Notes on Mathematical Logic

David W. Kueker

University of Maryland, College Park
E-mail address: dwk@math.umd.edu
URL: http://www-users.math.umd.edu/~dwk/

Contents

Chapter 0. Introduction: What Is Logic? 1

Part 1. Elementary Logic 5

Chapter 1. Sentential Logic 7
0. Introduction 7
1. Sentences of Sentential Logic 8
2. Truth Assignments 11
3. Logical Consequence 13
4. Compactness 17
5. Formal Deductions 19
6. Exercises 20
20

Chapter 2. First-Order Logic 23
0. Introduction 23
1. Formulas of First Order Logic 24
2. Structures for First Order Logic 28
3. Logical Consequence and Validity 33
4. Formal Deductions 37
5. Theories and Their Models 42
6. Exercises 46
46

Chapter 3. The Completeness Theorem 49
0. Introduction 49
1. Henkin Sets and Their Models 49
2. Constructing Henkin Sets 52
3. Consequences of the Completeness Theorem 54
4. Completeness Categoricity, Quantifier Elimination 57
5. Exercises 58
58

Part 2. Model Theory 59

Chapter 4. Some Methods in Model Theory 61
0. Introduction 61
1. Realizing and Omitting Types 61
2. Elementary Extensions and Chains 66
3. The Back-and-Forth Method 69

4. Exercises 71
71

Chapter 5. Countable Models of Complete Theories 73
0. Introduction 73
1. Prime Models 73
2. Universal and Saturated Models 75
3. Theories with Just Finitely Many Countable Models 77
4. Exercises 79
79

Chapter 6. Further Topics in Model Theory 81
0. Introduction 81
1. Interpolation and Definability 81
2. Saturated Models 84
3. Skolem Functions and Indescernables 87
4. Some Applications 91
5. Exercises 95
95

Appendix A. Appendix A: Set Theory 97
1. Cardinals and Counting 97
2. Ordinals and Induction 100

Appendix B. Appendix B: Notes on Validities and Logical Consequence 103
1. Some Useful Validities of Sentential Logic 103
2. Some Facts About Logical Consequence 104

Appendix C. Appendix C: Gothic Alphabet 105

Bibliography 107

Index 109

CHAPTER 0

Introduction: What Is Logic?

Mathematical logic is the study of mathematical reasoning. We do this by developing an *abstract model* of the process of reasoning in mathematics. We then study this model and determine some of its properties.

Mathematical reasoning is *deductive*; that is, it consists of drawing (correct) inferences from given or already established facts. Thus the basic concept is that of a statement being a *logical consequence* of some collection of statements. In ordinary mathematical English the use of "therefore" customarily means that the statement following it is a logical consequence of what comes before.

> Every integer is either even or odd; 7 is not even; *therefore* 7 is odd.

In our model of mathematical reasoning we will need to precisely define *logical consequence*. To motivate our definition let us examine the everyday notion. When we say that a statement σ is a logical consequence of ("follows from") some other statements $\theta_1, \ldots, \theta_n$, we mean, at the very least, that σ is true provided $\theta_1, \ldots, \theta_n$ are all true.

Unfortunately, this does not capture the essence of logical consequence. For example, consider the following:

> Some integers are odd; some integers are prime; therefore some integers are both odd and prime.

Here the hypotheses are both true and the conclusion is true, but the reasoning is not correct.

The problem is that for the reasoning to be logically correct it cannot depend on properties of odd or prime integers other than what is explicitly stated. Thus the reasoning would remain correct if *odd*, *prime*, and *integer* were changed to something else. But in the above example if we replaced *prime* by *even* we would have true hypotheses but a false conclusion. This shows that the reasoning is false, even in the original version in which the conclusion was true.

The key observation here is that in deciding whether a specific piece of reasoning is or is not correct we must consider alMathematical logic is the study of mathematical reasoning. We do this by developing an *abstract model* of the process of reasoning in mathematics. We then study this model and determine some of its properties.

Mathematical reasoning is *deductive*; that is, it consists of drawing (correct) inferences from given or already established facts. Thus the basic concept is that of a statement being a *logical consequence* of some collection of statements. In ordinary mathematical English the use of "therefore" customarily means that the statement following it is a logical consequence of what l ways of interpreting the undefined concepts—*integer*, *odd*, and *prime* in the above example. This is conceptually easier in a formal language in which the basic concepts are represented by symbols (like P, Q) without any standard or intuitive meanings to mislead one.

Thus the fundamental building blocks of our model are the following:

(1) a formal language $\mathcal{L}$,
(2) sentences of $\mathcal{L}$: $\sigma, \theta, \ldots$,
(3) interpretations for $\mathcal{L}$: $\mathfrak{A}, \mathfrak{B}, \ldots$,
(4) a relation $\models$ between interpretations for $\mathcal{L}$ and sentences of $\mathcal{L}$, with $\mathfrak{A} \models \sigma$ read as "σ is true in the interpretation $\mathfrak{A}$," or "$\mathfrak{A}$ is a model of σ."

Using these we can define logical consequence as follows:

Definition -1.1. Let $\Gamma = \{\theta_1, \ldots, \theta_n\}$ where $\theta_1, \ldots, \theta_n$ are sentences of $\mathcal{L}$, and let σ be a sentence of $\mathcal{L}$. Then σ is a *logical consequence* of Γ if and only if for every interpretation $\mathfrak{A}$ of $\mathcal{L}$, $\mathfrak{A} \models \sigma$ provided $\mathfrak{A} \models \theta_i$ for all $i = 1, \ldots, n$.

Our notation for logical consequence is $\Gamma \models \sigma$.

In particular note that $\Gamma \not\models \sigma$, that is, σ is not a logical consequence of Γ, if and only if there is some interpretation $\mathfrak{A}$ of $\mathcal{L}$ such that $\mathfrak{A} \models \theta_i$ for all $\theta_i \in \Gamma$ but $\mathfrak{A} \not\models \sigma$, $\mathfrak{A}$ is not a model of σ.

As a special limiting case note that $\emptyset \models \sigma$, which we will write simply as $\models \sigma$, means that $\mathfrak{A} \models \sigma$ for every interpretation $\mathfrak{A}$ of $\mathcal{L}$. Such a sentence σ is said to be *logically true* (or *valid*).

How would one actually show that $\Gamma \models \sigma$ for specific Γ and σ? There will be infinitely many different interpretations for $\mathcal{L}$ so it is not feasible to check each one in turn, and for that matter it may not be possible to decide whether a particular sentence is or is not true on a particular structure. Here is where another fundamental building block comes in, namely the formal analogue of mathematical proofs. A *proof* of σ from a set Γ of hypotheses is a finite sequence of statements $\sigma_0, \ldots, \sigma_k$ where σ is σ_k and each statement in the sequence is justified by some explicitly stated rule which guarantees that it is a logical consequence of Γ and the preceding statements. The point of requiring use only of rules which are explicitly stated and given in advance is that one should be able to check whether or not a given sequence $\sigma_0, \ldots, \sigma_k$ is a proof of σ from Γ.

The notation $\Gamma \vdash \sigma$ will mean that there is a formal proof (also called a *deduction* or *derivation*) of σ from Γ. Of course this notion only becomes precise when we actually give the rules allowed.

Provided the rules are correctly chosen, we will have the implication

$$\text{if } \Gamma \vdash \sigma \text{ then } \Gamma \models \sigma.$$

Obviously we want to know that our rules are adequate to derive all logical consequences. That is the content of the following fundamental result:

Theorem -1.1 (Completeness Theorem (K. Gödel)). *For sentences of a first-order language $\mathcal{L}$, we have $\Gamma \vdash \sigma$ if and only if $\Gamma \models \sigma$.*

First-order languages are the most widely studied in modern mathematical logic, largely to obtain the benefit of the Completeness Theorem and its applications. In these notes we will study first-order languages almost exclusively.

Part **??** is devoted to the detailed construction of our "model of reasoning" for first-order languages. It culminates in the proof of the Completeness Theorem and derivation of some of its consequences.

Part **??** is an introduction to Model Theory. If Γ is a set of sentences of $\mathcal{L}$, then $\mathrm{Mod}(\Gamma)$, the class of all models of Γ, is the class of all interpretations of $\mathcal{L}$ which make all sentences in Γ true. Model Theory discusses the properties such classes of interpretations have. One important result of model theory for first-order languages is the Compactness Theorem, which states that if $\mathrm{Mod}(\Gamma) = \emptyset$ then there must be some finite $\Gamma_0 \subseteq \Gamma$ with $\mathrm{Mod}(\Gamma_0) = \emptyset$.

Part **??** discusses the famous incompleteness and undecidability results of G'odel, Church, Tarski, et al. The fundamental problem here (the decision problem) is whether there is an effective procedure to decide whether or not a sentence is logically true. The Completeness Theorem does not automatically yield such a method.

Part **??** discusses topics from the abstract theory of computable functions (Recursion Theory).

Part 1

Elementary Logic

CHAPTER 1

Sentential Logic

0. Introduction

Our goal, as explained in Chapter 0, is to define a class of formal languages whose sentences include formalizations of the sttements commonly used in mathematics and whose interpretatins include the usual mathematical structures. The details of this become quite intricate, which obscures the "big picture." We therefore first consider a much simpler situation and carry out our program in this simpler context. The outline remains the same, and we will use some of the same ideas and techniques–especially the interplay of definition by recursion and proof by induction–when we come to first-order languages.

This simpler formal language is called *sentential logic*. In this system, we ignore the "internal" structure of sentences. Instead we imagine ourselves as given some collection of sentences and analyse how "compound" sentences are built up from them. We first see how this is done in English.

If A and B are (English) sentences then so are "A and B", "A or B", "A implies B", "if A then B", "A iff B", and the sentences which assert the opposite of A and B obtained by appropriately inserting "not" but which we will express as "not A" and "not B".

Other ways of connecting sentences in English, such as "A but B" or "A unless B", turn out to be superfluous for our purposes. In addition, we will consider "A implies B" and "if A then B" to be the same, so only one will be included in our formal system. In fact, as we will see, we could get by without all five of the remaining connectives. One important point to notice is that these constructions can be repeated ad infinitum, thus obtaining (for example):

"if (A and B) then (A implies B)",
"A and (B or C)",
"(A and B) or C".

We have improved on ordinary English usage by inserting parentheses to make the resulting sentences unambiguous.

Another important point to note is that the sentences constructed are longer than their component parts. This will have important consequences in our formal system.

In place of the English language connectives used above, we will use the following symbols, called sentential connectives.

English word	Symbol	Name
and	$\wedge$	conjunction
or	$\vee$	disjunction
implies	$\rightarrow$	implication
iff	$\leftrightarrow$	biconditional
not	$\neg$	negation

1. Sentences of Sentential Logic

To specify a formal language $\mathcal{L}$, we must first specify the set of *symbols* of $\mathcal{L}$. The *expressions* of $\mathcal{L}$are then just the finite sequences of symbols of $\mathcal{L}$. Certain distinguished subsets of the set of expressions are then defined which are studied because they are "meaningful" once the language is intepreted. The rules determining the various classes of meaningful expressions are sometimes referred to as the *syntax* of the language.

The *length* of an expression α, denoted $lh(\alpha)$, is the length of α as a sequence of symbols. Expressions α and β are *equal*, denoted by $\alpha = \beta$, if and only if α and β are precisely the same sequence–that is, they have the same length and for each i the i^{th} term of α is the same symbol as the i^{th} term of β. We normally write the sequence whose successive terms are $\varepsilon_0, \varepsilon_1, \ldots, \varepsilon_n$ as $\varepsilon_0\varepsilon_1 \ldots \varepsilon_n$. This is unambiguous provided no symbol is a finite sequence of other symbols, which we henceforth tacitly assume.

In the formal language $\mathcal{S}$ for sentential logic, we will need symbols (infinitely many) for the sentences we imagine ourselves as being given to start with. We will also need symbols for the connectives discussed in the previous section and parentheses for grouping. The only "meaningful" class of expressions of $\mathcal{S}$ we will consider is the set of *sentences*, which will essentially be those expressions built up in the way indicated in the previous section.

Thus we proceed as follows.

Definition 1.1. The symbols of the formal system $\mathcal{S}$ comprise the following:
1) a set of *sentence symbols*: $S_0, S_1, \ldots, S_n, \ldots$ for all $n \in \omega$
2) the sentential *connectives*: $\wedge, \vee, \rightarrow, \leftrightarrow$
3) parentheses: (,)

We emphasize that any finite sequence of symbols of $\mathcal{S}$ is an expression of $\mathcal{S}$. For example:

$$))(\neg S_{17}\neg$$

is an expression of length 6.

Definition 1.2. The set Sn of sentences of $\mathcal{S}$ is defined as follows:
1) $S_n \in Sn$ for all $n \in \omega$
2) if $\phi \in Sn$ then $(\neg\phi) \in Sn$
3) if $\phi, \psi \in Sn$ then $(\phi \star \psi) \in Sn$ where $\star$ is one of $\wedge, \vee, \rightarrow, \leftrightarrow$
4) nothing else is in Sn

To show that some expression is a sentence of $\mathcal{S}$ we can explicitly exhibit each step it its construction according to the definition. Thus

$$((S_3 \wedge (\neg S_1)) \rightarrow S_4) \in Sn$$

since it is constructed as follows:

$$S_4, S_1, (\neg S_1), S_3, (S_3 \wedge (\neg S_1)), ((S_3 \wedge (\neg S_1)) \rightarrow S_4).$$

Such a sequence exhibiting the formation of a sentence is called a *history* of the sentence. In general, a history is not unique since the ordering of (some) sentences in the sequence could be changed.

The fourth clause in the definition is really implicit in the rest of the definition. We put it in here to emphasize its essential role in determining properties of the set Sn. Thus it implies (for example) that every sentence satisfies one of clauses 1), 2), or 3). For example, if $\sigma \in Sn$ and $lh(\sigma) > 1$ then σ begins with (and ends with). So $\neg S_{17} \notin Sn$. Similarly, $(\neg S_{17}\neg) \notin Sn$ since if it were it would necessarily be $(\neg\phi)$ for some $\phi \in Sn$; this can only happen if $\phi = S_{17}\neg$, and $S_{17}\neg \notin Sn$ since it has length greater than 1, but has no parentheses.

The set Sn of sentences was defined as the closure of some explicitly given set (here the set of all sentence symbols) under certain operations (here the operations on expressions which lead from α, β to $(\alpha \wedge \beta)$, etc.). Such a definition is called a definition *by recursion*. Note also that in this definition the operations produce longer expressions. This has the important consequence that we can prove things about sentences *by induction* on their length. Our first theorem gives an elegant form of induction which has the advantage (or drawback, depending on your point of view) of obscuring the connection with length.

THEOREM 1.1. *Let $X \subseteq Sn$ and assume that (a) $S_n \in X$ for all $n \in \omega$, and (b) if $\phi, \psi \in X$ then $(\neg\phi)$ and $(\phi \star \psi)$ belong to X for each binary connective $\star$. Then $X = Sn$.*

PROOF. Suppose $X \neq Sn$. Then $Y = (Sn - X) \neq \emptyset$. Let $\theta_0 \in Y$ be such that $lh(\theta_0) \leq lh(\theta)$ for every $\theta \in Y$. Then $\theta_0 \neq S_n$ for any $n \in \omega$, by (a), hence $\theta_0 = (\neg\phi)$ or $\theta_0 = (\phi \star \psi)$ for sentences ϕ and ψ and some connective $\star$. But then $lh(\phi), lh(\psi) < lh(\theta_0)$ so by choice of θ_0, we have $\phi, \psi \in Y$, i.e. $\phi, \psi \in X$. But then (b) implies that $\theta_0 \in X$, a contradiction. $\square$

As a simple application we have the following.

COROLLARY 1.2. *A sentence contains the same number of left and right parentheses.*

PROOF. Let $p_l(\alpha)$ be the number of left parentheses in a α and let $p_r(\alpha)$ be the number of right parentheses in α. Let $X = \{\theta \in Sn |\ p_l(\theta) = p_r(\theta)\}$. Then $S_n \in X$ for all $n \in \omega$ since $p_l(S_n) = p_r(S_n) = 0$. Further, if $\phi \in X$ then $(\neg\phi) \in X$ since $p_l((\neg\phi)) = 1 + p_l(\phi)$, $p_r((\neg\phi)) = 1 + p_r(\phi)$, and $p_l(\phi) = p_r(\phi)$ since $\phi \in X$ (i.e. "by inductive hypothesis"). The binary connectives are handled similarly and so $X = Sn$. $\square$

The reason for using parentheses is to avoid ambiguity. We wish to prove that we have succeeded. First of all, what–in this abstract context–would be considered an ambiguity? If our language had no parentheses but were otherwise unchanged then $\neg S_0 \wedge S_1$ would be considered a "sentence." But there are two distinct ways to add parentheses to make this into a real sentence of our formal system, namely $((\neg S_0) \wedge S_1)$ and $(\neg(S_0 \wedge S_1))$. In the first case it would have the form $(\alpha \wedge \beta)$ and in the second the form $(\neg\alpha)$. Similarly, $S_0 \to S_1 \to S_2$ could be made into either of the sentences $((S_0 \to S_1) \to S_2)$ or $(S_0 \to (S_1 \to S_2))$. Each of these has the form $(\alpha \to \beta)$, but for different choices of α and β. What we mean by lack of ambiguity is that no such "double entendre" is possible, that we have instead *unique readability* for sentences.

THEOREM 1.3. *Every sentence of length greater than one has exactly one of the forms:* $(\neg\phi), (\phi \vee \psi), (\phi \wedge \psi), (\phi \to \psi), (\phi \leftrightarrow \psi)$ *for exactly one choice of sentences* ϕ, ψ *(or* ϕ *alone in the first form).*

This result will be proved using the following lemma, whose proof is left to the reader.

LEMMA 1.4. *No proper initial segment of a sentence is a sentence. (By a* proper initial segment *of a sequence* $\varepsilon_0\varepsilon_1 \ldots \varepsilon_{n-1}$ *is meant a sequence* $\varepsilon_0\varepsilon_1 \ldots \varepsilon_{m-1}$*, consisting of the first* m *terms for some* $m < n$*).*

PROOF. (of the Theorem from the Lemma) Every sentence of length greater than one has at least one of these forms, so we need only consider uniqueness. Suppose θ is a sentence and we have

$$\theta = (\alpha \star \beta) = (\alpha' \star' \beta')$$

for some binary connectives $\star, \star'$ and some sentences $\alpha, \beta, \alpha', \beta'$. We show that $\alpha = \alpha'$, from which it follows that $\star = \star'$ and $\beta = \beta'$. First note that if $lh(\alpha) = lh(\alpha')$ then $\alpha = \alpha'$ (explain!). If, say, $lh(\alpha) < lh(\alpha')$ then α is a proper initial segment of α', contradicting the Lemma. Thus the only possibility is $\alpha = \alpha'$. We leave to the reader the easy task of checking when one of the forms is $(\neg\phi)$. □

We in fact have more parentheses than absolutely needed for unique readability. The reader should check that we could delete parentheses around negations–thus allowing $\neg\phi$ to be a sentence whenever ϕ is–and still have unique readability. In fact, we could erase all right parentheses entirely–thus allowing $(\phi \wedge \psi$, $(\phi \vee \psi$, etc. to be sentences whenever ϕ, ψ are–and still maintain unique readability.

In practice, an abundance of parentheses detracts from readability. We therefore introduce some conventions which allow us to omit some parentheses when writing sentences. First of all, we will omit the outermost pair of parentheses, thus writing $\neg\phi$ or $\phi \wedge \psi$ in place of $(\neg\phi)$ or $(\phi \wedge \psi)$. Second we will omit the parentheses around negations even when forming further sentences–for example instead of $(\neg S_0) \wedge S_1$, we will normally write just $\neg S_0 \wedge S_1$. This convention does not cuase any ambiguity in practice because $(\neg(S_0 \wedge S_1))$ will be written as $\neg(S_0 \wedge S_1)$. The informal rule is that negation applies to as little as possible.

Building up sentences is not really a linear process. When forming $(\phi \to \psi)$, for example, we need to have both ϕ and ψ but the order in which they appear in a history of $(\phi \to \psi)$ is irrelevant. One can represent the formation of $(\phi \to \psi)$ uniquely in a two-dimensional fashion as follows:

!!!!!!!!!!!!!!!!!!!!!!!!!!!!

By iterating this process until sentence symbols are reached one obtains a *tree* representation of any sentence. This representation is unique and graphically represents the way in which the sentence is constructed.

For example the sentence

$$((S_7 \wedge (S_4 \to (\neg S_0))) \to (\neg(S_3 \wedge (S_0 \to S_2))))$$

is represented by the following tree:

!!!!!!!!!!!!!!!!!!!!!!!!!!!!

We have one final convention in writing sentences more readably. It is seldom important whether a sentence uses the sentence symbols $S_0,\ S_{13}$, and S_7 or $S_{23},\ S_6$, and S_{17}. We will use $A,\ B,\ C, \ldots$ (perhaps with sub- or superscripts) as variables standing for arbitrary sentence symbols (assumed distinct unless explicitly noted to the contrary). Thus we will normally refer to $A \to (B \to C)$, for example, rather than $S_0 \to (S_{17} \to S_{13})$.

2. Truth Assignments

An interpretation of a formal language $\mathcal{L}$must, at a minimum, determine which of the sentences of $\mathcal{L}$are true and which are false. For sentential logic this is all that could be expected. So an interpretation for $\mathcal{S}$ could be identified with a function mapping Sn into the two element set $\{T, F\}$, where T stands for "true" and F for "false."

Not every such function can be associated with an interpretation of $\mathcal{S}$, however, since a real interpretation must agree with the intuitive (or, better, the intended) meanings of the connectives. Thus $(\neg\phi)$ should be true iff ϕ is false and $(\phi \land \psi)$ shuld be true iff both ϕ and ψ are true. We adopt the inclusive interpretation of "or" and therefore say that $(\phi \lor \psi)$ is true if either (or both) of ϕ, ψ is true. We consider the implication $(\phi \to \psi)$ as meaning that ψ is true provided ϕ is true, and therefore we say that $(\phi \to \psi)$ is true unless ϕ is true and ψ is false. The biconditional $(\phi \leftrightarrow \psi)$ will thus be true iff ϕ, ψ are both true or both false.

We thus make the following definition.

Definition 2.1. An interpretation for $\mathcal{S}$ is a function $t : Sn \to \{T, F\}$ satisfying the following conditions for all $\phi, \psi \in Sn$:

(i) $t((\neg\phi)) = T$ iff $t(\phi) = F$,
(ii) $t((\phi \land \psi)) = T$ iff $t(\phi) = t(\psi) = T$,
(iii) $t((\phi \lor \psi)) = T$ iff $t(\phi) = T$ or $t(\psi) = T$ (or both),
(iv) $t((\phi \to \psi)) = F$ iff $t(\phi) = T$ and $t(\psi) = F$, and
(v) $t((\phi \leftrightarrow \psi))$ iff $t(\phi) = t(\psi)$.

How would one specify an interpretation in practice? The key is the following lemma, which is easily established by induction.

Lemma 2.1. *Assume t and t' are both interpretations for $\mathcal{S}$ and that $t(S_n) = t'(S_n)$ for all $n \in \omega$. Then $t(\sigma) = t'(\sigma)$ for all $\sigma \in Sn$.*

So an interpretation is determined completely once we know its values on the sentence symbols. One more piece of terminology is useful.

Definition 2.2. A truth assignment is a function $h : \{S_n \mid n \in \omega\} \to \{T, F\}$.

A truth assignment, then, can be extended to at most one interpretation. The obvious question is whether every truth assignment can be extended to an interpretation.

Given a truth assignment h, let's see how we could try to extend it to an interpretation t. Let $\sigma \in Sn$ and let $\phi_0, \ldots, \phi_n$ be a history of σ (so $\phi_n = \sigma$). We then can define t on each $\phi_i, 0 \le i \le n$, one step at a time, using the requirements in the definition of an interpretation; at the last step we will have defined $t(\sigma)$. Doing this for every $\sigma \in Sn$ we end up with what should be an interpretation t. The only way this could go wrong is if, in considering different histories, we were forced to assign different truth values to the same sentence ϕ. But this could only happen through a failure of unique readability.

This argument can be formalized to yield a proof of the remaining half of the following result.

THEOREM 2.2. *Every truth assignment can be extended to exactly one interpretation.*

PROOF. Let h be a truth assignment. We outline how to show that h can be extended to an interpretation t. The main fact to establish is:

> (*) assume that $h_k(S_n) = h(S_n)$ for all $n \in \omega$ and $h_k : \{\sigma \in Sn|\ lh(\sigma) \leq k\} \to \{T, F\}$ satisfies (i)-(v) in the definition of an interpretation for sentences in its domain; then h_k can be extended to h_{k+1} defined on $\{\sigma \in Sn|\ lh(\sigma) \leq k+1\}$ and which also satisfies (i)-(v) in the definition of an interpretation for all sentences in its domain.

Using this to define a chain

$$h = h_1 \subseteq h_2 \subseteq \ldots \subseteq h_k \ldots$$

and we see that $t = \bigcup\{h_k|\ k \in \omega\}$ is an interpretation, as desired. $\square$

In filling in the details of this argument the reader should be especially careful to see exactly where unique readability is used.

DEFINITION 2.3. For any truth assignment h its unique extension to an interpreteation is denoted by $\bar{h}$.

Given h and σ we can actually compute $\bar{h}(\sigma)$ by successively computing $\bar{h}(\phi_i)$ for each sentence ϕ_i in a history $\phi_0, \ldots, \phi_n$ of σ. Thus if $h(S_n) = F$ for all $n \in \omega$ we successively see that $\bar{h}(S_4) = F$, $\bar{h}(S_1) = F$, $\bar{h}(\neg S_1) = T$, $\bar{h}(S_3) = F$, $\bar{h}(S_3 \wedge S_1) = F$, and finally $\bar{h}((S_3 \wedge S_1) \to S_4) = T$. This process is particularly easy if σ is given in tree form–h tells you how to assign T, F to the sentence symbols at the base of the tree, and (i)-(v) of the definition of an interpretation tell you how to move up the tree, node by node.

There are many situations in which we are given some function f defined on the sentence symbols and want to extend it to all sentences satisfying certain conditions relating the values at $(\neg\phi)$, $(\phi \wedge \psi)$, etc. to its values at ϕ, ψ. Minor variations in the argument for extending truth assignments to interpretations establish that this can always be done. The resulting function is said to be defined *by recursion* , on the class of sentences.

THEOREM 2.3. *Let X be any set, and let $g_\neg : X \to X$ and $g_\star : X \times X \to X$ be given for each binary connective $\star$. Let $f : \{S_n|\ n \in \omega\} \to X$ be arbitrary. Then there is exactly one function $\bar{f} : Sn \to X$ such that*
$\bar{f}(S_n) = f(S_n)$ for all $n \in \omega$,
$\bar{f}(\neg\phi) = g_\neg(\bar{f}(\phi))$ for all $\phi \in Sn$,
$\bar{f}(\phi \star \psi) = g_\star(\bar{f}(\phi), \bar{f}(\psi))$ for all $\phi, \psi \in Sn$ and binary connectives $\star$.

Even when we have an informal definition of a function on the set Sn, it frequently is necessary to give a precise definition by recursion in order to study the properties of the function.

EXAMPLE 2.1. Let $X = \omega$, $f(S_n) = 0$ for all $n \in \omega$. Extend f to $\bar{f}$ on Sn via he recursion clauses

$\bar{f}((\neg\phi)) = \bar{f}(\phi) + 1$
$\bar{f}((\phi \star \psi)) = \bar{f}(\phi) + \bar{f}(\psi) + 1$ for binary connectives $\star$.

We can then interpret $\bar{f}(\theta)$ as giving any of the following:
the number of left parentheses in θ,
the number of right parentheses in θ,
the number of connectives in θ.

EXAMPLE 2.2. Let ϕ_0 be some fixed sentence. We wish to define $\bar{f}$ so that $\bar{f}(\theta)$ is the result of replacing S_0 throughout θ by ϕ_0. This is accomplished by recursion, by starting with f given by

$$f(S_n) = \begin{cases} \phi_0, & n = 0 \\ S_n, & n \neq 0 \end{cases}$$

and extending via the recursion clauses
$\bar{f}((\neg\phi)) = (\neg\bar{f}(\phi))$,
$\bar{f}((\phi \star \psi)) = (\bar{f}(\phi) \star \bar{f}(\psi))$ for binary connectives $\star$.

For the function $\bar{f}$ of the previous example, we note the following fact, established by induction.

LEMMA 2.4. *Given any truth assignment h define h^* by*

$$h^*(S_n) = \begin{cases} \bar{h}(\phi_0), & n = 0 \\ h(S_n), & n \neq 0 \end{cases}$$

Thus for any sentence θ we have $\bar{h^*}(\theta) = \bar{h}(\bar{f}(\theta))$.

PROOF. By definition of h^* and f we see that $h^*(S_n) = \bar{h}(f(S_n))$ for all n. The recursion clauses yielding $\bar{f}$ guarantees that this property is preserved under forming longer sentences. □

Note that the essential part in proving that a sentence has the same number of left parentheses as right parentheses was noting, as in Example 1.3.1, that these two functions satisfied the same recursion clauses.

As is common in mathematical practice, we will frequently not distinguish notationally between f and $\bar{f}$. Thus we will speak of defining f by recursion given the operation of f on $\{S_n | \, n \in \omega\}$ and certain recursion clauses involving f.

3. Logical Consequence

Since we now know that every truth assignment h extends to a unique interpretation, we follow the outline established in the Introduction using as our fundamental notion the truth of a sentence under a truth assignment.

DEFINITION 3.1. Let h be a truth assignment and $\theta \in Sn$. Then θ is true under h, written $h \models \theta$, iff $\bar{h}(\theta) = T$ where $\bar{h}$ is the unique extension of h to an interpretation.

Thus θ is *not true* under h, written $h \not\models \theta$, iff $\bar{h}(\theta) \neq T$. Thus $h \not\models \theta$ iff $\bar{h}(\theta) = F$ iff $h \models \neg\theta$.

We will also use the following terminology: h satisfies θ iff $h \models \theta$.

DEFINITION 3.2. A sentence θ is satisfiable iff it is satisfied by some truth assignment h.

We extend the terminology and notation to sets of sentences in the expected way.

DEFINITION 3.3. Let h be a truth assignment and $\Sigma \subseteq Sn$. Then Σ is true under h, or h satisfies Σ, written $h \models \Sigma$, iff $h \models \sigma$ for every $\sigma \in \Sigma$.

DEFINITION 3.4. A set Σ of sentences is satisfiable iff it is satisfied by some truth assignment h.

The definitions of logical consequence and (logical) validity now are exactly as given in the Introduction.

DEFINITION 3.5. Let $\theta \in Sn$ and $\Sigma \subseteq Sn$. Then θ is a logical consequence of Σ written $\Sigma \models \theta$, iff $h \models \theta$ for every truth assignment h which satisfies Σ.

DEFINITION 3.6. A sentence θ is (logically) valid, or a tautology, iff $\emptyset \models \theta$, i.e. $h \models \theta$ for every truth assignment h.

It is customary to use the word "tautology" in the context of sentential logic, and reserve "valid" for the corresponding notion in first order logic. Our notation in any case will be $\models \theta$, rather than $\emptyset \models \theta$.

The following lemma, translating these notions into satisfiability, is useful and immediate from the definitions.

LEMMA 3.1. *(a) θ is a tautology iff $\neg\theta$ is not satisfiable. (b) $\Sigma \models \theta$ iff $\Sigma \cup \{\neg\theta\}$ is not satisfiable.*

Although there are infinitely many (indeed uncountably many) different truth assignments, the process of checking validity or satisfiability is much simpler becdause only finitely many sentence symbols occur in any one sentence.

LEMMA 3.2. *Let $\theta \in Sn$ and let h, h^* be truth assignments such that $h(S_n) = h^*(S_n)$ for all S_n in θ. Then $\bar{h}(\theta) = \bar{h}^*(\theta)$, and thus $h \models \theta$ iff $h^* \models \theta$.*

PROOF. Let $A_1, \ldots, A_n$ be sentence symbols, and let h, h^* be truth assignments so that $h(A_i) = h^*(A_i)$ for all $i = 1, \ldots, n$. We show by induction that for every $\theta \in Sn$, $\bar{h}(\theta) = \bar{h}^*(\theta)$ provided θ uses no sentence symbols other than $A_1, \ldots, A_n$. The details are straightforward. $\square$

This yields a finite, effective process for checking validity and satisfiability of sentences, and also logical consequences of finite sets of sentences.

THEOREM 3.3. *Let $A_1, \ldots, A_n$ be sentence symbols. Then one can find a finite list $h_1, \ldots, h_m$ of truth assignments such that for every sentence θ using no sentence symbols other than $A_1, \ldots, A_n$ we have: (a) $\models \theta$ iff $h_j \models \theta$ for all $j = 1, \ldots, m$, and (b) θ is satisfiable iff $h_j \models \theta$ for some $j, 1 \leq j \leq m$. If further Σ is a set of sentences using no sentence symbols other than $A_1, \ldots, A_n$ then we also have: (c) $\Sigma \models \theta$ iff $h_j \models \theta$ whenever $h_j \models \Sigma$, for each $j = 1, \ldots, m$.*

PROOF. Given $A_1, \ldots, A_n$ we let $h_1, \ldots, h_m$ list all truth assignments h such that $h(S_k) = F$ for every S_k different from $A_1, \ldots, A_n$. There are exactly $m = 2^n$ such, and they work by the preceding lemma. $\square$

The information needed to check whether or not a sentence θ in the sentence symbols $A_1, \ldots, A_n$ is a tautology is conveniently represented in a table. Across the top of the table one puts a history of θ, beginning with $A_1, \ldots, A_n$, and each line of the table corresponds to a different assignment of truth values to $A_1, \ldots, A_n$.

For example, the following *truth table* shows that $(S_3 \wedge \neg S_1) \to S_4$ is not a tautology.

S_1	S_3	S_4	$\neg S_1$	$S_3 \wedge \neg S_1$	$(S_3 \wedge \neg S_1) \to S_4$
T	T	T	F	F	T
T	T	F	F	F	T
T	F	T	F	F	T
T	F	F	F	F	T
F	T	T	T	T	T
F	T	F	T	T	F
F	F	T	T	F	T
F	F	F	T	F	T

Writing down truth tables quickly becomes tedious. Frequently shortcuts are possible to reduce the drudgery. For example, if the question is to determine whether or not some sentence θ is a tautology, suppose that $\bar{h}(\theta) = F$ and work backwards to see what h must be. To use the preceding example, we see that

$$\begin{aligned} & \bar{h}((S_3 \wedge \neg S_1) \to S_4) = F \\ \text{iff} \quad & \bar{h}((S_3 \wedge \neg S_1)) = T \text{ and } h(S_4) = F \\ \text{and} \quad & \bar{h}((S_3 \wedge \neg S_1)) = T \\ \text{iff} \quad & h(S_1) = f \text{ and } h(S_3) = T. \end{aligned}$$

Thus this sentence is not a tautology since it is false for every h such that $h(S_1) = F$, $h(S_3) = T$, and $h(S_4) = F$.

As another example, consider $\theta = (A \to B) \to ((\neg A \to B) \to B)$. Then $\bar{h}(\theta) = F$ iff $\bar{h}(A \to B) = T$ and $\bar{h}((\neg A \to B) \to B = F$. And $\bar{h}((\neg A \to B) \to B) = F$ iff $\bar{h}(\neg A \to B) = T$ and $h(B) = F$. Now for $h(B) = F$ we have $\bar{h}(A \to B) = T$ iff $h(A) = F$ and $\bar{h}(\neg A \to B) = T$ iff $h(A) = T$. Since we can't have both $h(A) = T$ and $h(a) = F$ we may conclude that θ is a tautology.

Some care is needed in such arguments to ensure that the conditions obtained on h at the end are actually equivalent to $\bar{h}(\theta)$. Otherwise some relevant truth assignment may have escaped notice. Of course only the implications in one direction are needed to conclude θ is a tautology, and only the implications in the other direction to conclude that such an h actually falsifies θ. But until you know which conclusion holds, both implications need to be preserved.

An analogous process, except starting with the supposition $\bar{h}(\theta) = T$, can be used to determine the satisfiability of θ. If Σ is the finite set $\{\sigma_1, \ldots, \sigma_k\}$ of sentences then one can check whether or not $\Sigma \models \theta$ by supposing $\bar{h}(\theta) = F$ while $\bar{h}(\sigma_i) = T$ for all $i = 1, \ldots, k$ and working backwards from these hypotheses.

An important variation on logical consequence is given by logical equivalence.

Definition 3.7. Sentences ϕ, ψ are logically equivalent, written $\phi \vdash\dashv \psi$, iff $\{\phi\} \models \psi$ and $\{\psi\} \models \phi$.

Thus, logically equivalent sentences are satisfied by precisely the same truth assignments, and we will think of them as making the same assertion in different ways.

Some examples of particular interest to us invole writing one connective in terms of another.

LEMMA 3.4. *For any $\phi, \psi \in Sn$ we have:*

(a) $(\phi \to \psi) \quad \vdash\dashv \quad (\neg\phi \vee \psi)$
(b) $(\phi \vee \psi) \quad \vdash\dashv \quad (\neg\phi \to \psi)$
(c) $(\phi \vee \psi) \quad \vdash\dashv \quad \neg(\neg\phi \wedge \neg\psi)$
(d) $(\phi \wedge \psi) \quad \vdash\dashv \quad \neg(\neg\phi \vee \neg\psi)$
(e) $(\phi \wedge \psi) \quad \vdash\dashv \quad \neg(\phi \to \neg\psi)$
(f) $(\phi \leftrightarrow \psi) \quad \vdash\dashv \quad (\phi \to \psi) \wedge (\psi \to \phi)$

What we want to conclude, using parts (b), (e), and (f) of the above lemma is that every sentence θ is logically equivalent to a sentence θ^* using the same sentence symbols but only the connectives $\neg, \to$. This is indeed true, and we outline the steps needed to prove $\neg\theta$.

First of all, we must define (by recursion) the operation * on sentences described by saying that θ^* results from θ by replacing subexpressions $(\phi \vee \psi), (\phi \wedge \psi), (\phi \leftrightarrow \psi)$ of θ (for sentences ϕ, ψ) by their equivalents in terms of $\neg, \to$ given in the lemma.

Secondly, we must prove (by induction) that for every truth assignment h and every $\theta \in Sn$ we have $\bar{h}(\theta) = \bar{h}(\theta^*)$.

Details of this, and similar substitution facts, are left to the reader.

Due to the equivalence $(\phi \vee \psi) \vee \theta \vdash\dashv \phi \vee (\psi \vee \theta)$ and $(\phi \wedge \psi) \wedge \theta \vdash\dashv \phi \wedge (\psi \wedge \theta)$, we will omit the parentheses used for grouping conjunctions and disjunctions, thus writing $A \vee B \vee C \vee D$ instead of $((A \vee B) \vee C) \vee D$.

Sentences written purely in terms of $\neg, \to$ are not always readily understandable. Much preferable for some purposes are sentences written using $\neg, \vee, \wedge$– especially those in one of the following special forms:

DEFINITION 3.8. (a) A sentence θ is in disjunctive normal form iff it is a disjunction $(\theta_1 \vee \theta_2 \vee \ldots \vee \theta_n)$ in which each disjunct θ_i is a conjugation of sentence symbols and negations of sentence symbols. (b) A sentence θ is in conjunctive normal form iff it is a conjunction $(\theta_1 \wedge \theta_2 \wedge \ldots \wedge \theta_n)$ in which each conjunct θ_i is a disjunction of sentence symbols and negations of sentence symbols.

The advantage of having a sentence in disjunctive normal form is that it is easy to read off the truth assignments which statisfy it. For example

$$(A \wedge \neg B) \vee (A \wedge B \wedge \neg C) \vee (B \wedge C)$$

is satisfied by a truth assignment h iff either $h(A) = T$ and $h(B) = F$ or $h(A) = h(B) = T$ and $h(C) = F$ or $h(B) = h(C) = T$.

THEOREM 3.5. *Let θ be any sentence. Then there is a sentence θ^* in disjunctive normal form and there is a sentence θ^{**} in conjunctive normal form such that*

$$\theta \vdash\dashv \theta^*, \theta \vdash\dashv \theta^{**}.$$

PROOF. Let $A_1, \ldots, A_n$ be sentence symbols. For any $X \subseteq \{1, \ldots, n\}$ we define θ_X to be $(\phi_1 \wedge \ldots, \wedge\phi_n)$ where $\phi_i = A_i$ if $i \in x$ and $\phi_i = \neg A_i$ if $i \notin X$. It is then clear that a truth assignment h satisfies θ_X iff $h(A_i) = T$ for $i \in X$ and $h(A_i) = F$ for $i \notin X$. Now, given a sentence θ built up using no sentence symbols other than $A_1, \ldots, A_n$ let θ^* be the disjunction of all θ_X such that $(\theta \wedge \theta_X)$ is satisfiable– equivalently, such that $\models (\theta_X \to \theta)$. Then θ^* is, by construction, in disjunctive normal form and is easily seen to be equivalent to θ. If $(\theta \wedge \theta_X)$ is not satisfiable for any X then θ is not satisfiable, hence θ is equivalent to $(A_1 \wedge \neg A_1)$ which is in disjunctive normal form.

We leave the problem of finding θ^{**} to the reader. $\square$

Note that using θ_X's, without being given any θ to begin with, we can form sentences θ^* with *any* given truth table in $A_1, \ldots, A_n$. Thus there are no "new" connectives we could add to extend the expressive power of our system of sentential logic.

4. Compactness

If Σ is a finite set of sentences then the method of truth tables gives an effective, finite procedure for deciding whether or not Σ is satisfiable. Similarly one can decide whether or not $\Sigma \models \theta$ for finite $\Sigma \subseteq Sn$. The situation is much different for infinite sets of sentences. The Compactness Theorem does, however, reduce these questions to the corresponding questions for finite sets. The Compactness Theorem in first order logic will be one of our most important and useful results, and its proof in that setting will have some similarities to the arguments in this section.

THEOREM 4.1. *(Compactness) Let $\Sigma \subseteq Sn$. (a) Σ is satisfiable iff every finite $\Sigma_0 \subseteq \Sigma$ is satisfiable. (b) For $\theta \in Sn$, $\Sigma \models \theta$ iff there is some finite $\Sigma_0 \subseteq \Sigma$ such that $\Sigma_0 \models \theta$.*

Part (b) follows from part (a) using part (b) of Lemma 1.4.1. The implication from left to right in (a) is clear, so what needs to be shown is that Σ is satisfiable provided every finite $\Sigma_0 \subseteq \Sigma$ is satisfiable. The problem, of course, is that different finite subsets may be satisfied by different truth assignments and that, a priori, there is no reason to assume that a single truth assignment will satisfy *every* finite subset of Σ (equivalently, all of Σ).

Rather than taking the most direct path to this result, we will discuss in more generality correspondences between interpretatins and the sets of sentences they satisfy. In particular we look at the ways in which we could use a set Σ of sentences to define a truth assignment h which satisfies it.

Given Σ, if we wish to define a particular truth assignment h which satisfies Σ we must, for example, either set $h(S_0) = T$ or $h(S_0) = F$. If $S_0 \in \Sigma$ then we must make the first choice; if $\neg S_0 \in \Sigma$ we must make the second choice. The only case in which we may be in doubt is if neither S_0 nor $\neg S_0$ belongs in Σ. But even here we may be forced into one or the other choice, for example, if $(S_0 \wedge \neg S_3) \in \Sigma$ or $(\neg S_0 \wedge S_3) \in \Sigma$.

Our definition of a complete set of sentences is intended to characterize those for which we have no choice in defining a satisfying truth assignment and for which we are not forced into contradictory choices.

DEFINITION 4.1. A set $\Gamma \subseteq Sn$ is complete iff the following hold for all $\phi, \psi \in Sn$:

(i) $(\neg\phi) \in \Gamma$ iff $\phi \notin \Gamma$,
(ii) $(\phi \wedge \psi) \in \Gamma$ iff $\phi \in \Gamma$ and $\psi \in \Gamma$,
(iii) $(\phi \vee \psi) \in \Gamma$ iff $\phi \in \Gamma$ or $\psi \in \Gamma$,
(iv) $(\phi \to \psi) \in \Gamma$ iff $(\neg\phi) \in \Gamma$ or $\psi \in \Gamma$,
(v) $(\phi \leftrightarrow \psi) \in \Gamma$ iff either both $\phi, \psi \in \Gamma$ or both $\phi, \psi \notin \Gamma$.

DEFINITION 4.2. Given a truth assignment h, $T(h) = \{\sigma \in Sn \mid h \models \sigma\}$.

Complete sets of sentences are exactly what we are after, as shown by the next result.

Theorem 4.2. *A set Γ of sentences is complete iff $\Gamma = T(h)$ for some truth assignment h.*

Proof. From right to left is clear because the clauses in the definition of complete sets mimic the recursion clauses in extending h to $\bar{h}$.

Conversely, if Γ is complete we define h by $h(S_n) = T$ iff $S_n \in \Gamma$ and show by induction that a sentence θ belongs to Γ iff $\bar{h}(\theta) = T$. $\square$

Since clearly two truth assignments h_1, h_2 are equal iff $T(h_1) = T(h_2)$ we have a one-to-one correspondence between truth assignments and complete sets of sentences.

The relevance of this to proving the satisfiability of sets of sentences is the following consequence.

