $\mu \notin X$ (since ξ was \in -least), so $\mu \in \alpha$. Hence, $\xi \subseteq \alpha$.

Now, assume that $\xi \subsetneq \alpha$. Fix $\mu \in \alpha \setminus \xi$. Then $\mu, \xi \in \beta$, which is totally ordered by \in , so $\mu \notin \xi$ implies that either $\mu = \xi$ or $\xi \in \mu$. Note that $\xi \notin \alpha$ (since $\xi \in X$). Then $\mu = \xi$ is contradictory, since $\mu \in \alpha$. But $\xi \in \mu$ is also contradictory, since $\xi \in \mu \in \alpha \rightarrow \xi \in \alpha$ (since α is a transitive set). \square

Proof of Theorem I.8.5. (1) just restates the fact that γ is a transitive set: $\beta \in \gamma \rightarrow \beta \subseteq \gamma$. (2) uses the fact that \in is irreflexive on α : $x \notin x$ for all $x \in \alpha$, so $\alpha \in \alpha$ would be a contradiction.

For (3), let $\delta = \alpha \cap \beta$. Then δ is an ordinal by Lemma I.8.7, and $\delta \subseteq \alpha$ and $\delta \subseteq \beta$, so $\delta \in \alpha$ or $\delta = \alpha$, and $\delta \in \beta$ or $\delta = \beta$ by Lemma I.8.8. If either $\delta = \alpha$ or $\delta = \beta$, we are done. If not, then $\delta \in \alpha$ and $\delta \in \beta$, so $\delta \in \delta$, contradicting (2).

For (4), let X be any non-empty set of ordinals, and fix an $\alpha \in X$. If α is least, we're done. Otherwise, $\alpha \cap X = \{\xi \in X : \xi < \alpha\} \neq \emptyset$, and $\alpha \cap X$ has a least element, ξ , since α is well-ordered by \in . Then ξ is also the least element of X . \square

Theorem I.8.9 *ON is a proper class; that is, no set contains all the ordinals.*

Proof. If all ordinals were in X , then we would have the set of all ordinals, $ON = \{y \in X : y \text{ is an ordinal}\}$. By Lemma I.8.6 and Theorem I.8.5, ON is an ordinal, so $ON \in ON$, contradicting Theorem I.8.5.2. \square

This contradiction is sometimes known as the ‘‘Burali-Forti Paradox’’ (1897). Of course, this paradox predates the von Neumann ordinals (1923), but, working in Cantor’s naive set theory, Burali-Forti put together the set of all well-orderings to construct a well-ordering strictly longer than all well-orderings, including itself, which is a contradiction.

Exercise I.8.10 *If α, β are ordinals, then $\alpha \cup \beta$ and $\alpha \cap \beta$ are ordinals, with $\alpha \cup \beta = \max(\alpha, \beta)$ and $\alpha \cap \beta = \min(\alpha, \beta)$. If X is a non-empty set of ordinals, then $\bigcap X$ and $\bigcup X$ are ordinals, with $\bigcap X = \min(X)$ and $\bigcup X = \sup(X)$.*

Hint. We already know that X has a least element, and identifying it with $\bigcap X$ just uses the fact that \leq is \subseteq . One can check directly that $\bigcup X$ is an ordinal. X need not have a largest element, so $\bigcup X$ need not be in X . The statement ‘‘ $\bigcup X = \sup(X)$ ’’ is just shorthand for saying that $\bigcup X$ is the supremum, or least upper bound, of X ; that is, $\bigcup X$ is the smallest ordinal α such that $\alpha \geq \xi$ for all $\xi \in X$. \square

In Definition I.6.21, we called $S(x) = x \cup \{x\}$ the ‘‘ordinal successor function’’. We can now verify that $S(\alpha)$ really is the successor ordinal, or the least ordinal greater than α :

Exercise I.8.11 *If α is any ordinal, then $S(\alpha)$ is an ordinal, $\alpha \in S(\alpha)$, and for all ordinals γ : $\gamma < S(\alpha)$ iff $\gamma \leq \alpha$.*

Hint. Once you verify that $S(\alpha)$ is an ordinal, the rest is immediate when you replace ‘‘ $<$ ’’ by ‘‘ \in ’’ and ‘‘ \leq ’’ by ‘‘ \in or $=$ ’’. \square

We partition $ON \setminus \{0\}$ into successors and limits.

Definition I.8.12 *An ordinal β is*

- \Rightarrow a successor ordinal iff $\beta = S(\alpha)$ for some α .
- \Rightarrow a limit ordinal iff $\beta \neq 0$ and β is not a successor ordinal.

\Rightarrow a finite ordinal, or natural number, iff every $\alpha \leq \beta$ is either 0 or a successor.

So, 5 is finite because each of 5, 4, 3, 2, 1, 0 is either 0 or the successor of some ordinal. The natural numbers form an initial segment of the ordinals:

Exercise I.8.13 *If n is a natural number, then $S(n)$ is a natural number and every element of n is a natural number.*

Informally, ω is the set of all natural numbers, but we haven't yet proved that there is such a set.

Theorem I.8.14 (Principle of Ordinary Induction) *For any set X , if $\emptyset \in X$ and $\forall y \in X(S(y) \in X)$, then X contains all natural numbers.*

Proof. Suppose that n is a natural number and $n \notin X$. Let $Y = S(n) \setminus X$. Y is a set of natural numbers (by Exercise I.8.13), and is non-empty (since $n \in Y$), so it has a least element, $k \leq n$ (by Theorem I.8.5.4). Since k is a natural number, it is either 0 or a successor. But $k \neq 0$, since $0 \in X$, so $k = S(i)$ for some i . Then $i \notin Y$ (since k is least), so $i \in X$, and hence $k = S(i) \in X$, a contradiction. \square

The *Axiom of Infinity* (see Section I.2) says exactly that there exists an X satisfying the hypotheses of this theorem. Then X contains all natural numbers, so $\{n \in X : n \text{ is a natural number}\}$ is the set of all natural numbers, justifying:

Definition I.8.15 ω is the set of all natural numbers.

Our proof of Theorem I.8.14 was a bit ugly. It would have been more elegant to say that the least element of $\omega \setminus X$ would yield a contradiction, so $\omega \subseteq X$, but we could not say that because Theorem I.8.14 was used in our justification that the set ω exists.

Note that the Axiom of Infinity is equivalent to the assertion that ω exists. We have chosen as the “official” version of Infinity one which requires fewer defined notions to state — just \emptyset and the successor function, not the theory of ordinals.

Induction is often used as a proof method. To prove $\forall n \in \omega \varphi(n)$, it is sufficient to prove $\varphi(0)$ (the basis) and $\forall n \in \omega[\varphi(n) \rightarrow \varphi(S(n))]$ (the induction step); then, apply the Principle of Ordinary Induction to $X := \{n \in \omega : \varphi(n)\}$.

The set of natural numbers, ω , is also the least limit ordinal. To verify that it is an ordinal we use:

Lemma I.8.16 *Assume that X is a set of ordinals and is an initial segment of ON :*

$$\forall \beta \in X \forall \alpha < \beta [\alpha \in X] \quad . \quad (*)$$

Then $X \in ON$.

Proof. Theorem I.8.5 implies that X is well-ordered by \in , and $(*)$ just says that every $\beta \in X$ is a subset of X , so X is a transitive set. \square

Lemma I.8.17 ω is the least limit ordinal.

Proof. ω is an initial segment of ON by Exercise I.8.13, so ω is an ordinal. ω cannot be $0 = \emptyset$ (since $0 \in \omega$), and cannot be a successor ordinal by Exercise I.8.13. Every ordinal less than (i.e., in) ω is a natural number, and hence not a limit ordinal. \square

In elementary mathematics, there are two basic types of induction used to prove a theorem of the form $\forall \xi \in \omega \varphi(\xi)$. *Ordinary Induction* is specific to ω . *Transfinite Induction*, or the *Least Number Principle*, looks at the least ξ such that $\neg\varphi(\xi)$, and derives a contradiction. This is justified by the fact that ω is well-ordered, and is thus equally valid if you replace ω by any ordinal α , as we shall do in the proof of Lemma I.8.18. In fact, you can also replace ω by ON (see Theorem I.9.1).

Now that we have ω , we can count a little further:

$$0, 1, 2, 3, \dots, S(\omega), S(S(\omega)), S(S(S(\omega))), \dots .$$

$S(S(S(\omega)))$ is usually called $\omega + 3$, although we haven't defined $+$ yet. Informally, after we have counted through the $\omega + n$ for $n \in \omega$, the next ordinal we reach is $\omega + \omega = \omega \cdot 2$, so that $\omega \cdot 2$ will be the first ordinal past the sequence

$$0, 1, 2, \dots, n, \dots \omega, \omega + 1, \omega + 2, \dots, \omega + n, \dots \quad (n \in \omega) .$$

Formally, we need to define $+$ and \cdot . Actually, \cdot is easier; $\alpha \cdot \beta$ will be the ordinal isomorphic to $\beta \times \alpha$ given the lexicographic ordering, which is a well-order by Exercise I.7.24. For example, to get $\omega \cdot 2$, we look at the lexicographic well-ordering on $\{0, 1\} \times \omega$, and count it as we did in Section I.5:

$$\begin{array}{cccccccccccc} (0, 0) & (0, 1) & (0, 2) & \cdots & (0, n) & \cdots & (1, 0) & (1, 1) & (1, 2) & \cdots & (1, n) & \cdots \\ 0 & 1 & 2 & \cdots & n & \cdots & \omega & \omega + 1 & \omega + 2 & \cdots & \omega + n & \cdots \end{array}$$

In this way, we construct an ordinal γ such that γ (i.e., the set of ordinals less than γ) is well-ordered isomorphically to $\{0, 1\} \times \omega$, and this γ will be called $\omega \cdot 2$.

But now, to justify our intended definition of $\alpha \cdot \beta$, we must show (Theorem I.8.19) that every well-ordered set is isomorphic to an ordinal. First, observe:

Lemma I.8.18 If $f : \alpha \xrightarrow[\text{onto}]{} \beta$ is an isomorphism from $(\alpha; <)$ to $(\beta; <)$ then f is the identity map and hence $\alpha = \beta$.

Proof. Fix $\xi \in \alpha$; then $f(\xi) \in \beta$. Just thinking of α and β as totally ordered sets, the isomorphism f maps the elements below ξ in α onto the elements below $f(\xi)$ in β , so:

$$\{\nu : \nu \in \beta \wedge \nu < f(\xi)\} = \{f(\mu) : \mu \in \alpha \wedge \mu < \xi\} .$$

But since $<$ is membership and α, β are transitive sets, this simplifies to

$$f(\xi) = \{f(\mu) : \mu \in \xi\} . \quad (\dagger)$$

Now, we prove that $f(\xi) = \xi$ by *transfinite induction* on ξ . That is, let $X = \{\xi \in \alpha : f(\xi) \neq \xi\}$. If $X = \emptyset$, we are done. If $X \neq \emptyset$, let ξ be the least element of X . Then $\mu < \xi$ implies $f(\mu) = \mu$, so then (\dagger) tells us that $f(\xi)$ is $\{\mu : \mu \in \xi\}$, which is ξ , contradicting $\xi \in X$. \square

Note that this lemma applied with $\alpha = \beta$ implies that the ordering $(\alpha; <)$ has no automorphisms other than the identity map.

Theorem I.8.19 *If R well-orders A , then there is a unique $\alpha \in ON$ such that $(A; R) \cong (\alpha; \in)$.*

Proof. Uniqueness is immediate from Lemma I.8.18. Informally, we prove existence by counting off A , starting from the bottom, as in Example I.7.23. Since we have no formal notion of “counting”, we proceed in a slightly different way.

If $a \in A$, let $a \downarrow = \{x \in A : xRa\}$. Then $a \downarrow$ is well-ordered by R as well. Call $a \in A$ *good* iff the theorem holds for $a \downarrow$ — that is, $(a \downarrow; R) \cong (\xi; \in)$ for some ordinal ξ . Note that this ξ is then unique by Lemma I.8.18. Let G be the set of good elements of A . Let f be the function with domain G such that $f(a)$ is the (unique) ξ such that $(a \downarrow; R) \cong (\xi; \in)$; this definition is justified by the Replacement Axiom (see Lemma I.7.9). Lemma I.8.18 also implies that for $a \in G$, the isomorphism $(a \downarrow; R)$ onto $(f(a); \in)$ is unique; if there were two different isomorphisms, h and k , then $h \circ k^{-1}$ would be a non-trivial automorphism of $(f(a); \in)$. For $a \in G$, let h_a be this unique isomorphism.

For example, suppose that A has at least four elements, with the first four, in order, being a_0, a_1, a_2, a_3 . Then a_3 is good, with $f(a_3) = 3$, because the isomorphism $h_{a_3} = \{(a_0, 0), (a_1, 1), (a_2, 2)\}$ takes $a_3 \downarrow$ onto $3 = \{0, 1, 2\}$. Likewise, a_0 is good, with $f(a_0) = 0$, because the empty isomorphism $h_{a_0} = \emptyset$ takes $a_0 \downarrow = \emptyset$ onto $0 = \emptyset$. Likewise, a_1 and a_2 are good, with $f(a_1) = 1$ and $f(a_2) = 2$. Now note that each isomorphism h_{a_i} is the same as $f \upharpoonright (a_i \downarrow)$.

More generally:

$$\forall a \in G \forall c \in a \downarrow \left[c \in G \wedge h_c = h_a \upharpoonright (c \downarrow) \wedge f(c) = h_a(c) \right] . \quad (\otimes)$$

That is, if $a \in G$ and cRa , then $h_a \upharpoonright (c \downarrow)$ is an isomorphism from $c \downarrow$ onto $h_a(c)$, so that $c \in G$, with h_c the isomorphism $h_a \upharpoonright (c \downarrow)$ onto $h_a(c)$, which is then $f(c)$.

Hence, the map $f : G \rightarrow ON$ is order-preserving, that is, $cRa \rightarrow f(c) < f(a)$, since $f(c) = h_a(c)$ is a member of $f(a)$. Also, $\text{ran}(f)$ is an initial segment of ON , because if $\xi = f(a) = \text{ran}(h_a)$, then any $\eta < \xi$ must be of the form $h_a(c) = f(c)$ for some cRa . Thus, if $\alpha = \text{ran}(f)$, then α is an ordinal (by Lemma I.8.16), and f is an isomorphism from G onto α .

If $G = A$, we are done. If not, let e be the least element of $A \setminus G$. Then $e \downarrow = G$, because G is an initial segment of A by (\clubsuit) . Hence, $(e \downarrow; R) \cong (\alpha; \in)$, so $e \in G$, a contradiction. \square

Definition I.8.20 *If R well-orders A , then $\text{type}(A; R)$ is the unique $\alpha \in ON$ such that $(A; R) \cong (\alpha; \in)$. We also write $\text{type}(A)$ (when R is clear from context), or $\text{type}(R)$ (when A is clear from context).*

For example, we say $\text{type}\{n \in \omega : n > 0\} = \omega$, if it's understood that we're using the usual order on the natural numbers. However, if R well-orders ω by putting 0 on the top and ordering the positive natural numbers in the usual way, then we write $\text{type}(R) = \omega + 1$, if it's understood that we're discussing various ways to well-order ω .

Now, we define ordinal sum and product by:

Definition I.8.21 $\alpha \cdot \beta = \text{type}(\beta \times \alpha)$. $\alpha + \beta = \text{type}(\{0\} \times \alpha \cup \{1\} \times \beta)$.

In both cases, we're using lexicographic order to compare ordered pairs of ordinals.

For example, $\omega + \omega = \omega \cdot 2 = \text{type}(\{0, 1\} \times \omega)$. This well-ordering consists of a copy of ω followed by a second copy stacked on top of it.

Now that we have an "arithmetic" of ordinals, it is reasonable to ask what properties it has. These properties are summarized in Table I.1, p. 41, but it will take a bit of discussion, in this section and the next, to make everything in that table meaningful. After that is done, most of the proofs will be left as exercises.

Observe first that some of the standard facts about arithmetic on the natural numbers extend to the ordinals, but others do not. For example, $+$ and \cdot are both associative, but not commutative, since $1 + \omega = \omega$ but $\omega + 1 = S(\omega) > \omega$. Likewise, $2 \cdot \omega = \omega$ but $\omega \cdot 2 = \omega + \omega > \omega$.

Now, what exactly is the meaning of the "Recursive Computation" in Table I.1? The lines for $+$ and \cdot are simply facts about $+$ and \cdot which can easily be verified from the definition (I.8.21). Informally, they are also schemes for computing these functions, if we view computations as extending through the transfinite. Think of α as being fixed. To compute $\alpha + \beta$, we are told that $\alpha + 0 = \alpha$ and we are told how to obtain $\alpha + S(\beta)$ from $\alpha + \beta$, so we may successively compute $\alpha + 1, \alpha + 2, \dots$. Then, since ω is a limit, at stage ω of this process we know how to compute $\alpha + \omega$ as $\sup_{n < \omega} (\alpha + n)$. We may now proceed to compute $\alpha + (\omega + 1), \alpha + (\omega + 2), \dots$, and then, at stage $\omega + \omega$, we may compute $\alpha + (\omega + \omega)$. In general, we compute $\alpha + \beta$ at stage β of this process. Once we know how to compute $+$, the second line tells us how to compute $\alpha \cdot \beta$ by successively computing $\alpha \cdot 0, \alpha \cdot 1, \alpha \cdot 2, \dots$, until we get up to β .

Since we already have a perfectly good definition of $+$ and \cdot , we might be tempted to consider this discussion of computation to be purely informal. However, there are other functions on ordinals, such as exponentiation (α^β), where the most natural definitions are by recursion. In Section I.9, we take up recursion more formally, and in particular justify the definition of exponentiation.

Table I.1: Ordinal Arithmetic

Associative Laws

$$(\alpha + \beta) + \gamma = \alpha + (\beta + \gamma) \quad ; \quad (\alpha \cdot \beta) \cdot \gamma = \alpha \cdot (\beta \cdot \gamma)$$

– Commutative Laws

$$1 + \omega = \omega < \omega + 1 \quad ; \quad 2 \cdot \omega = \omega < \omega \cdot 2$$

Left Distributive Law: $\alpha \cdot (\beta + \gamma) = \alpha \cdot \beta + \alpha \cdot \gamma$ – Right Distributive Law: $(1 + 1) \cdot \omega = \omega < \omega + \omega = 1 \cdot \omega + 1 \cdot \omega$

0 and 1

$$S(\alpha) = \alpha + 1$$

$$\alpha + 0 = 0 + \alpha = \alpha \quad ; \quad \alpha \cdot 0 = 0 \cdot \alpha = 0 \quad ; \quad \alpha \cdot 1 = 1 \cdot \alpha = \alpha$$

$$\alpha^0 = 1 \quad ; \quad \alpha^1 = \alpha \quad ; \quad 1^\alpha = 1$$

$$\alpha > 0 \rightarrow 0^\alpha = 0$$

Left Cancellation

$$\alpha + \beta = \alpha + \gamma \rightarrow \beta = \gamma$$

$$\alpha \cdot \beta = \alpha \cdot \gamma \wedge \alpha \neq 0 \rightarrow \beta = \gamma$$

– Right Cancellation

$$1 + \omega = 2 + \omega = 1 \cdot \omega = 2 \cdot \omega = \omega$$

Subtraction

$$\alpha \leq \beta \rightarrow \exists! \gamma (\alpha + \gamma = \beta)$$

Division

$$\alpha > 0 \rightarrow \exists! \gamma \delta (\alpha \cdot \gamma + \delta = \beta \wedge \delta < \alpha)$$

Exponentiation

$$\alpha^{\beta+\gamma} = \alpha^\beta \cdot \alpha^\gamma \quad ; \quad \alpha^{\beta \cdot \gamma} = (\alpha^\beta)^\gamma$$

Logarithms

$$\alpha > 1 \wedge \beta > 0 \rightarrow \exists! \gamma \delta \xi (\beta = \alpha^\delta \cdot \xi + \gamma \wedge \xi < \alpha \wedge \gamma < \alpha^\delta \wedge \xi > 0)$$

$$\text{For finite ordinals, } \delta = \lfloor \log_\alpha \beta \rfloor$$

$$\text{Example: } 873 = 10^2 \cdot 8 + 73 \wedge 8 < 10 \wedge 73 < 10^2$$

Order

$$\beta < \gamma \rightarrow \alpha + \beta < \alpha + \gamma \wedge \beta + \alpha \leq \gamma + \alpha$$

$$\beta < \gamma \wedge \alpha > 0 \rightarrow \alpha \cdot \beta < \alpha \cdot \gamma \wedge \beta \cdot \alpha \leq \gamma \cdot \alpha$$

$$\beta < \gamma \wedge \alpha > 1 \rightarrow \alpha^\beta < \alpha^\gamma \wedge \beta^\alpha \leq \gamma^\alpha$$

$$2 < 3 \quad ; \quad 2 + \omega = 3 + \omega = 2 \cdot \omega = 3 \cdot \omega = 2^\omega = 3^\omega = \omega$$

Recursive Computation

$$\alpha + 0 = \alpha \quad ; \quad \alpha + S(\beta) = S(\alpha + \beta) \quad ; \quad \alpha + \gamma = \sup_{\beta < \gamma} (\alpha + \beta) \quad (\text{for } \gamma \text{ a limit})$$

$$\alpha \cdot 0 = 0 \quad ; \quad \alpha \cdot S(\beta) = \alpha \cdot \beta + \alpha \quad ; \quad \alpha \cdot \gamma = \sup_{\beta < \gamma} (\alpha \cdot \beta) \quad (\text{for } \gamma \text{ a limit})$$

$$\alpha^0 = 1 \quad ; \quad \alpha^{S(\beta)} = \alpha^\beta \cdot \alpha \quad ; \quad \alpha^\gamma = \sup_{\beta < \gamma} (\alpha^\beta) \quad (\text{for } \gamma \text{ a limit})$$

Continuity

The functions $\alpha + \beta$, $\alpha \cdot \beta$, α^β are continuous in β , but not in α .

Using ω and the finite ordinals, we can define lots of ordinals by taking sums and products. These will all be countable. For example, $\omega \cdot \omega$ is isomorphic to $\omega \times \omega$, and hence in 1-1 correspondence with $\omega \times \omega$, which is a countable set. We would expect that after we have counted past all the countable ordinals, we would reach some uncountable ones. The first uncountable ordinal will be called ω_1 . For details on countable and uncountable ordinals, see Section I.11.

Remark I.8.22 The Replacement Axiom (via the proof of Theorem I.8.19) is needed to prove that the ordinal $\omega + \omega$ exists. In Zermelo set theory, ZC (defined in Section I.2), one cannot prove that $\omega + \omega$ exists (see [20] or Exercise I.14.16), although one can produce the set $2 \times \omega$ and the lexicographic order on it. This did not cause a problem in Zermelo's day, since he worked before the definition of the von Neumann ordinals anyway; where we now use the ordinal $\omega + \omega$, he would simply use some (any) set, such as $\{0, 1\} \times \omega$, ordered in this type.

Exercise I.8.23 If R well-orders A and $X \subseteq A$, then R well-orders X (see Exercise I.7.22) and $\text{type}(X; R) \leq \text{type}(A; R)$.

Hint. Applying Theorem I.8.19, WLOG A is an ordinal, α , and R is $<$, the usual (membership) relation on the ordinals. Let f be the isomorphism from $(X; <)$ onto some ordinal, $\delta = \text{type}(X; <)$, and prove that $f(\xi) \leq \xi$ by transfinite induction on ξ . \square

I.9 Induction and Recursion on the Ordinals

These are related concepts, but they are not the same, although they are sometimes confused in the popular literature. Induction is a method of proof, whereas recursion is a method of defining functions. On ω , we already have the Principle of Ordinary Induction (Theorem I.8.14), whereas *recursion* would be used to justify the definition of the factorial (!) function by:

$$0! = 1 \quad (n + 1)! = n! \cdot (n + 1) \quad .$$

First, a few more remarks on induction. Ordinary induction is specific to ω , but *transfinite induction* is a generalization of the least number principle, which, as discussed in Section I.8, works on every ordinal. If α is an ordinal, then every non-empty subset of α has a least element; this is just the definition of well-order. Used as a proof procedure, we prove $\forall \xi < \alpha \varphi(\xi)$ by deriving a contradiction from the least $\xi < \alpha$ such that $\neg \varphi(\xi)$. This method is most familiar when $\alpha = \omega$, but we have used it for an arbitrary α in the proofs of Lemma I.8.18 and Exercise I.8.23. A "proper class" flavor of transfinite induction can be used to prove theorems about all the ordinals at once:

Theorem I.9.1 (Transfinite Induction on ON) For each formula ψ : if $\psi(\alpha)$ holds for some ordinal α , then there is a least ordinal ξ such that $\psi(\xi)$.

Proof. Fix α such that $\psi(\alpha)$. If α is least, then we are done. If not, then the set $X := \{\xi < \alpha : \psi(\xi)\}$ is non-empty, and the least ξ in X is the least ξ such that $\psi(\xi)$. \square

Note that this is really a theorem scheme, as discussed in subsection I.7.2. Formally, we are making an assertion in the metatheory that for each formula $\psi(\alpha)$, the universal closure of

$$\exists \alpha \psi(\alpha) \rightarrow \exists \alpha [\psi(\alpha) \wedge \forall \xi < \alpha [\neg \psi(\xi)]]$$

is provable from the axioms of set theory.

Now, recursion, like induction, comes in several flavors. In computing, a *recursion* is any definition of a function $f(x)$ which requires the evaluation of $f(y)$ for some other inputs y . This is discussed in more detail in *recursion theory* (see Chapter III). In set theory, we only need the special case of this, called *primitive recursion*, where the evaluation of $f(x)$ may require the evaluation of $f(y)$ for one or more y less than x . For example, everyone recognizes the following definition of the Fibonacci numbers as legitimate:

$$f(0) = f(1) = 1 \quad ; \quad f(x) = f(x-1) + f(x-2) \text{ when } x > 1 \quad .$$

Here the evaluation of $f(x)$ requires knowing two earlier values of the function when $x > 1$, and requires knowing zero earlier values when $x \leq 1$. Informally, this kind of definition is justified because we can fill out a table of values, working left to right:

x	0	1	2	3	4	5	⋯⋯⋯
$f(x)$	1	1	2	3	5	8	⋯⋯⋯

We shall prove a theorem (Theorem I.9.2) stating that definitions of “this general form” are legitimate, but first we have to state what “this general form” is. Roughly, one defines a function f by the recipe:

$$f(\xi) = G(f \upharpoonright \xi) \quad , \tag{*}$$

where G is a given function. Note that $f \upharpoonright \xi$ (Definition I.7.4) is the function f restricted to the set of ordinals less than ξ , and G tells us how to compute $f(\xi)$ from this. In the case of the Fibonacci numbers, ξ is always in ω , and we can define $G_{\text{fib}}(s)$ to be 1 unless s is a function with domain some natural number $x \geq 2$, in which case $G_{\text{fib}}(s) = s(x-1) + s(x-2)$. For example, the table above displays $f \upharpoonright 6$, and G_{fib} tells us that we should compute $f(6)$ from the table as $f(5) + f(4) = 8 + 5 = 13$.

Now, the Fibonacci numbers are defined only on ω , but we may wish to use the same scheme (*) to define a function on a larger ordinal, or even all the ordinals at once. For example, consider the recursive definition of ordinal exponentiation from Table I.1: $\alpha^0 = 1$; $\alpha^{S(\beta)} = \alpha^\beta \cdot \alpha$; $\alpha^\gamma = \sup_{\beta < \gamma} (\alpha^\beta)$ for γ a limit. If we fix α , we may consider the function $E_\alpha(\xi) = \alpha^\xi$ to be defined by (*); that is $E_\alpha(\xi) = G_\alpha(E_\alpha \upharpoonright \xi)$. Here, we may define $G_\alpha(s)$ to be 0 unless s is a function with domain an ordinal, ξ , in which case

$$G_\alpha(s) = \begin{cases} 1 & \text{if } \xi = 0 \\ s(\beta) \cdot \alpha & \text{if } \xi = \beta + 1 \\ \sup_{\beta < \xi} s(\beta) & \text{if } \xi \text{ is a limit} \end{cases}$$

Actually, our Theorem I.9.2 talks only about functions such as E_α defined on all the ordinals at once. To get a function on a specific ordinal, we can just restrict to that ordinal. For example, if G_{fib} is defined verbatim as above, the prescription $f(\xi) = G_{\text{fib}}(f \upharpoonright \xi)$ actually makes sense for all ξ , and makes $f(\xi) = 1$ when $\xi \geq \omega$, since then $f \upharpoonright \xi$ is not a function with domain some natural number.

Another well-known example: If $h : A \rightarrow A$ is a given function, then h^n (for $n \in \omega$) denotes the function h iterated n times, so $h^0(a) = a$ and $h^3(a) = h(h(h(a)))$. There is no standard meaning given to h^ξ for an infinite ordinal ξ . Consider h and A to be fixed, so h^n depends on n . If we denote h^n by $F(n)$, then we are defining F by recursion on n . In the recipe (*), we can let $G(s)$ be the identity function on A (formally, $\{(x, y) \in A \times A : x = y\}$) unless s is a function with domain some natural number $x \geq 1$, in which case $G(s) = h \circ s(x-1)$. With this specific G , our $h^\xi = F(\xi)$ will be the identity function on A whenever $\xi \geq \omega$. In practice, these h^ξ are never used, and we really only care about $F \upharpoonright \omega$, but we do not need a second theorem to handle the “only care” case; we can just prove one theorem.

Now, since our functions will be defined on all the ordinals, they cannot be considered sets of ordered pairs, but rather rules, given by formulas (see Subsection I.7.2). Thus, in the formal statement of the theorem, G is defined by a formula $\varphi(s, y)$ as before, and the defined function F is defined by another formula $\psi(\xi, y)$:

Theorem I.9.2 (Primitive Recursion on ON) *Suppose that $\forall s \exists! y \varphi(s, y)$, and define $G(s)$ to be the unique y such that $\varphi(s, y)$. Then we can define a formula ψ for which the following are provable:*

1. $\forall x \exists! y \psi(x, y)$, so ψ defines a function F , where $F(x)$ is the y such that $\psi(x, y)$.
2. $\forall \xi \in ON [F(\xi) = G(F \upharpoonright \xi)]$.

Remarks. In our formal statement of the theorem, $G(s)$ is defined for all s , even though the computation of $F(\xi)$ only uses $G(s)$ when s is relevant to the computation of a function on the ordinals (i.e., s is a function with domain some ordinal). This is no problem, since we can always let $G(s)$ take some default value for the “uninteresting” s , as we did with the examples above. Likewise, in the recursion, we are really “thinking” of F as defined only on the ordinals, but (1) says $\forall x \exists! y \psi(x, y)$, so $F(x)$ defined for all x ; this is no problem, since we can let $F(x)$ be some default value, say, 0, when x is not an ordinal.

Although F is not really a set, each $F \upharpoonright \delta = \{(\eta, F(\eta)) : \eta \in \delta\}$ is a set by the Replacement Axiom (see Lemma I.7.9). Thus, in applications where we really only care about F on δ , our theorem gives us a legitimate set-of-ordered-pairs function $F \upharpoonright \delta$.

Formally, this theorem is a scheme in the metatheory, saying that for each such formula φ , we can write another formula ψ and prove (1) and (2).

Proof. For any ordinal δ , let $\text{App}(\delta, h)$ (h is a δ -approximation to our function) say that h is a function (in the set-of-ordered-pairs sense), $\text{dom}(h) = \delta$, and $h(\xi) = G(h \upharpoonright \xi)$

for all $\xi < \delta$. Once the theorem is proved, it will be clear that $\text{App}(\delta, h)$ is equivalent to $h = F \upharpoonright \delta$. We shall prove the theorem by using this $\text{App}(\delta, h)$ to write down the following definition of ψ :

$$\psi(x, y) \iff \begin{aligned} & [x \notin ON \wedge y = 0] \vee \\ & [x \in ON \wedge \exists \delta > x \exists h [\text{App}(\delta, h) \wedge h(x) = y]] \end{aligned}$$

We now need to check that this definition works, which will be easy once we have verified the existence and uniqueness of these δ -approximations. Uniqueness means that all the δ -approximations agree wherever they are defined:

$$\delta \leq \delta' \wedge \text{App}(\delta, h) \wedge \text{App}(\delta', h') \rightarrow h = h' \upharpoonright \delta \quad . \quad (U)$$

In particular, with $\delta = \delta'$, this says that for each δ , there is at most one δ -approximation. Then, existence says:

$$\forall \delta \exists ! h \text{App}(\delta, h) \quad . \quad (E)$$

Given (U) and (E), the theorem follows easily: First, we note that they imply that $\forall x \exists ! y \psi(x, y)$, so ψ defines a function, F , as in (1). To verify (2), fix a ξ , and then fix a $\delta > \xi$. Applying (E), let h be the unique function satisfying $\text{App}(\delta, h)$. By the definition of ψ and F , we see that $F \upharpoonright \delta = h$ and hence $F \upharpoonright \xi = h \upharpoonright \xi$, so that, applying the definition of App , $F(\xi) = h(\xi) = G(h \upharpoonright \xi) = G(F \upharpoonright \xi)$, which yields (2).

To prove (U), we show by transfinite induction that $h(\xi) = h'(\xi)$ for all $\xi < \delta$. If this fails, let ξ be the least element in $\{\xi < \delta : h(\xi) \neq h'(\xi)\}$. But then, $h \upharpoonright \xi = h' \upharpoonright \xi$, so, applying $\text{App}(\delta, h)$ and $\text{App}(\delta', h')$, $h(\xi) = G(h \upharpoonright \xi) = G(h' \upharpoonright \xi) = h'(\xi)$, a contradiction.

To prove (E), we apply transfinite induction on ON (Theorem I.9.1). Since (U) will give us uniqueness, it is sufficient to prove that $\forall \delta \exists h \text{App}(\delta, h)$. If this fails, then fix $\delta \in ON$ to be least such that $\neg \exists h \text{App}(\delta, h)$. Since δ is least, whenever $\beta < \delta$, there is a unique (by (U)) function g_β with $\text{App}(\beta, g_\beta)$. There are now three cases:

Case 1. δ is a successor ordinal: Say $\delta = \beta + 1$. Let $f = g_\beta \cup \{\langle \beta, G(g_\beta) \rangle\}$. So, f is a function with $\text{dom}(f) = \beta \cup \{\beta\} = \delta$, and $f \upharpoonright \beta = g_\beta$. Observe that $f(\xi) = G(f \upharpoonright \xi)$ holds for all $\xi < \delta$: for $\xi < \beta$, use $\text{App}(\beta, g_\beta)$, and for $\xi = \beta$, use the definition of $f(\beta)$. But then $\text{App}(\delta, f)$, contradicting $\neg \exists h \text{App}(\delta, h)$.

Case 2. $\delta = 0$: Then $\text{App}(0, \emptyset)$.

Case 3. δ is a limit ordinal: Let $f = \bigcup \{g_\beta : \beta < \delta\}$. Then f is a function (by (U)) with $\text{dom}(f) = \delta$. Furthermore, observe that $\text{App}(\delta, f)$; this is proved by exactly the same argument which we used (above) to conclude the theorem from (U)+(E). But now again we contradict $\neg \exists h \text{App}(\delta, h)$. \square

Definition I.9.3 For $\alpha \in ON$, an α -sequence is a function s with domain α , and $s_\xi = s(\xi)$ for $\xi < \alpha$.

For example, the “infinite sequences” of real numbers studied in calculus are really functions $s : \omega \rightarrow \mathbb{R}$, and we usually write s_n rather than $s(n)$.

Remark I.9.4 *Recursive definitions of ω -sequences are common in mathematics, and are formally justified by Theorem I.9.2.* For example, we may say: Define a sequence of reals by $s_0 = 0.7$ and $s_{n+1} = \cos(s_n)$. Formally, we are defining a function $s : \omega \rightarrow \mathbb{R}$ by recursion, and the justification of the definition is the same as that of the Fibonacci numbers discussed above. Then s_n is just another notation for $s(n)$. The statement that the s_n converge to a value $t \sim 0.739$ such that $\cos(t) = t$ is a statement about the function s .

In set theory, it is often convenient to iterate the \cup operator. Given the set x , we can form the sequence

$$x = \cup^0 x, \cup x = \cup^1 x, \cup\cup x = \cup^2 x, \dots$$

Note that x is (trivially) the set of members of x , $\cup x$ is the set of members of members of x , $\cup^2 x = \cup\cup x$ is the set of members of members of members of x , and so forth. If we view \in as a directed graph, then $z \in \cup^n x$ iff there is a path from z to x of length exactly n . There may be multiple paths of different lengths. For example, $2 \in 9$ and also $2 \in \cup^3 9$ because $2 \in 4 \in 7 \in 9$.

To justify our notation: Formally, each x , we are defining a function f_x on ω by $f_x(0) = x$ and $f_x(n+1) = \cup f_x(n)$. Then, $\cup^n x$ is just another notation for $f_x(n)$.

Definition I.9.5 *Let $\cup^0 x = x$, $\cup^1 x = \cup x$, and, (recursively) $\cup^{n+1} x = \cup\cup^n x$. Then, $\text{trcl}(x) = \cup\{\cup^n x : n \in \omega\}$ is called the transitive closure of x .*

So, $z \in \text{trcl}(x)$ iff there is some finite \in -path from z to x ; as such, the relation “ $z \in \text{trcl}(x)$ ” is a special case of the general notion of the transitive closure of a directed graph. Also, $\text{trcl}(x)$ is the least transitive superset of x :

Exercise I.9.6 *For any set x :*

- $x \subseteq \text{trcl}(x)$.
- $\text{trcl}(x)$ is transitive.
- If $x \subseteq t$ and t is transitive, then $\text{trcl}(x) \subseteq t$.
- If $y \in x$ then $\text{trcl}(y) \subseteq \text{trcl}(x)$.

I.10 Power Sets

The Power Set Axiom says that for all x , there is a y such that $\forall z(z \subseteq x \rightarrow z \in y)$. Thus, applying comprehension, $\{z : z \subseteq x\}$ exists, justifying:

Definition I.10.1 *The power set of x is $\mathcal{P}(x) = \{z : z \subseteq x\}$*

We can now define function spaces. For example, if $f : A \rightarrow B$, then $f \in \mathcal{P}(A \times B)$, which justifies the following definition:

Definition I.10.2 $B^A = {}^A B$ is the set of functions f with $\text{dom}(f) = A$ and $\text{ran}(f) \subseteq B$.

We shall usually use B^A , which is the standard terminology in most of mathematics, but occasionally ${}^A B$ is used when A and B are ordinals. For example, 2^3 denotes the number 8, which is different from ${}^3 2$, which is a set of 8 functions. However, \mathbb{R}^3 can only mean the set of functions from $3 = \{0, 1, 2\}$ to \mathbb{R} . If $x \in \mathbb{R}^3$, then x is a sequence of three reals, x_0, x_1, x_2 , as in Definition I.9.3.

Definition I.10.3 $A^{<\alpha} = {}^{<\alpha} A = \bigcup_{\xi < \alpha} A^\xi$.

In particular, if we view A as an alphabet, then $A^{<\omega}$ is the set of all “words” (or strings of finite length) which can be formed from elements of A . We may use $\sigma = (x, y, z, t)$ to denote a point in \mathbb{R}^4 ; formally, this σ is a function from 4 into \mathbb{R} , and this notation is made precise by:

Definition I.10.4 If $a_0, \dots, a_{m-1} \in A$, then (a_0, \dots, a_{m-1}) denotes the $\sigma \in A^m$ defined so that $\sigma(i) = a_i$. If $\sigma \in A^m$ and $\tau \in A^n$, then $\sigma \frown \tau$ or $\sigma\tau$ denotes the concatenation $\pi \in A^{m+n}$ of σ and τ , defined so that $\pi(i) = \sigma(i)$ for $i < m$ and $\pi(m+i) = \tau(j)$ for $j < n$.

The concatenation of strings is frequently used in the theory of formal languages (see Sections I.15 and II.4). There is a possible ambiguity now in the notation for pairs, since (x, y) could denote a function with domain 2 or the pair $\langle x, y \rangle = \{\{x\}, \{x, y\}\}$. This will not cause a problem, since both notions of pair determine x, y uniquely, as in Exercise I.6.14.

Exercise I.10.5 Justify the definitions of $A \times B$ and A/R (see Definitions I.7.8 and I.7.15) using the Power Set Axiom but not using the Replacement Axiom (i.e., in the theory Z^- ; see Section I.2).

Hint. Note that $A/R \subseteq \mathcal{P}(A)$ and $A \times B \subseteq \mathcal{P}(\mathcal{P}(A \cup B))$. This exercise is of some historical interest because once the axiom system of Zermelo was introduced in 1908, it was clear that one could formalize these standard mathematical constructs in set theory, although the Axiom of Replacement was not introduced until 1922 (by Fraenkel). \square

I.11 Cardinals

Definition I.11.1

1. $X \preceq Y$ iff there is a function $f : X \xrightarrow{1-1} Y$.
2. $X \approx Y$ iff there is a function $f : X \xrightarrow[1-1]{\text{onto}} Y$.

So $X \approx Y$ says that X, Y have the *same size* in the sense that they can be put into 1-1 correspondence by some function, whereas $X \preceq Y$ means that there is a function which embeds X 1-1 into Y , so the size of X is equal to or less than that of Y .

Lemma I.11.2

1. \preceq is transitive and reflexive.
2. $X \subseteq Y \rightarrow X \preceq Y$.
3. \approx is an equivalence relation.

Proof. For (2), and in the proof of reflexivity in (1) and (3), use the identity function. To prove transitivity in (1) and (3), compose functions. \square

These notions of cardinalities for infinite sets were first studied in detail by Cantor, although the fact that an infinite set can have the same size as a proper subset is due to Galileo, who pointed out in 1638 that $\omega \approx \{n^2 : n \in \omega\}$ via the map $n \mapsto n^2$ (“we must say that there are as many squares as there are numbers” [13]).

Exercise I.11.3 Prove that $\mathbb{R} \times \mathbb{R} \approx (0, 1) \times (0, 1) \preceq (0, 1) \approx \mathbb{R}$.

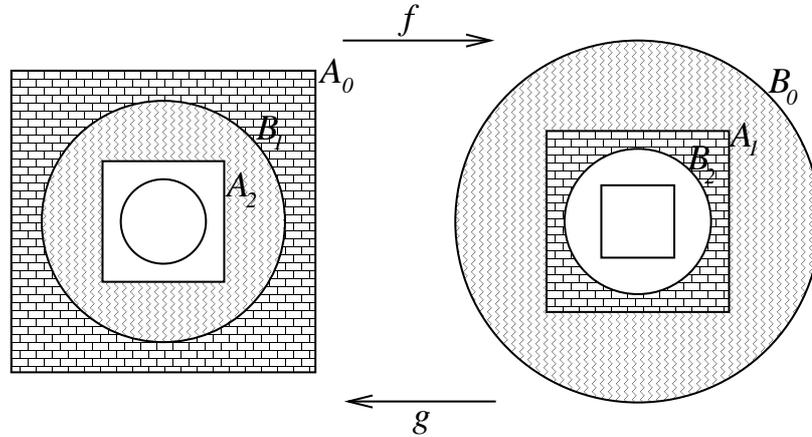
Hint. Of course, you have to use your knowledge of the real numbers here, since we haven’t formally developed them yet (see Section I.15). The tangent function maps an open interval onto \mathbb{R} . For $(0, 1) \times (0, 1) \preceq (0, 1)$, represent each $x \in (0, 1)$ by an infinite decimal, say, never ending in an infinite sequence of nines. Then, given x, y , map them to the real z obtained by shuffling the decimal representations of x, y . \square

It is often easier to demonstrate injections than bijections. For example, $\mathbb{R} \times \mathbb{R} \preceq \mathbb{R}$ by Exercise I.11.3 and $\mathbb{R} \preceq \mathbb{R} \times \mathbb{R}$ is trivial. The proof of Theorem I.11.4 below tells you how to write down an explicit definition of a bijection from \mathbb{R} onto $\mathbb{R} \times \mathbb{R}$, but it is not very “simple”.

Theorem I.11.4 (Schröder and Bernstein) $A \approx B$ iff $A \preceq B$ and $B \preceq A$.

Proof. Given $f : A \xrightarrow[1-1]{\text{onto}} B$, we have $A \preceq B$ using f , and $B \preceq A$ using f^{-1} .

For the nontrivial direction, we have $f : A \xrightarrow{1-1} B$ and $g : B \xrightarrow{1-1} A$, and we shall define a bijection $h : A \xrightarrow[1-1]{\text{onto}} B$.



Define (by recursion) sets A_n and B_n for $n \in \omega$ by: $A_0 = A$ and $B_0 = B$. If n is even, then $A_{n+1} = f(A_n)$ and $B_{n+1} = g(B_n)$. If n is odd, then $A_{n+1} = g(A_n)$ and $B_{n+1} = f(B_n)$. Observe that

$$A_0 \supseteq B_1 \supseteq A_2 \supseteq B_3 \supseteq \dots$$

$$B_0 \supseteq A_1 \supseteq B_2 \supseteq A_3 \supseteq \dots \quad .$$

Let

$$P = \bigcap_{k < \omega} A_{2k} = \bigcap_{k < \omega} B_{2k+1}$$

$$Q = \bigcap_{k < \omega} B_{2k} = \bigcap_{k < \omega} A_{2k+1} \quad .$$

P, Q might be empty, but in any case note that $f \upharpoonright P$ is a bijection from P onto Q . Also note that when n is even, $f \upharpoonright (A_n \setminus B_{n+1})$ is a bijection from $A_n \setminus B_{n+1}$ onto $A_{n+1} \setminus B_{n+2}$, and $g \upharpoonright (B_n \setminus A_{n+1})$ is a bijection from $B_n \setminus A_{n+1}$ onto $B_{n+1} \setminus A_{n+2}$.

Now, define $h : A \rightarrow B$ so that $h(x) = f(x)$ whenever either $x \in P$ or $x \in A_n \setminus B_{n+1}$ for some even n , whereas $h(x) = g^{-1}(x)$ whenever $x \in B_n \setminus A_{n+1}$ for some odd n . It follows from the above remarks that h is a bijection. □

We remark on the recursion in the above proof. Informally, we say that the A_n and B_n are being defined “simultaneously by recursion”. Formally, we are using our results on recursive definitions in Section I.9.2 to construct the map $n \mapsto (A_n, B_n)$ from ω into $\mathcal{P}(A \cup B) \times \mathcal{P}(A \cup B)$; see Remark I.9.4.

Cantor conjectured that Theorem I.11.4 held, but he didn’t have a proof. Cantor also conjectured that any two sets are comparable ($X \preceq Y$ or $Y \preceq X$). This is also true, but the proof requires the Axiom of Choice (see Section I.12).

Along with \preceq , there is a notion of strictly smaller in size:

Definition I.11.5 $X \prec Y$ iff $X \preceq Y$ and $Y \not\preceq X$.

In view of the Schröder-Bernstein Theorem, this is the same as saying that X can be mapped 1-1 into Y , but there is no bijection from X onto Y .

A famous theorem of Cantor states that $A \prec \mathcal{P}(A)$. We isolate the key idea of the proof, showing that one cannot map A onto $\mathcal{P}(A)$, as:

Lemma I.11.6 (Cantor's Diagonal Argument) *If f is a function, $\text{dom}(f) = A$, and $D = \{x \in A : x \notin f(x)\}$, then $D \notin \text{ran}(f)$.*

Proof. If $D = f(c)$ for some $c \in A$, then applying the definition of D with $x = c$ we would have $c \in D \leftrightarrow c \notin f(c)$, so $c \in D \leftrightarrow c \notin D$, a contradiction. \square

Theorem I.11.7 (Cantor) $A \prec \mathcal{P}(A)$.

Proof. $A \preceq \mathcal{P}(A)$ because the map $x \mapsto \{x\}$ defines an injection from A to $\mathcal{P}(A)$. $\mathcal{P}(A) \not\preceq A$ because if $\mathcal{P}(A) \preceq A$, then there would be a bijection from A onto $\mathcal{P}(A)$, contradicting Lemma I.11.6. \square

Applying this theorem with $A = V = \{x : x = x\}$, we get

Paradox I.11.8 (Cantor) $V \prec \mathcal{P}(V)$; but $\mathcal{P}(V) = V$, so $V \prec V$, a contradiction.