Corollary 4.3. *Let $\Sigma \subseteq Sn$. Then Σ is satisfiable iff there is some complete set Γ of sentences such that $\Sigma \subseteq \Gamma$.*

Thus our approach to showing that some set of sentences is satisfiable will be to extend it to a complete set. For the specific purposes of showing compactness we will need the following terminology.

Definition 4.3. A set $\Sigma \subseteq Sn$ is finitely satisfiable iff every finite $\Sigma_0 \subseteq \Sigma$ is satisfiable.

Thus our method in proving compactness will be to show that a finitely satisfiable set Σ of sentences can be extended to a complete set Γ. We will construct this extension step-by-step, using the following lemma at each step.

Lemma 4.4. *Assume Σ is finitely satisfiable and let θ be a sentence. Then at least one of $\Sigma \cup \{\theta\}, \Sigma \cup \{\neg\theta\}$ is fnitely satisfiable.*

At the end of the construction the verification that the resulting set Γ is complete will use the following two lemmas.

Lemma 4.5. *Assume that Σ_n is finitely satisfiable and $\Sigma_n \subseteq \Sigma_{n+1}$ for all $n \in \omega$. Let $\Gamma = \bigcup_{n\in\omega} \Sigma_n$. Then Γ is finitely satisfiable.*

Lemma 4.6. *Assume that Γ is finitely satisfiable and for all sentences ϕ either $\phi \in \Gamma$ or $(\neg\phi) \in \Gamma$. Then Γ is complete.*

We leave the proofs of these lemmas to the reader and proceed to give the construction.

First of all, since our formal system $\mathcal{S}$ has only countably many symbols and every sentence is a finite sequence of symbols, it follows that Sn is a countable set, so we may list it as $Sn = \{\phi_n |\ n \in \omega\}$.

Next we define, by recursion on $n \in \omega$ a chain $\{\Sigma_n\}_{n\in\omega}$ of finitely satisfiable sets of sentences as follows:

$$\Sigma_0 = \Sigma$$

$$\Sigma_{n+1} = \begin{cases} \Sigma_n \cup \{\phi_n\}, & \text{if this is finitely satisfiable} \\ \Sigma_n \cup \{\neg\phi_n\}, & \text{otherwise} \end{cases}$$

The first lemma above establishes that in either case Σ_{n+1} will be finitely satisfiable.

Finally, we let $\Gamma = \bigcup_{n\in\omega} \Sigma_n$. Γ is finitely satisfiable by the second lemma above. If $\phi \in Sn$ then there is some $n \in \omega$ such that $\phi = \phi_n$. Thus either $\phi \in \Sigma_{n+1} \subseteq \Gamma$

or $(\neg\phi) \in \Sigma_{n+1} \subseteq \Gamma$ by the construction. Thus we conclude that Γ is complete by the last lemma above.

To return to the question with which we opened this section, how does the Compactness Theorem help us decide whether or not $\Sigma \models \theta$? Assume that we are given some explicit listing of $\Sigma = \{\sigma_n | \, n \in \omega\}$. Then $\Sigma \models \theta$ iff $\Sigma_n = \{\sigma_0, \ldots, \sigma_n\} \models \theta$ for some $n \in \omega$. Thus we check each n in turn to see if $\Sigma_n \models \theta$. If in fact $\Sigma \models \theta$ then we will eventually find an $n \in \omega$ such that $\Sigma_n \models \theta$, and hence be able to conclude that $\Sigma \models \theta$. Unfortunately, if $\Sigma \not\models \theta$ this process never terminates and so we are unable to conclude that $\Sigma \not\models \theta$.

5. Formal Deductions

To complete the model of mathematical reasoning sketched in the Introduction we need to introduce the concept of a formal deduction. This does not play an important role in sentential logic because the method of truth tables enable us to determine which sentences are valid, so we only sketch the development in this section.

We will specify a set Λ_0 of validities to serve as *logical axioms* and a *rule* for deriving a sentence given certain others–both of these will be defined syntactically, that is purely in terms of the forms of the sentences involed.

The rule, called *modus ponens* (MP), states that ψ can be derived from ϕ and $(\phi \to \psi)$. Note that application of this rule preserves validity, and more generally, if $\Gamma \models \phi$ and $\Gamma \models (\phi \to \psi)$ then $\Gamma \models \psi$.

To minimize the set Λ_0 we restrict attention to sentences built using only the connectives $\neg, \to$. This entails no loss since every sentence of sentential logic is logically equivalent to such a sentence.

Definition 5.1. The set Λ_0 of axioms of sentential logic consists of all sentences of the following forms:

(a) $(\phi \to (\psi \to \phi))$
(b) $(\phi \to (\psi \to \chi)) \to ((\phi \to \psi) \to (\phi \to \chi))$
(c) $((\neg\psi \to \neg\phi) \to ((\neg\psi \to \phi) \to \psi))$

Definition 5.2. Let $\Gamma \subseteq Sn$. A deduction form Γ is a finite sequence $\phi_0, \ldots, \phi_n$ of sentences such that for each $i \leq n$ one of the following holds:

(i) $\phi_i \in \Lambda_0 \cup \Gamma$
(ii) there are $j, k < i$ such that $\phi_k = (\phi_j \to \phi_i)$.

We say ϕ is *deducible from* Γ, written $\Gamma \vdash \phi$, iff there is a deduction $\phi_0, \ldots, \phi_n$ from Γ with $\phi = \phi_n$.

The following is easily verified.

Lemma 5.1. *(Soundness) If* $\Gamma \vdash \phi$ *then* $\Gamma \models \phi$.

To prove the completeness of the system we assume that $\Gamma \nvdash \phi$ and show that $\Gamma \cup \{\neg\phi\} \subseteq \Gamma^*$ for some complete set Γ^*, and thus $\Gamma \cup \{\neg\phi\}$ is satisfiable and so $\Gamma \not\models \neg\phi$.

To explain what is going on in this argument we introduce the syntactical concept corresponding to satisfiability.

Definition 5.3. Let $\Sigma \subseteq Sn$. Σ is consistent iff there is *no* sentence ϕ such that $\Sigma \vdash \phi$ and $\Sigma \vdash \neg\phi$.

Soundness easily implies that a satisfiable set Σ is consistent. The converse is proved by showing that if Σ is consistent then $\Sigma \subseteq \Gamma$ for some complete set Γ. This is similar to the argument in the preceding section for compactness–the lemma needed is as follows:

LEMMA 5.2. *Assume Σ is consisten and let θ be any sentence. Then at least one of $\Sigma \cup \{\theta\}, \Sigma \cup \{\neg\theta\}$ is consistent.*

To see that this yields completeness, we need to show that $\Gamma \cup \{\neg\phi\}$ is consistent provided $\Gamma \nvdash \phi$. This uses the follwoing fact (the Deduction Theorem–also used in the preceding lemma):

PROPOSITION 5.3. *For any Γ, ϕ, ψ the follwoing are equivalent:*

$$\Gamma \vdash (\phi \to \psi), \ \Gamma \cup \{\phi\} \vdash \psi.$$

We will look more closely at deductions in the context of predicate logic.

6. Exercises

DEFINITION 6.1. A set Σ of sentences is *independent* iff there is no sentence $\sigma \in \Sigma$ such that $(\Sigma \smallsetminus \{\sigma\}) \models \sigma$.

DEFINITION 6.2. Sets Σ_1 and Σ_2 of sentences are *equivalent* iff $\Sigma_1 \models \Sigma_2$ and $\Sigma_2 \models \Sigma_1$.

(1) Let $\Sigma = \{(S_n \vee S_{n+1}) : n \in \omega\}$. Prove or disprove: Σ is independent.
(2) Let $\Sigma = \{(S_{n+1} \to S_n) : n \in \omega\}$. Decide whether or not Σ is independent.
(3) Prove or disprove (with a counterexample) each of the following, where the sentences belong to sentential logic:
   (a) if $\varphi \models \theta$ and $\psi \models \theta$ then $(\varphi \vee \psi) \models \theta$;
   (b) if $(\varphi \wedge \psi) \models \theta$ then $\varphi \models \theta$ or $\psi \models \theta$.
(4) For any expression α let $s(\alpha)$ be the number of occurences of sentence symbols in α and let $c(\alpha)$ be the number of occurences of binary connectives in α. Prove that for every $\sigma \in \mathrm{Sn}$ we have $s(\sigma) = c(\sigma) + 1$
(5) Prove Lemma 1.2.3 about proper initial segments of sentences. [Hint: *Why* will a proper initial segment of a sentence not be a sentence?]
(6) Decide, as efficiently as possible, whether or not
$$\{((C \to B) \to (A \to \neg D), ((B \to C) \to (D \to A))\} \models (B \to \neg D).$$
(7) Prove that every sentence σ in which no sentence symbol occurs more than once is satisfiable, but that no such sentence is a tautology.
(8) Assume Σ is a finite set of sentences. Prove that there is some $\Sigma' \subseteq \Sigma$ such that Σ' is independent and Σ and Σ' are equivalent.
(9) Let Σ be an arbitrary set of sentences. Prove that there is some Σ' such that Σ' is independent and Σ and Σ' are equivalent.
(10) Prove Lemma 1.5.3. [Since this is a lemma used to prove the Compactness Theorem, Theorem 1.5.1, you may not use this theorem in the proof.]
(11) Assume that $\sigma \models \varphi_k$ for all $k \in \omega$. Prove that there is some $n \in \omega$ such that $\varphi_0 \wedge \cdots \wedge \varphi_n \models \varphi_k$ for all $k \in \omega$.

(12) Give an example of a satisfiable sentence σ and sentences φ_k for $k \in \omega$ such that $\sigma \models \varphi_k$ for all $k \in \omega$ but there is no $n \in \omega$ such that $\varphi_0 \wedge \cdots \wedge \varphi_n \models \varphi_k$ for all $k \in \omega$.

(13) Assume that σ and φ_k are given so that for every assignment h we have

$$h \models \sigma \text{ iff } (h \models \varphi_k \text{ for every } k \in \omega).$$

Prove that there is some $n \in \omega$ such that $\varphi_0 \wedge \cdots \wedge \varphi_n \models \varphi_k$ for all $k \in \omega$.

CHAPTER 2

First-Order Logic

0. Introduction

In mathematics we investigate the properties of mathematical structures. A mathematical structure consists of some set A of objects (the *domain*, or universe, of the structure) together with some functions and/or relations on the domain–both must be specified to completely determine the structure. Thus the set $\mathbb{Z}$ of all integers can be the domain of many different structures on $\mathbb{Z}$ in which the functions $+$ and - are given; the ring structure in which also multiplication is considered; the (pure) order structure in which the relation $\leq$ is given, but no functions; the ordered group structure in which $\leq, +$, and $-$ are included; etc.

In all these possible structures one considers not just the functions and relations acutally listed, but also the functions and relations which are generated or *defined* from them in certain ways. In practice, the allowable ways of generating more functions and relations may be left vague, but in our formal systems we need to be precise on this point. Certainly, in all cases we would be allowed to form compositions of given functions obtaining, for example, polynomials like $x{\cdot}x{-}y{+}x{\cdot}z$ in the ring structure of $\mathbb{Z}$. Normally constant functions would also be allowed, tus obtaining all polynomials with integer coefficients in this example.

Similarly one can compose relations with functions obtaining, for example, relations like $(x + x) \leq y \cdot z$ in the ordered ring structure. Equality would also normally be used regardless of whether it was explicitly listed. Connectives like $\neg, \wedge, vee$ would enable us to form further relations. For example from binary relations $R(x, y), S(x, y)$ on A we define relations $\neg R(x, y)$, the relation which holds if R fails; $R(x, y) \wedge S(x, y)$, the relation which holds iff both R and S hold; etc.

In the ring structure on $\mathbb{Z}$ we would have, for example, the binary relation $R(x, y)$ which holds iff $x = y \cdot y$. Thus $R(1, 1), R(4, 2)$ would hold, $R(2, 1)$ would fail, etc. We would certainly also consider the new relation $P(x)$ which holds iff $R(x, y)$ holds for some y in the domain–$P(x)$ iff $x = y \cdot y$ for some $y \in \mathbb{Z}$ in this example. And from $\neg R(x, y)$ we can define $Q(x)$ which holds iff $\neg R(x, y)$ holds for all y in the domain–$Q(x)$ iff $x \neg y \cdot y$ for all $y \in \mathbb{Z}$ in this example.

Finally the statements made about a structure would be statements concerning the relations considered–for example, the statements that $P(x)$ holds for some x in the domain (true in this example) or that $P(x)$ holds for every x in the domain (flase in this example but true if the domain is enlarged from $\mathbb{Z}$ to the complex numbers). Normally we would also be allowed to refer to specific elements of the domain and make, for example, the statements that $P(4)$ holds or $Q(3)$ holds–both true in this example.

Our formal systems of first order logic are designed to mirror this process. Thus the *symbols* of a first order language will include symbols for functions, for relations, and for fixed elements ("constants") of a domain. Among the expressions we will pick out some which will define functions on a domain–and these functions will include the listed functions and be closed under composition. Similarly other expressions will define relations on a domain–and these relations will be closed under the operations outlined above. Finally, the *sentences* of the language will make assertions as indicated above about the definable relations.

Some important points to notice: first of all, there will be many different languages according to the selection of (symbols for) functions, relations, and constants made. Secondly, a given language may be interpreted on *any* domain, with *any* choice of functions, relations and elements consistent with the symbols–thus we will never have a language which *must* be interpreted on the domain $\mathbb{Z}$ or with a symbol which *must* be interpreted as $+$, for example.

1. Formulas of First Order Logic

We follow the outline in the previous section in defining the symbols of a first order language, the terms (which correspond to the functions) and the formulas (which correspond to the relations). In constructing formulas we use the symbols $\forall$ and $\exists$ for the quantifiers "for every" and "there is some" and we use $\equiv$ for the special relation of equality or identity which is in every first order language.

Definition 1.1. The symbols of a first order language $\mathcal{L}$comprise the following:

1) for each $m > 0$ some set (perhaps empty) of m-ary function symbols;
2) some set (perhaps empty) of individual constant symbols;
3) for each $m > 0$ some set (perhaps empty) of m-ary relation symbols;

3a) the binary relation symbol for equality: $\equiv$;

4) a (countably infinite) list of individual variables: $v_0, \dots, v_n, \dots$ for all $n \in \omega$;
5) the sentential connectives: $\neg, \wedge, \vee, \rightarrow$;
6) the quantifiers: $\forall, \exists$;
7) parentheses: $($, $)$.

We will use (perhaps with sub- or superscripts) letters like F, G for function symbols, c, d for constant symbols and R, S for relation symbols. Anticipating the formal definition of $\mathcal{L}$-structure in the next section, an interpretation of $\mathcal{L}$consists of a non-empty set A (the domain or universe of the interpretation) and for each m-ary function symbol F an m-ary function F^* on A, for each constant symbol c an element c^* of A, and for each m-ary relation symbol R an m-ary relation R^* on A–however $\equiv$ is always interpreted as actual equality on A. The variables will range over elements of A and quantification is over A.

The symbols listed in 3a)-7) are the same for all first order languages and will be called the *logical* symbols of $\mathcal{L}$. The symbols listed in 1)-3) will vary from language to language and are called the *non-logical* symbols of $\mathcal{L}$. We will write $\mathcal{L}^{nl}$for the set of non-logical symbols of $\mathcal{L}$. In specifying a language $\mathcal{L}$it suffices to specify $\mathcal{L}^{nl}$. Note that the smallest language will have $\mathcal{L}^{nl}= \emptyset$. Note also that to determine $\mathcal{L}$one cannot just specify the *set* $\mathcal{L}^{nl}$but must also specify what *type* of symbol each is, such as a binary function symbol.

The terms of $\mathcal{L}$will be those expressions of $\mathcal{L}$which will define functions in any interpretation. These functions are built from the (interpretations of the) function symbols by composition. In addition we can use any constant symbol of $\mathcal{L}$in defining these functions, and we consider a variable v_n standing alone as defining the identity function. We also allow the "limiting case" of a function of zero arguments as a function. We thus have the following definition.

DEFINITION 1.2. For any first order language $\mathcal{L}$the set $\mathrm{Tm}_{\mathcal{L}}$ of terms of $\mathcal{L}$is defined as follows: (1) $v_n \in \mathrm{Tm}_{\mathcal{L}}$ for every $n \in \omega$, $c \in \mathrm{Tm}_{\mathcal{L}}$ for every constant symbol of c of $\mathcal{L}$, (2) if F is an m-ary function symbol of $\mathcal{L}$and $t_1, \ldots, t_m \in \mathrm{Tm}_{\mathcal{L}}$ then $Ft_1 \ldots t_m \in \mathrm{Tm}_{\mathcal{L}}$.

This is, of course, a definition by recursion with the last clause "noting else is a term" understood. The reader may be surprised that we have not written $F(t_1, \ldots, t_m)$ but this is not required for unique readability (although it would certainly help practical readability at times).

Just as with sentences of sentential logic we have a theorem justifying proof by induction on terms, whose proof we leave to the reader.

THEOREM 1.1. *Let* $X \subseteq \mathrm{Tm}_{\mathcal{L}}$ *and assume that (a)* $v_n \in X$ *for all* $n \in \omega$*,* $c \in X$ *for every constant symbol* c *of* $\mathcal{L}$*, and (b) whenever* F *is an m-ary function symbol of* $\mathcal{L}$*and* $t_1, \ldots, t_m \in X$ *then* $Ft_1 \ldots t_m \in X$*. Then* $X = \mathrm{Tm}_{\mathcal{L}}$*.*

Even without parentheses every term is uniquely readable, as we leave to the reader to establish.

THEOREM 1.2. *For each* $t \in \mathrm{Tm}_{\mathcal{L}}$ *with* $lh(t) > 1$ *there is exactly one choice of* $m > 0$*, m-ary function symbol* F *of* $\mathcal{L}$*and* $t_1, \ldots, t_m \in \mathrm{Tm}_{\mathcal{L}}$ *such that* $t = Ft_1, \ldots, t_m$*.*

And finally, with unique readability we can define functions on $\mathrm{Tm}_{\mathcal{L}}$ by recursion. We leave the formulation and proof of this to the reader.

In defining the class of formulas of first order logic we start with the formulas obtained by "composing" the given relation (symbols) with terms.

DEFINITION 1.3. The atomic formulas of $\mathcal{L}$are the expressions of the form $Rt_1 \ldots t_m$ for m-ary relation symbols $R \in \mathcal{L}$ and $t_1, \ldots, t_m \in \mathrm{Tm}_{\mathcal{L}}$.

The atomic formulas are the basic building blocks for formulas, just as sentence symbols were the building blocks for sentences in sentential logic.

DEFINITION 1.4. For any first order language $\mathcal{L}$the set $\mathrm{Fm}_{\mathcal{L}}$ of formulas of $\mathcal{L}$is defined as follows:

1) if ϕ is an atomic formula of $\mathcal{L}$, then $\phi \in \mathrm{Fm}_{\mathcal{L}}$,
2) if $\phi \in \mathrm{Fm}_{\mathcal{L}}$ then $(\neg\phi) \in \mathrm{Fm}_{\mathcal{L}}$,
3) if $\phi, \psi \in \mathrm{Fm}_{\mathcal{L}}$ then $(\phi \star \psi) \in \mathrm{Fm}_{\mathcal{L}}$
   for any binary connective $\star$,
4) if $\phi \in \mathrm{Fm}_{\mathcal{L}}$ then $\forall v_n\phi, \exists v_n\phi \in \mathrm{Fm}_{\mathcal{L}}$ for every $n \in \omega$

Note that atomic formulas do not have length 1; in fact in some languages there will be arbitrarily long atomic formulas. Nevertheless induction on length yields the following principle of proof by induction in which the atomic formulas are the base case.

THEOREM 1.3. *Let* $X \subseteq \mathrm{Fm}_{\mathcal{L}}$ *and assume that: (a)* $\phi \in X$ *for every atomic formula* ϕ *of* $\mathcal{L}$*, (b)* $\phi \in X$ *implies* $(\neg\phi) \in X$*, (c)* $\phi, \psi \in X$ *implies that* $(\phi \star \psi) \in X$ *for binary connectives* $\star$*, (d)* $\phi \in X$ *implies* $\forall v_n\phi, \exists v_n\phi \in X$ *for every* $n \in \omega$*. Then* $X = \mathrm{Fm}_{\mathcal{L}}$*.*

As with terms, or sentences of sentential logic, both unique readability and a principle of definition by recursion hold for $\mathrm{Fm}_{\mathcal{L}}$. We leave both the formulation and proof of these to the reader.

We give here some examples of terms and formulas in particular first order languages.

(1) $\mathcal{L}^{nl}= \emptyset$. Here $\mathrm{Tm}_{\mathcal{L}} = \{v_n | \ n \in \omega\}$. Since $\equiv$, being a logical symbol, belongs to every first order language, the atomic formulas consist of the expressions $\equiv v_n v_k$ for $n, k \in \omega$. Specific formulas then include $(\neg \equiv v_0 v_1)$, $\exists v_1 (\neg \equiv v_0 v_1)$, $((\equiv v_0 v_1 \vee \equiv v_0 v_2) \vee (\equiv v_1 v_2))$, $\forall v_0 \exists v_1 (\neg \equiv v_0 v_1)$, $\forall v_0 \forall v_1 \forall v_2 ((\equiv v_0 v_1 \vee \equiv v_0 v_2) \vee (\equiv v_1 v_2))$.

An interpretation for this language will be determined by some $A \neq \emptyset$ as its domain. We will always interpret $\equiv$ as equality ("identity") on the domain. It is thus clear, for example, that the formula $(\neg \equiv v_0 v_1)$ will define the relation $R^*(x, y)$ on A such that $R^*(a, a')$ holds iff $a \neq a'$. Similarly the formula $\exists v_1 (\neg \equiv v_0 v_1)$ will define the unary relation $P^*(x)$ on A such that $P^*(a)$ holds iff there is some $a' \in A$ such that $R^*(a, a')$ holds, i.e. $a \neq a'$. Note that $P^*(a)$ will hold of no elements a of A.

(2) $\mathcal{L}^{nl}= \{R, F, c\}$ where R is a binary relation symbol, F a unary function symbol and c is a constant symbol. Now the terms of $\mathcal{L}$also include c, Fv_n, Fc, FFv_n, FFc, etc. The atomic formulas consist of all expressions $\equiv t_1 t_2$ and $Rt_1 t_2$ for $t_1, t_2 \in \mathrm{Tm}_{\mathcal{L}}$–for example $\equiv cFv_1$, $Rv_0 Fv_0$, Rcv_1. Further formulas will include $(\neg \equiv cFv_1)$, $Rv_0 v_1 \to RFv_0 Fv_1)$, $\exists v_1 \equiv v_0 Fv_1$, $\forall v_1 Rcv_1$.

One familiar interpretation for this language will have domain $A = \omega$, interpret R as $\leq$, F as immediate successor, and c as 0. That is $R^*(k, l)$ holds iff $k \leq l$, $F^*(k) = k + 1$, $c^* = 0$. The term FFv_n will ben define the function $(FFv_n)^*$ defined as $(FFv_n)^*(k) = F^*(F^*(k)) = k + 2$ for all $k \in \omega$. The term FFc will define the particular element $F^*(F^*(0)) = 2$ of ω The formula $\exists v_1 \equiv v_0 Fv_1$ will define the unary relation on ω which holds of k iff $k = F^*(l)$ for some $l \in \omega$, that is, iff $k = l + 1$ for some $l \in \omega$, thus iff $k \neq 0$.

Giving a precise definition of how terms and formulas are interpreted in complete generality is far from easy. One problem is that the relation defined, for example, by the formula $(\phi \vee \psi)$ is not just determined by the relations defined by ϕ and by ψ separately, but also depends on the variables used in ϕ and in ψ and on how they are related. Thus, we have pointed out that for any choice of distinct variables v_n, v_k the formula $(\neg \equiv v_n v_k)$ will define the binary relation $R^*(x, y)$ such that $R^*(a, a')$ holds iff $a \neq a'$. But the formula $((\neg \equiv v_n v_k) \vee (\neg \equiv v_m v_l))$ could define either a binary or ternary or 4-ary relation depending on the variables. The situation is even more complicated in our second example with the formulas $(Rv_0 v_1 \vee Rv_1 v_2)$, $(Rv_0 v_1 \vee Rv_2 v_1)$, $(Rv_0 v_2 \vee Rv_1 v_2)$ etc. all defining different ternary relations.

Our solution here is to realize that the interpretation of a term or formula depends not only on the term or formula itself but is also dependant on the choice of a particular list of variables in a specific order. Thus in addition to beig interpreted as the binary relation R^* on A, the formulas $Rv_0 v_1$ and $Rv_1 v_2$ can each be interpreted as ternary relations relative to the list v_0, v_1, v_2 of variables. $Rv_0 v_1$ would then be the relation S_0^* such that $S_0^*(a, a', a'')$ holds iff $R^*(a, a')$ holds, and $Rv_1 v_2$ would then be the relation S_1^* such that $S_1^*(a, a', a'')$ holds iff $R^*(a', a'')$ holds. We can then say that $(Rv_0 v_1 \vee Rv_1 v_2)$ is interpreted by the ternary relation $S_0^*(x, y, z) \vee S_1^*(x, y, z)$.

What variables *must* occur in a list so that a term t or a formula ϕ will define a function or relation relative to that list? Clearly for terms this would be just the list of all variables occurring in the term. The answer for formulas is less obvious. We have pointed out, for example, that the formula $\exists v_1 \equiv v_0 F v_1$ defines a unary relation on A, despite having two variables. The reason, of course, is that the variable v_1 is quantified and so the formula should express a property of v_0 alone. Unfortunately the same variable can be both quantified and not quantified in the same formula, as shown (for example) by

$$(Rv_1v_0 \to \exists v_1 \equiv v_0 F v_1).$$

This formula must be interpreted by (at least) a binary relation, since the first occurrence of v_1 is not bound by any quantifier.

We are thus lead to the following definition of the variables which occur *free* in a formula, and which must therefore be among the variables listed when considering the relation. The formula defines in an interpretation.

Definition 1.5. For any $\phi \in \mathrm{Fm}_{\mathcal{L}}$ the set $Fv(\phi)$ of variables occurring free in ϕ is defined as follows:

1) if ϕ is atomic then $Fv(\phi)$ is the set of all variables occuring in ϕ;
2) $Fv(\neg\phi) = Fv(\phi)$;
3) $Fv(\phi \star \psi) = Fv(\phi) \cup Fv(\psi)$;
4) $Fv(\exists v_n \phi) = Fv(\forall v_n \phi) = Fv(\phi) - \{v_n\}$;

Thus in any interpretation a formula ϕ will define a relation in the list of its free variables. If ϕ has no free variables then it will simply be either true or false in any interpretation, which justifies the following definition.

Definition 1.6. The set $\mathrm{Sn}_{\mathcal{L}}$ of sentences of $\mathcal{L}$is $\{\phi \in \mathrm{Fm}_{\mathcal{L}} |\ Fv(\phi) = \emptyset\}$.

We need to have a notation which will exhibit explicitly the list of variables considered in interpreting a term or formula.

Definition 1.7. 1) For any $t \in \mathrm{Tm}_{\mathcal{L}}$ we write $t = t(x_1, \ldots, x_n)$ provided $\{x_1, \ldots, x_n\}$ contains all variables occurring in t. 2) For any $\phi \in \mathrm{Fm}_{\mathcal{L}}$ we write $\phi = \phi(x_1, \ldots, x_n)$ provided $Fv(\phi) \subseteq \{x_1, \ldots, x_n\}$.

We emphasize that the term or formula in question does not determine the list of variables nor the order in which they occur. Thus, if ϕ is $\exists v_1 \equiv v_0 F v_1$ then we could have any of the following: $\phi = \phi(v_0)$, $\phi = \phi(v_0, v_3)$, $\phi = \phi(v_3, v_0)$, $\phi = \phi(v_0, v_1, v_2)$, etc. The list of variables will determine the *arity* of the function or relation defined in any interpretation, and the *order* in which the arguments are taken from the variables.

Consider $\phi(v_0) = \exists v_1 \equiv v_0 F v_1$. In any interpretation $\phi(v_0)$ will define the set (i.e. unary relation) consisting of all $a \in A$ for which $a = F^*(a')$ for some $a' \in A$. Let $\sigma = \exists v_1 \equiv cFv_1$. Then σ is a sentence and σ will be true in an interpretation iff c^* belongs to the set $(\phi(v_0))^*$ defined by $\phi(v_0)$. It is natural to express this by saying "c satisfies ϕ" and to write σ as $\phi(c)$. Our definition of substitution will justify this usage.

Definition 1.8. a) Let $t \in \mathrm{Tm}_{\mathcal{L}}$, x a variable and $s \in \mathrm{Tm}_{\mathcal{L}}$. Then t^x_s is the term formed by replacing all occurrences of x in t by s. b) Let $\phi \in \mathrm{Fm}_{\mathcal{L}}$, x a variable and $t \in \mathrm{Tm}_{\mathcal{L}}$. Then ϕ^x_t is the result of replacing all *free* occurrences of x in ϕ by the term t–formally:

1) ϕ^x_t is ϕ with all occurrences of x replaced by t if ϕ is atomic;
2) $(\neg\phi)^x_t = (\neg(\phi^x_t))$;
3) $(\phi \star \psi)^x_t = (\phi^x_t \star \psi^x_t)$ for $\star$ a binary connective;
4) $(\exists v_n \phi)^x_t = \exists v_n \phi$ if $x = v_n$ or $\exists v_n(\phi^x_t)$ if $x \neq v_n$;
5) Similarly for $(\forall v_n \phi)^x_t$.

In particular, if $t = t(x)$ we write $t(s)$ for t^x_s, and if $\phi = \phi(x)$ we will write $\phi(t)$ for ϕ^x_t.

More generally we can define $t^{x_1,\ldots,x_n}_{t_1,\ldots,t_n}$ and $\phi^{x_1,\ldots,x_n}_{t_1,\ldots,t_n}$ as the results of simultaneously substituting $t_1, \ldots, t_n$ for all (free) occurrences of $x_1, \ldots, x_n$ in t, ϕ respectively. Note that we may have

$$(\phi^{x_1}_{t_1})^{x_2}_{t_2} \neq \phi^{x_1 x_2}_{t_1 t_2} \neq (\phi^{x_2}_{t_2})^{x_1}_{t_1}.$$

If $t = t(x_1, \ldots, x_n)$ and $\phi = \phi(x_1, \ldots, x_n)$ then we will write $t(t_1, \ldots, t_n)$ for $t^{x_1,\ldots,x_n}_{t_1,\ldots,t_n}$ and $\phi(t_1, \ldots, t_n)$ for $\phi^{x_1,\ldots,x_n}_{t_1,\ldots,t_n}$.

2. Structures for First Order Logic

First order languages are interpreted on (mathematical) structures, among which we will find the usual structures studied in mathematics. The abstract definition, which is very close to the informal definition from the preceding section, is as follows.

Definition 2.1. A structure for a first order language $\mathcal{L}$is a pair $\mathfrak{A} = (A, \mathcal{I})$ where A is some non-empty set (called the *universe* or domain of the structure) and $\mathcal{I}$ is a function whose domain is $\mathcal{L}^{nl}$satisfying the following: (1) if F is an m-ary function symbol of $\mathcal{L}$then $\mathcal{I}(F) = F^{\mathfrak{A}}$ is an m-ary function defined on all of A and having values in A, (2) if c is a constant symbol of $\mathcal{L}$then $\mathcal{I}(c) = c^{\mathfrak{A}} \in A$, (3) if R is an m-ary relation symbol of $\mathcal{L}$then $\mathcal{I}(R) = R^{\mathfrak{A}}$ is an m-ary relation on A.

Note that $\equiv$ is not in the domain of $\mathcal{I}$ since it is a logical symbol, so it does not make sense to refer to $\mathcal{I}(\equiv)$ or $\equiv^{\mathfrak{A}}$. We also point out that the functions interpreting the function symbols are total–thus a binary function symbol cannot be interpreted, for example, as unary on ω.

We customarily use German script letters $\mathfrak{A}, \mathfrak{B}, \ldots$ to refer to structures, perhaps with sub- or superscripts. By convention the universe of a structure is denoted by the corresponding capital Latin letter, with the same sub- or superscript.

In practice we suppress reference to $\mathcal{I}$ and just give its values. Thus if $\mathcal{L}^{nl}= \{R, F, c\}$ where R is a binary relation symbol, F is a unary function symbol, and c is a constant symbol, we might specify a structure for $\mathcal{L}$as follows: $\mathfrak{A}$ is the structure whose universe is ω such that $R^{\mathfrak{A}}(k, l)$ holds iff $k \leq l$, $F^{\mathfrak{A}}(k) = k + 1$ for all k and $c^{\mathfrak{A}} = 0$. When the specific symbols involved are clear, we may just write the sequence of values of $\mathcal{I}$ in place of $\mathcal{I}$. Thus the preceding example could be written as $\mathfrak{A} = (\omega, \leq, s, 0)$ where $s : \omega \to \omega$ is the (immediate) successor function.

A structure is a structure for *exactly one* language $\mathcal{L}$. If $\mathcal{L}_1$ and $\mathcal{L}_2$ are different languages then *no* $\mathcal{L}_1$-structure can also be an $\mathcal{L}_2$-structure. Thus if $\mathcal{L}_1^{nl} = \{R, F, c\}$ as above and $\mathcal{L}_2^{nl} = \{S, G, d\}$ where S is a binary relation symbol, G is a unary function symbol and d is a constant symbol, then one $\mathcal{L}_1$-structure is $\mathfrak{A}$ given above. An $\mathcal{L}_2$-structure could be $\mathfrak{B}$ with universe ω, with S interpreted as $\leq$, G as the successor function and $d^{\mathfrak{B}} = 0$. Informally we could express $\mathfrak{B}$ as $\mathfrak{B} = (\omega, \leq, s, 0)$– but $\mathfrak{A}$ and $\mathfrak{B}$ are totally different structures since the symbols interpreted by $\leq, s$, and 0 are different. If $\mathcal{L}_3 = \mathcal{L}_1 \cup \mathcal{L}_2$, so $\mathcal{L}_3^{nl} = \mathcal{L}_1^{nl} \cup \mathcal{L}_2^{nl}$, then one $\mathcal{L}_3$-structure would be $\mathfrak{A}^*$ with universe ω, both R and S interpreted as $\leq$, both F and G as s and $c^{\mathfrak{A}^*} = d^{\mathfrak{A}^*} = 0$. It would be possible, but confusing, to write

$$\mathfrak{A}^* = (\omega, \leq, s, 0, \leq, s, 0).$$

There is, however, one very important relation between structures in different languages, in which one structure is a *reduct* of the other to a smaller language. In this case the structures are equivalent as far as the smaller language is concerned and can be used interchangeably.

DEFINITION 2.2. Let $\mathcal{L}_1$ and $\mathcal{L}_2$ be first order languages with $\mathcal{L}_1 \subseteq \mathcal{L}_2$ (equivalently $\mathcal{L}_1^{nl} \subseteq \mathcal{L}_2^{nl}$). Let $\mathfrak{A}$ be an $\mathcal{L}_1$-structure, $\mathfrak{B}$ an $\mathcal{L}_2$-structure. Then $\mathfrak{A}$ is the reduct of $\mathfrak{B}$ to $\mathcal{L}_1$, and $\mathfrak{B}$ is an expansion of $\mathfrak{A}$ to $\mathcal{L}_2$, iff $\mathfrak{A}$ and $\mathfrak{B}$ have the same universe and they interpret all symbols in $\mathcal{L}_1^{nl}$ precisely the same. We write $\mathfrak{A} = \mathfrak{B} \restriction \mathcal{L}_1$ if $\mathfrak{A}$ is the reduct of $\mathfrak{A}$ to $\mathcal{L}_1$.

Thus in the above examples of $\mathcal{L}_1$-, $\mathcal{L}_2$- and $\mathcal{L}_3$-structures, $\mathfrak{A} = \mathfrak{A}^* \restriction \mathcal{L}_1$ and $\mathfrak{B} = \mathfrak{A}^* \restriction \mathcal{L}_2$. Note that in spite of the terminology "expansion" that the *universe* of a structure remains fixed when passing to an expansion–only the language is expanded.

One of the most important special cases of an expansion of a structure occurs when we add (new) constant symbols so as to name some elements of the structure.

DEFINITION 2.3. Let $\mathcal{L}$be a first order language, let $\mathfrak{A}$ be an $\mathcal{L}$-structure and let $X \subseteq A$.

(a) $\mathcal{L}(X) = \mathcal{L} \cup \{c_a |\ a \in X\}$ is the language obtained from $\mathcal{L}$by adding a new constant symbol c_a for each $a \in X$.

(b) $\mathfrak{A}_X$ is the expansion of $\mathfrak{A}$ to an $\mathcal{L}(X)$-structure such that

$$c_a^{\mathfrak{A}_X} = a \text{ for all } a \in X.$$

In particular, $\mathfrak{A}_A$ is the expansion of $\mathfrak{A}$ obtained by adding constants for every element of A. In ordinary mathematical practice the structures $\mathfrak{A}$ and $\mathfrak{A}$ would not be distinguished–in talking about $\mathfrak{A}$ you would naturally want to talk about arbitrary elements of A, which means having constants for them in your language when you formalize.

We will also take the point of view that in talking about $\mathfrak{A}$ you will frequently wish to refer to specific elements of $\mathfrak{A}$, but we will always carefully distinguish $\mathfrak{A}_A$ from $\mathfrak{A}$.

We also emphasize that there is no way that we could–or would want to if we could–ensure at the outset that $\mathcal{L}$contained constants to name every element of every $\mathcal{L}$-structure. Since there are $\mathcal{L}$-structures $\mathfrak{A}$ with $|A| > |\mathcal{L}|$ the first point is clear. For the second, recall the language $\mathcal{L}$above with $\mathcal{L}^{nl}$= $\{R, F, c\}$ and the $\mathcal{L}$-structure $\mathfrak{A} = \{\omega, \leq, s, 0\}$. Another $\mathcal{L}$-structure one would naturally wish to consider would be $\mathfrak{B} = (\mathbb{Z}, \leq, s, 0)$. But if $\mathcal{L}$had constants to refer to every element of $\mathbb{Z}$ then those constants naming negative integers could not be interpreted in $\mathfrak{A}$, i.e. as elements of ω, in any natural way.

To recapitulate, a language $\mathcal{L}$determines the class of $\mathcal{L}$-structures, whose universes are arbitary (in particular arbitrarily large) non-empty sets. In studying any particular $\mathcal{L}$-structure $\mathfrak{A}$, we will customarily pass to the language $\mathcal{L}(A)$ and the expansion $\mathfrak{A}_A$). but in comparing two different $\mathcal{L}$-structures $\mathfrak{A}, \mathfrak{B}$ we must use properties expressible *in* $\mathcal{L}$since $\mathcal{L}(A)$ will not normally have any "natural" interpretation on $\mathfrak{B}$ nor will $\mathcal{L}(b)$ normally have any "natural" interpretation on $\mathfrak{A}$.

We now proceed to give a very intuitively natural definition of the truth of a sentence *of* $\mathcal{L}$on $\mathfrak{A}_A$. Since every sentence of $\mathcal{L}$is also a sentence of $\mathcal{L}(A)$ this definition will, in particular, determine when a sentence *of* $\mathcal{L}$is true on $\mathfrak{A}_A$. And since $\mathfrak{A}$ and $\mathfrak{A}_A$ are identical as far as $\mathcal{L}$is concerned, we will take this as the definition of the truth of a sentence of $\mathcal{L}$on the given $\mathcal{L}$-structure $\mathfrak{A}$.

An atomic formula $Rt_1 \dots t_m$ (or $\equiv t_1 t_2$) is a sentence iff the terms $t_1, \dots, t_m$ contain no variables. We will want to say that $Rt_1 \dots t_m$ is true on $\mathfrak{A}_A$ iff the relation $R^{\mathfrak{A}}$ (equivalently $R^{\mathfrak{A}_A}$) holds of the elements of A which are named by the terms $t_1, \dots, t_m$. If $\mathcal{L}$has function symbols we need to first give a definition by recursion stating how terms without variables (also called *closed terms*) are evaluated.

Definition 2.4. Given an $\mathcal{L}$-structure $\mathfrak{A}$ we define the interpretation $t^{\mathfrak{A}_A}$ of closed terms t of $\mathcal{L}(A)$ in $\mathfrak{A}_A$ as follows:

(1) if t is a constant symbol c of $\mathcal{L}(A)$ then $t^{\mathfrak{A}_A} = c^{\mathfrak{A}_A}$;
(2) if $t = Ft_1 \dots t_m$ for closed terms $t_1, \dots, t_m$ of $\mathcal{L}(A)$ then $t^{\mathfrak{A}_A} = F^{\mathfrak{A}}(t_1^{\mathfrak{A}_A}, \dots, t_m^{\mathfrak{A}_A})$.

Definition 2.5. Given an $\mathcal{L}$-structure $\mathfrak{A}$ we define the truth value $\theta^{\mathfrak{A}_A}$ of sentences θ of $\mathcal{L}(A)$ in $\mathfrak{A}_A$ so that $\theta^{\mathfrak{A}_A} \in \{T, F\}$ as follows:

1) if θ is $Rt_1 \dots t_m$ for closed terms $t_1, \dots, t_m$ and $R \in \mathcal{L}^{nl}$then $\theta^{\mathfrak{A}_A} = T$ iff $R^{\mathfrak{A}}(t_1^{\mathfrak{A}_A}, \dots, t_m^{\mathfrak{A}_A})$ holds;
2) if θ is $\equiv t_1 t_2$ for closed terms t_1, t_2 then $\theta^{\mathfrak{A}_A} = T$ iff $t_1^{\mathfrak{A}_A} = t_2^{\mathfrak{A}_A}$;
3) if $\theta = \neg\phi$ then $\theta^{\mathfrak{A}_A} = T$ iff $\phi^{\mathfrak{A}_A} = F$;
4) if $\theta = (\phi \wedge \psi)$ then $\theta^{\mathfrak{A}_A} = T$ iff $\phi^{\mathfrak{A}_A} = \psi^{\mathfrak{A}_A} = T$;
5) if $\theta = (\phi \vee \psi)$ then $\theta^{\mathfrak{A}_A} = F$ iff $\phi^{\mathfrak{A}_A} = \psi^{\mathfrak{A}_A} = F$;
6) if $\theta = (\phi \to \psi)$ then $\theta^{\mathfrak{A}_A} = F$ iff $\phi^{\mathfrak{A}_A} = T$ and $\psi^{\mathfrak{A}_A} = F$;
7) if $\theta = \forall v_n \phi$ then $\phi = \phi(v_n)$ and $\theta^{\mathfrak{A}_A} = T$ iff $\phi(c_a)^{\mathfrak{A}_A} = T$ for all $a \in A$;
8) if $\theta = \exists v_n \phi$ then $\phi = \phi(v_n)$ and $\theta^{\mathfrak{A}_A} = T$ iff $\phi(c_a)^{\mathfrak{A}_A} = T$ for some $a \in A$;

Notation 1. Let $\mathfrak{A}$ be an $\mathcal{L}$-structure. (a) If $\theta \in Sn_{\mathcal{L}(A)}$ then $\mathfrak{A}_A \models \theta$, read θ is *true on* $\mathfrak{A}_A$ or $\mathfrak{A}_A$ *satisfies* θ, iff $\theta^{\mathfrak{A}_A} = T$. (b) If $\theta \in \mathrm{Sn}_{\mathcal{L}}$ then $\mathfrak{A} \models \theta$, read θ is *true on* $\mathfrak{A}_A$ or $\mathfrak{A}_A$ *satisfies* θ or $\mathfrak{A}$ is a *model of* θ, iff $\mathfrak{A}_A \models \theta$.

The above definition is designed to capture the "common sense" idea that, say $\exists x \phi(x)$ is true on a structure iff ϕ holds of some element of the structure. We pass to the expanded language precisely so as to be able to express this "common sense" definition using sentences of a formal language.

We extend our notations $t^{\mathfrak{A}_A}, \theta^{\mathfrak{A}_A}$ to arbitrary terms and formulas of $\mathcal{L}(A)$ as follows.

Definition 2.6. Let $t(x_1, \dots, x_n) \in Tm_{\mathcal{L}(A)}$. Then $t^{\mathfrak{A}_A}$ is the function on A defined as follows: for any $a_1, \dots, a_n \in A$, $t^{\mathfrak{A}_A}(a_1, \dots, a_n) = t(c_{a_1}, \dots, c_{a_n})^{\mathfrak{A}_A}$. If t is actually a term of $\mathcal{L}$we write $t^{\mathfrak{A}}$ for the function $t^{\mathfrak{A}_A}$.

Definition 2.7. Let $\phi(x_1, \dots, x_n) \in Fm_{\mathcal{L}(A)}$. Then $\phi^{\mathfrak{A}_A}$ is the n-ary relation of A defined as follows: for any $a_1, \dots, a_n \in A$, $\phi^{\mathfrak{A}_A}(a_1, \dots, a_n)$ holds iff $\mathfrak{A}_A \models \phi(c_{a_1}, \dots, c_{a_n})$. If ϕ is actually a formula of $\mathcal{L}$we write $\phi^{\mathfrak{A}}$ for the relation $\phi^{\mathfrak{A}_A}$.

Just as in the informal discussion in the preceding section, the definitions of the functions $t^{\mathfrak{A}_A}$ and relations $\phi^{\mathfrak{A}_A}$ are relative to the list of variables used, but this ambiguity causes no problems.

DEFINITION 2.8. Given an $\mathcal{L}$-structure $\mathfrak{A}$, an $\mathcal{L}(A)$-formula $\phi(x_1, \ldots, x_n)$ and elements $a_1, \ldots, a_n \in A$ we say that ϕ is satisfied by $a_1, \ldots, a_n$ in $\mathfrak{A}_A$ iff $\mathfrak{A}_A \models \phi(c_{a_1}, \ldots, c_{a_n})$. If ϕ is in fact a formula of $\mathcal{L}$we will say it is satisfied by $a_1, \ldots, a_n$ in $\mathfrak{A}$ instead of $\mathfrak{A}_A$. In each case we will say ϕ is satisfiable in $\mathfrak{A}_A$ or $\mathfrak{A}$ to mean it is satisfied by some $a_1, \ldots, a_n$.

Note that if θ is a sentence of $\mathcal{L}(A)$ then either θ is satisfied by $\mathfrak{A}_A$ or $(\neg\theta)$ is satisfied by $\mathfrak{A}_A$, but not both.

In extreme cases it may make sense to talk of a formula (with free variables) being true on a stucture.

DEFINITION 2.9. Given an $\mathcal{L}$-structure $\mathfrak{A}$ and a formula $\phi = \phi(x_1, \ldots, x_n)$ of $\mathcal{L}(A)$ we say ϕ is true on $\mathfrak{A}_A$, written $\mathfrak{A}_A \models \phi$, iff $\mathfrak{A}_A \models \phi(c_{a_1}, \ldots, c_{a_n})$ for all $a_1, \ldots, a_n \in A$. If ϕ is a formula of $\mathcal{L}$we say ϕ is true on $\mathfrak{A}$ and write $\mathfrak{A} \models \phi$.

We thus see that the following are equivalent: $\mathfrak{A} \models \phi$, $\neg\phi$ is not satisfiable in $\mathfrak{A}$, $\mathfrak{A} \models \forall x_1 \cdots \forall x_n \phi$. At most one of $\mathfrak{A} \models \phi$, $\mathfrak{A} \models \neg\phi$ will hold but in general neither of them will hold.

We proceed to a series of examples, using the language $\mathcal{L}$whose non-logical symbols are precisely a binary relation symbol R and a unary function symbol F.

$\mathfrak{A} \models \equiv xx$ for all $\mathfrak{A}$, since $\equiv$ is interpreted by actual equality in every $\mathcal{L}$-structure. Hence also $\mathfrak{A} \models \forall x \equiv xx$ for all $\mathfrak{A}$.

If x, y are different variables then $\equiv xy$ is satisfiable in every $\mathfrak{A}$, since $\mathfrak{A}_A \models \equiv c_a c_a$ for all $a \in A$; hence $\mathfrak{A} \models \exists x \exists y \equiv xy$ for all $\mathfrak{A}$. Hwever $\equiv xy$ is true on $\mathfrak{A}$ iff A contains at most (and so exactly) one element; thus also $\mathfrak{A} \models \forall x \forall y \equiv xy$ iff $|A| = 1$.

Similarly $\neg \equiv xy$ (for different variables, x, y) is satisfiable on $\mathfrak{A}$ iff $\mathfrak{A} \models \exists x \exists y \neg \equiv xy$ iff $|A| \geq 2$. Analogously for x_1, x_2, x_3 all different variables the formula

$$\neg \equiv x_1x_2 \wedge \neg \equiv x_1x_3 \wedge \neg \equiv x_2x_3$$

is satisfiable in $\mathfrak{A}$ iff $|A| \geq 3$.

More gernerally, for each positive integer n we obtain a formula $\phi_n(x_1, \ldots, x_n)$ without quantifiers (hence called a quantifier-free formula) which is satisfiable in $\mathfrak{A}$ iff $|A| \geq n$. If we define θ_n to be the sentence $\exists x_1 \cdots \exists x_n \phi_n$ then $\mathfrak{A} \models \theta_n$ iff $|A| \geq n$. We then have $\mathfrak{A} \models (\theta_n \wedge \neg\theta_{n+1})$ iff $|A| = n$. Given integers k, l, n with $k \geq 2$, $k < l < n$ we could also, for example, write down a sentence σ such that $\mathfrak{A} \models \sigma$ iff either $|A| < k$ or $|A| = l$ or $|A| > n$. Note that these formulas and sentences use no non-logical symbols and thus will belong to every language.

We now consider two particular $\mathcal{L}$-structures: $\mathfrak{A} = (\omega, \leq, s)$ and $\mathfrak{B} = (\mathbb{Z}, \leq, s)$.

If $\phi_0(x)$ is $\exists y Rxy$ then $\phi_0^{\mathfrak{A}} = \omega$, $\phi_)^{\mathfrak{B}} = \mathbb{Z}$, hence both structures are models of the sentence $\forall x \exists y Rxy$.

If $\phi_1(x)$ is $\forall y Rxy$ then $\phi_1^{\mathfrak{A}} = \{0\}$ and $\phi_1^{\mathfrak{B}} = \emptyset$, hence $\mathfrak{A} \models \exists x \forall y Rxy$ by $\mathfrak{B} \models \neg\exists x \forall y Rxy$.

If $\phi_2(x)$ is $\exists y \equiv xFy$ then $\phi_2^{\mathfrak{A}} = \omega - \{0\}$ but $\phi_2^{\mathfrak{B}} = \mathbb{Z}$. Thus $\mathfrak{B} \models \forall x \exists y \equiv xFy$ but $\mathfrak{A} \models \neg\forall x \exists y \equiv xFy$, that $\mathfrak{A} \models \exists x \forall y \neg \equiv xFy$.

We noted above that $\phi_1(x)$ is such that $\phi_1^{\mathfrak{A}} = \{0\}$. If we now define $\phi_3(y)$ to be $\exists x(\phi_1(x) \wedge \equiv yFx)$ then $\phi_3^{\mathfrak{A}} = \{1\}$. In the same way we can find, for every

$k \in \omega$, a formula $\psi_k(y)$ such that $\psi_k^{\mathfrak{A}} = \{k\}$. Are there formulas χ_k for $k \in \mathbb{Z}$ such that $\chi_k^{\mathfrak{B}} = \{k\}$? Note that it would suffice to show that there is a formula χ_0 with $\chi_0^{\mathfrak{B}} = \{0\}$.

We conclude this section with three important facts about the truth or satisfiability of substitutions.

First, suppose $\mathcal{L}$is a language containing (among other things) an individual constant symbol d. Let $\mathfrak{A}$ be an $\mathcal{L}$-structure and let $a_0 = d^{\mathfrak{A}}$. Then in $\mathcal{L}(A)$ we will have also the constant symbol c_{a_0} and in $\mathfrak{A}_A$ both d and c_{a_0} will be interpreted as the element a_0. If $\phi(x)$ is a formula of $\mathcal{L}(A)$ then, by definition, we will have $\mathfrak{A}_A \models \forall x\phi(x)$ iff $\mathfrak{A}_A \models \phi(c_a)$ for all $a \in A$. *A priori* we could have $\mathfrak{A}_A \models \forall x\phi(x)$ even though $\mathfrak{A}_A \models \neg\phi(d)$, although this would clearly be undesirable. Luckily we can prove that this counter-intuitive state of affairs never occurs.

THEOREM 2.1. *Let $\mathfrak{A}$ be an $\mathcal{L}$-structure, let t be a closed term of $\mathcal{L}(A)$, let $a_0 = t^{\mathfrak{A}_a}$, and let $\phi(x)$ be any formula of $\mathcal{L}(A)$. Then*

$$\mathfrak{A}_A \models \phi(t) \textit{ iff } \mathfrak{A}_A \models \phi(c_{a_0}).$$

In particular

$$\mathfrak{A}_A \models \forall x\phi(x) \to \phi(t).$$

Our second fact attempts to generalize the first to the case in which the term need not be closed. That is, if $\mathfrak{A}$ is an $\mathcal{L}$-structure, $\phi(x)$ is and $\mathcal{L}(A)$-formula and t is a term of $\mathcal{L}$what can we say about the relation between $\phi(x)$ and $\phi(t)$? In particular will we still have $\mathfrak{A}_A \models (\forall x\phi(x) \to \phi(t))$? [Note that this will normall *not* be a sentence, due to the variables in t.]

As the simplest possibility, consider the case in which t is just another variable y. The desired result, then, is that $\phi = \phi(x)$ and $\phi(y) = \phi_y^x$ both define the same subset of A in $\mathfrak{A}_A$–that is, for every $a \in A$ we have $\mathfrak{A}_A \models \phi_{c_a}^x$ iff $\mathfrak{A}_A \models (\phi_y^x)_{c_a}^y$.

In this even we will certainly have $\mathfrak{A}_A \models \forall x\phi(x) \to \phi(y)$. Unfortunately there are certainly problems depending on how y occurs in ϕ. For example, let $\phi(x)$ be $\exists y \neg \equiv xy$. Then $\phi(y)$ is the *sentence* $\exists y \neg \equiv yy$, which is always false, and hence whenever $|A| \geq 2$ we will have $\mathfrak{A} \not\models \forall x\phi(x) \to \phi(y)$. What went wrong here is that, in passing from ϕ to ϕ_y^x some of the near occurrences of y became bound–if this did not happen there would be no problem. The formal definition of "no new occurrences of y become bound" is given in the following definition.

DEFINITION 2.10. For any $\mathcal{L}$and any variables x, y we define the property "y is substitutable for x in ϕ" for $\phi \in \mathrm{Fm}_{\mathcal{L}}$ as follows:

(1) if ϕ is atomic then y is substitutible for x in ϕ,

(2) if $\phi = (\neg\psi)$ then y is substitutible for x in ϕ iff y is substitutible for x in ψ,

(3) if $\phi = (\psi \star \chi)$ where $\star$ is a binary connective, then y is substitutable for x in ϕ iff y is substitutible for x in both ψ and χ,

(4) if $\phi = \forall v_n\psi$ or $\phi = \exists v_n\psi$ then y is substitutible for x in ϕ iff either $x \notin Fv(\phi)$ or $y \neq v_n$ and y is substitutible for x in ψ.

Note in particular that x is substitutible for x in any ϕ, and that y is substitutible for x in any ϕ in which y does not occur.

The following result states that this definition does weed out all problem cases.

THEOREM 2.2. *Let $\mathfrak{A}$ be an $\mathcal{L}$-structure. (1) Let $\phi(x) \in Fm_{\mathcal{L}(A)}$ and assume y is substitutible for x in ϕ. Then $\phi^{\mathfrak{A}_A} = (\phi_y^x)^{\mathfrak{A}_A}$. (2) Let $\phi \in Fm_{\mathcal{L}(A)}$ and assume y is substitutible for x in ϕ. Then $\mathfrak{A}_A \models (\forall x\phi \to \phi_y^x)$.*

In an entirely analogous fashion we can define, for arbitrary terms t of $\mathcal{L}$, the property t is substitutible for x in ϕ to mean (informally) no new occurrences in ϕ^x_t of any variable y occurring in t become bound. We leave the precise formulation of this to the reader. The resulting theorem is exactly what we were after.

THEOREM 2.3. *Let $\mathfrak{A}$ be an $\mathcal{L}$-structure, $\phi \in Fm_{\mathcal{L}(A)}$, $t \in Tm_{\mathcal{L}(A)}$. Assume t is substitutible for x in ϕ. Then $\mathfrak{A}_A \models (\forall x\phi \to \phi^x_t)$.*

We remark, finally, that we can extend our notion of substituition of a term for a variable x to a notioin of substitution of a term for a *constant* c. We leave to the reader the task of defining ϕ^c_t, and $\phi^{c_1 \dots c_n}_{t_1 \dots t_n}$. The main properties we will require are summarized in the following theorem.

THEOREM 2.4. *Let $\phi \in \mathrm{Fm}_{\mathcal{L}}$ and let y be a variable not occurring in ϕ. Then (1) c does not occur in ϕ^c_y, (2) $(\phi^c_y)^y_c = \phi$.*

3. Logical Consequence and Validity

The definitions of logically true formulas, and of logical consequences of sets of sentences, now are exacgly as expected. Some care, however, is needed in defining logical consequences of sets of formulas.

DEFINITION 3.1. Let ϕ be a formula of $\mathcal{L}$. (1) ϕ is logically true or valid, written $\models \phi$, iff $\mathfrak{A} \models \phi$ for every $\mathcal{L}$-structure $\mathfrak{A}$. (2) ϕ is satisfiable iff ϕ is satisfiable on some $\mathcal{L}$-structure $\mathfrak{A}$.

The basic connection between satisfiability and validity is just as in sentential logic. In addition the validity and satisfiability of formulas can be reduced to that of sentences.

LEMMA 3.1. *Let $\phi = \phi(x_1, \dots, x_n) \in \mathrm{Fm}_{\mathcal{L}}$.*

(1) $\models \phi$ iff $\neg\phi$ is not satisfiable
(2) $\models \phi$ iff $\models \forall x_1 \cdots \forall x_n \phi$
(3) ϕ is satisfiable iff $\exists x_1 \cdots \exists x_n \phi$ is satisfiable

Since there are infinitely many different $\mathcal{L}$-structures for any language $\mathcal{L}$one has no hope of checking them all to determine, for example, if some given formula is valid. Nevertheless, one can frequently figure this out, as a few examples will make clear.

EXAMPLE 3.1. Let $\mathcal{L}$be a language with unary relation symbols P and Q. Determine whether or not σ is valid where

$$\sigma = \forall x(Px \to Qx) \to (\forall xPx \to \forall xQx).$$

Suppose $\mathfrak{A} \not\models \sigma$, hence $\mathfrak{A} \models \neg\sigma$ since σ is a sentence. Then $\mathfrak{A} \models \forall x(Px \to Qx)$, $\mathfrak{A} \models \forall xPx$ but $\mathfrak{A} \not\models \forall xQx$. The last assertion means that $\mathfrak{A}_A \models \neg Qc_{a_0}$ for some $a_0 \in A$. But the other two assertions imply that $\mathfrak{A}_A \models Pc_{a_0}$ and $\mathfrak{A}_A \models (Pc_{a_0} \to Qc_{a_0})$, which contradict $\mathfrak{A}_A \models \neg Qc_{a_0}$. Thus we conclude σ is valid.

EXAMPLE 3.2. Determine whether or not θ is valid where

$$\theta = (\forall xPx \to \forall xQx) \to \forall x(Px \to Qx).$$

Suppose $\mathfrak{A} \not\models \theta$, hence $\mathfrak{A} \models \neg\theta$.Then $\mathfrak{A} \models (\forall xPx \to \forall xQx)$ but $\mathfrak{A} \not\models \forall x(Px \to Qx)$. The last assertion means that $\mathfrak{A}_A \models Pc_{a_0}$ and $\mathfrak{A}_A \models \neg Qc_{a_0}$ for some $a_0 \in A$. The first assertion breaks into two cases. In case 1, $\mathfrak{A} \not\models \forall xPx$ and in case 2,

$\mathfrak{A} \models \forall x Q x$. Case 2 is contradicted by the other information, but case 1 will hold provided $\mathfrak{A}_A \models \neg P c_{a_1}$ for some $a_1 \in A$. We thus conclude that θ is not valid since we will have $\mathfrak{A} \models \neg\theta$ whenever there are elements $a_0, a_1, \in A$ such that $a_0 \in P^{\mathfrak{A}}$, $a_0 \notin Q^{\mathfrak{A}}$, $a_1 \notin P^{\mathfrak{A}}$. For example, we can define $\mathfrak{A}$ by specifying that $A = \{0, 1\}$, $P^{\mathfrak{A}} = \{a_0\}$, $Q^{\mathfrak{A}} = \emptyset$.

We can generalize the result established in Example 2.4.1 as follows.

Example 3.3. For any formulas ϕ, ψ of any $\mathcal{L}$,

$$\models \forall x(\phi \to \psi) \to (\forall x\phi \to \forall x\psi).$$

Choose variables $y_1, \ldots, y_n$ such that $\phi = \phi(x, y_1, \ldots, y_n)$ and $\psi = \psi(x, y_1, \ldots, y_n)$. Suppose $\mathfrak{A}$ is an $\mathcal{L}$-structure such that $\mathfrak{A} \not\models \forall x(\phi \to \psi) \to (\forall x\phi \to \forall x\psi)$. Note that we cannot conclude that $\mathfrak{A} \models \forall x(\phi \to \psi)$ etc. since $\forall x(\phi \to \psi)$ is presumably not a sentence. We can, however, conclude that there are $a_1, \ldots, a_n \in A$ such that $\mathfrak{A}_A \not\models \theta(c_{a_1}, \ldots, c_{a_n})$ [where $\theta = \theta(y_1, \ldots, y_n) = (\forall x(\phi \to \psi) \to (\forall x\phi \to \forall x\psi))$] hence–since this is now a sentence–

$$\mathfrak{A}_A \models \forall x(\phi(x, c_{a_1}, \ldots, c_{a_n}) \to \psi(x, c_{a_1}, \ldots, c_{a_n})).$$

The rest of the argument proceeds as before.

Preparatory to defining logical consequence we extend some notations and terminology to sets of formulas and sentences.

Definition 3.2. If $\Gamma \subseteq \mathrm{Fm}_{\mathcal{L}}$ then we will write $\Gamma = \Gamma(x_1, \ldots, x_n)$ provided $Fv(\phi) \subseteq \{x_1, \ldots, x_n\}$ for all $\phi \in \Gamma$.

Definition 3.3. If $\Sigma \subseteq \mathrm{Sn}_{\mathcal{L}}$ and $\mathfrak{A}$ is an $\mathcal{L}$-structure then we say $\mathfrak{A}$ is a model of Σ, written $\mathfrak{A} \models \Sigma$, iff $\mathfrak{A} \models \sigma$ for every $\sigma \in \Sigma$.

Definition 3.4. (1) If $\Gamma \subseteq \mathrm{Fm}_{\mathcal{L}}$, $\Gamma = \Gamma(x_1, \ldots, x_n)$, and $a_1, \ldots, a_n$ are elements of an $\mathcal{L}$-structure $\mathfrak{A}$, then Γ is satisfied on $\mathfrak{A}$ by $a_1, \ldots, a_n$, written $\mathfrak{A}_A \models \Gamma(c_{a_1}, \ldots, c_{a_n})$, iff every formula in Γ is satisfied on $\mathfrak{A}$ by $a_1, \ldots, a_n$. (2) If $\Gamma = \Gamma(x_1, \ldots, x_n) \subseteq \mathrm{Fm}_{\mathcal{L}}$ and $\mathfrak{A}$ is an $\mathcal{L}$-structure then we say Γ is satisfiable in $\mathfrak{A}$ iff Γ is satisfied on $\mathfrak{A}$ by some $a_1, \ldots, a_n$.

Note that if Γ is satisfiable in $\mathfrak{A}$ then every $\phi \in \Gamma$ is satisfiable in $\mathfrak{A}$ but the converse may fail. A trivial example is given by $\Gamma = \{\equiv xy, \neg \equiv xy\}$ with $\mathfrak{A}$ any structure with at least two elements. A non-trivial example is given by $\Gamma = \{\phi_n(x) |\ 1 \leq n \in \omega\}$ where $\phi_1(x) = \exists y RFyx$, $\phi_2(x) = \exists y RFFyx$, etc. in the language $\mathcal{L}$whose none-logical symbols are a binary relation symbol R and a unary function symbol F. Consider the two $\mathcal{L}$-structures $\mathfrak{A} = (\omega, \leq, s)$ and $\mathfrak{B} = (\mathbb{Z}, \leq, s)$. Then $\phi_n^{\mathfrak{B}} = \mathbb{Z}$ for every $n \in \omega - \{0\}$, hence Γ is satisfiable in $\mathfrak{B}$. But $\phi_n^{\mathfrak{A}} = \{k \in \omega |\ n \leq k\}$ for each $n \in \omega - \{0\}$. Thus Γ is *not* satisfiable in $\mathfrak{A}$ although every formula in Γ–indeed every finite $\Gamma_0 \subseteq \Gamma$–is satisfiable in $\mathfrak{A}$.

Definition 3.5. (1) A set Σ of sentences is satisfiable iff it has a model. (2) A set $\Gamma = \Gamma(x_1, \ldots, x_n)$ of formulas is satisfiable iff Γ is satisfiable in some structure $\mathfrak{A}$.

Note that we have only defined satisfiability for sets Γ of formulas with only finitely many free variables total while we could extend these notions to arbitray sets of formulas, we will have no need for these extensions.

We finally can define logical consequence.

Definition 3.6. (1) Let $\Sigma \subseteq \mathrm{Sn}_{\mathcal{L}}$, $\phi \in \mathrm{Fm}_{\mathcal{L}}$. The ϕ is a logical consequence of Σ written $\Sigma \models \phi$, iff $\mathfrak{A} \models \phi$ for every $\mathcal{L}$-structure $\mathfrak{A}$ which is a model of Σ. (2) Let $\Gamma \subseteq \mathrm{Fm}_{\mathcal{L}}$, $\phi \in \mathrm{Fm}_{\mathcal{L}}$ and suppose that $\Gamma = \Gamma(x_1, \ldots, x_n)$, $\phi = \phi(x_1, \ldots, x_n)$. Then ϕ is a logical consequence of Γ written $\Gamma \models \phi$, iff $\mathfrak{A}_A \models \phi(c_{a_1}, \ldots, c_{a_n})$ for every $\mathcal{L}$-structure $\mathfrak{A}$ and every $a_!, \ldots, a_n \in A$ such that $\mathfrak{A}_A \models \Gamma(c_{a_1}, \ldots, c_{a_n})$.

Part (1) of the definition is as expected. We comment on part (2) and give some examples. First of all, the only restriction on Γ is that its formulas contain only finitely many free variables total–since then one can certainly find a single list $x_1, \ldots, x_n$ of variables which includes all variables occurring free either in ϕ or in formulas in Γ. The definition is also independent of the precise list used.

Next, the definition in part (1) is a special case of the definition in part (2). Thus if $\Sigma \subseteq \mathrm{Sn}_{\mathcal{L}}$ and $\phi(x_1, \ldots, x_n) \in \mathrm{Fm}_{\mathcal{L}}$ then also $\Sigma = \Sigma(x_1, \ldots, x_n)$. Now if $\mathfrak{A} \models \Sigma$ then in particular $\mathfrak{A} \models \Sigma(c_{a_1}, \ldots, c_{a_n})$ for all $a_1, \ldots, a_n \in A$. Thus the definition in part (2) yields $\mathfrak{A}_A \models \phi(c_{a_1}, \ldots, c_{a_n})$ for all $a_1, \ldots, a_n \in A$, and thus $\mathfrak{A} \models \phi$ as required for the definition in part (1). On the otherhand if $\mathfrak{A}$ is not a model of Σ then neither definition yields any conclusion about the satisfiability of ϕ in $\mathfrak{A}$.

The definition is formulated to make the following result hold.

Lemma 3.2. *For any* $\Gamma = \Gamma(x_1, \ldots, x_n)$, $\phi(x_1, \ldots, x_n)$, $\psi(x_1, \ldots, x_n)$ *we have*

$$\Gamma \cup \{\phi\} \models \psi \textit{ iff } \Gamma \models (\phi \to \psi).$$

Proof. $\Gamma \models (\phi \to \psi)$ iff there are $\mathfrak{A}$ and $a_1, \ldots, a_n$ such that $\mathfrak{A}_A \models \Gamma(c_{a_1}, \ldots, c_{a_n})$ but $\mathfrak{A}_A \not\models (\phi(c_{a_1}, \ldots, c_{a_n}) \to \psi(c_{a_1}, \ldots, c_{a_n}))$, that is, $a_1, \ldots, a_n$ satisfy $\Gamma \cup \{\phi\}$ in $\mathfrak{A}$ but do not satisfy ψ, thus iff $\mathfrak{A}$ and $a_1, \ldots, a_n$ also show that $\Gamma \cup \{\phi\} \not\models \psi$. $\square$

Thus we see, for example, that $\{Rxy\} \not\models \forall y Rxy$ since $(Rxy \to \forall y Rxy)$ is not valid. On the other hand $\{Rxy\} \models \exists y Rxy$ since $(Rxy \to \exists y Rxy)$ is valid.

In the remainder of this section we identify some classes of validities and we establish some further properties of the logical consequence relation. These validities and properties will then be used in the next section to establish a method which enables one to "mechanically" generate all the logical consequences of a given set Γ.

To begin, tautologies of sentential logic can be used to provide a large class of validities of first order logic. For example $(S_0 \to (S_1 \to S_0)) = \theta$ is a tautology. Of course it isn't even a formula of any first order language $\mathcal{L}$. But if $\phi_0, \phi_1 \in \mathrm{Fm}_{\mathcal{L}}$ then the result of replacing S_0 by ϕ_0 and S_1 by ϕ_1 throughout θ is the formula $\theta^* = (\phi_0 \to (\phi_1 \to \phi_0))$ of $\mathcal{L}$, *and* $\models \theta^*$ for the same reasons that θ is a tautology, as the reader should check. The same thing occurs regardless of what tautology one starts with; thus suggesting the following definition.

Definition 3.7. A formula ψ of $\mathcal{L}$is a tautology iff there is some tautology θ of sentential logic and some substitution of $\mathcal{L}$-formulas for the sentence symbols in $\mathcal{L}$which yields the formula ψ.

Despite the "existential" nature of this definition one can in fact check any given formula ψ of $\mathcal{L}$in a finite number of steps to decide if it is a tautology. The point is that there will only be finitely many sentences θ of sentential logic (except for the use of different sentence symbols) such that ψ can be obtained from θ by some such substitution, and each such θ can be checked to determine whether it is a tautology.

For example let σ be the sentence

$$(\forall v_0(Pv_0 \to Qv_0) \to (\forall v_0 Pv_0 \to \forall v_0 Qv_0)).$$

Then σ can be obtained *only* from the following sentences θ_i of sentential logic:

$$\theta_0 = A$$
$$\theta_1 = (A \to B)$$
$$\theta_2 = (A \to (B \to C)).$$

Since none of these is a tautology (for distinct sentence symbols A, B, C), σ is not a tautology either, although $\models \sigma$.

We leave the proof of the following result to the reader.

THEOREM 3.3. *If $\psi \in \mathrm{Fm}_{\mathcal{L}}$ is a tautology, then $\models \psi$.*

The following list of facts is left to the reader to establish.

THEOREM 3.4. *(1) $\models (\forall x\phi \to \phi^x_t)$ whenever t is substitutible for x in ϕ;*
(2) $\models (\phi \to \forall x\phi)$ if $x \notin Fv(\phi)$;
(3) $\models (\forall x\phi \to \forall y\phi^x_y)$ and $\models (\forall y\phi^x_y \to \forall x\phi)$ if y does not occur in ϕ;
(4) if $\Gamma \models \phi$ then $\gamma \models \forall x\phi$ provided x does not occur free in any formula in Γ;
(5) if $\phi \in \Gamma$ then $\Gamma \models \phi$;
(6) if $\Gamma \models \phi$ and $\Gamma \subseteq \Gamma'$ then $\Gamma' \models \phi$;
(7) if $\Gamma \models \phi$ and $\Gamma \models (\phi \to \psi)$ then $\Gamma \models \psi$;
(8) $\models \equiv xx$;
(9) $\models \equiv xy \to (\phi^z_x \to \phi^z_y)$ provided both x and y are substitutible for z in ϕ.

Logical equivalence is defined as in sentential logic.

DEFINITION 3.8. Formulas ϕ, ψ of $\mathcal{L}$are logically equivalent, written $\phi \vdash\dashv \psi$, iff $\{\phi\} \models \psi$ and $\{\psi\} \models \phi$.

Note, for example, that for any ϕ,

$$\exists x\phi \vdash\dashv \neg\forall x\neg\phi.$$

Together with equivalence from sentential logic this enables us to concludeL

THEOREM 3.5. *For any $\phi(x_1, \ldots, x_n) \in \mathrm{Fm}_{\mathcal{L}}$ there is some $\phi^*(x_1, \ldots, x_n) \in \mathrm{Fm}_{\mathcal{L}}$ such that $\phi \vdash\dashv \phi^*$ and ϕ^* is built using only the connectives $\neg, \to$ and only the quantifier $\forall$.*

For example, if ϕ is

$$\forall x\exists y(Rxy \vee Ryx)$$

then ϕ^* would be

$$\forall x\neg\forall y\neg(\neg Rxy \to Ryx).$$

We have been a little lax in one matter–technically, all our definitions are relative to a language $\mathcal{L}$. But of course a formula ϕ belongs to more than one language. That is: if $\mathcal{L}$, $\mathcal{L}'$ are first order languages and $\mathcal{L} \subseteq \mathcal{L}'$, then $\mathrm{Fm}_{\mathcal{L}} \subseteq Fm_{\mathcal{L}'}$. So we really have two different notions of validity here for $\mathcal{L}$-formulas ϕ:

$\models_{\mathcal{L}} \phi$ meaning $\mathfrak{A} \models \phi$ for all $\mathcal{L}$-structures $\mathfrak{A}$,

$\models_{\mathcal{L}'} \phi$ meaning $\mathfrak{A} \models \phi$ for all $\mathcal{L}'$-structures $\mathfrak{A}'$.

Happily these coincide due to the following easily established fact.

LEMMA 3.6. *Assume $\mathcal{L} \subseteq \mathcal{L}'$, $\mathfrak{A}'$ is an $\mathcal{L}$-structure, $\mathfrak{A} = \mathfrak{A} \restriction \mathcal{L}$. Let $\phi(x_1, \ldots, x_n)$ be a formula of $\mathcal{L}$, $a_1, \ldots, a_n \in A$. Then $\mathfrak{A}'_A \models \phi(c_{a_1}, \ldots, c_{a_n})$ iff $\mathfrak{A}_A \models \phi(c_{a_1}, \ldots, c_{a_n})$.*

4. Formal Deductions

The definition of validity given in the preceding section does not yield a method of deciding, in a finite number of steps, whether or not a given formula is valid. In this section, we describe a procedure for *generating* the validities. A set of formulas each known to be valid is picked out and called the set of *logical axioms*. A *rule* is stated which enables us to generate more formulas in a step-by-step fashion. A finite sequence of formulas showing exactly how a given formula is obtained by repeated applications of the rule beginning with the logical axioms is called a *deduction* of the given formula. Since the rule preserves logical validity all formulas which have deductions are valid. In the next chapter we will prove the converse, that all valid formulas have deductions.

This whole process is *syntactical* and capable of being automated. That is, whether or not a formula is a logical axiom can be determined, in a finite number of steps, by looking at the form of the formula, and whether or not the rule applies to yield a given formula from some other formulas is also determined just by looking at the form of the formulas in question. Thus looking at a given finite sequence of formulas one can determine, in a finite procedure, whether or not this sequence is a deduction. It follows that (for languages with just countably many formulas) one could program a computer to generate a listing of all deductions and thus of all formulas which have deductions. This does *not*, however, mean that we have a procedure which will decide, in a finite number of steps, whether or not a given formula *has a* deduction. So even with the theorem from the next chapter we will not have a procedure which determines, in a finite number of steps, whether or not a given formula is valid.

All of this generalizes to deductions *from* arbitrary sets Γ of formulas, and the theorem from the next chapter will state that ϕ is deducible from Γ iff ϕ is a logical consequence of Γ. This result will then become our main tool for studying properties of "logical consequence."

In particular, our goal in this section is not so much to develop techniques of showing that a specific ϕ is deducible from a specific Γ, but to develop properties *of* the relation of deducibility which will be of theoretical use to us later.

Before defining the logical axioms on piece of terminology is useful.

Definition 4.1. By a generalization of a formula ϕ is meant any formula of the form $\forall x_1, \ldots, \forall x_n \phi$, including ϕ itself.

Note that ϕ is valid iff every generalization of ϕ is valid.

We will simplify our deductive system by having it apply only to formulas built up using just the connectives $\neg, \rightarrow$ and the quantifier $\forall$. This is not a real restriction since every formula is logically equivalent to such a formula. We will continue to write formulas using $\wedge, \vee, \exists$ but these symbols will have to be treated as defined in terms of $\neg, \rightarrow, \forall$ in the context of deducibility.

Definition 4.2. For any first order language $\mathcal{L}$the set Λ of logical axioms of $\mathcal{L}$consists of all generalizations of formulas of $\mathcal{L}$of the following forms:

1) tautologies,
2) $(\forall x\phi \to \phi^x_t)$ where t is substitutible for x in ϕ,
3) $(\forall x(\phi \to \psi) \to (\forall x\phi \to \forall x\psi))$,
4) $(\phi \to \forall x\phi)$ where $x \notin Fv(\phi)$,
5) $\equiv xx$,
6) $(\equiv xy \to (\phi^z_x \to \phi^z_y))$ for atomic formulas ϕ.

We could restrict the tautologies allowed to just those of certain specified forms (see Chapter One Section Six). This would be preferable to certain purposes, but would require more effort in this section.

LEMMA 4.1. *Let $\phi \in \mathrm{Fm}_{\mathcal{L}}$. 1) If $\phi \in \Lambda$ then $\forall x\phi \in \Lambda$ for every variable x. 2) If $\phi \in \Lambda$ then $\models \phi$.*

Our only rule of inference is known by the Latin name "modus ponens" which we will abbreviate to MP. As used in a deduction it allows one to put down a formula ψ provided formulas ϕ and $(\phi \to \psi)$ precede it.

DEFINITION 4.3. Let $\Gamma \subseteq \mathrm{Fm}_{\mathcal{L}}$.

(a) A deduction from Γ is a finite sequence $\phi_0, \dots, \phi_n$ of formulas of $\mathcal{L}$such that for every $i \leq n$ we have either

(1) $\phi_i \in \Lambda \cup \Gamma$, or
(2) there are $j, k < i$ such that $\phi_k = (\phi_j \to \phi_i)$.

(b) The formula ϕ is deducible from Γ, written $\Gamma \vdash \phi$, iff there is a deduction $\phi_0, \dots, \phi_n$ from Γ with $\phi_n = \phi$.

(c) In particular a logical deduction is just a deduction from $\Gamma = \emptyset$, and ϕ is logically deducible, $\vdash \phi$, iff $\emptyset \vdash \phi$.

PROPOSITION 4.2. ***(Soundness)*** *If $\Gamma \vdash \phi$ then $\Gamma \models \phi$. In particular if ϕ is logically deducible then ϕ is valid.*

PROOF. Let $\phi_0, \dots, \phi_n$ be a deduction of ϕ from Γ. We show, by induction on i, that $\Gamma \models \phi_i$ for every $i \leq n$. Since $\phi_n = \phi$ this suffices to show $\Gamma \models \phi$. Let $i \leq n$ and suppose, as inductive hypothesis, that $\Gamma \models \phi_j$ for all $j < i$. If $\phi_i \in (\Lambda \cup \Gamma)$ then $\Gamma \models \phi_i$. In the other case there are $j, k < i$ such that $\phi_k = (\phi_j \to \phi_i)$. By the inductive hypothesis $\Gamma \models \phi_j$ and $\Gamma \models (\phi_j \to \phi_i)$, and so $\Gamma \models \phi_i$. □

LEMMA 4.3. *Let $\Gamma \subseteq \mathrm{Fm}_{\mathcal{L}}$.*

(1) If $\phi \in (\Lambda \cup \Gamma)$ then $\Gamma \vdash \phi$.

(2) If $\Gamma \vdash \phi$ and $\Gamma \vdash (\phi \to \psi)$ then $\Gamma \vdash \psi$.

PROOF. (of part 2) Let $\phi_0, \dots, \phi_n$ be a deduction of ϕ from Γ and let $\psi_0, \dots, \psi_m$ be a deduction of $(\phi \to \psi)$ from Γ. Then the sequence $\phi_0, \dots, \phi_n, \psi_0, \dots, \psi_m, \psi$ is a deduction of ψ from Γ. □

Clearly any formula in $(\Lambda \cup \Gamma)$ is deducible from Γ with a deduction of length one. The shortest possible deduction involving a use of MP will have length three. Here is an example:

$$(\forall x\neg\phi \to \neg\phi),$$
$$(\forall x\neg\phi \to \neg\phi) \to (\phi \to \neg\forall x\neg\phi),$$
$$(\phi \to \neg\forall x\neg\phi).$$

The first formula is a logical axiom since x is substitutible for x in any ϕ, and $\phi^x_x = \phi$. The second formula is a tautology, and the third follows by MP.

This example shows that $\vdash (\phi \to \neg\forall x\neg\phi)$ for every $\phi \in \mathrm{Fm}_{\mathcal{L}}$. Recalling our use of defined symbols, it may be more intelligibly expressed as $\vdash (\phi \to \exists x\phi)$. The reader should try to establish that $\vdash \forall x(\phi \to \exists x\phi)$.

Due to our restricting the connectives and quantifiers allowed in formulas, every non-atomic formula has either the form $\neg\phi$ or $(\phi \to \psi)$ or $\forall x\phi$. We proceed to give several results which characterize the conditions under which such formulas are deducible from Γ. These results can then be used to show deducibility of formulas.

LEMMA 4.4. ***Deduction Theorem*** *For any $\Gamma \subseteq \mathrm{Fm}_{\mathcal{L}}$, $\phi, \psi \in \mathrm{Fm}_{\mathcal{L}}$:*

$$\Gamma \cup \{\phi\} \vdash \psi \textit{ iff } \Gamma \vdash (\phi \to \psi).$$

PROOF. The implication from right to left is an easy consequence of Lemma 2.5.3 part 2 above–i.e. of MP.

For the other implication, supppose $\psi_0, \ldots, \psi_n$ is a deduction from $\Gamma \cup \{\phi\}$ of ψ. We show, by induction on i, that $\Gamma \vdash (\phi \to \psi_i)$ for all $i \leq n$. Since $\psi_i = \psi$ this will establish $\Gamma \vdash (\phi \to \psi)$. So let $i \leq n$ and assume as inductive hypothesis that $\Gamma \vdash (\phi \to \psi_j)$ for all $j < i$.

There are two cases. If $\psi_i \in \Lambda \cup \Gamma$ then $\Gamma \vdash \psi_i$ hence $\Gamma \vdash (\phi \to \psi_i)$ by MP since $(\psi_i \to (\phi \to \psi_i))$ is a tautology. If, on the other hand, ψ_i follows by MP then there are $j, k < i$ such that $\psi_k = (\psi_j \to \psi_i)$. By the inductive hypothesis, $\Gamma \vdash (\phi \to \psi_j)$ and $\Gamma \vdash (\phi \to (\psi_j \to \phi_i))$. Use of MP and the tautology $(\phi \to (\psi_j \to \psi_i)) \to ((\phi \to \psi_j) \to (\phi \to \psi_i))$ yields the conclusion $\Gamma \vdash (\phi \to \psi)$. $\square$

The use of the Deduction Theorem is to reduce the question of finding a deduction of $(\phi \to \psi)$ from Γ to that of finding a deduction of ψ from $\Gamma \cup \{\phi\}$. This second question will usually be easier since ψ is shorter than $(\phi \to \psi)$.

Our first reduction for universally quantified formulas is not completely satisfactory, but will be imporved later.

LEMMA 4.5. ***(Generalization)*** *Assume x does not occur free in any formula in Γ. Then $\Gamma \vdash \phi$ iff $\Gamma \vdash \forall x\phi$.*

PROOF. The implication from right to left is easily established. For the other direction, suppose $\phi_0, \ldots, \phi_n$ is a deduction from Γ of ϕ. We show that $\Gamma \vdash \forall x\phi_i$ for all $i \leq n$, by induction. So, let $i \leq n$ and suppose as induction hypothesis that $\Gamma \vdash \forall x\phi_j$ for all $j < i$. If $\phi_i \in \Lambda$ then also $\forall x\phi_i \in \Lambda$ and thus $\Gamma \vdash \forall x\phi_i$. If $\phi_i \in \Gamma$ then $x \notin Fv(\phi_i)$ hnce $(\phi_i \to \forall x\phi_i) \in \Lambda$ and thus $\Gamma \vdash \forall x\phi_i$ by MP. If ϕ_i follows by MP then there are $j, k < i$ such that $\phi_k = (\phi_j \to \phi_i)$. By the inductive hypothesis, $\Gamma \vdash \forall x\phi_j$ and $\Gamma \vdash \forall x(\phi_j \to \phi_i)$. Now $(\forall x(\phi_j \to \phi_i) \to (\forall x\phi_j \to \forall x\phi_i)) \in \Lambda$ so two uses of MP yield $\Gamma \vdash \forall x\phi_i$ as desired. $\square$

To remove the restriction in the statement of Generalization, we first prove a result about changing bound variables.

LEMMA 4.6. *Assume the variable y does not occur in ϕ. Then (1) $\vdash (\forall x\phi \to \forall y\phi^x_y)$ and (b) $\vdash (\forall y\phi^x_y \to \forall x\phi)$.*

PROOF. (a) Since y is substitutible for x in ϕ, $\forall y(\forall x\phi \to \phi^x_y) \in \Lambda$. Using an appropriate axiom of form 3) and MP we conclude $\vdash (\forall y\forall x\phi \to \forall y\phi^x_y)$. Since $y \notin Fv(\forall x\phi)$ we have $(\forall x\phi \to \forall y\forall x\phi) \in \Lambda$ and so $\vdash (\forall x\phi \to \forall y\phi^x_y)$ using MP and an appropriate tautology.

(b) One first proves that, since y does not occur in ϕ, we have that x is substitutible for y in ϕ_y^x and that $(\phi_y^x)_x^y = \phi$. This result then follows from (a). Details are left to the reader. $\square$

COROLLARY 4.7. ***(Generalization)*** *Assume y does not occur in ϕ and y does not occur free in any formula in Γ. Then $\Gamma \vdash \phi_y^x$ iff $\Gamma \vdash \forall x\phi$.*

PROOF. The implication from right to left is easy, so we just establish the other direction. If $\Gamma \vdash \phi_y^x$ then the first form of Generalization yields $\Gamma \vdash \forall y\phi_y^x$. But the above lemma implies $\Gamma \vdash (\forall y\phi_y^x \to \forall x\phi)$, so we conclude $\Gamma \vdash \forall x\phi$. $\square$

Thus to show $\forall x\phi$ is deducible from some Γ in which x occurs free, we first choose y not occurring in ϕ and not occurring free in Γ and then show $\Gamma \vdash \phi_y^x$. Since we virtually always are considering only sets Γ which have just finitely many free variables total, this choice of y is not a problem.

Before considering formulas of the form $\neg\phi$, we introduce the important notion of consistency and use it to characterize deducibility.

DEFINITION 4.4. (1) The set Γ of formulas is inconsistent iff there is some $\theta \in \mathrm{Fm}_{\mathcal{L}}$ such that $\Gamma \vdash \theta$ and $\Gamma \vdash \neg\theta$. (2) The set Γ is consistent iff Γ is not inconsistent.

We first note the following easy characterization of inconsistency.

LEMMA 4.8. *A set $\Gamma \subseteq \mathrm{Fm}_{\mathcal{L}}$ is inconsistent iff $\Gamma \vdash \phi$ for every $\phi \in \mathrm{Fm}_{\mathcal{L}}$.*

PROOF. The implication from right to left is clear. For the other direction, suppose $\Gamma \vdash \theta$ and $\Gamma \vdash \neg\theta$. For any $\phi \in \mathrm{Fm}_{\mathcal{L}}$, $(\theta \to (\neg\theta \to \phi))$ is a tautology, hence $\Gamma \vdash \phi$ with two uses of MP. $\square$

The following theorem enables us to reduce deducibility to (in-) consistency.

THEOREM 4.9. *Let $\Gamma \subseteq \mathrm{Fm}_{\mathcal{L}}$, $\phi \in \mathrm{Fm}_{\mathcal{L}}$. Then $\Gamma \vdash \phi$ iff $\Gamma \cup \{\phi\}$ is not consistent.*

PROOF. If $\Gamma \vdash \phi$ then $\Gamma \cup \{\phi\}$ is inconsistent since both ϕ and $\neg\phi$ are deducible from it. If $\Gamma \cup \{\phi\}$ is inconsistent then, by the preceding lemma, we see that in particular $\Gamma \cup \{\neg\phi\} \vdash \phi$ and so $\Gamma \vdash (\neg\phi \to \phi)$, by the Deduction Theorem. But $((\neg\phi \to \phi) \to \phi)$ is a tautology, and so we conclude $\Gamma \vdash \phi$. $\square$

In particular we derive a method of showing the deducibility of formulas of the form $\neg\phi$.

COROLLARY 4.10. ***(Proof by Contradiction)*** *$\Gamma \cup \{\phi\}$ is inconsistent iff $\Gamma \vdash \neg\phi$.*

This may not actually be very useful, since showing $\Gamma \cup \{\phi\}$ is inconsistent is completely open-ended–what contradiction θ, $\neg\theta$ you should try to derive is unspecified. As a prcatical matter of showing the deducibility of $\neg\theta$ it is usually better to use one of the following, if at all possible.

LEMMA 4.11. *(1) $\Gamma \vdash \phi$ iff $\Gamma \vdash \neg\neg\phi$. (2) $\Gamma \vdash \neg(\phi \to \psi)$ iff $\Gamma \vdash \phi$ and $\Gamma \vdash \neg\psi$. (3) $\Gamma \cup \{\phi\} \vdash \psi$ iff $\Gamma \cup \{\neg\phi\} \vdash \neg\phi$.*

The proofs are immediate consequences of appropriate tautologies and are left to the reader.