Of course, the upshot of this “contradiction” is that there is no universal set, which we already know (Theorem I.6.6), and proved much more simply using Russell's Paradox. Note that Russell's Paradox is just a special case of the diagonal argument: If we apply the proof of Lemma I.11.6 with $A = V = \mathcal{P}(V)$ and f the identity function (so $D = c$), we get $D = \{x : x \notin x\}$, yielding the contradiction $D \in D \leftrightarrow D \notin D$.

The power set operation thus yields larger and larger cardinalities. This is a special case of exponentiation.

Lemma I.11.9 $A^2 \approx \mathcal{P}(A)$.

Proof. $2 = \{0, 1\}$, so associate a subset of A with its characteristic function. \square

It is also easy to see:

Exercise I.11.10 *If $A \preceq B$ and $C \preceq D$ then $A^C \preceq B^D$. If $2 \preceq C$ then $A \prec \mathcal{P}(A) \preceq A^C$.*

The detailed study of cardinal exponentiation will be taken up in Section I.13, using the Axiom of Choice, but without Choice, it is easy to verify the analogs to the familiar laws of exponentiation for natural numbers: $(x^y)^z = x^{y \cdot z}$ and $x^{(y+z)} = x^y \cdot x^z$.

Lemma I.11.11

1. $C^{(B A)} \approx C^{\times B} A$
2. $(B \cup C) A \approx B A \times C A$ if B, C are disjoint.

Proof. For (1), define $\Phi : {}^C(BA) \xrightarrow[\text{onto}]{1-1} {}^{C \times B}A$ by saying that $\Phi(f)(c, b) = (f(c))(b)$. For (2), define $\Psi : ({}^{B \cup C}A) \xrightarrow[\text{onto}]{1-1} {}^B A \times {}^C A$ by saying that $\Psi(f) = (f \upharpoonright B, f \upharpoonright C)$. Of course, for both Φ, Ψ , one must check that their definitions make sense and that they really are bijections. \square

Φ and Ψ are actually “natural” in many settings, where A has some structure on it and exponentiation is defined. For example, if A is a group, and we use the standard group product, then the Φ and Ψ are group isomorphisms. Also, if A is a topological space and we use the standard (Tychonov) product topology, then Φ and Ψ are homeomorphisms.

Definition I.11.12 *A is countable iff $A \preceq \omega$. A is finite iff $A \preceq n$ for some $n \in \omega$. “infinite” means “not finite”. “uncountable” means “not countable”. A is countably infinite iff A is countable and infinite.*

By Theorem I.11.7, $\mathcal{P}(\omega)$ is uncountable. It is easy to prove that $\mathcal{P}(\omega) \approx \mathbb{R} \approx \mathbb{C}$, once \mathbb{R} and \mathbb{C} have been defined (see Section I.15). We say that the sets $\mathcal{P}(\omega), \mathbb{R}, \mathbb{C}$ have size *continuum*, since \mathbb{R} is sometimes referred to as the continuum of real numbers. The following continuation of Exercise I.8.23 is useful when dealing with sizes of sets:

Exercise I.11.13

1. $B \subseteq \alpha \rightarrow \text{type}(B; \in) \leq \alpha$.
2. If $B \preceq \alpha$ then $B \approx \delta$ for some $\delta \leq \alpha$.
3. If $\alpha \leq \beta \leq \gamma$ and $\alpha \approx \gamma$ then $\alpha \approx \beta \approx \gamma$.

By (3), the ordinals come in blocks of the same *size* or *cardinality*. The first ordinal in its block is called a *cardinal* (or a cardinal number). In English grammar, the words “two, three, four” are called cardinal numbers (denoting magnitudes), whereas the words “second, third, fourth” are called ordinal numbers (and are used in concepts involving ordering or sequences). In set theory, they just become the same 2, 3, 4. However, in the infinite case, one distinguishes between ordinals and cardinals:

Definition I.11.14 *A (von Neumann) cardinal is an ordinal α such that $\xi \prec \alpha$ for all $\xi < \alpha$.*

Since $\xi < \alpha \rightarrow \xi \subseteq \alpha \rightarrow \xi \preceq \alpha$, an ordinal α fails to be a cardinal iff there is some $\xi < \alpha$ with $\xi \approx \alpha$.

Theorem I.11.15

1. Every cardinal $\geq \omega$ is a limit ordinal.
2. Every natural number is a cardinal.
3. If A is a set of cardinals, then $\sup(A)$ is a cardinal.
4. ω is a cardinal.

Proof. For (1), assume that $\alpha \geq \omega$ and $\alpha = \delta + 1$ (so $\alpha > \delta \geq \omega$). Then α cannot be a cardinal because we can define a bijection f from $\alpha = \delta \cup \{\delta\}$ onto δ by: $f(\delta) = 0$, $f(n) = n + 1$ for $n \in \omega \subseteq \delta$, and $f(\xi) = \xi$ whenever $\omega \leq \xi < \delta$.

For (2), use ordinary induction (Theorem I.8.14): 0 is a cardinal because Definition I.11.14 is vacuous for $\alpha = 0$. So, we assume that n is a cardinal, and prove that $S(n)$ is a cardinal. If not, fix $\xi < S(n)$ such that $\xi \approx S(n)$. Let $f : \xi \xrightarrow[\text{onto}]{} S(n) = n \cup \{n\}$. If $\xi = \emptyset$, this is impossible, so let $\xi = S(m)$, where $m < n$. Then $f : m \cup \{m\} \xrightarrow[\text{onto}]{} n \cup \{n\}$. If $f(m) = n$, then $f \upharpoonright m : m \xrightarrow[\text{onto}]{} n$, contradicting the assumption that n is a cardinal. If $f(m) = j < n$, then fix $i < m$ such that $f(i) = n$. But then, if we define $g : m \rightarrow n$ so that $g(x) = f(x)$ unless $x = i$, in which case $g(x) = j$, then g is a bijection, so we have the same contradiction.

For (3), if $\sup A = \bigcup A$ fails to be a cardinal, fix some $\xi < \sup A$ such that $\xi \approx \sup A$. Now, fix $\alpha \in A$ such that $\xi < \alpha$. Then $\xi \approx \alpha$ by Exercise I.11.13, contradicting the assumption that all elements of A are cardinals.

(4) is immediate from (2) and (3), setting $A = \omega$. □

Now, we would like to define $|A|$ to be the cardinal κ such that $A \approx \kappa$. However, to prove that there *is* such a κ requires well-ordering A , which requires the Axiom of Choice (AC) (see Section I.12).

Definition I.11.16 *A set A is well-orderable iff there is some relation R which well-orders A .*

Exercise I.11.17 *A can be well-ordered in type α iff there is a bijection $f : A \xrightarrow[\text{onto}]{} \alpha$. Hence, A is well-orderable iff $A \approx \alpha$ for some ordinal α .*

Definition I.11.18 *If A is well-orderable, then $|A|$ is the least ordinal α such that $A \approx \alpha$.*

Clearly, every set of ordinals is well-orderable. In particular, $|\alpha|$ is defined for all α . It is easy to see directly that $|\omega + \omega| = \omega$; more generally, ordinal arithmetic applied to infinite ordinals does not raise the cardinality (see Exercise I.11.34).

Exercise I.11.19 *If A is well-orderable and $f : A \xrightarrow{\text{onto}} B$, then B is well-orderable and $|B| \leq |A|$.*

Exercise I.11.20 *If κ is a cardinal and B is a non-empty set, then $B \preceq \kappa$ iff there is a function $f : \kappa \xrightarrow{\text{onto}} B$.*

In particular, with $\kappa = \omega$, a non-empty B is countable by Definition I.11.12 (i.e., $B \preceq \omega$) iff there is a function $f : \omega \xrightarrow{\text{onto}} B$; that is, you can count B using the natural numbers.

Exercise I.11.21 *Assume that A, B are well-orderable. Then:*

- ♣ $|A|$ is a cardinal.
- ♣ $A \preceq B$ iff $|A| \leq |B|$.
- ♣ $A \approx B$ iff $|A| = |B|$.
- ♣ $A \prec B$ iff $|A| < |B|$.

Exercise I.11.22 A is finite iff A is well-orderable and $|A| < \omega$

Exercise I.11.23 If A and B are finite, then $A \cup B$ and $A \times B$ are finite.

Hint. For $m, n < \omega$, first show $m + n < \omega$ by induction on n , and then show $m \cdot n < \omega$ by induction on n . □

We have not yet produced an uncountable cardinal. We have produced uncountable sets, such as $\mathcal{P}(\omega)$. AC is equivalent to the statement that all sets are well-orderable (see Section I.12), so under AC, the *continuum* $|\mathcal{P}(\omega)|$ is an uncountable cardinal. By Cohen [6, 7] (or, see [18]), ZF does not prove that $\mathcal{P}(\omega)$ is well-orderable. However by the following argument, one can produce uncountable cardinals without using AC.

Theorem I.11.24 (Hartogs, 1915) For every set A , there is a cardinal κ such that $\kappa \not\preceq A$.

Proof. Let W be the set of pairs $(X, R) \in \mathcal{P}(A) \times \mathcal{P}(A \times A)$ such that $R \subseteq X \times X$ and R well-orders X . So, W is the set of all well-orderings of all subsets of A . Observe that $\alpha \preceq X$ iff $\alpha = \text{type}(X; R)$ for some $(X, R) \in W$ (see Exercise I.11.17). Applying the Replacement Axiom, we can set $\beta = \sup\{\text{type}(X, R) + 1 : (X, R) \in W\}$. Then $\beta > \alpha$ whenever $\alpha \preceq A$, so $\beta \not\preceq A$. Let $\kappa = |\beta|$. Then $\kappa \approx \beta$, so $\kappa \not\preceq A$. □

Definition I.11.25 $\aleph(A)$ is the least cardinal κ such that $\kappa \not\preceq A$. For ordinals, α , $\alpha^+ = \aleph(\alpha)$.

Exercise I.11.26 For ordinals, α , α^+ is the least cardinal greater than α .

Exercise I.11.27 The β occurring in the proof of Theorem I.11.24 is already a cardinal. So, $\beta = \kappa = \aleph(A)$, and $\aleph(A)$ is $|A| + 1$ when A is finite and $|A|^+$ when A is infinite and well-orderable.

This $\aleph(A)$ is the *Hartogs aleph function*. It is used most frequently when working in set theory without the Axiom of Choice. Under AC, $|A|$ is always defined, and $\aleph(A)$ is the same as $|A|^+$, which is the more standard notation.

Applying transfinite recursion:

Definition I.11.28 The $\aleph_\xi = \omega_\xi$ are defined by recursion on ξ by:

- $\aleph_0 = \omega_0 = \omega$.

- $\aleph_{\xi+1} = \omega_{\xi+1} = (\aleph_\xi)^+$.
- $\aleph_\eta = \omega_\eta = \sup\{\aleph_\xi : \xi < \eta\}$ when η is a limit ordinal.

Frequently, “ \aleph_ξ ” is used when talking about cardinalities and “ ω_ξ ” is used when talking about order types. The \aleph_ξ list the infinite cardinals in increasing order:

Exercise I.11.29 $\xi < \zeta \rightarrow \aleph_\xi < \aleph_\zeta$. κ is an infinite cardinal iff $\kappa = \aleph_\xi$ for some ξ .

AC is required to get a reasonable theory of cardinals, but some elementary facts can be proved just in ZF . For example, we already know that $\alpha \times \alpha$ can be well-ordered lexicographically. The type is, by definition, $\alpha \cdot \alpha$, which is greater than α whenever $\alpha \geq 2$. However, the cardinality is the same as that of α whenever α is infinite; in particular, $\omega \times \omega$ is countably infinite:

Theorem I.11.30 If $\alpha \geq \omega$ then $|\alpha \times \alpha| = |\alpha|$. Hence, if $\kappa \geq \omega$ is a cardinal, then $|\kappa \times \kappa| = \kappa$.

Proof. It is sufficient to prove the statement for cardinals, since then, with $\kappa = |\alpha|$, we would have $\alpha \times \alpha \approx \kappa \times \kappa \approx \kappa \approx \alpha$. Note that $\kappa \preceq \kappa \times \kappa$ via the map $\xi \mapsto (\xi, \xi)$, but we need to show that $|\kappa \times \kappa|$ can't be bigger than κ .

We shall define a relation \triangleleft on $ON \times ON$ which well-orders $ON \times ON$, and then prove that it well-orders $\kappa \times \kappa$ in order type κ whenever κ is an infinite cardinal. This will yield $\kappa \times \kappa \approx \kappa$.

Define $(\xi_1, \xi_2) \triangleleft (\eta_1, \eta_2)$ to hold iff either $\max(\xi_1, \xi_2) < \max(\eta_1, \eta_2)$ or we have both $\max(\xi_1, \xi_2) = \max(\eta_1, \eta_2)$ and (ξ_1, ξ_2) precedes (η_1, η_2) lexicographically. It is easy to verify that \triangleleft well-orders $ON \times ON$ (in the same sense that \in well-orders ON ; see Theorem I.8.5). If S is a subset of $ON \times ON$, then $\text{type}(S)$ denotes $\text{type}(S, \triangleleft)$. For example, $\text{type}((\omega + 1) \times (\omega + 1)) = \omega \cdot 3 + 1$.

We now need to verify that $\text{type}(\kappa \times \kappa) = \kappa$ whenever κ is an infinite cardinal. We proceed by transfinite induction (using Theorem I.9.1). So, assume that κ is the least infinite cardinal such that $\delta := \text{type}(\kappa \times \kappa) \neq \kappa$, and we derive a contradiction.

If $\alpha < \kappa$ is any ordinal, then $|\alpha \times \alpha| < \kappa$: If α is infinite, this follows because κ is least, so $\text{type}(|\alpha| \times |\alpha|) = |\alpha|$, and hence $|\alpha| \times |\alpha| \approx |\alpha|$; since $|\alpha| \approx \alpha$, we have $\alpha \times \alpha \approx \alpha < \kappa$. If α is finite, then $|\alpha \times \alpha| < \omega \leq \kappa$, by Exercise I.11.23.

Now, let $F : \delta \xrightarrow{\text{onto}} \kappa \times \kappa$ be the isomorphism from $(\delta; <)$ to $(\kappa \times \kappa; \triangleleft)$.

If $\delta > \kappa$, let $(\xi_1, \xi_2) = F(\kappa)$, and let $\alpha = \max(\xi_1, \xi_2) + 1$. Then $\alpha < \kappa$ because κ is a limit ordinal (Theorem I.11.15), and $F''\kappa \subseteq \alpha \times \alpha$ by the definition of \triangleleft , so we have $\kappa \preceq \alpha \times \alpha \prec \kappa$, a contradiction.

If $\delta < \kappa$, then $\kappa \preceq \kappa \times \kappa \approx \delta \prec \kappa$ (since κ is a cardinal), again a contradiction. Thus, $\delta = \kappa$. □

Exercise I.11.31 Every countable strict total ordering is isomorphic to a subset of \mathbb{Q} .

Exercise I.11.32 *The following are equivalent for an ordinal α :*

1. α is isomorphic to a subset of \mathbb{Q} .
2. α is isomorphic to a subset of \mathbb{R} .
3. α is countable.
4. α is isomorphic to a subset of \mathbb{Q} which is closed as a subset of \mathbb{R} .

Hint. Of course, we haven't officially defined \mathbb{Q} , \mathbb{R} yet. (4) \rightarrow (1) \rightarrow (2) is trivial. (3) \rightarrow (4) is done by induction on α (see Subsection I.7.3 for $\alpha = \omega^2$). (2) \rightarrow (3) holds because an uncountable well-ordered subset of \mathbb{R} would yield an uncountable family of disjoint open intervals, which is impossible because \mathbb{Q} is countable and dense in \mathbb{R} .

Regarding (4), say $\alpha \cong K \subset \mathbb{R}$, where K is closed. If α is a limit ordinal, then K is unbounded in \mathbb{R} , while if α is a successor, then K is bounded, and hence compact. A somewhat harder exercise is to show that if X is any compact Hausdorff space and $0 < |X| \leq \aleph_0$, then X is homeomorphic to some successor ordinal. \square

Exercise I.11.33 *Let $\delta = \delta_0$ be any ordinal, and let $\delta_{n+1} = \aleph_{\delta_n}$. Let $\gamma = \sup\{\delta_n : n \in \omega\}$. Show that $\aleph_\gamma = \gamma$. Furthermore, if $\delta_0 = 0$, then γ is the least ordinal ξ such that $\aleph_\xi = \xi$.*

Exercise I.11.34 *Ordinal arithmetic doesn't raise cardinality. That is, assume that α, β are ordinals with $2 \leq \min(\alpha, \beta)$ and $\omega \leq \max(\alpha, \beta)$. Then prove that $|\alpha + \beta| = |\alpha \cdot \beta| = |\alpha^\beta| = \max(|\alpha|, |\beta|)$.*

Hint. For $\alpha + \beta$ and $\alpha \cdot \beta$, this is easy from Theorem I.11.30 and the definitions of $+$ and \cdot . The proof is slightly tricky for exponentiation, since the natural proof by induction on β seems to use the Axiom of Choice (AC). Here, we take the recursive computation in Table I.1 to be the definition of α^β . Now, say we want to prove by induction on β that α^β is countable whenever α, β are countable. If β is a limit, we have $\alpha^\beta = \sup_{\xi < \beta} (\alpha^\xi) = \bigcup_{\xi < \beta} (\alpha^\xi)$ which (applying induction) is a countable union of countable sets. But the well-known fact that a countable union of countable sets is countable (Theorem I.12.12) uses AC. To avoid AC, first fix a $\delta < \omega_1$ and then fix an $f : \delta \xrightarrow{1-1} \omega$. Now, we may now define, for $\alpha, \beta < \delta$, an injection from α^β into ω ; the definition uses f , and is done by recursion on β ; the definition also uses a (fixed) injection from $\omega \times \omega$ into ω .

A similar method of proof shows in ZF that other ordinal arithmetic functions defined by recursion don't raise cardinality. For example, we may define a hyperexponential function h so that $h(\alpha, 3) = \alpha^{\alpha^{\alpha}}$; $h(\alpha, 0) = \alpha$; $h(\alpha, S(\beta)) = \alpha^{h(\alpha, \beta)}$; $h(\alpha^\gamma) = \sup_{\beta < \gamma} h(\alpha, \beta)$ for γ a limit. Then $h(\alpha, \beta)$ is countable when α, β are countable. \square

Exercise I.11.35 *Prove within the theory Z^- (see Section I.2) that there is an uncountable well-ordered set.*

Hint. Let $A = \omega$ and then form W as in the proof of Theorem I.11.24. Let $B = W/\cong$, where \cong is the isomorphism relation, and define an appropriate well-order on B . See Exercise I.10.5 for some historical remarks, and for how to construct a quotient without using the Axiom of Replacement. Note that Hartogs in 1915 could not have stated Theorem I.11.24 as we did here, since this was before von Neumann ordinals (1923) and the introduction of the Axiom of Replacement (1922). \square

I.12 The Axiom of Choice (AC)

There are many known equivalents to AC. For a detailed discussion, see Howard and Rubin [16]. We confine ourselves here to the versions of AC most likely to be useful in mathematics. This boils down to:

Theorem I.12.1 (ZF) *The following are equivalent*

1. *The Axiom of Choice (as in Section I.2).*
2. *Every set has a choice function.*
3. *Every set can be well-ordered.*
4. $\forall xy(x \preceq y \vee y \preceq x)$.
5. *Tukey's Lemma.*
6. *The Hausdorff Maximal Principle.*
7. *Zorn's Lemma.*

(2) (5) (6) (7) will be defined in this section, and proofs of equivalence given, although the equivalences involving (6) and (7) will be left to the exercises; these two are important and well-known, but they are never used in this book, except in the exercises.

The form of AC in Section I.2 was used because it requires few definitions to state:

$$\emptyset \notin F \wedge \forall x \in F \forall y \in F (x \neq y \rightarrow x \cap y = \emptyset) \rightarrow \exists C \forall x \in F (\text{SING}(C \cap x))$$

That is, whenever F is a disjoint family of non-empty sets, there is a set C such that $C \cap x$ is a singleton set for all $x \in F$. C is called a *choice set* for F because it *chooses* one element from each set in x . If you draw a picture, this principle will seem “obviously true”, although, since F may be infinite, it needs to be stated as a separate principle (as observed by Zermelo (1908), and proved by Cohen (1963)).

In practice, this version is not very useful, since one frequently needs to choose elements from sets which are not disjoint, in which case one must do the choosing by a *choice function*:

Definition I.12.2 *A choice function for a set A is a function $g : \mathcal{P}(A) \setminus \{\emptyset\} \rightarrow A$ such that $g(x) \in x$ for all $x \in \mathcal{P}(A) \setminus \{\emptyset\}$.*

Proof of Theorem I.12.1 (1) \leftrightarrow (2) .

For (2) \rightarrow (1): Given a disjoint family F of non-empty sets: Let $A = \bigcup F$, let g be a choice function for A , and let $C = \{g(x) : x \in F\}$.

For (1) \rightarrow (2): Given any set A , let $F = \{\{x\} \times x : x \in \mathcal{P}(A) \setminus \{\emptyset\}\}$. By this trick, we take each non-empty $x \subseteq A$ and form a copy of it, $\{x\} \times x = \{(x, i) : i \in x\}$. If $x \neq y$, then $\{x\} \times x$ and $\{y\} \times y$ are disjoint. Hence, by (1), there is a choice set C for F . But then $C : \mathcal{P}(A) \setminus \{\emptyset\} \rightarrow A$ is a choice function for A . \square

A choice function g for A gives you a way to well-order A . Informally, we list A in a transfinite sequence as described in Section I.5:

$$a_0, a_1, a_2, a_3, \dots$$

Get a_0 by choosing some element of A , then choose some a_1 different from a_0 , then choose some a_2 different from a_0, a_1 , then choose some a_3 different from a_0, a_1, a_2 , and so on. If A is infinite, this will require an infinite number of choices, and the “choosing” is justified formally by a choice function g :

Proof of Theorem I.12.1 (2) \leftrightarrow (3) \leftrightarrow (4) .

For (3) \rightarrow (2): If R well-orders A , we can define a choice function g by letting $g(x)$ be the R -least element of x .

For (2) \rightarrow (3): Conversely, let g be a choice function for A . Fix $\kappa = \aleph(A)$. Fix any $S \notin A$. Think of S as signifying “Stop”. Define $f : \kappa \rightarrow A \cup \{S\}$ so that $f(\alpha) = g(A \setminus \{f(\xi) : \xi < \alpha\})$ if $A \setminus \{f(\xi) : \xi < \alpha\}$ is non-empty; otherwise, $f(\alpha) = S$. Observe that $f(\xi) \neq f(\alpha)$ whenever $\xi < \alpha$ and $f(\alpha) \neq S$. Since $\kappa \not\preceq A$, there can be no $f : \kappa \xrightarrow{1-1} A$, so $f(\alpha) = S$ for some α . Now, let α be least such that $f(\alpha) = S$, and note that $f \upharpoonright \alpha : \alpha \xrightarrow{1-1}_{\text{onto}} A$; thus, A can be well-ordered in type α (see Exercise I.11.17).

For (4) \rightarrow (3): If $\kappa = \aleph(A)$, then $\kappa \not\preceq A$, so (4) implies that $A \preceq \kappa$, so A can be well-ordered (see Exercise I.11.13).

For (3) \rightarrow (4): If x and y can be well-ordered, then $|x|$ and $|y|$ are von Neumann cardinals, and $|x| \leq |y|$ or $|y| \leq |x|$. \square

Note that we have actually shown, for each fixed set A , that A is well-orderable iff A has a choice function.

We have now shown the equivalence of (1), (2), (3), (4) of Theorem I.12.1. The other three principles are all *maximal principles*. The most well-known among these is (7), *Zorn’s Lemma*. This is often quoted, for example, in the proof that every vector space has a basis. However, in many of the elementary applications, such as this one, Tukey’s Lemma (5) is a more convenient maximal principle to use, so we shall begin with (5), which involves the existence of a maximal set in a family of sets:

Definition I.12.3 If $\mathcal{F} \subseteq \mathcal{P}(A)$, then $X \in \mathcal{F}$ is maximal in \mathcal{F} iff it is maximal with respect to the relation \subsetneq (see Definition I.7.19); that is, X is not a proper subset of any set in \mathcal{F} .

Some examples: If $\mathcal{F} = \mathcal{P}(A)$, then A is maximal in \mathcal{F} . If \mathcal{F} is the set of finite subsets of A , and A is infinite, then \mathcal{F} has no maximal element. A less trivial example is:

Example I.12.4 *Let A be any vector space over some field. Define $\mathcal{F} \subseteq \mathcal{P}(A)$ so that $X \in \mathcal{F}$ iff X is linearly independent. Then X is maximal in \mathcal{F} iff X is a basis.*

The proof of this is an easy exercise in linear algebra, and does not use the Axiom of Choice. Note that linearly independent means that there is no set of vectors $\{x_1, \dots, x_n\} \subseteq X$ (where $1 \leq n < \omega$) and non-zero scalars a_1, \dots, a_n in the field such that $a_1x_1 + \dots + a_nx_n = 0$. If X is linearly independent and is not a basis (that is, fails to span A), then there is some $y \in A$ which is not in the linear span of X , which implies (by linear algebra) that $X \cup \{y\}$ is also linearly independent, so X is not maximal.

The feature which implies that this particular \mathcal{F} has a maximal element is that \mathcal{F} is of *finite character*.

Definition I.12.5 $\mathcal{F} \subseteq \mathcal{P}(A)$ is of finite character iff for all $X \subseteq A$: $X \in \mathcal{F}$ iff every finite subset of X is in \mathcal{F} .

Exercise I.12.6 If \mathcal{F} is of finite character, $X \in \mathcal{F}$, and $Y \subseteq X$, then $Y \in \mathcal{F}$.

Definition I.12.7 Tukey's Lemma is the assertion that whenever $\mathcal{F} \subseteq \mathcal{P}(A)$ is of finite character and $X \in \mathcal{F}$, there is a maximal $Y \in \mathcal{F}$ such that $X \subseteq Y$.

The \mathcal{F} in Example I.12.4 is of finite character because the notion of "linearly independent" is verified by looking only at finite subsets of X . Hence, Tukey's Lemma implies that every vector space has a basis. Here, as in many applications, X can be \emptyset . Applying Tukey's Lemma with an arbitrary X shows that every linearly independent set can be expanded to a basis.

This type of proof is common in undergraduate texts. The advantage of it is that you can follow the proof without having to know about ordinals and transfinite recursion; Tukey's Lemma or Zorn's Lemma can be taken as an axiom. The disadvantage is that these maximal principles are a bit complicated, and it's hard to explain why they should be axioms. The standard choice versions of AC, (1) or (2), have a much clearer intuitive motivation. Of course, once you know about ordinals and recursion, the proof of Tukey's lemma is quite simple:

Proof of Theorem I.12.1 (3) \rightarrow (5) \rightarrow (1) .

For (3) \rightarrow (5): Fix A, \mathcal{F}, X as in Definition I.12.7. Since A can be well-ordered, we can list A as $\{x_\alpha : \alpha < \kappa\}$, where $\kappa = |A|$. Recursively define $Y_\beta \subseteq \{x_\xi : \xi < \beta\}$, for $\beta \leq \kappa$, by:

- a. $Y_0 = X$.
- b. $Y_{\alpha+1}$ is $Y_\alpha \cup \{x_\alpha\}$ if $Y_\alpha \cup \{x_\alpha\} \in \mathcal{F}$; otherwise, $Y_{\alpha+1} = Y_\alpha$.

c. $Y_\gamma = \bigcup \{Y_\alpha : \alpha < \gamma\}$ if γ is a limit ordinal.

Now, check, inductively, that $Y_\beta \in \mathcal{F}$ for each $\beta \leq \kappa$. For the successor step, use (b). For limit β , use (c) and the fact that \mathcal{F} is of finite character. Let $Y = Y_\kappa$. Then $Y \in \mathcal{F}$. To see that Y is maximal, fix $x_\alpha \notin Y$. Then, by (b), $Y_\alpha \cup \{x_\alpha\} \notin \mathcal{F}$, so the superset $Y \cup \{x_\alpha\}$ is also not in \mathcal{F} (see Exercise I.12.6).

For (5) \rightarrow (1): Let F be any family of disjoint non-empty sets. We need to produce a choice set C , which intersects each $z \in F$ in a singleton. Let $A = \bigcup F$. Let \mathcal{G} be the set of all *partial choice sets* for F ; so $X \in \mathcal{G}$ iff $X \in \mathcal{P}(A)$ and $X \cap z$ is either a singleton or empty for all $z \in F$. \mathcal{G} is of finite character because if $X \subseteq A$ fails to be in \mathcal{G} , then some two-element subset of X fails to be in \mathcal{G} . Also, $\emptyset \in \mathcal{G}$, so applying Tukey's Lemma, fix a maximal $C \in \mathcal{G}$. If C is not a choice set for F , we can fix a $z \in F$ such that $C \cap z$ is not a singleton, and is hence empty (since $C \in \mathcal{G}$). But then, since $z \neq \emptyset$, we can fix $p \in z$ and note that $C \subsetneq C \cup \{p\} \in \mathcal{G}$, contradicting maximality. \square

We have now proved the equivalence of (1), (2), (3), (4), (5), in Theorem I.12.1. We shall now state (6) and (7), and leave the proofs of equivalence to the exercises. Both these equivalents involve partial orders. We shall phrase these in terms of strict partial orders $<$ (as in Definition I.7.2); then $x \leq y$ abbreviates $x < y \vee x = y$.

Definition I.12.8 Let $<$ be a strict partial order of a set A . Then $C \subseteq A$ is a chain iff C is totally ordered by $<$; C is a maximal chain iff in addition, there are no chains $X \supsetneq C$.

Definition I.12.9 The Hausdorff Maximal Principle is the assertion that whenever $<$ is a strict partial order of a set A , there is a maximal chain $C \subseteq A$.

Definition I.12.10 Zorn's Lemma is the assertion that whenever $<$ is a strict partial order of a set A satisfying

(\clubsuit) For all chains $C \subseteq A$ there is some $b \in A$ such that $x \leq b$ for all $x \in C$,

then for all $a \in A$, there is a maximal (see Definition I.7.19) $b \in A$ with $b \geq a$.

Exercise I.12.11 Finish the proof of Theorem I.12.1.

Hint. One way to do this is to show (5) \rightarrow (6) \rightarrow (7) \rightarrow (5); this stays within the maximal principles. For (5) \rightarrow (6), note that the family of all chains has finite character. For (6) \rightarrow (7), fix a maximal chain C containing a ; then the b we get in (\clubsuit) is a maximal element. For (7) \rightarrow (5), observe that if \mathcal{F} is of finite character, then \mathcal{F} , partially ordered by \subsetneq , satisfies the hypothesis (\clubsuit) of Zorn's Lemma. \square

In "pure" set theory, the most frequently used forms of the Axiom of Choice are (2), guaranteeing the existence of choice functions, and (3) (the well-ordering principle). In algebra, analysis, and model theory, frequent use is made of Zorn's Lemma (7) and/or

Tukey's Lemma (5) as well. Often, Tukey's Lemma is easier to apply, since one has a family which is obviously of finite character, as in Example I.12.4. However, there are some cases, such as Exercise I.12.14 below, where Zorn's Lemma is more useful.

Texts in algebra and analysis usually point out the use of AC when they prove that every vector space has a basis or that there is a subset of \mathbb{R} which is not Lebesgue measurable. However, some more elementary uses of AC are often glossed over without explicit mention. For example:

1. In elementary calculus, you learn that if 0 is a limit point of a set X of reals, then there is a sequence $\langle x_n : n \in \omega \rangle$ from X converging to 0. The usual proof *chooses* $x_n \in X \cap (-1/n, 1/n)$. It is known to be consistent with ZF to have an X dense in \mathbb{R} with $\omega \not\prec X$, so all ω -sequences from X are eventually constant.

2. A countable union of countable sets is countable. This is true under AC (see Theorem I.12.12 below). It is consistent with ZF that \mathbb{R} is a countable union of countable sets; it is also consistent that ω_1 is a countable union of countable sets; see [19] for more on consistency results involving $\neg AC$.

Theorem I.12.12 (AC) *Let κ be an infinite cardinal. If \mathcal{F} is a family of sets with $|\mathcal{F}| \leq \kappa$ and $|X| \leq \kappa$ for all $X \in \mathcal{F}$, then $|\bigcup \mathcal{F}| \leq \kappa$.*

Proof. Assume $\mathcal{F} \neq \emptyset$ (otherwise the result is trivial) and $\emptyset \notin \mathcal{F}$ (since removing \emptyset from \mathcal{F} does not change the union). Then (see Exercise I.11.20), fix $f : \kappa \xrightarrow{\text{onto}} \mathcal{F}$. Likewise, for each $B \in \mathcal{F}$, there are functions $g : \kappa \xrightarrow{\text{onto}} B$. By AC, choose $g_\alpha : \kappa \xrightarrow{\text{onto}} f(\alpha)$ for all $\alpha < \kappa$ (to do this, well-order $\mathcal{S} := {}^\kappa(\bigcup \mathcal{F})$, and let g_α be the least $g \in \mathcal{S}$ with $\text{ran}(g) = f(\alpha)$). This defines $h : \kappa \times \kappa \xrightarrow{\text{onto}} \bigcup \mathcal{F}$, where $h(\alpha, \beta) = g_{f(\alpha)}(\beta)$. Since $|\kappa \times \kappa| = \kappa$ (see Theorem I.11.30), we can map κ onto $\bigcup \mathcal{F}$, so $|\bigcup \mathcal{F}| \leq \kappa$ (see Exercise I.11.20). \square

Informally, it's "obvious" that if you have a map $f : A \xrightarrow{\text{onto}} B$, then A must be at least as large as B (i.e., $B \preceq A$), but producing a $g : B \xrightarrow{1-1} A$ requires AC. Concretely, ZF does not prove that there is an injection from ω_1 into $\mathcal{P}(\omega)$, but

Exercise I.12.13 *Prove, without using AC, that one can map $\mathcal{P}(\omega)$ onto ω_1 .*

Hint. First, define $f : \mathcal{P}(\omega \times \omega) \xrightarrow{\text{onto}} \omega_1 \setminus \omega$ so that $f(R) = \text{type}(R)$ whenever R is a well-order of ω . \square

Exercise I.12.14 *If X, Y are compact Hausdorff spaces, a continuous map $f : X \rightarrow Y$ is called irreducible iff $f : X \xrightarrow{\text{onto}} Y$, but $f(H) \neq Y$ for all proper closed subsets H of X . For example, if $Y = [0, 2] \subset \mathbb{R}$ and $X = [0, 1] \times \{0\} \cup [1, 2] \times \{1\} \subset \mathbb{R} \times \mathbb{R}$, then the usual projection map is irreducible. Now, fix any continuous $f : X \xrightarrow{\text{onto}} Y$, and prove that there is a closed $Z \subseteq X$ such that $f \upharpoonright Z : Z \xrightarrow{\text{onto}} Y$ and $f \upharpoonright Z$ is irreducible.*

Hint. Use Zorn's Lemma and get Z minimal in the set of all closed $Z \subseteq X$ such that $f \upharpoonright Z$ is onto. See Engelking's text [11] for more on the use of irreducible maps in general topology. \square

I.13 Cardinal Arithmetic

In this section, we assume AC throughout. Then, since every set can be well-ordered, $|x|$ is defined for all x . We define cardinal addition, multiplication, and exponentiation by:

Definition I.13.1 *If κ, λ are cardinals, then*

- $\boxed{\kappa + \lambda} = |\{0\} \times \kappa \cup \{1\} \times \lambda|$
- $\boxed{\kappa \cdot \lambda} = |\kappa \times \lambda|$
- $\boxed{\kappa^\lambda} = |{}^\lambda \kappa|$

The boxes are omitted when it is clear from context that cardinal arithmetic is intended. In particular, if a cardinal is listed as “ \aleph_α ”, then cardinal arithmetic is intended.

In the literature, boxes are never used, and the reader must determine from context whether ordinal or cardinal arithmetic is intended. The context is important because there are three possible meanings to “ κ^λ ”, and all are used: the ordinal exponent, the cardinal exponent (i.e., $\boxed{\kappa^\lambda}$), and the set of functions (i.e., ${}^\lambda \kappa$). To avoid too many boxes in our notation, we shall often phrase results by saying in English which operation is intended. This is done in the following two lemmas; the first says that the cardinal functions are monotonic in each argument; the second lists some elementary arithmetic facts which are well-known for finite cardinals.

Lemma I.13.2 *If $\kappa, \lambda, \kappa', \lambda'$ are cardinals and $\kappa \leq \kappa'$ and $\lambda \leq \lambda'$, then $\kappa + \lambda \leq \kappa' + \lambda'$, $\kappa \cdot \lambda \leq \kappa' \cdot \lambda'$, and $\kappa^\lambda \leq (\kappa')^{\lambda'}$ (unless $\kappa = \kappa' = \lambda = 0$), where cardinal arithmetic is meant throughout.*

Proof. For the first two inequalities, use $\{0\} \times \kappa \cup \{1\} \times \lambda \subseteq \{0\} \times \kappa' \cup \{1\} \times \lambda'$ and $|\kappa \times \lambda| \subseteq |\kappa' \times \lambda'|$. For the third, when $\kappa' > 0$, define $\varphi : {}^\lambda \kappa \xrightarrow{1-1} {}^{\lambda'} (\kappa')$ by: $\varphi(f) \upharpoonright \lambda = f$ and $(\varphi(f))(\xi) = 0$ for $\lambda \leq \xi < \lambda'$. When $\kappa = \kappa' = 0$, note that $0^0 = |{}^0 0| = |\{\emptyset\}| = 1$ (the empty function maps \emptyset to \emptyset), while for $\lambda > 0$, $0^\lambda = |{}^\lambda 0| = |\emptyset| = 0$. \square

Note that $0^0 = 1$ in ordinal exponentiation as well.

Lemma I.13.3 *If κ, λ, θ are cardinals, then using cardinal arithmetic throughout:*

1. $\kappa + \lambda = \lambda + \kappa$.
2. $\kappa \cdot \lambda = \lambda \cdot \kappa$.
3. $(\kappa + \lambda) \cdot \theta = \kappa \cdot \theta + \lambda \cdot \theta$.
4. $\kappa^{(\lambda \cdot \theta)} = (\kappa^\lambda)^\theta$.
5. $\kappa^{(\lambda + \theta)} = \kappa^\lambda \cdot \kappa^\theta$.

Proof. For (1,2,3), note that for any sets A, B, C : $A \cup B = B \cup A$, $A \times B \approx B \times A$, and $(A \cup B) \times C = A \times B \cup A \times C$. For (4,5), apply Lemma I.11.11. \square

We still need to use boxes in statements which involving both cardinal and ordinal operations. An example is the following, which is immediate from the definitions of the cardinal and ordinal sum and product.

Lemma I.13.4 *For any ordinals α, β : $|\alpha + \beta| = \boxed{|\alpha| + |\beta|}$ and $|\alpha \cdot \beta| = \boxed{|\alpha| \cdot |\beta|}$.*

An example when $\alpha = \beta = \omega = |\omega|$: $\omega, \omega + \omega$, and $\omega \cdot \omega$ are three different ordinals, but these three ordinals have the same cardinality. This lemma fails for exponentiation, since (see Exercise I.11.34) ω^ω is a countable ordinal but $\boxed{\omega^\omega} = (\aleph_0)^{\aleph_0} = 2^{\aleph_0}$ (by Lemma I.13.9), which is uncountable.

Lemma I.13.5 *If κ, λ are finite cardinals, then $\boxed{\kappa + \lambda} = \kappa + \lambda$, $\boxed{\kappa \cdot \lambda} = \kappa \cdot \lambda$, and $\boxed{\kappa^\lambda} = \kappa^\lambda$.*

Proof. Since finite ordinals are cardinals (by Theorem I.11.15), we have $\kappa + \lambda = |\kappa + \lambda| = \boxed{\kappa + \lambda}$ by Lemma I.13.4. The same argument works for product.

The fact that we know the lemma for sum and product lets us prove it for exponentiation by induction. The $\lambda = 0$ and $\lambda = 1$ cases are easy because $\kappa^0 = 1$ and $\kappa^1 = \kappa$ with both cardinal and ordinal exponentiation. Now, assume we know that $\boxed{\kappa^\lambda} = \kappa^\lambda$; in particular, $\boxed{\kappa^\lambda}$ is finite. Now, since we already know that the cardinal and ordinal sums and products agree for finite ordinals, Lemma I.13.3 gives us $\boxed{\kappa^{\lambda+1}} = \boxed{\kappa^\lambda} \cdot \kappa$. Since we also know that $\kappa^{\lambda+1} = \kappa^\lambda \cdot \kappa$, we can conclude that $\boxed{\kappa^{\lambda+1}} = \kappa^{\lambda+1}$. \square

Lemma I.13.6 *If κ, λ are cardinals and at least one of them is infinite, then $\boxed{\kappa + \lambda} = \max(\kappa, \lambda)$. Also, if neither of them are 0 then $\boxed{\kappa \cdot \lambda} = \max(\kappa, \lambda)$.*

Proof. Assume that $\kappa \leq \lambda$, so λ is infinite. It is clear from the definitions that $\lambda \preceq \boxed{\kappa + \lambda} \preceq \lambda \times \lambda$, so use the fact that $\lambda \times \lambda \approx \lambda$ (by Theorem I.11.30). The same argument works for product when $\kappa \neq 0$. \square

Because of these lemmas, cardinal sums and products are not very interesting, since they reduce to ordinal arithmetic for finite cardinals and they just compute the max for infinite cardinals; the sum and product notation is useful mainly for making general statements, such as Lemma I.13.3. Cardinal exponentiation, however, is of fundamental importance, since it is related to the Continuum Hypothesis. By Lemma I.11.9, and Theorem I.11.7, we have:

Lemma I.13.7 $2^\kappa = |\mathcal{P}(\kappa)|$ for every cardinal κ , and $2^{\aleph^\alpha} \geq \aleph_{\alpha+1}$ for every ordinal α . All exponentiation here is cardinal exponentiation.

Definition I.13.8 *The Continuum Hypothesis, CH, is the statement $2^{\aleph_0} = \aleph_1$. The General Continuum Hypothesis, GCH, is the statement $\forall \alpha [2^{\aleph_\alpha} = \aleph_{\alpha+1}]$. All exponentiation here is cardinal exponentiation.*

Knowing values of 2^λ for the infinite λ tells us a lot about all the powers κ^λ . In particular, the GCH implies a very simple formula for κ^λ (see Theorem I.13.14), which we shall obtain by examining the possible cases. Of course, $1^\lambda = 1$. Also:

Lemma I.13.9 *If $2 \leq \kappa \leq 2^\lambda$ and λ is infinite, then $\kappa^\lambda = 2^\lambda$. All exponentiation here is cardinal exponentiation.*

Proof. Applying Lemmas I.13.2, I.13.3, and I.13.6, we have $2^\lambda \leq \kappa^\lambda \leq (2^\lambda)^\lambda = 2^{\lambda \cdot \lambda} = 2^\lambda$. \square

So, infinite cardinal arithmetic is simpler than finite cardinal arithmetic. However, there are further complexities about κ^λ when $\lambda < \kappa$. Assuming GCH, Theorem I.13.14 will yield $(\aleph_3)^{\aleph_0} = \aleph_3$ and $(\aleph_{\omega_1})^{\aleph_0} = \aleph_{\omega_1}$, but $(\aleph_\omega)^{\aleph_0} = \aleph_{\omega+1}$. The key feature about \aleph_ω here is that it has countable cofinality:

Definition I.13.10 *If γ is any limit ordinal, then the cofinality of γ is*

$$\text{cf}(\gamma) = \min\{\text{type}(X) : X \subseteq \gamma \wedge \sup(X) = \gamma\} .$$

γ is regular iff $\text{cf}(\gamma) = \gamma$.

For example, consider $\gamma = \omega^2$. There are many $X \subseteq \gamma$ such that $\sup(X) = \gamma$ (that is, X is unbounded in γ). $\omega \cup \{\omega \cdot n : n \in \omega\}$ is unbounded in γ and has order type $\omega \cdot 2$. The set $\{\omega \cdot n : n \in \omega\}$ is also unbounded in γ , and this one has order type ω . This ω is the least possible order type since if $X \subseteq \gamma$ is finite, then X has a largest element, so it cannot be unbounded. Thus, $\text{cf}(\gamma) = \omega$. As an aid to computing cofinalities:

Lemma I.13.11 *For any limit ordinal γ :*

1. *If $A \subseteq \gamma$ and $\sup(A) = \gamma$ then $\text{cf}(\gamma) = \text{cf}(\text{type}(A))$.*
2. *$\text{cf}(\text{cf}(\gamma)) = \text{cf}(\gamma)$, so $\text{cf}(\gamma)$ is regular.*
3. *If γ is regular then γ is a cardinal.*
4. $\omega \leq \text{cf}(\gamma) \leq |\gamma| \leq \gamma$
5. *If $\gamma = \aleph_\alpha$ where α is either 0 or a successor, then γ is regular.*
6. *If $\gamma = \aleph_\alpha$ where α is a limit ordinal, then $\text{cf}(\gamma) = \text{cf}(\alpha)$.*

Proof.

For (1): Let $\alpha = \text{type}(A)$. Note that α is a limit ordinal since A is unbounded in γ . Let $f : \alpha \xrightarrow[\text{onto}]{1-1} A$ be the isomorphism from α onto A . To prove that $\text{cf}(\gamma) \leq \text{cf}(\alpha)$, note that if Y is an unbounded subset of α , then $f \upharpoonright Y$ is an unbounded subset of γ of

the same order type; in particular, taking Y to be of type $\text{cf}(\alpha)$ produces an unbounded subset of γ of type $\text{cf}(\alpha)$. To prove that $\text{cf}(\alpha) \leq \text{cf}(\gamma)$, fix an unbounded $X \subseteq \gamma$ of order type $\text{cf}(\gamma)$. For $\xi \in X$, let $h(\xi)$ be the least element of A which is $\geq \xi$. Note that $\xi < \eta \rightarrow h(\xi) \leq h(\eta)$. Let $X' = \{\eta \in X : \forall \xi \in X \cap \eta [h(\xi) < h(\eta)]\}$. Then $h \upharpoonright X' : X' \xrightarrow{1-1} A$ and $h \upharpoonright X'$ is order preserving. Also, $h(X')$ is unbounded in A , which has order type α , so $\text{cf}(\alpha) \leq \text{type}(X') \leq \text{type}(X) = \text{cf}(\gamma)$ (using Exercise I.8.23).

For (2): Let A be an unbounded subset of γ of order type $\text{cf}(\gamma)$, and apply (1).

For (3): Suppose $\kappa = |\gamma| < \gamma$, and let $f : \kappa \xrightarrow{\text{onto}} \gamma$; we show that $\text{cf}(\gamma) \leq \kappa$. Define (recursively) a function $g : \kappa \rightarrow ON$ so that $g(\eta) = \max(f(\eta), \sup\{g(\xi) + 1 : \xi < \eta\})$. Then $\xi < \eta \rightarrow g(\xi) < g(\eta)$, so g is an isomorphism onto its range. If $\text{ran}(g) \subseteq \gamma$, then $\text{ran}(g)$ is a subset of γ of order type κ , and $\text{ran}(g)$ is unbounded (since each $g(\eta) \geq f(\eta)$), so $\text{cf}(\gamma) \leq \kappa$. If $\text{ran}(g) \not\subseteq \gamma$, let η be the least ordinal such that $g(\eta) \leq \gamma$. Then η is a limit ordinal, since $g(\xi + 1) = \max(g(\xi) + 1, f(\xi + 1))$. Thus, $g \upharpoonright \eta$ is an unbounded subset of γ of order type η , so $\text{cf}(\gamma) \leq \eta < \kappa$.

In (4), $\omega \leq \text{cf}(\gamma) \leq \gamma$ and $|\gamma| \leq \gamma$ are clear from the definitions, and $\text{cf}(\gamma) \leq |\gamma|$ follows from (2) and (3)

For (5): $\aleph_0 = \omega$ is regular by (4). $\aleph_{\beta+1}$ is regular because if $A \subseteq \aleph_{\beta+1}$ and $|A| \leq \aleph_\beta$, then $\sup(A) = \bigcup A$ is the union of $\leq \aleph_\beta$ sets (ordinals), each of size $\leq \aleph_\beta$, so that $|\sup(A)| \leq \aleph_\beta$ by Theorem I.12.12

For (6): Apply (1), with $A = \{\aleph_\xi : \xi < \alpha\}$. □

By (1), $\text{cf}(\alpha + \beta) = \beta$ (let $A = \{\alpha + \xi : \xi < \beta\}$). By (4), every limit ordinal below ω_1 has cofinality ω . Likewise, every limit ordinal below ω_2 has cofinality either ω (for example, $\omega_1 + \omega$) or ω_1 (for example, $\omega_1 + \omega_1$).