As an example of showing deducibility using the established rules, we show that $\vdash (\exists x \forall y \phi \to \forall y \exists x \phi)$, that is $\vdash (\neg \forall x \neg \forall y \phi \to \forall y \neg \forall x \neg \phi)$. By the Deduction Theorem it suffices to show $\neg \forall x \neg \forall y \phi \vdash \forall y \neg \forall x \neg \phi$; by Generalization ($y$ not being free in $\neg \forall x \neg \forall y \phi$) it suffices to show $\neg \forall x \neg \forall y \phi \vdash \neg \forall x \neg \phi$. By Lemma 2.5.9 part 3 it suffices to show $\forall x \neg \phi \vdash \forall x \neg \forall y \phi$. By Generalization (since $x \notin Fv(\forall x \neg \phi)$) it suffices to show $\forall x \neg \phi \vdash \neg \forall y \phi$. Finally, by the corollary "Proof by Contradiction" (nothing else being applicable, it suffices to show $\Gamma = \{\forall x \neg \phi, \forall y \phi\}$ is inconsistent. But this is now easy, since $\Gamma \vdash \neg \phi$ and $\Gamma \vdash \phi$.

The "Soundness" result from earlier in this section has the following form applied to consistency.

COROLLARY 4.12. ***(Soundness)*** *Assume $\Gamma \subseteq \mathrm{Fm}_{\mathcal{L}}$ is satisfiable. Then Γ is consistent.*

PROOF. Suppose Γ is inconsistent. Then $\Gamma \vdash \forall x \neg \equiv xx$. So by Soundness we have $\Gamma \models \forall x \neg \equiv xx$. Thus, if Γ is satisfiable on $\mathfrak{A}$ then necessarily, $\mathfrak{A} \models \forall x \neg \equiv xx$– which is impossible. □

The Completeness Theorem proved in the next chapter will establish the converses of the two Soundness results, that is we will conclude the following equivalences.

THEOREM 4.13. *Let $\Gamma \subseteq \mathrm{Fm}_{\mathcal{L}}$. Then (1) for any $\phi \in \mathrm{Fm}_{\mathcal{L}}$, $\Gamma \vdash \phi$ iff $\Gamma \models \phi$; (2) Γ is consistent iff Γ is satisfiable.*

The importance of this result is that facts about deducibility and consistency can be translated into facts about logical consequence and satisfiability. The most important such general fact is the translation of the easy *finiteness* property of deductions.

LEMMA 4.14. *(1) $\Gamma \vdash \phi$ iff $\Gamma_0 \vdash \phi$ for some finite $\Gamma_0 \subseteq \Gamma$. (2) Γ is consistent iff every finite $\Gamma_0 \subseteq \Gamma$ is consistent.*

Both parts of the lemma are immediate from the fact that any specific deduction from Γ uses just finitely many formulas from Γ and thus is a deduction from a finite subset of Γ.

Using the Completeness Theorem the Finiteness Lemma becomes the highly important, and non-obvious, Compactness Theorem.

THEOREM 4.15. *Let $\Gamma \subseteq \mathrm{Fm}_{\mathcal{L}}$. Then (1) $\Gamma \models \phi$ iff $\Gamma_0 \models \phi$ for some finite $\Gamma_0 \subseteq \Gamma$; (2) Γ is satisfiable iff every finite $\Gamma_0 \subseteq \Gamma$ is satisfiable.*

For the proof of Completeness we will need two further facts about deducibility, both of which concern constant symbols. Recall that we defined ϕ^c_y as the result of replacing all occurrences of c in ϕ by the variable y. The resulting formula has no occurrences of c, and $(\phi^c_y)^y_c = \phi$ provided y does not occur in ϕ. The content of the following lemma is that this substitution preserves deducibility from sets Γ in which c does not occur.

LEMMA 4.16. *Let c be a constant symbol of $\mathcal{L}$not occurring in any formula in Γ. Let $\phi_0, \ldots, \phi_n$ be a deduction from Γ and let $\psi_i = (\phi_i)^c_y$ where y is a variable not occurring in any of $\phi_0, \ldots, \phi_n$. Then $\psi_0, \ldots, \psi_n$ is also a deduction from Γ.*

Proof. If $\phi_i \in \Gamma$ then $\psi_i = \phi_i$ since c does not occur in any formula in Γ. If ϕ_i follows by MP by ϕ_j, ϕ_k then it is easily checked that ψ_i likewise follows by MP from ψ_j, ψ_k. It thus suffices to show that $\psi_i \in \Lambda$ if $\phi_i \in \Lambda$. This is tedious–especially for tautologies–but not essentially difficult, so we leave it to the reader. $\square$

Our first corollary of this is yet another form of Generalization.

Corollary 4.17. ***(Generalization on Constants)*** *Assume c does not occur in Γ and $\Gamma \vdash \phi^x_c$. Then $\Gamma \vdash \forall x\phi$.*

Proof. Let $\phi_0, \ldots, \phi_n$ be a deduction of ϕ^x_c from Γ and let y be a variable not occurring in any of $\phi_0, \ldots, \phi_n$. Let $\psi_i = (\phi_i)^c_y$. Then $\psi_0, \ldots, \psi_n$ is a deduction from Γ, by the lemma, and hence from

$$\Gamma_0 = \{\phi_i |\ \phi_i \in \Gamma, i \leq n\} = \{\psi_i |\ \psi_i \in \Gamma, i \leq n\}.$$

Thus $\Gamma_0 \vdash (\phi^x_c)^c_y$. But $(\phi^x_c)^c_y = \phi^x_y$ since y does not occur in ϕ^x_c. Further y does not occur (free) in any formula of Γ_0, so the second form of Generalization yields $\Gamma_0 \vdash \forall x\phi$, and so $\Gamma \vdash \forall x\phi$. $\square$

The second consequence of this lemma concerns the result of changing languages. Suppose $\mathcal{L} \subseteq \mathcal{L}'$, $\Gamma \subseteq \mathrm{Fm}_{\mathcal{L}}$, $\phi \in \mathrm{Fm}_{\mathcal{L}}$. We then really have two different definitions of "ϕ is deducible from Γ" according to whether the deduction consists only of formulas of $\mathcal{L}$on whether formulas of $\mathcal{L}'$ are allowed. Let us express these as $\Gamma \vdash_{\mathcal{L}} \phi$, $\Gamma \vdash \mathcal{L}'\phi$. Clearly, if $\Gamma \vdash_{\mathcal{L}} \phi$ then $\Gamma \vdash_{\mathcal{L}'} \phi$. The converse is much less clear. We are, however, able to prove this now provided that $\mathcal{L}' - \mathcal{L}$ consists entirely of constant symbols.

Theorem 4.18. *Assume $\mathcal{L} \subseteq \mathcal{L}'$ and $\mathcal{L}' - \mathcal{L}$ consists entirely of constant symbols. Let $\Gamma \subseteq \mathrm{Fm}_{\mathcal{L}}$. Then (1) for any $\phi \in \mathrm{Fm}_{\mathcal{L}}$, $\Gamma \vdash_{\mathcal{L}} \phi$ iff $\Gamma \vdash_{\mathcal{L}'} \phi$; (2) Γ is consistent with respect to $\mathcal{L}$-deductions iff Γ is consistent with respect to $\mathcal{L}'$-deductions.*

Proof. (1) Let $\phi_0, \ldots, \phi_n$ be an $\mathcal{L}'$-deduction from Γ of ϕ. Let $c_0, \ldots, c_m$ list all constants from $\mathcal{L}' - \mathcal{L}$ appearing in this deduction–so this is an $(\mathcal{L} \cup \{c_0, \ldots, c_m\}$-deduction. Let $\psi^0_i = (\phi_i)^{c_0}_{y_0}$ for each $i = 0, \ldots, n$ where y_0 is a variable not occurring in any of $\phi_0, \ldots, \phi_n$. Then by the lemma $\psi^0_0, \ldots, \psi^0_n$ is a deduction from Γ consisting of formulas of $\mathcal{L} \cup \{c_1, \ldots, c_m\}$. Since $\phi_n = \phi \in \mathrm{Fm}_{\mathcal{L}}$ we have $\psi^0_n = \phi_n$, so this is still a deduction of ϕ. Repeating this for $c_1, \ldots, c_m$ we eventually arrive at a deduction from Γ of ϕ consisting just of formulas of $\mathcal{L}$.

(2) This follows immediately from (1). $\square$

5. Theories and Their Models

There are two different paradigms for doing mathematics. One is to study all structures in some class defined by certain properites. The other is to study some specific structure. An example of the first would be group theory, which investigates the class of all structures satisfying the group axioms. An example of the second would be real analysis, which studies the particular structure of the real numbers.

Both of these paradigms have counterparts in logic. What is charateristic of the logical approach in both cases is that the properties used to define the class of structures, and the properties of the structures themselves, should be expressible in first order logic. To begin with we concentrate on the first paradigm.

First some terminology.

Definition 5.1. Let $\Sigma \subseteq \mathrm{Sn}_{\mathcal{L}}$. Then the set of consequences of Σ is

$$Cn_{\mathcal{L}}(\Sigma) = \{\theta \in \mathrm{Sn}_{\mathcal{L}} |\ \Sigma \models \theta\}.$$

Note that if $\mathcal{L} \subset \mathcal{L}'$ and $\Sigma \subseteq \mathrm{Sn}_{\mathcal{L}}$ then $Cn_{\mathcal{L}}(\Sigma) \subset Cn_{\mathcal{L}'}(\Sigma)$. Nevertheless we will frequently omit the subscript $\mathcal{L}$if there is no chance of confusion.

Definition 5.2. Let $\Sigma \subseteq \mathrm{Sn}_{\mathcal{L}}$. Then the class of models of Σ is

$$Mod_{\mathcal{L}}(\Sigma) = \{\mathcal{L}\text{-structures } \mathfrak{A} |\ \mathfrak{A} \models \Sigma\}.$$

We note the following easy fact.

Lemma 5.1. *Let $\Sigma_1, \Sigma_2 \subseteq \mathrm{Sn}_{\mathcal{L}}$. Then $Cn(\Sigma_1) = Cn(\Sigma_2)$ iff $Mod(\Sigma_1) = Mod(\Sigma_2)$.*

We think of a set Σ of sentences as the axioms of a theory. Then $Mod(\Sigma)$ is the calss of models of the theory, and $Cn(\Sigma)$ is the set of theorems of the theory, that is the set of sentences true on all models of the theory. By the above lemma, two sets of sentences have the same models iff they have the same consequences. In this case we will consider them to both define the same theory. This is conveniently captured in the following definition.

Definition 5.3. (1) By a theory of $\mathcal{L}$is meant any set of sentences of $\mathcal{L}$of the form $T = Cn_{\mathcal{L}}(\Sigma)$, $\Sigma \subseteq \mathrm{Sn}_{\mathcal{L}}$. (2) If T is a theory of $\mathcal{L}$then any set $\Sigma \subseteq \mathrm{Sn}_{\mathcal{L}}$ such that $T = Cn_{\mathcal{L}}(\Sigma)$ is called a set of axioms for T.

Lemma 5.2. *Let $T \subseteq \mathrm{Sn}_{\mathcal{L}}$. Then T is a theory of $\mathcal{L}$iff $T = Cn_{\mathcal{L}}(T)$.*

The largest thoery of $\mathcal{L}$is $T = \mathrm{Sn}_{\mathcal{L}}$. This theory has no models and can be axiomatized by the negation of any logically valid sentence of $\mathcal{L}$, for example $\forall x \neg \equiv xx$.

The smallest theory of $\mathcal{L}$is $T = \{\theta \in \mathrm{Sn}_{\mathcal{L}} |\ \models \theta\}$. This theory is equal to $Cn_{\mathcal{L}}(\emptyset)$ and every structure of $\mathcal{L}$is a model of it.

In between these extremes are the theories of T which have models but which are not satisfied by every $\mathcal{L}$-structure. One important kind of example is given by the (complete) theory of an $\mathcal{L}$-structure.

Definition 5.4. Let $\mathfrak{A}$ be an $\mathcal{L}$-structure. Then the (complete) theory of $\mathfrak{A}$ is

$$Th(\mathfrak{A}) = \{\sigma \in \mathrm{Sn}_{\mathcal{L}} |\ \mathfrak{A} \models \sigma\}.$$

Definition 5.5. Let T be a theory of $\mathcal{L}$. Then T is complete iff T has a model and for every $\sigma \in \mathrm{Sn}_{\mathcal{L}}$ either $\sigma \in T$ or $\neg\sigma \in T$.

The following fact is easily verified.

Lemma 5.3. *A set $T \subseteq \mathrm{Sn}_{\mathcal{L}}$ is a complete theory of $\mathcal{L}$iff $T = Th(\mathfrak{A})$ for some $\mathcal{L}$-structure $\mathfrak{A}$. In this case $T = Th(\mathfrak{A})$ for every $\mathfrak{A} \models T$.*

The complete theory of $\mathfrak{A}$ tells you *everything* about $\mathfrak{A}$ that can be expressed by first order sentences of $\mathcal{L}$. Having the *same* complete theory defines a very natural equivalence relation on the class of $\mathcal{L}$-structures.

Definition 5.6. $\mathcal{L}$-structures $\mathfrak{A}$ and $\mathfrak{B}$ are elementarily equivalent written $\mathfrak{A} \equiv \mathfrak{B}$ iff $Th(\mathfrak{A}) = Th(\mathfrak{B})$.

We will see later that elementarily equivalent structures may look very different–to begin with, they may have universes of different cardinalities. In fact, we will prove in Chapter 3 that whenver $\mathfrak{A}$ is infinite (meaning A, the universe of $\mathfrak{A}$, is infinite) then there is a $\mathfrak{B}$ such that $\mathfrak{A} \equiv \mathfrak{B}$ and $|A| < |B|$.

The natural question is how much "alike" must elementary equivalent structures be? This is vague, but we will interpret this to mean what can we prove about the models of a complete theory? This will in fact be the central topic of much of Part B, in partiuclar of Chapter 5.

Even more fundamental is the question of how we would show that two structures $\mathfrak{A}, \mathfrak{B}$ are elementarily equivalent. We won't be able to prove directly that for every $\theta \in \mathrm{Sn}_{\mathcal{L}}$ we have $\mathfrak{A} \models \sigma$ iff $\mathfrak{B} \models \sigma$. If for a given $\mathfrak{A}$ we could explicitly determine ("write down") a set Σ of axioms for $Th(\mathfrak{A})$ then we could conclude that $\mathfrak{B} \equiv \mathfrak{A}$ iff $\mathfrak{B} \models \Sigma$. But determining whether or not Σ axiomatizes a complete theory is of the same level of difficulty–we are not going to be able to prove directly that for every $\theta \in \mathrm{Sn}_{\mathcal{L}}$ we have either $\Sigma \models \theta$ or $\Sigma \models \neg\theta$ but not both.

We will in fact develop some techniques for showing that a theory given axiomatically is complete, although they will be of restricted applicability. More importantly, we will develop techniques for showing that a theory–including one given in the form $Th(\mathfrak{A})$–will have models with certain properties. These techniques will not yield a complete description of the structures proved to exist, but they yield a great deal of information about the models of a theory.

As a beginning step we introduce *isomorphisms* between $\mathcal{L}$-structures and prove that isomorphic $\mathcal{L}$-structures are elementarily equivalent. Roughly speaking, two structures are isomorphic provided there is a one-to-one correspondence between their universes which "translates" one structure into the other.

Definition 5.7. Given $\mathcal{L}$-structures $\mathfrak{A}, \mathfrak{B}$ a function h is an isomorphism of $\mathfrak{A}$ onto $\mathfrak{B}$, written $h : \mathfrak{A} \cong \mathfrak{B}$, iff h is a function mapping A one-to-one onto B such that (i) $h(c^{\mathfrak{A}}) = c^{\mathfrak{B}}$ for all constants $c \in \mathcal{L}$, (ii) $h(F^{\mathfrak{A}}(a_1, \dots, a_m)) = F^{\mathfrak{B}}(h(a_1), \dots, h(a_m))$ for all m-ary function symbols $F \in \mathcal{L}$ and all $a_1, \dots, a_m \in A$, (iii) $R^{\mathfrak{A}}(a_1, \dots, a_m)$ holds iff $R^{\mathfrak{B}}(h(a_1), \dots, h(a_m))$ for all m-ary relation symbols $R \in \mathcal{L}$ and all $a_1, \dots, a_m \in A$.

The reader should note that this definition agrees with the familiar algebraic definition on algebraic structures like groups, rings, etc. Since isomorphic structures are "the same" except for the identity of the elements of their universes it is not surprising that they will be elementarily equivalent. In fact, we prove something stronger.

Theorem 5.4. *Let $\mathfrak{A}, \mathfrak{B}$ be $\mathcal{L}$-structures and assume $h : \mathfrak{A} \cong \mathfrak{B}$. Then for every $\phi(x_0, \dots, x_{n-1}) \in \mathrm{Fm}_{\mathcal{L}}$ and for all $a_0, \dots, a_{n-1} \in A$ we have $\mathfrak{A}_A \models \phi(c_{a_0}, \dots, c_{a_{n-1}})$ iff $\mathfrak{B}_B \models \phi(c_{b_0}, \dots, c_{b_{n-1}})$ where $b_i = h(a_i)$, $i = 0, \dots, n-1$.*

Proof. One first shows by induction on $\mathrm{Tm}_{\mathcal{L}}$ that for every $t(x_0, \dots, x_{n-1}) \in \mathrm{Tm}_{\mathcal{L}}$ and every $a_0, \dots, a_{n-1} \in A$,

$$h(t^{\mathfrak{A}}(a_0, \dots, a_{n-1})) = t^{\mathfrak{B}}(h(a_0), \dots, h(a_{n-1})).$$

This argument is left to the reader. One next shows the equivalence in the statement of the theorem by induction on $\mathrm{Fm}_{\mathcal{L}}$. We do two parts of the argument and leave the rest to the reader.

If ϕ is the atomic formula $\equiv t_1 t_2$, then the following are equivalent:

$$\mathfrak{A}_A \models \phi(c_{a_0}, \ldots, c_{a_{n-1}}),$$

$$t_1^{\mathfrak{A}}(a_0, \ldots, a_{n-1}) = t_2^{\mathfrak{A}}(a_0, \ldots, a_{n-1}),$$

$$h(t_1^{\mathfrak{A}}(a_0, \ldots, a_{n-1})) = h(t_2^{\mathfrak{A}}(a_0, \ldots, a_{n-1})),$$

$$t_1^{\mathfrak{B}}(h(a_0), \ldots, h(a_{n-1})) = t_2^{\mathfrak{B}}(h(a_0), \ldots, h(a_{n-1})),$$

$$\mathfrak{B}_B \models \phi(c_{b_0}, \ldots, c_{b_{n-1}}).$$

The equivalence of the second and third lines follows since h is one-to-one, and the equivalence of the third and fourth lines follows from the preliminary lemma on $\mathrm{Tm}_{\mathcal{L}}$.

Suppose $\phi(x_0, \ldots, x_{n-1}) = \forall y \psi(x_0, \ldots, x_{n-1}, y)$ and, as inductive hypothesis, that the equivalence holds for ψ and for all $a_0, \ldots, a_{n-1}, a \in A$. Fixing $a_0, \ldots, a_{n-1} \in A$ the following are equivalent:

$$\mathfrak{A}_A \models \phi(c_{a_0}, \ldots, c_{a_{n-1}}),$$

$$\mathfrak{A}_A \models \psi(c_{a_0}, \ldots, c_{a_{n-1}}, c_a) \text{ for all } a \in A,$$

$$\mathfrak{B}_B \models \psi(c_{b_0}, \ldots, c_{b_{n-1}}, c_{h(b)}) \text{ for all } a \in A,$$

$$\mathfrak{B}_B \models \psi(c_{b_0}, \ldots, c_{b_{n-1}}, c_b) \text{ for all } b \in B,$$

$$\mathfrak{B}_B \models \phi(c_{b_0}, \ldots, c_{b_{n-1}}).$$

The equivalence of the second and third lines follows from the inductive hypothesis, and the equivalence of the third and fourth lines follows since h maps A onto B. $\square$

As usual we say $\mathfrak{A}$ and $\mathfrak{B}$ *are isomorphic*, written $\mathfrak{A} \cong \mathfrak{B}$, iff there is an isomorphism h of $\mathfrak{A}$ onto $\mathfrak{B}$. Further an *automorphism* of $\mathfrak{A}$ is an isomorphism h of $\mathfrak{A}$ onto itself.

EXAMPLE 5.1. Let $\mathcal{L}$be the language whose only non-logical symbol is a binary relation symbol R. Let $\mathfrak{A} = (\omega, \leq)$ and $\mathfrak{B} = (B, \leq)$ where $B = \{2k |\ k \in \omega\}$. Then $\mathfrak{A}$ and $\mathfrak{B}$ are isomorphic via the mapping $h : A \to B$ defined by $h(k) = 2k$ for all $k \in \omega$. All that needs to be checked is that h maps A one-to-one onto B and that $R^{\mathfrak{A}}(k, l)$ holds iff $R^{\mathfrak{B}}(h(k), h(l))$ holds, that is $k \leq l$ iff $2k \leq 2l$.

EXAMPLE 5.2. With $\mathcal{L}$as in the previous example, et $\mathfrak{A} = (\mathbb{Z}, \leq)$. Then for every $k_0, l_0 \in A$ there is an automorphism h of $\mathfrak{A}$ such that $h(k_0) = l_0$. We leave the reader to check that h defined by $h(k) = k + (l_0 - k_0)$ works.

Note that it follows, in this example, that for every $\phi(x) \in \mathrm{Fm}_{\mathcal{L}}$ and every $k_0, l_0 \in A$ we have

$$\mathfrak{A}_A \models \phi(c_{k_0}) \text{ iff } \mathfrak{A}_A \models \phi(c_{l_0}).$$

It follows that either $\phi^{\mathfrak{A}} = A$ or $\phi^{\mathfrak{A}} = \emptyset$.

EXAMPLE 5.3. Let $\mathcal{L}$be the language whose only non-logical symbol is a constant symbol c. Let $\mathfrak{A}, \mathfrak{B}$ be any two $\mathcal{L}$-structures with $|A| = |B|$. Then $\mathfrak{A} \cong \mathfrak{B}$.

Let $A_0 = A - \{c^{\mathfrak{A}}\}$, $B_0 = B - \{c^{\mathfrak{B}}\}$. Then $|A_0| = |B_0|$, so there is some one-to-one function h_0 mapping A_0 onto B_0. Define $h : A \to B$ by $h(a) = h_0(a)$ for $a \in A_0$ and $h(a) = c^{\mathfrak{B}}$ otherwise. Then $h : \mathfrak{A} \cong \mathfrak{B}$.

EXAMPLE 5.4. Let $\mathcal{L}$have as its only non-logical symbols the constants c_n for $n \in \omega$. Let $T = Cn_{\mathcal{L}}(\{\neg \equiv c_n c_m | \ n < m \in \omega\})$. Let $\mathfrak{A}, \mathfrak{B}$ both be models of T with $|A| = |B| > \omega$. Then $\mathfrak{A} \cong \mathfrak{B}$.

Let $A_0 = A - \{c_n^{\mathfrak{A}} | \ n \in \omega\}$ and let $B_0 = B - \{c_n^{\mathfrak{B}} | \ n \in \omega\}$. Then $|A_0| = |A| = |B| = |B_0|$ since A and B are uncountable. Thus there is some one-to-one h_0 mapping A_0 onto B_0. Define $h : A \to B$ by $h(a) = h_0(a)$ if $a \in A_0$, $h(c_n^{\mathfrak{A}}) = c_n^{\mathfrak{B}}$, all $n \in \omega$. Then h is well-defined and one-to-one, since both $\mathfrak{A}, \mathfrak{B} \models T$, hence $h : \mathfrak{A} \cong \mathfrak{B}$.

The reader should check that this theory has exactly ω many non-isomorphic countable models–one for each cardinality for A_0.

The statement of the theorem above on isomorphisms does *not* say that $\mathfrak{A}_A \equiv \mathfrak{B}_B$. This usually wouldn't even make sense since these would be structures for different languages. However if $h : \mathfrak{A} \cong \mathfrak{B}$ then $\mathfrak{B}$ can be expanded to an $\mathcal{L}(A)$-structure $\mathfrak{B}^*$ such that $\mathfrak{A}_A \equiv \mathfrak{B}^*$ by defining $(c_a)^{\mathfrak{B}^*} = h(a)$, and if we add the condition that $B = \{c^{\mathfrak{B}^*} | \ c \in \mathcal{L}(A)\}$ then this is equivalent to $\mathfrak{A} \cong \mathfrak{B}$.

6. Exercises

(1) Show directly from the definitions or produce a counterexample to each of the following.
   (a) $\models (\exists x Px \to \forall x Qx) \to \forall x(Px \to Qx)$
   (b) $\models (Px \to Qy) \to (\exists x Px \to \exists y Qy)$

(2) Establish the following, using the rules from the notes.
$$\vdash \forall x(Px \to Qy) \to (\exists x Px \to Qy)$$

(3) (a) Let $\mathcal{L}^{nl} = \{R, d\}$ where R is a binary relation symbol and d is an individual constant symbol, and let $\mathfrak{A} = (\mathbb{Z}, <, 0)$.
      (i) Find an $\mathcal{L}$-formula $\varphi(x)$ such that $\varphi^{\mathfrak{A}} = \{-1\}$
      (ii) Determine $\psi^{\mathfrak{A}}$ where $\psi(x, y)$ is $\forall z(Rxz \to \neg Rzy)$.
   (b) Let $\mathcal{L}^{nl} = \{F\}$ where F is a binary function symbol. Let $\mathfrak{A} = (\omega, +)$ and let $\mathfrak{B} = (\mathbb{Q}, +)$. Find a setence σ of $\mathcal{L}$ such that $\mathfrak{A} \models \sigma$ and $\mathfrak{B} \models \neg\sigma$.

(4) Let $\mathcal{L}$ be any language and let $\Phi = \{\varphi_n : n \in \omega\} \subseteq \mathrm{Fm}_{\mathcal{L}}$. Assume that Φ is inconsistent. Prove that there is some $n \in \omega$ such that
$$\vdash (\neg\varphi_0 \vee \cdots \vee \neg\varphi_n)$$

(5) Establish the following:
$$\vdash (\exists x \forall y Rxy \to \exists y Ryy).$$

(6) Prove Theorem 2.2.2 about unique readability of terms.

(7) Let $\mathcal{L}^{nl} = \{R, E, P\}$ where R is binary and E and P are both unary relation symbols. Let $\mathfrak{A}$ be the $\mathcal{L}$-structure with $A = \omega$ where $R^{\mathfrak{A}}$ is $\leq$, $E^{\mathfrak{A}}$ is the set of even numbers and $P^{\mathfrak{A}}$ is the set of prime numbers. Give sentences of $\mathcal{L}$ which 'naturally' express the following facts about $\mathfrak{A}$.
   (a) There is no largest prime number.
   (b) There is exactly one even prime number.
   (c) The smallest prime number is even.
   (d) The immediate successor of an even number is not even.

(8) Let $\mathcal{L}^{nl} = \{R\}$ where R is a binary relation symbol, and let $\mathfrak{A} = (\omega, <)$. Determine $\varphi^{\mathfrak{A}}$ where:

(a) φ is $\forall y(Ryx \to \neg Ryx)$
(b) φ is $\forall z(Rzx \vee Ryz)$

(9) Let $\mathcal{L}^{nl} = \{F, G\}$ where F and G are both binary function symbols, and let $\mathfrak{B} = (\mathbb{R}, +, \cdot)$. Find $\mathcal{L}$-formulas with the following properties:
(a) $\varphi_1^{\mathfrak{B}} = \{1\}$
(b) $\varphi_2^{\mathfrak{B}} = \{a \in \mathbb{R} : 0 \leq a\}$
(c) $\varphi_3^{\mathfrak{B}} = \{(a_1, a_2) \in \mathbb{R}^2 : a_1 < a_2\}$
(d) $\varphi_4^{\mathfrak{B}} = \{(a_1, a_2) \in \mathbb{R}^2 : |a_1| = |a_2|\}$

(10) Let $\mathcal{L}^{nl} = \{R\}$ where R is a binary relation symbol, and let $\mathfrak{A}_1 = (\omega, \leq), \mathfrak{A}_2 = (\mathbb{Z}, \leq), \mathfrak{A}_3 = (\mathbb{Q}, \leq)$ and $\mathfrak{A}_4 = (\omega \smallsetminus \{0\}, |)$. For each $1 \leq i < j \leq 4$ find a sentence $\sigma_{i,j}$ of $\mathcal{L}$ such that $\mathfrak{A}_i \models \sigma_{i,j}$ but $\mathfrak{A}_j \models \neg\sigma_{i,j}$.

(11) Fix a sequence $\{\theta_i\}_{i \in \omega}$ of sentences of $\mathcal{L}$. For any sentence σ of *sentential* logic let σ^* be the result of replacing all occurences of S_i in σ by θ_i for each $i \in \omega$.
(a) Give a definition by recursion of σ^*.
(b) Define, for each $\mathcal{L}$-structure $\mathfrak{A}$, a truth assignment $h_{\mathfrak{A}}$ such that for every σ of sentential logic we have $h_{\mathfrak{A}} \models \sigma$ iff $\mathfrak{A} \models \sigma^*$. (Note that it follows that σ^* is valid whenever σ is a tautology.)

(12) Given $\Sigma_1, \Sigma_2 \subseteq \mathrm{Sn}_{\mathcal{L}}$ define $\Sigma^* = \{\theta \in \mathrm{Sn}_{\mathcal{L}} : \Sigma_1 \models \theta \text{ and } \Sigma_2 \models \theta\}$. Prove that for every $\mathcal{L}$-structure $\mathfrak{A}, \mathfrak{A} \models \Sigma^*$ iff either $\mathfrak{A} \models \Sigma_1$ or $\mathfrak{A} \models \Sigma_2$.

CHAPTER 3

The Completeness Theorem

0. Introduction

In this chapter we prove the most fundamental result of first order logic, the Completeness Theorem. The two forms of Soundness, from Chapter 2, give one direction of the biconditionals, and thus we need to establish the following.

THEOREM 0.1. *Let* $\Gamma \subseteq \mathrm{Fm}_{\mathcal{L}}$. *Then: (1) for any* $\phi \in \mathrm{Fm}_{\mathcal{L}}$*, if* $\Gamma \models \phi$ *then* $\Gamma \vdash \phi$*; (2) if* Γ *is consistent then* Γ *is satisfiable.*

Further, due to the equivalences $\Gamma \vdash \phi$ iff $\Gamma \cup \{\neg\phi\}$ is not satisfiable, it suffices to establish (2) above. As we will argue in more detail later, the full version of (2) for sets Γ of formulas with free variables follows from the version for sentences. We therefore concentrate on establishing the following Model Existence Theorem.

THEOREM 0.2. *Every consistent set* Σ *of sentences of* $\mathcal{L}$*has a model.*

In our proof of this result we will first define a special sort of sets of sentences called Henkin sets–these sets will look roughly like sets of the form $Th(\mathfrak{A}_A)$. We will show that Henkin sets determine structures which are their models, in essentially the same way in which we could recover $\mathfrak{A}_A$ from $Th(\mathfrak{A}_A)$. This will be done in the next section. In Section 3.3, we will show that if Σ is any consistent set of sentences of $\mathcal{L}$, then $\Sigma \subseteq \Gamma$ for some Henkin set Γ of sentences of $\mathcal{L}'$, where $\mathcal{L}'$ is a language obtained from $\mathcal{L}$by adding new constant symbols. This will finish the proof of the Completeness Theorem.

In Section 4, we derive two of the most important consequences of the Completeness Theorem. These are the Compactness Theorem and Löwenheim-Skolem Theorem. We then give several applications of these results to the study of theories and their models, as initiated in the preceding chapter.

1. Henkin Sets and Their Models

How could we use a set of sentences to define a structure which will be a model of the set? In defining a structure you need to specify the elements of the (universe of the) structure and specify the interpretations of the non-logical symbols of the language on the universe. For this information to be provided by a set of sentences, the language should contain constants for all elements of the structure and contain all the sentences with these constants true on the structure. Thus it should be roughly like $Th(\mathfrak{A}_A)$ in the language $\mathcal{L}(A)$.

The trick here is to decide on the necessary properties the set of sentences should have without knowing $\mathfrak{A}$ to begin with. The following definitions collect all the properties we will need.

DEFINITION 1.1. Let $\Gamma \subseteq \mathrm{Sn}_{\mathcal{L}}$.

(a) Γ is a complete set of sentences of $\mathcal{L}$iff

(i) for every sentence θ of $\mathcal{L}\neg\theta \in \Gamma$ iff $\theta \notin \Gamma$, and

(ii) for all $\phi, \psi \in \mathrm{Sn}_{\mathcal{L}}$, $(\phi \to \psi) \in \Gamma$ iff either $\neg\phi \in \Gamma$ or $\psi \in \Gamma$.

(b) Γ has witnesses iff for every $\phi(x) \in \mathrm{Fm}_{\mathcal{L}}$ we have $\forall x\phi \in \Gamma$ iff $\phi(c) \in \Gamma$ for all constants c of $\mathcal{L}$.

(c) Γ respects equality iff the following hold for all closed terms $t, t', t_1, \ldots, t_m, t'_1, \ldots, t'_m$ of $\mathcal{L}$:

(i) $\equiv tt \in \Gamma$,

(ii) if $\equiv tt' \in \Gamma$ then $\equiv t't \in \Gamma$,

(iii) if $\equiv t_1t_2 \in \Gamma$ and $\equiv t_2t_3 \in \Gamma$ then $\equiv t_1t_3 \in \Gamma$,

(iv) if $\equiv t_it'_i \in \Gamma$ for all $i = 1, \ldots, m$ then $\equiv Ft_1 \ldots t_mFt'_1 \ldots t'_m \in \Gamma$ for every m-ary function symbol F of $\mathcal{L}$,

(v) if $\equiv t_it'_i \in \Gamma$ for all $i = 1, \ldots, m$ and if $Rt_1 \ldots t_m$ holds where R is an m-ary relation symbol of $\mathcal{L}$then $Rt'_1 \ldots t'_m \in \Gamma$,

(vi) there is some constant c with $\equiv tc \in \Gamma$.

Note that we continue in this chapter to assume that all sentences are written using just the connectives $\neg, \to$ and the quantifier $\forall$. A set Γ can be a complete set of sentences without being a complete theory, since it need not have a model. The additional properties are needed to guarantee that the set has a model.

If $\mathfrak{A}$ is any $\mathcal{L}$-structure and we define $\Gamma = Th(\mathfrak{A}_A)$ then Γ is a complete set of $\mathcal{L}(A)$-sentences which has witnesses and respects equality. Our first goal is to see that every complete set of sentences with witnesses which respects equality determines a structure which is a model of the set.

Definition 1.2. Let $\Gamma \subseteq \mathrm{Sn}_{\mathcal{L}}$ and let $\mathfrak{A}$ be an $\mathcal{L}$-structure. Then $\mathfrak{A}$ is a canonical structure determined by Γ iff the following hold:

(i) $a = \{c^{\mathfrak{A}} |\ c \text{ is a constant of } \mathcal{L}\}$,

(ii) for each m-ary relation symbol R of $\mathcal{L}$and all $c_1, \ldots, c_m \in \mathcal{L}$, $R^{\mathfrak{A}}(c_1^{\mathfrak{A}}, \ldots, c_m^{\mathfrak{A}})$ holds iff $Rc_1 \ldots c_m \in \Gamma$.,

(iii) for all c_1, c_2 constants in $\mathcal{L}$, $c_1^{\mathfrak{A}} = c_2^{\mathfrak{A}}$ iff $\equiv c_1c_2 \in \Gamma$,

(iv) for each m-ary function symbol F of $\mathcal{L}$and all $c_1, \ldots, c_m, d \in \mathcal{L}$, $F^{\mathfrak{A}}(c_1^{\mathfrak{A}}, \ldots, c_m^{\mathfrak{A}}) = d^{\mathfrak{A}}$ iff $\equiv Fc_1 \ldots c_md \in \Gamma$.

Note that a given Γ may not determine any canonical structure, but any two canonical structures determined by Γ are isomorphic. Thus we will speak of *the* canonical structure for Γ. $\mathfrak{A}_A$ is the canonical structure determined by $Th(\mathfrak{A}_A)$.

Theorem 1.1. *Let Γ be a complete set of sentences of $\mathcal{L}$which has witnesses and respects equality. Then Γ determines a canonical structure.*

Proof. Let C be the set of constant symbols in $\mathcal{L}$. We define the relation $\sim_\Gamma$ on C by $c \sim_\Gamma d$ iff $\equiv cd \in \Gamma$. Then $\sim_\Gamma$ is an equivalence relation on C, since Γ respects equality.

For $c \in C$ we define $c/\!\sim_\Gamma = \{d \in C|\ c \in C\}$.
We now proceed to define the structure $\mathfrak{A}$. Let A be any set such that $|A| = |\{c/\!\sim_\Gamma \mid c \in C\}|$, and let $h : \{c/\!\sim_\Gamma \mid c \in C\} \to A$ be one-to-one and onto.

For any constant $c \in \mathcal{L}$ (hence in C), define $c^{\mathfrak{A}}$ by $c^{\mathfrak{A}} = h(c/\!\sim_\Gamma)$.
For any m-ary relation symbol R of $\mathcal{L}$and any $c_1, \ldots, c_m \in C$, define $R^{\mathfrak{A}}$ by saying $R^{\mathfrak{A}}(c_1^{\mathfrak{A}}, \ldots, c_m^{\mathfrak{A}})$ holds iff $Rc_1 \ldots c_m \in \Gamma$.
For any m-ary function symbol F of $\mathcal{L}$and any $c_1, \ldots, c_m, d \in C$ define $F^{\mathfrak{A}}$ by $F^{\mathfrak{A}}(c_1^{\mathfrak{A}}, \ldots, c_m^{\mathfrak{A}}) = d^{\mathfrak{A}}$ iff $\equiv Fc_1 \ldots c_md \in \Gamma$.

Then $R^{\mathfrak{A}}, F^{\mathfrak{A}}$ are well-defined since Γ respects equality, and also $c_1^{\mathfrak{A}} = c_2^{\mathfrak{A}}$ iff $\equiv c_1c_2 \in \Gamma$. Thus $\mathfrak{A}$ is a canonical structure determined by Γ, as desired. $\square$

THEOREM 1.2. *Let Γ be a complete set of sentences of $\mathcal{L}$which has witnesses and respects equality. Let $\mathfrak{A}$ be the canonical structure determined by Γ. Then $\mathfrak{A} \models \Gamma$.*

PROOF. We first must show that for every closed term t of $\mathcal{L}$and every constant symbol c of $\mathcal{L}$, if $\equiv ct \in \Gamma$ then $t^{\mathfrak{A}} = c^{\mathfrak{A}}$. We leave this proof, by induction on $lh(t)$ and using the assumption that Γ respects equality, to the reader.

We now show, by induction on $lh(\theta)$ for $\theta \in \mathrm{Sn}_{\mathcal{L}}$, that $\mathfrak{A} \models \theta$ iff $\theta \in \Gamma$.

As inductive hypothesis suppose this equivalence holds for all $\theta \in \mathrm{Sn}_{\mathcal{L}}$ with $lh(\theta) < n$. Now consider θ with $lh(\theta) = n$. There are several cases.

If θ is $\equiv t_1t_2$ for closed terms t_1, t_2 then (since Γ respects equality) there are constants $c_1, c_2 \in \mathcal{L}$ such that $\equiv c_1t_1, \equiv c_2t_2 \in \Gamma$. By the preliminary result on terms, $t_1^{\mathfrak{A}} = c_1^{\mathfrak{A}}$ and $t_2^{\mathfrak{A}} = c_2^{\mathfrak{A}}$. Since Γ respects equality we know $\equiv c_1c_2 \in \Gamma$, hence $c_1^{\mathfrak{A}} = c_2^{\mathfrak{A}}$ by the definition of canonical structure. Therefore $t_1^{\mathfrak{A}} = t_2^{\mathfrak{A}}$, hence $\mathfrak{A} \models \theta$.

The argument is similar when θ is $Rt_1 \ldots t_m$ and therefore left to the reader.

The cases in which $\theta = \neg\phi$ or $\theta = (\phi \to \psi)$ are easy using the hypothesis that Γ is a complete set of sentences.

We conclude by considering the case $\theta = \forall x\phi(x)$. Then, since $\mathfrak{A}$ is a canonical structure we have $\mathfrak{A} \models \theta$ iff $\mathfrak{A} \models \phi(c)$ for all $c \in \mathcal{L}$. Now

$$lh(\phi(c)) = lh(\phi) < lh(\theta)$$

so the inductive hypothesis holds for every $\phi(c)$. Therefore we see $\mathfrak{A} \models \theta$ iff $\phi(c) \in \Gamma$ for all $c \in \mathcal{L}$. But Γ has witnesses, so $\phi(c) \in \Gamma$ for all $c \in \mathcal{L}$ iff $\forall x\phi \in \Gamma$. Thus $\mathfrak{A} \models \theta$ iff $\theta \in \Gamma$, which finishes the proof. $\square$

Note that we have nowhere assumed that the set Γ is consistent, although it follows from having a model. When we add consistency, the list of other properties Γ must have for the preceding two theorems to hold can be considerably shortened.

DEFINITION 1.3. Let $\Gamma \subseteq \mathrm{Sn}_{\mathcal{L}}$. Then Γ is a Henkin set of sentences of $\mathcal{L}$iff the following hold:

(i) Γ is consistent,

(ii) for every $\theta \in \mathrm{Sn}_{\mathcal{L}}$ either $\theta \in \Gamma$ or $\neg\theta \in \Gamma$,

(iii) for every formula $\phi(x)$ of $\mathcal{L}$if $\neg\forall x\phi(x) \in \Gamma$ then $\neg\phi(c) \in \Gamma$ for some $c \in \mathcal{L}$.

LEMMA 1.3. *Let Γ be a Henkin set of sentences of $\mathcal{L}$. Then for every $\phi \in \mathrm{Sn}_{\mathcal{L}}$, $\phi \in \Gamma$ iff $\Gamma \vdash \phi$.*

PROOF. From left to right is clear. For the other direction, if $\phi \notin \Gamma$ then $\neg\phi \in \Gamma$ (by condition (ii) in the definition of Henkin sets), so $\Gamma \vdash \neg\phi$ and therefore $\Gamma \nvdash \phi$ by the consistency of Γ. $\square$

We thus establish the following theorem.

THEOREM 1.4. *Let $\Gamma \subseteq \mathrm{Sn}_{\mathcal{L}}$. Then Γ is a Henkin set iff Γ is complete, has witnesses, and respects equality.*

PROOF. If Γ is complete, has witnesses and respects equaltiy then we have shown that Γ has a model, so Γ is consistent and therefore a Henkin set.

For the other direction, assume Γ is a Henkin set. Then Γ is easily shown to be complete, using the lemma. Similarly Γ is easily shown to have witnesses, using the lemma and the fact that

$$\Gamma \vdash (\forall x \phi(x) \to \phi(c)$$

for all constants c of $\mathcal{L}$. Further, Γ is easily seen to respect equality, using the lemma and noting (for condition (vi)) that $\vdash \neg\forall x \neg \equiv tx$ hence $\neg\forall x \neg \equiv tx \in \Gamma$, and so $\equiv tc \in \Gamma$ for some c. $\square$

To summarize, we have shown that every Henkin set of sentences has a model, which is in fact the canonical structure it determines. In the next section we show that consistent sets of sentences can be extended to Henkin sets of sentences–in some larger language.

2. Constructing Henkin Sets

The work in the preceding section reduces the Model Existence Theorem at the beginning of section 3.1, to the following, purely syntactical result.

THEOREM 2.1. *Let Σ be a consistent set of sentences of $\mathcal{L}$. Then there is some $\mathcal{L}' \supseteq \mathcal{L}$ and some Henkin set Γ of sentences of $\mathcal{L}'$ such that $\Sigma \subseteq \Gamma$.*

It is clear that we will normally be forced to have $\mathcal{L} \neq \mathcal{L}'$ in this theorem (for example, $\mathcal{L}$may not have any constant symbols). On the other hand, we will need to guarantee that the set Σ we begin with remains consistent with respect to $\mathcal{L}'$-deductions. Thus we will need to have that $(\mathcal{L}' - \mathcal{L})$ consists just of constant symbols.

Given $\mathcal{L}$, the question is: how many new constant symbols must we add to $\mathcal{L}$to obtain a language $\mathcal{L}'$ such that the above extension result can be established? The answer–which is easier to justify after the fact–is that $|\mathcal{L}|$ new constants will suffice.

We give the argument in detail for the case in which $\mathcal{L}$is countable, after which we indicate how the general case could be proved.

So, fix a countable language $\mathcal{L}$. Let C be a countably infinite set of individual constant symbols not in $\mathcal{L}$and let $\mathcal{L}' = \mathcal{L} \cup C$. Then $\mathcal{L}'$ is also countable thus in particular $|Sn_{\mathcal{L}'}| = \omega$, so we may list the sentences of $\mathcal{L}'$ as $Sn_{\mathcal{L}'} = \{\sigma_n | \ n \in \omega\}$.

Now let Σ be a consistent set of sentences of $\mathcal{L}$. Then Σ remains consistent with respect to $\mathcal{L}'$-deductions, as established by Theorem 2.4.5. We wish to define a set Γ of sentences of $\mathcal{L}'$ which is a Henkin set and contains Σ. We will do this by defining, by recursion on ω, a chain $\{\Gamma_n\}_{n\in\omega}$ of subsets of $Sn_{\mathcal{L}'}$ whose union is the desired Henkin set. This chain will have the following properties: $\Gamma_0 = \Sigma$ and, for each $n \in \omega$:

(0n) $\Gamma_n \subseteq \Gamma_{n+1}$, Γ_{n+1} is consistent, and $(\Gamma_{n+1} - \Gamma_n)$ is finite;

(1n) either $\sigma_n \in \Gamma_{n+1}$ or $\neg\sigma_n \in \Gamma_{n+1}$;

(2n) if $\sigma_n = \forall x \phi(x)$ for some $\phi(x)$ and if $\neg\sigma_n \in \Gamma_{n+1}$ then $\neg\phi(c) \in \Gamma_{n+1}$ for some $c \in C$.

Assuming $\{\Gamma_n\}_{n\in\omega}$ is such a chain we show that $\Gamma = \bigcup_{n\in\omega} \Gamma_n$ is a Henkin set of sentences of $\mathcal{L}'$.

We first show Γ is consistent. If not then some finite $\Gamma' \subseteq \Gamma$ is inconsistent by the Finiteness Lemma. But then $\Gamma' \subseteq \Gamma_{n+1}$ for some $n \in \omega$ and so Γ_{n+1} would be inconsistent, contradicting (0n).

Next, let $\theta \in Sn_{\mathcal{L}'}$. Then $\theta = \sigma_n$ for some $n \in \omega$ and so, by (1n), either

$\theta = \sigma_n \in \Gamma_{n+1} \subseteq \Gamma$ or
$\neg\theta = \neg\sigma_n \in \Gamma_{n+1} \subseteq \Gamma$.

Finally, if $\phi(x) \in Fm_{\mathcal{L}'}$ and if $\neg\forall x\phi(x) \in \Gamma$ then $\forall x\phi(x) = \sigma_n$ for some $n \in \omega$ and necessarily $\neg\sigma_n \in \Gamma_{n+1}$ (since $\sigma_n \in \Gamma_{n+1} \subseteq \Gamma$ would contradict the consistency of Γ), and so, by (2n), $\neg\phi(c) \in \Gamma_{n+1} \subseteq \Gamma$ for some constant c.

We now show how to construct such a chain $\{\Gamma_n\}_{n\in\omega}$ by recursion. Suppose we have $\Sigma = \Gamma_0 \subseteq \Gamma_1 \subseteq \cdots \subseteq \Gamma_n$ satisfying (0i), (1i), (2i) for all $i < n$. Then Γ_n is consistent and $|\Gamma_n - \Sigma| < \omega$. We show how to define Γ_{n+1} so that (0n), (1n), (2n) hold. We obtain Γ_{n+1} in two steps. The first will guarantee (1n) and the second will guarantee (2n). To ensure (0n) we will only need to preserve consistency in performing these steps.

We define $\Gamma_{n+\frac{1}{2}}$ as follows:
$\Gamma_{n+\frac{1}{2}} = \Gamma_n \cup \{\sigma_n\}$ if $\Gamma_n \cup \{\sigma_n\}$ is consistent;
$\Gamma_{n+\frac{1}{2}} = \Gamma_n \cup \{\neg\sigma_n\}$ if $\Gamma_n \cup \{\sigma_n\}$ is inconsistent.

We claim that in either case $\Gamma_{n+\frac{1}{2}}$ is consistent. This is clear in the first case. If, on the other hand, $\Gamma_n \cup \{\sigma_n\}$ is inconsistent then $\Gamma_n \vdash \neg\sigma_n$ and so $\Gamma_{n+\frac{1}{2}} = \Gamma_n \cup \{\sigma_n\}$ is consistent because Γ_n is consistent.

From $\Gamma_{n+\frac{1}{2}}$ we define Γ_{n+1} as follows $\Gamma_{n+1} = \Gamma_{n+\frac{1}{2}}$ if either $\sigma_n \in \Gamma_{n+\frac{1}{2}}$ or σ_n does not have the form $\forall x\phi$; $\Gamma_{n+1} = \Gamma_{n+\frac{1}{2}} \cup \{\neg\phi(c)\}$ where $c \in C$ does not occur in any sentence in $\Gamma_{n+\frac{1}{2}}$ if $\sigma_n = \forall x\phi(x)$ and $\neg\sigma_n \in \Gamma_{n+\frac{1}{2}}$. We must show Γ_{n+1} is consistent in the second case. If $\Gamma_{n+\frac{1}{2}} \cup \{\neg\phi(c)\}$ is inconsistent then $\Gamma_{n+\frac{1}{2}} \vdash \phi(c)$. But c does not occur in $\Gamma_{n+\frac{1}{2}}$ nor in $\phi(x)$, and so Generalization on Constants yields $\Gamma_{n+\frac{1}{2}} \vdash \forall x\phi(x)$. But this contradicts the consistency of $\Gamma_{n+\frac{1}{2}}$. Finally note that there will be constants in C not occurring in $\Gamma_{n+\frac{1}{2}}$ since no constants in C occur in Σ, $|\Gamma_{n+\frac{1}{2}} - \Sigma| < \omega$, and C is infinite.

This completes the proof of the theorem stated at the beginning of this section for countable languages $\mathcal{L}$. The same outline may be used for uncountable languages, but now transfinite recursion is needed.

Suppose $|\mathcal{L}| = \kappa$ and let $\mathcal{L}' = \mathcal{L} \cup C$ where C is a set of new individual constant symbols, $|C| = \kappa$. Then $|\mathcal{L}'| = \kappa$ hence $|Sn_{\mathcal{L}'}| = \kappa$ and so $Sn_{\mathcal{L}'}$ can be listed as $Sn_{\mathcal{L}'} = \{\sigma_\xi |\ \xi \in \kappa\}$. We define, by recursion on κ, a chain $\{\Gamma_\xi\}_{\xi\in\kappa}$ of subsets of $Sn_{\mathcal{L}'}$ so that $\Sigma = \Gamma_0$ and for all $\xi \in \kappa$ we have

(0ξ) $\Gamma_\xi \subseteq \Gamma_{\xi+1}$, $\Gamma_{\xi+1}$ is consistent, and $(\Gamma_{\xi+1} - \Gamma_\xi)$ is finite;
(1ξ) either $\sigma_\xi \in \Gamma\xi + 1$ on $\neg\sigma \in \Gamma_{\xi+1}$;
(2ξ) if $\sigma_\xi = \forall x\phi(x)$ and $\neg\sigma_\xi \in \Gamma_{\xi+1}$ then $\neg\phi(c) \in \Gamma_{\xi+1}$ for some $c \in C$;
and further if ξ is a limit ordinal then $\Gamma_\xi = \bigcup_{\nu<\xi} \Gamma_\nu$.

The construction proceeds exactly as before noting, for (2ξ), that is $\xi \in \kappa$ and (0ν) holds for all $\nu < \xi$ then $|\Gamma_\xi - \Sigma| < \kappa = |C|$ and hence there will be constants in C not occurring in $\Gamma_\xi \cup \{\neg\sigma_\xi\}$.

Since a Henkin set Γ completely determines a canonical structure up to isomorphism, we can control some properties of the canonical structure by constructing a Henkin set with specific properties. We will exploit this later, especially in deriving the Omitting Types Theorem.

Note that if Σ is a consistent set of sentences of $\mathcal{L}$then our proof yields an $\mathcal{L}$-structure which is a model of Σ–namely the reduct to $\mathcal{L}$of the canonical $\mathcal{L}'$-structure determined by the Henkin set Γ of $\mathcal{L}'$-sentences we construct containing Σ.

3. Consequences of the Completeness Theorem

We very seldom will use the Completeness Theorem directly, thus we will not prove that ta set Σ of sentences has a model by showing that it is consistent. Instead we will derive two purely semantic (or, model theoretic) consequences which will give us virtually everything we will need from Completeness, for the purposes of model theory. (A third consequence will be given much later in the context of decidability and the decision problem).

The first of these consequences is the Compactness Theorem, a direct translation of the Finiteness Lemma. We will argue for the full version later. Since we have established Completeness for sets of sentences we can at this point conclude the following version.

THEOREM 3.1. ***(Compactness)*** *Let $\Sigma \subseteq \mathrm{Sn}_{\mathcal{L}}$.*
(1) For any $\phi \in \mathrm{Fm}_{\mathcal{L}}$, $\Sigma \models \phi$ iff $\Sigma_0 \models \phi$ for some finite $\Sigma_0 \subseteq \Sigma$.
(2) Σ has a model iff every finite $\Sigma_0 \subseteq \Sigma$ has a model.

To allow formulas ϕ in part (1) we use the fact that if $\phi = \phi(x_0, \ldots, x_n)$ then $\Sigma \models \phi$ iff $\Sigma \models \forall x_0 \cdots \forall x_n \phi$.

The force of the Compactness Theorem in form (2), as we will see in many examples, is that you may be able to produce a model $\mathfrak{A}_{\Sigma_0}$ of each finite $\Sigma_0 \subseteq \Sigma$ without knowing how to define a model of the entire set Σ.

The second consequence is actually a consequence of our proof of the Completeness Theorem using canonical structures determined by Henkin sets.

THEOREM 3.2. ***(Löwenheim-Skolem)*** *Let $\kappa = |\mathcal{L}|$ and assume that $\Sigma \subseteq \mathrm{Sn}_{\mathcal{L}}$ has a model. Then Σ has a model $\mathfrak{A}$ with $|A| \leq \kappa$.*

PROOF. Since Σ has a model, and is therefore consistent, our proof of the Completeness Theorem produced an $\mathcal{L}$-structure $\mathfrak{A}$ which is a model of Σ and which is the reduct to $\mathcal{L}$of the canonical structure $\mathfrak{A}'$ determined by a Henkin set of sentences of some $\mathcal{L}'$ where $|\mathcal{L}'| = \kappa$. But

$$A = A' = \{c^{\mathfrak{A}'} \mid c \in \mathcal{L}' \text{ is a constant}\}.$$

Therefore $|A| \leq |\mathcal{L}'| = \kappa$. $\square$

As a first application of the Löwenheim-Skolem Theorem note that if $\mathfrak{B}$ is any $\mathcal{L}$-structure, $|\mathcal{L}| = \kappa$, then there is some $\mathcal{L}$-structure $\mathfrak{A}$ with $|A| \leq \kappa$ such that $\mathfrak{A} \equiv \mathfrak{B}$. This is immediate by considering $\Sigma = Th(\mathfrak{B})$. The reader should consider what such a structure $\mathfrak{A}$ would be if, for example, $\mathfrak{B}$ is $(\mathbb{R}, \leq)$ or $(\mathbb{R}, \leq, +, \cdot)$– both structures for countable languages, hence elementarily equivalent to countable structures.

The Compactness Theorem is one of the most important tools we have in model theory. We give here some easy, but typical, examples of its use.

To begin with the Compactness Theorem, in its second form stated above, is a model existence theorem–it asserts the existence of a model of Σ, given the existence of models of every finite subset of Σ. We will frequently use it to prove the existence of structures with certain specified properties. To do so we attempt to express the properties the structure should have by a set of sentences, and then use Compactness (or, later, other similar results) to show the set of sentences has a model.

As a first example of this procedure we prove the following result.

THEOREM 3.3. *Let $\Sigma \subseteq \mathrm{Sn}_{\mathcal{L}}$ and assume that for every $n \in \omega$, Σ has a model $\mathfrak{A}_n$ with $|A_n| \geq n$. Then Σ has an infinite model.*

PROOF. Our conclusion asks for an $\mathcal{L}$-structure $\mathfrak{A}$ with two properties:
$\mathfrak{A} \models \Sigma$ and
$|A|$ is infinite.
Recall that we know a set Θ_ω of sentences (with no non-logical symbols) such for every $\mathcal{L}$-structure $\mathfrak{A}$, $|A|$ is infinite iff $\mathfrak{A} \models \Theta_\omega$. Thus what we want to prove is *precisely* that $(\Sigma \cup \Theta_\omega)$ has a model. By compactness it suffices to show that every finite $\Sigma_0 \subseteq \Sigma \cup \Theta_\omega$ has a model. Let $\Sigma_0 \subseteq \Sigma \cup \Theta_\omega$ be finite. Then $\Sigma_0 \subseteq \Sigma \cup \{\theta_k | \ 2 \leq k < n\}$ for some $n \in \omega$. But the hypotheses to this theorem provide a model $\mathfrak{A}_n$ of Σ with at least n elements. Therefore $\mathfrak{A}_n \models \Sigma \cup \{\theta_k | \ 2 \leq k < n\}$ and so $\mathfrak{A}_n \models \Sigma_0$. Thus $\Sigma \cup \Theta_\omega$ has a model, as desired. □

This theorem enables us to answer some questions we raised earlier.

COROLLARY 3.4. *There is no set Θ_f of sentences (of any $\mathcal{L}$) whose models are precisely the finite $\mathcal{L}$-structures. In particular, there is no sentence θ_ω whose models are precisely the infinite $\mathcal{L}$-structures.*

PROOF. Such a Θ_f would have arbitrarily large finite models, hence it would have infinite models by the theorem–a contradiction. □

We give now a second proof of the above theorem, indicating another technique–changing languages.

PROOF. We show another way to express "A is infinite" with a set of sentences. Let c_n, $n \in \omega$ be distinct constant symbols not in $\mathcal{L}$. Let $\mathcal{L}' = \mathcal{L} \cup \{c_n | \ n \in \omega\}$ and let $\Sigma' = \{\neg \equiv c_n c_m | \ n < m \in \omega\}$. Then an $\mathcal{L}$-structure is infinite iff it can be expanded to an $\mathcal{L}'$-structure which is a model of Σ'. Thus to show that the set Σ of sentences of $\mathcal{L}$has an infinite model it suffices to show that $(\Sigma \cup \Sigma')$ has a model, since the reduct to $\mathcal{L}$of such a model will be as desired. By compactness it suffices to show that every finite $\Sigma_0 \subseteq \Sigma \cup \Sigma'$ has a model. If $\Sigma_0 \subseteq \Sigma \cup \Sigma'$ is finite then $\Sigma_0 \subseteq \Sigma \cup \{\neg \equiv c_n c_m | \ n < m \leq k\}$ for some $k \in \omega$. Consider $\mathfrak{A}_{k+1}$–the model of Σ we are given which has at least $k+1$ elements. We show how to expand it to an $\mathcal{L}'$-structure which will be a model of Σ_0. Pick $a_0, \ldots, a_k \in A_k$ all different. Define $\mathfrak{A}'_{k+1}$ to be the expansion of $\mathfrak{A}_{k+1}$ such that $c_n^{\mathfrak{A}'_{k+1}} = a_0$ for all $n > k$. Then clearly $\mathfrak{A}'_{k+1}$ is as desired. □

The advantage of this second proof is that it generalizes to yield models of cardinality $\geq \kappa$ for every κ. All you need to do is add κ new constants and consider the set of sentences asserting they are all different. By combining this with the Löwenheim-Skolem Theorem we obtain an important stregthening of that result.

THEOREM 3.5. *Let $|\mathcal{L}| = \kappa$ and assume that $\Sigma \subseteq \mathrm{Sn}_{\mathcal{L}}$ has an infinite model. Let λ be a cardinal with $\kappa \leq \lambda$. Then Σ has a model of cardinality equal to λ.*

PROOF. Let C be a set of constant symbols not in $\mathcal{L}$, $|C| = \lambda$. Let $\mathcal{L}' = \mathcal{L} \cup C$. Note that $|\mathcal{L}'| = \lambda$ since $\kappa \leq \lambda$. Let $\Sigma' = \{\neg \equiv cd | \ c, d \in C, c \neq d\}$ and define $\Sigma^* = \Sigma \cup \Sigma'$.

Then, as described above, Σ^* has a model since every finite $\Sigma_0 \subseteq \Sigma^*$ has a model. Further, every model of Σ^* has cardinality $\geq \lambda$. Since $|\mathcal{L}'| = \lambda$ the original Löwenheim-Skolem result implies that Σ^* has a model $\mathfrak{A}'$ of cardinality $\leq \lambda$, and so of cardinality exactly $= \lambda$ due to the property of Σ^* mentioned previously. Therefore $\mathfrak{A} = \mathfrak{A}' \restriction \mathcal{L}$ is an $\mathcal{L}$-structure which is a model of Σ and which has cardinality exactly $= \lambda$. $\square$

COROLLARY 3.6. *Let $\mathfrak{A}$ be any infinite structure. Then there is some $\mathfrak{B}$ such that $\mathfrak{A} \equiv \mathfrak{B}$ but $|A| \neq |B|$, in particular, $\mathfrak{A} \not\cong \mathfrak{B}$.*

PROOF. Let $\Sigma = Th(\mathfrak{A})$ and let λ be such that $\lambda \geq |\mathcal{L}|$ and $\lambda \neq |A|$. Then any model $\mathfrak{B}$ of Σ of cardinality equal to λ suffices. $\square$

Given an infinite $\mathfrak{A}$ can we find some $\mathfrak{B} \equiv \mathfrak{A}$ such that $\mathfrak{A} \not\cong \mathfrak{B}$ *and* $|A| = |B|$? Clearly this is not always possible–for example if $\mathcal{L}^{nl} = \emptyset$ and $\mathfrak{A}$ is any $\mathcal{L}$-structure then $\mathfrak{A} \cong \mathfrak{B}$ for every $\mathcal{L}$-structure $\mathfrak{B}$ with $|A| = |B|$. We will see less trivial examples of this phenomenon later.

As a final application of Compactness in this section we show that there are structures $\mathfrak{A}$, $\mathfrak{B}$ for a countable language sucdh that $|A| = |B| = \omega$ but $\mathfrak{A} \not\cong \mathfrak{B}$.

THEOREM 3.7. *Let $\mathcal{L}$be the language whose only non-logical symbol is a binary relation symbol R. Let $\mathfrak{A}$ be the $\mathcal{L}$-structure $(\omega, \leq)$. Then there is some $\mathcal{L}$-structure $\mathfrak{B}$, $|B| = \omega$, such that $\mathfrak{A} \equiv \mathfrak{B}$ but $\mathfrak{A} \not\cong \mathfrak{B}$.*

PROOF. As a model existence result, this asserts the existence of an $\mathcal{L}$-structure $\mathfrak{B}$ having the following three properties: $\mathfrak{B} \equiv \mathfrak{A}$, $\mathfrak{A} \not\cong \mathfrak{B}$, and $|B| = \omega$. The first property is expressed by the first order condition $\mathfrak{B} \models Th(\mathfrak{A})$. Provided that we can guarantee the second property by some set of sentences in a countable language, then the last can be guaranteed using the Löwenheim-Skolem theorem.

To assert, with first order sentences, that $\mathfrak{A} \not\cong \mathfrak{B}$ we first recall that, for each $n \in \omega$, there are $\mathcal{L}$-formulas $\phi_n(x)$ such that $\phi_n^{\mathfrak{A}} = \{n\}$. Thus to be non-isomorphic to $\mathfrak{A}$ means having an element satisfying *no* ϕ_n.

So let $\mathcal{L}' = \mathcal{L} \cup \{d\}$ and let $\Sigma' = \{\neg\phi_n(d) |\ n \in \omega\}$. Then if $\mathfrak{B}' \models \Sigma'$ we must have

$$\mathfrak{B} = \mathfrak{B}' \restriction \mathcal{L} \not\cong \mathfrak{A}$$

since if $h : \mathfrak{B} \cong \mathfrak{A}$ then $\mathfrak{A}_A \models \neg\phi_n(c_{a^*})$ for all $n \in \omega$ where $a^* = h(d^{\mathfrak{B}'})$. Since no $a^* \in A$ has this property we must have $\mathfrak{B} \not\cong \mathfrak{A}$.

Thus to show that some countable $\mathfrak{B} \equiv \mathfrak{A}$ is not isomorphic to $\mathfrak{A}$ it suffices to show $(Th(\mathfrak{A}) \cup \Sigma')$ has a model.

By Compactness it suffices to show that every finite subset of $(Th(\mathfrak{A}) \cup \Sigma')$ has a model. Let $\Sigma_0 \subseteq Th(\mathfrak{A}) \cup \Sigma'$ be finite. Then

$$\Sigma_0 \subseteq Th(\mathfrak{A}) \cup \{\neg\phi_n(d) |\ n < k\}$$

for some $k \in \omega$. We show how to expand $\mathfrak{A}$ to an $\mathcal{L}'$-structure which is a model of Σ_0. Define $\mathfrak{A}'$ by $\mathfrak{A}' \restriction \mathcal{L} = \mathfrak{A}$ and $d^{\mathfrak{A}'} = k$. Then clearly

$$\mathfrak{A}' \models Th(\mathfrak{A}) \cup \{\neg\phi_n(d) |\ n < k\}$$

as desired. $\square$

A careful examination of this proof will reveal that it applies much more generally than these particular circumstances. We will give a very general version in the next chapter.

Note that this proof does not provide any description of what such a $\mathfrak{B}$ might look like. The reader should attempt to decide what a good candidate for such a structure would be.

4. Completeness Categoricity, Quantifier Elimination

In the preceding section we have proved that for any infinite $\mathcal{L}$-structure $\mathfrak{A}$ there are $\mathcal{L}$-structures $\mathfrak{B}$ such that $\mathfrak{A} \equiv \mathfrak{B}$ but $\mathfrak{A} \not\cong \mathfrak{B}$. In particular, if $\mathfrak{A}$ is $(\omega, \leq)$ there is a countable $\mathfrak{B}$ with this property. These proofs do not give a construction of $\mathfrak{B}$ or describe $\mathfrak{B}$ completely. The question thus remains of whether we can exhibit such a $\mathfrak{B}$ explicitly, at least given a specific $\mathfrak{A}$.

This leads us to the following, even more fundamental problem:

given non-isomorphic $\mathcal{L}$-structures $\mathfrak{A}, \mathfrak{B}$ how can we show that $\mathfrak{A} \equiv \mathfrak{B}$?

In this section we give a simple method that can be applied in some cases. Although not widely applicable, it will give us several important examples. We will improve the method later, but the problem in general is intractibly difficult. For example, let $\mathcal{L}$be the language for groups, let $\mathfrak{A}$ be (the $\mathcal{L}$-structure which is) a free group on 2 generators, and let $\mathfrak{B}$ be a free group on 3 generators. It has long been conjectured that $\mathfrak{A}$ and $\mathfrak{B}$ are elementarily equivalent, but no one has been able to decide the question.

We actually will consider a variant on this problem, namely:

given $\Sigma \subseteq \mathrm{Sn}_{\mathcal{L}}$ how can we show that $T = Cn(\Sigma)$ is complete?

This will apply to the preceding problem in the case in which Σ is an explicitly given set of sentences, and the method would be to show both $\mathfrak{A}$ and $\mathfrak{B}$ are models of Σ.

Certainly if any two models of Σ are isomorphic, then $T = Cn(\Sigma)$ is complete. But this will never be the case unless Σ only has models of a fixed finite cardinality. The method we give here, and will imporve on later, will only require that any two models of Σ, satisfying some other requirements, be isomorphic. These methods will apply to (some) theories with infinite models.

Definition 4.1. Let κ be a cardinal number and let T be a theory of some language $\mathcal{L}$. Then T is κ-categorical iff T has a model of cardinality κ and any two models of T of cardinality κ are isomorphic.

Theorem 4.1. ***(Łoŝ-Vaught Test)*** *Let T be a theory in a language $\mathcal{L}$and let κ be a cardinal number, $|\mathcal{L}| \leq \kappa$. Assume that T has no finite models and T is κ-categorical. Then T is complete.*

Proof. If T, satisfying the hypotheses, is not complete then there must be a sentence σ of h such that both $T \cup \{\sigma\}$ and $T \cup \{\neg\sigma\}$ have models. In fact, they must both have infinite models, so by the general form of the Löwenheim-Skolem Theorem they both have models of cardinality equal to κ. But since these models must be non-elementarily equivalent this contradicts the hypothesis of κ-categoricity. □

We give, without much detail, several applications of this test.

(1) $\mathcal{L}^{nl}= \emptyset$, $T = Cn(\emptyset)$. Then T is κ-categorical for all $\kappa \geq 1$, but T is not complete since it has finite models.

(2) $\mathcal{L}^{nl}= \emptyset$, $T = Cn(\Theta_\omega)$. Then T is κ-categorical for all infinite cardinals κ and has no finite models, hence T is complete.

(3) $\mathcal{L}^{nl}= \{P\}$ where P is a unary relation symbol. Let Σ be the set of all sentences asserting "P has at least n elements" and "$\neg P$ has at least n elements" for all positive $n \in \omega$. Then an $\mathcal{L}$-structure $\mathfrak{A}$ is a model of $T = Cn(\Sigma)$ iff $|P^{\mathfrak{A}}| \geq \omega$ and $|\neg P^{\mathfrak{A}}| \geq \omega$. It follows that T has no finite models and is ω-categorical, hence

is complete; but T is not κ-categorical for any uncountable κ. For example, T has exactly three non-isomorphic models of cardinality ω_1.

(4) $\mathcal{L}^{nl} = \{c_n |\ n \in \omega\}$, $T = Cn(\Sigma)$ where $\Sigma = \{\neg \equiv c_i c_j |\ i \neq j\}$. For any $\mathcal{L}$-structure $\mathfrak{A}$ we define $A_0 = \{c_n^{\mathfrak{A}} | n \in \omega\}$ and $A_1 = A - A_0$. We claim that, for models $\mathfrak{A}$ and $\mathfrak{B}$ of T, $\mathfrak{A} \cong \mathfrak{B}$ iff $|A_1| = |B_1|$. Supposing that $\mathfrak{A}, \mathfrak{B} \models T$ and $|A_1| = |B_1|$, we choose some h_1 mapping A_1 one-to-one onto B_1. We then define h on A by $h(c_n^{\mathfrak{A}}) = c_n^{\mathfrak{B}}$ for all $n \in \omega$, $h(a) = h_1(a)$ for all $a \in A_1$. Then $h : \mathfrak{A} \cong \mathfrak{B}$ since both structures are models of T. Now, for a model $\mathfrak{A}$ of T, A_1 can have any cardinality, finite (including 0) or infinite, but $|A_0| = \omega$. It follows that T is κ-categorical for all $\kappa > \omega$ but not ω-categorical.

5. Exercises

(1) Let T be a theory of $\mathcal{L}$ and $\sigma_n \in \mathrm{Sn}_{\mathcal{L}}$ for all $n \in \omega$. Assume that $T \models (\sigma_{n+1} \to \sigma_n)$ for all $n \in \omega$, but that $T \not\models (\sigma_n \to \sigma_{n+1})$ for all $n \in \omega$. Prove that there is some model $\mathfrak{A}$ of T such that $\mathfrak{A} \models \sigma_n$ for every $n \in \omega$.

(2) Give an example of a language $\mathcal{L}$ (with just finitely many non-logical symbols) and some finite set $\Sigma \subseteq \mathrm{Sn}_{\mathcal{L}}$ such that the theory $T = \mathrm{Cn}(\Sigma)$ has models and all models of T are infinite.

(3) Let $\mathcal{L}^{nl} = \{R\}$ where R is a binary relation symbol and let $\mathfrak{A} = (\omega, <)$. Let $\mathfrak{B}$ be such that $\mathfrak{A} \equiv \mathfrak{B}$ but $\mathfrak{A}$ is not isomorphic to $\mathfrak{B}$. Prove that there is some infinite sequence $\{b_n\}_{n\in\omega}$ of elements of B which is strictly decreasing, that is $R^{\mathfrak{B}}(b_{n+1}, b_n)$ holds for all $n \in \omega$.

(4) Let T_1 and T_2 be theories of $\mathcal{L}$, and assume that there is *no* sentence θ of $\mathcal{L}$ such that $T_1 \models \theta$ and $T_2 \models \neg\theta$. Prove that $(T_1 \cup T_2)$ has a model. [Warning: $(T_1 \cup T_2)$ need not be a theory.]

(5) Let T_1 and T_2 be theories of $\mathcal{L}$. Assume that for every $\mathcal{L}$-structure $\mathfrak{A}$ we have
$$\mathfrak{A} \models T_1 \text{ iff } \mathfrak{A} \not\models T_2.$$
Prove that there is some sentence σ of $\mathcal{L}$ such that $T_1 = \mathrm{Cn}(\sigma)$.

Part 2

Model Theory

CHAPTER 4

Some Methods in Model Theory

0. Introduction

In model theory one investigates the classes of models of (first order) theories. The questions largely concern the "variety" of possible models of a theory, and how different models of a theory are *related*. The most basic and important results of model theory yield the *existence* of models of a theory with specified properties. The Compactness and Löwenheim-Skolem theorems are both such results, and form the basis for many important applications, some of which were covered in Chapter 3, section 4.

In sections 1 and 2 of this chapter we present further model-existence results and some of their applications, which are pursued much further in the next chapter. Of especial importance is the concept of a *type* (of elements in a model) and the theorem on *omitting types* (in a model of a theory). Two natural relations of *inclusion* between models are introduced in section 3 and the method of *chains* (especially under elementary inclusion) to construct models is introduced.

In section 4, we introduce the "back-and-forth" method, widely used to show two models are isomorphic. This method will also be used in the next chapter.

1. Realizing and Omitting Types

If two models are non-isomorphic, how might one recognize the fact? Certainly if they are not elementarily equivalent or if they have different cardinalities they are non-isomorphic. If they are elementarily equivalent of the same cardinality, however, the question becomes harder, particularly for countable models. We have seen two models be non-isomorphic since there was a set $\{\psi_i(x) |\ i \in I\}$ of formulas such that in one model there was an element simultaneously satisfying every $\psi_i(x)$ while there was no such element in the other model. In the terminology we are about to introduce we could conclude the models were non-isomorphic since there was a *type* which was *realized* in one model, but *omitted* in the other model.

DEFINITION 1.1. (a) A type, in the variables $x_0, \ldots, x_n$, is a set Γ of formulas, with at most the variables $x_0, \ldots, x_n$ free. We write $\Gamma(x_0, \ldots, x_n)$ to mean that only $x_0, \ldots, x_n$ occur free in Γ.

(b) $\Gamma(x_0, \ldots, x_n)$ is realized in $\mathfrak{A}$ by $a_0, \ldots, a_n \in A$ iff $\mathfrak{A}_A \models \phi(\overline{a_0}, \ldots, \overline{a_n})$ for all $\phi(x_0, \ldots, x_n) \in \Gamma$. Γ is realized in $\mathfrak{A}$ if it is realized in $\mathfrak{A}$ by some $a_0, \ldots, a_n \in A$.

(c) $\Gamma(x_0, \ldots, x_n)$ is omitted in $\mathfrak{A}$ iff Γ is not realized in $\mathfrak{A}$, that is, for all $a_0, \ldots, a_n \in A$ there is some $\phi \in \Gamma$ such that $\mathfrak{A}_A \models \neg\phi(\overline{a_0}, \ldots, \overline{a_n})$.

The connection of types with isomorphisms is contained in this easy result.

PROPOSITION 1.1. *Assume h is an isomorphism of $\mathfrak{A}$ onto $\mathfrak{B}$. If $a_0, \ldots, a_n$ realize the type $\Gamma(x_0, \ldots, x_n)$ in $\mathfrak{A}$ then $h(a_0), \ldots, h(a_n)$ realize Γ in $\mathfrak{B}$.*

PROOF. This is immediate from the fact that

$$\mathfrak{A}_A \models \phi(\overline{a_0}, \ldots, \overline{a_n}) \text{ iff } \mathfrak{B}_B \models \phi(\overline{h(a_0)}, \ldots, \overline{h(a_n)}).$$

□

COROLLARY 1.2. *If $\mathfrak{A} \cong \mathfrak{B}$ then $\mathfrak{A}$ and $\mathfrak{B}$ realize precisely the same types.*

Thus, showing that two models realize different collections of types is one way of showing they are non-isomorphic, even if they are elementarily equivalent and of the same cardinality. One can ask if this will always work. That is: if $\mathfrak{A} \equiv \mathfrak{B}$, $|A| = |B| = \omega$ and $\mathfrak{A}$ and $\mathfrak{B}$ realize precisely the same types, must $\mathfrak{A} \cong \mathfrak{B}$?

More generally (or, rather, more vaguely) we want to use types to investigate the countable models of (complete) theories–how they help us distinguish between models, etc. For example, we know that $Th((\omega, \leq, +, \cdot, 0, 1))$ has at least 2 non-isomorphic countable models. Precisely how many non-isomorphic countable models does this theory have? We will be able to answer this question exactly by looking at the types realized in models of this theory.

The two questions we want to attack immediately are: when does a theory T have a model realizing (omitting) a type Γ?

If $\Gamma(x_0, \ldots, x_n)$ is *finite*, say $\Gamma = \{\phi_0(\vec{x}), \ldots, \phi_k(\vec{x})\}$, then obviously $\mathfrak{A}$ realizes Γ iff $\mathfrak{A} \models \exists x_0 \cdots \exists x_n(\phi_0 \wedge \cdots \wedge \phi_n)$. Thus T has a model realizing this finite Γ iff $T \cup \{\exists x_0 \cdots \exists x_n(\phi_0 \wedge \cdots \wedge \phi_n)\}$ is consistent. We will normally say simply that $T \cup \{\phi_0, \ldots, \phi_k\}$ *is consistent*, or $\{\phi_0 \ldots, \phi_n\}$ is *consistent with* T.

If $\Gamma(x_0, \ldots, x_n)$ is an infinite set of formulas we cannot write down a sentence saying that Γ is realized. At the very least, if Γ is realized on some model of T then every finite $\Gamma_0 \subseteq \Gamma$ is consistent with T. We will also say Γ is consistent with T, or Γ is *finitely satisfiable* in T, to mean that every finite $\Gamma_0 \subseteq \Gamma$ is consistent with T, or equivalently, every finite $\Gamma_0 \subseteq \Gamma$ is realized in some model of T. A simple compactness argument yields the following:

THEOREM 1.3. *T has a model realizing $\Gamma(x_0, \ldots, x_n)$ iff Γ is finitely satisfiable in T.*

PROOF. We only need to show the implication from right to left. First let $\mathcal{L}' = \mathcal{L} \cup \{c_0, \ldots, c_n\}$ where $c_0, \ldots, c_n$ are new constants. Let

$$\Sigma = T \cup \{\phi(c_0, \ldots, c_n)|\ \phi \in \Gamma\}.$$

If $\mathfrak{A}' \models \Sigma$ then $\mathfrak{A}' \restriction \mathcal{L}$ is a model of T realizing Γ. We show Σ has a model by compactness. Let $\Sigma_0 \subseteq \Sigma$ be finite. Then

$$\Sigma_0 \subseteq T \cup \{\phi(c_0, \ldots, c_n)|\ \phi \in \Gamma_0\}$$

for some finite $\Gamma_0 \subseteq \Gamma$. By assumption, T has some model $\mathfrak{A}$ in which Γ_0 is realized, say by $a_0, \ldots, a_n \in A$. Let $\mathfrak{A}'$ be the expansion of $\mathfrak{A}$ to $\mathcal{L}'$ such that $c_i^{\mathfrak{A}'} = a_i$. Then $\mathfrak{A}' \models \Sigma_0$. □

Warning: This result does *not* say that a (fixed) model $\mathfrak{A}$ realizes Γ provided every finite subset of Γ is realized by $\mathfrak{A}$. This is certainly false, as easy examples show.

As a consequence of the Löwenheim-Skolem Theorem applied to the set Σ in the preceding proof we obtain:

COROLLARY 1.4. *If T has an infinite model realizing Γ, then T has a model of cardinality κ realizing Γ for each $\kappa \geq |\mathcal{L}|$.*

As an example of the use of types we have the following:

EXAMPLE 1.1. Let $T = Th((\mathbb{Z}, \leq))$ and let $\Gamma(x, y) = \{\phi_n(x, y) |\ n \in \omega\}$, where $\phi_n(x, y)$ says "$x < y$ and there are at least n elements between x and y." Then Γ is omitted in $(\mathbb{Z}, \leq)$ but every finite $\Gamma_0 \subseteq \Gamma$ is realized in $(\mathbb{Z}, \leq)$, hence Γ is realized in some countable model $\mathfrak{A}$ of T, which must then not be isomorphic to $(\mathbb{Z}, \leq)$.

Let us turn to the question of when a theory T has a model omitting Γ. The way to approach this question is to look at the negation and ask instead when *every* model of T realizes Γ. One will then be led to the following concept:

DEFINITION 1.2. T locally realizes $\Gamma(x_0, \ldots, x_n)$ iff there is some formula $\theta(x_0, \ldots, x_n)$ consistent with T such that

$$T \models \forall x_0 \cdots \forall x_n(\theta \to \phi)$$

for all $\phi \in \Gamma$.

We can easily see:

PROPOSITION 1.5. *Assume T is complete. If T locally realizes Γ then every model of T realizes Γ.*

PROOF. Let $\mathfrak{A} \models T$. Since T is complete and θ is consistent with T, we must have $\mathfrak{A} \models \exists x_0 \cdots \exists x_n \theta$. If $\mathfrak{A}_A \models \theta(\overline{a_0}, \ldots, \overline{a_n})$ then $a_0, \ldots, a_n$ will realize Γ in $\mathfrak{A}$. □

Our goal is to obtain the converse of the preceding proposition. We will not need the assumption that T is complete, but we will need to assume that our language is countable. The fact that we do not know any good criterion for a theory in an uncountable language to have a model omitting a type is one of the primary difficulties in a satisfactory development of the model theory of uncountable languages, and also accounts for many of the problems in dealing with uncountable models even for countable languages.

DEFINITION 1.3. T locally omits $\Gamma(x_0, \ldots, x_n)$ iff T does not locally realize Γ, that is, iff for every formula $\theta(x_0, \ldots, x_n)$ consistent with T there is some $\phi(x_0, \ldots, x_n) \in \Gamma$ such that $(\theta \wedge \neg\phi)$ is consistent with T.

The theorem we are after is the following fundamental result:

THEOREM 1.6. ***(Omitting Types Theorem)*** *Assume $\mathcal{L}$is countable. Assume that T is consistent and that T locally omits $\Gamma(x_0, \ldots, x_n)$. Then T has a countable model which omits Γ.*

Our proof of this result will be via the Henkin method and follow closely the proof of the completeness theorem. That is, we will expand our language by adding new individual constant symbols and expand T to a Henkin Set in the new language. This expansion will be done in such a way that the canonical model of the resulting Henkin set will omit the type Γ.

One point needs to be considered first–the definitions of locally omit and locally realize depend on the language that the formula $\theta(x_0, \ldots, x_n)$ is allowed to come from. Just because T locally omits Γ *in* $\mathcal{L}$ (i.e. considering just θ's in $\mathcal{L}$) is no guarantee that it still does so in a larger language $\mathcal{L}'$ (i.e. allowing θ's in $\mathcal{L}'$). The following lemma says that we can add constants to our language without harming things, and furthermore we can add finitely many axioms (in the new language) to

T and still have Γ be locally omitted. This is the essential fact needed to build up our Henkin Set.

LEMMA 1.7. *Assume that T locally omits $\Gamma(x_0, \ldots, x_n)$. Let $\mathcal{L}' \supseteq \mathcal{L}$ add only individual constant symbols to $\mathcal{L}$. Let $\Sigma' = T \cup \{\psi_1, \ldots, \psi_k\}$ be consistent, where $\psi_i \in Sn_{\mathcal{L}'}$. Then Σ' locally omits Γ (in $\mathcal{L}'$).*

PROOF. Let $\theta'(x_0, \ldots, x_n)$ of $\mathcal{L}'$ be consistent with Σ'. Let $c_0, \ldots, c_m$ list all constants of $\mathcal{L}' - \mathcal{L}$ occurring in any of $\theta', \psi_1, \ldots, psi_k$. Let $y_0, \ldots, y_m$ be variables not occurring in any of $\theta', \psi_1, \ldots, \psi_k$. Thus

$$\chi(x_0, \ldots, x_n, y_0, \ldots, y_m) = (\theta' \wedge \psi_1 \wedge \cdots \wedge \psi_k)^{c_0, \ldots, c_m}_{y_0, \ldots, y_m}$$

is a formula of $\mathcal{L}$, and χ is consistent with T. Finally, let

$$\theta(x_0, \ldots, x_n) = \exists y_0 \cdots \exists y_m \chi.$$

Then θ is also consistent with T, hence there is some $\phi \in \Gamma$ such that $(\theta \wedge \neg\phi)$ is consistent with T. Therefore $(\chi \wedge \neg\phi)$ is consistent with T, hence $\chi(x_0, \ldots, x_n, c_0, \ldots, c_m) \wedge \phi$ is consistent with T, i.e. $(\theta' \wedge \psi_1 \wedge \cdots \wedge \psi_k \wedge \neg\phi)$ is consistent with T, which says precisely that $(\theta' \wedge \neg\phi)$ is consistent with Σ'. $\square$

Having this lemma we can now proceed to prove the Omitting Types Theorem.

PROOF. Let $\mathcal{L}^* = \mathcal{L} \cup C_0$, where C_0 is a set of new individual constant symbols with $|C_0| = \omega$. Then T is still consistent and still locally omits Γ in $\mathcal{L}^*$. For simplicity in notation we assume that Γ has just one free variable, so Γ is $\Gamma(x)$. We will extent T to a Henkin Set Σ^* of $\mathcal{L}^*$ such that the following holds:

(3) for every constant c of $\mathcal{L}^*$ there is some $\phi(x) \in \Gamma$ such that $\neg\phi(c) \in \Sigma^*$.

If $\mathfrak{A}^*$ is the canonical model of a Σ^* satisfying (3), then for every $a \in A^*$, $a = c^{\mathfrak{A}^*}$ for some c, hence $\mathfrak{A}^*_{A^*} \models \neg\phi(\overline{a})$–i.e. $\mathfrak{A}^*$ omits $\Gamma(x)$.

Let $\{\sigma_n | \ n \in \omega\}$ be a listing of all the sentences of $\mathcal{L}^*$, and let $\{c_n | \ n \in \omega\}$ list all constant symbols of $\mathcal{L}^*$. We define recursively a set Σ_n of sentences of $\mathcal{L}^*$ such that $\Sigma_0 = T$ and for each $n \in \omega$ we have:

(0n) $\Sigma_n \subseteq \Sigma_{n+1}$ and Σ_{n+1} is consistent,

($\frac{1}{2}$n) $\Sigma_{n+1} - \Sigma_n$ is finite,

(1n) $\sigma_n \in \Sigma_{n+1}$ or $\neg\sigma_n \in \Sigma_{n+1}$,

(2n) if $\sigma_n \notin \Sigma_{n+1}$ and $\sigma_n = \forall x \phi(x)$ for some $\phi(x)$, then $\neg\phi(c) \in \Sigma_{n+1}$ for some c,

(3n) $\neg\phi(c_n) \in \Sigma_{n+1}$ for some $\phi(x) \in \Gamma$.

We pass from Σ_n to Σ_{n+1} in three steps: first we obtain $\Sigma_{n+\frac{1}{2}}$ so that $\sigma_n \in \Sigma_{n+\frac{1}{2}}$ or $\neg\sigma_n \in \Sigma_{n+\frac{1}{2}}$; then we obtain $\Sigma_{n+\frac{3}{4}}$ such that $\neg\phi(c) \in \Sigma_{n+\frac{3}{4}}$ for some c, provided $\sigma_n = \forall x \phi(x) \notin \Sigma_{n+\frac{1}{2}}$–these are just as in the proof of the Completeness Theorem.

Given $\Sigma_{n+\frac{3}{4}}$–a finite, consistent extension of T–we know by the lemma that $\Sigma_{n+\frac{3}{4}}$ locally omits $\Gamma(x)$, in $\mathcal{L}^*$. Look at the consistent formula $\theta(x) = x \equiv c_n$. There must be some $\phi(x)$ in Γ such that $x \equiv c_n \wedge \neg\phi(x)$ is consistent with $\Sigma_{n+\frac{3}{4}}$. Therefore we define $\Sigma_{n+1} = \Sigma_{n+\frac{3}{4}} \cup \{\neg\phi(c_n)\}$, and this is as desired.

$\Sigma^* = \bigcup_{n \in \omega} \Sigma_n$ is the desired Henkin set, whose existence completes the proof. $\square$

COROLLARY 1.8. *Let T be a complete theory in a countable language. The following are equivalent:*

(i) T has a model omitting $\Gamma(x_1, \ldots, x_n)$,
(ii) T has a countable model omitting Γ,
(iii) T locally omits Γ.

Even if T is not complete, (i) is equivalent to (ii), but (ii)$\Rightarrow$(iii) fails.

The most natural examples of types consistent with a theory T are the sets of formulas satisfied by a fixed tuple of elements in a fixed model of T. These types are also "maximal."

DEFINITION 1.4. (1) Given a model $\mathfrak{A}$ and $a_0, \ldots, a_n \in A$, the (complete) type of $a_0, \ldots, a_n \in \mathfrak{A}$ is

$$tp_{\mathfrak{A}}(a_0, \ldots, a_n) = \{\phi(x_0, \ldots, x_n) |\ \mathfrak{A}_A \models \phi(\overline{a_0}, \ldots, \overline{a_n})\}.$$

(2) A type $\Gamma(x_0, \ldots, x_n)$ is a complete type (of T) iff Γ is realized on some model (of T) and for every $\phi(x_0, \ldots, x_n)$ either $\phi \in \Gamma$ or $\neg\phi \in \Gamma$.

LEMMA 1.9. *$\Gamma(x_0, \ldots, x_n)$ is a complete type of T iff $\Gamma = tp_{\mathfrak{A}}(a_0, \ldots, a_n)$ for some $\mathfrak{A} \models T$ and $a_0, \ldots, a_n \in A$.*

A consequence of the realizing types theorem is the following:

COROLLARY 1.10. *If $\Gamma(x_0, \ldots, x_n)$ is consistent with T then there is some complete type $\Gamma^*(x_0, \ldots, x_n)$ of T with $\Gamma \subseteq \Gamma^*$.*

We will sometimes say that $\Gamma(x_0, \ldots, x_n)$ generates a complete type of T to mean that Γ is contained in *exactly one* complete type $\Gamma^*(x_0, \ldots, x_n)$ of T–i.e. whenever $\mathfrak{A}, \mathfrak{B} \models T$, $a_0, \ldots, a_n$ realize Γ in $\mathfrak{A}$, and $b_0, \ldots, b_n$ realize Γ in $\mathfrak{B}$ then $tp_{\mathfrak{A}}(a_0, \ldots, a_n) = tp_{\mathfrak{B}}(b_0, \ldots, b_n)$.