By (3), regular ordinals are cardinals. The following fact about unions generalizes Theorem I.12.12 in the case that $\theta = \kappa^+$.

Theorem I.13.12 *Let θ be any cardinal.*

1. *If θ is regular and \mathcal{F} is a family of sets with $|\mathcal{F}| < \theta$ and $|S| < \theta$ for all $S \in \mathcal{F}$, then $|\bigcup \mathcal{F}| < \theta$.*
2. *If $\text{cf}(\theta) = \lambda < \theta$, then there is a family \mathcal{F} of subsets of θ with $|\mathcal{F}| = \lambda$ and $\bigcup \mathcal{F} = \theta$, and $|S| < \theta$ for all $S \in \mathcal{F}$.*

Proof. For (1), let $X = \{|S| : S \in \mathcal{F}\}$. Then $X \subseteq \theta$ and $|X| < \theta$, so $\text{type}(X) < \theta$, and hence $\sup(X) < \theta$. Let $\kappa = \max(\sup(X), |\mathcal{F}|) < \theta$. If κ is infinite, then by Theorem I.12.12, $|\bigcup \mathcal{F}| \leq \kappa$. If κ is finite, then $\bigcup \mathcal{F}$ is finite. In either case, $|\bigcup \mathcal{F}| < \theta$.

For (2), let \mathcal{F} be a subset of θ of order type λ such that $\sup \mathcal{F} = \bigcup \mathcal{F} = \lambda$. □

By generalizing Cantor's diagonal argument for proving $2^\lambda > \lambda$, we get:

Theorem I.13.13 (König, 1905) *If $\kappa \geq 2$ and λ is infinite, then $\text{cf}(\kappa^\lambda) > \lambda$.*

Proof. Let $\theta = \kappa^\lambda$, which must be infinite. Note that $\theta^\lambda = \kappa^{\lambda \cdot \lambda} = \kappa^\lambda = \theta$. List ${}^\lambda\theta$ as $\{f_\alpha : \alpha < \theta\}$. If $\text{cf}(\theta) \leq \lambda$, we can write the ordinal θ as $\theta = \bigcup_{\xi < \lambda} A_\xi$, where $|A_\xi| < \theta$. But now define $g : \lambda \rightarrow \theta$ so that $g(\xi) = \min(\theta \setminus \{f_\alpha(\xi) : \alpha \in A_\xi\})$. Then $g \in {}^\lambda\theta$ and g differs from every f_α , a contradiction. \square

For example, $\text{cf}(2^{\aleph_0}) > \omega$, so that 2^{\aleph_0} cannot be \aleph_ω . Roughly, it is consistent for 2^{\aleph_0} to be any cardinal of uncountable cofinality, such as \aleph_1 (*CH*), or \aleph_7 , or $\aleph_{\omega+1}$, or \aleph_{ω_1} ; see [6, 7, 18, 20].

The following lemma tells how to compute κ^λ under *GCH* when at least one of κ, λ are infinite. It shows that it is the smallest possible value, given König's Theorem. We omit listing some trivial cases, when one of them is 0 or 1:

Theorem I.13.14 *Assume GCH, and let κ, λ be cardinals with $\max(\kappa, \lambda)$ infinite.*

1. *If $2 \leq \kappa \leq \lambda^+$, then $\kappa^\lambda = \lambda^+$.*
2. *If $1 \leq \lambda \leq \kappa$, then κ^λ is κ if $\lambda < \text{cf}(\kappa)$ and κ^+ if $\lambda \geq \text{cf}(\kappa)$.*

Proof. Part (1) is immediate from Lemma I.13.9. For (2), we have, applying *GCH*: $\kappa \leq \lambda \leq \kappa \leq \kappa^\kappa = 2^\kappa = \kappa^+$, so κ^λ is either κ or κ^+ . If $\lambda \geq \text{cf}(\kappa)$, then κ^λ cannot be κ by König's Theorem, so it must be κ^+ . If $\lambda < \text{cf}(\kappa)$, then every $f : \lambda \rightarrow \kappa$ is bounded, so that ${}^\lambda\kappa = \bigcup_{\alpha < \kappa} {}^\lambda\alpha$. By *GCH*, each $|{}^\lambda\alpha| \leq (\max(|\alpha|, \lambda))^+ \leq \kappa$, so that $|{}^\lambda\kappa| \leq \kappa$. \square

We remark that (1) and (2) overlap slightly; if κ is either λ or λ^+ , then either applies to show that $\kappa^\lambda = \lambda^+$.

Analogously to Definition I.11.28:

Definition I.13.15 *The \beth_ξ are defined by recursion on ξ by:*

- $\beth_0 = \aleph_0 = \omega$.
- $\beth_{\xi+1} = 2^{\beth_\xi}$.
- $\beth_\eta = \sup\{\beth_\xi : \xi < \eta\}$ when η is a limit ordinal.

So, *CH* is equivalent to the statement that $\beth_1 = \aleph_1$ and *GCH* is equivalent to the statement that $\beth_\xi = \aleph_\xi$ for all ξ .

Definition I.13.16 *A cardinal κ is weakly inaccessible iff $\kappa > \omega$, κ is regular, and $\kappa > \lambda^+$ for all $\lambda < \kappa$. κ is strongly inaccessible iff $\kappa > \omega$, κ is regular, and $\kappa > 2^\lambda$ for all $\lambda < \kappa$.*

It is clear from the definitions that all strong inaccessibles are weak inaccessibles, and that the two notions are equivalent under *GCH*. One cannot prove in *ZFC* that weak inaccessibles exist (see [18, 20]). By modifying Exercise I.11.33:

Exercise I.13.17 *If κ is weakly inaccessible, then it is the κ^{th} element of $\{\alpha : \alpha = \aleph_\alpha\}$. If κ is strongly inaccessible, then it is the κ^{th} element of $\{\alpha : \alpha = \beth_\alpha\}$.*

However, just in *ZFC*, the method of Exercise I.11.33 yields:

Exercise I.13.18 *Prove that there is an α such that $\alpha = \beth_\alpha$.*

Following Erdős, we define:

Definition I.13.19 $[A]^\kappa = \{x \subseteq A : |x| = \kappa\}$ and $[A]^{<\kappa} = \{x \subseteq A : |x| < \kappa\}$.

Exercise I.13.20 *If λ is an infinite cardinal and $\kappa \leq \lambda$ is a cardinal, then $||[\lambda]^\kappa| = \lambda^\kappa$. If $0 < \kappa \leq \lambda$ then $||[\lambda]^{<\kappa}| = \sup\{\lambda^\theta : \theta = |\theta| < \kappa\}$. In particular, $||[\lambda]^{<\omega}| = \lambda$.*

Exercise I.13.21 *Let λ be an infinite cardinal. Prove that there is a non-abelian group of size λ and a field of size λ .*

Hint. For example, you could use the free group on λ generators. You need something like $||[\lambda]^{<\omega}| = \lambda$ to show that this group has size λ . This exercise is trivial if you quote the Löwenheim–Skolem Theorem (Theorem II.7.16), whose proof uses similar set theory, plus a lot of model theory. \square

I.14 The Axiom of Foundation

The Axiom of Foundation was stated in Section I.2 just in terms of \in and $=$. However, it is clearly equivalent to:

$$\forall x[x \neq \emptyset \rightarrow \exists y \in x(y \cap x = \emptyset)] .$$

This axiom is completely irrelevant to the development of mathematics within set theory. Thus, it is reasonable to ask: What does it say and why is it on the list?

A brief answer is that it avoids “pathological sets” such as cycles in the \in relation. For example, if $a \in a$, then $x = \{a\}$ is a counter-example to Foundation, since the only member of x is a , and $x \cap a = x \neq \emptyset$. More generally, the \in relation can have no cycles, since if $a_1 \in a_2 \in \cdots \in a_n \in a_1$, then $x = \{a_1, \dots, a_n\}$ contradicts Foundation.

The Axiom of Extensionality excluded non-sets, such as badgers and ducks, from the set-theoretic universe under consideration (see Figure I.1, page 18). Then, the Axiom of Foundation excludes sets such that $a \in a$ or $b \in c \in b$ from the set-theoretic universe under consideration (see Figure I.2, page 68). Neither Extensionality nor Foundation makes any philosophical comment as to whether badgers or ducks or such sets a, b, c really exist in the “real world”, whatever that is.

Now, *it turns out* that sets such as these a, b, c never arise in the construction of mathematical objects, so their existence has the same philosophical status as do ducks and badgers and trolls – they may exist or they may not exist, but their existence doesn’t affect the mathematics. The purpose of this section is to make the informal “it

turns out” into part of the theory. We define a subclass WF of V , the class of *well-founded sets*, and show that all mathematics takes place within WF . Then, Foundation is equivalent to the statement $V = WF$ (see Theorem I.14.9). Foundation is never used in the development of mathematics because mathematical objects, such as \mathbb{Q} and \mathbb{R} , are well-founded anyway.

Roughly, the well-founded sets are those sets which can be obtained from nothing, \emptyset , by iterating *collection* that is, taking a bunch of sets and putting them between brackets, $\{\dots\dots\}$. Thinking of this as a construction, we start with \emptyset at stage 0. At stage 1, we can put brackets around it to form $\{\emptyset\} = 1$. At stage 2, we already have constructed 0, 1, and we can form any of the four subsets of $\{0, 1\}$. Two of these, $\{0\} = 1$ and $\{\} = \emptyset = 0$ we already have. The two new ones are the *rank 2* sets, shown in Table I.2.

Table I.2: The First Sixteen Sets

rank	sets
0	$\emptyset = 0$
1	$\{\emptyset\} = 1$
2	$\{\{\emptyset\}\} = \{1\}$, $\{\emptyset, \{\emptyset\}\} = 2$
3	$\{\{1\}\}$, $\{0, \{1\}\}$, $\{1, \{1\}\}$, $\{0, 1, \{1\}\}$, $\{2\}$, $\{0, 2\}$, $\{1, 2\}$, $\{0, 1, 2\} = 3$, $\{\{1\}, 2\}$, $\{0, \{1\}, 2\}$, $\{1, \{1\}, 2\}$, $\{0, 1, \{1\}, 2\}$

At stage 3, we already have constructed 0, 1, 2, $\{1\}$, and we can form any of the $2^4 = 16$ subsets of $\{0, 1, 2, \{1\}\}$. Four of these we already have, and the 12 new ones are the rank 3 sets. The 16 sets displayed in Table I.2 make up $R(4)$, the collection of sets whose ranks are 0, 1, 2, or 3. This construction is generalized by:

Definition I.14.1 *By recursion on $\alpha \in ON$, define $R(\alpha)$ by:*

1. $R(0) = \emptyset$.
2. $R(\alpha + 1) = \mathcal{P}(R(\alpha))$.
3. $R(\gamma) = \bigcup_{\alpha < \gamma} R(\alpha)$ for limit ordinals γ .

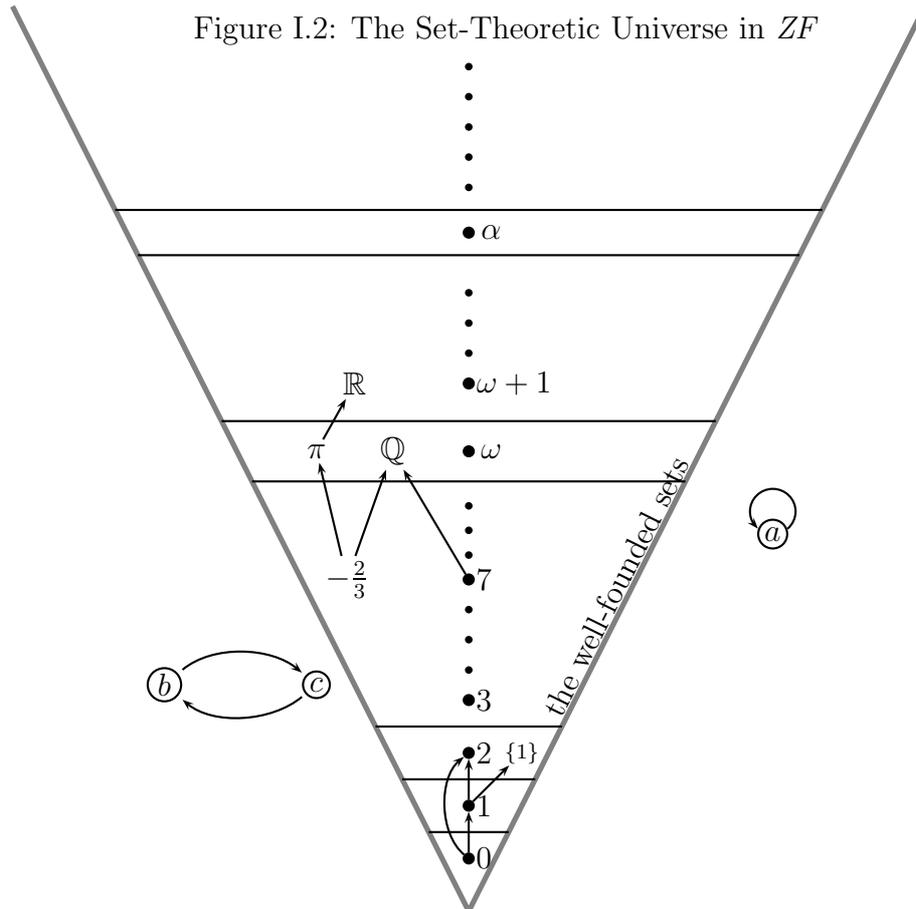
Then:

4. $WF = \bigcup_{\delta \in ON} R(\delta) =$ the class of all well-founded sets.
5. The set x is well-founded iff $\exists \delta[x \in R(\delta)]$.
6. For $x \in WF$: $\text{rank}(x)$ is the least α such that $x \in R(\alpha + 1)$.

WF is a proper class, since it contains all the ordinals (see Lemma I.14.5). Thus, WF does not really exist, but as with ON , the notation is useful. That is “ $x \in WF$ ” is shorthand for “ x is well-founded” and “ $x \subseteq WF$ ” is shorthand for “ $\forall y \in x \exists \delta[y \in R(\delta)]$ ”. Actually, item (6) of Definition I.14.1 needs some justification, given by:

Lemma I.14.2 *If $x \in WF$, the least δ such that $x \in R(\delta)$ is a successor ordinal.*

Proof. $\delta \neq 0$ by item (1) and δ cannot be a limit ordinal by item (3). □



Note that $|R(0)| = 0$ and $R(n + 1)$ has size $2^{|R(n)|}$, since $R(n + 1)$ is the power set of $R(n)$. Thus, $|R(1)| = 2^0 = 1$, $|R(2)| = 2^1 = 2$, $|R(3)| = 2^2 = 4$, $|R(4)| = 2^4 = 16$, $|R(5)| = 2^{16} = 65536$, etc. The four elements of $R(3)$ are the four elements of ranks 0, 1 or 2. Of the 16 elements of $R(4)$, four of them occur already in $R(3)$; the other 12 do not, so they have rank 3.

Exercise I.14.3 *List the 65520 sets of rank 4; this is easy to do with a computer program implementing Exercise I.14.12.*

Lemma I.14.4

1. Every $R(\beta)$ is a transitive set.
2. $\alpha \leq \beta \rightarrow R(\alpha) \subseteq R(\beta)$.

3. $R(\alpha + 1) \setminus R(\alpha) = \{x \in WF : \text{rank}(x) = \alpha\}$.
4. $R(\alpha) = \{x \in WF : \text{rank}(x) < \alpha\}$.
5. If $x \in y$ and y in WF , then $x \in WF$ and $\text{rank}(x) < \text{rank}(y)$.

Proof. For (1), induct on β . If $\beta = 0$, then $R(\beta) = 0$, which is transitive. At limits, use the fact that any union of transitive sets is transitive. Finally, assume $R(\beta)$ is transitive; we shall prove that $R(\beta + 1)$ is transitive. Note that $R(\beta) \subseteq \mathcal{P}(R(\beta)) = R(\beta + 1)$ since every element of $R(\beta)$ is a subset of $R(\beta)$. Thus, if $x \in R(\beta + 1) = \mathcal{P}(R(\beta))$, we have $x \subseteq R(\beta) \subseteq R(\beta + 1)$, so $R(\beta + 1)$ is transitive.

For (2), fix α and prove that $R(\alpha) \subseteq R(\beta)$ by induction on $\beta \geq \alpha$. It is trivial for $\beta = \alpha$, and the successor step follows from the fact that $R(\beta) \subseteq R(\beta + 1)$, which we just proved. If $\beta > \alpha$ is a limit, then $R(\beta) = \bigcup_{\xi < \beta} R(\xi) \supseteq R(\alpha)$.

(3) and (4) are immediate from (2) and the definition of rank. (5) is immediate from (3) and (4). \square

Some of the elements of WF are shown in Figure I.2. These have been placed at the proper level to indicate their rank, and a few arrows have been inserted to indicate the \in relation. Lemma I.14.4(5) says that WF is a transitive class and that the \in -arrows all slope upwards.

Table I.2 shows that the ordinals 0, 1, 2, 3 have ranks 0, 1, 2, 3, respectively. This generalizes:

Lemma I.14.5

1. $ON \cap R(\alpha) = \alpha$ for each $\alpha \in ON$.
2. $ON \subseteq WF$
3. $\text{rank}(\alpha) = \alpha$ for each $\alpha \in ON$.

Proof. (1) is proved by induction on α . The cases where α is 0 or a limit are easy. Now, assume $ON \cap R(\alpha) = \alpha$. We show that $ON \cap R(\alpha + 1) = \alpha + 1 = \alpha \cup \{\alpha\}$. To see that $\alpha \cup \{\alpha\} \subseteq R(\alpha + 1)$, note that $\alpha \subseteq R(\alpha) \subseteq R(\alpha + 1)$ and $\alpha \in \mathcal{P}(R(\alpha)) = R(\alpha + 1)$. Now, let δ be any ordinal in $R(\alpha + 1) = \mathcal{P}(R(\alpha))$. Then $\delta \subseteq R(\alpha) \cap ON = \alpha$, so $\delta \leq \alpha$. Thus, $ON \cap R(\alpha + 1) = \{\delta : \delta \leq \alpha\} = \alpha + 1$.

(1) implies that $\alpha \in R(\alpha + 1) \setminus R(\alpha)$, which yields (2) and (3). \square

To compute the rank of a set other than an ordinal, the following formula is useful:

Lemma I.14.6 For any set y : $y \in WF \leftrightarrow y \subseteq WF$, in which case:

$$\text{rank}(y) = \sup\{\text{rank}(x) + 1 : x \in y\}$$

Proof. If $y \in WF$ then $y \subseteq WF$ by (5) of Lemma I.14.4. If $y \subseteq WF$, then let $\beta = \sup\{\text{rank}(x) + 1 : x \in y\}$. Then $y \subseteq R(\beta)$, so $y \in R(\beta + 1)$, so $y \in WF$.

This also shows that $\text{rank}(y) \leq \beta$, and we need to show that $\text{rank}(y) = \beta$. Assume that $\alpha := \text{rank}(y) < \beta$. Then for all $x \in y$ we have $\text{rank}(x) < \alpha$ (by (5) of Lemma I.14.4), so that $\text{rank}(x) + 1 \leq \alpha$. Hence, $\beta \leq \alpha$, a contradiction. \square

Using this lemma, if you can display a set explicitly, then you can compute its rank. For example, 2 and 5 are ordinals, so $\text{rank}(2) = 2$ and $\text{rank}(5) = 5$, and then $\text{rank}(\{2, 5\}) = \max\{3, 6\} = 6$, and $\text{rank}(\{2\}) = 3$, so $\text{rank}(\langle 2, 5 \rangle) = \text{rank}(\{\{2\}, \{2, 5\}\}) = \max\{4, 7\} = 7$. Some more general facts about ranks are given by the next two lemmas. The first is immediate from Lemma I.14.6.

Lemma I.14.7 *If $z \subseteq y \in WF$ then $z \in WF$ and $\text{rank}(z) \leq \text{rank}(y)$.*

Lemma I.14.8 *Suppose that $x, y \in WF$. Then:*

1. $\{x, y\} \in WF$ and $\text{rank}(\{x, y\}) = \max(\text{rank}(x), \text{rank}(y)) + 1$.
2. $\langle x, y \rangle \in WF$ and $\text{rank}(\langle x, y \rangle) = \max(\text{rank}(x), \text{rank}(y)) + 2$.
3. $\mathcal{P}(x) \in WF$ and $\text{rank}(\mathcal{P}(x)) = \text{rank}(x) + 1$.
4. $\bigcup x \in WF$ and $\text{rank}(\bigcup x) \leq \text{rank}(x)$.

Proof. (1) and (2) are immediate from Lemma I.14.6. For (3), if $x \in WF$, then Lemma I.14.6 implies that all $y \subseteq x$ are in WF , with $\text{rank}(y) \leq \text{rank}(x)$. Applying Lemma I.14.6 again yields $\mathcal{P}(x) \in WF$ and $\text{rank}(\mathcal{P}(x)) = \sup\{\text{rank}(y) + 1 : y \subseteq x\} = \text{rank}(x) + 1$. (4) is proved by a similar use of Lemma I.14.6. \square

We now know that WF contains all the ordinals and is closed under standard set-theoretic operations, such as those given by Lemma I.14.8. It is perhaps not surprising then that WF contains everything, which is true assuming the Axiom of Foundation:

Theorem I.14.9 (ZF^-) *The Axiom of Foundation is equivalent to the assertion that $V = WF$.*

Proof. For \leftarrow , note that if $x \in WF$ and $x \neq \emptyset$, and $y \in x$ has least rank among the members of x , then $y \cap x = \emptyset$.

For \rightarrow , assume Foundation and fix x ; we prove that $x \in WF$. Let $t = \text{trcl}(x)$ (see Exercise I.9.6). Then t is transitive. If $t \subseteq WF$, then $x \subseteq t \subseteq WF$, so $x \in WF$ by Lemma I.14.6. Now, assume $t \not\subseteq WF$. Then $t \setminus WF \neq \emptyset$, so by Foundation, we can fix $y \in t \setminus WF$ with $y \cap (t \setminus WF) = \emptyset$. But $y \subseteq t$, since t is transitive, so that $y \subseteq WF$, so that $y \in WF$, a contradiction. \square

We note that the occurrence of $\text{trcl}(x)$ is natural here. If we view Foundation informally as saying that each x is obtained from nothing by a transfinite sequence of collections, then $\text{trcl}(x)$ traces all the sets constructed in the process of obtaining x . The cardinality of $\text{trcl}(x)$ counts how many sets are needed to obtain x . Of particular importance are the x for which this number is finite:

Lemma I.14.10 $x \in R(\omega)$ iff $x \in WF$ and $\text{trcl}(x)$ is finite.

Proof. For \rightarrow : If $x \in R(\omega)$, then $x \in R(n)$ for some finite n . Since $R(n)$ is transitive, we have $\text{trcl}(x) \subseteq R(n)$, and each $R(n)$ is finite by induction on n .

For \leftarrow : Apply transfinite induction (Theorem I.9.1). Let E be the class of all $\alpha \in ON$ such that $\alpha = \text{rank}(x)$ for some $x \in WF$ with $\text{trcl}(x)$ finite. Then the \leftarrow direction is equivalent to asserting that $E \subseteq \omega$. If this fails, let α be the least ordinal in $E \setminus \omega$ and fix $x \in WF$ with $\text{rank}(x) = \alpha$ and $\text{trcl}(x)$ finite. Now

$$\alpha = \text{rank}(x) = \sup\{\text{rank}(y) + 1 : y \in x\} .$$

But for each $y \in x$, $\text{trcl}(y)$ is finite (by Exercise I.9.6), and $\text{rank}(y) < \text{rank}(x)$, so that $\text{rank}(y) \in \omega$ (since α is least). Thus, α is a sup of a finite set of finite ordinals, so $\alpha < \omega$, a contradiction. \square

In the \leftarrow , we cannot drop the assumption that $x \in WF$ because it is consistent with ZFC^- to have a set x such that $x = \{x\}$; then $x = \text{trcl}(x)$, which is finite, but $x \notin WF$ (since $x \in x$ implies $\text{rank}(x) < \text{rank}(x)$). As we pointed out in the beginning of this section, the Axiom of Foundation implies that no x satisfies $x \in x$.

Definition I.14.11 $HF = R(\omega)$ is called the set of hereditarily finite sets.

If you think of \in as a directed graph and draw a picture of x together with $\text{trcl}(x)$, then x is the set of children of x , while $\text{trcl}(x)$ is the set of all descendents of x ; hence the name “hereditarily finite”.

Equivalently, HF is the set of all well-founded sets of finite rank. Informally, all of finite mathematics lives in HF . It is clear from Lemma I.14.6 that any set which can be displayed explicitly as a finite expression using \emptyset and brackets (i.e., $\{, \}$) is in HF . This includes the finite ordinals, plus, e.g., ${}^m n$ whenever $m, n \in \omega$, plus anything else in finite combinatorics. For example, every finite group is isomorphic to one of the form $\langle n, \cdot \rangle$ where $n \in \omega$ and $\cdot \subseteq (n \times n) \times n$ (since the product, \cdot is a function from $n \times n \rightarrow n$). The rank of this $\langle n, \cdot \rangle$ is finite (in fact, $n + 4$).

Similarly, *all* mathematics lives in WF . For example, all ordinals are in WF , and it is easily seen that $\alpha\beta \in WF$ whenever α, β are ordinals. Likewise, (using the Axiom of Choice), every group is isomorphic to one of the form $\langle \kappa, \cdot \rangle$ for some cardinal κ , and this $\langle \kappa, \cdot \rangle$ is in WF . By the same reasoning, within WF we have isomorphic copies of every other kind of abstract algebraic structure. Likewise, concrete objects, such as \mathbb{R} and \mathbb{C} , are in WF and we can compute their rank (see Section I.15). This also explains why the Axiom of Foundation ($V = WF$) is irrelevant for the development of mathematics in set theory; all mathematics lives in WF anyway, so the question of whether or not $V \setminus WF$ is empty has no bearing on the mathematics.

The development of WF is due to von Neumann. He also made these informal remarks into a theorem expressing the consistency of the Axiom of Foundation by using

WF as a model. For example, he showed that if ZFC^- is consistent, then so is ZFC , since any proof of a contradiction $\varphi \wedge \neg\varphi$ from ZFC can be translated into a contradiction from ZFC^- by showing that φ is both true and false within the model WF . For details, see a set theory text [18, 20].

Exercise I.14.12 Define $E \subseteq \omega \times \omega$ by: $nEm \leftrightarrow 2 \nmid \lfloor m2^{-n} \rfloor$ (equivalently, there is a 1 in place n in the binary representation of m). Prove that $(\omega; E) \cong (HF; \in)$. Remark. This gives a very explicit enumeration of HF in type ω . The first 16 sets are listed in order in Table I.2.

Exercise I.14.13 Let K be any class such that for all sets y , if $y \subseteq K$ then $y \in K$. Then $WF \subseteq K$. Remark. WF has this property by Lemma I.14.6.

Exercise I.14.14 Prove that $HF = \{x : |\text{trcl}(x)| < \aleph_0\}$.

Hint. Prove that $\text{trcl}(x)$ is not finite whenever $x \neq R(\omega)$. You need Foundation here. In ZFC^- , one cannot even prove that HF is a set, since $\{x : x = \{x\}\}$ might form a proper class. \square

Exercise I.14.15 For any infinite cardinal κ , define $H(\kappa) = \{x : |\text{trcl}(x)| < \kappa\}$. Prove:

1. $H(\kappa) \subseteq R(\kappa)$, so that $H(\kappa)$ is a set.
2. $|H(\kappa)| = 2^{<\kappa} := \sup\{2^\theta : \theta < \kappa\}$.

Hint. For (2): Show that whenever $x \neq y$, the \in relations on $\{x\} \cup \text{trcl}(x)$ and on $\{y\} \cup \text{trcl}(y)$ cannot be isomorphic. Then, when $|\text{trcl}(x)| = \theta$, there are at most 2^θ possibilities for the isomorphism type of \in on $\{x\} \cup \text{trcl}(x)$. Note that $|H(\kappa)| \geq 2^{<\kappa}$ is easy because $\mathcal{P}(\theta) \subseteq H(\kappa)$ whenever $\theta < \kappa$. \square

Exercise I.14.16 Prove that $R(\omega + \omega)$ is a model for ZC (that is, all the ZFC axioms except the Replacement Axiom), and that some instance of the Replacement Axiom is false in $R(\omega + \omega)$.

Exercise I.14.17 Prove that for regular $\kappa > \omega$, $H(\kappa)$ is a model for all the ZFC axioms except the possibly the Power Set Axiom, which holds in $H(\kappa)$ iff κ is strongly inaccessible. Also, $HF = H(\omega) = R(\omega)$ is a model for all the ZFC axioms except the Axiom of Infinity.

These last two exercises form the beginning of the discussion of models of set theory; this discussion eventually leads to models of ZFC in which statements such as CH are either true or false; see [18, 20].

I.15 Real Numbers and Symbolic Entities

Once one has the set ω of natural numbers, one can construct the sets of rationals (\mathbb{Q}), reals (\mathbb{R}), and complex numbers (\mathbb{C}) in ZF^- by standard set-theoretic arguments. The details are a bit tedious (see [21]), since along with these sets, one must also define their addition, multiplication, and order, and verify the usual algebraic laws. We only outline the development here. We also explain how to regard symbolic entities, such as polynomials over a field, within the framework of axiomatic set theory. All the objects we construct will lie within the well-founded sets, and we shall see what their ranks are (see Section I.14). In particular, all “essentially finite” objects will have finite rank – that is, will be members of HF . These “essentially finite” objects include the kind of objects one can (and does) enter into a computer program: natural numbers (we already have $\omega \subseteq HF$), as well as rational numbers, polynomials over the rationals, and finite boolean expressions. We begin with the rationals.

Informally, one can get \mathbb{Q} by adding to ω objects such as $2/3$ or $-2/3$ or -7 . One need not answer philosophical questions, such as what is the “true essence” of $-2/3$; one may view it purely as a “symbolic entity”, which may be represented in set theory as follows:

Definition I.15.1 \mathbb{Q} is the union of ω with the set of all $\langle i, \langle m, n \rangle \rangle \in \omega \times (\omega \times \omega)$ such that:

1. $m, n \geq 1$
2. $i \in \{0, 1\}$
3. $\gcd(m, n) = 1$
4. If $i = 0$ then $n \geq 2$

With this formalism, $2/3, -2/3, -7$ are, respectively, $\langle 0, \langle 2, 3 \rangle \rangle, \langle 1, \langle 2, 3 \rangle \rangle, \langle 1, \langle 7, 1 \rangle \rangle$. So, the i in $\langle i, \langle m, n \rangle \rangle$ is a sign bit, with 0 signifying $+$ and 1 signifying $-$. The point of (3) is to avoid multiple representations of the same number. The point of (4) is to avoid entities such as $\langle 0, \langle 7, 1 \rangle \rangle$, which would represent 7, which is already in ω .

Exercise I.15.2 $\mathbb{Q} \subseteq HF$ and $\text{rank}(\mathbb{Q}) = \omega + 1$.

Definition I.15.3 $+, \cdot,$ and $<$ are defined on \mathbb{Q} in the “obvious way”, to make \mathbb{Q} into an ordered field containing ω .

Obviously, this definition must be expanded to fill in the details.

Algebraically, it would be more elegant to begin by defining \mathbb{Z} , the ring of positive and negative integers. Then \mathbb{Z} is an integral domain and \mathbb{Q} is its quotient field. With this approach, each element of \mathbb{Q} is an equivalence class of pairs of integers. For example, $2/3$ would be the countably infinite set $\{\langle 2, 3 \rangle, \langle 4, 6 \rangle, \langle 6, 9 \rangle \cdots\}$. This definition is preferred in algebra because it is a general construction and does not rely on a special trick for

picking representatives of classes, which works only in the case of the integers. On the other hand, with this definition, \mathbb{Q} would not be contained in HF . Our Definition I.15.1 approximates the finite symbolic expression you would use to enter a rational number into a computer program.

Given \mathbb{Q} and its order, we define the real and complex numbers by:

Definition I.15.4 \mathbb{R} is the set of all $x \in \mathcal{P}(\mathbb{Q})$ such that $x \neq \emptyset$, $x \neq \mathbb{Q}$, x has no largest element, and

$$\forall p, q \in \mathbb{Q}[p < q \in x \rightarrow p \in x] \quad . \quad (*)$$

$$\mathbb{C} = \mathbb{R} \times \mathbb{R}.$$

Informally, if $x \in \mathbb{R}$, then one can define its *lower Dedekind cut*, $C_x = \{q \in \mathbb{Q} : q < x\}$; then C_x is a subset of \mathbb{Q} satisfying (*). Formally, we identify x with C_x and simply *define* \mathbb{R} to be the collection of all sets satisfying (*). Of course, we have to define appropriate $+$, \cdot , and $<$ on \mathbb{R} . Then, a complex number is a pair of reals $\langle x, y \rangle$ (representing $x + iy$).

Exercise I.15.5 $\text{rank}(x) = \omega$ for each $x \in \mathbb{R}$, and $\text{rank}(\mathbb{R}) = \omega + 1$. $\text{rank}(\mathbb{C}) = \omega + 3$.

Of course, real numbers and complex numbers are not “essentially finite” objects, and there is no way to get $\mathbb{R} \subseteq HF$ because HF is countable and \mathbb{R} is not. The sets one encounters in elementary mathematics, such as \mathbb{R} , \mathbb{C} , and Lebesgue measure on \mathbb{R} are all in $R(\omega + \omega)$, a natural model for ZC (see Exercise I.14.16 and [18, 20]).

There are many different ways to construct the real numbers. It is important to note that these various ways all lead to the same thing, up to isomorphism, so that mathematicians can talk about *the* real numbers, without referring to the specific way they were constructed. To formalize this,

Definition I.15.6 An ordered field $(F; +, \cdot, <)$ satisfies the usual rules of manipulation you used in high school algebra for \mathbb{R} and \mathbb{Q} . It is *Dedekind-complete* iff it satisfies the *least upper bound axiom* — that is, whenever $X \subseteq F$ is non-empty and bounded above, the least upper bound, $\sup X$, exists.

Proposition I.15.7 All Dedekind-complete ordered fields are isomorphic.

Because of this proposition, elementary analysis texts often assume as an axiom that \mathbb{R} forms a Dedekind-complete ordered field, and do not construct \mathbb{R} from more basic set-theoretic objects. We shall not prove this proposition here, so we shall not list all the axioms for ordered fields; besides the field axioms (see Example II.8.23), there are a few axioms involving the order, such as the statement that the sum and product of positive elements are positive.

The sets arising in elementary analysis are either countable or of sizes 2^{\aleph_0} or $2^{2^{\aleph_0}}$. These cardinalities are easily computed using basic cardinal arithmetic; for example:

Exercise I.15.8 $|\mathbb{R}| = C(\mathbb{R}, \mathbb{R}) = 2^{\aleph_0}$ and $\mathbb{R}^{\mathbb{R}} = 2^{2^{\aleph_0}}$, where $C(\mathbb{R}, \mathbb{R})$ is the set of all continuous functions in $\mathbb{R}^{\mathbb{R}}$. There are 2^{\aleph_0} Borel subsets of \mathbb{R} and $2^{2^{\aleph_0}}$ Lebesgue measurable subsets of \mathbb{R} .

Hint. For the collection \mathcal{B} of Borel sets: If you define \mathcal{B} as the least σ -algebra containing the open sets, as is often done in analysis texts, then you have no idea what $|\mathcal{B}|$ is. Instead, prove that $\mathcal{B} = \mathcal{B}_{\omega_1}$ can be obtained by a transfinite process, where \mathcal{B}_0 is the family of all sets which are either open or closed, $\mathcal{B}_{\alpha+1}$ is the family of all countable unions and intersections of sets in \mathcal{B}_α , and $\mathcal{B}_\gamma = \bigcup\{\mathcal{B}_\alpha : \alpha < \gamma\}$ for limit γ . \square

A deeper fact is that every Borel subset of \mathbb{R} is either countable or of size 2^{\aleph_0} . This was proved independently around 1915 by Hausdorff and Aleksandrov; hence, as they both clearly realized, a counter-example to the Continuum Hypotheses could not come from elementary analysis.

Finally, we consider how to handle symbolic expressions within axiomatic set theory. These expressions occur frequently in mathematics, especially in algebra and logic.

In algebra, consider polynomials. It is important to regard polynomials as symbolic expressions, not functions. For example, if K is a field, the polynomials $x^2 + y^4$ and $x^4 + y^2$ always denote distinct polynomials over K , whereas the corresponding functions might be the same (when K has 2 or 3 elements).

Another algebra example: in group theory, when we describe the free group F on 2 generators, we build it from all the set of all words in two letters, say, x, y . For example, $xxxy^{-1}x^{-1}x^{-1}$ is such a word. It (or its equivalence class, depending on the exposition) is a member of F .

For both these examples, what are the symbols x and y ? In our development of *ZFC*, we have not yet encountered any such entity called a “symbol”. These two examples are handled like the following example from logic, which we shall discuss in more detail since it is closer to the theme of this book.

Say we want to discuss propositional logic and truth tables. So, our objects are boolean expressions such as $\neg[p \wedge q]$. Call this expression σ . It is a sequence of six symbols. We have already discussed sequences (see Definition I.9.3); σ can be considered to be a function with domain $6 = \{0, 1, 2, 3, 4, 5\}$, where $\sigma(0)$ is the symbol ‘ \neg ’, $\sigma(1)$ is the symbol ‘ $]$ ’, $\sigma(2)$ is the symbol ‘ p ’, etc. But, *what is a symbol?* Humans use their visual processing ability to recognize the symbol ‘ \neg ’ by its shape. However, the mathematics of boolean expressions should not depend on psychology or on vision or on what the object ‘ \neg ’ really is, as long as this object is well defined. Working in *ZF*, we only have sets, so we must define ‘ \neg ’ to be some specific set. To be definite, we shall take it to be some natural number. If we choose to represent all our symbols as natural numbers, then all our syntactical objects will lie within *HF*, in line with our expectation that finite mathematics is done within *HF*.

Definition I.15.9 P_n is the number $2n + 2$. Let the symbols $]$, $[$, \neg , \vee , \wedge be shorthand for the numbers 1, 3, 5, 7, 9, respectively.

In this particular application, we are thinking of P_0, P_1, \dots as proposition letters or boolean variables. Let $A = \{1, 3, 5, 7, 9\} \cup \{2n + 2 : n \in \omega\} \subseteq \omega$. Then A is our *alphabet*, which consists of the proposition letters plus the five other “symbols”, $]$, $[$, \neg , \vee , \wedge .

In the theory of formal languages, we need to discuss strings of symbols and concatenation of strings. These were defined in Definitions I.10.3 and I.10.4.

Exercise I.15.10 $HF^{<\omega} \subseteq HF$.

Of course, it is not essential that we use positive integers for symbols. It is useful to assume that $A \cap A^{<\omega} = \emptyset$, so that no symbol is also a string of symbols. This holds for the *positive* integers; we avoid $0 = \emptyset$, which is also the empty sequence:

Exercise I.15.11 If $A \subseteq \omega \setminus \{0\}$, then $A \cap A^{<\omega} = \emptyset$.

For example, $\varphi := \neg[P_2 \wedge P_3]$ is really the sequence $\varphi = (5, 3, 6, 9, 8, 1)$. Let ψ be $[P_1 \wedge P_3]$, which is really $(3, 4, 9, 8, 1)$. Then $[\psi \wedge \phi]$ denotes the concatenation of the “[” symbol, the symbols in ψ , the “^” symbol, the symbols in ψ , and the “]” symbol, which is the string $(3, 3, 4, 9, 8, 1, 9, 5, 3, 6, 9, 8, 1, 1)$, or $[[P_1 \wedge P_3] \wedge \neg[P_2 \wedge P_3]]$.

Note that our notation “[$\psi \wedge \phi$]” for a concatenation is an example of the following, which is more convenient than the “raw” terminology of Definition I.10.4:

Definition I.15.12 Assume that $A \cap A^{<\omega} = \emptyset$, and fix $\tau_0, \dots, \tau_{m-1} \in A \cup A^{<\omega}$. Let σ_i be τ_i if $\tau_i \in A^{<\omega}$, and the sequence of length 1, (τ_i) , if $\tau_i \in A$. Then $\tau_0, \dots, \tau_{m-1}$ denotes the string $\sigma_0 \widehat{\ } \dots \widehat{\ } \sigma_{m-1} \in A^{<\omega}$.

This is just the beginning of formal logic. We must now define precisely which elements of $A^{<\omega}$ are well-formed boolean expressions; these well-formed expressions form a *formal language*; that is, a language defined artificially with mathematically precise rules of syntax. We must also define notions such as “truth table” and “tautology”; for example, the $\psi \wedge \varphi$ above is logically equivalent to $[[P_1 \wedge P_3] \wedge \neg P_2]$ because the two expressions have the same truth table. We also need to extend propositional logic to predicate logic, which has variables and quantifiers along with boolean connectives. The axioms of *ZFC* are written in predicate logic. For details, see Chapter II.

Note that we have been careful to use $]$ and $[$ instead of the more common $)$ and $($ when discussing the formal language, since we wanted to avoid confusion with the $)$ and $($ used in writing sequences of numbers. In Chapter II, as well as in our statement of the *ZFC* axioms in Section I.2, we use $)$ and $($ instead, which is much more common in the literature.

In any discussion of formal languages, one must distinguish between the symbols of the language and meta-variables. For example, we may say:

$$[p \vee \neg p] \tag{1}$$

is a tautology. In (1), the four symbols $[\vee, \neg, \wedge, \rightarrow]$ denote specific symbols of our formal language, 3, 1, 7, 5. However, p is meta-variable; it does not denote any specific proposition letter; and the assertion that (1) is a tautology is just a convenient shorthand for saying that for every $n \in \omega$, $[P_n \vee \neg P_n]$ is a tautology. Besides being shorter, this terminology yields more easily recognizable statements. The reader may safely forget that our “official” proposition letters are the $P_n = 2n + 2$, since this definition is found only in this section of this book; other books will use other definitions. However, everyone will recognize (1).

The same concept is familiar from discussions of computer languages. For example we may say that

$$id_1 = \text{sqrt} (id_2); \quad (2)$$

is an assignment statement in C or C++ or Java whenever id_1 and id_2 are legal identifiers. The “= sqrt (” and “);” are part of the formal language – that is, you literally type these into your program, whereas the id_1 and id_2 are meta-variables.

In model theory, one frequently needs to study languages which use uncountably many symbols; for example, we may wish to have available uncountably many boolean variables. In that case, we could, for example, let P_α be the ordinal $2 \cdot \alpha + 2$; of course, then our boolean expressions will no longer all lie in HF .

As we shall see, some areas of model theory are very set-theoretic and make frequent use of the material in this chapter. However, there are no further foundational problems. Now that we have explained how to treat symbolic entities within ZF , the rest of the subject is developed using the type of “ordinary mathematical reasoning” which is obviously formalizable within ZFC .

Chapter II

Model Theory and Proof Theory

II.1 Plan

At the elementary level of this book, model theory and proof theory are very closely related, and we shall treat these subjects together. In more advanced work, which we shall only mention briefly, the subjects diverge.

II.2 Historical Introduction to Proof Theory

As soon as people became aware that one might base mathematics and other sciences on logical deduction, it became natural to ask whether one could say precisely exactly what constitutes a correct deduction. As far as we know, this was first investigated in detail by Aristotle (384 BC – 322 BC), who described various forms of syllogism, such as:

If
 every Greek is a person and
 every person is mortal
then
 every Greek is mortal.

which we would phrase in modern logic as:

$$\forall x(G(x) \rightarrow P(x)) , \forall x(P(x) \rightarrow M(x)) \vdash \forall x(G(x) \rightarrow M(x)) .$$

As in Section 0.4, this turnstile symbol “ \vdash ” is read “proves”. Aristotle was aware of the axiomatic foundations of geometry, which was studied by Eudoxus (408 BC – 355 BC) and others; these foundations were, somewhat after Aristotle, expounded in detail in the famous *Elements* [12] of Euclid (~ 325 BC – ~ 265 BC). Scholastic philosophers and theologians in the Middle Ages carried on in the spirit of Aristotle. We realize today that syllogistic reasoning captures only a small part of logical deduction. Further

progress was made by Boole, around 1850, who studied boolean, or propositional logic. Frege's *Begriffsschrift* (1879) described a version of full predicate logic, although the syntax differed quite a bit from that of modern predicate logic.

The subject took on more urgency after Cantor. Previously, one might take the position that questions about axiomatics and deductions were of interest only in philosophy, since one could use physical intuition as a guide when rigor was lacking. After all, calculus was developed and applied in physics for about 200 years after its discovery in the 1600s, before it was finally put on a rigorous basis in the 1800s. However, Cantor's set theory led to branches of mathematics which were far removed from physical reality, so that it became more important to say precisely what is acceptable as correct mathematics.

At the present time, it is conventional to say that mathematics consists of anything provable from *ZFC*. The *ZFC* axioms were given in Chapter I, and were written there (in Section I.2) using the symbolism of formal logic, but the proofs were presented informally, in ordinary English. As we explained in Section 0.4, to complete the discussion of the foundations of mathematics, we need to give a formal definition of what a proof is, which we shall do in this chapter. Then, we shall define (see Section II.10) the notion $\Sigma \vdash \varphi$ to mean that there is a proof of φ from Σ .

Of course, it is not enough just to write down some definition of \vdash . We must also prove some theorems about this notion, saying that it has the properties expected of a notion of provability. The main result in this direction is the Completeness Theorem, which was mentioned in Section 0.4 and which will be proved in Section II.12.

In fact, there are many different approaches to proof theory, and the definition of \vdash that you use will depend on your reason for studying it. We list three possible goals that you might have in mind. The first two are mentioned only peripherally in this book; some further remarks are in Section II.17.

Goal 1. Easy transcription of informal mathematics into the formal system. In a sense, this motivation goes back to Aristotle, who wanted to show that reasoning could be transcribed into syllogisms. At present, this is of interest because the formal proofs can then be input into a computer. In particular, the systems Isabelle/ZF [17] and Mizar [25] work in a form of axiomatic set theory and allow the user to enter theorems and proofs into the computer, which then verifies that the proof is correct. Both systems have verified a significant body of abstract mathematics. In addition, systems ACL2 [1] and Coq [8] are optimized for verifying theorems about finite objects.

Goal 2. Having the computer *discover* formal proofs. This is different from Goal 1, since the system must be optimized for efficient search. McCune's systems OTTER [26] and Prover9 [27] are universally recognized as the state-of-the-art here.

Goal 3. Easy development of the *theory* of the formal system. This theory will include the basic definitions and the proof of the Completeness Theorem (Theorem II.12.1). This goal is fulfilled very well by using a Hilbert-style (see [15]) proof theory (see Section II.10). These Hilbert-style systems have an extremely simple definition of \vdash , which makes it easier to prove theorems about the notion. It is also of some philosophical interest

that in principle, all uses of language and reasoning in mathematics can be reduced to a few very simple elements. The extreme simplicity of our \vdash will make it very difficult to display actual formal proofs of any statements of mathematical interest, which makes our \vdash useless for Goals 1 and 2.

II.3 NON-Historical Introduction to Model Theory

This section outlines the modern view of model theory, not the way it arose historically.

If you continue in the spirit of Chapter I, you will work in *ZFC* and go on to develop algebra, analysis, and other areas of mathematics. In particular, you will find it useful to define various classes of structures, such as:

1. groups
2. rings
3. fields
4. ordered fields
5. totally ordered sets
6. well-ordered sets
7. Dedekind-complete ordered fields
8. cyclic groups

All these classes are closed under isomorphism; for example every relation isomorphic (in the sense of Definition I.7.14) to a well-order is a well-order and every group isomorphic to a cyclic group is cyclic (see Definition II.8.18 for the general definition of isomorphism).

You will then notice that some of these classes are *first-order* classes. This means that they are defined by quantification only over the elements of a structure. Thus, a *group* $(G; \cdot)$ must satisfy the axioms γ_1, γ_2 described in Section 0.4, which talk only about elements of G . For example, γ_1 was $\forall xyz[x \cdot (y \cdot z) = (x \cdot y) \cdot z]$, and the meaning of the $\forall xyz$ was for all x, y, z in G . Likewise, a total order $(A; <)$ must satisfy the properties of Definition I.7.2, all of which refer only to elements of A ; for example, the transitivity of $<$ is expressed by $\forall xyz[x < y \wedge y < z \rightarrow x < z]$, meaning, for all x, y, z in A . We shall presently make the notion of “first order” more precise, but, roughly, a first-order class will be the class of models of a set Σ of axioms in ordinary first-order logic. This first-order logic was discussed informally in Section 0.2, and was used to state the *ZFC* axioms in Section I.2. We shall then see that classes (1)(2)(3)(4)(5) are all first-order classes.