Just as the sets of sentences one normally writes down do not usually axiomatize complete theories, so the types one writes down do not usually generate complete types. For example, let $T = Th((\omega, \leq, +, \cdot, 0, 1))$ and let $\Gamma_0(x) = \{\overline{n} < x |\ n \in \omega\}$. Then Γ_0 is realized on all models of T which are not isomorphic to $\mathfrak{N} = (\omega, \leq, +, \cdot, 0, 1)$. The elements of any $\mathfrak{A} \models T$ which realize $\Gamma_0(x)$ are "infinite", i.e. have infinitely many elements of A less than them. But $\Gamma_0(x)$ does not generate a complete type since it does not decide for example whether or not "$\overline{p} \mid x$" should hold (for primes $p \in \omega$). In fact *any* way of adding "$\overline{p} \mid x$" or "$\neg\overline{p} \mid x$" for any primes $p \in \omega$ leads to a consistent type extending $\Gamma_0(x)$.

Let $\mathbb{P}$ be the set of prime numbers in ω. Let $X \subseteq \mathbb{P}$. We define the type

$$\Gamma_X(x) = \Gamma_0(x) \cup \{\overline{p} \mid x |\ p \in X\} \cup \{\overline{p} \not| x |\ p \in \mathbb{P} - X\}.$$

Then each Γ_X is a consistent type with T. Although Γ_X does not generate a complete type, we can say that the same complete type cannot contain both Γ_X and Γ_Y, for instance $X \neq Y$, both subsets of $\mathbb{P}$. Thus–*there are 2^ω different complete types* in x consistent with T. An element of a model realizes *exactly one* complete type, hence a countable model can realize just countably many complete types. Since every type consistent with T is realized on *some* countable model of T we can conclude that *T has (at least) 2^ω non-isomorphic countable models.* In fact an elementary lemma shows that there are *at most* 2^ω non-isomorphic countable models for any countable language $\mathcal{L}$–just count the number of different models with universe $= \omega$. We thus have:

PROPOSITION 1.11. *$T = Th(\mathfrak{N})$ has precisely 2^ω non-isomorphic countable models.*

In applications (as in the next chapter) we frequently need to realize or omit an infinite number of types simultaneously. The relevant results are as follows:

THEOREM 1.12. *Let T be a complete theory (in a countable $\mathcal{L}$). Assume each type $\Gamma_k(x_0, \ldots, x_{n_k})$ is consistent with T, $k \in \omega$. Then T has a (countable) model realizing* every *Γ_k, $k \in \omega$.*

THEOREM 1.13. ***(Extended Omitting Types Theorem)*** *Let T be a theory in a countable $\mathcal{L}$. Assume that T locally omits each type $\Gamma_k(x_0, \ldots, x_{n_k})$. Then T has a countable model which omits* every *Γ_k, $k \in \omega$.*

2. Elementary Extensions and Chains

To discuss how two different models of a theory "compare" on frequently wants to talk about one model being included in another. We have two notions of inclusion, one of which is the straightforward analogue of a subalgebra (as subgroup).

DEFINITION 2.1. $\mathfrak{A}$ is a submodel of $\mathfrak{B}$ (or $\mathfrak{B}$ is an extension of $\mathfrak{A}$), written $\mathfrak{A} \subseteq \mathfrak{B}$, iff $A \subseteq B$ and the interpretations in $\mathfrak{A}$ of the nonlogical symbols are just the restrictions to A of their interpretations in $\mathfrak{B}$, that is:
$c^{\mathfrak{A}} = c^{\mathfrak{B}}$ for any constant of $\mathcal{L}$;
$F^{\mathfrak{A}}(a_1, \ldots, a_m) = F^{\mathfrak{B}}(a_1, \ldots, a_m)$ for all $a_1, \ldots, a_m \in A$ and every function symbol F of $\mathcal{L}$;
$R^{\mathfrak{A}}(a_1, \ldots, a_m)$ holds iff $R^{\mathfrak{B}}(a_1, \ldots, a_m)$ holds, for all $a_1, \ldots, a_m \in A$ and every relation symbol R of $\mathcal{L}$.

A submodel of $\mathfrak{B}$ is uniquely determined by its universe, and there is a simple test for checking whether a subset of B is the universe of a submodel.

LEMMA 2.1. *(1) If $\mathfrak{A}_1, \mathfrak{A}_2 \subseteq \mathfrak{B}$, $A_1 = A_2$ then $\mathfrak{A}_1 = \mathfrak{A}_2$. (2) $X \subseteq B$ is the universe of a submodel of $\mathfrak{B}$ iff $X \neq \emptyset$, $c^{\mathfrak{B}} \in X$ for all constants c of $\mathcal{L}$, and $F^{\mathfrak{B}}(a_1, \ldots, a_m) \in X$ for all $a_1, \ldots, a_m \in X$ and every function symbol F of $\mathcal{L}$.*

EXAMPLE 2.1. Here are examples of chains of submodels:

(1) $(\{0,3\}, \{(0,0), (0,3), (3,3)\}) \subseteq (\omega, \leq) \subseteq (\mathbb{Z}, \leq) \subseteq (\mathbb{Q}, \leq) \subseteq (\mathbb{R}, \leq)$
(2) $(\omega, +) \subseteq (\mathbb{Z}, +)$

Clearly as the above examples can be used to show, $\mathfrak{A} \subseteq \mathfrak{B}$, $\mathfrak{B} \models T$ need not imply $\mathfrak{A} \models T$. A more serious problem from our point of view is shown by the following example:

EXAMPLE 2.2. Let $\mathfrak{B} = (\omega, \leq), \mathfrak{A} = (A, \leq)$ where $A = \{2k+1 |\ k \in \omega\}$. Then $\mathfrak{A} \subseteq \mathfrak{B}$ by definition. Further $\mathfrak{A} \equiv \mathfrak{B}$ (since $\mathfrak{A} \cong \mathfrak{B}$). If we let $\phi_0(x)$ be $\forall y(x \leq y)$, which defines the first element in the ordering, then $\mathfrak{B}_B \models \phi_0(\overline{0}) \wedge \neg\phi_0(\overline{1})$. Now $0 \notin A$ and in fact $\mathfrak{A}_A \models \phi_0(\overline{1})$. That is: the element 1 which belongs to both universes, satisfies a formula in one model which it does not satisfy in the other model. This is an undersirable property, and our second concept of extension is designed precisely to avoid it.

DEFINITION 2.2. $\mathfrak{A}$ is an elementary submodel of $\mathfrak{B}$ ($\mathfrak{B}$ is an elementary extension of $\mathfrak{A}$), written $\mathfrak{A} \prec \mathfrak{B}$, iff $\mathfrak{A} \subseteq \mathfrak{B}$ and for every formula $\phi(x_0, \ldots, x_n)$ of $\mathcal{L}$and every $a_0, \ldots, a_n \in A$ we have $\mathfrak{A}_A \models \phi(\overline{a_0}, \ldots, \overline{a_n})$ iff $\mathfrak{B}_A \models \phi(\overline{a_0}, \ldots, \overline{a_n})$.

Clearly, $\mathfrak{A} \prec \mathfrak{B} \Rightarrow \mathfrak{A} \equiv \mathfrak{B}$ but the above example shows that $\mathfrak{A} \equiv \mathfrak{B}$, $\mathfrak{A} \subseteq \mathfrak{B}$ does not imply that $\mathfrak{A} \prec \mathfrak{B}$.

Combining submodel with isomorphism yields two hybrid notions:

DEFINITION 2.3. (1)$\mathfrak{A} \cong\subseteq \mathfrak{B}$ ($\mathfrak{A}$ is embeddable in $\mathfrak{B}$) means $\mathfrak{A} \cong \mathfrak{A}' \subseteq \mathfrak{B}$ for some $\mathfrak{A}'$. (2) $\mathfrak{A} \cong\prec \mathfrak{B}$ ($\mathfrak{A}$ is elementarily embeddable in $\mathfrak{B}$) means $\mathfrak{A} \cong \mathfrak{A}' \prec \mathfrak{B}$ for some $\mathfrak{A}'$.

The next lemma enables us to "reverse the order" of isomorphism and inclusion in the previous definitions. It is important because in trying to show that $\mathfrak{A}$ has an elementary extension with a certain property, our methods yield a model into which $\mathfrak{A}$ can be elementarily *embedded* first of all.

LEMMA 2.2. *(1) $\mathfrak{A} \cong\subseteq \mathfrak{B}$ iff $\mathfrak{A} \subseteq \mathfrak{B}' \cong \mathfrak{B}$ for some $\mathfrak{B}'$. (2) $\mathfrak{A} \cong\prec \mathfrak{B}$ iff $\mathfrak{A} \prec \mathfrak{B}' \cong \mathfrak{B}$ for some $\mathfrak{B}'$.*

PROOF. From right to left is clear, so assume $h : \mathfrak{A} \cong \mathfrak{A}'$ where $\mathfrak{A}' \subseteq \mathfrak{B}$. Let B_0 be such that $A \cap B_0 = \emptyset$ and $|B_0| = |B - A'|$, say $h_0 : B_0 \cong B - A'$. Let $B' = A \cup B_0$ and interpret the symbols of $\mathcal{L}$on B' so that $g = h \cup h_0$ is an isomorphism of $\mathfrak{B}'$ onto $\mathfrak{B}$. Since h is an isomorphism of $\mathfrak{A}$ onto $\mathfrak{A}'$, we will automatically have $\mathfrak{A} \subseteq \mathfrak{B}'$ ($\mathfrak{A} \prec \mathfrak{B}'$ for (2)). $\square$

Extensions and elementary extensions of $\mathfrak{A}$ can be characterized by sets of sentences of $\mathcal{L}(A)$.

DEFINITION 2.4. (1) A basic formula is a formula which is either atomic or the negation of an atomic formula. (2) The (basic) diagram of $\mathfrak{A}$ is the set

$$\Delta_{\mathfrak{A}} = \{\beta|\ \beta \text{ is a basic sentence of } \mathcal{L}(A) \text{ s.t. } \mathfrak{A}_A \models \beta\}.$$

LEMMA 2.3. *(1) $\mathfrak{A} \subseteq \mathfrak{B}$ iff $\mathfrak{B}_A \models \Delta_{\mathfrak{A}}$. (2) $\mathfrak{A} \prec \mathfrak{B}$ iff $\mathfrak{B}_A \equiv \mathfrak{A}_A$, i.e. $\mathfrak{B}_A \models Th(\mathfrak{A}_A)$.*

The lemma is immediate from the definitions. $Th(\mathfrak{A}_A)$ is sometimes called the *elementary diagram* of $\mathfrak{A}$.

Of course, $\Delta_{\mathfrak{A}}$ and $Th(\mathfrak{A}_A)$ can have models in which the constant $\overline{a}$ is not interpreted as the element a of A. Tese other models turn out to be precisely the models into which $\mathfrak{A}$ can be embedded or elementarily embedded. In the first case this is quite clear.

LEMMA 2.4. $\mathfrak{B}$ *(an $\mathcal{L}$-structure) can be expanded to a model $\mathfrak{B}^*$ of $\Delta_{\mathfrak{A}}$ iff $\mathfrak{A}$ is embeddable in $\mathfrak{B}$.*

PROOF. If $h : \mathfrak{A} \cong \mathfrak{A}'$ where $\mathfrak{A}' \subseteq \mathfrak{B}$ then defining $\mathfrak{B}^*$ for $\mathcal{L}(A)$ by setting $\overline{a}^{\mathfrak{B}^*} = h(a)$ yields a model of $\Delta_{\mathfrak{A}}$. Conversely, if $\mathfrak{B}^* \models \Delta_{\mathfrak{A}}$ then defining $h(a) = \overline{a}^{\mathfrak{B}^*}$ yields an isomorphism of $\mathfrak{A}$ onto the submodel $\mathfrak{A}' \subseteq \mathfrak{B}$ whose universe is $A' = \{\overline{a}^{\mathfrak{B}^*}\}$. $\square$

The corresponding result for elementary embedding is less obvious, because the requirement that $\mathfrak{A}' \prec \mathfrak{B}$ involves satisfiability *in* $\mathfrak{A}'$ rather than just $\mathfrak{B}$. It is convenient to have a lemma which tells us when a submodel of $\mathfrak{B}$ is an elementary submodel which refers only to satisfiability in $\mathfrak{B}$.

LEMMA 2.5. *Let $\mathfrak{A} \subseteq \mathfrak{B}$. Then $\mathfrak{A} \prec \mathfrak{B}$ iff for every $\phi(x_1, \ldots, x_n, y)$ of $\mathcal{L}$and all $a_1, \ldots, a_n \in A$, if $\mathfrak{B}_A \models \exists y \phi(\overline{a_1}, \ldots, \overline{a_n}, y)$ then there is some $b \in A$ such that $\mathfrak{B}_A \models \phi(\overline{a_1}, \ldots, \overline{a_n}, \overline{b})$.*

PROOF. We show by induction on formulas $\psi(x_1, \ldots, x_n)$ of $\mathcal{L}$that for all $a_1, \ldots, a_n \in A$ $\mathfrak{A}_A \models \psi(\overline{a_1}, \ldots, \overline{a_n})$ iff $\mathfrak{B}_A \models \psi(\overline{a_1}, \ldots, \overline{a_n})$. All steps are clear except for the quantifiers. So assume $\phi(x_1, \ldots, x_n, y)$ satisfies the inductive hypothesis and consider $\psi(x_1, \ldots, x_n) = \exists y \phi$. If $\mathfrak{A}_A \models \psi(\overline{a_1}, \ldots, \overline{a_n})$ then $\mathfrak{A}_A \models \phi(\overline{a_1}, \ldots, \overline{a_n}, \overline{b})$ for some $b \in A$, hence $\mathfrak{B}_A \models \phi(\overline{a_1}, \ldots, \overline{a_n}, \overline{b})$ and so $\mathfrak{B}_A \models \psi(\overline{a_1}, \ldots, \overline{a_n})$. Conversely, if $\mathfrak{A}_A \models \psi(\overline{a_1}, \ldots, \overline{a_n})$ then the hypothesis of the lemma yields $\mathfrak{B}_A \models \phi(\overline{a_1}, \ldots, \overline{a_n}, \overline{b})$ for some $b \in A$, so $\mathfrak{A}_A \models \phi(\overline{a_1}, \ldots, \overline{a_n}, \overline{b})$ and $\mathfrak{A}_A \models \psi(\overline{a_1}, \ldots, \overline{a_n})$. $\square$

With the lemma it is now easy to prove the characterization of the models into which $\mathfrak{A}$ is elementarily embeddable.

PROPOSITION 2.6. *$\mathfrak{A}$ is elementarily embeddable in $\mathfrak{B}$ iff $\mathfrak{B}$ can be expanded to a model $\mathfrak{B}^*$ of $Th(\mathfrak{A}_A)$.*

PROOF. We need just to prove the direction from right to left. If $\mathfrak{B}^* \models Th(\mathfrak{A}_A)$ then we define, as before, $\mathfrak{A}'$ to be the submodel of $\mathfrak{B}$ with universe $A' = \{\overline{a}^{\mathfrak{B}^*} | \ a \in A\}$. Then $\mathfrak{A} \cong \mathfrak{A}'$, so it remains to show $\mathfrak{A}' \prec \mathfrak{B}$, for which we use the lemma. Given $\phi(x_1, \ldots, x_n, y)$ of $\mathcal{L}$and $a'_1, \ldots, a'_n \in A'$, suppose $\mathfrak{B}_{A'} \models \exists y \phi(x_1, \ldots, x_n, y)$. Well, $a'_i = \overline{a}_i^{\mathfrak{B}^*}$ for $a_i \in A$, so $\mathfrak{B}^* \models \exists y \phi(\overline{a_1}, \ldots, \overline{a_n}, y)$. Therefore $\mathfrak{A}_A \models \exists y \phi(\overline{a_1}, \ldots, \overline{a_n}, y)$, so $\mathfrak{A}_A \models \phi(\overline{a_1}, \ldots, \overline{a_n}, \overline{b})$ for some $b \in A$. Therefore $\mathfrak{B}^* \models \phi(\overline{a_1}, \ldots, \overline{a_n}, \overline{b})$, i.e. $\mathfrak{B}_{A'} \models \phi(\overline{a'_1}, \ldots, \overline{a'_n}, \overline{b'})$ where $b' = \overline{b}^{\mathfrak{B}^*} \in A'$. $\square$

In this proof, the models $\mathfrak{B}^*$ and $\mathfrak{B}_{A'}$ are different, since they are for different languages, but the correspondence $\overline{a} \leftrightarrow \overline{(\overline{a}^{\mathfrak{B}^*})}$ sets up a "translation" between them.

Intuitively, an elementary extension of $\mathfrak{A}$ is just like $\mathfrak{A}$ but *larger*, that is, has more elements and (perhaps) more *types* of elements. The following is easy from the definition.

LEMMA 2.7. *Assume $\mathfrak{A} \prec \mathfrak{B}$. Then $\mathfrak{B}$ realizes every type realized in $\mathfrak{A}$.*

PROOF. If $a_0, \ldots, a_n \in A$ realize $\Gamma(x_0, \ldots, x_n)$ in $\mathfrak{A}$ then they also realize $\Gamma(x_0, \ldots, x_n)$ in $\mathfrak{B}$ since $\mathfrak{A} \prec \mathfrak{B}$. $\square$

More importantly, we can now realize types in elementary extensions of given models.

THEOREM 2.8. *Assume $\Gamma(x_0, \ldots, x_n)$ is consistent with $Th(\mathfrak{A})$. Then there is some $\mathfrak{B}$ such that $\mathfrak{A} \prec \mathfrak{B}$ and $\mathfrak{B}$ realizes Γ.*

PROOF. Let $\mathcal{L}^* = \mathcal{L}(A) \cup \{c_0, \ldots, c_n\}$, where $c_0, \ldots, c_n$ are constants not in $\mathcal{L}(A)$. Let

$$\Sigma = Th(\mathfrak{A}_A) \cup \{\phi(c_0, \ldots, c_n) | \ \phi \in \Gamma\}.$$

If $\mathfrak{B}^* \models \Sigma$ then $\mathfrak{A}$ is elementarily embeddable in $\mathfrak{B}' = \mathfrak{B}^* \restriction \mathcal{L}$ which realizes Γ. By Lemma 4.3.2, we obtain $\mathfrak{B}$ such that $\mathfrak{A} \prec \mathfrak{B}$ and $\mathfrak{B} \cong \mathfrak{B}'$, so $\mathfrak{B}$ also realizes Γ. To show Σ has a model we use compactness. If $\Sigma_0 \subseteq \Sigma$ is finite, then

$$\Sigma_0 \subseteq Th(\mathfrak{A}_A) \cup \{\phi(c_0, \ldots, c_n) | \ \phi \in \Gamma_0\}$$

for some finite $\Gamma_0 \subseteq \Gamma$. The hypothesis yields $a_0^*, \ldots, a_n^* \in A$ such that $(\mathfrak{A}_A, a_0^*, \ldots, a_n^*) \models \Sigma_0$. $\square$

Elementary embeddings give us a precise way to say that $(\omega, \leq)$ is the *smallest* model of its complete theory.

PROPOSITION 2.9. *Let $T = Th((\omega, \leq))$ and let $\mathfrak{B} \models T$. Then $(\omega, \leq) \cong\prec \mathfrak{B}$.*

PROOF. For each $n \in \omega$ there is a formula $\phi_n(x)$ which defines $\{n\}$ in $(\omega, \leq)$. Therefore $T \models \exists x\phi_n(x)$ for each $n \in \omega$, and since $\mathfrak{B} \models T$ we must have $\mathfrak{B}_B \models \phi_n(\overline{b_n})$ for $b_n \in B$, for each $n \in \omega$. Let $\mathfrak{B}^* = (\mathfrak{B}, b_n)_{n\in\omega}$ be the expansion of $\mathfrak{B}$ to $\mathcal{L}(\omega)$ in which $\overline{n}$ is interpreted by b_n. We claim that $\mathfrak{B}^* \models Th((\omega, \leq)_\omega)$. Let $\psi(x_0, \ldots, x_k)$ of $\mathcal{L}$be given and suppose $(\omega, \leq)_\omega \models \psi(\overline{n_0}, \ldots, \overline{n_k})$. Then, since ϕ_n's *define* $\{n\}$, we actually have

$$(\omega, \leq) \models \forall x_0 \cdots \forall x_k[\phi_{n_0}(x_0) \wedge \cdots \wedge \phi_{n_k}(x_k) \rightarrow \psi].$$

Therefore this sentence is also true in $\mathfrak{B}$, whence it follows that $\mathfrak{B}^* \models \psi(\overline{n_0}, \ldots, \overline{n_k})$. Proposition 4.3.6 thus yields the conclusion. $\square$

Also, of course, it follows from this that a type realized on $(\omega, \leq)$ is realized in *every* $\mathfrak{B} \equiv (\omega, \leq)$. Further, if $\mathfrak{B} \equiv (\omega, \leq)$ but $\mathfrak{B} \ncong (\omega, \leq)$ then $\mathfrak{B}$ must realize the type $\{\neg\phi_n(x) | \ n \in \omega\}$, which is omitted in $(\omega, \leq)$. Thus, $(\omega, \leq)$ is the *only* model of T which realizes *only* the types which T locally realizes.

We shall also consider *chains* under inclusion.

DEFINITION 2.5. Let I be linearly ordered by $\leq$. A family $\{\mathfrak{A}_i\}_{i\in I}$ is a chain iff $\mathfrak{A}_i \subseteq \mathfrak{A}_j$ whenever $i, j \in I$, $i \leq j$. $\{\mathfrak{A}_i\}_{i\in I}$ is an elementary chain iff $\mathfrak{A}_i \prec \mathfrak{A}_j$ whenever $i, j \in I$, $i \leq j$.

DEFINITION 2.6. Let $\{\mathfrak{A}_i\}_{i\in I}$ be a chain. The union of the chain $\{\mathfrak{A}_i\}_{i\in I}$ is the model $\mathfrak{B}$ defined by $B = \bigcup_{i\in I} A_i$, $c^{\mathfrak{B}} = c^{\mathfrak{A}_i}$, $F^{\mathfrak{B}}(a_1, \ldots, a_m) = F^{\mathfrak{A}_i}(a_1, \ldots, a_m)$ if $a_1, \ldots, a_m \in A_i$, and $R^{\mathfrak{B}}(a_1, \ldots, a_m)$ holds iff $R^{\mathfrak{A}_i}(a_1, \ldots, a_m)$ holds if $a_1, \ldots, a_m \in A_i$. We write $\mathfrak{B} = \bigcup_{i\in I} \mathfrak{A}_i$.

LEMMA 2.10. *The union $\mathfrak{B}$ of the chain $\{\mathfrak{A}_i\}_{i\in I}$ is well defined, and $\mathfrak{A}_i \subseteq \mathfrak{B}$ for every $i \in I$.*

THEOREM 2.11. *Let $\{\mathfrak{A}_i\}_{i\in I}$ be an elementary chain and $\mathfrak{B} = \bigcup_{i\in I} \mathfrak{A}_i$. Then $\mathfrak{A}_i \prec \mathfrak{B}$ for every $i \in I$.*

PROOF. We prove, by induction on $\phi(x_0, \ldots, x_n)$, that for every $i \in I$ and for every $a_0, \ldots, a_n \in A_i$ we have $(\mathfrak{A}_i)_{A_i} \models \phi(\overline{a_0}, \ldots, \overline{a_n})$ iff $\mathfrak{B}_{\mathfrak{A}_i} \models \phi(\overline{a_0}, \ldots, \overline{a_n})$. The only case of difficulty occurs when $\phi(\vec{x}) = \exists y\psi(\vec{x}, y)$. Let $a_0, \ldots, a_n \in A_i$ and suppose $\mathfrak{B}_{A_i} \models \phi(\overline{a_0}, \ldots, \overline{a_n})$. Then $\mathfrak{B}_B \models \psi(\overline{a_0}, \ldots, \overline{a_n}, \overline{b})$ for some $b \in B$– but then $b \in A_j$ for some $j \in I$, $i \leq j$. So by inductive hypothesis we know $(\mathfrak{A}_j)_{A_j} \models \phi(\overline{a_0}, \ldots, \overline{a_n})$ as desired. $\square$

We will use this theorem to construct a model in an infinite number of steps, each "approximating" the result.

3. The Back-and-Forth Method

If two models *are* isomorphic, how might one prove this? In this section, we explain a general method which we will use, both here and in the next chapter, to show models are isomorphic. Although this method has extensions to uncountable models, we will restrict our attention to just countable models.

To motivate the method, consider two countable models $\mathfrak{A}, \mathfrak{B}$. Then we can list the elements of both their universes, say $A = \{a_n | \ n \in \omega\}$ and $B = \{b_n | \ n \in \omega\}$. To define an isomorphism, say h, of $\mathfrak{A}$ onto $\mathfrak{B}$ we need to specify for each $n \in \omega$ an element $d_n \in B$ such that $h(a_n) = d_n$. We must also specify for each $n \in \omega$ an element $c_n \in A$ such that $h(c_n) = b_n$. [Of course, each $d_n$is some $b_{n'}$ and each c_n is some $a_{n''}$ but it is notationally easier to introduce new designations c_n, d_n.] We want to do this in such a way that the resulting h is an isomorphism, that is for every *atomic formula* $\phi(x_0, \ldots, x_n)$ and $a'_0, \ldots, a'_n \in A$

$$\mathfrak{A}_A \models \phi(\overline{a'_0}, \ldots, \overline{a'_n}) \text{ iff } \mathfrak{B}_B \models \phi(\overline{h(a'_0)}, \ldots, \overline{h(a_n)'}).$$

Our procedure is to do this recursively–that is, knowing $d_0, ldots, d_{n-1}$ and $c_0, \ldots, c_{n-1}$ we will define d_n and c_n. The problem is what *inductive hypothesis* on $d_0, \ldots, d_{n-1}$ and $c_0, \ldots, c_{n-1}$ will enable us to choose appropriate d_n and c_n? At the very least we will need to know that the choices already made will make the piece of h already defined behave like an isomorphism–i.e. for any atomic $\alpha(x_0, \ldots, x_k)$ and $a'_0, \ldots, a'_k \in \{a_0, \ldots, a_{n-1}, c_0, \ldots, c_{n-1}\}$ we have $\mathfrak{A}_A \models \alpha(\overline{a'_0}, \ldots, \overline{a'_k})$ iff $\mathfrak{B}_B \models \alpha(\overline{h(a'_0)}, \ldots, \overline{h(a'_k)})$. This may not be enough, however. In general we need to have some notion of "similarity" between tuples from A and tuples from B which *does* guarantee that we can continue to build up our isomorphism.

THEOREM 3.1. *Let $\mathfrak{A}$ and $\mathfrak{B}$ be countable. Assume there is some relation $\sim$ between n-tuples of elements of A and n-tuples of elements of B, for all n, such that: $\forall a'_0, \ldots, a'_{n-1} \in A$ $\forall b'_0, \ldots, b'_{n-1} \in B$*

1) $(a'_0, \ldots, a'_{n-1}) \sim (b'_0, \ldots, b'_{n-1}) \Rightarrow$ for all atomic $\alpha(x_0, \ldots, x_{k-1})$ of $\mathcal{L}$, $\mathfrak{A}_A \models \alpha(\overline{a'_0}, \ldots, \overline{a'_{n-1}})$ iff $\mathfrak{B}_B \models \alpha(\overline{b'_0}, \ldots, \overline{b'_{n-1}})$.

2) $(a'_0, \ldots, a'_{n-1}) \sim (b'_0, \ldots, b'_{n-1}) \Rightarrow \forall a'_n \in A \exists b'_n \in B$ such that $(a'_0, \ldots, a'_n) \sim (b'_0, \ldots, b'_n)$ and $\forall b'_n \in B \exists a'_n \in A$ such that $(a'_0, \ldots, a'_n) \sim (b'_0, \ldots, b'_n)$.

3) $\emptyset \sim \emptyset$.

Then $\mathfrak{A} \cong \mathfrak{B}$, and in fact whenever $(a'_0, \ldots, a'_{n-1}) \sim (b'_0, \ldots, b'_{n-1})$ there is some isomorphism h of $\mathfrak{A}$ onto $\mathfrak{B}$ with $h(a'_i) = b'_i$ for all $i = 0, \ldots, n-1$.

PROOF. Let $A = \{a_n | \ n \in \omega\}$, $B = \{b_n | \ n \in \omega\}$. We define, by recursion on n, elements $d_n \in B$, $c_n \in A$ such that for each $n \in \omega$ we have

$$(a_0, c_0, a_1, c_1, \ldots, a_n, c_n) \sim (d_0, b_0, d_1, b_1, \ldots, d_n, c_n).$$

This is clear, using condition 3) to start and condition 2) to continue. By condition 1) the resulting map h is an isomorphism. Given to begin with that $(a'_0, \ldots, a'_{n-1}) \sim (b'_0, \ldots, b'_{n-1})$ we could have started with that rather than $\emptyset \sim \emptyset$ to obtain an isomorphism taking a'_i to b'_i for each $i = 0, \ldots, n-1$. $\square$

The trick in using this result is in coming up with a relation $\sim$ of similarity which has properties 1)-3).

EXAMPLE 3.1. $\mathcal{L}^{nl} = \{\leq\}$, $\Sigma =$ the axioms for dense linear order without endpoints. Let $\mathfrak{A}, \mathfrak{B} \models \Sigma$. Define $(a_0, \ldots, a_n) \sim (b_0, \ldots, b_n)$ iff $(a_i \leq^{\mathfrak{A}} a_j \Leftrightarrow b_i \leq^{\mathfrak{B}} b_j$ all $i, j = 0, \ldots, n)$. Then 1)-3) of the theorem are satisfied, therefore $\mathfrak{A} \cong \mathfrak{B}$ provided $|A| = \omega = |B|$.

(Even in this example, one can find $\mathfrak{A}, \mathfrak{B} \models \Sigma$ such that $|A| = |B| \neq \omega$ but $\mathfrak{A} \not\cong \mathfrak{B}$.) Note that this example shows that $T = Cn(\Sigma)$, "the theory of dense linear order without endpoints," is ω-categorical, hence complete by the Łoš-Vaught test. It follows, for example, that $(\mathbb{Q}, \leq) \equiv (\mathbb{R}, \leq)$.

4. Exercises

(1) Let T be a complete theory of $\mathcal{L}$, and let $\Phi(x)$ and $\Psi(x)$ be complete $\mathcal{L}$-types consistent with T. Assume that every element of every model of T realizes either Φ or Ψ. Prove that both Φ and Ψ contain complete formulas.

(2) Let $\mathcal{L}$ be a countable language containing (at least) the binary relation symbol E. Let T be a consistent theory of $\mathcal{L}$ such that for every $\mathfrak{A} \models T$, $E^{\mathfrak{A}}$ is an equivalence relation on A which has at least one infinite equivalence class. Prove that there is some $\varphi(x) \in \mathrm{Fm}_{\mathcal{L}}$, consistent with T, such that for every $\mathfrak{A} \models T$ every $a \in A$ which satisfies φ belongs to an infinite $E^{\mathfrak{A}}$-equivalence class.

(3) Prove or disprove:
   (a) $(\mathbb{Q}, <)$ has a proper elementary submodel.
   (b) $(\mathbb{Z}, <)$ has a proper elementary submodel.

(4) (a) Let T be a complete theory of $\mathcal{L}$, and let $\Phi(x)$ and $\Psi(y)$ be types each of which is realized on some model of T. Prove that T has some models realizing both Φ and Ψ.
   (b) Give an example to show that the result in part (a) can fail if the theory T is not complete.

(5) Let $\mathfrak{A} = (\mathbb{Q}, +, \cdot, 0, 1)$. Prove that $T = \mathrm{Th}(\mathfrak{A})$ is not ω-categorical.

(6) Let $\Phi(x)$ be a type consistent with the complete theory T but which is realized by at most one element in every model of T. Prove that there is some formula $\psi(x)$ consistent with T such that

$$T \models \forall x(\psi(x) \to \varphi(x)) \text{ for all } \varphi \in \Phi.$$

(7) Let T be a complete theory of $\mathcal{L}$, and let $\mathfrak{A}$ and $\mathfrak{B}$ be models of T. Prove that there is some $\mathfrak{C} \models T$ such that both $\mathfrak{A}$ and $\mathfrak{B}$ can be elementarily embedded in $\mathfrak{C}$.

(8) (a) Assume that $\mathfrak{A} \prec \mathfrak{B}$ and $\mathfrak{A} \neq \mathfrak{B}$. Prove that there is no formula $\varphi(x)$ of $\mathcal{L}$ such that $\varphi^{\mathfrak{B}} = A$.
   (b) Give an example of $\mathfrak{A}, \mathfrak{B}$ and $\varphi(x)$ where $\mathfrak{A} \subseteq \mathfrak{B}, \mathfrak{A} \neq \mathfrak{B}$, but $\varphi^{\mathfrak{B}} = A$.

(9) Let $T = \mathrm{Th}((\mathbb{Z}, <))$. Let $\psi(x, y) = (x < y) \wedge \neg\exists z(x < z \wedge z < y)$. Prove that $\psi(x, y)$ is a complete formula with respect to T.

CHAPTER 5

Countable Models of Complete Theories

0. Introduction

In this chapter we give detailed attention to the question of what the collection of countable models of a complete theory T can look like. This includes consideration of *how many* countable models (up to isomorphism) T can have, and also how the models are related to one another, particularly by elementary embedding. We are also interested in how the properties of the collection of countable models of T are related to syntactical properties of T. Some particular questions we will consider (and answer) are the following:

(1) What cardinals κ can be the number of non-isomorphic countable models of some T? We have seen examples where $\kappa = 1, \omega, 2^\omega$, and we know $\kappa \le 2^\omega$, so this comes down to: (a) there are complete theories T such that T has exactly n non-isomorphic countable models, for $2 \le n \in \omega$? (b) if $\omega < \kappa < 2^\omega$ is there a T with exactly κ non-isomorphic countable models? The question in (b) is complicated by the fact that even the existence of such a cardinal κ is independent of the axioms of set theory. Although we will discuss this later we will not be able to solve this question.

(2) Characterize (syntactically) the theories T with exactly one countable model, i.e., which are ω-categorical.

(3) Does $\mathfrak{A} \cong\prec \mathfrak{B}$ and $\mathfrak{B} \cong\prec \mathfrak{A}$ imply that $\mathfrak{A} \cong \mathfrak{B}$, for countable $\mathfrak{A}, \mathfrak{B}$?

(4) If $\mathfrak{A}, \mathfrak{B}$ are both countable models of T must we have either $\mathfrak{A} \cong\prec \mathfrak{B}$ or $\mathfrak{B} \cong\prec \mathfrak{A}$?

Positive answers to both (3) and (4) would say that $\cong\prec$ is a linear order on isomorphism types of countable models of T.

(5) Does every T have a *smallest* model, i.e. a model $\mathfrak{A}$ such that $\mathfrak{A} \cong\prec \mathfrak{B}$ for every models $\mathfrak{B}$ of T? Must such an $\mathfrak{A}$ be unique? (Recall that previously we showed that $(\omega, \le)$ was such a model of $T = Th((\omega, \le))$).

(6) Does every T have a *largest* countable model, $\mathfrak{B}$, i.e., a countable model $\mathfrak{B}$ such that $\mathfrak{A} \cong\prec \mathfrak{B}$ for every countable model $\mathfrak{A}$ of T? Must such a $\mathfrak{B}$ be unique (up to isomorphism)?

In each of (5), (6) if the answer is no, one would want to characterize the theories with such models.

(7) If $\mathfrak{A} \equiv \mathfrak{B}$ are countable and realize the same types must $\mathfrak{A} \cong \mathfrak{B}$?

Similarly, we know that $\mathfrak{A} \cong\prec \mathfrak{B}$ implies that $\mathfrak{B}$ realizes every type that $\mathfrak{A}$ does.

(8) Does the converse hold, for countable $\mathfrak{A}, \mathfrak{B}$?

1. Prime Models

Throughout, T is a complete theory in a countable language.

DEFINITION 1.1. A model $\mathfrak{A}$ of T is *prime* iff for every $\mathfrak{B} \models T$ $\mathfrak{A} \cong\prec \mathfrak{B}$.

In this section we will characterize prime models, characterize the complete theories which have prime models, and show the uniqueness of the prime model of a theory.

Suppose that $\mathfrak{A}$ is a prime model of T. Then, obviously, $\mathfrak{A}$ is countable. Further, a type realized in $\mathfrak{A}$ is realized in *every* model of T and thus must be *locally realized* in T (by the omitting types theorem). In fact we will see that these two properties characterize prime models.

Let $\Gamma(x_0, \ldots, x_n)$ be a complete type of T. If Γ is locally realized in T then there is a formula $\alpha(x_0, \ldots, x_n)$ consistent with T such that

$$T \models \forall x_0 \cdots \forall x_n(\alpha \to \phi)$$

for every $\phi(x_0, \ldots, x_n) \in \Gamma$. Since Γ is complete we must have

$$\alpha(x_0, \ldots, x_n) \in \Gamma.$$

A complete type generated (with respect to T) by a single formula is called *principal*, and the generating formula is a *complete* as in the following definition.

DEFINITION 1.2. A formula $\alpha(x_0, \ldots, x_n)$ is a complete formula (or atom) of T iff $T \models \exists x_0 \cdots \exists x_n \alpha$, and for every $\phi(x_0, \ldots, x_n)$ either $T \models (\alpha \to \phi)$ or $T \models (\alpha \to \neg\phi)$.

DEFINITION 1.3. A model $\mathfrak{A}$ of T is atomic iff for every $a_0, \ldots, a_n \in A$ there is some complete formula $\alpha(x_0, \ldots, x_n)$ of T such that

$$\mathfrak{A}_A \models \alpha(\overline{a_0}, \ldots, \overline{a_n}).$$

We have essentially shown the easy half of the following characterization theorem.

THEOREM 1.1. *Let $\mathfrak{A} \models T$. Then $\mathfrak{A}$ is prime iff $\mathfrak{A}$ is countable and atomic.*

PROOF. ($\Rightarrow$) Assume $\mathfrak{A}$ is prime and let $a_0, \ldots, a_n \in A$. Then $tp_{\mathfrak{A}}(a_0, \ldots, a_n)$ is a complete type realized in $\mathfrak{A}$, hence by the above argument it contains a complete formula. Therefore $\mathfrak{A}$ is atomic.

($\Leftarrow$) Let $\mathfrak{A}$ be countable and atomic. Let $\mathfrak{B}$ be any model of T. We must show that $\mathfrak{A} \cong\prec \mathfrak{B}$. We will do this by a variation on the back-and-forth method. Let $A = \{a_n | \, n \in \omega\}$. We will define, by recursion on n, elements $b_n \in B$ such that for each n we have

$$(\mathfrak{A}, a_0, \ldots, a_n) \equiv (\mathfrak{B}, b_0, \ldots, b_n).$$

That is: $(\mathfrak{B}, b_0, \ldots, b_n) \models Th(\mathfrak{A}_{\{a_0,\ldots,a_n\}})$ in the language $\mathcal{L} \cup \{\overline{a_0}, \ldots, \overline{a_n}\}$ where $\overline{a_i}$ is interpreted as $b_i \in B$. It will then follow that $\mathfrak{B}^* = (\mathfrak{B}, b_n)_{n\in\omega}$ is a model of $Th(\mathfrak{A}_A)$, hence $\mathfrak{A} \cong\prec \mathfrak{B}$ as desired.

So, we first pick b_0. Since $\mathfrak{A}$ is atomic, a_0 must satisfy some complete formula $\alpha_0(x_0)$. Then $\mathfrak{B}_B \models \alpha_0(b_0)$ for some $b_0 \in B$. Since α_0 is complete, we have $T \models \forall x_0(\alpha_0(x_0) \to \phi(x_0))$, for every $\phi \in tp_{\mathfrak{A}}(a_0)$. Therefore $\mathfrak{B}_B \models \phi(\overline{b_0})$ whenever $\mathfrak{A}_A \models \phi(\overline{a_0})$, so $(\mathfrak{A}, a_0) \equiv (\mathfrak{B}, b_0)$ as desired. Now, given $b_0, \ldots, d_n$ such that $(\mathfrak{A}, a_0, \ldots, a_n) \equiv (\mathfrak{B}, b_0, \ldots, b_n)$ we show how to choose b_{n+1}. Let $\alpha_n(x_0, \ldots, x_n)$ be a complete formula satisfied by $a_0, \ldots, a_n$ and let $\alpha_{n+1}(x_0, \ldots, x_{n+1})$ be a complete formula satisfieed by $a_0, \ldots, a_{n+1}$. Then we must have

$$T \models \forall x_0 \cdots \forall x_n(\alpha_n \to \exists x_{n+1}\alpha_{n+1})$$

since $\exists x_{n+1}\alpha_{n+1} \in tp_{\mathfrak{A}}(a_0, \ldots, a_n)$. By our inductive hypothesis we must have $\mathfrak{B}_B \models \alpha_n(\overline{b_0}, \ldots, \overline{b_n})$, so there is some $b_{n+1} \in B$ with $\mathfrak{B}_B \models \alpha_{n+1}(\overline{b_0}, \ldots, \overline{b_{n+1}})$. Then as before we see that

$$(\mathfrak{A}, a_0, \ldots, a_{n+1}) \equiv (\mathfrak{B}, b_0, \ldots, b_{n+1})$$

which completes the proof. □

For each $n = 1, 2, \ldots$ define

$$\Gamma_n(x_0, \ldots, x_{n-1}) = \{\neg\alpha(x_0, \ldots, x_{n-1}) | \ \alpha \text{ is an atom}\}.$$

Then $\mathfrak{A}$ is atomic iff for each n $\mathfrak{A}$ omits Γ_n. Thus, by the preceding theorem, T has a prime model iff it has a countable model omitting each Γ_n. By the extended omitting types theorem, this happens iff T locally omits each Γ_n.

Definition 1.4. T is atomic (or atomistic) iff for every formula $\phi(x_0, \ldots, x_n)$ consistent with T there is some complete formula $\alpha(x_0, \ldots, x_n)$ such that $(\alpha \wedge \phi)$ is consistent with T.

We leave the reader to check that T is atomistic iff T locally omits each Γ_n. The preceding discussion then shows:

Theorem 1.2. *T has a prime model iff T is atomic.*

It is not immediately clear how restrictive the condition that T is atomic really is. There *are* theories that are not atomic, but most examples are specially contrived–$Th((\mathbb{Z}, +))$ is a natural example of a theory without a prime model, but the proof uses a lot of information about this theory.

The next result gives a large class of atomic theories.

Theorem 1.3. *Assume there are just* countably *many different* complete *types consistent with T. Then T has a prime model.*

Proof. For each $n \in \omega$, let $\Phi_k^n(x_0, \ldots, x_n)$ list all *non-principal* complete types in $x_0, \ldots, x_n$ consistent with T. Then each Φ_k^n is locally omitted in T. So, by the extended omitting types theorem, T has a countable model omitting every Φ_k^n; this model must then be atomic, hence prime. □

Theorem 1.4. *If $\mathfrak{A}, \mathfrak{B}$ are both prime models of T then $\mathfrak{A} \cong \mathfrak{B}$.*

2. Universal and Saturated Models

Definition 2.1. A model $\mathfrak{A}$ of T is countably universal iff $\mathfrak{A}$ is countable and whenever $\mathfrak{B}$ is a countable model of T then $\mathfrak{B} \cong\prec \mathfrak{A}$.

The following lemma is clear:

Lemma 2.1. *(1) If $\mathfrak{A}$ is a countably universal model of T then $\mathfrak{A}$ realizes every type consistent with T.*

(2) If T has a countably universal model then there are just countably many complete types consistent with T.

The obvious conjecture is that the converses of both parts of the lemma hold. First note that we *do* have the following.

Proposition 2.2. *If there are just countably many complete types consistent with T then T has a countable model realizing every type consistent with T.*

Thus all that remains is the converse to part (1) of the lemma. In analogy with the proof that countable atomic models are prime, we would proceed as follows: let $\mathfrak{A}$ (be countable and) realize every type consistent with T. Let $\mathfrak{B}$ be an arbitrary countable model of T, say $B = \{b_n | \ n \in \omega\}$. We try to define, by recursion on n, elements $a_n \in A$ such that for every n we have

$$(\mathfrak{B}, b_0, \ldots, b_n) \equiv (\mathfrak{A}, a_0, \ldots, a_n).$$

For $n = 0$ let $\Gamma_0(x_0) = tp_{\mathfrak{B}}(b_0)$. By hypothesis Γ_0 is realized in $\mathfrak{A}$, say by a_0, so we have $(\mathfrak{B}, b_0) \equiv (\mathfrak{A}, a_0)$. Now for $n = 1$ let $\Gamma_1(x_0, x_1) = tp_{\mathfrak{B}}(b_0, b_1)$. Once again there are $a_0', a_1' \in A$ realizing $\Gamma_1(x_0, x_1)$, hence such that $(\mathfrak{B}, b_0', b_1') \equiv (\mathfrak{A}, a_0', a_1')$. But we have already chosen a_0 so we need an a_1 such that (a_0, a_1) realize Γ_1. There is no obvious reason why such an a_1 should exist. We are thus forced into the formulating of the following stronger property:

Definition 2.2. $\mathfrak{A}$ is ω-saturated iff for every $n \in \omega$ and all $a_0, \ldots, a_{n-1} \in A$, the structure $(\mathfrak{A}, a_0, \ldots, a_{n-1})$ [for the language $\mathcal{L}(a_0, \ldots, a_n) = \mathcal{L} \cup \{\overline{a_0}, \ldots, \overline{a_{n-1}}]$ realizes every $\mathcal{L}(a_0, \ldots, a_n)$-type $\Gamma(y)$ consistent with it.