However, classes (6)(7)(8) are not first-order. To define the notion of well-order, you have to say that all non-empty *subsets* have a least element. Likewise, to define Dedekind-complete (see Definition I.15.6), you have to say that all non-empty *subsets* which are bounded above have a least upper bound. So, the definition talks about subsets of the structure, not just elements of the structure. For (8), the natural definition of

“cyclic” says that the group is generated by one element, which you could write as $\exists x \in G \forall y \in G \exists n \in \mathbb{Z} [y = x^n]$, which is not first-order because it refers to the set \mathbb{Z} , which is external to the group G .

Of course, we have only explained why the *standard definitions* of classes (6)(7)(8) are not first-order. Perhaps one might find some other equivalent definitions which are first-order. However, we shall see shortly that this is not the case.

This distinction between first-order and non-first-order is important, and not just a curiosity, because many of the basic theorems in model theory apply only to first-order classes. These theorems give you some powerful tools, but you can only use them if the class is first-order, so the tools apply to many, but not all, of the classes of structures occurring naturally in mathematics.

One of the basic results of model theory is the *Löwenheim–Skolem Theorem* (Theorem II.7.16), which states that if set Σ has an infinite model, then Σ has models of all infinite cardinalities. By this theorem, class (8) cannot be first-order, because every infinite cyclic group is countable (and is isomorphic to \mathbb{Z}). Likewise, class (7) cannot be first-order, because every Dedekind-complete ordered field has size 2^{\aleph_0} (and is isomorphic to \mathbb{R} ; see Proposition I.15.7).

Class (6) is also not first-order, although there are well-ordered sets of all cardinalities. To prove that (6) is not first-order, suppose that we had a set Σ of sentences in ordinary first-order logic such that any $(A; <)$ is a model for Σ iff $<$ is a well-order of A . Let Σ^+ be Σ with the additional first-order axiom:

$$\forall x (\neg \exists y (y < x) \vee \exists z (z < x \wedge \neg \exists y (z < y < x))) \quad .$$

This says that every element $x \in A$ is either the first element of the order or has an immediate predecessor. This is true of well-orders of type ω or less, but not of any well-order of type greater than ω . Thus, Σ^+ has a countably infinite model but no uncountable models, contradicting the Löwenheim–Skolem Theorem.

Elementary model theory, as is described here and in basic texts [5, 24], studies the mathematical structure of models for first-order theories. However, the structural properties themselves are usually not first-order. The most basic structural property of a model is its cardinality, and this is addressed by the Löwenheim–Skolem Theorem, which implies that cardinality is not a first-order property. Other theorems give a more refined account of the structural properties, but before we prove any theorems, we should pause and do things a bit more rigorously, starting with properties of the formal language.

II.4 Polish Notation

Recall our discussion of goals in Section II.2. The syntax used in ordinary mathematics is quite complex. If your goal is to be able to type theorems and proofs into a computer program, your formal language must be able to handle at least some of this complexity, or your system will not be useful in practice. Thus, the syntax of your formal logic will

be fairly complex, and will start to resemble the syntax of a programming language, such as C or Java or Python. It will then be a rather lengthy process to write down a formal definition of the syntax, and even more lengthy to prove non-trivial theorems about the logic. Furthermore, at the present time, no formal language captures all of informal mathematical usage, so no matter what formal language you use, you will have to learn by experience how to translate ordinary mathematical terminology into the formal system.

In this book, we take the opposite extreme. We start with standard mathematical usage and simplify it as much as possible without losing expressive power. This will make our formal notation look a bit ugly to most mathematicians, but it will also enable us to give fairly short proofs of basic results *about* the formal logic.

As an example, consider the various ways we write functions of one or two variables. The function symbol can come before the variables (*prefix notation*, e.g. $f(x)$ or $f(x, y)$), or afterwards (*postfix notation*, e.g. $x!$), or between them (*infix notation*, e.g., $x + y$ or $x \cdot y$). It can also be missing and inferred from context; e.g., xy (written horizontally) means $x \cdot y$, whereas x^y (written diagonally) denotes the exponential function. To understand mathematical writing, one must also know the standard conventions on precedence; e.g., $x + yz$ means $x + (y \cdot z)$, not $(x + y) \cdot z$.

In *Polish Notation* (developed by Jan Łukasiewicz in the 1920s), we write everything uniformly in prefix. For example, we write $+xy$ for $x + y$, $+x \cdot yz$ for $x + (y \cdot z)$, and $\cdot +xyz$ for $(x + y) \cdot z$. Note that the meaning of expressions is unambiguous, without requiring either parentheses or any knowledge of precedence (see Lemma II.4.3 below). However, we do need to know the *arity* of each symbol. For example, $+$ and \cdot are *binary*, or have arity 2, meaning they apply to two expressions. The factorial symbol, $!$, is *unary*, or has arity 1. The symbols x, y, z have arity 0. In general, once we have a set of symbols with designated arities, we may define the Polish expressions of these symbols, and prove that our grammar is unambiguous. Polish notation is defined by:

Definition II.4.1 *A lexicon for Polish notation is a pair (\mathcal{W}, α) where \mathcal{W} is a set of symbols and $\alpha : \mathcal{W} \rightarrow \omega$. Let $\mathcal{W}_n = \{s \in \mathcal{W} : \alpha(s) = n\}$. We say that the symbols in \mathcal{W}_n have arity n . As in Definition I.10.3, $\mathcal{W}^{<\omega}$ denotes the set of all finite sequences of symbols in \mathcal{W} . The (well-formed) expressions of (\mathcal{W}, α) are all sequences constructed by the following rule:*

(\star) *If $s \in \mathcal{W}_n$ and τ_i is an expression for each $i < n$, then $s\tau_0 \cdots \tau_{n-1}$ is an expression.*

In the “standard” applications, most of the \mathcal{W}_n are empty. For example, we can let $\mathcal{W} = \{x, y, z, !, +, \cdot\}$, with $\mathcal{W}_0 = \{x, y, z\}$, $\mathcal{W}_1 = \{!\}$, $\mathcal{W}_2 = \{+, \cdot\}$. Then the following shows 9 expressions of this lexicon:

$$x \quad y \quad z \quad +xy \quad \cdot yz \quad +xx \quad +x \cdot yz \quad \cdot +xyz \quad ! \cdot +xyz \quad (\heartsuit)$$

For the first 3, note that when $n = 0$, the rule (\star) says that every element of \mathcal{W}_0 forms an expression. Also note that the empty sequence is not an expression by our definition.

Notation II.4.2 If $\tau \in \mathcal{W}^{<\omega}$, then $|\tau|$ denotes the length of τ . If $j \leq |\tau|$ then $\tau \upharpoonright j$ is the sequence consisting of the first j elements of τ .

For example, if τ is $!\cdot+xyz$, then $|\tau| = 6$ and $\tau \upharpoonright 4$ is $!\cdot+x$. This is really just standard set-theoretic notation, since τ as an element of \mathcal{W}^6 is a function and a set of 6 ordered pairs, and $\tau \upharpoonright \{0, 1, 2, 3\}$ is the restriction of the function τ to 4.

We remark on the formal meaning of Definition II.4.1 within *ZFC*, following the comments in Section I.15. A symbol is really just a set, since everything is a set. As in Section I.15, we should assume that $\mathcal{W} \cap \mathcal{W}^{<\omega} = \emptyset$, and (\star) uses the terminology of concatenation from Definition I.15.12. We should also distinguish between the symbol $x \in \mathcal{W}_0$ and the expression of length 1 consisting of x appearing as the first item in (\mathfrak{B}) , which is set-theoretically $(x) = \{\langle 0, x \rangle\} \in \mathcal{W}^1 = \mathcal{W}^{\{0\}}$. Note that the rule (\star) is essentially a recursive definition of a formal language, but such recursions were not actually discussed in Section I.9. To formalize this definition, define a *construction sequence* to be a finite sequence $(\sigma_0, \dots, \sigma_k)$ such that each σ_ℓ is of the form $s\tau_0 \cdots \tau_{n-1}$, where $n \geq 0$, $s \in \mathcal{W}_n$, and $\{\tau_i : i < n\} \subseteq \{\sigma_m : m < \ell\}$; then, τ is an *expression* iff τ occurs in some construction sequence. For example, (\mathfrak{B}) displays a construction sequence of length 9.

If σ is an expression and s is the first symbol of σ , then σ must be of the form $s\tau_0 \cdots \tau_{n-1}$, where n is the arity of s , since σ must be formed using Rule (\star) . It is important to know that σ is *uniquely* of this form. To continue our example with $\mathcal{W} = \{x, y, z, !, +, \cdot\}$, suppose that σ is $\cdot+x!y+!zy$, which is the Polish way of writing $(x + y!) \cdot (z! + y)$. Then the first symbol is the \cdot , and τ_1 is $+x!y$ and τ_2 is $+!zy$. Of course, one can write σ in different ways in the form $\cdot\tau'_1\tau'_2$; for example, τ'_1 can be $+x!$ and τ'_2 can be $y+!zy$; but then τ'_1 and τ'_2 will not both be expressions. This *unique readability* (Lemma II.4.3) is important, because it implies that σ has a unique *meaning* (or semantics). In this algebraic example, the meaning of σ is the numeric value we compute for it if we assign numbers to x, y, z ; we write σ *uniquely* as $\cdot\tau_1\tau_2$, compute (recursively) values for τ_1, τ_2 , and then multiply them. In model theory, we shall write logical formulas in Polish notation, and the semantics will consist of a truth value of the formula in a given structure.

For the example $\cdot+x!y+!zy$, unique readability can be verified by inspection. We now prove it now, together with a related useful fact:

Lemma II.4.3 (unique readability) *Let σ be an expression of the lexicon (\mathcal{W}, α) . Then*

1. *No proper initial segment of σ is an expression.*
2. *If σ has first symbol s of arity n , then there exist unique expressions $\tau_0, \dots, \tau_{n-1}$ such that σ is $s\tau_0 \cdots \tau_{n-1}$.*

Proof. We prove (1)(2) simultaneously by induction on $|\sigma|$, so assume that they hold for all shorter expressions. Also, note, as remarked above, that the existence part of (2) is immediate from the definition of “expression”.

Now, let σ' be any expression which is an initial segment (possibly not proper) of σ . Since the empty string is not an expression, we must have $\sigma' = s\tau'_0 \cdots \tau'_{n-1}$, where the τ'_i are all expressions. Then τ_0 must be the same as τ'_0 , since otherwise one would be a proper initial segment of the other, contradicting (1) (applied inductively). Likewise, we prove $\tau_i = \tau'_i$ by induction on i : If $\tau_j = \tau'_j$ for all $j < i$, then τ_i and τ'_i begin with the same symbol of σ , so $\tau_i = \tau'_i$ because otherwise one would be a proper initial segment of the other. But now we know that $\sigma' = \sigma$, and we have established both (1) and (2). \square

We shall also need to deal with subexpressions.

Definition II.4.4 *If σ is an expression of the lexicon (\mathcal{W}, α) , then a subexpression of σ is a consecutive sequence from σ which is also an expression.*

For example, say σ is $++xy+zu$, which is Polish for $(x + y) + (z + u)$. Then $+xy$ is a subexpression, as is the one-symbol expression x . $+xu$ is not a subexpression; it is an expression taken from the symbols in σ , but it is not consecutive. $+xy+$ is not a subexpression; it is consecutive, but it is not an expression. In fact, if we focus on the second $+$ in σ , we see that $+xy$ is the only subexpression beginning with that $+$. More generally:

Lemma II.4.5 *If σ is an expression of the lexicon (\mathcal{W}, α) , then every occurrence of a symbol in σ begins a unique subexpression.*

Proof. Uniqueness is immediate from Lemma II.4.3, and existence is easily proved by induction from the definition (II.4.1) of “expression”. \square

Definition II.4.6 *If σ is an expression of the lexicon (\mathcal{W}, α) , then the scope of an occurrence of a symbol in σ is the unique subexpression which it begins.*

If σ is $++xy+zu$, then the scope of the first $+$ is σ itself, the scope of the second $+$ is $+xy$, and the scope of the third $+$ is $+zu$. We remark that formally σ is a function on a finite ordinal, and the somewhat informal word “occurrence” in the last lemma and definition could be made more formal by referring to some $\sigma(i)$. For example, if σ is $++xy+zu$, then $\text{dom}(\sigma) = 7$, and the three $+$ signs are respectively $\sigma(0), \sigma(1), \sigma(4)$.

We conclude this section with a few additional remarks.

Working in set theory, the set \mathcal{W} of symbols can have arbitrary cardinality, but if \mathcal{W} is finite or countable, it is conventional to assume that $\mathcal{W} \subseteq HF$. Then, by Exercise I.15.10, all expressions of (\mathcal{W}, α) will lie in HF also, in line with our expectation, expressed in Section I.15, that finite mathematics lives within HF .

A remark on the decidability of parsing Polish notation: Say \mathcal{W} is finite, so that we may enter a string $\sigma \in \mathcal{W}^{<\omega}$ into a computer program. A *parsing algorithm* will decide whether or not σ is an expression, and, if it is, return the unique expressions $\tau_0, \dots, \tau_{n-1}$ such that σ is $s\tau_0 \cdots \tau_{n-1}$. The *existence* of such an algorithm is clear from the definitions:

Simply list all possible ways of writing σ as $s\tau_0 \cdots \tau_{n-1}$ (with the τ_i arbitrary non-empty elements of $\mathcal{W}^{<\omega}$), and for each one of these ways, call the algorithm (recursively) to decide whether each τ_i is an expression. This algorithm is clearly horribly inefficient. A much more efficient procedure is given by:

Exercise II.4.7 Given a lexicon (\mathcal{W}, α) and $\sigma = (s_0, \dots, s_{k-1}) \in \mathcal{W}^{<\omega}$, let $\text{count}(\sigma) = \sum_{j < k} (\alpha(s_j) - 1)$. Let $\text{count}(\epsilon) = 0$, where ϵ is the empty sequence. Then σ is an expression iff $\text{count}(\sigma) = -1$ and $\text{count}(\sigma \upharpoonright \ell) \geq 0$ whenever $\ell < |\sigma|$.

Using this result, we can decide very quickly whether σ is an expression, and then, if it is, compute where τ_0 ends, and then where τ_1 ends, and so forth.

Polish notation has some use in computing practice, as does its twin, Reverse Polish Notation (postfix, or RPN), where the symbols are written the end (e.g., $++xy+zu$ would be written as $xy+zu++$). The programming language Lisp is written in a variant of our Polish notation. RPN is used for input to stack machines; in particular a number of calculators designed by Hewlett-Packard, starting in 1972. If x, y, z, u represent numbers, you enter the computation in the order $xy+zu++$. Many compilers translate code into RPN as a step before producing machine code, and RPN is used in the language Postscript, which is used to describe text documents.

Those readers who are familiar with computer languages will realize that the discussion of this section is very primitive. We have chosen Polish notation because we could easily give a formal mathematical definition of it and a proof of unique readability. See a text on compiler design for a discussion of the syntax and parsing of computer languages. In a computer language, the really basic “symbols” are the characters, and strings of them are used to form words. For example, we might write `(alice + bob) + (bill + mary)` in C or Java or Python. Writing this in Polish would give us `+ + alice bob + bill mary`. Note that now we need a space character to separate the words, or *tokens*; it is these tokens which form the elements our \mathcal{W} , and a preliminary stage of lexical analysis is needed to parse the string of characters into a string of 7 tokens. Then, this string would be sent to a parser for Polish notation, which would recognize it as a well-formed expression, assuming that `+` has arity 2 and that `alice`, `bob`, `bill`, and `mary` all have arity 0.

II.5 First-Order Logic Syntax

We begin in the spirit of Section 0.2, which described logical notation as it is used informally in mathematics. The distinction between syntax and semantics of formal logic was discussed in Section 0.4. In this section, we give a precise description of logical syntax, together with some *informal* discussion of semantics. Consider three logical sentences, which we display first in informal mathematical notation and then in the

corresponding official Polish:

$$\begin{array}{ll}
 SQ : & \forall x(0 < x \rightarrow \exists y(x = y \cdot y)) & \forall x \rightarrow < 0x \exists y = x \cdot yy \\
 EXT : & \forall x, y (\forall z(z \in x \leftrightarrow z \in y) \rightarrow x = y) & \forall x \forall y \rightarrow \forall z \leftrightarrow \in z x \in z y = xy \\
 EM : & \exists y \forall x(x \notin y) & \exists y \forall x \neg \in xy
 \end{array}$$

It is part of the *syntax* to say that these are strings made up of symbols such as the implication sign \rightarrow , some variables x, y , some quantifiers \forall, \exists , etc. Also, the syntax will tell us that these strings really are logical sentences, whereas the string $xx\forall\rightarrow$, which is made up of some of the same symbols, isn't a logical sentence. The definition of formal provability, $\Sigma \vdash \varphi$, is also part of syntax (see Section II.10). For example, we saw in Section I.6 that one can prove in ZF that there is an empty set. Our formal proof theory will yield $ZF \vdash EM$. Of course, EXT is the Extensionality Axiom of ZF .

It is part of the *semantics* to attach a *meaning* to logical sentences. For example, SQ has a definite truth value, T or F , in every ordered field. Since it asserts that every positive element has a square root, it is T in \mathbb{R} and F in \mathbb{Q} . Later, we shall say precisely how the truth value is defined. Exercise I.2.1 provided some informal examples of finite models in which EXT and EM had various truth values. However, the string $xx\forall\rightarrow$ is meaningless — that is, we do not define a truth value for it.

The reader will see from these examples that the informal notation is much easier to understand than the Polish notation. That will not be a problem in this book; most of the time, we shall continue to use the informal notation as an abbreviation for the formal notation. Abbreviations are discussed further in Section II.6. For model theory, it is important only to know that there is *some* way of defining the syntax of a formal language so that unique readability holds; and Polish notation yields an easy way of providing such a syntax. We are not attempting to rewrite all of mathematics into this formal notation. As a practical matter, it is quite easy to define a formal grammar in which the informal renditions of SQ, EXT, EM are formally correct; one just has to spell out the rules for infix operators and for the use of parentheses. However, it is quite complex to design a formal grammar which would allow one to write the Continuum Hypothesis as a formal sentence in the language of set theory. Written just using \in and $=$, CH would be enormously long and incomprehensible, regardless of whether or not we use Polish notation. Our statement of it in Definition I.13.8 presupposed a long chain of definitions. Computer verification languages such as Isabelle/ZF [17] and Mizar [25] allow the user to build such definitions as part of the formal language (just as one defines functions in programming languages), but then it becomes very lengthy to give a precise definition of the syntax and semantics of these languages.

Now, to define our syntax, we begin by specifying what the symbols are. They are partitioned into two types, *logical symbols* and *nonlogical symbols*. The *logical symbols* will be fixed throughout this entire chapter. The ones occurring in SQ, EXT, EM are $=, \forall, \exists, x, y, \rightarrow$, but there are others, listed below. The *nonlogical symbols* vary with context. SQ , which uses $0, <, \cdot$, would be appropriate if we're talking about ordered

fields or ordered rings. *EXT*, *EM* use \in , and might be appropriate if we're discussing models for some of the set theory axioms, as we did in Exercise I.2.1.

Definition II.5.1 *Our logical symbols are the eight symbols:*

$$\wedge \quad \vee \quad \neg \quad \rightarrow \quad \leftrightarrow \quad \forall \quad \exists \quad =$$

together with a countably infinite set VAR of variables. We'll usually use u, v, w, x, y, z , perhaps with subscripts, for variables.

These symbols will be fixed in this book, but many variations on these occur in other books. On a trivial level, some people use “&” for “and” rather than “ \wedge ”. Also, many books use infix notation rather than Polish, so that the logical symbols would have to include parentheses, to distinguish between $(\varphi \vee \psi) \wedge \chi$ and $\varphi \vee (\psi \wedge \chi)$. In our Polish notation, these are, respectively, $\wedge \forall \varphi \psi \chi$ and $\vee \varphi \wedge \psi \chi$.

Somewhat less trivially, some authors use other collections of propositional connectives. For example, one might use only \vee and \neg , since other ones may be expressed in terms of these; e.g., $\rightarrow \varphi \psi$ is logically equivalent to $\vee \neg \varphi \psi$. We shall define logical equivalence precisely later (see Definition II.8.2).

On a more basic level, note that $=$ is a logical symbol. This is the usual convention in modern mathematical logic, but not in some older works. The effect of this on the syntax is that sentences just using $=$, such as $\forall x = xx$, will always be present regardless of the nonlogical symbols. The effect of this on the semantics is that the meaning of $=$ will be fixed, so that $\forall x = xx$ will be logically valid (see Definition II.8.1) — that is, true in all models. Thus, as pointed out in Section I.2, this, and other valid statements about $=$, are not listed when we list the axioms of a theory, such as the axioms for set theory or for group theory.

Definition II.5.2 *A lexicon for predicate logic consists of a set \mathcal{L} (of nonlogical symbols), partitioned into disjoint sets $\mathcal{L} = \mathcal{F} \cup \mathcal{P}$ (of function and predicate symbols). \mathcal{F} and \mathcal{P} are further partitioned by arity: $\mathcal{F} = \bigcup_{n \in \omega} \mathcal{F}_n$, and $\mathcal{P} = \bigcup_{n \in \omega} \mathcal{P}_n$. Symbols in \mathcal{F}_n are called n -place or n -ary function symbols. Symbols in \mathcal{P}_n are called n -place or n -ary predicate symbols. Symbols in \mathcal{F}_0 are called constant symbols. Symbols in \mathcal{P}_0 are called proposition letters.*

In most of the elementary uses of logic, \mathcal{L} is finite, so most of the \mathcal{F}_n and \mathcal{P}_n are empty. For example, in axiomatizing set theory, $\mathcal{L} = \mathcal{P}_2 = \{\in\}$, with all the other \mathcal{F}_n and \mathcal{P}_n empty. In axiomatizing group theory, the choice of \mathcal{L} varies with the presentation. In Section 0.4, we wrote the axioms as $GP = \{\gamma_1, \gamma_2\}$:

$$\begin{aligned} \gamma_1. & \forall xyz[x \cdot (y \cdot z) = (x \cdot y) \cdot z] \\ \gamma_2. & \exists u[\forall x[x \cdot u = u \cdot x = x] \wedge \forall x \exists y[x \cdot y = y \cdot x = u]] \end{aligned}$$

Here, $\mathcal{L} = \mathcal{F}_2 = \{\cdot\}$, with all the other \mathcal{F}_n and \mathcal{P}_n empty. It is easy to rewrite γ_1 and γ_2 into Polish notation (see Section II.6 on abbreviations). More importantly, note that many books will write the axioms as: $\{\gamma_1, \gamma_{2,1}, \gamma_{2,2}\}$, replacing γ_2 by:

$$\begin{aligned} \gamma_{2,1}. & \forall x[x \cdot 1 = 1 \cdot x = x] \\ \gamma_{2,2}. & \forall x[x \cdot i(x) = i(x) \cdot x = 1] \end{aligned}$$

Now, \mathcal{L} has become $\{\cdot, i, 1\}$, with $\mathcal{F}_2 = \{\cdot\}$, $\mathcal{F}_1 = \{i\}$, and $\mathcal{F}_0 = \{1\}$. Most people would write the inverse, $i(x)$, as x^{-1} . This is a bigger lexicon, but it makes the axioms simpler. In particular, with this lexicon, the class of groups forms an *equational variety*; that is, it is defined by a list of universally quantified equations. Equational varieties are a special kind of first-order class which have some additional interesting properties; see Section II.14. The fact that the two ways of axiomatizing groups are “essentially equivalent” is taken up in Section II.15; one needs the fact that on the bases of $\{\gamma_1, \gamma_2\}$, one can prove that inverses exist and are unique, so that it is “harmless” to introduce a function symbol denoting inverses. This is related to the fact that it is “harmless” to introduce defined functions, such as $x \cup y$, when we developed *ZF*.

When discussing abelian groups, it is conventional to write the axioms additively, using $\mathcal{L} = \{+, -, 0\}$, where $\mathcal{F}_2 = \{+\}$, $\mathcal{F}_1 = \{-\}$, and $\mathcal{F}_0 = \{0\}$. Note that $-$ is unary here, in analogy with the inverse $i(x)$ or x^{-1} in multiplicative groups, so that $x - y$ is really an abbreviation for $x + (-y)$ (our Polish syntax does not allow a symbol to be both binary and unary). More on abbreviations in Section II.6. When discussing ordered rings or fields (such as $\mathbb{Z}, \mathbb{Q}, \mathbb{R}$), one often takes

$$\mathcal{L}_{OR} = \{<, +, \cdot, -, 0, 1\} \quad ,$$

where $\mathcal{F}_2 = \{+, \cdot\}$, $\mathcal{F}_1 = \{-\}$, $\mathcal{F}_0 = \{0, 1\}$, and $\mathcal{P}_2 = \{<\}$; the sentence *SQ* is expressed in this lexicon.

In the above examples, we have been using familiar mathematical symbols to denote familiar algebraic functions and relations. In abstract discussions, we shall often use ordinary letters to denote functions and relations. For example, we might write the logical sentence $\forall x(p(x, x) \rightarrow \exists y q(x, f(y, x), g(f(g(x), g(y))))))$ (which abbreviates the Polish $\forall \rightarrow p x x \exists y q x f y x g f g x g y$); this makes some meaningful (perhaps uninteresting) assertion as long as $p \in \mathcal{P}_2$, $q \in \mathcal{P}_3$, $f \in \mathcal{F}_2$, and $g \in \mathcal{F}_1$.

Before we define the notion of logical sentence, we first define the more elementary notion of *term*. Informally, the terms of \mathcal{L} denote objects; for example, if $f \in \mathcal{F}_2$ and $g \in \mathcal{F}_1$, then $g f g x g y$ will be a term. This is a sequence of 6 symbols. The g, f are non-logical symbols, and the variables x, y are logical symbols.

Definition II.5.3 *Given a lexicon $\mathcal{L} = \mathcal{F} \cup \mathcal{P} = \bigcup_{n \in \omega} \mathcal{F}_n \cup \bigcup_{n \in \omega} \mathcal{P}_n$, as in Definition II.5.2:*

1. *The terms of \mathcal{L} are the well-formed expressions of the Polish lexicon $\mathcal{F} \cup \text{VAR}$, as defined in Definition II.4.1, where symbols in *VAR* have arity 0 and symbols in \mathcal{F}_n have arity n .*

2. The atomic formulas of \mathcal{L} are those sequences of symbols of the form $p\tau_1 \cdots \tau_n$, where $n \geq 0$, τ_1, \dots, τ_n are terms of \mathcal{L} , and either $p \in \mathcal{P}_n$ or p is the symbol $=$ and $n = 2$.
3. The formulas of \mathcal{L} are those sequences of symbols constructed by the rules:
 - a. All atomic formulas are formulas.
 - b. If φ is a formula and $x \in \text{VAR}$, then $\forall x\varphi$ and $\exists x\varphi$ are formulas.
 - c. If φ is a formula then so is $\neg\varphi$.
 - d. If φ and ψ are formulas then so are $\forall\varphi\psi$, $\wedge\varphi\psi$, $\rightarrow\varphi\psi$, and $\leftrightarrow\varphi\psi$.

We needed to make a special case for “=” in (2) because “=” is a logical symbol, not a member of \mathcal{P}_2 . Observe that all the formulas and terms are well-formed expressions of the Polish lexicon $\mathcal{F} \cup \mathcal{P} \cup \text{VAR} \cup \{\wedge, \vee, \neg, \rightarrow, \leftrightarrow, \forall, \exists, =\}$, where \neg has arity 1 and the members of $\{\wedge, \vee, \rightarrow, \leftrightarrow, \forall, \exists, =\}$ have arity 2. However, many well-formed expressions are neither formulas nor terms (e.g., $\forall xy$). This means that our unique readability Lemma II.4.3 tells us *more* than what we need, not *less*. For example, say χ is $\forall\varphi\psi$, with φ and ψ formulas. When we assign a truth value to χ (see Section II.7), it will be important to know that the same χ cannot be written in a different way, as $\forall\varphi'\psi'$, with φ' and ψ' also both formulas. In fact, Lemma II.4.3 says that this is impossible even if φ' and ψ' are arbitrary well-formed expressions.

The discussion of *scope* for Polish expression (see Definition II.4.6) applies to give us the definition of free and bound variables. First, observe, by induction,

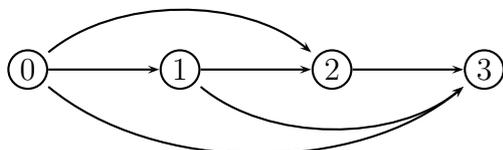
Lemma II.5.4 *In a formula φ , the scope of any occurrence in φ of any of the symbols in $\mathcal{P} \cup \text{VAR} \cup \{\wedge, \vee, \neg, \rightarrow, \leftrightarrow, \forall, \exists, =\}$ is a formula.*

This scope is often called a *subformula* of φ .

Definition II.5.5 *An occurrence of a variable y in a formula φ is bound iff it lies inside the scope of a \forall or \exists acting on (i.e., followed by) y . An occurrence is free iff it is not bound. The formula φ is a sentence iff no variable is free in φ .*

For example, if EXT is the formula $\forall x\forall y\rightarrow\forall z\leftrightarrow\in zx\in zy=xy$ described at the beginning of this section, then the scope of the first \forall is all of EXT , the scope of the second \forall is all of EXT except for the beginning “ $\forall x$ ” and the scope of the third \forall is the subformula $\forall z\leftrightarrow\in zx\in zy$; call this ψ . This third \forall acts on (i.e., is followed by) z , and the three occurrences of z in EXT lie inside ψ , so all occurrences of z in EXT are bound. Likewise, all occurrences of x and y in EXT are bound, so that EXT is a sentence. ψ is not a sentence, since x and y are free in ψ .

In the semantics (to be defined in Section II.7), a sentence will have a definite truth value (true or false) in a given model, although this value might depend on the model (for EXT , see Exercise I.2.1). A formula expresses a property of its free variables. For example, in the model



from Exercise I.2.1, EXT is true, and the formula ψ expresses a property of pairs of elements (which is true iff both elements of the pair are the same).

Axioms of a theory are always sentences. For brevity of notation, these are frequently displayed as formulas with the understanding that the free variables are to be universally quantified; this same convention is common throughout mathematics. For example, algebra texts will often write the associative law as $x \cdot (y \cdot z) = (x \cdot y) \cdot z$; call this formula χ , which is $=x \cdot yz \cdot xyz$; it is understood that the axiom is really the sentence γ_1 above, namely $\forall x \forall y \forall z \chi$, which is a *universal closure* of χ :

Definition II.5.6 *If φ is a formula, a universal closure of φ is any sentence of the form $\forall x_1 \forall x_2 \cdots \forall x_n \varphi$, where $n \geq 0$.*

So, if φ is already a sentence, it is a universal closure of φ . Note that we are not specifying any ordering on the variables; this will not matter, since it will be easy to see (Lemma II.8.3) (once we have defined the semantics) that all universal closures are logically equivalent anyway. So, the above χ also has $\forall z \forall y \forall x \chi$, $\forall z \forall x \forall y \chi$, and even $\forall x \forall z \forall x \forall y \forall y \chi$ as universal closures. In listing the axioms of set theory in Section I.2, we said “Axioms stated with free variables are understood to be universally quantified”, meaning that each of the listed axioms should really be replaced by one of its universal closures.

Our definition of syntax allows the same variable to occur both free and bound in the same formula, although some readers might find such usage confusing. For example, with $\mathcal{L} = \{\in\}$, let φ be the formula $\wedge \exists y \in y x \in xy$, which say “ x is non-empty and $x \in y$ ”. The first two occurrences of y are inside the subformula $\exists y \in y x$ and are hence bound, whereas the third occurrence of y is free. φ is logically equivalent (see Definition II.8.2) to the formula $\varphi' : \wedge \exists z \in z x \in xy$, obtained by changing the name of the bound (or dummy) variable y to z . Most people would find φ' easier to read than φ . The same issue arises in calculus; for example we could define

$$f(x, y) = \sin(xy) + \int_1^2 \cos(xy) dy = \sin(xy) + \int_1^2 \cos(xt) dt .$$

Both forms are correct, but most people would prefer the second form, using t as the dummy (or bound) variable of integration.

Remark. We have been using the term *lexicon* for the set \mathcal{L} of non-logical symbols. In the model theory literature, it is more common to say “the *language* \mathcal{L} ”, whereas in works on the theory of formal languages, a *language* is a set of strings made up of the basic symbols (such as the set of formulas of \mathcal{L}). Our terminology here is closer to the common English meaning of “lexicon” as the collection of words of a language; e.g.

“cat”, “hat”, etc. whereas a sentence in the English *language* is a string of these words, such as “The cat wore a hat”.

II.6 Abbreviations

It is always difficult to translate informal mathematics into a formal logical system. This is especially true of our Polish notation, which the reader undoubtedly has already found a bit painful to use. Since the goal of formal logic is its applications to mathematics, we need to introduce some abbreviations so that we may express statements of mathematical interest without too much pain. We shall classify such abbreviations roughly as *low level*, *middle level*, and *high level*:

Low Level: The true (unabbreviated) formula or term is determined uniquely by standard mathematical conventions. This was the case, for example, in the sentences SQ , EXT , EM from Section II.5, where the Polish notation was simply the obvious translation of the more standard notation. The “standard mathematical conventions” include the usual precedence relations in algebra. For example, $x + y \cdot z$ and $x + yz$ both abbreviate $+x \cdot yz$, not $\cdot +xyz$ because the standard convention is that \cdot binds more tightly than $+$, and that the \cdot may be omitted in products. The propositional connectives are usually given the precedence $\{\neg, \wedge, \vee, \rightarrow, \leftrightarrow\}$, so, for example, the propositional sentence $\neg p \vee q \rightarrow r$ abbreviates $\rightarrow \vee \neg pq$. This particular ordering of the connectives is not completely universal in the literature, and we shall frequently insert parentheses if there is a danger of confusion. It is only important to remember that \neg binds the most tightly, and that both \wedge and \vee bind more tightly than either \rightarrow or \leftrightarrow .

Middle Level: The true (unabbreviated) formula or term is clear only up to logical equivalence (see Definition II.8.2 for precisely what this means). Here, there is no “standard mathematical convention”, and often ordinary English is mixed in with the logical notation. For example, it is important in model theory that for each finite n , one can write a sentence δ_n which says that the universe has size at least n . We might display δ_4 as $\exists w, x, y, z[\textit{they're all different}]$. It never matters which one of a large number of logically equivalent ways we choose to write this in the formal system. One such way is:

$$\exists w, x, y, z[w \neq x \wedge w \neq y \wedge w \neq z \wedge x \neq y \wedge x \neq z \wedge y \neq z] .$$

Before you translate this sentence into Polish notation, you will have to decide whether \wedge associates left or right (that is, $p \wedge q \wedge r$ might abbreviate $\wedge \wedge pqr$ or $\wedge p \wedge qr$). It is not necessary to make a convention here, since both translations are logically equivalent.

High Level: The true (unabbreviated) formula or term is clear only up to equivalence with respect to some theory (see Definition II.8.4). This is common in algebra. For example, say we are using $\mathcal{L} = \{+, \cdot, -, 0, 1\}$ to discuss rings with a unity (or 1 element). Then, it is important to know that polynomials with integer coefficients “are” terms in our formal logic. Thus, $3x$ can abbreviate $x + (x + x)$; but it could also abbreviate $(x + x) + x$. The equivalence $\forall x[x + (x + x) = (x + x) + x]$ is not logically valid, since

it fails when $+$ is not associative, but it is valid in rings. Also, $3x$ might abbreviate $(1 + 1 + 1) \cdot x$. As long as one is working in rings with unity, it is never necessary to spell out which formal term is really meant by $3x$.

This equivalence with respect to a theory was used extensively in our development of set theory. For example, we said when discussing the axioms in Section I.2 that logical formulas with defined notions are viewed as abbreviations for formulas in the lexicon $\mathcal{L} = \{\in\}$. Then in Section I.6, we defined \emptyset to be the (unique) y such that $\text{emp}(y)$, where $\text{emp}(y)$ abbreviates $\forall x[x \notin y]$. But exactly what formula of \mathcal{L} does “ $\emptyset \in z$ ” abbreviate? Two possibilities are $\varphi_1 : \exists y[\text{emp}(y) \wedge y \in z]$ and $\varphi_2 : \forall y[\text{emp}(y) \rightarrow y \in z]$.

The formulas φ_1 and φ_2 are not logically equivalent, since the sentence $\forall z[\varphi_1 \leftrightarrow \varphi_2]$ is false in the model:



from Exercise I.2.1, since emp is false of both elements, making φ_1 false of both elements and φ_2 true of both elements. However, $\forall z[\varphi_1 \leftrightarrow \varphi_2]$ is a logical consequence of the axioms of ZF (which imply that there is a unique empty set), so in developing ZF , it is never necessary to spell out which abbreviation is meant. In Section I.2, when we displayed the Axiom of Infinity as an official sentence of \mathcal{L} , we chose to use φ_1 .

Of course, the underlined terms in the previous paragraph still need to be given precise definitions; this will be done in the next two sections.

On a still higher level, the statement “ x is countable” could in principle be written as a formula just using $\mathcal{L} = \{\in\}$, but it is not made clear, in this book or in any other book on set theory, exactly which formula we have in mind. It is of fundamental importance that there is *some* such formula, because the Comprehension Axiom (see Sections I.2 and I.6) asserts that $\{x \in z : Q(x)\}$ exists for properties $Q(x)$ *expressible in \mathcal{L}* , so we need that “ x is countable” to be so expressible to justify forming the set $\{x \in z : x \text{ is countable}\}$. A more detailed discussion of this issue of defined notions in the development of an axiomatic theory is taken up in Section II.15.

II.7 First-Order Logic Semantics

We said in Section II.5 that the *semantics* will attach a *meaning* to logical sentences. Although in the concrete examples we discussed, the meaning should be clear informally, it is important to give this meaning a precise mathematical definition. Consider, for example, the sentence SQ :

$$SQ : \quad \forall x(0 < x \rightarrow \exists y(x = y \cdot y)) \quad \forall x \rightarrow < 0x \exists y = x \cdot yy$$

We said that SQ asserts that every positive element has a square root, so SQ is T in \mathbb{R} and F in \mathbb{Q} . More generally, we can evaluate the truth or falsity of SQ in an

arbitrary abstract structure \mathfrak{A} . This structure should consist of a domain of discourse A (a non-empty set over which the variables “range”), together with a binary product operation $\cdot_{\mathfrak{A}}$, a binary relation $<_{\mathfrak{A}}$, and a distinguished element $0_{\mathfrak{A}}$. We shall write $\mathfrak{A} = (A; \cdot_{\mathfrak{A}}, <_{\mathfrak{A}}, 0_{\mathfrak{A}})$. We are following here the usual convention in model theory of using italic letters, A, B, C, D, \dots for the domain of discourse of a structure and the corresponding gothic or fraktur letters, $\mathfrak{A}, \mathfrak{B}, \mathfrak{C}, \mathfrak{D}, \dots$ for the full structure. To be really formal, this subscripted notation $(\cdot_{\mathfrak{A}}, <_{\mathfrak{A}}, 0_{\mathfrak{A}})$ indicates the presence of a function which assigns to symbols of \mathcal{L} a semantic entity of the correct type:

Definition II.7.1 *Given a lexicon for predicate logic, $\mathcal{L} = \mathcal{F} \cup \mathcal{P} = \bigcup_{n \in \omega} \mathcal{F}_n \cup \bigcup_{n \in \omega} \mathcal{P}_n$ (see Definitions II.5.2 and II.5.3), a structure for \mathcal{L} is a pair $\mathfrak{A} = (A, \mathcal{I})$ such that A is a non-empty set and \mathcal{I} is a function with domain \mathcal{L} with each $\mathcal{I}(s)$ a semantic entity of the correct type; specifically, writing $s_{\mathfrak{A}}$ for $\mathcal{I}(s)$:*

- ☞ If $f \in \mathcal{F}_n$ with $n > 0$, then $f_{\mathfrak{A}} : A^n \rightarrow A$.
- ☞ If $p \in \mathcal{P}_n$ with $n > 0$, then $p_{\mathfrak{A}} \subseteq A^n$.
- ☞ If $c \in \mathcal{F}_0$, then $c_{\mathfrak{A}} \in A$.
- ☞ If $p \in \mathcal{P}_0$, then $p_{\mathfrak{A}} \in 2 = \{0, 1\} = \{F, T\}$.

Note the special case for $n = 0$. Symbols in \mathcal{F}_0 are constant symbols, so they denote an element of the universe of the structure. Symbols in \mathcal{P}_0 are proposition letters, so they denote a truth value, F or T ; for these, the universe of the structure is irrelevant. We follow the usual convention of using 0 to denote “false” and 1 to denote “true”.

We are also following the usual convention in model theory of requiring $A \neq \emptyset$, since allowing the empty structure leads to some pathologies later (see Remark II.8.16).

Unless there is danger of confusion, we shall be fairly informal in denoting structures. For example, if $\mathcal{L} = \{p, f\}$ with $p \in \mathcal{P}_2$ and $f \in \mathcal{F}_1$, we might say, “let $\mathfrak{A} = (\mathbb{R}; <, \cos)$ ”, since this makes it clear that $p_{\mathfrak{A}}$ is the $<$ relation and $f_{\mathfrak{A}}$ is the cosine function. But if $\mathcal{L} = \mathcal{F}_1 = \{f, g\}$, the definition, “let $\mathfrak{A} = (\mathbb{R}; \sin, \cos)$ ” might be ambiguous, and we need to say, e.g., that $f_{\mathfrak{A}}$ is the cosine function and $g_{\mathfrak{A}}$ is the sine function. If $\mathcal{L} = \{+, \cdot, 0, 1\}$, we might simply say, “let $\mathfrak{A} = \mathbb{R}$ ”, since it is usually clear from context that the symbols $+, \cdot, 0, 1$ denote the expected addition, multiplication, zero, and one in the real numbers. We are perpetuating here the standard abuse of notation employed in algebra, where the same symbols $+, \cdot$ are used to denote the addition and multiplication functions of some (any) ring, as well as to denote symbols in formal expressions such as the term (polynomial) $x \cdot y + z + 1$.

We shall rarely use the “ (A, \mathcal{I}) ” terminology from Definition II.7.1. We gave it mainly to make it clear what \mathfrak{A} is set-theoretically; it is also useful in discussing reducts and expansions (see Definition II.8.14).

We shall next define the notion “ $\mathfrak{A} \models \varphi$ ” (φ is true in \mathfrak{A}), where \mathfrak{A} is a structure for \mathcal{L} and φ is a sentence of \mathcal{L} . Roughly, this is done by recursion on φ . Consider, for example, the above sentence SQ , where $\mathcal{L} = \{<, 0, \cdot\}$, and $\mathfrak{A} = \mathbb{Q}$ (with the usual order,

zero, and product understood). That this is false in \mathfrak{A} will be expressed symbolically in one of the following three equivalent ways:

$$1. \mathfrak{A} \not\models SQ \quad 2. \mathfrak{A} \models \neg SQ \quad 3. \text{val}_{\mathfrak{A}}(SQ) = F$$

The double turnstile symbol “ \models ” here is usually read “models” or “satisfies”. Forms (1) and (2) are the more standard terminologies, but we use (3) to emphasize that given \mathfrak{A} , we are defining a function $\text{val}_{\mathfrak{A}}$ on the sentences of \mathcal{L} . Our precise definition of $\text{val}_{\mathfrak{A}}(SQ)$ will unwind the syntax of SQ ; it begins with a $\forall x$; informally, the statement isn’t true for all $x \in \mathbb{Q}$, for example,

$$\text{val}_{\mathfrak{A}}(0 < x \rightarrow \exists y(x = y \cdot y)) [2] = F \quad . \quad (A)$$

Since the formula $0 < x \rightarrow \exists y(x = y \cdot y)$ is not a sentence, but has x as a free variable, the “[2]” is needed to say that we are interpreting the x as the rational number 2, whereas

$$\text{val}_{\mathfrak{A}}(0 < x \rightarrow \exists y(x = y \cdot y)) [4] = T \quad . \quad (B)$$

The reason for (A) is that 2 is positive but does not have a square root in \mathbb{Q} ; formally:

$$\text{val}_{\mathfrak{A}}(0 < x) [2] = T \quad , \quad (C)$$

but

$$\text{val}_{\mathfrak{A}}(\exists y(x = y \cdot y)) [2] = F \quad , \quad (D)$$

so that (A) follows by the usual truth table for \rightarrow (i.e., $(T \rightarrow F) = F$). We explain (C) and (D) by using, respectively, the meaning of $<_{\mathfrak{A}}$ and the meaning of \exists . If we use the official Polish notation, then the way we (recursively) compute the value of $\text{val}_{\mathfrak{A}}(\varphi)$ is determined by the first (leftmost) symbol of φ . Because of this, we present the official definition of val (Definitions II.7.4, II.7.6, II.7.8) using the official Polish notation.

Before giving a precise definition, one more remark on the “[2]” terminology: This notation easily extends to formulas with several free variables, but there is a danger of ambiguity. For example, with $\mathfrak{A} = \mathbb{Q}$, the meaning of $\text{val}_{\mathfrak{A}}(y = x \cdot z)[2, 6, 3]$ is not clear, since it does not specify which of y, x, z get replaced by which of 2, 6, 3. In most cases, this is clear from context (one is either ordering the variables alphabetically or by their order of occurrence in the formula), but when it is not, we shall use the notation:

$$\text{val}_{\mathfrak{A}}(y = x \cdot z) \left[\begin{array}{c} x \ y \ z \\ 2 \ 6 \ 3 \end{array} \right] = T \quad \text{val}_{\mathfrak{A}}(y = x \cdot z) \left[\begin{array}{c} y \ x \ z \\ 2 \ 6 \ 3 \end{array} \right] = F \quad .$$

The array notation $\left[\begin{array}{c} x \ y \ z \\ 2 \ 6 \ 3 \end{array} \right]$ really denotes a *function* σ whose domain is the set of variable symbols $\{x, y, z\}$, and $\sigma(x) = 2, \sigma(y) = 6, \sigma(z) = 3$. So, we need to define $\text{val}_{\mathfrak{A}}(\varphi)[\sigma]$; for a fixed \mathfrak{A} , this is a function of φ and σ , and will take values in $2 = \{0, 1\} = \{F, T\}$. It will be defined by recursion on φ . Now, to compute $\text{val}_{\mathfrak{A}}(y = x \cdot z)[\sigma]$, we need to use the function $\cdot_{\mathfrak{A}}$ to compute $\text{val}_{\mathfrak{A}}(x \cdot z) \in A$ and see if that is the same as $\sigma(y)$. More generally, before defining the truth value of formulas, we need to define $\text{val}_{\mathfrak{A}}\tau[\sigma] \in A$ for terms τ .

Definition II.7.2 For terms τ , let $V(\tau)$ be the set of variables which occur in τ . For formulas φ , let $V(\varphi)$ be the set of variables which have a free occurrence in φ .

Definition II.7.3 If α is either a term or a formula, an assignment for α in A is a function σ such that $V(\alpha) \subseteq \text{dom}(\sigma) \subseteq \text{VAR}$ and $\text{ran}(\sigma) \subseteq A$.

Definition II.7.4 If \mathfrak{A} is a structure for \mathcal{L} , then we define $\text{val}_{\mathfrak{A}}(\tau)[\sigma] \in A$ whenever τ is a term of \mathcal{L} and σ is an assignment for τ in A as follows:

1. $\text{val}_{\mathfrak{A}}(x)[\sigma] = \sigma(x)$ when $x \in \text{dom}(\sigma)$.
2. $\text{val}_{\mathfrak{A}}(c)[\sigma] = c_{\mathfrak{A}}$ when $c \in \mathcal{F}_0$.
3. $\text{val}_{\mathfrak{A}}(f\tau_1 \cdots \tau_n)[\sigma] = f_{\mathfrak{A}}(\text{val}_{\mathfrak{A}}(\tau_1)[\sigma], \dots, \text{val}_{\mathfrak{A}}(\tau_n)[\sigma])$ when $f \in \mathcal{F}_n$ and $n > 0$.