If $\mathfrak{A}$ is countable and ω-saturated, we may also say that $\mathfrak{A}$ is *countably saturated*. The above argument can be carried out for ω-saturated models, so we obtain the following theorem.

Theorem 2.3. *Assume $\mathfrak{A}$ is an ω-saturated model of T. Then every countable model of T can be elementarily embedded in $\mathfrak{A}$. Therefore $\mathfrak{A}$ is countably universal provided $\mathfrak{A}$ is countable.*

We thus have (i)$\Rightarrow$(ii)$\Rightarrow$(iii), where:
(i) $\mathfrak{A}$ is countably saturated
(ii) $\mathfrak{A}$ is countably universal
(iii) $\mathfrak{A}$ is countable and realizes all types consistent with T.

Do either, or both, of these arrows reverse? We cannot answer these questions at the moment, but we can show that the existence of a countable ω-saturated model is equivalent to the existence of the other two types of models.

Theorem 2.4. *The following are equivalent:*
(i) T has a countable ω-saturated model
(ii) T has a countable universal model
(iii) There are just countably many complete types consistent with T.

Proof. It suffices to show (iii)$\Rightarrow$(i). Assume that (iii) holds. Then in fact, we have the following:

(†) for every $\mathfrak{A} \models T$, every $n \in \omega$ and every $a_0, \ldots, a_{n-1} \in A$ there are just countably many different complete 1-types $\Gamma(y)$ of $\mathcal{L} \cup \{\overline{a_0}, \ldots, \overline{a_{n-1}}\}$ consistent with $Th((\mathfrak{A}, a_0, \ldots, a_{n-1}))$.

The reason for this is that every such $\Gamma(y)$ is $\Phi(\overline{a_0}, \ldots, \overline{a_{n-1}}, y)$ for some complete $\mathcal{L}$-type $\Phi(x_0, \ldots, x_{n-1}, y)$ consistent with $Th(\mathfrak{A}) = T$, and (iii) asserts there are just countably many such Φ.

(‡) for every countable $\mathfrak{A} \models T$ there is some countable $\mathfrak{A}'$ such that $\mathfrak{A} \prec \mathfrak{A}'$ and for every $n \in \omega$ and every $a_0, \ldots, a_n \in A$, $\mathfrak{A}'$ realizes every type $\Gamma(y)$ of $\mathcal{L} \cup \{\overline{a_0}, \ldots, \overline{a_n}\}$ consistent with $Th(\mathfrak{A}, a_0, \ldots, a_n)$.

This holds by applying Theorem 4.2.8 to the theory $T^* = Th(\mathfrak{A}_A)$ where the Γ_k's list the complete 1-types in just finitely many new constants–the list being countable since A is countable and by (†).

We now can show T has a countable ω-saturated model. Let $\mathfrak{A}_0$ be any countable model of T an, for every $n \in \omega$, let $\mathfrak{A}_{n+1} = (\mathfrak{A}_n)'$ be the model given by (‡) starting with $\mathfrak{A}_n$. Thus $\{\mathfrak{A}_n\}_{n\in\omega}$ is an elementary chain of countable models. Let $\mathfrak{A} = \bigcup_{n\in\omega} \mathfrak{A}_n$. Then $\mathfrak{A}$ is a countable model of T, which can be shown to be ω-saturated, since every finite subset of A is a finite subset of some A_n, hence the types over that finite subset are all realized in $\mathfrak{A}_{n+1}$, so in particular in $\mathfrak{A}$. □

If T has uncountably many complete types, then we no longer get countable ω-saturated models but similar arguments show:

Proposition 2.5. *T has ω-saturated models of power κ for every $\kappa \geq 2^\omega$.*

Definition 2.3. $\mathfrak{A}$ is ω-homogeneous iff for every $n \in \omega$ and every $a_0, \dots, a_n$ and $b_0, \dots, b_n \in A$

if $(\mathfrak{A}, a_0, \dots, a_n) \equiv (\mathfrak{A}, b_0, \dots, b_n)$ then $(\mathfrak{A}, a_0, \dots, a_n) \cong (\mathfrak{A}, b_0, \dots, b_n)$.

Theorem 2.6. *A countable model is ω-saturated iff it is both ω-homogeneous and universal.*

3. Theories with Just Finitely Many Countable Models

Definition 3.1. A model $\mathfrak{A}$ of T is countably universal iff $\mathfrak{A}$ is countable and whenever $\mathfrak{B}$ is a countable model of T then $\mathfrak{B} \cong\prec \mathfrak{A}$.

The following lemma is clear:

Lemma 3.1. *(1) If $\mathfrak{A}$ is a countably universal model of T then $\mathfrak{A}$ realizes every type consistent with T.*

(2) If T has a countably universal model then there are just countably many complete types consistent with T.

The obvious conjecture is that the converses of both parts of the lemma hold. First note that we *do* have the following.

Proposition 3.2. *If there are just countably many complete types consistent with T then T has a countable model realizing every type consistent with T.*

Thus all that remains is the converse to part (1) of the lemma. In analogy with the proof that countable atomic models are prime, we would proceed as follows: let $\mathfrak{A}$ (be countable and) realize every type consistent with T. Let $\mathfrak{B}$ be an arbitrary countable model of T, say $B = \{b_n |\ n \in \omega\}$. We try to define, by recursion on n, elements $a_n \in A$ such that for every n we have

$$(\mathfrak{B}, b_0, \dots, b_n) \equiv (\mathfrak{A}, a_0, \dots, a_n).$$

For $n = 0$ let $\Gamma_0(x_0) = tp_{\mathfrak{B}}(b_0)$. By hypothesis Γ_0 is realized in $\mathfrak{A}$, say by a_0, so we have $(\mathfrak{B}, b_0) \equiv (\mathfrak{A}, a_0)$. Now for $n = 1$ let $\Gamma_1(x_0, x_1) = tp_{\mathfrak{B}}(b_0, b_1)$. Once again there are $a_0', a_1' \in A$ realizing $\Gamma_1(x_0, x_1)$, hence such that $(\mathfrak{B}, b_0', b_1') \equiv (\mathfrak{A}, a_0', a_1')$. But we have already chosen a_0 so we need an a_1 such that (a_0, a_1) realize Γ_1. There is no obvious reason why such an a_1 should exist. We are thus forced into the formulating of the following stronger property:

Definition 3.2. $\mathfrak{A}$ is ω-saturated iff for every $n \in \omega$ and all $a_0, \dots, a_{n-1} \in A$, the structure $(\mathfrak{A}, a_0, \dots, a_{n-1})$ [for the language $\mathcal{L}(a_0, \dots, a_n) = \mathcal{L} \cup \{\overline{a_0}, \dots, \overline{a_{n-1}}]$ realizes every $\mathcal{L}(a_0, \dots, a_n)$-type $\Gamma(y)$ consistent with it.

If $\mathfrak{A}$ is countable and ω-saturated, we may also say that $\mathfrak{A}$ is *countably saturated*. The above argument can be carried out for ω-saturated models, so we obtain the following theorem.

THEOREM 3.3. *Assume $\mathfrak{A}$ is an ω-saturated model of T. Then every countable model of T can be elementarily embedded in $\mathfrak{A}$. Therefore $\mathfrak{A}$ is countably universal provided $\mathfrak{A}$ is countable.*

We thus have (i)$\Rightarrow$(ii)$\Rightarrow$(iii), where:
(i) $\mathfrak{A}$ is countably saturated
(ii) $\mathfrak{A}$ is countably universal
(iii) $\mathfrak{A}$ is countable and realizes all types consistent with T.
Do either, or both, of these arrows reverse? We cannot answer these questions at the moment, but we can show that the existence of a countable ω-saturated model is equivalent to the existence of the other two types of models.

THEOREM 3.4. *The following are equivalent:*
(i) T has a countable ω-saturated model
(ii) T has a countable universal model
(iii) There are just countably many complete types consistent with T.

PROOF. It suffices to show (iii)$\Rightarrow$(i). Assume that (iii) holds. Then in fact, we have the following:
(†) for every $\mathfrak{A} \models T$, every $n \in \omega$ and every $a_0, \ldots, a_{n-1} \in A$ there are just countably many different complete 1-types $\Gamma(y)$ of $\mathcal{L} \cup \{\overline{a_0}, \ldots, \overline{a_{n-1}}\}$ consistent with $Th((\mathfrak{A}, a_0, \ldots, a_{n-1}))$.
The reason for this is that every such $\Gamma(y)$ is $\Phi(\overline{a_0}, \ldots, \overline{a_{n-1}}, y)$ for some complete $\mathcal{L}$-type $\Phi(x_0, \ldots, x_{n-1}, y)$ consistent with $Th(\mathfrak{A}) = T$, and (iii) asserts there are just countably many such Φ.
(‡) for every countable $\mathfrak{A} \models T$ there is some countable $\mathfrak{A}'$ such that $\mathfrak{A} \prec \mathfrak{A}'$ and for every $n \in \omega$ and every $a_0, \ldots, a_n \in A$, $\mathfrak{A}'$ realizes every type $\Gamma(y)$ of $\mathcal{L} \cup \{\overline{a_0}, \ldots, \overline{a_n}\}$ consistent with $Th(\mathfrak{A}, a_0, \ldots, a_n)$.
This holds by applying Theorem 4.2.8 to the theory $T^* = Th(\mathfrak{A}_A)$ where the Γ_k's list the complete 1-types in just finitely many new constants–the list being countable since A is countable and by (†).
We now can show T has a countable ω-saturated model. Let $\mathfrak{A}_0$ be any countable model of T an, for every $n \in \omega$, let $\mathfrak{A}_{n+1} = (\mathfrak{A}_n)'$ be the model given by (‡) starting with $\mathfrak{A}_n$. Thus $\{\mathfrak{A}_n\}_{n\in\omega}$ is an elementary chain of countable models. Let $\mathfrak{A} = \bigcup_{n\in\omega} \mathfrak{A}_n$. Then $\mathfrak{A}$ is a countable model of T, which can be shown to be ω-saturated, since every finite subset of A is a finite subset of some A_n, hence the types over that finite subset are all realized in $\mathfrak{A}_{n+1}$, so in particular in $\mathfrak{A}$. □

If T has uncountably many complete types, then we no longer get countable ω-saturated models but similar arguments show:

PROPOSITION 3.5. *T has ω-saturated models of power κ for every $\kappa \geq 2^\omega$.*

DEFINITION 3.3. $\mathfrak{A}$ is ω-homogeneous iff for every $n \in \omega$ and every $a_0, \ldots, a_n$ and $b_0, \ldots, b_n \in A$

$$\text{if } (\mathfrak{A}, a_0, \ldots, a_n) \equiv (\mathfrak{A}, b_0, \ldots, b_n) \text{ then } (\mathfrak{A}, a_0, \ldots, a_n) \cong (\mathfrak{A}, b_0, \ldots, b_n).$$

THEOREM 3.6. *A countable model is ω-saturated iff it is both ω-homogeneous and universal.*

4. Exercises

(1)

CHAPTER 6

Further Topics in Model Theory

0. Introduction

In this chapter we cover a number of additional topics needed for further work in model theory. Proofs are only given in outline, with the reader left to fill in the details. We begin with another application of the Henkin method, to prove Craig's interpolation Theorem, a fundamental property of first-order theories. From it we derive the definability result of Beth. In 6.2 we define the κ-saturated models, for arbitrary κ, generalizing the results in 5.2. Skolem functions are introduced in 6.3 and used to derive stronger forms of the Löwenheim-Skolem theorem. We also introduce the notion of indiscernible sequence, although our proof of the existence of such sequences is even sketchier than usual and will depend on a combinatorial theorem which we will not prove. Various applications are briefly discussed in 6.4, including preservation theorems, model-completeness and 2-cardinal theorems. We also exhibit a finite axiomatization of $\mathrm{Th}(\langle\omega, \leq\rangle)$ and study the countable models of this theory which provide counterexamples to the conjectures of the previous chapter.

1. Interpolation and Definability

Let $\mathcal{L}_1$ and $\mathcal{L}_2$ be two languages, let $\mathcal{L} = \mathcal{L}_1 \cap \mathcal{L}_2$, and let σ_1, σ_2 be sentences of $\mathcal{L}_1, \mathcal{L}_2$ respectively. Suppose that $\models (\sigma_1 \to \sigma_2)$ — that is, whenever $\mathfrak{A}$ is a structure for $(\mathcal{L}_1 \cup \mathcal{L}_2)$ and $\mathfrak{A} \models \sigma_1$ then $\mathfrak{A} \models \sigma_2$. Now $\mathfrak{A} \models \sigma_1$ iff $\mathfrak{A} \restriction \mathcal{L}_1 \models \sigma_1$, so if we take *any* $\mathfrak{A}'$ with $\mathfrak{A}' \restriction \mathcal{L}_1 = \mathcal{A} \restriction \mathcal{L}_1$ we will still have $\mathfrak{A}' \models \sigma_1$, and so $\mathfrak{A}' \models \sigma_2$ (despite the fact that $\mathfrak{A}' \restriction \mathcal{L}_2 \neq \mathfrak{A} \restriction \mathcal{L}_2$). Similarly, if $\mathfrak{A}''$ is such that $\mathfrak{A}'' \restriction \mathcal{L}_2 = \mathfrak{A}' \restriction \mathcal{L}_2$ then we still have $\mathfrak{A}'' \models \sigma_2$. That is, we have shown the following: if $\mathfrak{A} \models \sigma_1$ and $\mathfrak{A}'' \restriction \mathcal{L} = \mathfrak{A} \restriction \mathcal{L}$ then $\mathfrak{A}'' \models \sigma_2$. It is reasonable to expect, then, that the validity of the *implication* depends just on the language $\mathcal{L}$. This is in fact true, and is the content of Craig's well-known theorem.

Theorem 1.1. ***(Interpolation)*** *Let $\mathcal{L} = \mathcal{L}_1 \cap \mathcal{L}_2$, $\sigma_1 \in \mathrm{Sn}_{\mathcal{L}_1}, \sigma_2 \in \mathrm{Sn}_{\mathcal{L}_2}$. Suppose that $\models (\sigma_1 \to \sigma_2)$. Then there is some sentence θ of $\mathcal{L}$ such that*

$$\models (\sigma_1 \to \theta) \text{ and } \models \theta \to \sigma_2.$$

Proof. First note that we may assume that $\mathcal{L}_1$ and $\mathcal{L}_2$ are both countable (since only finitely many symbols can occur in σ_1 and σ_2).

We may assume that there is *no* sentence θ of $\mathcal{L}$ such that $\models (\sigma_1 \to \theta)$ and $\models (\theta \to \sigma_2)$, or, equivalently, there is no sentence θ of $\mathcal{L}$ such that $\sigma_1 \models \theta$ and $\neg\sigma_2 \models \neg\theta$. We will show that $\{\sigma_1, \neg\sigma_2\}$ has a model and so $\not\models (\sigma_1 \to \sigma_2)$.

Let C_0 be a countably infinite set of individual constant symbols not in $\mathcal{L}_1 \cup \mathcal{L}_2$. Let $\mathcal{L}^*, \mathcal{L}_1^*, \mathcal{L}_2^*$ be the results of adding the constants in C_0 to $\mathcal{L}, \mathcal{L}_1, \mathcal{L}_2$ respectively.

We will prove that there is a Henkin set Σ_1^* of $\mathcal{L}_1^*$-sentences and a Henkin set Σ_2^* of $\mathcal{L}_2^*$-sentences such that $\sigma_1 \in \Sigma_1^*$, $\neg\sigma_2 \in \Sigma_2^*$ and (†) for every sentence θ of $\mathcal{L}^*$, $\theta \in \Sigma_1^*$ iff $\theta \in \Sigma_2^*$.

Once we have done this, let $\mathfrak{A}_1$ be an $\mathcal{L}_1^*$-structure which is a canonical model for Σ_1^*, and let $\mathfrak{A}_2$ be an $\mathcal{L}_2^*$-structure which is a canonical model for Σ_2^*. Then $\mathfrak{A}_1 \restriction \mathcal{L}^* \cong \mathfrak{A}_2 \restriction \mathcal{L}^*$ by (†), say under the isomorphism h. We can then expand $\mathfrak{A}_1$ to an $(\mathcal{L}_1^* \cup \mathcal{L}_2^*)$-structure $\mathfrak{A}$ so that h is an isomorphism of $\mathfrak{A}_1 \restriction \mathcal{L}_2^*$ onto $\mathfrak{A}_2$. Therefore $\mathfrak{A}$ is a model of $\Sigma_1^* \cup \Sigma_2^*$, in particular of $\{\sigma_1, \neg\sigma_2\}$ as desired.

Thus we are done once we have defined Henkin sets Σ_1^*, Σ_2^* as above. This we do by carrying out the Henkin construction *simultaneously* starting with $\Sigma_1 = \{\sigma_1\}$ and $\Sigma_2 = \{\neg\sigma_2\}$, with an added *joint* condition to guarantee that we end up with (†) holding.

Let $\mathrm{Sn}_{\mathcal{L}_1^*} = \{\theta_n^1 : n \in \omega\}$ and $\mathrm{Sn}_{\mathcal{L}_2^*} = \{\theta_n^2 : n \in \omega\}$. We define, by recursion on n, sets $\Gamma_n^1 \subseteq \mathrm{Sn}_{\mathcal{L}_1^*}$ and $\Gamma_n^2 \subseteq \mathrm{Sn}_{\mathcal{L}_2^*}$ starting with $\Gamma_0^i = \Sigma_i$ $(i = 1, 2)$ and such that the following hold for all $n \in \omega$ and $i = 1, 2$:

(0_n^i) $\Gamma_n^i \subseteq \Gamma_{n+1}^i$ and $\Gamma_{n+1}^i \smallsetminus \Gamma_n^i$ is finite,

(1_n^i) $\theta_n^i \in \Gamma_{n+1}^i$ or $\neg\theta_n^i \in \Gamma_{n+1}^i$,

(2_n^i) if $\theta_n^i \notin \Gamma_{n+1}^i$ and $\theta_n^i = \forall x \varphi(x)$ for some $\varphi(x)$, then $\neg\varphi(c) \in \Gamma_{n+1}^i$ for some constant c,

and the following joint condition

(3_n) there is no sentence θ of $\mathcal{L}^*$ such that $\Gamma_{n+1}^* \models \theta$ and $\Gamma_{n+1}^2 \models \neg\theta$.

Note that (3_n) implies that both Γ_{n+1}^1 and Γ_{n+1}^2 are consistent. Further (3_n) for all n guarantees that (†) holds for $\Sigma_i^* = \bigcup_{n\in\omega} \Gamma_n^i$, and thus we are done once we show the construction can be carried out.

First of all, note that our beginning assumption is that there is no sentence θ of $\mathcal{L}$ such that $\Gamma_0^i \models \theta$ and $\Gamma_0^2 \models \neg\theta$. Since Γ_0^i do not contain constants from C_0 we see that there is no sentence θ of $\mathcal{L}^*$ such that $\Gamma_0^i \models \theta^*$ and $\Gamma_0^2 \models \neg\theta^*$ (by generalization on constants). [We refer to this as (3_{-1})]

So, given Γ_n^i, finite sets of sentences of $\mathcal{L}_i^*$ satisfying (3_{n-1}), we wish to define Γ_{n+1}^i satisfying $(0_n^i), (1_n^i), (2_n^i)$ and (3_n). We follow the proof of the corresponding steps in the Completeness Theorem (pp 50-51) to define in turn $\Gamma_{n+1/2}^1, \Gamma_{n+1}^1, \Gamma_{n+1/2}^2, \Gamma_{n+1}^2$ where the requirements that each of these sets is consistent is replaced by the stronger requirements leading to (3_n) — that is, that there is no sentence θ of $\mathcal{L}^*$ such that

(a) $\Gamma_{n+1/2}^1 \models \theta$ and $\Gamma_n^2 \models \neg\theta$,
(b) $\Gamma_{n+1}^1 \models \theta$ and $\Gamma_n^2 \models \neg\theta$,
(c) $\Gamma_{n+1}^1 \models \theta$ and $\Gamma_{n+1/2}^2 \models \neg\theta$, and finally (3_n) itself.

Obviously, the hypothesis that (3_{n-1}) holds is used to start. All the sets are then similar to those of the Completeness Theorem, and are left to the reader.

□

Interpolation also holds for formulas with free variables in place of sentences σ_1, σ_2. To see this, replace the preceding Theorem to obtain an interpolating θ' with these new constants, and replace the constants with the variable to get a formula θ of $\mathcal{L}$. Thus:

THEOREM 1.2. *Let* $\mathcal{L} = \mathcal{L}_1 \cap \mathcal{L}_2$, $\varphi_i(\overline{x}) \in \mathrm{Fm}_{\mathcal{L}_i}$ $(i = 1, 2)$. *Suppose that* $\models (\varphi_1 \to \varphi_2)$. *Then there is a* $\theta(\overline{x})$ *of* $\mathcal{L}$ *such that* $\models (\varphi_1 \to \theta)$ *and* $\models (\theta \to \varphi_2)$.

As an application of Craig's Interpolation Theorem (in the form applying to formulas) we derive another classic result, which concerns definability of a relation in all models of a theory.

Let $\mathcal{L}$ be a language, R a predicate symbol of n arguments no in $\mathcal{L}$. Let $\mathcal{L}_1 = \mathcal{L} \cup \{R\}$, and let T be a theory of $\mathcal{L}_1$. We say that R is *($\mathcal{L}$ -) definable in T* iff there is a formula $\varphi(x_, \dots, x_{n-1})$ of $\mathcal{L}$ such that

$$T \models \forall x_0 \cdots \forall x_{n-1}[Rx_0 \cdots x_{n-1} \leftrightarrow \varphi(\overline{x})].$$

(Reference to $\mathcal{L}$ is sometimes suppressed).

Clearly if R is definable in T then the interpretation of R in any model of T is uniquely determined by the interpretations of the other symbols (since it must be whatever relation φ defines) — more formally, the following must hold:

($*$) if $\mathfrak{A}, \mathfrak{B} \models T$ and $\mathfrak{A} \restriction \mathcal{L} = \mathfrak{B} \restriction \mathcal{L}$, then $R^{\mathfrak{A}} = R^{\mathfrak{B}}$ (and so $\mathfrak{A} = \mathfrak{B}$).

The definability theorem of Beth states that the converse also holds.

THEOREM 1.3. *Let $\mathcal{L}_1 = \mathcal{L} \cup \{R\}$ where R is an n-ary predicate symbol not in $\mathcal{L}$. Then R is $\mathcal{L}$-definable in T iff ($*$) holds for T, where T is any theory of $\mathcal{L}$.*

PROOF. It suffices to show that ($*$) implies the definability of R, so let us suppose that ($*$) holds. We wish to formulate ($*$) in a syntactical manner. Let S be another n-ary predicate symbol not in $\mathcal{L}_1$. Let R' be the theory of $\mathcal{L}_2 = \mathcal{L} \cup \{S\}$ obtained from T by replacing R everywhere by S. Then T and T' are "the same", except one uses the symbol R where the other uses S. ($*$) then precisely says that

$$T \cup T' \models \forall \overline{x}[R\overline{x} \leftrightarrow S\overline{x}].$$

In particular, then

$$T \cup T' \models (R\overline{x} \to S\overline{x}).$$

By compactness there are sentences $\sigma \in T, \sigma' \in T'$ such that

$$\{\sigma, \sigma'\} \models (R\overline{x} \to S\overline{s}),$$

i.e.,

$$\models (\sigma \wedge R\overline{x}) \to (\sigma' \to S\overline{x}).$$

Interpolation applied to $\varphi_1 = (\sigma \wedge R\overline{x})$ and $\varphi_2 = (\sigma' \to S\overline{x})$ yields a formula $\theta(\overline{x})$ such that

$$\models (\sigma \wedge R\overline{x}) \to \theta(\overline{x}) \text{ and } \models \theta(\overline{x}) \to (\sigma' \to S\overline{x}).$$

Therefore $T \models (R\overline{x} \to \theta)$ and $T' \models (\theta \to S\overline{x})$. Replacing S everywhere by R in the last consequence yields $T \models (\theta \to R\overline{x})$.

Thus we have $T \models (R\overline{x} \leftrightarrow \theta(\overline{x}))$, and so R is $\mathcal{L}$-definable in T, as desired. $\square$

Beth's Definability Theorem also applies to functions — simply consider an n-ary function as an $(n+1)$-ary predicate. The defining formula then defines the graph of the function.

The following easy fact is sometimes useful.

LEMMA 1.4. *Assume R is $\mathcal{L}$-definable in T. Then every $\mathcal{L}$-structure $\mathfrak{A}$ which is a model of $T \cap \mathrm{Sn}_{\mathcal{L}}$ can be expanded to an $\mathcal{L} \cup \{R\}$-structure $\mathfrak{A}_1$ which is a model of T.*

PROOF. Simply define $R^{\mathfrak{A}_1} = \varphi^{\mathfrak{A}}$, where φ is the $\mathcal{L}$-formula defining R in T. This works since for any sentence θ of $\mathcal{L}_1$ we have $T \models \theta$ iff $T \models \theta'$, where θ' is an $\mathcal{L}$-sentence resulting from θ by replacing R throughout by φ. $\square$

Specific, direct applications of Beth's theorem are hard to come by. The most striking applications are of the contrapositive — i.e., from the *non*-definability of R in T to conclude $(*)$ fails. The non-definability might be shown directly or use the contrapositive of the above lemm.

EXAMPLE 1.1. $\mathcal{L} = \{\leq\}$, $\mathcal{L}_1 = \mathcal{L} \cup \{+\}$, $T = \mathrm{Th}(\langle\omega, \leq, +\rangle)$. Then $\mathfrak{A} = \langle\omega + \mathbb{Z}, \leq\rangle \models T \cap \mathrm{Sn}_{\mathcal{L}}$ but $\mathfrak{A}$ cannot be expanded to a model of T, hence there is some $\mathfrak{B} = \langle B, \leq^{\mathfrak{B}}\rangle \equiv \langle\omega, \leq\rangle$ such that there are 2 different functions $+_1$ and $+_2$ on B such that $\langle B, \leq^{\mathfrak{B}}, +_1\rangle \equiv \langle B, \leq^{\mathfrak{B}}, +_2\rangle \equiv \langle\omega, \leq, +\rangle$.

EXAMPLE 1.2. Similarly, for some $\langle B, \leq^{\mathfrak{B}}, +^{\mathfrak{B}}\rangle$ there are 2 different $\cdot_1, \cdot_2$ with $\langle B, \leq^{\mathfrak{B}}, +^{\mathfrak{B}}, \cdot_1\rangle \equiv \langle B, \leq^{\mathfrak{B}}, +^{\mathfrak{B}}, \cdot_2\rangle \equiv \langle\omega, \leq, +, \cdot\rangle$.

2. Saturated Models

This section concerns generalizations of the material in 5.2 to uncountable cardinals. For simplicity, *we consider only theories in countable languages in this section.* This restriction is not essential, but without it we would have to make reference to the cardinality of the language in our results, and proofs.

The fundamental definitions are as follows:

DEFINITION 2.1. (a) Let $\kappa \geq \omega$. A model $\mathfrak{A}$ is *κ-saturated* iff for every $X \subseteq A$, if $|X| < \kappa$ then $\mathfrak{A}_X$ realizes all $\mathcal{L}(X)$ types (in one variable) consistent with $\mathrm{Th}(\mathfrak{A}_X)$.

(b) Let T be a complete theory and $\kappa > \omega$. $\mathfrak{A}$ is a *κ-universal* model of T iff $\mathfrak{A} \models T$ and whenever $\mathfrak{B} \models T$, $|B| < \kappa$ then $\mathfrak{B} \cong\prec \mathfrak{A}$.

In this terminology, the theorem on pg. 92 states that on ω-saturated model of T is ω_1-universal. The generalization we obtain replaces ω by any $\kappa \geq \omega$ and ω_1 by κ^+.

EXAMPLE 2.1. $\mathcal{L}$ has just a unary predicate P, T is the complete theory stating that both P and $\neg P$ are infinite. Then all models of T are ω-saturated, but of the e models of T of cardinality ω_1, only the one in which both P and $\neg P$ are uncountable is ω_1-saturated.

EXAMPLE 2.2. $\langle\mathbb{R}, \leq\rangle$ is ω-saturated but not ω_1-saturated.

We first require a better notation for dealing with models of the form $\mathfrak{A}_X$.

NOTATION 2. Let $X = \{a_\xi : \xi < \alpha\} \subseteq A$. Then $(\mathfrak{A}, a_\xi)_{\xi<\alpha}$ is the expansion of $\mathfrak{A}$ to the language $\mathcal{L} \cup \{c_\xi : \xi < \alpha\}$, where c_ξ's are distinct constants symbols not in $\mathcal{L}$. *By convention*, if in the same context we also have $(\mathfrak{B}, b_\xi)_{\xi<\alpha}$ this is also a structure for the *same* language, and thus the statements $(\mathfrak{A}, a_\xi)_{\xi<\alpha} \equiv (\mathfrak{B}, b_\xi)_{\xi<\alpha}$ and $(\mathfrak{A}, a_\xi)_{\xi<\alpha} \cong (\mathfrak{B}, b_{xi})_{\xi<\alpha}$ are unambiguously defined.

We can now follow the argument in 5.2, although using *transfinite* recursion, to establish the following:

THEOREM 2.1. *Let $\mathfrak{A} \models T$ be κ-saturated ($\kappa \geq \omega$). Then $\mathfrak{A}$ is a κ^+-universal model of T.*

PROOF. Let $\mathfrak{B} \models T$, $|\mathfrak{B}| \leq \kappa$. Then $B = \{b_\xi : \xi < \kappa\}$. We define, by recursion on $\xi < \kappa$, elements $a_\xi \in A$ such that for every $\alpha \leq \kappa$ we have:

$$(\mathfrak{B}, b_\xi)_{\xi<\alpha} \equiv (\mathfrak{A}, a_\xi)_{\xi<\alpha}.$$

Then the map sending b_ξ to a_ξ (all $\xi < \kappa$) will be an elementary embedding of $\mathfrak{B}$ into $\mathfrak{A}$.

The base of our induction $\alpha = 0$, is the statement $\mathfrak{B} \equiv \mathfrak{A}$, which holds since T is complete.

We break the inductive step into two cases — successor and limit ordinals.

First case: $\alpha = \beta+1$. The inductive hypothesis is that $(\mathfrak{B}, b_\xi)_{\xi<\beta} \equiv (\mathfrak{A}, a_\xi)_{\xi<\beta}$. We need to define a_β so that $(\mathfrak{B}, b_\xi)_{\xi\leq\beta} \equiv (\mathfrak{A}, a_\xi)_{\xi\leq\beta}$, which we do by picking a_β to realize in $(\mathfrak{A}, a_\xi)_{\xi<\beta}$ the complete $(\mathcal{L} \cup \{c_\xi : \xi < \beta\})$-type of b_β in $(\mathfrak{B}, b_\xi)_{\xi<\beta}$. This can be done since $\mathfrak{A}$ is κ-saturated and $|\{a_\xi : \xi < \beta\}| \leq |\beta| \leq \beta < \kappa$.

Second case: If α is a limit ordinal then our inductive hypothesis is that for every $\beta < \alpha$ $(\mathfrak{B}, b_\xi)_{\xi<\beta} \equiv (\mathfrak{A}, a_\xi)_{\xi<\beta}$. It follows that $(\mathfrak{B}, b_\xi)_{\xi<\alpha} \equiv (\mathfrak{A}, a_\xi)_{\xi<\alpha}$, since a sentence of $\mathcal{L} \cup \{c_\xi : \xi < \beta\}$ can only contain finitely many constants and hence really is a sentence of $\mathcal{L} \cup \{c_\xi : \xi < \beta\}$ for some $\beta < \alpha$. This completes the limit ordinal step. □

This proof is really half of a back-and-forth argument (of length κ rather than ω). Adding the other half of the argument we obtain this uniqueness result, details of which are left to the reader.

Theorem 2.2. *If $\mathfrak{A}, \mathfrak{B}$ are both κ-saturated, $|A| = |B| = \kappa$ and $\mathfrak{A} \equiv \mathfrak{B}$, then $\mathfrak{A} \cong \mathfrak{B}$.*

Terminology 1. If $\mathfrak{A}$ is κ-saturated when $\kappa = |A|$, we say that $\mathfrak{A}$ is *saturated.*

In 5.2 we characterized the theories having countable saturated models. The existence of uncountable saturated models is complicated by set theoretical considerations. The following lemma is clear from the definition.

Lemma 2.3. *If T has a κ^+-saturated model of cardinality κ^+ then for every model $\mathfrak{A}$ of T and every $X \subseteq A$ with $|X| \leq \kappa$, there are at most κ^+ complete types of $\mathcal{L}(X)$ consistent with* $\mathrm{Th}(\mathfrak{A}_X)$.

The problem is that there may be 2^κ complete types consistent with a theory in a language of cardinality κ. In this fashion we see, for example:

Corollary 2.4. *If* $\mathrm{Th}(\langle \mathbb{Q}, \leq \rangle)$ *has a saturated model of cardinality of ω_1, then $2^\omega = \omega_1$.*

We are in even worse shape with singular cardinals since one can easily show, for example, that $\mathrm{Th}(\langle \mathbb{Q}, \leq \rangle)$ has no saturated model of cardinality λ, for any singular λ. Hint to the reader: find a set $X \subseteq A$ where $|X| = \mathrm{cf}(\lambda) < \lambda$ and a type of $\mathcal{L}(X)$ not realized in $\mathfrak{A}_X$, for any $\mathfrak{A} \equiv \langle \mathbb{Q}, \leq \rangle, |A| = \lambda$.

The set-theoretical problem pointed out above is the only difficulty with κ^+-saturated models.

Theorem 2.5. *Every model of cardinality $\leq 2^\kappa$ has a κ^+-saturated elementary extension of cardinality $\leq 2^\kappa$.*

Proof. Here too we follow the proof of the corresponding existence result in 5.2, with additional cardinality arguments. We first show:

(1) Let $|B| \leq 2^\kappa$, $X \subseteq B$, $|X| \leq \kappa$. Then $\mathfrak{B} \prec \mathfrak{B}'$ for some $\mathfrak{B}'$ such that $|B'| \leq 2^\kappa$ and $\mathfrak{B}'_X$ realizes every $\mathcal{L}(X)$-type consistent with $\mathrm{Th}(\mathfrak{B}_X) = \mathrm{Th}(\mathfrak{B}'_X)$.

First note that $|\mathcal{L}(X)| \leq \kappa$ hence there are at most 2^κ $\mathcal{L}(X)$-types consistent with $\mathrm{Th}(\mathfrak{B}_X)$. By Compactness and the Löwenheim-Skolem Theorem we can find a model $\mathfrak{B}'$ such that $\mathfrak{B} \prec \mathfrak{B}'$, $\mathfrak{B}'_X$ realizes each of the $\mathcal{L}(X)$-types consisten with $\mathrm{Th}(\mathfrak{B}_X)$, and $|B'| \leq 2^\kappa$ (cf. the Theorem on pg 73), thus establishing (1). Next we derive:

(2) Let $|B| \leq 2^\kappa$. Then there is some $\mathfrak{B}^*$ such that $\mathfrak{B} \prec \mathfrak{B}^*$, $|B^*| \leq 2^\kappa$ and for *every* $X \subseteq B$, $|X| \leq \kappa$, $\mathfrak{B}^*_X$ realizes every $\mathcal{L}(X)$-type consistent with $\mathrm{Th}(\mathfrak{B}_X) = \mathrm{Th}(\mathfrak{B}^*_X)$.

We first establish a necessary fact about cardinality, namely:

$$|\{X : X \subseteq B, |X| \leq \kappa\}| \leq 2^\kappa.$$

This is true since $\{X : X \subseteq B, |X| \leq \kappa\} \preceq {}^\kappa B$ and $|{}^\kappa B| \leq (2^\kappa)^\kappa = 2^\kappa$. Thus we can write $\{X : X \subseteq B, |X| \leq \kappa\} = \{X_\xi : \xi < 2^\kappa\}$. Finally, we define an elementary chain by: $\mathfrak{B}_0 = \mathfrak{B}, \mathfrak{B}_{\xi+1} = (\mathfrak{B}_\xi)', \mathfrak{B}_\eta = \bigcup_{\xi<\eta} \xi$ for limit $\eta \leq 2^\kappa$, where $(\mathfrak{B}_\xi)'$ is the result applying (1) to $\mathfrak{B}_\xi$, for the set X_ξ. One easily establishes by induction that $|B_\xi| \leq 2^\kappa$ for all $\xi \leq 2^\kappa$ (the limit ordinal cases are the only ones requiring attention), and so $\mathfrak{B}^* = \mathfrak{B}_{2^\kappa}$is as desired.

Finally, let $\mathfrak{A}$ be given with $|A| \leq 2^\kappa$. We define an elementary chain $\mathfrak{A}_\xi, \xi \leq \kappa^+$, by $\mathfrak{A}_0 = \mathfrak{A}, \mathfrak{A}_{\xi+1} = (\xi)^*$, $\mathfrak{A}_\eta = \bigcup_{\xi<\eta} \mathfrak{A}_\xi$ for limit ordinals $\eta \leq \kappa^+$. Once again, it is easy to see that $\mathfrak{A}^\# = \mathfrak{A}_{\kappa^+}$ has cardinality $\leq 2^\kappa$. If $X \subseteq A^\#$, $X \leq \kappa$ then $X \subseteq A_\xi$ for some $\xi < \kappa^+$ *because κ^+ is regular*, hence every $\mathcal{L}(X)$-type consistent with $\mathrm{Th}(\mathfrak{A}^\#_X)$ is realized in $(\mathfrak{A}_{\xi+1})_X$. Thus $\mathfrak{A}^\#$ is κ^+-saturated. □

The reader should try to find an ω_1-saturated elementary extension of $\langle \mathbb{R}, \leq \rangle$ of cardinality 2^ω.

One more reason for the seductive appeal of the GCH is the following obvious consequence of the preceding:

Corollary 2.6. *Assume GCH. Then every theory has a saturated model of cardinality κ^+ for every $\kappa \geq \omega$.*

Knowing that a theory has κ^+-universal models is not terribly useful. But knowing that it has saturated models is very useful. One can give, for example, a simple proof of the Interpolation Theorem of the preceding section assuming the existence of saturated models. As the reader is left to check, it suffices to establish the following lemma (due to A. Robinson who used it to derive Beth's Theorem):

Lemma 2.7. *Let $\mathcal{L} = \mathcal{L}_1 \cap \mathcal{L}_2$, let T_1, T_2 be theories of $\mathcal{L}_1, \mathcal{L}_2$ r'espectively, and suppose $T = (T_1 \restriction \mathcal{L} \cup T_2 \restriction \mathcal{L})$ is consistent. Then $T_1 \cup T_2$ has a model.*

Proof. One first shows (using compactness) that there is some *complete* $\mathcal{L}$-theory $T^* \supseteq T$ such that $(T_1 \cup T^*)$ and $(T_2 \cup T^*)$ both have models. *Assume $(T \cup T^*)$ and $(T_2 \cup T^*)$ both have saturated models of cardinality κ.* Let these saturated models be $\mathfrak{A}_1$ and $\mathfrak{A}_2$ respectively. Then $\mathfrak{A}_1 \restriction \mathcal{L}$ and $\mathfrak{A}_2 \restriction \mathcal{L}$ are also saturated and both models of T^*, hence $\mathfrak{A}_1 \restriction \mathcal{L} \cong \mathfrak{A}_2 \restriction \mathcal{L}$ by the uniqueness theorem. One then defines $\mathfrak{A} \models (T_1 \cup T_2 \cup T^*)$ such that $\mathfrak{A} \restriction \mathcal{L}_1 = \mathfrak{A}_1$ and $\mathfrak{A} \restriction \mathcal{L}_2 \cong \mathfrak{A}_2$ as in the proof in the preceding section. □

Thus, we have another proof of Interpolation, under the set theoretic assumption that $\kappa^+ = 2^\kappa$ for some $\kappa \geq \omega$. There are ways of using essentially this same proof and avoiding the additional assumption, but we will not pursue them here.

The spirit in which saturated models were introduced originally involved uniqueness less than their acting as a sort of "universal domain" for a theory T. That is, fixing a sufficiently large saturated model of T, one could study the models of

T just by looking at elementary submodels of this fixed model. Some arguments require the existence of many automorphisms.

DEFINITION 2.2. $\mathfrak{A}$ is *strongly κ-homogeneous* ($\kappa \geq \omega$) iff: for every $\alpha < \kappa$ and for all $a_\xi, a'_\xi \in A$ for $\xi < \alpha$, if $(\mathfrak{A}, a_\xi)_{\xi<\alpha} \equiv (\mathfrak{A}, a'_\xi)_{\xi<\alpha}$ then $(\mathfrak{A}, a_\xi)_{\xi<\alpha} \cong (\mathfrak{A}, a'_\xi)_{\xi<\alpha}$.

DEFINITION 2.3. $\mathfrak{A}$ is *strongly κ-saturated* iff $\mathfrak{A}$ is κ-saturated and strongly κ-homogeneous.

As a consequence of the proof of uniqueness we see:

COROLLARY 2.8. *If $\mathfrak{A}$ is saturated and $|A| = \kappa$ then $\mathfrak{A}$ is strongly κ-saturated.*

The κ^+-saturated models of cardinality 2^κ we have constructed are not automatically strongly κ^+-saturated. We can make them strongly κ^+-saturated by use of the following result in our construction.

LEMMA 2.9. *Assume that $(\mathfrak{A}, a_\xi)_{\xi<\alpha} \equiv (\mathfrak{A}, a'_\xi)_{\xi<\alpha}$. Then there is some $\mathfrak{B}$ such that $\mathfrak{A} \prec \mathfrak{B}$, $|B| = |A|$ and $(\mathfrak{B}, a_\xi)_{\xi<\alpha} \cong (\mathfrak{B}, a'_\xi)_{\xi<\alpha}$.*

We leave the proof of this lemma to the reader with the following hint: define an elementary chain $\mathfrak{A}_n$, $n \in \omega$, by $\mathfrak{A}_0 = \mathfrak{A}$ and $\mathfrak{A}_{n+1}$ is an elementary extension of $\mathfrak{A}_n$ such that $(\mathfrak{A}_{n+1}, a)_{a\in A_n} \equiv (\mathfrak{A}_{n+1}, h(a))_{a\in A_n}$ for some map h with $h(a_\xi) = a'_\xi$ for all $\xi < \alpha$.

The existence result easily follows:

THEOREM 2.10. *Every model of cardinality $\leq 2^\kappa$ has a strongly κ^+-saturated elementary extension of cardinality $\leq 2^\kappa$.*

3. Skolem Functions and Indescernables

For any model $\mathfrak{A}$ and any $X \subseteq A$ there is a well-defined submodel of $\mathfrak{A}$ *generated* by X (provided either $X \neq \emptyset$ or the language contains individual constants), that is a submodel $\mathfrak{B} \subseteq \mathfrak{A}$ such that $X \subseteq B$ and $\mathfrak{B} \subseteq \mathfrak{B}'$ for every $\mathfrak{B}' \subseteq \mathfrak{A}$ with $X \subseteq B'$. This is not in general true for *elementary* submodels, however. It would be true if $\mathfrak{A}$ had the property that $\mathfrak{B} \subseteq \mathfrak{A}$ implies $\mathfrak{B} \prec \mathfrak{A}$, since then the submodel generated by X would in fact be the elementary submodel generated by X. In this section we show that any theory T of a language $\mathcal{L}$ can be expanded in a natural way to a theory T^* in a larger language $\mathcal{L}^*$ formed by adding "Skolem functions" such that all models of T^* have the above stated property.

DEFINITION 3.1. Let $\mathfrak{A}$ be a model for some language $\mathcal{L}$, and let $X \subseteq A$.

(a) $H(X)$, the *hull of X in* $\mathfrak{A}$, is the subset of A obtained by closing $X \cup \{c^{\mathfrak{A}} : c \text{ in } \mathfrak{L}\}$ under $F^{\mathfrak{A}}$ for every F of $\mathcal{L}$.
(b) $\mathfrak{H}(X)$ is the submodel of $\mathfrak{A}$ whose universe $= H(X)$, provided $H(X) \neq \emptyset$.

Note that the model $\mathfrak{H}(X)$ exists by the Lemma on page 74, when $H(X) \neq \emptyset$. The following Lemma is easily established.

LEMMA 3.1. *Given $\mathfrak{A}$ and $X \subseteq A$:*

(1) $|H(X)| \leq \max(|X|, |\mathcal{L}|)$.
(2) $\mathfrak{H}(X)$ *is the smallest* $\mathfrak{B} \subseteq \mathfrak{A}$ *with* $X \subseteq B$, *provided* $H(X) \neq \emptyset$.

Our goal is to establish the following:

Theorem 3.2. *Let T be a theory in a language $\mathcal{L}$. Then there is a theory T^* in a language $\mathcal{L}^*$ such that*

(i) $\mathcal{L} \subseteq \mathcal{L}^*$, $T \subseteq T^*$.
(ii) $|\mathcal{L}^*| = |\mathcal{L}|$
(iii) every model of T can be expanded to a model of T^*.
(iv) if $\mathfrak{A}^* \models T^*$ and $\mathfrak{B}^* \subseteq \mathfrak{A}^*$ then $\mathfrak{B}^* \prec \mathfrak{A}^*$.