If $V(\tau) = \emptyset$, then $\text{val}_{\mathfrak{A}}(\tau)$ abbreviates $\text{val}_{\mathfrak{A}}(\tau)[\sigma]$.

Again, with $\mathfrak{A} = \mathbb{Q}$:

$$\text{val}_{\mathfrak{A}}(x \cdot y) \left[\begin{array}{l} x \ y \\ 2 \ 6 \end{array} \right] = \text{val}_{\mathfrak{A}}(x) \left[\begin{array}{l} x \ y \\ 2 \ 6 \end{array} \right] \cdot_{\mathfrak{A}} \text{val}_{\mathfrak{A}}(y) \left[\begin{array}{l} x \ y \\ 2 \ 6 \end{array} \right] = 2 \cdot_{\mathfrak{A}} 6 = 12 \ .$$

We are using successively clauses (3) and (1) of the definition and the meaning of $\cdot_{\mathfrak{A}}$. Note that in the definition of “assignment”, we are allowing $\text{dom}(\sigma)$ to be a proper superset of $V(\alpha)$; otherwise clauses such as (3) would be rather awkward to state. However, one easily verifies by induction:

Exercise II.7.5 $\text{val}_{\mathfrak{A}}(\tau)[\sigma]$ only depends on $\sigma \upharpoonright V(\tau)$; that is, if $\sigma' \upharpoonright V(\tau) = \sigma \upharpoonright V(\tau)$ then $\text{val}_{\mathfrak{A}}(\tau)[\sigma'] = \text{val}_{\mathfrak{A}}(\tau)[\sigma]$.

In particular, when $V(\sigma) = \emptyset$, the notation $\text{val}_{\mathfrak{A}}(\tau)$ for $\text{val}_{\mathfrak{A}}(\tau)[\sigma]$ is unambiguous.

Definition II.7.4 was by recursion on the length of terms, while the definition of the truth value for an atomic formula is explicit:

Definition II.7.6 If \mathfrak{A} is a structure for \mathcal{L} , then we define $\text{val}_{\mathfrak{A}}(\varphi)[\sigma] \in \{0, 1\} = \{F, T\}$ whenever φ is an atomic formula of \mathcal{L} and σ is an assignment for φ in A as follows:

1. $\text{val}_{\mathfrak{A}}(p)[\sigma] = p_{\mathfrak{A}}$ when $p \in \mathcal{P}_0$.
2. $\text{val}_{\mathfrak{A}}(p\tau_1 \cdots \tau_n)[\sigma] = T$ iff $(\text{val}_{\mathfrak{A}}(\tau_1)[\sigma], \dots, \text{val}_{\mathfrak{A}}(\tau_n)[\sigma]) \in p_{\mathfrak{A}}$ when $p \in \mathcal{P}_n$ and $n > 0$.
3. $\text{val}_{\mathfrak{A}}(=\tau_1\tau_2)[\sigma] = T$ iff $\text{val}_{\mathfrak{A}}(\tau_1)[\sigma] = \text{val}_{\mathfrak{A}}(\tau_2)[\sigma]$.

Note that clause (3) is needed because $=$ is a logical symbol, not a symbol \mathcal{L} , and \mathfrak{A} does not assign a meaning to $=$; rather, $\tau_1 = \tau_2$ always means that τ_1, τ_2 are the same object.

The following definition will be useful when defining the value of formulas:

Definition II.7.7 $\sigma + (y/a) = \sigma \upharpoonright (VAR \setminus \{y\}) \cup \{(y, a)\}$.

That is, we assign y value a , if necessary discarding the value σ gives to y . For example,

$$\begin{bmatrix} x \\ 4 \end{bmatrix} + (y/2) = \begin{bmatrix} x & y \\ 4 & 2 \end{bmatrix} \quad \begin{bmatrix} x & y & z \\ 4 & 5 & 6 \end{bmatrix} + (y/2) = \begin{bmatrix} x & y & z \\ 4 & 2 & 6 \end{bmatrix} .$$

The truth value of a formula φ is, like the value of a term, computed recursively:

Definition II.7.8 If \mathfrak{A} is a structure for \mathcal{L} , then we define $\text{val}_{\mathfrak{A}}(\varphi)[\sigma] \in \{0, 1\} = \{F, T\}$ whenever φ is a formula of \mathcal{L} and σ is an assignment for φ in A as follows:

1. $\text{val}_{\mathfrak{A}}(\neg\varphi)[\sigma] = 1 - \text{val}_{\mathfrak{A}}(\varphi)[\sigma]$.
2. $\text{val}_{\mathfrak{A}}(\wedge\varphi\psi)[\sigma]$, $\text{val}_{\mathfrak{A}}(\vee\varphi\psi)[\sigma]$, $\text{val}_{\mathfrak{A}}(\rightarrow\varphi\psi)[\sigma]$, and $\text{val}_{\mathfrak{A}}(\leftrightarrow\varphi\psi)[\sigma]$, are obtained from $\text{val}_{\mathfrak{A}}(\varphi)[\sigma]$ and $\text{val}_{\mathfrak{A}}(\psi)[\sigma]$ using the truth tables (Table 1, page 4) for $\wedge, \vee, \rightarrow, \leftrightarrow$.
3. $\text{val}_{\mathfrak{A}}(\exists y\varphi)[\sigma] = T$ iff $\text{val}_{\mathfrak{A}}(\varphi)[\sigma + (y/a)] = T$ for some $a \in A$.
4. $\text{val}_{\mathfrak{A}}(\forall y\varphi)[\sigma] = T$ iff $\text{val}_{\mathfrak{A}}(\varphi)[\sigma + (y/a)] = T$ for all $a \in A$.

$\mathfrak{A} \models \varphi[\sigma]$ means $\text{val}_{\mathfrak{A}}(\varphi)[\sigma] = T$. If $V(\varphi) = \emptyset$ (that is, φ is a sentence), then $\text{val}_{\mathfrak{A}}(\varphi)$ abbreviates $\text{val}_{\mathfrak{A}}(\varphi)[\sigma]$, and $\mathfrak{A} \models \varphi$ means $\text{val}_{\mathfrak{A}}(\varphi) = T$.

Of course, clause (1) is equivalent to saying that we are using the usual truth table for \neg . Definition II.7.7 is needed because σ may give values to variables which are not free in φ ; then those values are irrelevant, and they may need to be discarded in computing the truth value of φ . For example, in the rationals:

$$\text{val}_{\mathfrak{A}}(\exists y(x = y \cdot y)) \begin{bmatrix} x & y & z \\ 4 & 5 & 6 \end{bmatrix} = T \quad \text{because} \quad \text{val}_{\mathfrak{A}}(x = y \cdot y) \begin{bmatrix} x & y & z \\ 4 & 2 & 6 \end{bmatrix} = T .$$

As with terms, one easily verifies by induction:

Exercise II.7.9 $\text{val}_{\mathfrak{A}}(\varphi)[\sigma]$ only depends on $\sigma \upharpoonright V(\varphi)$; that is, if $\sigma' \upharpoonright V(\varphi) = \sigma \upharpoonright V(\varphi)$ then $\text{val}_{\mathfrak{A}}(\varphi)[\sigma'] = \text{val}_{\mathfrak{A}}(\varphi)[\sigma]$.

In particular, when φ is a sentence, the notation $\text{val}_{\mathfrak{A}}(\varphi)$ for $\text{val}_{\mathfrak{A}}(\varphi)[\sigma]$ is unambiguous.

We now describe some related semantic notions and then state the two main model theory results to be proved in this chapter.

Definition II.7.10 If \mathfrak{A} is a structure for \mathcal{L} and Σ is a set of sentences of \mathcal{L} , then $\mathfrak{A} \models \Sigma$ iff $\mathfrak{A} \models \varphi$ for each $\varphi \in \Sigma$.

The symbols “ $\mathfrak{A} \models \varphi$ ” are usually read “ \mathfrak{A} satisfies Σ ” or “ \mathfrak{A} is a model for Σ ”. For example, as we said in Section 0.4, a group is a model for the axioms GP of group theory.

Definition II.7.11 If Σ is a set of sentences of \mathcal{L} and ψ is a sentence of \mathcal{L} , then $\Sigma \models \psi$ holds iff $\mathfrak{A} \models \psi$ for all \mathcal{L} -structures \mathfrak{A} such that $\mathfrak{A} \models \Sigma$.

In English, we say that ψ is a *semantic consequence* or *logical consequence* of Σ . Note the overloading of the symbol \models ; it has been given two different meanings in Definitions II.7.10 and II.7.12. This never causes an ambiguity because it is always clear from context whether the object on the left side of the \models (\mathfrak{A} or Σ) is a structure or a set of sentences. In all its uses, the double turnstile \models always refers to semantic notions, whereas the single turnstile \vdash refers to the syntactic notion of provability; $\Sigma \vdash \psi$ means that there is a formal proof of ψ from Σ (to be defined in Section II.10). By the Completeness Theorem (Theorem II.12.1), $\Sigma \vdash \psi$ iff $\Sigma \models \psi$.

Definition II.7.12 *If Σ is a set of sentences of \mathcal{L} then Σ is semantically consistent ($\text{Con}_{\models}(\Sigma)$) iff there is some \mathfrak{A} such that $\mathfrak{A} \models \Sigma$. “inconsistent” means “not consistent”.*

There is also a syntactic notion $\text{Con}_{\vdash}(\Sigma)$, which asserts that Σ cannot prove a contradiction in the formal proof theory (see Section II.10). The Completeness Theorem will also imply that $\text{Con}_{\vdash}(\Sigma)$ iff $\text{Con}_{\models}(\Sigma)$. After that, we drop the subscripts and just write $\text{Con}(\Sigma)$.

The usual axiomatic theories discussed in algebra texts (e.g., groups, rings, and fields) are clearly consistent, since these axiom sets are usually presented together with sample models of them. It is easy to write down “artificial” examples of inconsistent sets of sentences, but the notion of inconsistency occurs naturally in the following lemma, whose proof is immediate from the definition of \models :

Lemma II.7.13 (reductio ad absurdum) *If Σ is a set of sentences of \mathcal{L} and ψ is a sentence of \mathcal{L} , then*

- a. $\Sigma \models \psi$ iff $\Sigma \cup \{\neg\psi\}$ is semantically inconsistent.
- b. $\Sigma \models \neg\psi$ iff $\Sigma \cup \{\psi\}$ is semantically inconsistent.

The proof theory version of this is also true (see Lemma II.11.4), and corresponds to a common step in informal mathematical reasoning: To prove ψ , we reduce $\neg\psi$ to an absurdity; that is, we assume that ψ is false and derive a contradiction. The use of Latin in phrases such as “reductio ad absurdum” and “modus ponens” (see Definition II.10.3) originates with the Scholastic philosophers in the Middle Ages, although the concepts involved go back to Aristotle.

Theorem II.7.14 (Compactness Theorem) *If Σ is a set of sentences of \mathcal{L} :*

1. *If every finite subset of Σ is semantically consistent, then Σ is semantically consistent.*
2. *If $\Sigma \models \psi$, then there is a finite $\Delta \subseteq \Sigma$ such that $\Delta \models \psi$.*

In view of Lemma II.7.13, (1) and (2) are equivalent statements. We shall prove them in Section II.12.

The Compactness Theorem involves only \models , so it is a theorem of “pure model theory”, whereas the Completeness Theorem is a result which relates model theory (\models) to proof theory (\vdash). We shall actually prove the Completeness Theorem first (see Theorem II.12.1). From this, Compactness is easy, since if one replaces \models by \vdash in (2), the result will follow from the fact that formal proofs are finite.

The Löwenheim-Skolem Theorem is another result of “pure model theory” involving the cardinalities of models.

Definition II.7.15 *If \mathfrak{A} is a structure for \mathcal{L} with universe A , then $|\mathfrak{A}|$ denotes $|A|$.*

Likewise, other statements about the size of \mathfrak{A} really refer to $|A|$; for example “ \mathfrak{A} is an infinite model” means that $|A|$ is infinite.

Theorem II.7.16 (Löwenheim–Skolem Theorem) *Let Σ be a set of sentences of \mathcal{L} such that for all finite n , Σ has a (finite or infinite) model of size $\geq n$. Then for all $\kappa \geq \max(|\mathcal{L}|, \aleph_0)$, Σ has a model of size κ .*

Here, $|\mathcal{L}|$ means literally the number of nonlogical symbols. In all examples up to now, this has been finite. As long as $|\mathcal{L}| \leq \aleph_0$, this theorem says, informally, that first-order logic cannot distinguish between infinite cardinalities, since if Σ has an infinite model, it has models of all infinite sizes.

If \mathcal{L} is uncountable, then we really do need $\kappa \geq |\mathcal{L}|$ for the theorem to hold. For example, suppose $\mathcal{L} = \mathcal{F}_0 = \{c_\alpha : \alpha < \lambda\}$, where λ is an infinite cardinal, and $\Sigma = \{c_\alpha \neq c_\beta : \alpha < \beta < \kappa\}$. Then Σ has a model of size κ iff $\kappa \geq \lambda$. This example may seem a bit artificial, but uncountable languages are useful in model theory; for example, they occur in the proof of the Löwenheim–Skolem Theorem in Section II.12; this proof will occur after our proof of the Compactness Theorem provides us with a technique for constructing a model of a desired cardinality.

Note that finite sizes are special. For example, if Σ consists of the one sentence $\forall x, y, z(x = y \vee y = z \vee z = x)$, then Σ has models of sizes 1 and 2, but no other sizes.

II.8 Further Semantic Notions

We collect here a few auxiliary semantic notions.

Definition II.8.1 *If ψ is a formula of \mathcal{L} , then ψ is logically valid iff $\mathfrak{A} \models \psi[\sigma]$ for all \mathcal{L} -structures \mathfrak{A} and all assignments σ for ψ in \mathfrak{A} .*

A sentence ψ is logically valid iff $\emptyset \models \psi$, where \emptyset is the empty set of sentences, since Definition II.7.10 implies that $\mathfrak{A} \models \emptyset$ for all \mathfrak{A} . The formula $x = x$ and the sentence

$\forall x(x = x)$ are logically valid because our definition of \models always interprets the logical symbol $=$ as true identity. Many formulas, such as $\forall xp(x) \rightarrow \neg\exists x\neg p(x)$, are obviously logically valid, and many others, such as $p(x) \rightarrow \forall yp(y)$, are obviously not logically valid. There are many such trivial examples, but by a famous theorem of Church (see Chapter III), there is no algorithm which can decide in general which formulas are logically valid. A subset of the logically valid formulas, the *propositional tautologies* (such as $p(x) \rightarrow \neg\neg p(x)$), is decidable (using truth tables); see Section II.9.

Definition II.8.2 *If φ, ψ are formulas of \mathcal{L} , then φ, ψ are logically equivalent iff the formula $\varphi \leftrightarrow \psi$ is logically valid.*

This is the same as saying that $\mathfrak{A} \models \varphi[\sigma]$ iff $\mathfrak{A} \models \psi[\sigma]$ for all \mathfrak{A} and all σ . For example, $p(x) \vee q(x)$ and $q(x) \vee p(x)$ are logically equivalent. All universal closures of a formula (see Definition II.5.6) are logically equivalent:

Lemma II.8.3 *If φ is a formula, and the sentences ψ and χ are both universal closures of φ , then ψ, χ are logically equivalent.*

Proof. Say y_1, \dots, y_k are the free variables of φ , where $k \geq 0$. Then ψ is of the form $\forall x_1 \forall x_2 \cdots \forall x_n \varphi$, where each y_i is listed at least once in x_1, x_2, \dots, x_n . Note that $\mathfrak{A} \models \psi$ iff $\mathfrak{A} \models \varphi[a_1, \dots, a_k]$ for all $a_1, \dots, a_k \in A$. Since the same is also true for χ , we have $\mathfrak{A} \models \psi$ iff $\mathfrak{A} \models \chi$. \square

In particular, if φ is a sentence, then it is a universal closure of itself and all universal closures of φ are logically equivalent to φ . In view of Lemma II.8.3, if φ is any formula, we shall usually say “*the universal closure of φ* ” to refer to some (any) universal closure of φ , since it usually will not matter which one we use.

There is also a *relative* notion of logical equivalence:

Definition II.8.4 *If φ, ψ are formulas of \mathcal{L} and Σ is a set of sentences of \mathcal{L} , then φ, ψ are equivalent with respect to Σ iff the universal closure of $\varphi \leftrightarrow \psi$ is true in all models of Σ . If τ_1 and τ_2 are terms, then τ_1, τ_2 are equivalent with respect to Σ iff for all $\mathfrak{A} \models \Sigma$ and all assignments σ for τ_1, τ_2 in A , $\text{val}_{\mathfrak{A}}(\tau_1)[\sigma] = \text{val}_{\mathfrak{A}}(\tau_2)[\sigma]$.*

This notion came up in the discussion of abbreviations (see Section II.6). For example, if Σ contains the associative law, then the terms $x \cdot (y \cdot z)$ and $(x \cdot y) \cdot z$ are equivalent with respect to Σ , so as long as we are discussing only models of Σ , it is safe to use xyz as an abbreviation, without having to remember which of these two terms it abbreviates. One could define τ_1, τ_2 to be *logically equivalent* iff they are equivalent with respect to \emptyset , but this is uninteresting by:

Exercise II.8.5 *If $\text{val}_{\mathfrak{A}}(\tau_1)[\sigma] = \text{val}_{\mathfrak{A}}(\tau_2)[\sigma]$ for all \mathfrak{A} and σ , then τ_1 and τ_2 are the same term.*

We now consider the notion of substitution:

Definition II.8.6 *If β and τ are terms and x is a variable, then $\beta(x \rightsquigarrow \tau)$ is the term which results from β by replacing all free occurrences of x by τ .*

Of course, one must verify that $\beta(x \rightsquigarrow \tau)$ really is a term, but this is easily done by induction on β .

For example, using the language of ordered rings $\mathcal{L}_{OR} = \{<, +, \cdot, -, 0, 1\}$ as in Section II.5, if β is the term (polynomial) $x \cdot y$ then $\beta(x \rightsquigarrow x + z)$ is $(x + z) \cdot y$. The parentheses are needed here, but in the official Polish, where β is $\cdot xy$, one literally replaces x by $+xz$ to see that $\beta(x \rightsquigarrow x + z)$ is $\cdot +xzy$. This substitution does the “right thing” in the semantics, when we compute the value of terms as in Section II.7. In the rationals,

$$\text{val}_{\mathfrak{A}}(\beta(x \rightsquigarrow x + z)) \begin{bmatrix} x & y & z \\ 1 & 2 & 5 \end{bmatrix} = \text{val}_{\mathfrak{A}}(\beta) \begin{bmatrix} x & y & z \\ 6 & 2 & 5 \end{bmatrix} = 12 \quad .$$

In the second expression, we changed the value of x to $6 = \text{val}_{\mathfrak{A}}(x + z)[1, 5]$. More generally, to evaluate $\beta(x \rightsquigarrow \tau)$ given an assignment σ , we change the value of $\sigma(x)$ to $\text{val}_{\mathfrak{A}}(\tau)[\sigma]$, and then evaluate β . Using the terminology of Definition II.7.7,

Lemma II.8.7 *If \mathfrak{A} is a structure for \mathcal{L} , and σ is both an assignment for β in A and an assignment for τ in A (see Definition II.7.3), and $a = \text{val}_{\mathfrak{A}}(\tau)[\sigma]$, then*

$$\text{val}_{\mathfrak{A}}(\beta(x \rightsquigarrow \tau))[\sigma] = \text{val}_{\mathfrak{A}}(\beta)[\sigma + (x/a)] \quad .$$

Proof. Induct on β . □

There is a similar discussion for formulas:

Definition II.8.8 *If φ is a formula, x is a variable, and τ is a term, then $\varphi(x \rightsquigarrow \tau)$ is the formula which results from φ by replacing all free occurrences of x by τ .*

Of course, one must verify that $\varphi(x \rightsquigarrow \tau)$ really is a formula, but this is easily done by induction on φ .

Roughly, $\varphi(x \rightsquigarrow \tau)$ says about τ what φ says about x . For example, again using \mathcal{L}_{OR} , if φ is the formula $\exists y(y \cdot y = x + z)$, asserting “ $x + z$ has a square root”, then $\varphi(x \rightsquigarrow 1)$ is $\exists y(y \cdot y = 1 + z)$ asserting “ $1 + z$ has a square root”. But $\varphi(y \rightsquigarrow 1)$ is φ ; since y is only a bound (dummy) variable, φ doesn’t say anything about y .

One must use some care when τ contains variables. For example, let φ be $\exists y(x < y)$. Then $\forall x\varphi$ is true in \mathbb{R} , so one would expect the universal closure of each $\varphi(x \rightsquigarrow \tau)$ to be true. For example, if τ is, respectively, 1 and $z + z$, then $\exists y(1 < y)$ and $\forall z\exists y(z + z < y)$ are both true in \mathbb{R} . However, if τ is $y + 1$, then $\varphi(x \rightsquigarrow \tau)$ is the sentence $\exists y(y + 1 < y)$, which is false in \mathbb{R} . The problem is that the variable y in τ got “captured” by the $\exists y$, changing its meaning. To make this problem precise,

Definition II.8.9 A term τ is free for x in a formula φ iff no free occurrence of x is inside the scope of a quantifier $\exists y$ or $\forall y$ where y is a variable which occurs in τ .

If the substitution is free, then it has the intended meaning, made formal by:

Lemma II.8.10 Assume that \mathfrak{A} is a structure for \mathcal{L} , φ is a formula of \mathcal{L} , τ is a term of \mathcal{L} , and σ is both an assignment for φ in A and an assignment for τ in A (see Definition II.7.3), and $a = \text{val}_{\mathfrak{A}}(\tau)[\sigma]$. Assume that τ is free for x in φ . Then

$$\mathfrak{A} \models \varphi(x \rightsquigarrow \tau)[\sigma] \quad \text{iff} \quad \mathfrak{A} \models \varphi[\sigma + (x/a)] \quad .$$

Proof. Induct on φ . The basis, where φ is atomic, uses Lemma II.8.7. Also note that if x is not free in φ , then $\varphi(x \rightsquigarrow \tau)$ is φ , and the value assigned by σ to x is irrelevant, so the lemma in this case is immediate and does not use the inductive hypothesis. The propositional cases for the induction are straightforward. Now, consider the the quantifier step, where φ is $\exists y\psi$ or $\forall y\psi$. Assume that x really has a free occurrence in φ (otherwise the result is immediate). Then the variables y and x must be distinct, and x has a free occurrence in ψ , so that y cannot occur in τ (since τ is free x in φ). The induction is now straightforward, using the definition of \models (see Definition II.7.8). This definition requires that we consider various $\sigma + (y/b)$, and we observe that $a = \text{val}_{\mathfrak{A}}(\tau)[\sigma] = \text{val}_{\mathfrak{A}}(\tau)[\sigma + (y/b)]$ because y does not occur in τ . \square

We frequently use the following simpler notation for substitution:

Notation II.8.11 $\varphi(\tau)$ abbreviates $\varphi(x \rightsquigarrow \tau)$ when it is clear from context that it is the variable x which is being replaced. To that end, one often refers to φ as “ $\varphi(x)$ ” during the discussion. Likewise if one refers to φ as “ $\varphi(x_1, \dots, x_n)$ ”, and τ_1, \dots, τ_n are terms, then $\varphi(\tau_1, \dots, \tau_n)$ denotes the formula obtained by simultaneously replacing each free occurrence of x_i in φ by τ_i .

As an example using this convention, we mention that Lemma II.8.10 implies:

Corollary II.8.12 If τ is free for x in $\varphi(x)$, then the formulas $\forall x\varphi(x) \rightarrow \varphi(\tau)$ and $\varphi(\tau) \rightarrow \exists x\varphi(x)$ are logically valid.

Proof. By Lemma II.8.10 and the definitions (II.8.1 and II.7.8) of “logically valid” and “ \models ”. \square

For another example, say we are talking about the real numbers, using the lexicon $\mathcal{L} = \{+, \cdot, -, 0, 1\}$, and we say, “let $\varphi(x, y)$ be $x + 1 = y$ ”. Then $\varphi(1 + 1, 1 + (1 + 1))$ is $(1 + 1) + 1 = 1 + (1 + 1)$, which is true in \mathbb{R} , while $\varphi(1 + (1 + 1), 1 + 1)$ is false in \mathbb{R} . The structure \mathfrak{A} here is $(\mathbb{R}; +, \cdot, -, 0, 1)$. Note that the two statements, $\mathfrak{A} \models \varphi(1+1, 1+(1+1))$ and $\mathfrak{A} \models \varphi[2, 3]$ say essentially the same thing, but are formally different. The first says that the sentence $\varphi(1 + 1, 1 + (1 + 1))$ is true in \mathfrak{A} , while the second says that the formula $x + 1 = y$ is true in \mathfrak{A} if we assign x value 2 and y value 3. The fact that these have the same meaning generalizes to:

Lemma II.8.13 *Assume that \mathfrak{A} is a structure for \mathcal{L} , $\varphi(x_1, \dots, x_n)$ is a formula of \mathcal{L} with no variables other than x_1, \dots, x_n free. and τ_1, \dots, τ_n are terms of \mathcal{L} with no free variables. Let $a_i = \text{val}_{\mathfrak{A}}(\tau_i)$. Then $\mathfrak{A} \models \varphi(\tau_1, \dots, \tau_n)$ iff $\mathfrak{A} \models \varphi[a_1, \dots, a_n]$.*

Proof.

$$\begin{aligned} \mathfrak{A} \models \varphi(\tau_1, \tau_2, \dots, \tau_n) &\text{ iff } \mathfrak{A} \models \varphi(x_1, \tau_2, \dots, \tau_n)[a_1] &\text{ iff } \mathfrak{A} \models \varphi(x_1, x_2, \dots, \tau_n)[a_1, a_2] \\ &\text{ iff } \dots \dots \text{ iff } \mathfrak{A} \models \varphi(x_1, x_2, \dots, x_n)[a_1, \dots, a_n] . \end{aligned}$$

Each of the n ‘iff’s uses Lemma II.8.10. □

So far, the lexicon \mathcal{L} has been fixed for each structure under discussion. But one frequently considers a fixed domain of discourse and varies the language. For example, we may consider the real numbers as a field, $\mathfrak{A} = (\mathbb{R}; +, \cdot, -, 0, 1)$, so that our language is $\mathcal{L}_1 = \{+, \cdot, -, 0, 1\}$. But we may also wish to consider \mathbb{R} just as an abelian group, using $\mathcal{L}_0 = \{+, -, 0\}$, and write $\mathfrak{A} \upharpoonright \mathcal{L}_0 = (\mathbb{R}; +, -, 0)$. Then we say that $\mathfrak{A} \upharpoonright \mathcal{L}_0$ is a *reduct* of \mathfrak{A} , and that \mathfrak{A} is an *expansion* of $\mathfrak{A} \upharpoonright \mathcal{L}_0$. In the terminology of category theory, we would say that we are describing *forgetful functor* from the category of fields to the category of abelian groups, since in the group $\mathfrak{A} \upharpoonright \mathcal{L}_0$, we *forget about* the product operation.

The terminology $\mathcal{L}_0 \subseteq \mathcal{L}_1$ implies that all the symbols have the same types in \mathcal{L}_0 and \mathcal{L}_1 ; we never, in one discussion, use the same name for symbols of different types. We give the general definition of reduct and expansion following the terminology of Definition II.7.1:

Definition II.8.14 *If $\mathcal{L}_0 \subseteq \mathcal{L}_1$ and $\mathfrak{A} = (A, \mathcal{I})$ is a structure for \mathcal{L}_1 then $\mathfrak{A} \upharpoonright \mathcal{L}_0$ denotes $(A, \mathcal{I} \upharpoonright \mathcal{L}_0)$. $\mathfrak{A} \upharpoonright \mathcal{L}_0$ is called a reduct of \mathfrak{A} and \mathfrak{A} is called an expansion of $\mathfrak{A} \upharpoonright \mathcal{L}_0$.*

Note that in Definition II.7.1, \mathcal{I} was really a function with domain \mathcal{L}_1 , and we are literally restricting this function to \mathcal{L}_0 .

Often, we start with an \mathcal{L}_0 structure and ask about its expansions. For example, if $(A; +, -, 0)$ is an abelian group, we might ask when it is the additive group of a field. This is really a (fairly easy) algebra question; using our model-theoretic terminology, we are asking whether $(A; +, -, 0)$ has an expansion of the form $(A; +, \cdot, -, 0, 1)$ which satisfies the field axioms.

The next lemma shows that notions such as $\Sigma \models \psi$ (Definition II.7.11) and $\text{Con}_{\models}(\Sigma)$ (Definition II.7.12) do not change if we expand the language. Thus, we did not mention \mathcal{L} explicitly and write something like $\mathfrak{A} \models_{\mathcal{L}} \psi$ or $\text{Con}_{\models, \mathcal{L}}(\Sigma)$.

Lemma II.8.15 *Suppose that Σ is a set of sentences of \mathcal{L}_0 and ψ is a sentence of \mathcal{L}_0 and suppose that $\mathcal{L}_0 \subseteq \mathcal{L}_1$. Then the following are equivalent:*

- α . $\mathfrak{A}_0 \models \psi$ for all \mathcal{L}_0 -structures \mathfrak{A}_0 such that $\mathfrak{A}_0 \models \Sigma$.
- β . $\mathfrak{A}_1 \models \psi$ for all \mathcal{L}_1 -structures \mathfrak{A}_1 such that $\mathfrak{A}_1 \models \Sigma$.

Also, the following are equivalent:

- a. There is an \mathcal{L}_0 -structure \mathfrak{A}_0 such that $\mathfrak{A}_0 \models \Sigma$.
- b. There is an \mathcal{L}_1 -structure \mathfrak{A}_1 such that $\mathfrak{A}_1 \models \Sigma$.

Proof. For (b) \rightarrow (a): If $\mathfrak{A}_1 \models \Sigma$ then also $\mathfrak{A}_1 \upharpoonright \mathcal{L}_0 \models \Sigma$, since the truth of \mathcal{L}_0 sentences is the same in \mathfrak{A}_1 and $\mathfrak{A}_1 \upharpoonright \mathcal{L}_0$.

For (a) \rightarrow (b): Let \mathfrak{A}_0 be any \mathcal{L}_0 -structure such that $\mathfrak{A}_0 \models \Sigma$. Expand \mathfrak{A}_0 arbitrarily to an \mathcal{L}_1 -structure \mathfrak{A}_1 . Then we still have $\mathfrak{A}_1 \models \Sigma$.

(α) \leftrightarrow (β) is similar. □

Remark II.8.16 The above proof is essentially trivial, but it does rely on the fact that structures are non-empty by Definition II.7.1. If we allowed the domain of discourse A to be empty, then all the basic definitions could still be made, but this lemma would fail. For example, if ψ is $\forall xp(x) \rightarrow \exists xp(x)$, then ψ is false in the empty set (where $\forall xp(x)$ is true and $\exists xp(x)$ is false), but ψ is true in every other structure. If $\mathcal{L}_0 = \{p\}$ and $\mathcal{L}_1 = \{p, c\}$ with c a constant symbol, we would have the somewhat pathological situation that $\{\neg\psi\}$ would be consistent as an \mathcal{L}_0 -sentence but not as an \mathcal{L}_1 -sentence. In the proof of (a) \rightarrow (b), there would be no way to expand the empty structure to an \mathcal{L}_1 -structure because constant symbols must be interpreted as elements of A . This sort of pathology explains why the universe is always assumed to be non-empty in model theory.

In reduct/expansion, we fix A and decrease/increase \mathcal{L} . This should not be confused with submodel/extension, where we fix \mathcal{L} and decrease/increase A . The notion of submodel generalizes the notions of subgroup, subring, etc., from algebra:

Definition II.8.17 Suppose that $\mathfrak{A} = (A, \mathcal{I})$ and $\mathfrak{B} = (B, \mathcal{J})$ are structures for \mathcal{L} . Then $\mathfrak{A} \subseteq \mathfrak{B}$ means that $A \subseteq B$ and the functions and predicates of \mathfrak{A} are the restrictions of the corresponding functions and predicates of \mathfrak{B} . Specifically:

- ☞ If $f \in \mathcal{F}_n$ with $n > 0$, then $f_{\mathfrak{A}} = f_{\mathfrak{B}} \upharpoonright A^n$.
- ☞ If $p \in \mathcal{P}_n$ with $n > 0$, then $p_{\mathfrak{A}} = p_{\mathfrak{B}} \cap A^n$.
- ☞ If $c \in \mathcal{F}_0$, then $c_{\mathfrak{A}} = c_{\mathfrak{B}}$.
- ☞ If $p \in \mathcal{P}_0$, then $p_{\mathfrak{A}} = p_{\mathfrak{B}} \in 2 = \{0, 1\} = \{F, T\}$.

\mathfrak{A} is called a submodel of \mathfrak{B} and \mathfrak{B} is called an extension of \mathfrak{A} .

In the case of constants and functions, observe that $c_{\mathfrak{A}}$ must be an element of A and $f_{\mathfrak{A}}$ must map into A . So, if we start with \mathfrak{B} and an arbitrary non-empty $A \subseteq B$, it is not true in general that A can be made into a submodel of \mathfrak{B} . For example, suppose $\mathfrak{B} = (B; \cdot, i, 1)$ is a group, where, as in Section II.5, we are taking the language of group theory to be $\mathcal{L} = \{\cdot, i, 1\}$. If $A \subseteq B$, it cannot be made into a submodel of \mathfrak{B} unless it

is closed under product and inverse and contains 1, that is, unless it is a subgroup. Also note that which subsets form submodels changes if we go to reducts or expansions. For example, if we reduct to the language $\mathcal{L}_0 = \{\cdot\}$, we can still express the group axioms (as in Section 0.4), but submodels of $(B; \cdot)$ are subsemigroups; that is, closed under product but not necessarily inverse.

There is a notion of isomorphism between groups or rings; this generalizes easily to arbitrary structures.

Definition II.8.18 Suppose that $\mathfrak{A} = (A, \mathcal{I})$ and $\mathfrak{B} = (B, \mathcal{J})$ are structures for the same language \mathcal{L} . Φ is an isomorphism from \mathfrak{A} onto \mathfrak{B} iff $\Phi : A \xrightarrow[\text{onto}]{1-1} B$ and Φ preserves the structure. Specifically:

- ☞ If $f \in \mathcal{F}_n$ with $n > 0$, then $f_{\mathfrak{B}}(\Phi(a_1), \dots, \Phi(a_n)) = \Phi(f_{\mathfrak{A}}(a_1, \dots, a_n))$.
- ☞ If $p \in \mathcal{P}_n$ with $n > 0$, then $(\Phi(a_1), \dots, \Phi(a_n)) \in p_{\mathfrak{B}}$ iff $(a_1, \dots, a_n) \in p_{\mathfrak{A}}$.
- ☞ If $c \in \mathcal{F}_0$, then $c_{\mathfrak{B}} = \Phi(c_{\mathfrak{A}})$.
- ☞ If $p \in \mathcal{P}_0$, then $p_{\mathfrak{B}} = p_{\mathfrak{A}} \in 2 = \{0, 1\} = \{F, T\}$.

\mathfrak{A} and \mathfrak{B} are isomorphic ($\mathfrak{A} \cong \mathfrak{B}$) iff there exists an isomorphism from \mathfrak{A} onto \mathfrak{B} .

This definition also generalizes Definition I.7.14, which was given for the special case of ordered structures.

A set of axioms Σ is *complete* iff it decides all possible statements:

Definition II.8.19 If Σ is a set of sentences of \mathcal{L} , then Σ is complete (with respect to \mathcal{L}) iff Σ is semantically consistent and for all sentences φ of \mathcal{L} , either $\Sigma \models \varphi$ or $\Sigma \models \neg\varphi$.

If we just say “ Σ is complete”, it is understood that \mathcal{L} is the set of symbols actually used in Σ . Σ will usually *not* be complete with respect to a larger \mathcal{L} :

Exercise II.8.20 Suppose that Σ is a set of sentences of \mathcal{L} and $\mathcal{L}' \supsetneq \mathcal{L}$, with $\mathcal{L}' \setminus \mathcal{L}$ containing at least one predicate symbol. Then Σ cannot be complete with respect to \mathcal{L}' .

Hint. Consider φ of the form $\exists x_1, \dots, x_n p(x_1, \dots, x_n)$. □

A (perhaps artificial) example of a complete Σ is the theory of a given structure.

Definition II.8.21 If \mathfrak{A} is a structure for \mathcal{L} , then the theory of \mathfrak{A} , $\text{Th}(\mathfrak{A})$ is the set of all \mathcal{L} -sentences φ such that $\mathfrak{A} \models \varphi$.

Lemma II.8.22 $\text{Th}(\mathfrak{A})$ is complete (with respect to \mathcal{L}).

Proof. $\text{Th}(\mathfrak{A})$ is semantically consistent because $\mathfrak{A} \models \text{Th}(\mathfrak{A})$, and for all sentences φ of \mathcal{L} , either $\varphi \in \Sigma$ or $(\neg\varphi) \in \Sigma$. \square

There are many natural examples of complete theories in algebra. We shall describe a few in Section II.13. Further examples may be found in model theory texts, such as [5, 24].

The following example from algebra illustrates one additional point about formalizing algebraic theories.

Example II.8.23 Let $\mathcal{L} = \{0, 1, +, \cdot, -, i\}$, where “ $-$ ” denotes the unary additive inverse and “ i ” denotes the unary multiplicative inverse (or reciprocal). Let Σ in \mathcal{L} be the axioms for fields, expressed by:

1. The axioms for groups, written in $+, 0, -$ (see $\gamma_1, \gamma_{2,1}, \gamma_{2,2}$ in Section II.5).
2. The associative and identity laws, written in $\cdot, 1$ (see $\gamma_1, \gamma_{2,1}$ in Section II.5).
3. The commutative laws: $\forall x, y [x \cdot y = y \cdot x]$ and $\forall x, y [x + y = y + x]$.
4. The distributive law: $\forall x, y, z [x \cdot (y + z) = (x \cdot y) + (x \cdot z)]$.
5. The multiplicative inverse law: $\forall x [x \neq 0 \rightarrow x \cdot i(x) = 1]$.
6. $i(0) = 0$.
7. $0 \neq 1$.

Axiom (5) states the existence of a reciprocal for every *non-zero* element. Informally, in algebra, we say “ $1/0$ is undefined”. Formally, since our model theory does not allow for partially defined function symbols, $i(x)$ is defined for all x , and we just assert that it denotes the reciprocal of x when the reciprocal exists (i.e., when $x \neq 0$). The value of $i(0)$ is “irrelevant”, but we include axiom (6) specifying its value so that the usual notion of field isomorphism in algebra corresponds to the notion of isomorphism in model theory, where Definition II.8.18 requires $\Phi(i_{\mathfrak{A}}(0_{\mathfrak{A}})) = i_{\mathfrak{B}}(0_{\mathfrak{B}})$. If we dropped axiom (6), then there would be three non-isomorphic fields of order three; depending on whether $i(0)$ is 0, 1, or 2. Axiom (7) disallows the “trivial” 1-element field.

II.9 Tautologies

Informally, a *propositional tautology* (or, just *tautology*) is a formula whose logical validity is apparent just from the meaning of the propositional connectives, without reference to the meaning of $=, \forall, \exists$. For example, $p(x) \rightarrow p(x)$ is a tautology, whereas $\forall x p(x) \rightarrow \forall y p(y)$ and $x = x$ are not, since you have to understand the meaning of \forall and $=$, respectively, to see that they are logically valid.

Definition II.9.1 A formula is basic iff (in its Polish notation) it does not begin with a propositional connective.

For example, $\forall xp(x) \rightarrow \forall yp(y)$ (i.e. $\rightarrow \forall xp(x)\forall yp(y)$) is not basic, but it is an implication between the two basic formulas $\forall xp(x)$ and $\forall yp(y)$. In the definition of “tautology”, we consider these basic formulas as distinct un-analyzed atoms. Note that every formula is obtained from basic formulas by using propositional connectives.

Definition II.9.2 *A truth assignment for \mathcal{L} is a function v from the set of basic formulas of \mathcal{L} into $\{0, 1\} = \{F, T\}$. Given such a v , we define (recursively) $\bar{v}(\varphi) \in \{F, T\}$ as follows:*

1. $\bar{v}(\neg\varphi) = 1 - \bar{v}(\varphi)$.
2. $\bar{v}(\wedge\varphi\psi)$, $\bar{v}(\vee\varphi\psi)$, $\bar{v}(\rightarrow\varphi\psi)$, and $\bar{v}(\leftrightarrow\varphi\psi)$, are obtained from $\bar{v}(\varphi)$ and $\bar{v}(\psi)$ using the truth tables (Table 1, page 4) for $\wedge, \vee, \rightarrow, \leftrightarrow$.

φ is a propositional tautology iff $\bar{v}(\varphi) = T$ for all truth assignments v .

There is a similarity between this definition and Definition II.7.8 (of \models), but here we are only studying the meaning of the propositional connectives. To test whether φ is a tautology, you just check all possible assignments of T or F to the basic formulas out of which φ is built. For example, if φ is $\forall xp(x) \rightarrow \forall yp(y)$, then one such v has $v(\forall xp(x)) = T$ and $v(\forall yp(y)) = F$; this is allowed because $\forall xp(x)$ and $\forall yp(y)$ are distinct formulas (even though they are logically equivalent); this v makes $\bar{v}(\varphi) = F$, so that φ is not a tautology. However, $p(x) \rightarrow p(x)$ is built out of the one basic formula $p(x)$, which v may make either T or F , but in either case $\bar{v}(p(x) \rightarrow p(x)) = T$, so that $p(x) \rightarrow p(x)$ is a tautology.

Comparing Definitions II.9.2 and II.7.8 and the definition (II.8.1) of logical validity, we see:

Exercise II.9.3 *Every propositional tautology is logically valid.*

II.10 Formal Proofs

We now give a presentation of formal proof theory, as we promised in Sections 0.4 and II.2. As mentioned in Section II.2, we are striving for a system which is easy to define and analyze mathematically, not one which is easy to use in a practical settings. Some remarks about “practical” proof theories and references to the literature are in Section II.17.

In this proof theory, we have one rule of inference:

$$\text{MODUS PONENS: } \frac{\varphi \quad \varphi \rightarrow \psi}{\psi}$$

Informally, this means that if we have proved both φ and $\varphi \rightarrow \psi$, then we can conclude ψ . Formally, Modus Ponens is embedded in our definition (II.10.3) of “formal proof”. First, we single out some “obviously valid” statements and call them “logical axioms”:

Definition II.10.1 A logical axiom of \mathcal{L} is any sentence of \mathcal{L} which is a universal closure (see Definition II.5.6) of a formula of one the types listed below. Here, x, y, z , possibly with subscripts, denote arbitrary variables.

1. propositional tautologies.
2. $\varphi \rightarrow \forall x\varphi$, where x is not free in φ .
3. $\forall x(\varphi \rightarrow \psi) \rightarrow (\forall x\varphi \rightarrow \forall x\psi)$.
4. $\forall x\varphi \rightarrow \varphi(x \rightsquigarrow \tau)$, where τ is any term which is free for x in φ .
5. $\varphi(x \rightsquigarrow \tau) \rightarrow \exists x\varphi$, where τ is any term which is free for x in φ .
6. $\forall x\neg\varphi \leftrightarrow \neg\exists x\varphi$.
7. $x = x$.
8. $x = y \leftrightarrow y = x$.
9. $x = y \wedge y = z \rightarrow x = z$.
10. $x_1 = y_1 \wedge \dots \wedge x_n = y_n \rightarrow (f(x_1 \dots x_n) = f(y_1 \dots y_n))$, whenever $n > 0$ and f is an n -place function symbol of \mathcal{L} .
11. $x_1 = y_1 \wedge \dots \wedge x_n = y_n \rightarrow (p(x_1 \dots x_n) \leftrightarrow p(y_1 \dots y_n))$, whenever $n > 0$ and p is an n -place predicate symbol of \mathcal{L} .

Exercise II.10.2 All the logical axioms are logically valid.

Hint. The hard ones have already been done. For the tautologies, see Exercise II.9.3. For axioms of types (4),(5), see Corollary II.8.12. \square

Definition II.10.3 If Σ is a set of sentences of \mathcal{L} , then a formal proof from Σ is a finite, non-empty sequence of sentences of \mathcal{L} , $\varphi_0, \dots, \varphi_n$, such that for each i , either $\varphi_i \in \Sigma$ or φ_i is a logical axiom or for some $j, k < i$, φ_i follows from φ_j, φ_k by Modus Ponens (that is, φ_k is $(\varphi_j \rightarrow \varphi_i)$).

Definition II.10.4 If Σ is a set of sentences of \mathcal{L} , and φ is a sentence of \mathcal{L} , then $\Sigma \vdash_{\mathcal{L}} \varphi$ iff there is a formal proof from Σ whose last sentence is φ .

Lemma II.10.5 (Soundness) If $\Sigma \vdash_{\mathcal{L}} \varphi$ then $\Sigma \models \varphi$.

Proof. Assume that $\Sigma \vdash_{\mathcal{L}} \varphi$ and $\mathfrak{A} \models \Sigma$. we need to show that $\mathfrak{A} \models \varphi$.

Let $\varphi_0, \dots, \varphi_n$ be a formal proof of φ from Σ ; then φ_n is φ . By induction on i , show that $\mathfrak{A} \models \varphi_i$. There are three cases. If $\varphi_i \in \Sigma$ use $\mathfrak{A} \models \Sigma$. If φ_i is a logical axiom, use Exercise II.10.2. These two cases don't use the inductive hypothesis. If Modus Ponens is used, then note that $\mathfrak{A} \models \varphi_i$ follows from $\mathfrak{A} \models \varphi_j \rightarrow \varphi_i$ and $\mathfrak{A} \models \varphi_j$. \square

Note that our definition of formal proof is very simple, *except* for the list of logical axioms. The choice of exactly which statements to put on this list is a bit arbitrary, and differs in different texts. There are only three *important* things about this list:

1. Every logical axiom is logically valid, so that *Soundness* is true.
2. We have listed enough logical axioms to verify *Completeness*.
3. When \mathcal{L} is finite, the set of logical axioms is *decidable*.

Soundness is Lemma II.10.5 above. *Completeness* (see Theorem II.12.1), the converse statement, asserts that if $\Sigma \models \varphi$ then $\Sigma \vdash_{\mathcal{L}} \varphi$. When we quoted the Completeness Theorem before, in Sections 0.4 and II.7, we said that $\Sigma \models \varphi$ iff $\Sigma \vdash_{\mathcal{L}} \varphi$, but as we have just seen, the Soundness direction of this “iff” is very easy. The other direction requires much more work.

When \mathcal{L} is finite, we may view syntactic objects as possible inputs into a computer. By (3), a computer can check whether or not a sequence of formulas is a formal proof, and the computer can in principle generate its own formal proofs. In practice, computer programs which manipulate formal proofs use proof systems which differ significantly from the one described here (see Section II.17).

It might seem more elegant to define the logical axioms to be exactly the set of logically valid sentences. That would simplify the definition, make (1) obvious, and would make the proof of (2) somewhat easier, but by Church’s theorem (see Chapter III) we would lose (3).

We are writing “ $\Sigma \vdash_{\mathcal{L}} \varphi$ ” rather than “ $\Sigma \vdash \varphi$ ” because *conceivably* this notion could depend on \mathcal{L} . Suppose that Σ and φ are in \mathcal{L} , and $\mathcal{L}' \supset \mathcal{L}$. Perhaps $\Sigma \vdash_{\mathcal{L}'} \varphi$, and the formal proof uses symbols of $\mathcal{L}' \setminus \mathcal{L}$. It is true, but not immediately obvious, that we can always get another formal proof just using symbols of \mathcal{L} , so that $\Sigma \vdash_{\mathcal{L}'} \varphi$ iff $\Sigma \vdash_{\mathcal{L}} \varphi$. A direct proof of this (see Exercise II.11.13) is bit tedious. Our official proof of this (see Lemma II.12.21) will be from the Completeness Theorem, since the notion “ $\Sigma \models \varphi$ ” doesn’t depend on \mathcal{L} (see Lemma II.8.15).

As we have remarked when listing the axioms of set theory in Section I.2, we are following the usual convention in modern algebra and logic that basic facts about $=$ are logical facts, and need not be stated when axiomatizing a theory. For example, $\emptyset \vdash \forall x(x = x)$, since this is a logical axiom of type 7. Also the converse to Extensionality is a logical fact; $\emptyset \vdash \forall x, y(x = y \rightarrow \forall z(z \in x \leftrightarrow z \in y))$; see Exercise II.11.14. Also, $\emptyset \vdash \exists x(x = x)$, the universe is non-empty; see Exercise II.10.6; here, we listed this explicitly as an axiom of set theory to avoid possible confusion, since in algebra one does sometimes allow an empty structure (e.g., the empty semigroup).

Note that formal proofs only involve sentences, not arbitrary formulas. In informal mathematical reasoning, when you see a free variable (i.e., a letter) in an argument, it is left to the reader to decide from context whether it is universally or existentially quantified.

We conclude this section with two examples of formal proofs. To show $p \wedge q \vdash p$:

0. $p \wedge q \rightarrow p$ tautology
1. $p \wedge q$ given
2. p 0, 1, modus ponens

Some remarks. Formally, we are showing $\Sigma \vdash_{\mathcal{L}} \varphi$ where $\Sigma = \{p \wedge q\}$ and \mathcal{L} contains at least the proposition letters p and q . Note that following Definition II.10.3, the formal proof itself is just the sequence of three sentences, $(p \wedge q \rightarrow p, p \wedge q, p)$, not the commentary. Given any sequence of sentences, $(\varphi_0, \dots, \varphi_n)$, without any commentary, it is decidable whether it forms a formal proof, since we may, for each φ_i , check all possible justifications for φ_i being legitimate.

It is often tedious to write out formal proofs of trivial things. Figure II.1 shows that $\forall x[p(x) \wedge q(x)] \vdash \forall y p(y)$. Deriving $\forall x p(x)$ only requires 5 lines; we use a type 3 axiom to do the modus ponens step from the previous proof inside a universal quantifier. An additional 6 lines is required to change the “ y ” to an “ x ”.

Figure II.1: $\forall x[p(x) \wedge q(x)] \vdash \forall y p(y)$

0.	$\forall x[p(x) \wedge q(x) \rightarrow p(x)]$	tautology
1.	$\forall x[p(x) \wedge q(x) \rightarrow p(x)] \rightarrow (\forall x[p(x) \wedge q(x)] \rightarrow \forall x p(x))$	type 3 axiom
2.	$\forall x[p(x) \wedge q(x)] \rightarrow \forall x p(x)$	1, 0, modus ponens
3.	$\forall x[p(x) \wedge q(x)]$	given
4.	$\forall x p(x)$	2, 3, modus ponens
5.	$\forall y[\forall x p(x) \rightarrow p(y)]$	type 4 axiom
6.	$\forall y[\forall x p(x) \rightarrow p(y)] \rightarrow (\forall y \forall x p(x) \rightarrow \forall y p(y))$	type 3 axiom
7.	$\forall y \forall x p(x) \rightarrow \forall y p(y)$	6, 5, modus ponens
8.	$\forall x p(x) \rightarrow \forall y \forall x p(x)$	type 2 axiom
9.	$\forall y \forall x p(x)$	8, 4, modus ponens
10.	$\forall y p(y)$	7, 9, modus ponens

Lines 0 and 5 illustrate the fact that the logical axioms are actually *closures* of the formulas listed in Definition II.10.1. Informally, one would prove $\forall y p(y)$ from $\forall x[p(x) \wedge q(x)]$ trivially by:

Assume $\forall x[p(x) \wedge q(x)]$. Fix any object c . Then $p(c) \wedge q(c)$ holds, so $p(c)$ follows tautologically. Since c was arbitrary, we have $\forall y p(y)$.

In Section II.11, we shall introduce some proof rules which will allow one to construct a formal proof directly from this informal proof.

Exercise II.10.6 Write out a formal proof of $\exists x(x = x)$ from \emptyset . As in the above examples, you may use the standard abbreviations for the sentences occurring in the proof, but don't skip steps in the proof.

Hint. Observe that $\forall x(x = x \rightarrow \exists x(x = x))$ and $\forall x \exists x(x = x) \rightarrow \exists x(x = x)$ are logical axioms of types 5 and 4 respectively. □

II.11 Some Strategies for Constructing Proofs

As we have indicated before, our proof theory is really not suited to the task of formalizing large bodies of mathematics; see Section II.17 for a description of some proof theories which are better suited. However, we shall, in this section, establish a few general principles which show how informal mathematical arguments can be replicated in the formal proof theory. The results here will be useful later as lemmas in the proof of the Competeness Theorem.

First, we consider the informal rule that to prove $\varphi \rightarrow \psi$, we may assume that φ is true and derive ψ . This becomes:

Lemma II.11.1 (The Deduction Theorem) $\Sigma \vdash_{\mathcal{L}} \varphi \rightarrow \psi$ iff $\Sigma \cup \{\varphi\} \vdash_{\mathcal{L}} \psi$.

Proof. For \rightarrow , just use Modus Ponens. That is, given a proof of $\varphi \rightarrow \psi$ from Σ , we may add two lines to get a proof of $\varphi \rightarrow \psi$ from $\Sigma \cup \{\varphi\}$: First write down φ , and then apply Modus Ponens to write down ψ .

For \leftarrow , assume that ψ_0, \dots, ψ_n is a formal proof of ψ from $\Sigma \cup \{\varphi\}$; so ψ_n is ψ . We shall prove by induction on i that $\Sigma \vdash_{\mathcal{L}} \varphi \rightarrow \psi_i$. So, assume, inductively, that $\Sigma \vdash_{\mathcal{L}} \varphi \rightarrow \psi_j$ for all $j < i$. There are now three cases; the first two do not use the inductive hypothesis.

Case 1. ψ_i is either a logical axiom or in Σ . Then in a proof from Σ , we can just write ψ_i down, so we have a 3-line proof of $\varphi \rightarrow \psi_i$ from Σ :

0. ψ_i
1. $\psi_i \rightarrow (\varphi \rightarrow \psi_i)$ tautology
2. $\varphi \rightarrow \psi_i$ 1, 0, modus ponens

Case 2. ψ_i is φ , so $\varphi \rightarrow \psi_i$ is a tautology, so it has a 1-line proof.

Case 3. For some $j, k < i$, ψ_i follows from ψ_j, ψ_k by Modus Ponens, so ψ_k is $(\psi_j \rightarrow \psi_i)$. Then

0. $\Sigma \vdash_{\mathcal{L}} \varphi \rightarrow \psi_j$ induction
1. $\Sigma \vdash_{\mathcal{L}} \varphi \rightarrow (\psi_j \rightarrow \psi_i)$ induction
2. $\Sigma \vdash_{\mathcal{L}} (\varphi \rightarrow \psi_j) \rightarrow [[\varphi \rightarrow (\psi_j \rightarrow \psi_i)] \rightarrow (\varphi \rightarrow \psi_i)]$ tautology
3. $\Sigma \vdash_{\mathcal{L}} [\varphi \rightarrow (\psi_j \rightarrow \psi_i)] \rightarrow (\varphi \rightarrow \psi_i)$ 2, 0, modus ponens
4. $\Sigma \vdash_{\mathcal{L}} \varphi \rightarrow \psi_i$ 3, 1, modus ponens

Note that in Case 1, we have explicitly displayed a formal proof (after one strips off the comments and line numbers), whereas in Case 3, we are really showing how to construct a formal proof of $\varphi \rightarrow \psi_i$ from formal proofs of $\varphi \rightarrow \psi_j$ and $\varphi \rightarrow \psi_k$. \square

Next, we consider *proof by contradiction*, or *reductio ad absurdum*. That is, to prove φ , we may assume that φ is false and derive a contradiction. Before stating this formally, we must say what “contradiction” means in our formal proof theory.

Definition II.11.2 *If Σ is a set of sentences of \mathcal{L} then Σ is syntactically inconsistent ($\neg\text{Con}_{+, \mathcal{L}}(\Sigma)$) iff there is some sentence φ of \mathcal{L} such that $\Sigma \vdash_{\mathcal{L}} \varphi$ and $\Sigma \vdash_{\mathcal{L}} \neg\varphi$. “consistent” means “not inconsistent”.*

This definition should be compared with Definition II.7.12, which defined a semantic notion of consistency. We shall soon prove the Completeness Theorem (Theorem II.12.1), which will imply that $\text{Con}_{+, \mathcal{L}}(\Sigma)$ iff $\text{Con}_{\models}(\Sigma)$ whenever \mathcal{L} is any lexicon large enough to include all the symbols of Σ . After that, we drop the subscripts and just write $\text{Con}(\Sigma)$. Right now, we point out the equivalence of a minor variant of “consistent”.

Lemma II.11.3 *If Σ is a set of sentences of \mathcal{L} , then the following are equivalent:*

- a. $\neg\text{Con}_{+, \mathcal{L}}(\Sigma)$.
- b. $\Sigma \vdash_{\mathcal{L}} \psi$ for all sentences ψ of \mathcal{L}

Proof. (b) \rightarrow (a) is trivial. For (a) \rightarrow (b), use the fact that $\varphi \rightarrow (\neg\varphi \rightarrow \psi)$ is a tautology, and apply Modus Ponens twice. \square

Cantor probably felt that his set theory was only mildly inconsistent, since the paradoxes it derived (see page 50) did not involve “ordinary” sets. But, in formal logic, there is no notion of “mildly inconsistent”; once we have derived an inconsistency, we can prove everything.

The next lemma is the proof theory version of reductio ad absurdum (see Lemma II.7.13):

Lemma II.11.4 (Proof by Contradiction) *If Σ is a set of sentences of \mathcal{L} and φ is a sentence of \mathcal{L} , then*

1. $\Sigma \vdash_{\mathcal{L}} \varphi$ iff $\neg\text{Con}_{+, \mathcal{L}}(\Sigma \cup \{\neg\varphi\})$.
2. $\Sigma \vdash_{\mathcal{L}} \neg\varphi$ iff $\neg\text{Con}_{+, \mathcal{L}}(\Sigma \cup \{\varphi\})$.

Proof. For (1), \rightarrow is immediate from the definition of “ $\neg\text{Con}$ ”. For \leftarrow , we have $\Sigma \cup \{\neg\varphi\} \vdash_{\mathcal{L}} \varphi$ (applying Lemma II.11.3), so that $\Sigma \vdash_{\mathcal{L}} \neg\varphi \rightarrow \varphi$ by the Deduction Theorem (Lemma II.11.1). But $(\neg\varphi \rightarrow \varphi) \rightarrow \varphi$ is a tautology, so $\Sigma \vdash_{\mathcal{L}} \varphi$ by Modus Ponens.

(2) is similar and is left as an exercise. \square

In this last argument, and a few times earlier, we have written down a tautology and then applied Modus Ponens. This can be generalized to a statement about tautological reasoning (Lemma II.11.6).

Definition II.11.5 ψ follows tautologically from $\varphi_1, \dots, \varphi_n$ iff $(\varphi_1 \wedge \dots \wedge \varphi_n) \rightarrow \psi$ is a propositional tautology.

Lemma II.11.6 (Tautological Reasoning) *If $\psi, \varphi_1, \dots, \varphi_n$ are sentence of \mathcal{L} and ψ follows tautologically from $\varphi_1, \dots, \varphi_n$, then $\{\varphi_1, \dots, \varphi_n\} \vdash_{\mathcal{L}} \psi$.*

Proof. Note that $\varphi_1 \rightarrow (\varphi_2 \rightarrow (\dots \rightarrow (\varphi_n \rightarrow \psi) \dots))$ is a tautology, and use Modus Ponens n times. \square

This is often used in conjunction with the following fact, which is easily demonstrated by pasting together formal proofs:

Lemma II.11.7 (Transitivity of \vdash) *If $\{\varphi_1, \dots, \varphi_n\} \vdash_{\mathcal{L}} \psi$, and $\Sigma \vdash_{\mathcal{L}} \varphi_i$ for $i = 1, \dots, n$, then $\Sigma \vdash_{\mathcal{L}} \psi$.*

So, for example, in proving Case 3 of Lemma II.11.1, we could have just said that from $\Sigma \vdash_{\mathcal{L}} \varphi \rightarrow \psi_j$ and $\Sigma \vdash_{\mathcal{L}} \varphi \rightarrow (\psi_j \rightarrow \psi_i)$, we get $\Sigma \vdash_{\mathcal{L}} \varphi \rightarrow \psi_i$ because $\varphi \rightarrow \psi_i$ follows tautologically from $\varphi \rightarrow (\psi_j \rightarrow \psi_i)$ and $\varphi \rightarrow \psi_j$ (and applying Lemma II.11.6 and II.11.7). Note that every use of Modus Ponens can be subsumed under Lemma II.11.6, since ψ follows tautologically from φ and $\varphi \rightarrow \psi$.

Next, we come to some rules for handling quantifiers. In informal mathematical reasoning, quantifiers are often not written explicitly, but they are handled implicitly by the informal analog of the following:

Lemma II.11.8 (Quantifier Rules)

$$\begin{array}{ll} \text{UI:} & \forall x\varphi(x) \vdash_{\mathcal{L}} \varphi(\tau) \\ \text{EI:} & \frac{\Sigma \cup \{\varphi(c)\} \vdash_{\mathcal{L}'} \psi}{\Sigma \cup \{\exists x\varphi(x)\} \vdash_{\mathcal{L}} \psi} \\ \text{UG:} & \frac{\Sigma \vdash_{\mathcal{L}'} \varphi(c)}{\Sigma \vdash_{\mathcal{L}} \forall x\varphi(x)} \\ \text{EG:} & \varphi(\tau) \vdash_{\mathcal{L}} \exists x\varphi(x) \end{array}$$

Here, Σ is a set of sentences of \mathcal{L} , and $\varphi(x)$ is a formula of \mathcal{L} with at most the variable x free. In UI and EG, τ is a term of \mathcal{L} with no variables, so that $\varphi(\tau)$ is a sentence. In UG and EI, c is a constant symbol which is not in \mathcal{L} , and $\mathcal{L}' = \mathcal{L} \cup \{c\}$. In EI, ψ is a sentence of \mathcal{L} .

Some explanation, before the proof: “U” stands for “Universal” and “E” stands for “Existential”. “I” stands for “Instantiation” and “G” stands for “Generalization”. The UI (Universal Instantiation) rule corresponds to the informal step that if a universal statement $\forall x\varphi(x)$ is true, we can conclude a specific instance $\varphi(\tau)$. Likewise, EG (Existential Generalization) corresponds to the informal step that if we can prove φ holds of a specific τ , then we know $\exists x\varphi(x)$. Informally, UI and EG are “obviously correct”, since $\forall x\varphi(x) \rightarrow \varphi(\tau)$ and $\varphi(\tau) \rightarrow \exists x\varphi(x)$ are logically valid (see Corollary II.8.12).

The horizontal line stands for an “if \dots then \dots ”, so UG (Universal Generalization) asserts that if $\Sigma \vdash_{\mathcal{L}'} \varphi(c)$ then $\Sigma \vdash_{\mathcal{L}} \forall x\varphi(x)$. UG is trickier than UI, since from an instance $\varphi(c)$ one cannot usually generalize to $\forall x\varphi(x)$. UG corresponds to the informal

words “but c was arbitrary”. Say we’re talking about real numbers and we want to prove $\forall x\varphi(x)$, where $\varphi(x)$ is $\exists y(y^3 = x)$. We would say: let c be an arbitrary real. Apply the Intermediate Value Theorem to prove that $y^3 - c$ has a root, concluding $\varphi(c)$. But c was arbitrary, so $\forall x\varphi(x)$. Because everyone accepts this informal means of proof, you rarely see explicit quantifiers written in an elementary analysis text, although beginning students are sometimes confused about which symbols denote “arbitrary” numbers. Note that in the official statement of the UG rule, we assumed that Σ was in \mathcal{L} , so that the axioms Σ do not mention the constant c . In our informal example, Σ denotes basic facts about the real numbers which we are assuming to be known already. Most likely, Σ uses the constant 0 and $\Sigma \vdash 0 + 0 = 0$, but we cannot conclude from this that $\Sigma \vdash \forall x(x + x = x)$, since the constant 0 is explicitly mentioned in Σ .

The EI (Existential Instantiation) rule corresponds to the informal words “fix c such that $\dots\dots$ ”. To illustrate all four quantifier rules, say we are proving $\exists x\forall y p(x, y) \rightarrow \forall y\exists x p(x, y)$. Informally, assume that $\exists x\forall y p(x, y)$ is true and fix (EI) c such that $\forall y p(c, y)$. Consider any object d . Then $p(c, d)$ (UI), so $\exists x p(x, d)$ (EG). But d was arbitrary (UG), so $\forall y\exists x p(x, y)$. When writing out the steps more formally, the order of application of the rules gets permuted:

- | | |
|---|----------------------|
| 0. $p(c, d) \vdash_{\mathcal{L}''} p(c, d)$ | tautology |
| 1. $p(c, d) \vdash_{\mathcal{L}''} \exists x p(x, d)$ | 0, EG |
| 2. $\forall y p(c, y) \vdash_{\mathcal{L}''} \exists x p(x, d)$ | 1, UI |
| 3. $\forall y p(c, y) \vdash_{\mathcal{L}'} \forall y\exists x p(x, y)$ | 2, UG |
| 4. $\exists x\forall y p(x, y) \vdash_{\mathcal{L}} \forall y\exists x p(x, y)$ | 3, EI |
| 5. $\emptyset \vdash_{\mathcal{L}} \exists x\forall y p(x, y) \rightarrow \forall y\exists x p(x, y)$ | 4, Deduction Theorem |

Here, $\mathcal{L} = \{p\}$, $\mathcal{L}' = \{p, c\}$, $\mathcal{L}'' = \{p, c, d\}$. In step (2), we are implicitly using the transitivity of \vdash . In step (0), we could be quoting Lemma II.11.6, but it is also trivial by the definition of \vdash . Lines (0–5) do not constitute a formal proof, but rather a demonstration that there is a formal proof. Our justifications for the quantifier rules (see below) and for the Deduction Theorem (see above) are all constructive, in that they tell you how to write down an explicit formal proof of $\exists x\forall y p(x, y) \rightarrow \forall y\exists x p(x, y)$, although it will have many more than 6 lines and will look a bit ugly.

Exercise II.11.9 Write a formal proof of $\exists x\forall y p(x, y) \rightarrow \forall y\exists x p(x, y)$ from \emptyset .

Natural deduction systems (see Section II.17) let one write something like (0–5) as part of the formal proof theory. This has the advantage of making it easier to display real formal proofs in the system. It has the disadvantage of having a more complex definition of what a formal proof is.

Here’s another example. We have $\forall x\neg\varphi(x) \leftrightarrow \neg\exists x\varphi(x)$ as a logical axiom (of type 6). Similar quantifier manipulation can now be derived from this using Lemma II.11.8:

Example II.11.10 $\emptyset \vdash_{\mathcal{L}} \neg\forall x\psi(x) \rightarrow \exists x\neg\psi(x)$

Proof.

- | | |
|---|----------------------|
| 1. $\emptyset \vdash_{\mathcal{L}} \forall x \neg \neg \psi(x) \leftrightarrow \neg \exists x \neg \psi(x)$ | type 6 axiom |
| 2. $\neg \neg \psi(c) \vdash_{\mathcal{L}'} \psi(c)$ | tautology |
| 3. $\forall x \neg \neg \psi(x) \vdash_{\mathcal{L}'} \psi(c)$ | 2, UI |
| 4. $\forall x \neg \neg \psi(x) \vdash_{\mathcal{L}} \forall x \psi(x)$ | 3, UG |
| 5. $\emptyset \vdash_{\mathcal{L}} \forall x \neg \neg \psi(x) \rightarrow \forall x \psi(x)$ | 4, Deduction Theorem |
| 6. $\emptyset \vdash_{\mathcal{L}} \neg \forall x \psi(x) \rightarrow \exists x \neg \psi(x)$ | 1, 5, tautology |

\mathcal{L}' is $\mathcal{L} \cup \{c\}$. In lines (2) and (6), we are really quoting Lemma II.11.6. □

Proof of Lemma II.11.8. For the UI and EG rules, just use Modus Ponens and the fact that the sentences $\forall x \varphi(x) \rightarrow \varphi(\tau)$ and $\varphi(\tau) \rightarrow \exists x \varphi(x)$ are logically axioms (of types 4, 5, respectively).

For UG, let $\varphi_0, \dots, \varphi_n$ be a formal proof in \mathcal{L}' of $\varphi(c)$ from Σ ; so, φ_n is $\varphi(c)$. Let y be a variable which does not occur anywhere in the proof, and let $\psi_i(y)$ be the formula which results from φ_i by replacing all occurrences of c by y ; so $\psi_n(y)$ is $\varphi(y)$. We shall prove by induction on i that $\Sigma \vdash_{\mathcal{L}} \forall y \psi_i(y)$; For $i = n$, our induction will establish $\Sigma \vdash_{\mathcal{L}} \forall y \varphi(y)$. We are then done because $\forall y \varphi(y) \vdash_{\mathcal{L}} \forall x \varphi(x)$ by:

- | | |
|--|--------------------|
| 0. $\forall y \varphi(y)$ | given |
| 1. $\forall x [\forall y \varphi(y) \rightarrow \varphi(x)]$ | type 4 axiom |
| 2. $\forall x [\forall y \varphi(y) \rightarrow \varphi(x)] \rightarrow (\forall x \forall y \varphi(y) \rightarrow \forall x \varphi(x))$ | type 3 axiom |
| 3. $\forall x \forall y \varphi(y) \rightarrow \forall x \varphi(x)$ | 2, 1, modus ponens |
| 4. $\forall y \varphi(y) \rightarrow \forall x \forall y \varphi(y)$ | type 2 axiom |
| 5. $\forall x \forall y \varphi(y)$ | 4, 0, modus ponens |
| 6. $\forall x \varphi(x)$ | 3, 5, modus ponens |

Now, for the induction itself there are three cases. *Case 1* shows why we required a new variable y here, which necessitated this ugly formal proof. The first two cases do not use the inductive hypothesis.

Case 1. φ_i is a logical axiom. Then it is easily checked that $\forall y \psi_i(y)$ is a logical axiom of the same type. In checking this, we use the fact that y does not occur in φ_i . For example, say φ_i is $\forall z \exists x p(z, x) \rightarrow \exists x p(c, x)$, which is a logical axiom of type 4. Then $\psi_i(y)$ is $\forall z \exists x p(z, x) \rightarrow \exists x p(y, x)$, and $\forall y \psi_i(y)$ is a universal closure of $\psi_i(y)$, which is also a logical axiom of type 4. We could not have used the same variable x here; replacing c by x in φ_i yields $\forall z \exists x p(z, x) \rightarrow \exists x p(x, x)$, which is not a logical axiom, and which is not even logically valid.

Case 2. $\varphi_i \in \Sigma$. Then φ_i does not use the constant c , so $\psi_i(y)$ is just the sentence φ_i , and $\Sigma \vdash_{\mathcal{L}} \forall y \psi_i(y)$ because $\varphi_i \rightarrow \forall y \varphi_i$ is logical axiom of type 2.

Case 3. For some $j, k < i$, φ_i follows from φ_j, φ_k by Modus Ponens, so φ_k is

$(\varphi_j \rightarrow \varphi_i)$. Then

- | | |
|--|--------------------|
| 0. $\Sigma \vdash_{\mathcal{L}} \forall y \psi_j(y)$ | induction |
| 1. $\Sigma \vdash_{\mathcal{L}} \forall y (\psi_j(y) \rightarrow \psi_i(y))$ | induction |
| 2. $\Sigma \vdash_{\mathcal{L}} \forall y (\psi_j(y) \rightarrow \psi_i(y)) \rightarrow (\forall y \psi_j(y) \rightarrow \forall y \psi_i(y))$ | type 3 axiom |
| 3. $\Sigma \vdash_{\mathcal{L}} \forall y \psi_j(y) \rightarrow \forall y \psi_i(y)$ | 2, 1, modus ponens |
| 4. $\Sigma \vdash_{\mathcal{L}} \forall y \psi_i(y)$ | 3, 0, modus ponens |

Finally, we verify the EI rule by translating it to an application of UG. Assuming $\Sigma \cup \{\varphi(c)\} \vdash_{\mathcal{L}'} \psi$ we have $\neg \text{Con}_{\vdash, \mathcal{L}'}(\Sigma \cup \{\varphi(c), \neg \psi\})$, so that by proof by contradiction (Lemma II.11.4), we have $\Sigma \cup \{\neg \psi\} \vdash_{\mathcal{L}'} \neg \varphi(c)$. Then, by UG, $\Sigma \cup \{\neg \psi\} \vdash_{\mathcal{L}} \forall x \neg \varphi(x)$. Thus, $\neg \text{Con}_{\vdash, \mathcal{L}}(\Sigma \cup \{\neg \psi, \neg \forall x \neg \varphi(x)\})$, so that $\Sigma \cup \{\neg \forall x \neg \varphi(x)\} \vdash_{\mathcal{L}} \psi$ (using proof by contradiction again).

Now, $\forall x \neg \varphi(x) \leftrightarrow \neg \exists x \varphi(x)$ is a logical axiom (of type 6), and $\neg \forall x \neg \varphi(x)$ follows tautologically from this axiom and $\exists x \varphi(x)$. Thus, using tautological reasoning (Lemma II.11.6) and transitivity of \vdash , we have $\Sigma \cup \{\exists x \varphi(x)\} \vdash_{\mathcal{L}} \psi$ □

When we showed (see Figure II.1, page 109) that $\forall x [p(x) \wedge q(x)] \vdash \forall y p(y)$ by explicitly writing down a formal proof, we remarked informally that we used a type 3 axiom to do a modus ponens step inside a universal quantifier. Now, in the proof of the UG rule, we used a type 3 axiom to justify the modus ponens step in *Case 3*. Note that, using our rules, $\forall x [p(x) \wedge q(x)] \vdash \forall y p(y)$ has become trivial; echoing our informal proof (see page 109), we say:

- | | | |
|---|-----------|--|
| 0. $p(c) \wedge q(c) \vdash_{\mathcal{L}'} p(c)$ | tautology | |
| 1. $\forall x [p(x) \wedge q(x)] \vdash_{\mathcal{L}'} p(c)$ | 1, UI | |
| 2. $\forall x [p(x) \wedge q(x)] \vdash_{\mathcal{L}} \forall y p(y)$ | 2, UG | |

The UG got used where we said informally “since c was arbitrary”.

Exercise II.11.11 Let $\mathcal{L} = \{\cdot\}$, and let $GP = \{\gamma_1, \gamma_2\}$ be the axioms of group theory as stated in Section 0.4. Show that from GP one can prove the cancellation rule:

$$GP \vdash_{\mathcal{L}} \forall xyz [x \cdot y = x \cdot z \rightarrow y = z] .$$

A special case of UG is worth pointing out

Lemma II.11.12 Assume that Σ and φ are in \mathcal{L} , and $\mathcal{L}' = \mathcal{L} \cup \{c\}$, where c is a constant symbol. Assume that $\Sigma \vdash_{\mathcal{L}'} \varphi$. Then $\Sigma \vdash_{\mathcal{L}} \varphi$.

Proof. UG gives us $\Sigma \vdash_{\mathcal{L}} \forall x \varphi$. But since φ is a sentence, $\forall x \varphi \rightarrow \varphi$ is a logical axiom of type 4. □

This can be generalized to:

Exercise II.11.13 Assume that Σ and φ are in \mathcal{L} , $\mathcal{L}' \supset \mathcal{L}$, and $\Sigma \vdash_{\mathcal{L}'} \varphi$. Prove constructively that $\Sigma \vdash_{\mathcal{L}} \varphi$.

Hint. Let $\varphi_0, \dots, \varphi_n$ be a formal proof of φ from Σ ; so φ_n is φ . This proof may use some symbols in $\mathcal{L}' \setminus \mathcal{L}$. Here, “constructively” means that you should give explicit instructions for constructing another proof just using the symbols of \mathcal{L} . Non-constructively, the result will be trivial once the Completeness Theorem is proved (see Lemma II.12.21). There are results about proofs (such as $\text{Con}(ZC)$) which are easy non-constructively (see Exercise I.14.16), but are not provable constructively by the Gödel Incompleteness Theorem.

It may be simpler to do this in two steps. First, assume that \mathcal{L} contains some constant symbol c . Get φ'_i from φ by replacing all atomic formulas $p(\tau_1, \dots, \tau_m)$ with $p \in \mathcal{L}' \setminus \mathcal{L}$ by $\forall x(x = x)$, and by replacing all terms $f(\tau_1, \dots, \tau_m)$ with $f \in \mathcal{L}' \setminus \mathcal{L}$ by c . Then φ'_n is still φ because φ is in \mathcal{L} . Then $\varphi'_0, \dots, \varphi'_n$ isn't exactly a correct formal proof, but it is easy to show by induction that $\Sigma \vdash_{\mathcal{L}} \varphi'_i$.

Now, it is enough to consider the case where $\mathcal{L}' = \mathcal{L} \cup \{c\}$, where one can use Lemma II.11.12. \square

We now end our discussion of proof theory. We have done enough to prove the Completeness Theorem. From the point of view of model theory, a detailed discussion of \vdash is irrelevant, and facts such as Exercises II.11.13 and II.11.9 are uninteresting, since they become trivial when you replace \vdash by \models . If you are interested in actually generating formal proofs, perhaps by a computer program, see Section II.17.

Exercise II.11.14 Show that $\emptyset \vdash \forall x, y (x = y \rightarrow \forall z (z \in x \leftrightarrow z \in y))$.

Hint. $\forall xyz [z = z \wedge x = y \rightarrow (z \in x \leftrightarrow z \in y)]$ is a logical axiom, of type 11 \square

Exercise II.11.15 Show that $\emptyset \not\vdash \forall y \exists x p(x, y) \rightarrow \exists x \forall y p(x, y)$.

Hint. Find a two-element model in which the statement is false, and apply Soundness (Lemma II.10.5). \square

Remark. We often verify $\Sigma \vdash \varphi$ by using some of the above strategies for constructing proofs, but verifications of $\Sigma \not\vdash \varphi$ are almost always model-theoretic, as in Exercise II.11.15, not via some combinatorial analysis of the definition of formal proof. Exercise II.11.15 only requires finitistic reasoning, since we can refute the conclusion in a finite model. Similarly, the axioms for groups do not prove the statement $\forall xy [xy = yx]$ because there is a (six-element) non-abelian group. In many cases, the verification of $\Sigma \not\vdash \varphi$ requires an infinite model, in which case we must consider the entire argument to be formalized within *ZFC*. For example, working in *ZFC*, we can show that the Power Set Axiom is not provable from the other axioms of *ZFC* because the Power Set Axiom is false in the model $H(\aleph_1)$, which satisfies all the other axioms (see Exercise I.14.17).

II.12 The Completeness Theorem

This result relates the semantic (\models) notions to the syntactic (\vdash) notions. It may be viewed either as a result about consistency, or as a result about provability:

Theorem II.12.1 (Completeness Theorem) *Let Σ be a set of sentences of \mathcal{L} . Then*

1. $\text{Con}_{\models}(\Sigma)$ iff $\text{Con}_{+, \mathcal{L}}(\Sigma)$.
2. For every sentence φ of \mathcal{L} , $\Sigma \models \varphi$ iff $\Sigma \vdash_{\mathcal{L}} \varphi$.

Once this is proved, we can drop the subscripts on the “ \vdash ” and on the “Con”.

Actually, the two “iff” statements in Theorem II.12.1 are easily seen to be equivalent, and we have already proved one direction of them. To focus on what we still need to prove, we state:

Lemma II.12.2 (Main Lemma) *Let Σ be a set of sentences of \mathcal{L} , and assume that $\text{Con}_{+, \mathcal{L}}(\Sigma)$. Then $\text{Con}_{\models}(\Sigma)$.*

Lemma II.12.3 *Lemma II.12.2 implies Theorem II.12.1.*

Proof. We have already proved $\Sigma \vdash_{\mathcal{L}} \varphi \Rightarrow \Sigma \models \varphi$, the *soundness* direction of (2) (see Lemma II.10.5). This proves the soundness direction of (1), $\text{Con}_{\models}(\Sigma) \Rightarrow \text{Con}_{+, \mathcal{L}}(\Sigma)$, since if $\neg \text{Con}_{+, \mathcal{L}}(\Sigma)$ then there is some φ such that $\Sigma \vdash_{\mathcal{L}} \varphi$ and $\Sigma \vdash_{\mathcal{L}} \neg \varphi$ (see Definition II.11.2); but then $\Sigma \models \varphi$ and $\Sigma \models \neg \varphi$, so by definition of \models (Definition II.7.11), there can be no model of Σ , so $\neg \text{Con}_{\models}(\Sigma)$.

Assuming Lemma II.12.2, we have both directions of (1), but then we have (2), since

$$\Sigma \vdash_{\mathcal{L}} \varphi \Leftrightarrow \neg \text{Con}_{+, \mathcal{L}}(\Sigma \cup \{\neg \varphi\}) \Leftrightarrow \neg \text{Con}_{\models}(\Sigma \cup \{\neg \varphi\}) \Leftrightarrow \Sigma \models \varphi .$$

Here, the first “ \Leftrightarrow ” uses Lemma II.11.4 (on proof by contradiction), the second “ \Leftrightarrow ” uses (1), and the third “ \Leftrightarrow ” is clear from the meaning of \models . \square

Now, we turn to the proof of the Main Lemma. The Completeness Theorem was first proved by Gödel in 1929. Independently, in 1929, Herbrand described a method for constructing models using the terms of \mathcal{L} . In 1949, Henkin [14] realized that one might use Herbrand’s ideas as the basis for an exposition of the Completeness Theorem. Following (roughly) [14], there are three basic steps, which in logical order are:

- Step 1. Add witnessing constants.
- Step 2. Extend to a maximal consistent set.
- Step 3. Write down the Herbrand model and prove that it works.

We shall in fact start with (3), before even explaining what (1) and (2) mean. Once it is clear what the Herbrand model is, (1) and (2) will arise naturally in an attempt to prove that the Herbrand model works.

To prove the Main Lemma, we are assuming $\text{Con}_{+, \mathcal{L}}(\Sigma)$, which is a purely syntactic assumption about Σ . We must prove that $\text{Con}_{\models}(\Sigma)$, which means that we must build a model $\mathfrak{A} \models \Sigma$. Where will our model come from? Since all we are given is syntax, we must use our syntactic objects to form the model. Thus, our construction generalizes

some constructions in algebra, such as free groups and polynomial rings, where a ring or group or field is constructed from syntactic objects.

We start by describing the universe of the model. If we had a model, the *terms* of the language would denote elements of the model. Since we don't have a model yet, it is natural to let the terms themselves be the objects in the model. Actually, we only use the terms with no variables, since these denote "fixed" elements of the model:

Definition II.12.4 *A term τ of \mathcal{L} is closed iff τ contains no variables. Let $\mathbf{CT}_0(\mathcal{L})$ be the set of all closed terms of \mathcal{L} .*

For such a τ , its value $\text{val}_{\mathfrak{A}}(\tau)$ depends only on τ and \mathfrak{A} , not on any variable assignment (see Definition II.7.4 and Exercise II.7.5). Since we are not allowing the empty structure, we can only use $\mathbf{CT}_0(\mathcal{L})$ as the universe of the model when \mathcal{L} has some closed terms; equivalently when $\mathcal{F}_0 \neq \emptyset$, where, as in Definition II.5.2, \mathcal{F}_0 is the set of constant symbols of \mathcal{L} . Then, we build a structure for \mathcal{L} in the sense of Definition II.7.1 as follows:

Definition II.12.5 *Let Σ be a set of sentences of \mathcal{L} and assume that $\mathcal{F}_0 \neq \emptyset$. Define the closed term model $\mathfrak{A}_0 = \mathfrak{CT}_0(\mathcal{L}, \Sigma)$ to be the \mathcal{L} -structure whose universe is $\mathbf{CT}_0(\mathcal{L})$ so that:*

- ☞ If $f \in \mathcal{F}_n$ with $n > 0$, and $\tau_1, \dots, \tau_n \in \mathbf{CT}_0(\mathcal{L})$, then $f_{\mathfrak{A}_0}(\tau_1, \dots, \tau_n)$ is the closed term $f(\tau_1, \dots, \tau_n)$.
- ☞ If $p \in \mathcal{P}_n$ with $n > 0$, and $\tau_1, \dots, \tau_n \in \mathbf{CT}_0(\mathcal{L})$, then $(\tau_1, \dots, \tau_n) \in p_{\mathfrak{A}_0}$ iff $\Sigma \vdash_{\mathcal{L}} p(\tau_1, \dots, \tau_n)$.
- ☞ If $c \in \mathcal{F}_0$, then $c_{\mathfrak{A}_0} = c$.
- ☞ If $p \in \mathcal{P}_0$, then $p_{\mathfrak{A}_0} = 1(\text{true})$ iff $\Sigma \vdash_{\mathcal{L}} p$.

An easy induction shows that the value of a closed term (see Definition II.7.4) is itself:

Lemma II.12.6 *If $\tau \in \mathbf{CT}_0(\mathcal{L})$ then $\text{val}_{\mathfrak{CT}_0(\mathcal{L}, \Sigma)}(\tau) = \tau$.*

Note that the definition of $\mathfrak{CT}_0(\mathcal{L}, \Sigma)$ makes sense even if Σ is inconsistent, so we cannot possibly claim that $\mathfrak{CT}_0(\mathcal{L}, \Sigma) \models \Sigma$ in general. Also note that the interpretations of the functions do not depend on Σ . For example, say $\mathcal{L} = \{+, 0\}$ and Σ contains the axiom $\forall x[x + 0 = x]$. The domain $\mathbf{CT}_0(\mathcal{L})$ is countably infinite, and is formed just using \mathcal{L} , not Σ . The domain contains, for example, the closed terms $0, 0 + 0, 0 + (0 + 0), (0 + 0) + 0$, etc. These are all distinct objects, so that, for example, $\mathfrak{CT}_0(\mathcal{L}, \Sigma) \models 0 + 0 \neq 0$, even though $\Sigma \vdash 0 + 0 = 0$. This example indicates that the elements of our domain should not really be the closed terms, but rather the *equivalence classes* of closed terms, where two terms are equivalent iff Σ proves them to be equal.

Definition II.12.7 Define a relation \sim (actually, $\sim_{\mathcal{L}, \Sigma}$) on $\mathbf{CT}_0(\mathcal{L})$ by:

$$\tau \sim \sigma \text{ iff } \Sigma \vdash_{\mathcal{L}} \tau = \sigma \text{ .}$$

Lemma II.12.8 \sim is an equivalence relation on $\mathbf{CT}_0(\mathcal{L})$.

Proof. We must verify that \sim is reflexive, symmetric, and transitive. (see Definition I.7.2). To prove that \sim is symmetric, we assume that $\Sigma \vdash_{\mathcal{L}} \tau = \sigma$ and we must prove that $\Sigma \vdash_{\mathcal{L}} \sigma = \tau$. Now, $\forall x, y [x = y \leftrightarrow y = x]$ is a logical axiom of type 8, so $\emptyset \vdash_{\mathcal{L}} \tau = \sigma \leftrightarrow \sigma = \tau$ by UI (Lemma II.11.8). Thus, $\Sigma \vdash_{\mathcal{L}} \sigma = \tau$ because $\sigma = \tau$ follows tautologically from $\tau = \sigma$ and $\tau = \sigma \leftrightarrow \sigma = \tau$ (see Lemma II.11.6). The proofs that \sim is reflexive and transitive are similar, using logical axioms of types 7 and 9. \square

We can now form the quotient set $\mathbf{CT}_0(\mathcal{L})/\sim$ (see Definition I.7.15), but we also need to define an appropriate \mathcal{L} -structure on it:

Definition II.12.9 . Let Σ be a set of sentences of \mathcal{L} and assume that $\mathcal{F}_0 \neq \emptyset$. Define $\mathbf{CT}(\mathcal{L}, \Sigma) = \mathbf{CT}_0(\mathcal{L})/\sim$, and define the Herbrand model $\mathfrak{A} = \mathfrak{CT}(\mathcal{L}, \Sigma)$ to be the \mathcal{L} -structure whose universe is $\mathbf{CT}(\mathcal{L}, \Sigma)$ so that:

- \Leftrightarrow If $f \in \mathcal{F}_n$ with $n > 0$, and $[\tau_1], \dots, [\tau_n] \in \mathbf{CT}(\mathcal{L}, \Sigma)$, then $f_{\mathfrak{A}}([\tau_1], \dots, [\tau_n])$ is the equivalence class $[f(\tau_1, \dots, \tau_n)]$.
- \Leftrightarrow If $p \in \mathcal{P}_n$ with $n > 0$, and $[\tau_1], \dots, [\tau_n] \in \mathbf{CT}(\mathcal{L}, \Sigma)$, then $([\tau_1], \dots, [\tau_n]) \in p_{\mathfrak{A}}$ iff $\Sigma \vdash_{\mathcal{L}} p(\tau_1, \dots, \tau_n)$.
- \Leftrightarrow If $c \in \mathcal{F}_0$, then $c_{\mathfrak{A}} = [c]$.
- \Leftrightarrow If $p \in \mathcal{P}_0$, then $p_{\mathfrak{A}} = 1(\text{true})$ iff $\Sigma \vdash_{\mathcal{L}} p$.

Justification. When $n > 0$, we must check that our definitions of $f_{\mathfrak{A}}$ and $p_{\mathfrak{A}}$ are independent of the chosen representatives of the equivalence classes. Specifically, say $\tau_i \sim \sigma_i$ for $i = 1, \dots, n$. Then each $[\tau_i] = [\sigma_i]$. Our definition of $f_{\mathfrak{A}}$ would be ambiguous unless we can check that $[f(\tau_1, \dots, \tau_n)] = [f(\sigma_1, \dots, \sigma_n)]$; that is, $f(\tau_1, \dots, \tau_n) \sim f(\sigma_1, \dots, \sigma_n)$.

Now, $\forall x_1, \dots, x_n, y_1, \dots, y_n [x_1 = y_1 \wedge \dots \wedge x_n = y_n \rightarrow (f(x_1 \dots x_n) = f(y_1 \dots y_n))]$ is a logical axiom of type 10, so $\emptyset \vdash_{\mathcal{L}} [\tau_1 = \sigma_1 \wedge \dots \wedge \tau_n = \sigma_n \rightarrow (f(\tau_1 \dots \tau_n) = f(\sigma_1 \dots \sigma_n))]$ by UI (Lemma II.11.8). Since $\Sigma \vdash_{\mathcal{L}} \tau_i = \sigma_i$ for each i by our definition of \sim , we have $\Sigma \vdash_{\mathcal{L}} f(\tau_1, \dots, \tau_n) = f(\sigma_1, \dots, \sigma_n)$ because it follows tautologically (see Lemma II.11.6). Thus, $f(\tau_1, \dots, \tau_n) \sim f(\sigma_1, \dots, \sigma_n)$.

Likewise, our definition of $p_{\mathfrak{A}}$ would be ambiguous unless we can check that $\Sigma \vdash_{\mathcal{L}} p(\tau_1, \dots, \tau_n)$ iff $\Sigma \vdash_{\mathcal{L}} p(\sigma_1, \dots, \sigma_n)$. But this follows by a similar argument, since $\forall x_1, \dots, x_n, y_1, \dots, y_n [x_1 = y_1 \wedge \dots \wedge x_n = y_n \rightarrow (p(x_1 \dots x_n) \leftrightarrow p(y_1 \dots y_n))]$ is a logical axiom of type 11. \square

Note that in this proof and the proof of Lemma II.12.8, we “just happened” to have the required statements in our list of logical axioms. There is nothing magical about this.

As we pointed out in Section II.10, the specific list of logical axioms is a bit arbitrary; we chose our list so that our intended proof of the Completeness Theorem would work.

In the quotient structure $\mathfrak{C}\mathfrak{I}(\mathcal{L}, \Sigma)$, terms which “should be” equal are equal. To continue the previous example, if $\mathcal{L} \supseteq \{+, 0\}$ and Σ contains $\forall x[x + 0 = x]$, then $\mathfrak{C}\mathfrak{I}_0(\mathcal{L}, \Sigma) \models 0 + 0 \neq 0$, but $\mathfrak{C}\mathfrak{I}(\mathcal{L}, \Sigma) \models 0 + 0 = 0$ because $[0 + 0] = [0]$. More generally, the following lemma spells out the extent to which $\mathfrak{C}\mathfrak{I}(\mathcal{L}, \Sigma)$ “works” without further restrictions on Σ :

Lemma II.12.10 *Let Σ be a set of sentences of \mathcal{L} and assume that $\mathcal{F}_0 \neq \emptyset$. Let $\mathfrak{A} = \mathfrak{C}\mathfrak{I}(\mathcal{L}, \Sigma)$. Then*

1. *If τ is any closed term of \mathcal{L} , then $\text{val}_{\mathfrak{A}}(\tau) = [\tau]$ (see Definition II.7.4).*
2. *If φ is a sentence of \mathcal{L} of the form $\forall x_1, \dots, x_n \psi(x_1, \dots, x_n)$, then $\mathfrak{A} \models \varphi$ iff $\mathfrak{A} \models \psi(\tau_1, \dots, \tau_n)$ for all closed terms τ_1, \dots, τ_n .*
3. *If φ is an atomic sentence, then $\Sigma \vdash_{\mathcal{L}} \varphi$ iff $\mathfrak{A} \models \varphi$.*
4. *If $\Sigma \vdash_{\mathcal{L}} \varphi$, where φ is a closure of an atomic formula, then $\mathfrak{A} \models \varphi$.*

Proof. (1) is easily proved by induction on τ . For (2), use (1) plus Lemma II.8.13, plus the fact that every element of A is of the form $[\tau]$ for some closed term τ .

For (3), there are two cases. If φ is $p(\tau_1, \dots, \tau_n)$, where τ_1, \dots, τ_n are closed terms, then

$$\mathfrak{A} \models \varphi \Leftrightarrow ([\tau_1], \dots, [\tau_n]) \in p_{\mathfrak{A}} \Leftrightarrow \Sigma \vdash_{\mathcal{L}} p(\tau_1, \dots, \tau_n) ;$$

here, we are using the definitions of \models and $p_{\mathfrak{A}}$ and the fact that each $\text{val}_{\mathfrak{A}}(\tau_i) = [\tau_i]$ by (1). If φ is $\tau_1 = \tau_2$, then $\mathfrak{A} \models \varphi$ iff $[\tau_1] = [\tau_2]$ iff $\Sigma \vdash_{\mathcal{L}} \tau_1 = \tau_2$.

For (4), say φ is $\forall x_1, \dots, x_k \psi(x_1, \dots, x_k)$, where ψ is atomic. Then by (2), we need to show $\mathfrak{A} \models \psi(\tau_1, \dots, \tau_k)$ for all closed terms τ_1, \dots, τ_k . Since $\Sigma \vdash_{\mathcal{L}} \varphi$, we have $\Sigma \vdash_{\mathcal{L}} \psi(\tau_1, \dots, \tau_k)$ by UI (Lemma II.11.8). Since $\psi(\tau_1, \dots, \tau_k)$ is an atomic sentence, the result follows by (3). \square

In particular, if all the sentences of Σ happen to be closures of atomic formulas, then $\mathfrak{C}\mathfrak{I}(\mathcal{L}, \Sigma) \models \Sigma$. This situation does occur occasionally. For example, say $\mathcal{L} = \{\cdot, i, 1, a, b\}$, where a, b are constant symbols, and Σ is the set of axioms for groups $\{\gamma_1, \gamma_{2,1}, \gamma_{2,2}\}$ on page 88. Then Σ does not mention a, b , but does imply that various closed terms involving a, b which “should” be equal are – e.g., by the associative law γ_1 , and UI, $\Sigma \vdash a \cdot (b \cdot a) = (a \cdot b) \cdot a$. More generally, Lemma II.12.10 implies that $\mathfrak{C}\mathfrak{I}(\mathcal{L}, \Sigma)$ is a group, since the three axioms of Σ are universally quantified equations. It is easily seen to be the free group on two generators. Further applications to algebra are described in Section II.14.

In many cases, $\mathfrak{C}\mathfrak{I}(\mathcal{L}, \Sigma)$ will fail to be a model for Σ . For a trivial example, Σ might be inconsistent, so it proves everything and has no model. Now all closed terms are equivalent, so that $\mathfrak{C}\mathfrak{I}(\mathcal{L}, \Sigma)$ is a 1-element model, and all predicates are true of this element. As we expect from (3) of Lemma II.12.10, all atomic sentences are true in this model.

The following less trivial example illustrates the two main reasons that we might have $\mathfrak{C}\mathfrak{T}(\mathcal{L}, \Sigma) \not\models \Sigma$ when Σ is consistent. Let $\mathcal{L} = \{<, a, b\}$, where a, b are constant symbols, and let Σ say that $<$ is a strict total order (see Definition I.7.2) with no largest element ($\forall y \exists x (y < x)$). Then Σ does not mention a, b . The only closed terms are a and b ; and $a \not\sim b$, since $\Sigma \not\vdash a = b$. Thus, $\mathfrak{A} = \mathfrak{C}\mathfrak{T}(\mathcal{L}, \Sigma)$ has two elements, $[a] = \{a\}$ and $[b] = \{b\}$, and $<_{\mathfrak{A}}$ is the empty relation, since $\Sigma \not\vdash \tau < \sigma$ for any τ, σ in $\{a, b\}$.

- Problem 1. $\Sigma \vdash \exists x (b < x)$, but $\mathfrak{A} \not\models \exists x (b < x)$.
- Problem 2. Σ contains the total order axiom trichotomy (see Definition I.7.2), so by UI, $\Sigma \vdash (a < b \vee b < a \vee a = b)$, but this is false in \mathfrak{A} because $<_{\mathfrak{A}}$ is empty.

These two problems are cured by Steps 1 and 2 on page 117. In Step 1, we add a new “witnessing constant” to name something that should exist; in this example, we could add a constant c plus the axiom $b < c$. This process must be repeated infinitely many times, since any total order with no largest element is infinite. In Step 2, we extend the consistent Σ to a maximally consistent set; in this example, we wind up choosing one of the three disjuncts, $a < b$, $b < a$, and $a = b$, and adding them to Σ . These two steps may be discussed in either order. We start with Step 2, which is a little quicker to explain:

Definition II.12.11 *A set of sentences Σ in \mathcal{L} is maximally (\vdash, \mathcal{L}) consistent iff*

1. $\text{Con}_{\vdash, \mathcal{L}}(\Sigma)$, and
2. There is no set of sentences Π in \mathcal{L} such that $\text{Con}_{\vdash, \mathcal{L}}(\Pi)$ and $\Sigma \subsetneq \Pi$.

As mentioned before, once the Completeness Theorem is proved, we shall just write $\text{Con}(\Sigma)$. Then, one usually just says that Σ is “maximally consistent”, but it is still important that we have a specific \mathcal{L} in mind, since every consistent Σ will have proper supersets which are consistent if we are allowed to expand \mathcal{L} .

Lemma II.12.12 *If Δ is a set of sentences in \mathcal{L} and $\text{Con}_{\vdash, \mathcal{L}}(\Delta)$, then there is a Σ in \mathcal{L} such that $\Sigma \supseteq \Delta$ and Σ is maximally (\vdash, \mathcal{L}) consistent.*

Proof. Let S be the set of all sentences of \mathcal{L} ; then $\Delta \in \mathcal{P}(S)$. Let $\mathcal{F} = \{\Pi \in \mathcal{P}(S) : \text{Con}_{\vdash, \mathcal{L}}(\Pi)\}$. Then \mathcal{F} is of finite character (see Definition I.12.5) because every formal proof from a $\Pi \in \mathcal{P}(S)$ only uses finitely many sentences of Π . It follows by Tukey’s Lemma (Definition I.12.7) that there is a maximal $\Sigma \in \mathcal{F}$ such that $\Sigma \supseteq \Delta$. \square

One could equally well use Zorn’s Lemma or transfinite recursion to prove that Σ exists. In the proof of the Completeness Theorem, if we are trying to produce a model for Δ , we shall first get a maximal $\Sigma \supseteq \Delta$ and then get an $\mathfrak{A} \models \Sigma$. Of course, we then have $\mathfrak{A} \models \Delta$. The closed term model $\mathfrak{C}\mathfrak{T}(\mathcal{L}, \Sigma)$ will be easier to deal with than $\mathfrak{C}\mathfrak{T}(\mathcal{L}, \Delta)$ because of the following properties which follow from maximality:

Lemma II.12.13 *Assume that Σ in \mathcal{L} is maximally (\vdash, \mathcal{L}) consistent. Then for any sentence φ, ψ of \mathcal{L} :*

1. $\Sigma \vdash_{\mathcal{L}} \varphi$ iff $\varphi \in \Sigma$.
2. $(\neg\varphi) \in \Sigma$ iff $\varphi \notin \Sigma$.
3. $(\varphi \vee \psi) \in \Sigma$ iff $\varphi \in \Sigma$ or $\psi \in \Sigma$.

Proof. For (1): \Leftarrow is clear. For \Rightarrow , using $\text{Con}_{\vdash, \mathcal{L}}(\Sigma)$ and $\Sigma \vdash_{\mathcal{L}} \varphi$ we get $\text{Con}_{\vdash, \mathcal{L}}(\Sigma \cup \{\varphi\})$; so $\varphi \in \Sigma$ by maximality.

For (2): \Rightarrow is clear from $\text{Con}_{\vdash, \mathcal{L}}(\Sigma)$. For \Leftarrow , $\varphi \notin \Sigma$ implies $\neg\text{Con}_{\vdash, \mathcal{L}}(\Sigma \cup \{\varphi\})$ by maximality. But then $\Sigma \vdash_{\mathcal{L}} \neg\varphi$ using proof by contradiction (Lemma II.11.4), so $(\neg\varphi) \in \Sigma$ by (1).

For (3) \Leftarrow : If Σ contains either φ or ψ , then $\Sigma \vdash_{\mathcal{L}} \varphi \vee \psi$ because $\varphi \vee \psi$ follows tautologically (see Lemma II.11.6). Then $(\varphi \vee \psi) \in \Sigma$ by (1).

For (3) \Rightarrow : Suppose $(\varphi \vee \psi) \in \Sigma$, and $\varphi \notin \Sigma$ and $\psi \notin \Sigma$. Then Σ contains $\neg\varphi$ and $\neg\psi$ by (2), contradicting $\text{Con}_{\vdash, \mathcal{L}}(\Sigma)$ \square

To continue the above example, where Σ in $\mathcal{L} = \{<, a, b\}$ says that that $<$ is a strict total order with no largest element, let $\Sigma' \supset \Sigma$ be maximal. Using Σ' , Lemma II.12.13 cures Problem 2. That is, now let $\mathfrak{A} = \mathfrak{CT}(\mathcal{L}, \Sigma')$. $\Sigma' \vdash (a < b \vee b < a \vee a = b)$, so $(a < b \vee b < a \vee a = b) \in \Sigma'$, so Σ' contains one of $a < b$, $b < a$, $a = b$; which one depends on Σ' (the maximal extension is not unique), but here Σ' cannot contain more than one, since any two of them contradict the strict total order axioms. If Σ' contains $a = b$, then $[a] = [b] = \{a, b\}$, so \mathfrak{A} is a 1-element total order, in which $<_{\mathfrak{A}}$ is empty, as it should be. If Σ contains $a < b$, then \mathfrak{A} is a 2-element total order, with a below b .

This example can be generalized; maximally consistent sets handle all sentences which do not use quantifiers:

Lemma II.12.14 *Let Σ be a set of sentences of \mathcal{L} . Assume that $\mathcal{F}_0 \neq \emptyset$ and that Σ is maximally (\vdash, \mathcal{L}) consistent. Let $\mathfrak{A} = \mathfrak{CT}(\mathcal{L}, \Sigma)$. Let φ be a sentence of \mathcal{L} which does not use any quantifiers. Then $\varphi \in \Sigma$ iff $\mathfrak{A} \models \varphi$.*

Proof. To simplify the notation, recall, from Definition II.7.8, that $\text{val}_{\mathfrak{A}}(\varphi)$ denotes the truth value (T or F) of the sentence φ in \mathfrak{A} . Let us now define $\text{val}_{\Sigma}(\varphi)$ to be T iff $\varphi \in \Sigma$ and F iff $\varphi \notin \Sigma$ (equivalently, iff $\neg\varphi \in \Sigma$, by Lemma II.12.13). Then Lemma II.12.14 may be restated as saying that $\text{val}_{\mathfrak{A}}(\varphi) = \text{val}_{\Sigma}(\varphi)$ whenever φ uses no quantifiers.

We now induct on φ . Since φ uses no quantifiers, it must be obtained from atomic sentences using propositional connectives.

For the basis of the induction, assume that φ is an atomic sentence. Then $\varphi \in \Sigma$ iff $\Sigma \vdash_{\mathcal{L}} \varphi$ by Lemma II.12.13, and $\Sigma \vdash_{\mathcal{L}} \varphi$ iff $\mathfrak{A} \models \varphi$ by Lemma II.12.10.

For the induction step, we assume that the lemma holds for shorter sentences (the inductive hypothesis), and we prove the lemma for φ . There are five cases, depending on how φ is constructed from shorter sentences.

If φ is $\neg\psi$, then $\neg\psi \in \Sigma$ iff $\psi \notin \Sigma$ iff $\mathfrak{A} \not\models \psi$ iff $\mathfrak{A} \models \neg\psi$. The three “iff”s used, respectively, Lemma II.12.13, the inductive hypothesis, and the definition of \models .

For the other four cases, φ is $\psi \odot \chi$, where \odot is one of $\vee, \wedge, \rightarrow, \leftrightarrow$. By Definition II.7.8, $\text{val}_{\mathfrak{A}}(\varphi)$ is computed from $\text{val}_{\mathfrak{A}}(\psi)$ and $\text{val}_{\mathfrak{A}}(\chi)$ using the truth table for \odot (see Table 1, page 4). Our induction is completed by noting that $\text{val}_{\Sigma}(\varphi)$ is also determined from $\text{val}_{\Sigma}(\psi)$ and $\text{val}_{\Sigma}(\chi)$ using the same truth tables. To verify this, we must examine the four possibilities for \odot .

If \odot is \vee , this is immediate from Lemma II.12.13. Now, say \odot is \leftrightarrow . If $\text{val}_{\Sigma}(\psi) = \text{val}_{\Sigma}(\chi)$, then either $\{\psi, \chi\} \subseteq \Sigma$ or $\{\neg\psi, \neg\chi\} \subseteq \Sigma$, so $\Sigma \vdash_{\mathcal{L}} \varphi$ (because φ follows tautologically; see Lemma II.11.6), so $\varphi \in \Sigma$ by Lemma II.12.13, so $\text{val}_{\Sigma}(\varphi) = T$. If $\text{val}_{\Sigma}(\psi) \neq \text{val}_{\Sigma}(\chi)$, then either $\{\psi, \neg\chi\} \subseteq \Sigma$ or $\{\neg\psi, \chi\} \subseteq \Sigma$, so $\Sigma \vdash_{\mathcal{L}} \neg\varphi$ (because $\neg\varphi$ follows tautologically), so $\neg\varphi \in \Sigma$, so $\text{val}_{\Sigma}(\varphi) = F$. In all four cases for $\text{val}_{\Sigma}(\psi), \text{val}_{\Sigma}(\chi)$, we have $\text{val}_{\Sigma}(\varphi)$ determined from $\text{val}_{\Sigma}(\psi)$ and $\text{val}_{\Sigma}(\chi)$ using the truth tables for \leftrightarrow . Of course, none of this is surprising, since the same truth tables were used to define “follows tautologically”. The other two cases for \odot are analyzed similarly. \square

Now that we have cured Problem 2, we cure Problem 1, which involves sentences with quantifiers. As mentioned above, we must make sure that whenever the axioms imply that something exists, we have a closed term which names it. This is made formal by Definition II.12.16.

Definition II.12.15 *An existential sentence is a sentence of the form $\exists x\varphi(x)$.*

Here, no variable besides x can be free in φ , so $\varphi(\tau)$ is a sentence for each closed term τ .

Definition II.12.16 *If Σ is a set of sentences of \mathcal{L} and $\exists x\varphi(x)$ is an existential sentence, then τ is a witnessing term for $\exists x\varphi(x)$ (with respect to Σ, \mathcal{L}) iff $\tau \in \mathbf{CT}_0(\mathcal{L})$ and $\Sigma \vdash_{\mathcal{L}} (\exists x\varphi(x) \rightarrow \varphi(\tau))$. Then Σ has witnesses in \mathcal{L} iff every existential sentence has some witnessing term.*

Many “natural” sets of axioms have witnessing terms for some existential sentences and not for others. For example, let Σ be the axioms for fields, expressed in $\mathcal{L} = \{0, 1, +, \cdot, -, i\}$, where “ $-$ ” denotes the unary additive inverse and “ i ” denotes the unary multiplicative inverse (or reciprocal); see Example II.8.23. Let $\varphi(x)$ be $x + 1 = 0$ and let $\psi(x)$ be $(1 + 1) \cdot x = 1$. Then -1 is a witnessing term for $\exists x\varphi(x)$ and $i(1 + 1)$ is a witnessing term for $\exists x\psi(x)$. In the case of φ , both $\exists x\varphi(x)$ and $\varphi(-1)$ are provable from Σ . In the case of ψ , neither $\exists x\psi(x)$ nor $\psi(i(1 + 1))$ is provable from Σ (since a field may have characteristic 2), but the implication $\exists x\psi(x) \rightarrow \psi(i(1 + 1))$ is provable (using axiom (5) of Example II.8.23). However, there is no symbol in the language for $\sqrt{2}$, so if $\chi(x)$ is $x \cdot x = 1 + 1$, then there is no witnessing term for $\exists x\chi(x)$; for example, in the field of real numbers, $\exists x\chi(x)$ is true but each closed term τ denote a rational, so $\chi(\tau)$ is false, hence the implication $\exists x\chi(x) \rightarrow \chi(\tau)$ is false for all closed terms τ . Informally, it would be consistent to add a new constant symbol c plus the axiom $\exists x\chi(x) \rightarrow \chi(c)$. This axiom does not assert that 2 must have a square root, but only that if it does, then “ c ”

denotes a square root of 2. Thus, informally, not only is this new axiom consistent, but it says essentially nothing new; it actually yields a *conservative extension* in the sense that if θ is any statement about fields which does not mention “ c ”, and θ is provable using the new axiom, then θ is provable without using the new axiom. This assertion will be made formal by Lemmas II.12.18 and II.12.20 and Exercise II.12.23, which show that every consistent Σ can be extended to have witnesses for all existential sentences if enough constants are added. First, we prove that this really will cure Problem 1:

Lemma II.12.17 *Let Σ be a set of sentences of \mathcal{L} . Assume that Σ is maximally (\vdash, \mathcal{L}) consistent and that Σ has witnesses in \mathcal{L} . Let $\mathfrak{A} = \mathfrak{CT}(\mathcal{L}, \Sigma)$. Then $\mathfrak{A} \models \Sigma$.*

Proof. We prove that

$$\varphi \in \Sigma \leftrightarrow \mathfrak{A} \models \varphi \quad (*)$$

for every sentence φ of \mathcal{L} . Let $S(\varphi)$ be the total number of occurrences of the symbols $\wedge, \vee, \neg, \rightarrow, \leftrightarrow, \forall, \exists$ in φ . We induct on $S(\varphi)$.

If $S(\varphi) = 0$ (i.e., φ is atomic), then $(*)$ holds by Lemma II.12.14.

Now, assume that $S(\varphi) > 0$ and that $(*)$ holds for sentences with smaller S . There are two main cases to consider.

In the propositional cases, φ is either of the form $\neg\psi$ or of the form φ is $\psi \odot \chi$, where \odot is one of $\vee, \wedge, \rightarrow, \leftrightarrow$. We then derive $(*)$ for φ from $(*)$ for ψ (or $(*)$ for ψ and χ) exactly as in the proof of Lemma II.12.14.

In the quantifier cases, φ is either $\exists x\psi(x)$ or $\forall x\psi(x)$.

If φ is $\exists x\psi(x)$, then fix a witnessing term τ such that $\Sigma \vdash_{\mathcal{L}} (\exists x\psi(x) \rightarrow \psi(\tau))$. Then $S(\psi(\tau)) = S(\varphi) - 1$, so $(*)$ holds for $\psi(\tau)$. If $\varphi \in \Sigma$, then $\Sigma \vdash_{\mathcal{L}} \psi(\tau)$, so $\psi(\tau) \in \Sigma$ by maximality (see Lemma II.12.13), so that $\mathfrak{A} \models \psi(\tau)$ and hence $\mathfrak{A} \models \varphi$. Conversely, if $\mathfrak{A} \models \varphi$ then by definition of “ \models ”, $\mathfrak{A} \models \psi[a]$ for some $a \in A$. By definition of $\mathfrak{CT}(\mathcal{L}, \Sigma)$, this a is of the form $[\pi]$ for some closed term π , so that $\mathfrak{A} \models \psi(\pi)$ (using Lemmas II.12.10.1 and II.8.10). By $(*)$ for $\psi(\pi)$, we have $\psi(\pi) \in \Sigma$, so $\Sigma \vdash_{\mathcal{L}} \varphi$ by EG (Lemma II.11.8), and then $\varphi \in \Sigma$ by maximality.

Note that unlike in the proof of Lemma II.12.14, we cannot induct on the length of φ , since $\psi(\pi)$ could be longer than φ .

Finally, say φ is $\forall x\psi(x)$. If $\varphi \in \Sigma$, then $\psi(\pi) \in \Sigma$ for every closed term π (now using UI plus maximality). Then, as in the \exists case, we have $\mathfrak{A} \models \psi(\pi)$ for all π , so $\mathfrak{A} \models \varphi$. Conversely, suppose that $\varphi \notin \Sigma$. Then $\neg\varphi \in \Sigma$ by maximality. Fix a witnessing term τ such that $\Sigma \vdash_{\mathcal{L}} (\exists x\neg\psi(x) \rightarrow \neg\psi(\tau))$. Since $\emptyset \vdash_{\mathcal{L}} \neg\forall x\psi(x) \rightarrow \exists x\neg\psi(x)$ (see Example II.11.10), we have $\Sigma \vdash_{\mathcal{L}} \neg\psi(\tau)$ (since it follows tautologically), so $\psi(\tau) \notin \Sigma$ (since Σ is consistent), so that $\mathfrak{A} \not\models \psi(\tau)$ by $(*)$ for $\psi(\tau)$, and hence $\mathfrak{A} \not\models \varphi$. \square

We still need to prove that we can construct sets of sentences with witnesses. We begin with:

Lemma II.12.18 *Assume that Σ is a set of sentences of \mathcal{L} , $\exists x\varphi(x)$ is an existential sentence of \mathcal{L} , and $\mathcal{L}' = \mathcal{L} \cup \{c\}$, where c is a constant symbol and $c \notin \mathcal{L}$. Let $\Sigma' = \Sigma \cup \{\exists x\varphi(x) \rightarrow \varphi(c)\}$. Assume that $\text{Con}_{\vdash, \mathcal{L}}(\Sigma)$. Then $\text{Con}_{\vdash, \mathcal{L}'}(\Sigma')$.*

Proof. We use the proof theory facts from Section II.11. Assume that $\neg\text{Con}_{+, \mathcal{L}'}(\Sigma')$. Then we get

- | | |
|---|------------------------|
| 0. $\Sigma \vdash_{\mathcal{L}'} \neg(\exists x\varphi(x) \rightarrow \varphi(c))$ | proof by contradiction |
| 1. $\Sigma \vdash_{\mathcal{L}'} \exists x\varphi(x)$ | 0, tautology |
| 2. $\Sigma \vdash_{\mathcal{L}'} \neg\varphi(c)$ | 0, tautology |
| 3. $\Sigma \vdash_{\mathcal{L}} \forall x\neg\varphi(x)$ | 2, UG |
| 4. $\emptyset \vdash_{\mathcal{L}} \forall x\neg\varphi(x) \leftrightarrow \neg\exists x\varphi(x)$ | type 6 axiom |
| 5. $\Sigma \vdash_{\mathcal{L}} \neg\exists x\varphi(x)$ | 3, 4, tautology |
| 6. $\Sigma \vdash_{\mathcal{L}} \exists x\varphi(x)$ | 1, Lemma II.11.12 |

Using (5)(6), we see that $\neg\text{Con}_{+, \mathcal{L}}(\Sigma)$. In steps (1)(2)(5), we are really quoting Theorem II.11.6. In Step (0), we are quoting Lemma II.11.4. \square

We can iterate this lemma and add witnessing constants for several sentences by adding one new constant symbol for each sentence:

Lemma II.12.19 *Assume that Σ is a set of sentences of \mathcal{L} , κ is any ordinal, $\exists x\varphi_\alpha(x)$ is an existential sentence of \mathcal{L} for each $\alpha < \kappa$, and $\mathcal{L}_\kappa = \mathcal{L} \cup \{c_\alpha : \alpha < \kappa\}$, where the c_α are new constant symbols. Let $\Sigma_\kappa = \Sigma \cup \{\exists x\varphi_\alpha(x) \rightarrow \varphi(c_\alpha) : \alpha < \kappa\}$, and suppose that $\text{Con}_{+, \mathcal{L}}(\Sigma)$. Then $\text{Con}_{+, \mathcal{L}_\kappa}(\Sigma_\kappa)$.*

Proof. Induct on κ . The case $\kappa = 0$ is trivial, and the successor case is handled by Lemma II.12.18. For limit κ , use the fact that every formal proof from Σ_κ only uses finitely many sentences. \square

In this lemma, we could let the sentences $\exists x\varphi_\alpha(x)$ enumerate all the existential sentences of \mathcal{L} , but we could not claim that Σ_κ has witnesses in \mathcal{L}_κ , since there might be existential sentences of \mathcal{L}_κ without witnessing terms. However, if we repeat this procedure ω times, we can construct a set of sentences with witnesses. In this construction, all the witnessing terms turn out to be constant symbols:

Lemma II.12.20 *Assume that Σ is a set of sentences of \mathcal{L} and $\text{Con}_{+, \mathcal{L}}(\Sigma)$. Let $\kappa = \max(|\mathcal{L}|, \aleph_0)$. Then there exist Σ' and \mathcal{L}' such that*

1. $\mathcal{L}' = \mathcal{L} \cup \mathcal{C}$, where \mathcal{C} is a set of constant symbols.
2. $|\mathcal{L}'| = |\mathcal{C}| = \kappa$.
3. $\Sigma \subseteq \Sigma'$ and Σ' is a set of sentences of \mathcal{L}' .
4. $\text{Con}_{+, \mathcal{L}'}(\Sigma')$.
5. Σ' has witnesses in \mathcal{L}' .

Proof. Let $\mathcal{C} = \{c_\alpha^n : \alpha < \kappa \wedge n < \omega\}$; so (1) and (2) are clear. Let $\Sigma^0 = \Sigma$. Let $\mathcal{L}^n = \mathcal{L} \cup \{c_\alpha^j : \alpha < \kappa \wedge j < n\}$, so $\mathcal{L}^0 = \mathcal{L}$. We shall now get:

$$\begin{array}{ccccccc} \Sigma^0 & \subset & \Sigma^1 & \subset & \Sigma^2 & \subset & \dots\dots \\ \text{in} & & \text{in} & & \text{in} & & \dots\dots \\ \mathcal{L}^0 & \subset & \mathcal{L}^1 & \subset & \mathcal{L}^2 & \subset & \dots\dots \end{array}$$

There are exactly κ existential sentences of \mathcal{L}^n , so we list them all as $\{\exists x\varphi_\alpha^n(x) : \alpha < \kappa\}$. Then, let $\Sigma^{n+1} = \Sigma^n \cup \{\exists x\varphi_\alpha^n(x) \rightarrow \varphi_\alpha^n(c_\alpha^n) : \alpha < \kappa\}$. Let $\mathcal{L}' = \bigcup_{n < \omega} \mathcal{L}^n$ and let $\Sigma' = \bigcup_{n < \omega} \Sigma^n$. Each existential sentence of \mathcal{L}' is really in \mathcal{L}^n for some n , so (5) holds. By Lemma II.12.19 and induction, we have $\text{Con}_{\mathcal{L}^n}(\Sigma^n)$ for each n . Then (4) holds because every formal proof from Σ' only uses finitely many sentences. \square

As noted above (Lemma II.12.3), to prove the Completeness Theorem II.12.1, it is enough to prove the Main Lemma II.12.2, which we now do:

Proof of Lemma II.12.2 and the Completeness Theorem. Let Σ be a set of sentences of \mathcal{L} , and assume that $\text{Con}_{\mathcal{L}}(\Sigma)$. We must show that Σ has a model. To do this, we follow the three steps listed on page 117.

Step 1: Applying Lemma II.12.20, there are Σ' and \mathcal{L}' such that Σ' is a set of sentences of \mathcal{L}' , $\text{Con}_{\mathcal{L}'}(\Sigma')$, and Σ' has witnesses in \mathcal{L}' .

Step 2: Applying Lemma II.12.12, let $\Sigma^* \supseteq \Sigma'$ such that Σ^* is a set of sentences of \mathcal{L}' , and Σ^* is maximally (\vdash, \mathcal{L}') consistent. Since the definition (II.12.16) of “having witnesses” just requires that Σ' contain certain types of sentences, Σ^* also has witnesses in \mathcal{L}' .

Step 3: Now, the Herbrand model, $\mathfrak{A}' := \mathfrak{CT}(\mathcal{L}', \Sigma^*)$ is a model of Σ^* by Lemma II.12.17. Of course, \mathfrak{A}' is an \mathcal{L}' -structure. If \mathfrak{A} is the reduct, $\mathfrak{A}' \upharpoonright \mathcal{L}$, then \mathfrak{A} is an \mathcal{L} -structure and $\mathfrak{A} \models \Sigma$. \square

Now that we have proved the Completeness Theorem, we shall drop subscripts on our “Con”, and just write “Con(Σ)”, and for logical consequence, we just write, interchangeably, $\Sigma \models \varphi$ or $\Sigma \vdash \varphi$. As mentioned on page 108, we can also drop the ugly subscripts on our \vdash :

Lemma II.12.21 *Suppose that Σ is a set of sentences of \mathcal{L}_0 and ψ is a sentence of \mathcal{L}_0 and suppose that $\mathcal{L}_0 \subseteq \mathcal{L}_1$. Then $\Sigma \vdash_{\mathcal{L}_0} \varphi$ iff $\Sigma \vdash_{\mathcal{L}_1} \varphi$.*

Proof. Both are equivalent to $\Sigma \models \varphi$, which does not depend on \mathcal{L} (Lemma II.8.15). \square

The Compactness Theorem (page 97) is now easy:

Proof of Theorem II.7.14. As pointed out in Section II.7, the two parts of this theorem are equivalent, so we just prove the second part. We must show that $\Sigma \models \psi$, then there is a finite $\Delta \subseteq \Sigma$ such that $\Delta \models \psi$. But this is obvious if we replace “ \models ” by “ \vdash ”, since formal proofs are finite. \square

Since we were careful to keep track of the size of the model, we can now prove the Löwenheim–Skolem Theorem II.7.16. Our proof of Lemma II.12.2 shows the “downward” part of this theorem, that if Σ has a model, then Σ has a small model:

Lemma II.12.22 *Let Σ be a set of sentences of \mathcal{L} , and assume that $\text{Con}(\Sigma)$. Let $\kappa = \max(|\mathcal{L}|, \aleph_0)$. Then Σ has a model \mathfrak{A} with $|\mathfrak{A}| \leq \kappa$.*

Proof. Build a model for Σ as in proof of Lemma II.12.2. First, we got witnesses in an expanded \mathcal{L}' by applying Lemma II.12.20, which gives us $|\mathcal{L}'| = \kappa$. We then extended Σ' to a maximal Σ^* in the same \mathcal{L}' . We then formed the Herbrand model $\mathfrak{CT}(\mathcal{L}', \Sigma^*)$ whose universe was $\mathfrak{CT}(\mathcal{L}', \Sigma) = \mathfrak{CT}_0(\mathcal{L}')/\sim$. Now $\mathfrak{CT}_0(\mathcal{L}')$, the set of all closed terms, has size precisely $|\mathcal{L}'| = \kappa$, so that $|\mathfrak{CT}(\mathcal{L}', \Sigma)| \leq |\mathfrak{CT}_0(\mathcal{L}')| = \kappa$. \square

In this proof, it is possible that $|\mathfrak{CT}(\mathcal{L}', \Sigma)| < \kappa$. This cannot be avoided, since it is possible that Σ has only finite models. However, if Σ has infinite models, then we get models in any infinite size $\geq |\mathcal{L}|$ by the Löwenheim–Skolem Theorem (page 98), which we now prove:

Proof of Theorem II.7.16. Fix $\kappa \geq \max(|\mathcal{L}|, \aleph_0)$. We need to produce a model of Σ of size κ . Let $\mathcal{L}^* = \mathcal{L} \cup \mathcal{C}$, where $\mathcal{C} = \{c_\alpha : \alpha < \kappa\}$ is a set of κ new constant symbols; then $|\mathcal{L}^*| = \kappa$. Let $\Sigma^* = \Sigma \cup \{c_\alpha \neq c_\beta : \alpha < \beta < \kappa\}$. Clearly, every model of Σ^* has size at least κ , and Lemma II.12.22 implies that if $\text{Con}(\Sigma^*)$, then Σ^* has a model \mathfrak{A} with $|\mathfrak{A}| \leq \kappa$, and hence $|\mathfrak{A}| = \kappa$.

So, we are done if we can prove $\text{Con}(\Sigma^*)$. By the Compactness Theorem (II.7.14), $\text{Con}(\Sigma^*)$ follows if we can show $\text{Con}(\Delta)$ for every finite $\Delta \subseteq \Sigma^*$. Such a Δ consists of some sentences of Σ plus some sentences of the form $c_\alpha \neq c_\beta$ for $\alpha, \beta \in F$, where F is a finite subset of κ . Let \mathfrak{B} be any model of Σ with $|\mathfrak{B}| \geq |F|$. Then \mathfrak{B} is an \mathcal{L} -structure. Expand \mathfrak{B} to an \mathcal{L}^* -structure, \mathfrak{B}^* , by interpreting the c_α for $\alpha \in F$ to be distinct elements of B ; the c_α for $\alpha \notin F$ can be interpreted arbitrarily. Then $\mathfrak{B}^* \models \Delta$, so $\text{Con}(\Delta)$. \square

The notion of a maximal consistent set of sentences, as obtained in Lemma II.12.12, was of key importance in the proof of the Completeness Theorem. In general, these sets are obtained non-constructively, using the Axiom of Choice. However, there are some natural examples of maximal sets. First, if \mathfrak{A} is any structure for \mathcal{L} , then the theory of \mathfrak{A} , $\text{Th}(\mathfrak{A})$ (see Definition II.8.21), is maximal. Second, if Σ is complete (see Definition II.8.19), then the set of all \mathcal{L} -sentences φ such that $\Sigma \models \varphi$ is maximal. A number of complete theories occur naturally in algebra, as we shall explain in the next section.

Exercise II.12.23 Prove that the Σ' obtained in Lemma II.12.20 is actually a conservative extension of Σ ; that is, if θ is any sentence of \mathcal{L} , then $\Sigma \vdash \theta$ iff $\Sigma' \vdash \theta$.

Hint. If $\Sigma' \vdash \theta$ but $\Sigma \not\vdash \theta$, apply the proof of Lemma II.12.20 starting from $\Sigma \cup \{\neg\theta\}$ to obtain a contradiction. \square

II.13 More Model Theory

Now that we have the basics of model theory, we indicate briefly some applications to standard algebraic systems. We begin with the notions of *elementary equivalence*:

Definition II.13.1 If $\mathfrak{A}, \mathfrak{B}$ are structures for \mathcal{L} , then $\mathfrak{A} \equiv \mathfrak{B}$ ($\mathfrak{A}, \mathfrak{B}$ are elementarily equivalent) iff for all \mathcal{L} -sentences φ : $\mathfrak{A} \models \varphi$ iff $\mathfrak{B} \models \varphi$.

This is equivalent to saying that $\text{Th}(\mathfrak{A}) = \text{Th}(\mathfrak{B})$ (see Definition II.8.21). We can rephrase the notion of a complete set of axioms (Definition II.8.19) in terms of \equiv :

Lemma II.13.2 If Σ is a set of sentences of \mathcal{L} , then Σ is complete iff Σ is consistent and $\mathfrak{A} \equiv \mathfrak{B}$ whenever \mathfrak{A} and \mathfrak{B} are models of Σ .

If $\mathfrak{A}, \mathfrak{B}$ are isomorphic ($\mathfrak{A} \cong \mathfrak{B}$; see Definition II.8.18), then they are “essentially the same”, so not surprisingly, $\mathfrak{A} \equiv \mathfrak{B}$. Formally, this is proved by:

Lemma II.13.3 If $\mathfrak{A} \cong \mathfrak{B}$ then $\mathfrak{A} \equiv \mathfrak{B}$.

Proof. Let $\Phi : A \xrightarrow{1-1} B$ be an isomorphism. Then, show by induction on all formulas ψ that $\mathfrak{A} \models \psi[a_1, \dots, a_n]$ iff $\mathfrak{B} \models \psi[\Phi(a_1), \dots, \Phi(a_n)]$ for all $a_1, \dots, a_n \in A$. Then applying this with sentences, where $n = 0$, we get $\mathfrak{A} \equiv \mathfrak{B}$. \square

It is easy to give examples where $\mathfrak{A} \not\equiv \mathfrak{B}$. For example, the group $(\mathbb{Z}; +)$ is not elementarily equivalent to the group $(\mathbb{Q}; +)$ since the sentence $\forall x \exists y (y + y = x)$ is true in \mathbb{Q} but false in \mathbb{Z} . It is harder to give non-trivial *specific* examples where $\mathfrak{A} \equiv \mathfrak{B}$. The following two easy lemmas provides many *abstract* examples:

Lemma II.13.4 If \mathfrak{B} is infinite, \mathcal{L} is countable, and $\kappa \geq \aleph_0$, then there is a structure \mathfrak{A} for \mathcal{L} with $|\mathfrak{A}| = \kappa$ and $\mathfrak{A} \equiv \mathfrak{B}$.

Proof. Apply the Löwenheim–Skolem Theorem (II.7.16) with $\Sigma = \text{Th}(\mathfrak{B})$. \square

However, given *specific* structures, such as the groups $(\mathbb{Q}; +)$ and $(\mathbb{R}; +)$, it is not always easy to tell if they are elementarily equivalent. Here, we shall prove (see Lemma II.13.7 below) that $(\mathbb{Q}; +) \equiv (\mathbb{R}; +)$. To do this, we shall identify a “natural” set of axioms Σ such that $(\mathbb{Q}; +)$ and $(\mathbb{R}; +)$ are both models of Σ , and then prove Σ is complete, so that $(\mathbb{Q}; +) \equiv (\mathbb{R}; +)$ follows by Lemma II.13.2.

We shall return to this example after describing a method for proving a given Σ is complete. One well-known method is called *quantifier elimination*, one proves “ $\Sigma \models \varphi$ or $\Sigma \models \neg\varphi$ ” by induction on the logical complexity of φ ; this is described in model theory texts — e.g., [5, 24]. We shall describe another method, called the *Łoś – Vaught test* which is easier to apply in some cases. It involves the notion of *categoricity*, and provides a simple application of the Löwenheim–Skolem Theorem.

Definition II.13.5 Suppose that Σ is a set of sentences of \mathcal{L} and κ any cardinal. Then Σ is κ -categorical iff all models of Σ of size κ are isomorphic.

Theorem II.13.6 (Łoś – Vaught test) Let Σ be a set of sentences of \mathcal{L} , and assume:

1. Σ is consistent.
2. All models of Σ are infinite.
3. \mathcal{L} is countable.
4. Σ is κ -categorical for some infinite κ .

Then Σ is complete.

Proof. If Σ is not complete, then there is a sentence φ of \mathcal{L} such that $\Sigma \not\models \varphi$ and $\Sigma \not\models \neg\varphi$. Fix $\mathfrak{A}, \mathfrak{B}$ with $\mathfrak{A} \models \Sigma \cup \{\neg\varphi\}$ and $\mathfrak{B} \models \Sigma \cup \{\varphi\}$. Since $\mathfrak{A}, \mathfrak{B}$ are infinite (by (2)), apply Lemma II.13.4 to get $\mathfrak{A}', \mathfrak{B}'$ with $\mathfrak{A}' \equiv \mathfrak{A}$ and $\mathfrak{B}' \equiv \mathfrak{B}$ and $|\mathfrak{A}'| = |\mathfrak{B}'| = \kappa$. Then $\mathfrak{A}' \cong \mathfrak{B}'$ by (4), so $\mathfrak{A}' \equiv \mathfrak{B}'$ by Lemma II.13.3, but φ is true in \mathfrak{B}' and false in \mathfrak{A}' , a contradiction. \square

For a trivial example, let $\mathcal{L} = \emptyset$ and $\Sigma = \{\psi_n : 1 \leq n < \omega\}$ be the theory of infinite sets. Here, ψ_n says that there are n distinct elements; for example, the sentence ψ_3 is $\exists x, y, z [x \neq y \wedge y \neq z \wedge x \neq z]$. Now, structures are just sets, and every bijection is an isomorphism, so Σ is κ -categorical for all infinite κ . It follows that Σ is complete.

To apply model-theoretic methods such as the Łoś – Vaught test to more interesting algebraic theories, one must have a detailed knowledge of the algebra. For example, as claimed above, we show that $(\mathbb{Q}; +) \equiv (\mathbb{R}; +)$. These are both models of the theory Σ of torsion-free divisible abelian groups. To obtain the theory of divisible abelian groups, we add to the theory of abelian groups the axioms $\forall x \exists y [x = ny]$ whenever $0 < n \leq \omega$; here ny just abbreviates the term $y + y + \dots + y$ with n copies of y ; for example, $4y$ abbreviates $y + y + y + y$. Then we get Σ by adding the further axioms $\forall y [y \neq 0 \rightarrow ny \neq 0]$ for $0 < n \leq \omega$. If \mathcal{L} is just $\{+\}$, rather than $\{+, -, 0\}$, we can consider $x = 0$ to be an abbreviation for $x + x = x$. Now, $(\mathbb{Q}; +) \equiv (\mathbb{R}; +)$ follows from:

Lemma II.13.7 *The theory of torsion-free divisible abelian groups is κ -categorical for all $\kappa > \aleph_0$, and hence complete.*

Proof. Completeness follows from categoricity by the Łoś – Vaught test.

To prove categoricity, note that if G is a torsion-free divisible abelian group, then we may regard G as a vector space over the field \mathbb{Q} . Here if m/n is a “scalar” in \mathbb{Q} , where $m, n \in \mathbb{Z}$ and $n \neq 0$, we can define $(m/n)\vec{v}$ to be the unique $\vec{w} \in G$ such that $n\vec{w} = m\vec{v}$. It is easy to see that this multiplication makes G into a vector space. Then the dimension, $\dim(G)$, is the size of a basis. Note that since $|\mathbb{Q}| = \aleph_0$, we have $|G| = \max(\dim(G), \aleph_0)$. Thus, if G, H have the same uncountable size κ , they also have dimension κ , and hence are isomorphic. \square

A related example:

Exercise II.13.8 *Let p be a prime and let Σ be the theory of infinite abelian groups of exponent p (i.e., satisfying $\forall x [x^p = 1]$). Prove that Σ is κ -categorical for all $\kappa \geq \aleph_0$.*

Hint. View a model of Σ as a vector space over the finite field \mathbb{Z}_p . Note that it doesn't matter here whether we consider \mathcal{L} to be $\{\cdot\}$ or $\{\cdot, i, 1\}$ or $\{+\}$ or $\{+, -, 0\}$. \square

It is also known that the theory of algebraically closed fields of characteristic p (where p is either zero or a prime) is κ -categorical for all $\kappa \geq \aleph_1$ but not for $\kappa = \aleph_0$. Here, the isomorphism type of such a field K is determined by the size of a transcendence basis; that is a set B of mutually transcendental elements such that all elements of K are algebraic over B . If $|B| \leq \aleph_0$, then $|K| = \aleph_0$ (yielding \aleph_0 countable models), while if $|B| = \kappa \geq \aleph_1$, then $|K| = \kappa$.

Some further remarks on categoricity: Suppose that Σ in a countable \mathcal{L} is consistent and has only infinite models (this is (1)(2)(3) of Theorem II.13.6). By a theorem of Morley, Σ is κ -categorical for *some* $\kappa > \aleph_0$ iff Σ is κ -categorical for *all* $\kappa > \aleph_0$; such Σ are usually called \aleph_1 -categorical. The previous lemma and exercise show that an \aleph_1 -categorical Σ may or may not be \aleph_0 -categorical. Furthermore, by a theorem of Baldwin and Lachlan, if such a Σ fails to be \aleph_0 -categorical, then it has exactly \aleph_0 non-isomorphic countable models, arranged in a chain of type $\omega + 1$; for example, in Lemma II.13.7, the countable models consist of the groups of dimension α for $\alpha \leq \omega$. Proofs of the Morley and Baldwin–Lachlan theorems are in the texts [5, 24].

An \aleph_0 -categorical theory need not be \aleph_1 -categorical, as we see from the following example due to Cantor. A total order $<$ on a set A is called *dense* (or dense in itself) iff it satisfies $\forall x, y [x < y \rightarrow \exists z [x < z < y]]$. Four examples of dense total orders are given by the intervals $(0, 1)$, $(0, 1]$, $[0, 1)$, $[0, 1]$ in \mathbb{R} . These models are not elementarily equivalent, so that the theory of dense total orders is not complete. However, if we specify the existence or nonexistence of a least and greatest element, then we get four complete theories. Define $\Delta^{(\cdot)}$ to be the axioms for dense total orders plus the statements $\exists x \forall y [x \leq y]$ and $\neg \exists x \forall y [y \leq x]$. Likewise, define $\Delta^{(\cdot)}$, $\Delta^{[}$, and $\Delta^{]}$ in the obvious way.

Exercise II.13.9 *Each of the four theories, $\Delta^{(\cdot)}$, $\Delta^{(\cdot)}$, $\Delta^{[}$, and $\Delta^{]}$ in $\mathcal{L} = \{<\}$ is \aleph_0 -categorical and not \aleph_1 -categorical.*

Hint. To prove, say, that $\Delta^{(\cdot)}$ is \aleph_0 -categorical, fix countable models A, B , and construct an isomorphism f between them. As a set of ordered pairs, $f \subset A \times B$, and we'll have $f = \bigcup_n f_n$, where f_n is a set of n ordered pairs, and f_{n+1} is f_n plus one more pair. So, $f_0 = \emptyset$. Make sure that each f_n is order-preserving where it is defined. Before starting, list A and B as $\{a_i : i < \omega\}$ and $\{b_i : i < \omega\}$. Make sure that $a_i \in \text{dom}(f_{2i+1})$ and $b_i \in \text{ran}(f_{2i+2})$.

$\Delta^{(\cdot)}$ is not 2^{\aleph_0} -categorical because \mathbb{R} is not isomorphic to the irrationals. To prove that it isn't \aleph_1 -categorical without using Morley's Theorem (or the *CH*), note that the "long rationals" (consisting of ω_1 copies of \mathbb{Q}) is not isomorphic to its reverse order. \square

It follows by the Łoś – Vaught test that these four theories are complete. So, if φ is a sentence of \mathcal{L} , either $\Delta^{(\cdot)} \vdash \varphi$ or $\Delta^{(\cdot)} \vdash \neg \varphi$, and not both. Note that we now have an algorithm to decide whether or not $\Delta^{(\cdot)} \vdash \varphi$ (equivalently, whether or not φ is true in

\mathbb{R} , \mathbb{Q} , etc.), but it is horribly inefficient. Namely, our algorithm lists all formal proofs from $\Delta^{(\cdot)}$ as $\Pi_0, \Pi_1, \Pi_2, \dots$, and stops when it finds a proof of either φ or $\neg\varphi$; such a proof will be found eventually because $\Delta^{(\cdot)}$ is complete. There is no problem *in theory* with listing all formal proofs in type ω since we can list all of $R(\omega)$ in type ω (Exercise I.14.12). In practice, this algorithm is not feasible because there are too many formal proofs. A feasible algorithm is provided by the method of *quantifier elimination*, which work by induction on the complexity of φ . Roughly, it uses facts about dense orders to reduce a sentence φ to an equivalent sentence with one fewer quantifier. This process is repeated until φ is reduced to T or F ; see [5, 24] for details.

Many theories are complete but not κ -categorical for any κ , so the Łoś – Vaught test does not apply. An example of such a theory is real-closed fields. Here, $\mathcal{L} = \{0, 1, +, \cdot, -, i, <\}$. A real-closed field is an ordered field in which every positive element has a square root and every polynomial of odd degree has a root; so \mathbb{R} is a model but \mathbb{Q} is not. Tarski showed that the theory is complete by quantifier elimination.

A much simpler example of a complete theory which is not κ -categorical for any κ is given by the following example of Ehrenfeucht:

Exercise II.13.10 *Let \mathcal{L} contain $<$ plus constant symbols c_n for $n \in \omega$. Let Σ be the axioms for dense total order without endpoints ($\Delta^{(\cdot)}$ above), together with the axioms $c_n < c_{n+1}$ for each $n \in \omega$. Prove that Σ is complete and not κ -categorical for any κ . Furthermore, show that Σ has precisely three models of size \aleph_0 .*

Hint. For completeness, note that the reduct of Σ to $\{<, c_0, c_1, \dots, c_k\}$ is \aleph_0 -categorical for each k . For the three countable models, identify a countable model with \mathbb{Q} . Then $\sup_n c_n$ can be either ∞ or in \mathbb{Q} or in $\mathbb{R} \setminus \mathbb{Q}$. \square

Now, let Σ be any complete theory in a countable \mathcal{L} , and let $n(\Sigma)$ be the number of countably infinite models of Σ (up to isomorphism). Ehrenfeucht also gave examples showing that $n(\Sigma)$ can be any finite number except 2, and Vaught showed that $n(\Sigma)$ can never equal 2. Some examples where $n(\Sigma) = \aleph_0$ were described above. It is easy to see that $n(\Sigma) \leq 2^{\aleph_0}$, and there are many examples of $n(\Sigma) = 2^{\aleph_0}$, such as:

Exercise II.13.11 *Let \mathcal{L} contain $<$ plus constant symbols c_q for $q \in \mathbb{Q}$. Let Σ be the axioms for dense total order without endpoints ($\Delta^{(\cdot)}$ above), together with the axioms $c_p < c_q$ for each $p, q \in \mathbb{Q}$ such that $p < q$. Prove that Σ is complete and that Σ has 2^{\aleph_0} models of size \aleph_0 .*

Vaught's Conjecture is the statement that for all such Σ , $n(\Sigma) \geq \aleph_1$ implies $n(\Sigma) = 2^{\aleph_0}$. This is trivial under *CH*, but it is unknown whether the conjecture is provable in *ZFC*. Morley showed that $n(\Sigma) \geq \aleph_2$ does imply $n(\Sigma) = 2^{\aleph_0}$; of course, this is vacuous unless *CH* is false. See [5, 24] for more details.

II.14 Equational Varieties and Horn Theories

In our proof of the Completeness Theorem, we built a model out of the terms of the language. The use of symbolic expressions as objects has its roots in the algebra of the 1800s. One familiar example is the ring $F[x]$ of polynomials over a field F , and the use of this ring for obtaining algebraic extensions of F . Another familiar example is the notion of a group given by generators and relations; a special case of this, with the empty set of relations, was described on page 120. These examples are really special cases of the Herbrand model $\mathfrak{CT}(\mathcal{L}, \Sigma)$. We now single out the features of these examples which allow us to use $\mathfrak{CT}(\mathcal{L}, \Sigma)$ without first adding witnesses and extending to a maximal consistent set of axioms.

Definition II.14.1 *A positive literal is an atomic formula. A negative literal is the negation of an atomic formula. A Horn clause is a disjunction of one or more literals such that no more than one of them is positive. A set Σ of sentences is universal Horn iff every sentence in Σ is a closure of a Horn clause.*

These are named after Alfred Horn (1918 – 1988), who first studied this kind of sentence. Further results on Horn sentences are in Chang–Keisler [5].

A special case of universal Horn theories is:

Definition II.14.2 *A set Σ of sentences of \mathcal{L} is an equational variety iff \mathcal{L} has no predicate symbols every sentence in Σ is a closure of an equation (of form $\tau = \sigma$).*

For example, the group axioms GP , written in $\mathcal{L} = \{\cdot, i, 1\}$ as on page 88 form an equational variety, and hence are universal Horn. Likewise, the natural axioms for rings form an equational variety, as do the additional axioms for some common varieties of rings, such as commutative rings, rings with unity (1), etc.

The axioms for strict partial order (see Definition I.7.2) are logically equivalent to a universal Horn set. Irreflexivity, $\forall x \neg R(x, x)$, is the closure of a Horn clause with zero positive literals and one negative literal. Transitivity is equivalent to the statement $\forall xyz[\neg R(x, y) \vee \neg R(y, z) \vee R(x, z)]$, which is the closure of a Horn clause with one positive literal and two negative literals.

However, the axioms for total order include the trichotomy axiom, which has the disjunction of three positive literals, and hence is not Horn. Also, the axioms for fields (see Example II.8.23) includes the axiom that non-zero elements have inverses. Written as a disjunction, $\forall x[x = 0 \vee x \cdot i(x) = 1]$, it has two positive literals. Of course, conceivably these axioms could be rewritten in an equivalent way just using universal Horn sentences, but this is refuted by:

Exercise II.14.3 *If $\mathfrak{A}, \mathfrak{B}$ are structures for \mathcal{L} , then define $\mathfrak{A} \times \mathfrak{B}$ to be the structure with universe $A \times B$, with functions and relations defined coordinatewise in the natural way. Prove that if Σ is universal Horn, $\mathfrak{A} \models \Sigma$, and $\mathfrak{B} \models \Sigma$, then $\mathfrak{A} \times \mathfrak{B} \models \Sigma$.*

Since the product of fields is never a field and the product of total orders with more than one element is never a total order, there is no way to axiomatize fields or total orders by universal Horn sentences.

Theories involving an order are trivially not equational varieties, since \mathcal{L} contains a predicate symbol, but there are natural examples universal Horn theories involving just functions which are not equational varieties. For example, let Σ be the group axioms GP , written in $\mathcal{L} = \{\cdot, i, 1\}$ plus the additional axiom $\forall x[x^2 = 1 \rightarrow x = 1]$ (no element has order 2). Then Σ is not an equational variety, but, expressing the implication as a disjunction, we see that Σ is universal Horn. Of course, we again have the possibility that there is some equivalent way of axiomatizing Σ which is equational. To see that this is in fact not possible, note, switching to additive notation for groups, that $\mathbb{Z} \models \Sigma$, but its quotient $\mathbb{Z}_2 = \mathbb{Z}/2\mathbb{Z}$ is not a model for Σ ; then, apply

Exercise II.14.4 *Prove that if Σ is an equational variety and $\mathfrak{A} \models \Sigma$, then every homomorphic image of \mathfrak{A} is a model for Σ .*

Observe that since \mathcal{L} has no predicate symbols, the notion of homomorphic image is defined almost verbatim as it is defined for groups and rings.

We now consider closed term models.

Exercise II.14.5 *Assume that Σ is universal Horn and consistent, and that $\mathcal{F}_0 \neq \emptyset$. Prove that $\mathfrak{CT}(\mathcal{L}, \Sigma) \models \Sigma$.*

That is, in this case, we don't need witnesses or maximality. Note that this formalizes in model-theoretic language the standard algebraic construction of a structure presented by generators and relations. For example, in group theory, with $\mathcal{L} = \{\cdot, i, 1, a, b\}$, and $\Sigma = GP \cup \{a^2 = b^3 = 1, aba = b^2\}$, our $\mathfrak{CT}(\mathcal{L}, \Sigma)$ is the group with 2 generators modulo these relations (which is the 6-element nonabelian group). It is easy to see that our description of this as $\mathfrak{CT}(\mathcal{L}, \Sigma)$ is equivalent to the description in group theory texts (the free group with generators a, b modulo the normal subgroup generated by $\{a^2, b^3, abab^{-2}\}$).

Now, consider polynomials and field extensions. For example, $\mathbb{C}[z, w]$ denotes the ring of polynomials in two "variables" with complex coefficients. One such polynomial is $2\pi w z^3 + 6zw + 2iw^2$. Everyone recognizes that $2\pi w z^3 + 6zw + 2iw^2$ and $2w(\pi z^3 + 3z + iw)$ are two ways of writing the "same" polynomial, so there is implicitly an equivalence relation between polynomials here.

To express this notion in our formalism, let $\mathcal{L} = \{0, 1, +, \cdot, -\}$ be the language of ring theory. Let Σ be the axioms for commutative rings with unity; specifically, axioms (1)(2)(3)(4) in Example II.8.23. Now, \mathcal{L} does not contain symbols for i or π , or for 2^{\aleph_0} other elements of \mathbb{C} which are legal coefficients for polynomials over \mathbb{C} . To express these, we need to add some constant symbols:

Definition II.14.6 *For any lexicon \mathcal{L} and any set A , \mathcal{L}_A denotes \mathcal{L} augmented by a set of new constant symbols, $\{c_a : a \in A\}$.*

In practice, the set-theoretic identity of the c_a is never important, as long as they are different from each other and from the other symbols in \mathcal{L} and the logical symbols. Often, if it does not cause confusion, we use a itself; this is common, for example, when discussing the *diagram*:

Definition II.14.7 *For any lexicon \mathcal{L} and any structure \mathfrak{A} for \mathcal{L} , the natural expansion of \mathfrak{A} is the expansion of \mathfrak{A} to the \mathcal{L}_A structure called \mathfrak{A}_A obtained by interpreting the symbol c_a as the element $a \in A$. Then the elementary diagram, $\text{eDiag}(\mathfrak{A})$, is $\text{Th}(\mathfrak{A}_A)$; that is, the set of all \mathcal{L}_A -sentences true in \mathfrak{A}_A , and the diagram, $\text{Diag}(\mathfrak{A})$, is the set of all quantifier-free sentences in $\text{eDiag}(\mathfrak{A})$.*

For example, let $\mathcal{L} = \{0, 1, +, \cdot, -\}$ and let \mathfrak{A} be \mathbb{C} , with the standard interpretation of the symbols of \mathcal{L} . Then $c_\pi \neq c_2 \cdot c_2$ is in $\text{Diag}(\mathfrak{A})$, whereas $\exists x [c_\pi = x \cdot x]$ is in $\text{eDiag}(\mathfrak{A})$, but not in $\text{Diag}(\mathfrak{A})$. In most applications, we would write these sentences simply as $\pi \neq 2 \cdot 2$ and $\exists x [\pi = x \cdot x]$. There is a slight danger of confusion here, since one should really distinguish between the complex number 1 and the symbol 1 of \mathcal{L} , so that $c_1 + c_1 = c_2$ is not really the same logical sentence as $c_1 + 1 = c_2$; but since $[c_1 = 1] \in \text{Diag}(\mathfrak{A})$, this abuse of notation rarely causes any confusion in practice; we simply say that $1 + 1 = 2$ is true in \mathbb{C} .

Now, to form the coefficients of $\mathbb{C}[z, w]$, we use $\mathcal{L}_{\mathbb{C}}$. But also, the z, w are really constant symbols, not variables, so that the polynomials are really closed terms:

Definition II.14.8 *Let Σ be the axioms for commutative rings with unity, in $\mathcal{L} = \{0, 1, +, \cdot, -\}$. Let $\mathfrak{R} = (R; 0, 1, +, \cdot, -) \models \Sigma$, and fix a positive integer n . Then the ring of polynomials $R[z_1, \dots, z_n]$ is $\mathfrak{CT}(\mathcal{L}_R \cup \{z_1, \dots, z_n\}, \Sigma \cup \text{Diag}(\mathfrak{R}))$, where z_1, \dots, z_n are new constant symbols.*

Note that $\mathfrak{CT}(\mathcal{L}_R \cup \{z_1, \dots, z_n\}, \Sigma \cup \text{Diag}(\mathfrak{R})) \models \Sigma$ by Exercise II.14.5. Σ remains unchanged, in the original \mathcal{L} .

A polynomial is really an equivalence class. To return to our example with $\mathbb{C}[z, w]$, the terms $2\pi w z^3 + 6z w + 2i w^2$ and $2w(\pi z^3 + 3z + i w)$ are different closed terms, but they are provably equal, so they denote the same equivalence class in the closed term model. The proof uses the commutative ring axioms Σ (such as the distributive law), together with facts from $\text{Diag}(\mathfrak{C})$, such as $2 \cdot 3 = 6$. Of course, to be completely pedantic, the $2, 3, 6, \pi, i$ should really be $c_2, c_3, c_6, c_\pi, c_i$; also, “ $6z w$ ” could mean either $c_6 \cdot (z \cdot w)$ or $(c_6 \cdot z) \cdot w$, but these two expressions are provably equal using the associative law from Σ .

Regarding $\Sigma \cup \text{Diag}(\mathfrak{R})$, note that we are using the axioms of Σ (including the ones with quantifiers), along with the (quantifier-free) sentences from $\text{Diag}(\mathfrak{R})$. We do not want to use the larger set $\text{eDiag}(\mathfrak{R})$ in forming the polynomial ring. For example, in algebra, the polynomials z and z^3 are *always* different polynomials, although they may happen to denote the same polynomial function over a particular ring; that is, the

sentence $\forall x[x = x^3]$ may be in $\text{eDiag}(\mathfrak{A})$ (this happens when \mathfrak{A} is the 2-element field or the 3-element field).

The elementary diagram, $\text{eDiag}(\mathfrak{A})$, is important in the discussion of elementary submodels; see Section II.16.

The abstract study of notions such as varieties, free algebras, and quotients is called *universal algebra*; it is halfway between model theory and conventional algebra. For more on this subject, see the text of Burris and Sankappanavar [3].

II.15 Extensions by Definitions

We discussed two common ways of presenting the axioms for groups. One, using $\mathcal{L} = \{\cdot\}$ in Section 0.4 had axioms $GP = \{\gamma_1, \gamma_2\}$:

$$\begin{aligned} \gamma_1. & \forall xyz[x \cdot (y \cdot z) = (x \cdot y) \cdot z] \\ \gamma_2. & \exists u[\forall x[x \cdot u = u \cdot x = x] \wedge \forall x \exists y[x \cdot y = y \cdot x = u]] \end{aligned}$$

The other, using $\mathcal{L}' = \{\cdot, i, 1\}$ in Section II.5, had axioms $GP' = \{\gamma_1, \gamma_{2,1}, \gamma_{2,2}\}$:

$$\begin{aligned} \gamma_1. & \forall xyz[x \cdot (y \cdot z) = (x \cdot y) \cdot z] \\ \gamma_{2,1}. & \forall x[x \cdot 1 = 1 \cdot x = x] \\ \gamma_{2,2}. & \forall x[x \cdot i(x) = i(x) \cdot x = 1] \end{aligned}$$

Although GP' is an equational variety while GP is not, there is a sense in which GP and GP' are equivalent; we make this sense precise here; we say that GP' is an *extension by definitions* of GP .

This discussion is important especially when we have a theory such as *ZFC*, where the development of that theory involves thousands of definitions, not just one or two. That is, any mathematical terminology beyond \in and $=$ defined anywhere is an extension by definitions over *ZFC*.

Definition II.15.1 Assume that $\mathcal{L} \subseteq \mathcal{L}'$ and Σ is a set of sentences of \mathcal{L} .

If $p \in \mathcal{L}' \setminus \mathcal{L}$ is an n -ary predicate symbol, then a definition of p over \mathcal{L}, Σ is a sentence of the form $\forall x_1, \dots, x_n [p(x_1, \dots, x_n) \leftrightarrow \theta(x_1, \dots, x_n)]$, where θ is a formula of \mathcal{L} .

If $f \in \mathcal{L}' \setminus \mathcal{L}$ is an n -ary function symbol, then a definition of f over \mathcal{L}, Σ is a sentence of the form $\forall x_1, \dots, x_n [\theta(x_1, \dots, x_n, f(x_1, \dots, x_n))]$, where θ is a formula of \mathcal{L} and $\Sigma \vdash \forall x_1, \dots, x_n \exists! y \theta(x_1, \dots, x_n, y)$.

A set of sentences Σ' of \mathcal{L}' is an extension by definitions of Σ iff $\Sigma' = \Sigma \cup \Delta$, where $\Delta = \{\delta_s : s \in \mathcal{L}' \setminus \mathcal{L}\}$ and each δ_s is a definition of s over \mathcal{L}, Σ .

Note that in the case of predicate symbols, Σ is irrelevant. Also, the $n = 0$ case of the above definition frequently occurs; here the x_1, \dots, x_n is missing. A 0-ary function symbol is a constant symbol. For example, in developing set theory with $\mathcal{L} = \{\in\}$, we introduced (in Section I.6) the constant \emptyset by letting θ_\emptyset be the formula $\text{emp}(y)$; that is,

$\forall z(z \notin y)$. If Σ denotes ZF (or, just the axioms of Extensionality and Comprehension), then we proved from Σ that $\exists!y \theta_\emptyset(y)$. Then $\Sigma' = \Sigma \cup \theta(\emptyset)$ in $\{\in, \emptyset\}$ is an extension of Σ by definitions. A 0-ary predicate symbol is a proposition letter. An example of this in set theory is CH ; so θ_{CH} is a sentence just using \in and $=$ which is equivalent to CH . Of course, we have not come close to writing such a sentence. It is better to think of CH as arising from a chain of extensions by definitions, as we describe below.

For the group example, with $\mathcal{L} = \{\cdot\}$ and $\mathcal{L}' = \{\cdot, i, 1\}$, we note that GP proves that the identity element and the inverses, which exist by γ_2 , are in fact unique. We may then, as suggested by $\gamma_{2,1}$, let $\theta_1(y)$ be $\forall x[x \cdot y = y \cdot x = x]$; although $\theta_1(y)$ could also be just $y \cdot y = y$. To define inverse using \cdot , but not 1, we could let $\theta_i(x, y)$ by $y \cdot (x \cdot x) = x$. Then GP' is GP plus the axioms $\theta_1(1)$ and $\forall x \theta_i(x, i(x))$. This is easily seen to be equivalent to the axioms $\{\gamma_1, \gamma_{2,1}, \gamma_{2,2}\}$.

The idea that extensions by definitions add nothing essentially new is made precise by:

Theorem II.15.2 *Assume that $\mathcal{L} \subseteq \mathcal{L}'$, Σ is a set of sentences of \mathcal{L} , and Σ' in \mathcal{L}' is an extension by definitions of Σ in the sense of Definition II.15.1. Let $\forall\forall\chi$ denote some universal closure of χ (see Definition II.5.6 and Lemma II.8.3). Then*

1. *If φ is any sentence of \mathcal{L} , then $\Sigma \vdash \varphi$ iff $\Sigma' \vdash \varphi$.*
2. *If φ is any formula of \mathcal{L}' , then there is a formula $\widehat{\varphi}$ of \mathcal{L} with the same free variables such that $\Sigma' \vdash \forall\forall[\widehat{\varphi} \leftrightarrow \varphi]$.*
3. *If τ is any term of \mathcal{L}' , then there is a formula $\zeta_\tau(y)$ of \mathcal{L} using the same variables as τ plus a new variable y such that $\Sigma \vdash \forall\forall \exists!y \zeta_\tau(y)$ and $\Sigma' \vdash \forall\forall \zeta_\tau(\tau)$.*

Item (1) alone says that Σ' is a *conservative extension* of Σ (see also Exercise II.12.23). In the case of groups, for example, say we are proving cancellation: $\forall u, v, w [u \cdot w = v \cdot w \rightarrow u = v]$. It might be somewhat easier to do this within GP' , since we just multiply on the right by w^{-1} , but then (1) of the theorem says that we may find such a proof within GP . By items (2)(3), Σ' has no new expressive power — anything that we can express using \mathcal{L}' can be expressed using just \mathcal{L} . This is of importance in developing ZFC . We frequently apply the Comprehension Axiom to form $\{x \in w : \varphi(x)\}$, where φ involves various defined notions. However, in the original statement of the axiom (see the discussion in Section I.6), the formulas used only \in and $=$. This application of Comprehension is legitimate because we may replace φ by some equivalent $\widehat{\varphi}$ which uses only \in and $=$.

Proof of Theorem II.15.2. For (1): $\Sigma \subseteq \Sigma'$, so $\Sigma \vdash \varphi$ implies $\Sigma' \vdash \varphi$. Now assume that we had $\Sigma' \vdash \varphi$ but $\Sigma \not\vdash \varphi$, where φ is a sentence of \mathcal{L} . Applying the Completeness Theorem, let \mathfrak{A} be a structure for \mathcal{L} such that $\mathfrak{A} \models \Sigma$ and $\mathfrak{A} \models \neg\varphi$. Now, we may expand \mathfrak{A} to a structure for \mathcal{L}' , where for each symbol $s \in \mathcal{L}' \setminus \mathcal{L}$, we define $s_{\mathfrak{A}'}$ so that its definition δ_s is true. But then $\mathfrak{A}' \models \Sigma'$ and $\mathfrak{A}' \models \neg\varphi$, contradicting $\Sigma' \vdash \varphi$.

For (3), we induct on τ . If τ is just a variable, x , let $\zeta_\tau(y)$ be $y = x$. If τ is $f(\sigma_1, \dots, \sigma_n)$, where $n \geq 0$ and $f \in \mathcal{L}' \setminus \mathcal{L}$, let $\zeta_\tau(y)$ be

$$\exists z_1, \dots, z_n [\delta_f(z_1, \dots, z_n, y) \wedge \zeta_{\sigma_1}(z_1) \wedge \dots \wedge \zeta_{\sigma_n}(z_n)] \cdot$$

When $f \in \mathcal{L}$, replace the “ $\delta_f(z_1, \dots, z_n, y)$ ” by “ $f(z_1, \dots, z_n) = y$ ”.

For (2), we induct on φ , but the induction is essentially trivial except for the basis, where φ is $p(\sigma_1, \dots, \sigma_n)$, in which case we use (3). If $p \in \mathcal{L}' \setminus \mathcal{L}$, let $\widehat{\varphi}$ be

$$\exists z_1, \dots, z_n [\delta_p(z_1, \dots, z_n) \wedge \zeta_{\sigma_1}(z_1) \wedge \dots \wedge \zeta_{\sigma_n}(z_n)] .$$

When $p \in \mathcal{L}$, replace the “ $\delta_p(z_1, \dots, z_n)$ ” by “ $p(z_1, \dots, z_n)$ ”; likewise when p is =. \square

In developing an axiomatic theory, we often have a chain of extensions by definitions. For example, in group theory, we frequently use the commutator, defined by $[x_1, x_2] = x_1^{-1}x_2^{-1}x_1x_2$. Of course, the “[x_1, x_2]” is just syntactic sugar for a 2-ary function symbol which we could have written $c(x_1, x_2)$. But now we have $\mathcal{L} = \{\cdot\}$ and $\mathcal{L}' = \{\cdot, i, 1\}$ and $\mathcal{L}'' = \{\cdot, i, 1, c\}$, and $GP'' = GP' \cup \{\theta_c\}$ in \mathcal{L}'' is a definitional extension of GP' , which in turn was a definitional extension of GP . The next lemma says that such chains of definitional extensions can always be obtained in one step. We also remark that the introduction of the new function $c(x_1, x_2)$ as a composition of old functions is really just a special case of what is described by Definition II.15.1. Here, $\theta_c(x_1, x_2, y)$ is just $y = x_1^{-1}x_2^{-1}x_1x_2$, and in this case the required $\forall x_1, x_2 \exists! y \theta_c(x_1, x_2, y)$ is logically valid, and hence trivially provable.

Lemma II.15.3 *Assume that $\mathcal{L} \subseteq \mathcal{L}' \subseteq \mathcal{L}''$, Σ is a set of sentences of \mathcal{L} , Σ' is a set of sentences of \mathcal{L}' , and Σ'' is a set of sentences of \mathcal{L}'' . Assume that Σ' is an extension by definitions of Σ and Σ'' is an extension by definitions of Σ' . Then Σ'' is equivalent to an extension by definitions of Σ .*

Proof. Applying Definition II.15.1, $\Sigma' = \Sigma \cup \Delta$ and $\Sigma'' = \Sigma \cup \Delta \cup \Delta'$, where Δ contains a definition for each symbol in $\mathcal{L}' \setminus \mathcal{L}$, while Δ' contains a definition for each symbol in $\mathcal{L}'' \setminus \mathcal{L}'$. So, Δ is in \mathcal{L} , while $\Delta' = \{\delta_s : s \in \mathcal{L}'' \setminus \mathcal{L}'\}$ is a set of sentences in \mathcal{L}' . We shall define $\widetilde{\Delta}' = \{\widetilde{\delta}_s : s \in \mathcal{L}'' \setminus \mathcal{L}'\}$, where $\widetilde{\delta}_s$ in \mathcal{L} is obtained as follows:

If $p \in \mathcal{L}'' \setminus \mathcal{L}'$ is an n -ary predicate symbol, then δ_p is some sentence of the form $\forall x_1, \dots, x_n [p(x_1, \dots, x_n) \leftrightarrow \theta(x_1, \dots, x_n)]$, where θ is a formula of \mathcal{L}' . Apply Theorem II.15.2 to get a formula $\widehat{\theta}$ of \mathcal{L} such that $\Sigma' \vdash \forall \forall [\widehat{\theta} \leftrightarrow \theta]$, and let $\widetilde{\delta}_p$ be the sentence $\forall x_1, \dots, x_n [p(x_1, \dots, x_n) \leftrightarrow \widehat{\theta}(x_1, \dots, x_n)]$.

Likewise, if $f \in \mathcal{L}'' \setminus \mathcal{L}'$ is an n -ary predicate symbol, we obtain $\widetilde{\delta}_f$ by replacing the θ occurring in the δ_f by a suitable $\widehat{\theta}$.

Then $\Sigma \cup \Delta \cup \widetilde{\Delta}'$ is an extension of Σ by definitions and is equivalent to Σ'' . \square

For the above group example, it is *slightly* more elegant to take $\mathcal{L} = \{\cdot\}$ and $\mathcal{L}'' = \{\cdot, i, 1\}$, with $\mathcal{L}' = \{\cdot, 1\}$. Then $\theta_1(y)$ is as described above, but now $\theta_i(x, y)$ can simply be the \mathcal{L}' formula $x \cdot y = 1$.

Lemma II.15.3 is much more important when developing a theory such as *ZFC*. For example, in defining the proposition letter *CH*, we would express θ_{CH} (formalizing

Definition I.13.8) not in the original $\mathcal{L} = \{\in\}$, but in some \mathcal{L}' which is itself obtained by a long chain of definitions.

The notion of “equational variety” can change if we pass to a definitional extension. For example, the theory of groups is an equational variety if it is expressed using $\{\cdot, i, 1\}$, but not if it is expressed using $\{\cdot\}$. Observe that no definitional extension of the theory of fields can be an equational variety, or even universal Horn, since we still would not have closure under products (see Exercise II.14.3).

Exercise II.15.4 Show that the theory of lattices is an equational variety, using $\mathcal{L} = \{\vee, \wedge\}$.

Hint. Often, a lattice is defined to be a partial order $(A; <)$ in which every two elements x, y have a least upper bound $x \vee y$ and a greatest lower bound $x \wedge y$; so this way, \vee and \wedge are defined notions. To axiomatize the lattice with $\mathcal{L} = \{\vee, \wedge\}$, use the equations:

$$\begin{array}{ll} x \vee (y \vee z) = (x \vee y) \vee z & x \wedge (y \wedge z) = (x \wedge y) \wedge z \\ x \vee y = y \vee x & x \wedge y = y \wedge x \\ x \vee x = x & x \wedge x = x \\ x \wedge (x \vee y) = x & x \vee (x \wedge y) = x \end{array}$$

Then define $x \leq y$ to be $y = x \vee y$; this is equivalent to $x = x \wedge y$ (as one should check from the axioms); $x < y$ means $x \leq y$ & $x \neq y$. Of course, one has to check that $<$, as defined in this way, is really a partial order with \vee and \wedge the corresponding least upper bound and greatest lower bound functions. \square

Note that a total order is now a special kind of lattice (in which $x \wedge y$ is always either x or y); there is no way to express this using universal Horn sentences, since closure under products fails (see Exercise II.14.3).

II.16 Elementary Submodels

Recall (Definition II.8.17) the notion of *substructure*: $\mathfrak{A} \subseteq \mathfrak{B}$ means that $A \subseteq B$ and the functions and predicates of \mathfrak{A} are the restrictions of the corresponding functions and predicates of \mathfrak{B} . A stronger notion, *elementary substructure*, requires the same to hold for all definable properties:

Definition II.16.1 Let \mathfrak{A} and \mathfrak{B} be structures for \mathcal{L} with $\mathfrak{A} \subseteq \mathfrak{B}$. If φ is a formula of \mathcal{L} , then $\mathfrak{A} \preceq_{\varphi} \mathfrak{B}$ means that $\mathfrak{A} \models \varphi[\sigma]$ iff $\mathfrak{B} \models \varphi[\sigma]$ for all assignments σ into A . $\mathfrak{A} \preceq \mathfrak{B}$ (elementary substructure) means that $\mathfrak{A} \preceq_{\varphi} \mathfrak{B}$ for all formulas φ of \mathcal{L} .

Lemma II.16.2 If $\mathfrak{A} \subseteq \mathfrak{B}$, then $\mathfrak{A} \preceq_{\varphi} \mathfrak{B}$ whenever φ is quantifier-free.

For example, let $\mathcal{L} = \{<\}$, and let $A = [0, 1] \subseteq B = [0, 2] \subset \mathbb{R}$, and let $\mathfrak{A}, \mathfrak{B}$ use the natural interpretation of $<$. Then $\mathfrak{A} \cong \mathfrak{B}$, so that $\mathfrak{A} \equiv \mathfrak{B}$ (elementarily equivalent – see Definition II.13.1), so that $\mathfrak{A} \preceq_{\varphi} \mathfrak{B}$ for all sentences φ . However, $\mathfrak{A} \not\preceq \mathfrak{B}$ because $\mathfrak{A} \not\preceq_{\varphi} \mathfrak{B}$ where $\varphi(x)$ is the formula $\exists y [x < y]$, since $\mathfrak{B} \models \varphi[1]$ and $\mathfrak{A} \models \neg\varphi[1]$.

It is true that $(0, 1) \preceq (0, 2)$ (with the same \mathcal{L}), although this is not obvious, since to prove this one must analyze the meaning of an arbitrary logical formula in these models. In fact, the theory $\Delta^{()}$ (see Exercise II.13.9) is *model-complete*; this means that $\mathfrak{A} \preceq \mathfrak{B}$ whenever $\mathfrak{A} \subseteq \mathfrak{B}$ and $\mathfrak{A}, \mathfrak{B}$ are models for $\Delta^{()}$. The example with $[0, 1] \subseteq [0, 2]$ shows that $\Delta^{[]}$ is not model-complete; neither are $\Delta^{[]}$ or $\Delta^{(]}$ by similar examples.

There are now two basic kinds of results about elementary submodels. First, one may establish $\mathfrak{A} \preceq \mathfrak{B}$ using some facts (such as model-completeness) specific to the theories of $\mathfrak{A}, \mathfrak{B}$. Second, the *Löwenheim–Skolem–Tarski Theorem* below involves just the cardinalities of the models; for example, if \mathcal{L} is countable and \mathfrak{B} is arbitrary then there is always a countable \mathfrak{A} with $\mathfrak{A} \preceq \mathfrak{B}$. This is related to the Löwenheim–Skolem Theorem, which yields a countable \mathfrak{A} with $\mathfrak{A} \equiv \mathfrak{B}$ (see Lemma II.13.4). The improvement to $\mathfrak{A} \preceq \mathfrak{B}$ uses the notion of elementary submodel due to Tarski and Vaught. To prove this, we use the following “Tarski–Vaught criterion”, which expresses $\mathfrak{A} \preceq \mathfrak{B}$ as a closure property of \mathfrak{A} .

Lemma II.16.3 *Let \mathfrak{A} and \mathfrak{B} be structures for \mathcal{L} with $\mathfrak{A} \subseteq \mathfrak{B}$. Then the following are equivalent:*

1. $\mathfrak{A} \preceq \mathfrak{B}$.
2. For all existential formulas $\varphi(\vec{x})$ of \mathcal{L} (so $\varphi(\vec{x})$ is of the form $\exists y\psi(\vec{x}, y)$), and all $\vec{a} \in A$, if $\mathfrak{B} \models \varphi[\vec{a}]$, then there is some $b \in A$ such that $\mathfrak{B} \models \psi[\vec{a}, b]$.

We remark that it is common to use this vector notation $\varphi(\vec{x})$ for the longer notation $\varphi(x_1, \dots, x_n)$, where $n \geq 0$ and x_1, \dots, x_n lists the free variables of φ . Then $\vec{a} \in A$ means that \vec{a} denotes an n -tuple of elements of A . Now, the definition (II.7.8) of $\mathfrak{B} \models \varphi[\vec{a}]$ is that there is some $b \in B$ such that $\mathfrak{B} \models \varphi[\vec{a}, b]$. Item (2) is asserting that we can in fact find b in A . Note that (2), unlike (1), only refers to “ $\mathfrak{B} \models$ ”, not “ $\mathfrak{A} \models$ ”. If we are given a large \mathfrak{B} and want to construct a small $\mathfrak{A} \preceq \mathfrak{B}$, then we may view (2) as a closure property of the set A , and we construct a small A to satisfy this property. Until the construction is done, we don’t yet have A , so we can’t talk about “ $\mathfrak{A} \models$ ”.

Proof of Lemma II.16.3. (1) \rightarrow (2): If $\vec{a} \in A$, $\mathfrak{B} \models \varphi[\vec{a}]$, then $\mathfrak{A} \models \varphi[\vec{a}]$ by $\mathfrak{A} \preceq \mathfrak{B}$, so by the definition of \models , there is some $b \in A$ such that $\mathfrak{A} \models \psi[\vec{a}, b]$. But then $\mathfrak{B} \models \psi[\vec{a}, b]$ by $\mathfrak{A} \preceq \mathfrak{B}$.

(2) \rightarrow (1): Assume (2), and now prove $\mathfrak{A} \preceq_{\varphi} \mathfrak{B}$ by induction on the number of quantifiers in φ . If φ has no quantifiers, then this is Lemma II.16.2. Now, assume that φ has one or more quantifiers, and assume the result for formulas with fewer quantifiers. If φ begins (in Polish notation) with a propositional connective, then the induction is elementary, so assume that φ begins with either an \exists or a \forall .

If $\varphi(\vec{x})$ is $\exists y\psi(\vec{x}, y)$, then we prove $\mathfrak{A} \preceq_{\varphi} \mathfrak{B}$ by fixing $\vec{a} \in A$ and noting:

$$\mathfrak{A} \models \varphi[\vec{a}] \text{ iff } \exists b \in A [\mathfrak{A} \models \psi[\vec{a}, b]] \text{ iff } \exists b \in A [\mathfrak{B} \models \psi[\vec{a}, b]] \text{ iff } \mathfrak{B} \models \varphi[\vec{a}] .$$

The first “iff” used the definition of \models and the second “iff” used (inductively) $\mathfrak{A} \preceq_{\psi} \mathfrak{B}$. The \rightarrow of the third “iff” used the definition of \models , while the \leftarrow of the third “iff” used (2).

If φ is $\forall y\psi$, then φ is logically equivalent to $\neg\exists y\neg\psi$, and we use the above inductive argument for \exists . \square

Theorem II.16.4 (Downward Löwenheim–Skolem–Tarski Theorem) *Let \mathfrak{B} be any structure for \mathcal{L} . Fix κ with $\max(|\mathcal{L}|, \aleph_0) \leq \kappa \leq |B|$, and then fix $S \subseteq B$ with $|S| \leq \kappa$. Then there is an $\mathfrak{A} \preceq \mathfrak{B}$ such that $S \subseteq A$ and $|A| = \kappa$.*

Proof. Following the terminology in Lemma II.16.3, for each existential formula φ of \mathcal{L} , choose a Skolem function f_{φ} as follows: Say φ has n free variables, which we list as $\vec{x} = (x_1, \dots, x_n)$ and write φ as $\varphi(\vec{x})$. Then $\varphi(\vec{x})$ is of the form $\exists y\psi(\vec{x}, y)$. Applying the Axiom of Choice, let $f_{\varphi} : B^n \rightarrow B$ be such that for $\vec{a} \in B$, if $\mathfrak{B} \models \varphi[\vec{a}]$, then $\mathfrak{B} \models \psi[\vec{a}, f_{\varphi}(\vec{a})]$. So, if there is a $b \in B$ such that $\mathfrak{B} \models \psi[\vec{a}, b]$, then $f_{\varphi}(\vec{a})$ chooses one such b ; if not, then $f_{\varphi}(\vec{a})$ can be chosen arbitrarily. Since $n = n_{\varphi}$ depends on φ , we really have $f_{\varphi} : B^{n_{\varphi}} \rightarrow B$.

Call $A \subseteq B$ Skolem-closed iff for each existential φ of \mathcal{L} , $f_{\varphi}(A^{n_{\varphi}}) \subseteq A$. Our A will be of this form.

First, note what the above says when $n = n_{\varphi} = 0$. A function of 0 variables “is essentially” a constant. More formally, φ is a sentence, $\exists y\psi(y)$. $B^0 = \{\emptyset\}$, and $f_{\varphi}(\emptyset) = c_{\varphi}$ for some $c_{\varphi} \in A$; if $\mathfrak{B} \models \varphi$, then $\mathfrak{B} \models \psi[c_{\varphi}]$. If A is Skolem-closed, then $c_{\varphi} = f_{\varphi}(\emptyset) \in A$. In particular, letting φ be $\exists y[y = y]$, we see that $A \neq \emptyset$.

Next, a Skolem-closed A is closed under all the functions of \mathcal{L} . That is, if $g \in \mathcal{L}$ is an n -ary function symbol, so that $g_{\mathfrak{B}} : B^n \rightarrow B$, then $g_{\mathfrak{B}}(A^n) \subseteq A$. To prove this, let $\varphi(\vec{x})$ be $\exists y[g(\vec{x}) = y]$, and note that the Skolem function f_{φ} must be the function $g_{\mathfrak{B}}$.

We may now define an \mathcal{L} structure \mathfrak{A} with universe A by declaring that $g_{\mathfrak{A}}$ is the function $g_{\mathfrak{B}} \upharpoonright A^n$ whenever $g \in \mathcal{L}$ is an n -ary function symbol, and $p_{\mathfrak{A}} = p_{\mathfrak{B}} \cap A^n$ $p \in \mathcal{L}$ is an n -ary predicate symbol. In the case of functions, we are using $g_{\mathfrak{B}}(A^n) \subseteq A$ to show that we are defining a legitimate structure; that is, $g_{\mathfrak{A}} : A^n \rightarrow A$.

In the special case $n = 0$: If $g \in \mathcal{L}$ is a constant symbol, then $g_{\mathfrak{A}} = g_{\mathfrak{B}}$, which is in A by the above closure argument. If $p \in \mathcal{L}$ is a proposition letter, we let $p_{\mathfrak{A}} = p_{\mathfrak{B}} \in \{T, F\}$.

Now, given any Skolem-closed A , we have a structure $\mathfrak{A} \subseteq \mathfrak{B}$, and then $\mathfrak{A} \preceq \mathfrak{B}$ is immediate from Lemma II.16.3. We are thus done if we can construct such an A with $S \subseteq A$ and $|A| = \kappa$. Observe first that we may assume that $|S| = \kappa$, since if $|S| < \kappa$, we may replace S with a larger superset.

Let \mathcal{E} be the set of all existential formulas of \mathcal{L} . If $T \subseteq B$, let $T^* = \bigcup_{\varphi \in \mathcal{E}} f_{\varphi}(T^{n_{\varphi}})$. Then $T \subseteq T^*$, since if $\varphi(x)$ is $\exists y[x = y]$, then $n_{\varphi} = 1$ and f_{φ} is the identity function.

Also, if $|T| = \kappa$, then $|T^*| = \kappa$, since $|T^*| \leq \kappa$ follows from the fact that T^* is the union of $|\mathcal{E}| \leq \kappa$ sets, and each set is of the form $f_\varphi(T^{n_\varphi})$, which has size no more than $|T^{n_\varphi}| \leq \kappa$. Now, let $S = S_0 \subseteq S_1 \subseteq S_2 \subseteq \dots$, where $S_{i+1} = (S_i)^*$ for $i \in \omega$, and let $A = \bigcup_i S_i$. Then each $|S_i| = \kappa$, so that $|A| = \kappa$. Furthermore, A is Skolem-closed, since each $\vec{a} \in A^{n_\varphi}$ lies in some $S_i^{n_\varphi}$, so that $f_\varphi(\vec{a}) \in S_{i+1} \subseteq A$. \square

As an example of the above proof, let $\mathcal{L} = \mathcal{L}_{OR} = \{<, +, \cdot, -, 0, 1\}$, as on page 88, and let \mathfrak{B} be the real numbers, with the usual interpretations of the symbols of \mathcal{L} . Start with $S = \emptyset$, so that we produce a countable $\mathfrak{A} \preceq \mathfrak{B}$. Exactly what reals A contains will depend on the choice of the Skolem functions. A must contain all real algebraic numbers, since each such number is the root of a polynomial with integer coefficients. For example, the cubic polynomial $x^3 - 3x + 1$ has exactly three real roots; call them p, q, r , with $p < q < r$. To see that $p, q, r \in A$, let $\tau(x)$ be the term $x^3 - 3x + 1$ (which is easily expressed in \mathcal{L}), let $\varphi(x, y, z)$ be $[x < y < z \wedge \tau(x) = 0 \wedge \tau(y) = 0 \wedge \tau(z) = 0]$, and let ψ be $\exists x \exists y \exists z \varphi(x, y, z)$. Then $\mathfrak{B} \models \psi$, so $\mathfrak{A} \models \psi$ (using $\mathfrak{A} \preceq \mathfrak{B}$), so there must be $a, b, c \in A$ such that $\mathfrak{A} \models \varphi[a, b, c]$, and then $\mathfrak{B} \models \varphi[a, b, c]$ (using $\mathfrak{A} \preceq \mathfrak{B}$). But then $a = p, b = q, c = r$, so $p, q, r \in A$.

In the above, we applied the Löwenheim–Skolem–Tarski Theorem, which is a result of general model theory, to obtain some countable $\mathfrak{A} \preceq \mathfrak{B}$. But in fact, in the specific case where \mathfrak{B} is the real numbers, a well-known result of Tarski implies that we may take A to be precisely the real algebraic numbers. One proof of this is to show that the theory of *real-closed fields* is model-complete, so $\mathfrak{A} \preceq \mathfrak{B}$ follows from $\mathfrak{A} \subseteq \mathfrak{B}$. The axioms for real-closed fields contains the usual axioms for ordered fields, plus the statements that every polynomial of odd degree has a root (this is an infinite list of axioms), plus the statement that every positive element has a square root. The real numbers and the real algebraic numbers are clearly models for these axioms, but proving model-completeness is non-trivial. The proof uses the fact that the axioms allow us to analyze exactly which polynomials have roots; for details, see [5, 24].

As this example illustrates, model theory contains a mixture of general results about general axiomatic theories and detailed facts about specific algebraic systems.

Theorem II.16.4 is called the *downward* Löwenheim–Skolem–Tarski Theorem because it produces elementary submodels. The *upward* theorem, producing elementary extensions, is easy by compactness:

Theorem II.16.5 *Let \mathfrak{B} be any infinite structure for \mathcal{L} . Fix $\kappa \geq \max(|\mathcal{L}|, |B|)$. Then there is a structure \mathfrak{C} for \mathcal{L} with $\mathfrak{B} \preceq \mathfrak{C}$ and $|\mathfrak{C}| = \kappa$.*

Proof. As in Definition II.14.7, let $\text{eDiag}(\mathfrak{B})$ be the elementary diagram of \mathfrak{B} , written in $\mathcal{L}_B = \mathcal{L} \cup \{c_b : b \in B\}$. Then $\text{eDiag}(\mathfrak{B})$ has an infinite model (namely \mathfrak{B}), and $\kappa \geq \max(|\mathcal{L}_B|, \aleph_0)$, so the standard Löwenheim–Skolem Theorem (II.7.16) implies that $\text{eDiag}(\mathfrak{B})$ has a model \mathfrak{C} of size κ . We may assume that \mathfrak{C} interprets the constant c_b as the actual object b ; then $\mathfrak{C} \models \text{eDiag}(\mathfrak{B})$ implies that $\mathfrak{B} \preceq \mathfrak{C}$. \square

In particular, every infinite structure has a proper elementary extension. A simple application of that fact is given in:

Exercise II.16.6 Throughout, $\mathcal{L} = \{<, S\}$ where S is a unary function symbol, and $\mathfrak{B} = (\omega; <, S)$, where S is the successor function.

1. If $\mathfrak{B} \preceq \mathfrak{C}$, then ω is an initial segment of C .
2. If $\mathfrak{B} \preceq \mathfrak{C}$, define $\Phi : C \rightarrow C$ so that $\Phi(c)$ is c if $c \in \omega$ and $S(c)$ if $c \notin \omega$. Then Φ is an automorphism of \mathfrak{C} (that is, an isomorphism from \mathfrak{C} onto \mathfrak{C}).
3. Let $E := \{2n : n \in \omega\}$. Then E is not definable in \mathfrak{B} ; that is, there is no formula $\varphi(x)$ of \mathcal{L} such that $\mathfrak{B} \models \varphi[a]$ iff a is even.

Hint. For (2), note that every element other than 0 has an immediate predecessor. For (3): Let $\mathfrak{B} \preceq \mathfrak{C}$ with $\mathfrak{B} \neq \mathfrak{C}$. If there were a φ as in (3), then $\mathfrak{B} \models \forall x [\varphi(x) \leftrightarrow \neg\varphi(S(x))]$, so the same sentence holds in \mathfrak{C} . But this yields a contradiction, using the automorphism Φ in (2). \square

II.17 Other Proof Theories

This section doesn't exist yet.

Chapter III

Recursion Theory

This page is under construction.

Bibliography

- [1] ACL2 Version 2.9 home page, <http://www.cs.utexas.edu/~moore/acl2/>
- [2] J. Barwise, *Admissible Sets and Structures*, Springer-Verlag, 1975.
- [3] S Burris and H. P. Sankappanavar, *A Course in Universal Algebra*, Springer-Verlag, 1981; see also <http://www.math.uwaterloo.ca/~snburris/htdocs/ualg.html>
- [4] A. Cayley, *The Collected Mathematical Papers of Arthur Cayley*, 13 vols., A. R. Forsyth, ed. The University Press, Cambridge, 1889-1897.
- [5] C. C. Chang and H. J. Keisler, *Model Theory*, Third Edition, North-Holland Publishing Co., 1990.
- [6] P. J. Cohen, The independence of the continuum hypothesis, *Proc. Nat. Acad. Sci. U.S.A.* 50 (1963) 1143-1148.
- [7] P. J. Cohen, The independence of the continuum hypothesis, II, *Proc. Nat. Acad. Sci. U.S.A.* 51 (1964) 105-110.
- [8] Coq home page, <http://coq.inria.fr/>
- [9] N. Cutland, *Computability, An introduction to recursive function theory*, Cambridge University Press, 1980.
- [10] J. W. Dauben, *Georg Cantor, His Mathematics and Philosophy of the Infinite*, Princeton University Press, 1979.
- [11] R. Engelking, *General Topology*, Revised Edition, Heldermann Verlag, 1989.
- [12] Euclid, *Elements*, ~300 BC; English translation:
<http://aleph0.clarku.edu/~djoyce/java/elements/elements.html>
- [13] Galileo Galilei, *Dialogue Concerning Two New Sciences*, 1638; English translation:
http://galileoandeinstein.physics.virginia.edu/tns_draft/index.html
- [14] L. Henkin, The completeness of the first-order functional calculus, *J. Symbolic Logic* 14 (1949) 159-166.

- [15] D. Hilbert and W. Ackermann, *Grundzüge der theoretischen Logik*, Springer, 1928.
- [16] P. Howard and J. E. Rubin, *Consequences of the Axiom of Choice*, American Mathematical Society, 1998.
- [17] Isabelle home page, <http://isabelle.in.tum.de/index.html>
- [18] T. Jech, *Set Theory*, 2nd Edition, Springer, 1996.
- [19] T. Jech, *The Axiom of Choice* North-Holland Pub. Co., 1973.
- [20] K. Kunen, *Set Theory*, North-Holland Pub. Co., 1980.
- [21] E. Landau, *Grundlagen der Analysis* (das Rechnen mit ganzen, rationalen, irrationalen, komplexen Zahlen), Chelsea Publishing Co., 1960.
- [22] N. Luzin, Function, I (Translated by A. Shenitzer), *Amer. Math. Monthly* 105 (1998) 59-67.
- [23] N. Luzin, Function, II (Translated by A. Shenitzer), *Amer. Math. Monthly* 105 (1998) 263-270.
- [24] D. Marker *Model Theory, An Introduction*, Springer-Verlag, 2002.
- [25] Mizar home page, <http://mizar.org/>
- [26] W. W. McCune, *OTTER 3.3 Reference Manual*, Argonne National Laboratory Technical Memorandum ANL/MCS-TM-263, 2003, available at: <http://www.cs.unm.edu/~mccune/otter/>
- [27] W. W. McCune, *Prover9 Manual*, available at: <http://www.cs.unm.edu/~mccune/prover9/manual/Aug-2007/>
- [28] RPN, an introduction to reverse polish notation, web site by Hewlett-Packard, <http://h411111.www4.hp.com/calculators/uk/en/articles/rpn.html>
- [29] J. C. Shepherdson and H. E. Sturgis, Computability of recursive functions, *J. Assoc. Comput. Mach.* 10 (1963) 217-255.
- [30] D. Varberg, E. J. Purcell, and S. E. Rigdon, *Calculus, Instructor's Edition*, 8th edition, Prentice Hall, 2000.
- [31] R. Zach, Hilbert's Program, in *The Stanford Encyclopedia of Philosophy* (Fall 2003 Edition), E. N. Zalta editor, <http://plato.stanford.edu/archives/fall12003/entries/hilbert-program/>

Index

- atomic formula, 89
- axioms
 - of field theory, 105
 - of group theory, 6, 88
 - of set theory, 10
 - Axiom of Choice, 11, 56
 - Comprehension Axiom, 11, 19
 - Extensionality Axiom, 10, 17
 - Foundation Axiom, 11, 66
 - Infinity Axiom, 11, 37
 - Pairing Axiom, 11, 22
 - Power Set Axiom, 11, 46
 - Replacement Axiom, 11, 27
 - Union Axiom, 11, 22
- bound variable, 89
- cardinal, 48
 - arithmetic, 61
- cartesian product, 26
- choice function, 56
- choice set, 56
- cofinality, 63
- complete theory, 104
- Completeness Theorem, 117–126
- conservative extension, 124, 127, 136
- consistent, 97, 111
- Continuum Hypothesis, 8, 14, 16, 63
- counting, 15, 22, 23
- Dedekind-complete, 74, 80, 81
- ducks, 15, 17, 66
- empty set (\emptyset), 20
- empty structure, 103
- equational variety, 88
- falsity (in a model), 96
- field, 105
- finitistic, 29
- formal theory, 29
- formula, 4, 89
- free variable, 4, 89
- function, 26, 30
 - bijection ($\overset{1-1}{\text{onto}}$), 26
 - composition ($G \circ F$), 28
 - injection ($\overset{1-1}{\rightarrow}$), 26
 - restriction of (\upharpoonright), 26
 - surjection (onto), 26
- Hartogs, 53
- Hausdorff Maximal Principle, 59
- inaccessible cardinal, 65
- inconsistent, 97, 111
- induction
 - ordinary, 37
 - transfinite, 39, 42
- isomorphism, 28, 104
- König, 64
- Löwenheim–Skolem Theorem, 66, 81, **98**, **126**, 127
- Löwenheim–Skolem–Tarski Theorem, 7, 139, **140**
- lattice, 138
- lexicographic order, 28
- logical consequence (\models), 97
- logical symbols, 86
- logically equivalent, 99
- maximal, 31

- meta-variable, 76, 77
- metatheory, 29
- minimal, 31
- Modus Ponens, 6, 106
- nonlogical symbols, 86
- ordinal, 16, **33**
 - arithmetic, 41
 - limit, 36
 - successor, 36
- paradox
 - Burali-Forti's, 36
 - Cantor's, 50
 - Russell's, 19, 50
- Polish notation, 81
- precedence, 91
- proper class, 10, 20, 29, 34
- recursion, 42
- relation, 24, 29
 - equivalence, 25
 - inverse (R^{-1}), 27
 - irreflexive, 25
 - partial order, 25
 - reflexive, 25
 - total order, 25
 - transitive, 25
 - well-founded, 31
 - well-order, 32
- Schröder–Bernstein Theorem, 48
- scope, 84, 89
- semantic consequence (\models), 97
- sentence, 4, 89
- structure, 93
- substitution, 100–102
- successor
 - function, 11, 24
 - ordinal, 36
- tautology, 105–106
- transitive set, 33
- truth (in a model), 96
- truth table, 4
- Tukey's Lemma, 58
- turnstile, 6, 78, 94, 97
- unique readability, 83
- universal closure, 90
- universal set (V), 19
- universe
 - of a model, 3
 - of set theory (V), 19
- Venn diagram, 23
- well-founded
 - relation, *see* relation, well-founded
 - set, 67
- well-order, *see* relation, well-order
- Zorn's Lemma, 59