The language $\mathcal{L}^*$ and theory T^* of this theorem will be defined as unions of chains $\{\mathcal{L}_n\}_{n\in\omega}$, $\{T_n\}_{n\in\omega}$ which "approximate" the property in (iv) more and more closely. We go from $\mathcal{L}_n$ and T_n to $\mathcal{L}_{n+1}$ and T_{n+1} by the following result.

Proposition 3.3. *Let T be a theory of $\mathcal{L}$. Then there is a theory T' in $\mathcal{L}'$ such that:*

(i) $\mathcal{L} \subseteq \mathcal{L}'$, $T \subseteq T'$
(ii) $|\mathcal{L}'| = |\mathcal{L}|$
(iii) every model of T can be expanded to a model of T'
(iv) if $\mathfrak{A}' \models T'$ and $\mathfrak{B}' \subseteq \mathfrak{A}'$ then $\mathfrak{B}' \restriction \mathcal{L} \prec \mathfrak{A} \restriction \mathcal{L}$.

Assuming the Proposition the Theorem easily follows: we define $\mathcal{L}_0 = \mathcal{L}$, $T_0 = T$, $\mathcal{L}_{n+1} = (\mathcal{L}_n)'$, $T_{n+1} = (T_n)'$, $\mathcal{L}^* = \bigcup_{n\in\omega} \mathcal{L}_n$ and $T^* = \bigcup_{n\in\omega} T_n$. (i)-(iv) of the Theorem follow from the corresponding properties of the Proposition, making use of the elementary fact that every formula of $\mathcal{L}^*$ will be a formula of some $\mathcal{L}_n$, $n \in \omega$, since it can contain just finitely many symbols. Thus, we now turn to the:

Proof of the Proposition. $\mathcal{L}'$ is the language which adds to $\mathcal{L}$ a new n-place function symbol $F_{\exists y\varphi}$ for every formula $\exists y\varphi(y, x_0, \ldots, x_{n-1})$ of $\mathcal{L}$ and a new constant symbol $c_{\exists y\varphi}$ for every sentence $\exists y\varphi(y)$ of $\mathcal{L}$.

We let Σ' consist of all $\mathcal{L}$-sentences of the following forms:

$$\exists y\varphi(y) \to \varphi(c_{\exists y\varphi})$$

$$\forall x_0 \cdots\cdots x_{n-1}[\exists y\varphi \to \varphi(F_{\exists y\varphi}(x_0, \ldots, x_{n-1}), x_0, \ldots, x_{n-1})],$$

for all formulas $\exists y\varphi$ of $\mathcal{L}$. Finally, $T' = \mathrm{Cn}_{\mathcal{L}'}(T \cup \Sigma')$. The requirements of the Proposition are easily verified, using the Axiom of Choice for (iii) and the Lemma on page 77 for (iv). □

The functions (and constants) added to $\mathcal{L}$ to obtain $\mathcal{L}'$ are called *Skolem functions* for $\mathcal{L}$, and the set Σ' is the set of *Skolem axioms* for $\mathcal{L}$. Thus the language $\mathcal{L}^*$ constructed in the proof of the Theorem has Skolem functions *for* $\mathcal{L}^*$, and $T^* = T \cup \Sigma^*$ where Σ^* is the set of Skolem axioms *for* $\mathcal{L}^*$.

A theory T' having the property that whenever $\mathfrak{A} \models T'$ and $\mathfrak{B} \subseteq \mathfrak{A}$ then $\mathfrak{B} \prec \mathfrak{A}$ is said to *have Skolem functions*, even if T' is not constructed explicitly by adding "Skolem functions" to some theory in a smaller language. If $\mathfrak{A}$ is a model of a theory having Skolem functions then $\mathfrak{H}(X) \prec \mathfrak{A}$ for any $X \subseteq A$ and $H(X)$ is called the *Skolem hull* of X in $\mathfrak{A}$. Also, $X \subseteq Y \subseteq A$ implies $\mathfrak{H}(X) \prec \mathfrak{H}(Y)$.

As a consequence of the existence of expansions with Skolem functions we note an improved Löwenheim-Skolem result.

Theorem 3.4. *Given any $\mathfrak{A}$, any $X \subseteq A$, and any infinite cardinal κ such that $\max(|\mathcal{L}|, |X|) \leq \kappa \leq |A|$ there is some $\mathfrak{B} \prec \mathfrak{A}$ such that $X \subseteq B$ and $|X| = \kappa$.*

PROOF. Let $T = \mathrm{Th}(\mathfrak{A})$ and let $\mathcal{L}^*, T^*$ be as in the Theorem, and let $\mathfrak{A}^*$ be an expansion of $\mathfrak{A}$ to a model of T^*. We may suppose $|X| = \kappa$. Thus $|H^*(X)| = \kappa$ (since $|\mathcal{L}^*| = |\mathcal{L}| \leq \kappa$), and so $\mathfrak{B} = H^*(X) \restriction \mathcal{L}$ as desired. □

In the remainder of this section we study Skolem hulls of *indiscernible* sets of elements, and derive several important consequences concerning models of arbitrary theories.

DEFINITION 3.2. Let $X \subseteq A$ and let $\leq$ be a linear order of X. Then X is a set of *indiscernibles* (with respect to $\leq$) for $\mathfrak{A}$ iff for all $a_1, \ldots, a_n, a'_1, \ldots, a'_n \in X$ such that $a_1 < a_2 < \cdots < a_n$ and $a'_1 < a'_2 < \cdots < a'_n$ we have $(\mathfrak{A}, a_1, \ldots, a_n) \equiv (\mathfrak{A}, a'_1, \ldots, a'_n)$.

Note that the ordering of X is *not* assumed to be one of the relations of $\mathfrak{A}$, or even definable in $\mathfrak{A}$. Also nothing is assumed about the relative positions in the ordering of a_i's and a'_i's. Thus, if $a_1 < a_2 < a_3 < a_4$, all in X, indiscernibles for $\mathfrak{A}$, we have, e.g.,

$$(\mathfrak{A}, a_1) \equiv (\mathfrak{A}, a_2) \equiv (\mathfrak{A}, a_3) \equiv (\mathfrak{A}, a_4)$$
$$(\mathfrak{A}, a_1, a_2) \equiv (\mathfrak{A}, a_1, a_3) \equiv (\mathfrak{A}, a_1, a_4) \equiv (\mathfrak{A}, a_2, a_3)$$

etc. However, one does *not* necessarily have $(\mathfrak{A}, a_1, a_2) \equiv (\mathfrak{A}, a_4, a_3)$.

Probably the simplest non-trivial example of an infinite set of indiscernibles is $X = A$ in any model $\mathfrak{A} \equiv (\mathbb{Q}, \leq)$, where (in this case, though not in general) the ordering on X is $\leq^{\mathfrak{A}}$. For another example, let $\mathfrak{L}$ have just a binary predicate E, and let T be the ω-categorical theory asserting that E is an equivalence relation with infinitely many equivalence classes, all of them infinite. Then in any model $\mathfrak{A}$ of T one can choose two radically different sets of indiscernibles, namely: X containing just elements from one equivalence class or Y containing only one element from each equivalence class. In either case the ordering is arbitrary. Note that $\mathfrak{A}_A \models \neg E\overline{a}_1\overline{a}_2$ for all $a_1, a_2 \in Y$ with $a_1 \neq a_2$.

The following, possibly surprising, result has many important consequences.

THEOREM 3.5. *Let T be a complete theory with infinite models. Let $(B, \leq)$ be any linear ordering. Then there is some model $\mathfrak{A}$ of T such that $B \subseteq A$ and B is a set of indiscernibles for $\mathfrak{A}$.*

PROOF. We expand the language $\mathcal{L}$ of $\mathfrak{A}$ to $\mathcal{L}(B) = \mathcal{L} \cup \{\overline{b} : b \in B\}$. Let Σ^* be the result of adding to T all sentences of $\mathcal{L}(B)$ of the form

$$(\varphi(\overline{b}_1, \ldots, \overline{b}_n) \leftrightarrow \varphi(\overline{b}'_1, \ldots, \overline{b}'_n))$$

where $\varphi(x_1, \ldots, x_n)$ is a formula of $\mathcal{L}$ and $b_1 < \cdots < b_n, b'_1 < \cdots < b'_n$, all from B. Any model of Σ^* will be a model of T with the required indiscernibles. We use compactness to show Σ^* has a model. Let $\Sigma_0 \subseteq \Sigma^*$ be finite. Then Σ_0 contains only finitely many of the added sentences. Thus there is a finite $B_0 \subseteq B$ and a finite set Φ_0 of $\mathcal{L}$-formulas such that Σ_0 is contained in T together iwth only the setences $(\varphi(\overline{b}_1, \ldots, \overline{b}_n) \leftrightarrow \varphi(\overline{b}'_1, \ldots, \overline{b}'_n))$ for $\varphi \in \Phi_0$ and $b_1 < \cdots < b_n, b'_1 < \cdots < b'_n$ all in B_0. We may assume (by adding dummy variables) that all formulas in Φ_0 have the same number of free variables, say n. Let $\mathfrak{A}$ be any model of T. We show how to interpret the constants in B_0 as elements of A so as to obtain a model of Σ_0. First let $\leq$ be any linear order of A. Let $a_2, \ldots, a_n, a'_1, \ldots, a'_n \in A$ where $a_1 < a_2 < \cdots < a_n$ and $a'_1 < a'_2 < \cdots < a'_n$. We define $\{a_1, \ldots, a_n\} \sim \{a'_1, \ldots, a'_n\}$ to hold iff $\mathfrak{A} \models [\varphi(\overline{a}_1, \ldots, \overline{a}_n) \leftrightarrow \varphi(\overline{a}'_1, \ldots, \overline{a}'_n)]$ for every $\varphi \in \Phi_0$. Then $\sim$ is an equivalence relation on the collection $A^{[n]}$ of n-element subsets of A. Further (since Φ_0 is finite), this relation divides $A^{[n]}$ into *finitely* many equivalence classes. By a classical result of combinatorics, "Ramsey's Theorem", there must (since A is infinite) be an infinite set $A_0 \subseteq A$ such that *all* n-element subsets of A_0 belong to the *same* equivalence class. Thus, if we interpret the (finitely many) elements of B_0, listed in increasing order, by any elements of A_0, also listed in increasing order, we obtain a model of Σ_0, as desired. $\square$

Note that a set of indiscernibles X in $\mathfrak{A}$ determines, for every $n \in \omega \setminus \{0\}$, a unique complete n-type $\Phi(x_1, \ldots, x_n)$ which is the complete type of any increasing sequence of n elements of X in $\mathfrak{A}$. If Y is a set of indiscernibles in $\mathfrak{B}$, then we say that X and Y have the *same type*, written $\mathrm{tp}_{\mathfrak{A}}(X) = \mathrm{tp}_{\mathfrak{B}}(Y)$, iff X and Y determine exactly the same complete n-types for every n. The following result is easily established, and thus is left to the reader:

THEOREM 3.6. *Let X be an infinite set of indiscernibles in $\mathfrak{A}$ and let $(Y, \leq)$ be any infinite linear ordering. Then there is some $\mathfrak{B}$ such that Y is a set of indiscernibles in $\mathfrak{B}$ and $\mathrm{tp}_{\mathfrak{A}}(X) = \mathrm{tp}_{\mathfrak{B}}(Y)$. [In particular, $\mathfrak{A} \equiv \mathfrak{B}$].*

In theories with Skolem functions, the Skolem hulls of sets of indiscernibles are particularly well-behaved. The following summarises some of their most important properties.

THEOREM 3.7. *Let T be a complete theory (of $\mathcal{L}$) with Skolem functions. Let $\mathfrak{A}, \mathfrak{B} \models T$ and suppose X, Y are infinite sets of indiscernibles in $\mathfrak{A}, \mathfrak{B}$ respectively (with respect to $\leq$). Further, suppose that $\mathrm{tp}_{\mathfrak{A}}(X) = \mathrm{tp}_{\mathfrak{B}}(Y)$.*

(1) $\mathfrak{H}(X)$ and $\mathfrak{H}(Y)$ are models of T realizing precisely the same n-types of $\mathcal{L}$, for all n.

(2) If h is an order-preserving map of X into Y then h extends to a unique isomorphism h^* of $\mathfrak{H}(X)$ onto some elementary submodel of $\mathfrak{H}(Y)$ — in fact, onto $\mathfrak{H}(\{h(a) : a \in X\})$.

(3) If h is an order-preserving map of X onto itself, then h extends to a unique automorphism of $\mathfrak{H}(X)$.

PROOF. First note that $H(X) = \{t^{\mathfrak{A}} : t \in \mathrm{Tm}_{\mathcal{L}}, a_1 < a_2 < \cdots < a_n \text{ all in } X\}$ and there is a corresponding representation for $H(Y)$. Thus, let us suppose that $\mathfrak{H}(X)$ realizes some 1-type $\Phi(x)$. This means that for some t and $a_1 < \cdots < a_n$ in X we have $\mathfrak{A}_x \models \varphi(t^{\mathfrak{A}}(\overline{a}_1, \ldots, \overline{a}_n))$ for every $\varphi \in \Phi$. That is, $\varphi(t(x_1, \ldots, x_n))$ belongs to the n-type of increasing sequences from X. Since we assumed that $\mathrm{tp}_{\mathfrak{A}}(X) = \mathrm{tp}_{\mathfrak{B}}(Y)$ it follows that $t^{\mathfrak{B}}(b_1, \ldots, b_n)$ realizes $\Phi(x)$ in $\mathfrak{H}(Y)$ for any $b_1 < b_2 < \cdots < b_n$ in Y. The same argument for types in more than one variable establishes (1).

In a similar way, the extension h^* needed in (2) can be seen to be the map taking $t^{\mathfrak{A}}(a_1, \ldots, a_n)$ to $t^{\mathfrak{B}}(h(a_1), \ldots, h(a_n))$, for $a_1 < \cdots < a_n$ in X.

(3) follows immediately from (2). $\square$

We can now establish the following very important result.

THEOREM 3.8. *Let T be a complete theory, in a countable language, having infinite models. Then there is a* countable *collection $\mathcal{T}$ of complete types (in all numbers of variables) such that for every $\kappa \geq \omega$ T has a model of power κ which realizes* precisely *the complete types in $\mathcal{T}$.*

Proof. We first choose an extension T^* of T in $\mathcal{L}^*$ as given by the Theorem on pg 120. Note that $\mathcal{L}^*$ is still countable. Let X be a countably infinite set of indiscernibles in some model $\mathfrak{A}^*$ of T^*. Let $\mathcal{T}$ be the collection of complete types of $\mathcal{L}$ realized in $\mathfrak{H}(X)\restriction\mathcal{L}$. Then $\mathcal{T}$ is countable since $H^*(X)$ is. Now given any $\kappa > \omega$ let $(Y, \leq)$ be any linear ordering of cardinality κ. By the Theorem on pg 125, there is some model $\mathfrak{B}^* \equiv \mathfrak{A}^*$ such that Y is a set of indiscernibles in $\mathfrak{B}^*$ and $\mathrm{tp}_{\mathfrak{B}^*}(Y) = \mathrm{tp}_{\mathfrak{A}^*}(X)$. By Theorem (1) on pg 126, $\mathfrak{H}^*(Y)$ realizes precisely the same complete $\mathcal{L}^*$-types as $\mathfrak{H}^*(X)$. In particular, $\mathfrak{H}^*(Y)\restriction\mathcal{L}$ is a model of T which realizes precisely the complete types in $\mathcal{T}$, and has power κ. $\square$

The reader should note that these are the first results which gives us any control over building uncountable models analogous to the control which the omitting types theorem yields over countable models. We are thus able to derive our first non-trivial result on the number of uncountable models of a theory.

Corollary 3.9. *Let T be a complete theory in a countable language, and assume there are 2^ω different complete types consistent with T. Then for every $\kappa \geq \omega$ T has at least 2^ω non-isomorphic models of cardinality κ.*

Another important application of indiscernibles is the following:

Theorem 3.10. *Let T be a complete theory, with infinite models, in a countable language. Then for every $\kappa \geq \omega$ T has a model of cardinality κ with 2^κ automorphisms.*

Proof. As in the preceding Theorem we first expand T to a theory T^* with Skolem functions. If $(X, \leq)$ is a linear ordering of power κ with 2^κ automorphisms, and X is indiscernible in $\mathfrak{B}^* \models T^*$, then $\mathfrak{H}^*(X)\restriction\mathcal{L}$ is a model of T of cardinality κ with 2^κ automorphisms, by Theorem (3) on page 126. $\square$

4. Some Applications

In this section we introduce no new techniques, but use previous material to derive some standard results, which are interesting both in themselves and as examples of applications. Various further results in similar lines are mentioned briefly.

Our first topic concerns results connecting the syntactical form of a sentence with properties of the class of models of the sentence. These are examples of *preservation theorems*.

Definition 4.1. (1) A sentence θ is a *universal* sentence (or, $\forall$-*sentence*) iff θ is $\forall x_0 \cdots \forall x_n \alpha$ where α is open (i.e., has no quantifiers).

(2) A set Σ os sentences if *preserved under substructures* iff for every $\mathfrak{A} \models \Sigma$, if $\mathfrak{B} \subseteq \mathfrak{A}$ then $\mathfrak{B} \models \Sigma$.

It is easily verified that any set of universal sentences is preserved under substructures. The theorem we are after states that the converse holds "up to equivalence". More precisely:

Theorem 4.1. *A theory T is preserved under substructures iff $T = \mathrm{Cn}(\Sigma)$ for some set Σ of universal sentences.*

We will derive this from a result stating when a model can be embedded in some model of T.

The following notation is useful:

DEFINITION 4.2. For any theory T we define $T_\forall = \{\forall\text{-sentences } \theta : T \models \theta\}$; in particular, we write $\mathrm{Th}_\forall(\mathfrak{A})$ for $(\mathrm{Th}(\mathfrak{A}))_\forall = \{\forall\text{-sentences } \theta : \mathfrak{A} \models \theta\}$.

The Theorem on the preceding page is an immediate consequence (with $\Sigma = T_\forall$) of the following result:

THEOREM 4.2. *$\mathfrak{B} \models T_\forall$ iff $\mathfrak{B} \subseteq \mathfrak{A}$ for some $\mathfrak{A} \models T$.*

PROOF. From right to left follows by our original remark that $\forall$-sentences are preserved under substructures. For the other direction, suppose $\mathfrak{B} \models T_\forall$. To show that $\mathfrak{B}$ can be embedded in some model of T is suffices (by the top Lemma on page 77) to show that $T \cup \Delta_{\mathfrak{B}}$ has a model. If not then, by compactness, $T \models \neg\alpha(\overline{b}_0, \ldots, \overline{b}_n)$ for some open $\mathcal{L}(B)$-sentence $\alpha(\overline{b}_0, \ldots, \overline{b}_n)$ true on $\mathfrak{B}$. But then $T \models \forall x_0 \cdots \forall x_n \neg\alpha(x_0, \ldots, x_n)$, so $\forall x_0 \cdots \forall x_n \neg\alpha(x_0, \ldots, x_n) \in T_\forall$, a contradiction to our hypothesis that $\mathfrak{B} \models T_\forall$. $\square$

As a consequence, using compactness, we have the following:

COROLLARY 4.3. *$\{\sigma\}$ is preserved under substructures iff $\sigma \vdash\dashv \theta$ for some universal sentence θ.*

The following is an important consequence of the above Theorem:

COROLLARY 4.4. *If $\mathfrak{B}_A \models \mathrm{Th}_\forall(\mathfrak{A}_A)$ then $\mathfrak{B} \subseteq \mathfrak{A}'$ for some $\mathfrak{A}'$ with $\mathfrak{A} \prec \mathfrak{A}'$ (where $\mathfrak{A} \subseteq \mathfrak{B}$).*

Note that a theory with Skolem functions is preserved under substructures, hence has a set of universal axioms. Similar results hold for preservation relative to another theory, where:

DEFINITION 4.3. Σ is *preserved under substructures relative to T* iff for every $\mathfrak{A} \models T \cup \Sigma$, if $\mathfrak{B} \subseteq \mathfrak{A}$ and $\mathfrak{B} \models T$ then $\mathfrak{B} \models \Sigma$.

Corresponding to the first Corollary on the preceding page we have, for example:

THEOREM 4.5. *$\{\sigma\}$ is preserved under substructures relative to T iff $T \models (\sigma \leftrightarrow \theta)$ for some $\forall$-sentence θ.*

We may also speak of a formula $\varphi(x_1, \ldots, x_n)$ being preserved, meaning that $\varphi(c_1, \ldots, c_n)$ is preserved where $c_1, \ldots, c_n$ are new constants. As a consequence of the preceding Theorem we thus obtain:

COROLLARY 4.6. *$\varphi(\overline{x})$ is preserved under substructures relative to T iff $T \models \forall\overline{x}[\varphi(\overline{x}) \leftrightarrow \theta(\overline{x})]$ for some* $\forall$-formula $\theta(\overline{x})$ *(where the notion of $\forall$-formula has the obvious definition).*

Many other varieties of preservation results are known, for example concerning the following:

DEFINITION 4.4. (1) A sentence θ is an *$\forall\exists$-sentence* iff θ is $\forall x_1 \cdots \forall x_n \exists y_1 \cdots \exists y_m \alpha(\overline{x}, \overline{y})$ where α is open.

(2) Σ is *preserved under unions of chains* iff whenever $\{\mathfrak{A}_i\}_{i \in I}$ is a chain, under $\subseteq$, of models of Σ then $\bigcup_{i \in I} \mathfrak{A}_i \models \Sigma$.

The corresponding preservation theorem is:

THEOREM 4.7. *T is preserved under unions of chains iff $T = \mathrm{Cn}(\Sigma)$ for some set Σ of $\forall\exists$-sentences.*

The proof of the above Theorem, and the statements and proofs of its consequences and variations, are left to the reader. We also do not mention the numerous other preservation results concerning sentences preserved under homomorphism, direct products, etc.

The next topic concerns theories for whose models substructure implies *elementary* substructure.

DEFINITION 4.5. T is *model complete* iff for all $\mathfrak{A}, \mathfrak{B} \models T$ if $\mathfrak{A} \subseteq \mathfrak{B}$ then $\mathfrak{A} \prec \mathfrak{B}$.

There are numerous natural examples of model complete theories, e.g. $\mathrm{Th}(\langle \mathbb{Q}, \leq \rangle)$; certainly if T has Skolem functions then T is model complete but the converse fails.

Note that a model complete theory need not be complete, but there is a simple sufficient condition for a model complete theory to be complete.

LEMMA 4.8. *Assume T is model complete and that there is some $\mathfrak{A} \models T$ such that $\mathfrak{A}$ can be embedded in every model of T. Then T is complete.*

PROOF. If $\mathfrak{B} \models T$ then $\mathfrak{A} \cong\prec \mathfrak{B}$, hence $\mathfrak{A} \cong\prec \mathfrak{B}$ by hypothesis, so $\mathfrak{B} \equiv \mathfrak{A}$. $\square$

Every theory can be expanded to a theory in a larger language which is model complete; more importantly, this expansion can be done in such a way (roughly, by adding only *definable* relations) that the new theory is "equivalent" to the old as far as most properties of their models are concerned. This was *not* the case with Skolem expansions.

DEFINITION 4.6. Given any language $\mathcal{L}$ we define $\mathcal{L}^{\#}$ as $\mathcal{L}$ together with a new n-ary predicate symbol R_φ for every formula $\varphi(x_1, \dots, x_n)$ of $\mathcal{L}$ with n free variables.

(1) $\Sigma^{\#} = \{\forall \overline{x}[\varphi(\overline{x}) \leftrightarrow R_\varphi(\overline{x})] : \varphi(\overline{x}) \text{ of } \mathcal{L}\}$.

LEMMA 4.9. *(1)* $|\mathcal{L}| = |\mathcal{L}^{\#}|$.
(2) Every $\mathcal{L}$-structure $\mathfrak{A}$ has exactly one *expansion to an $\mathcal{L}^{\#}$-structure $\mathfrak{A}^{\#} \models \Sigma^{\#}$.*
(3) $\mathrm{Cn}(\Sigma^{\#} \cup T)$ is model complete, for any theory T of $\mathcal{L}$.
(4) The following are equivalent: $\mathfrak{A} \prec \mathfrak{B}, \mathfrak{A}^{\#} \subseteq \mathfrak{B}^{\#}, \mathfrak{A}^{\#} \prec \mathfrak{B}^{\#}$.
(5) $\mathfrak{A} \cong \mathfrak{B}$ iff $\mathfrak{A}^{\#} \cong \mathfrak{B}^{\#}$.

We leave the proof of the Lemma to the reader. The only part requiring much argument is (3), for which one first proves (by induction) that for every formula $\psi(\overline{x})$ of $\mathcal{L}^{\#}$ there is a formula $\varphi(\overline{x})$ of $\mathcal{L}$ such that $\Sigma^{\#} \models (\psi \leftrightarrow \varphi)$, and hence $\Sigma^{\#} \models \forall \overline{x}[\psi(\overline{x}) \leftrightarrow R_\varphi(\overline{x})]$.

The Lemma implies that we can replace any T by the model complete theory $T^{\#} = \mathrm{Cn}(T \cup \Sigma^{\#})$ and not change, for example, the number of models in any power, on the relations of elementary embeddability between models.

There is a simple characterization of model complete theories which is useful in determining whether a specific theory is model complete.

THEOREM 4.10. *Given a theory T, the following are equivalent:*
(1) T is model complete,
(2) For any $\mathfrak{A}, \mathfrak{B} \models T$ if $\mathfrak{A} \subseteq \mathfrak{B}$ then $\mathfrak{B}_A \models \mathrm{Th}_\forall(\mathfrak{A}_A)$ — i.e., whenever $\varphi(x_0, \dots, x_{n-1})$ is a universal *formula of $\mathcal{L}$ and $\mathfrak{A}_A \models \varphi(\overline{a}_0, \dots, \overline{a}_{n-1})$ then $\mathfrak{B}_A \models \varphi(\overline{a}_0, \dots, \overline{a}_{n-1})$.*

PROOF. (1) clearly implies (2). For the converse, suppose (2) holds. We define two elementary chains $\{\mathfrak{A}_n\}_{n\in\omega}$ and $\{\mathfrak{B}_n\}_{n\in\omega}$ which "alternate" as follows:

Let $\mathfrak{A}_0 = \mathfrak{A}$ and $\mathfrak{B}_0 = \mathfrak{B}$, where $\mathfrak{A}, \mathfrak{B} \models T$ are given and $\mathfrak{A} \subseteq \mathfrak{B}$. By (2), $(\mathfrak{B}_0)_{A_0} \models \mathrm{Th}_\forall(\mathfrak{A}_{0A_0})$, hence there is some $\mathfrak{A}_1$ such that $\mathfrak{A}_0 \prec \mathfrak{A}_0$ and $\mathfrak{B}_0 \subseteq \mathfrak{A}_1$ — by the Corollary on the bottom of pg 130. Then, by (2) again, $(\mathfrak{A}_1)_{B_0} \models \mathrm{Th}_\forall(\mathfrak{B}_{0B_0})$, hence there is some $\mathfrak{B}_1$ such that $\mathfrak{B}_0 \prec \mathfrak{B}_1$ and $\mathfrak{A}_1 \subseteq \mathfrak{B}_1$. Continuing in this way we obtain elementary chains $\{\mathfrak{A}_n\}_{n\in\omega}$ and $\{\mathfrak{B}_n\}_{n\in\omega}$ such that $\mathfrak{A}_n \subseteq \mathfrak{B} \subseteq \mathfrak{A}_{n+1}$ for all n. Thus we have $\mathfrak{A}^* = \bigcup_{n\in\omega} \mathfrak{A}_n = \bigcup_{n\in\omega} \mathfrak{B}_n = \mathfrak{B}^*$. But $\mathfrak{A} \prec \mathfrak{A}^*$ and $\mathfrak{B} \prec \mathfrak{B}^*$, so we must have $\mathfrak{A} \prec \mathfrak{B}$, since $\mathfrak{A} \subseteq \mathfrak{B}$. □

In the next result, whose proof is left to the reader, we list two other conditions equivalent to being model complete — the first of which is the source of the name.

THEOREM 4.11. *The following are equivalent:*

(1) T is model complete,
(2) For any $\mathfrak{A} \models T$, $(T \cup \Delta_{\mathfrak{A}})$ is a complete theory in $\mathcal{L}(A)$,
(3) For every formula $\varphi(\overline{x})$ there is some universal formula $\theta(\overline{x})$ such that $T \models (\varphi \leftrightarrow \theta)$.

Since a model complete theory is preserved under unions of chains, it must have a set of $\forall\exists$ axioms, by the Theorem on the bottom of page 131. Easy examples show that not all such theories are model complete, but the following result gives an interesting sufficient condition for model completeness.

THEOREM 4.12. *Assume T is preserved under unions of chains and that all uncountable models of T are ω-saturated. Then T is model complete.*

It is known that if T, in a countable $\mathcal{L}$, is κ-categorical for some $\kappa \geq \omega$ then all uncountable models of T are ω-saturated. The converse has been conjectured — although this has been established for several important classes of theories, the question has still not been completely settled.

This last theorem can be used to show that the theory of algebraically closed fields (in $\{+, \cdot, 0, 1\}$) is model complete.

Finally, we study the class of models of $\mathrm{Th}(\langle \omega, \leq, 0, S\rangle)$, where s is the immediate successor function on ω. This theory is an expansion by definitions of $\mathrm{Th}(\langle \omega, \leq\rangle)$ so all our results translate over to the theory in the smaller language, but we leave this translation to the reader.

We will exhibit a finite set of axioms for $\mathrm{Th}(\langle \omega, 0, s\rangle)$ and show this theory is model complete. This will enable us to determine the number of models in each cardinality of this theory and determine how they can be related under elementary embedding.

Let Σ^ω be the set of sentences (in $\leq, \overline{0}, s$) saying that $\leq$ is a linear order, $\overline{0}$ is the first element in the order, sx is the immediate successor in the order of x for any x, and $\overline{0}$ is the only element which does not have an immediate successor.

Then every model of Σ^ω begins with a copy of ω and is then followed by some number of copies of $\mathbb{Z}$, linearly ordered in any way, and $\overline{0}, s$ have their intended interpretations. Conversely, any such structure is a model of Σ^ω. Note the following:

LEMMA 4.13. *Assume $\mathfrak{A} \models \Sigma^\omega$, $a_1, a_2 \in A$ and $\{a \in A : a_1 \leq^{\mathfrak{A}} a \leq^{\mathfrak{A}} \leq a_2\}$ is infinite. Then $\mathfrak{A} \prec \mathfrak{B}$ for some $\mathfrak{B}$ such that $\{b \in B \smallsetminus A : a_1 \leq^{\mathfrak{B}} b \leq^{\mathfrak{B}} a_2\}$ is infinite.*

Proof. $\mathfrak{B}$ is any elementary extension of $\mathfrak{A}$ realizing the type which says $\overline{a}_1 < x, x < \overline{a}_2$, and there are infinitely many elements between $\overline{a}_1$ and x and between x and $\overline{a}_2$. $\square$

We can now show that Σ^ω is the desired set of axioms.

Theorem 4.14. *(1)* $\mathrm{Cn}(\Sigma^\omega)$ *is model complete.*
(2) $\mathrm{Cn}(\Sigma^\omega) = \mathrm{Th}(\langle\omega, \leq, 0, s\rangle)$.

Proof. (1) Let $\mathfrak{A}, \mathfrak{B} \models \Sigma^\omega$, $\mathfrak{A} \subseteq \mathfrak{B}$. By the Theorem on page 134 it suffices to show that $\mathfrak{B}_A \models \mathrm{Th}_\forall(\mathfrak{A}_A)$. So let $\alpha(x_0, \ldots, x_{n-1}, y_0, \ldots, y_{m-1})$ be an open formula of $\mathcal{L}$ and let $a_0, \ldots, a_{n-1} \in A$. We need to show that if $\mathfrak{B}_A \models \exists\overline{y}\neg\alpha(\overline{a}_0, \ldots, \overline{a}_{n-1}, \overline{y})$ then also $\mathfrak{A}_A \models \exists\overline{y}\neg\alpha(\overline{a}_0, \ldots, \overline{a}_{n-1}, \overline{y})$. So, let $b_0, \ldots, b_{m-1} \in B$ be such that $\mathfrak{B}_B \models \neg\alpha(\overline{a}_0, \ldots, \overline{a}_{n-1}, \overline{b}_0, \ldots, \overline{b}_{m-1})$. If $b_i \notin A$ then b must be "infinitely far" (in the sense of $\leq^{\mathfrak{B}}$) from every element of A. By repeated application of the preceding Lemma we obtain $\mathfrak{A} \prec \mathfrak{A}'$ which has added elements in every "gap" inhabited by $b_0, \ldots, b_{m-1}$. Thus, there are elements $a'_0, \ldots, a'_{m-1}$ in A' such that $\mathfrak{A}'_{A'} \models \neg\alpha(\overline{a}_0, \ldots, \overline{a}_{n-1}, \overline{a}'_0, \ldots, \overline{a}'_{m-1})$, since all α can specify is the ordering and distances (finite) between elements. Thus in particular, $\mathfrak{A}_A \models \exists\overline{y}\neg\alpha(\overline{a}_0, \ldots, \overline{a}_{n-1}, \overline{b})$ as desired.

(2) follows from the Lemma on page 132, using $\mathfrak{A} = \langle\omega, \leq, 0, 2\rangle$. $\square$

What this tells us is the following:

(a) every linear ordering (including the empty order) determines a model of $\mathrm{Th}(\langle, \leq, 0, s\rangle)$ in which the copies of $\mathbb{Z}$ are so ordered;
(b) non-isomorphic linear orderings determine non-isomorphic models of $\mathrm{Th}(\langle\omega, \leq, 0, s\rangle)$;
(c) one model of $\mathrm{Th}(\langle\omega, \leq, 0, s\rangle)$ can be elementarily embedded in another iff the linear ordering of the $\mathbb{Z}$ blocks in the first can be (isomorphically) embedded in the corresponding linear ordering in the second.

Thus, by looking at appropriate linear orderings we derive results about the models of $\mathrm{Th}(\langle\omega, 0, s\rangle)$ under elementary embedding. The main results we obtain are as follows:

(1) for every $\kappa \geq \omega$, $\mathrm{Th}(\langle\omega, \leq, 0, s\rangle)$ has 2^κ non-isomorphic models of cardinality κ;
(2) $\mathrm{Th}(\langle\omega, \leq, 0, s\rangle)$ has 2^ω non-isomorphic countable ω_1-universal models;
(3) $\mathrm{Th}(\langle\omega, \leq, 0, s\rangle)$ has two countable models, realizing precisely the same types, neither of which can be elementarily embedded in the other.

Note that it follows from [(1) and] (2) that $\mathrm{Th}(\langle\omega, \leq, 0, s\rangle)$ has 2^ω non-isomorphic countable models despite the fact that there are only countably many complete types consistent with $\mathrm{Th}(\langle\omega, \leq, 0, s\rangle)$.

5. Exercises

(1) Let $\mathcal{L}^{nl} = \{E\}$ where E is a binary relation symbol. Let $\mathfrak{A}$ be the countable $\mathcal{L}$-structure in which $E^{\mathfrak{A}}$ is an equivalence relation on A with exactly one $E^{\mathfrak{A}}$ class of size n for every positive integer n but with no infinite $E^{\mathfrak{A}}$-classes. Define $T = \mathrm{Th}(\mathfrak{A})$.
 (a) Prove or disprove: T is model complete.

(b) Prove that $\mathfrak{A}$ is a prime model of T.
(c) Prove that T is not ω-categorical.

APPENDIX A

Appendix A: Set Theory

The natural numbers (that is, non-negative integers) are used in two very different ways. The first way is to *count* the number of elements in a (finite) set. The second way is to *order* the elements in a set–in this way one can prove things about all the elements in the set *by induction.* These two roles of natural numbers both are generalized to infinite numbers, but these are split into two groups according to the function they perform: cardinal numbers (to count) and ordinal numbers (to order). The basic facts and concepts are surveyed in this Appendix.

1. Cardinals and Counting

It was noted as long ago as Galileo that (some) infinite sets can be put into one-to-one correspondence with proper subsets of themselves, and thus, for example, that the set of integers may be considered as having the same "size" as the set of even integers. However no serious investigation into the "sizes" of infinite sets, on comparing them, was undertaken until Cantor, in the second half of the 19^{th} century, created set theory, including both cardinal and ordinal numbers.

The basic definitions about comparing the sizes of sets (including finite) are as follows:

Definition 1.1. (i) $X \sim Y$ (X and Y are *equivalent*, or have the "same number" of elements) iff there is some function mapping X one-to-one onto Y. (ii) $X \preceq Y$ (X has "at most as many" elements as Y) iff there is some function mapping X one-to-one into Y. (iii) $X \prec Y$ (X has strictly "fewer" elements than Y) iff $X \preceq Y$ but not $X \sim Y$.

The following proposition is certainly essential if one is to think of $\preceq$ as some sort of ordering on sets. It was not proved, however, until the end of the 19^{th} century.

Proposition 1.1. *If $X \preceq Y$ and $Y \preceq X$ then $X \sim Y$.*

All the basic facts about size comparisons could be expressed by the above notation but this would be quite clumsy. Instead certain sets, called *cardinal numbers*, are picked out so that for every set X there is *exactly one* cardinal number κ such that $X \sim \kappa$. We then call κ the *cardinality* of X and write $|X| = \kappa$. $|X| \leq |Y|$ means $X \preceq Y$ and $|X| < |Y|$ means $X \prec Y$. Notice that $|\kappa| = \kappa$ if, and only if, κ is a cardinal number.

The first cardinal numbers are defined as follows:

$$\begin{aligned} 0 &= \emptyset, \\ 1 &= \{0\}, \\ 2 &= 1 \cup \{1\}, \\ &\vdots \\ n+1 &= n \cup \{n+1\} \\ &\vdots \end{aligned}$$

$\omega = \{0, 1, \ldots, n, \ldots\}$ is defined as the smallest set containing 0 and such that if $x \in \omega$ then $x \cup \{x\} \in \omega$.

Notice that $\omega + 1 = \omega \cup \{\omega\}$ cannot also be a cardinal number since $\omega \sim \omega \cup \{\omega\}$.

Definition 1.2. (a) X is *finite* iff $X \sim n$ for some $n \in \omega$ ($|X| \in \omega$). (b) X is *countable* iff $X \preceq \omega$ (i.e. $|X| \leq \omega$).

Lemma 1.2. *(i) X is finite iff $|X| < \omega$. (ii) X is countable and infinite iff $|X| = \omega$.*

[The essential content of this lemma is that all cardinals less than ω in fact belong to ω].

One of Cantor's fundamental discoveries is that there are infinite sets which are not equivalent, and in fact that there can be no biggest cardinal number.

Definition 1.3. The *power set* of X, is defined by $\mathcal{P}(X) = \{Y |\ Y \subseteq X\}$.

Theorem 1.3. *For ever X, $X \prec \mathcal{P}(X)$.*

Proof. Obviously, $X \preceq \mathcal{P}(X)$. Suppose that $X \sim \mathcal{P}(X)$, say that h maps X bijectively onto $\mathcal{P}(X)$. Let $D = \{x \in X |\ x \notin h(x)\}$. Then $D = h(d)$ for some $d \in X$. But $d \in D$ iff $d \notin h(d) = D$, which is a contradiction. $\square$

Thus, there must be cardinals κ such that $\omega < \kappa$. We must put off defining them, however, until after we introduce ordinal numbers–we also should admit that we will need the Axiom of Choice (AC) to define cardinal numbers in general. Recall that AC states that if X is a set of non-empty sets, then there is a function f defined on X such that $f(x) \in x$ for every $x \in X$.

It is important to know that many set-theoretic operations lead from countable sets to countable sets.

Definition 1.4. (a) Y_X is the set of all functions f with domain Y and range a subset of X. (b) $\overset{\omega}{\smile}X = \bigcup_{n \in \omega} n_X$ is the set of all finite sequences of elements of X (thinking of a sequence of length n as a function defined on n); an alternate notation is ${}^{\omega>}X$.

Theorem 1.4. *If X is countable then so is $\overset{\omega}{\smile}X$.*

Proof. It suffices to show that $\overset{\omega}{\smile}\omega \preceq \omega$, which follows by using the one-to-one map which sends $(k_0, \ldots, k_{n-1})$ to $2^{k_0+1} \cdot 3^{k_1+1} \cdot \cdots \cdot p_{n-1}^{k_{n-1}+1}$ where p_j is the j^{th} odd prime. $\square$

Corollary 1.5. *If X, Y are countable so are $X \cup Y$ and $X \times Y$.*

THEOREM 1.6. *(AC) If X_n is countable for every $n \in \omega$ then $\bigcup_{n\in\omega} X_n$ is also countable.*

ω is, in fact, the the *smallest* infinite cardinal, although the proof requires the axiom of choice.

PROPOSITION 1.7. *(AC) If X is infinite then $\omega \preceq X$.*

The analogues of $+$ and $\cdot$ trivialize on infinite cardinals because of the preceding corollary, but exponentiation is important.

NOTATION 3. If κ, λ are cardinal numbers then κ^λ is the cardinal $|\lambda_\kappa|$.

LEMMA 1.8. *For any X, $\mathcal{P}(X) \sim X_2$.*

Hence from Cantor's Theorem we see the following corollary.

COROLLARY 1.9. *If $|X| = \kappa$ then $|\mathcal{P}(X)| = 2^\kappa$, and so $\kappa < 2^\kappa$ for every cardinal κ.*

However, increasing the base does not yield still larger cardinals.

LEMMA 1.10. *$2^\omega = n^\omega = \omega^\omega = (2^\omega)^\omega$, any $n \in \omega$.*

PROOF. It suffices to show $(2^\omega)^\omega \leq 2^\omega$, which follows since

$$\omega_{(\omega_2)} \sim (\omega \times \omega)_2 \sim \omega_2.$$

$\square$

Without proof we list some facts about (uncountable) cardinalities, all depending on AC.

(1) If X is infinite then $|X| = |\overset{\omega}{\smile}X|$.
(2) If X, Y are infinite then $|X \cup Y| = |X \times Y| = max(|X|, |Y|)$.
(3) If $|I| \leq \kappa$ and $|X_i| \leq \kappa$ for all $i \in I$, then $|\bigcup_{i\in I} X_i| \leq \kappa$, for $\kappa \geq \omega$.
(4) $(\mu^\kappa)^\lambda = \mu^{max(\kappa,\lambda)}$, for $\kappa, \lambda \geq \omega$, $\mu \geq 2$.
(5) For any cardinal κ there is a unique *next* cardinal called κ^+, but there is no set X such that $\kappa \preceq X \preceq \kappa^+$.
(6) If X is a non-empty set of cardinal numbers, then $\bigcup X$ is a cardinal number and it is the first cardinal $\leq$ all cardinals in X.
(7) $(\kappa^+)^\lambda = max(\kappa^\lambda, \kappa^+)$ for $\kappa, \lambda \geq \omega$.
(8) For any sets X, Y either $X \preceq Y$ or $Y \preceq X$, hence for any cardinals κ, λ either $\kappa \leq \lambda$ or $\lambda \leq \kappa$.

Some notation, based on (5), is the following which we will extend in the next section: $\omega_1 = \omega^+, \omega_{n+1} = \omega_n^+, \omega_\omega = \bigcup_{n\in\omega} \omega_n$–writing also $\omega_0 = \omega$. An alternate notation is to use the Hebrew letter "aleph"–thus $\aleph_0, \aleph_1, \ldots, \aleph_\omega, \ldots$.

Note that $\omega_1 \leq 2^\omega$ and, in general, $\kappa^+ \leq 2^\kappa$ for each $\kappa \geq \omega$. It is natural to enquire about whether equality holds or not.

CONJECTURE 1.1. **Continuum Hypothesis** (CH): $2^\omega = \omega_1$

CONJECTURE 1.2. **Generalized Continuum Hypothesis** (GCH): For every infinite cardinal κ, $2^\kappa = \kappa^+$.

CH and GCH are consistent with, but independent of, the usual axiioms of set theory. In fact, each of the following is consistent with the usual axioms:

$$2^\omega = \omega_1, 2^\omega = \omega_2, 2^\omega = \omega_n \quad \text{for any} n \in \omega$$
$$2^\omega = \omega_\omega)+, 2^\omega = (\omega_\omega)^{++}, \ldots$$

We can, however, prove that $2^\omega \neq \omega_\omega$ since we have $(\omega_\omega)^\omega > \omega_\omega$.

Some further facts about cardinals will be presented at the end of the next section.

2. Ordinals and Induction

The principles of proof by induction and definition by recursion on the natural numbers are consequences just of the fact that ω is well-ordered by the usual order.

DEFINITION 2.1. $(X, \leq)$ is a well-ordering iff $\leq$ is a linear order of X and every non-empty subset $Y \subseteq X$ contains a least element, i.e. there is some $a_0 \in Y$ such that $a_0 \leq a$ for all $a \in Y$.

THEOREM 2.1. *(Proof by Induction) Let $(X, \leq)$ be a well-ordering. Let $A \subseteq X$ have the property that for every $a \in X$, if $b \in A$ for all $b < a$ then $a \in A$. Then $A = X$.*

PROOF. If not, consider $Y = X - A$ and obtain a contradiction to the definition. □

The way this is used if one wants to prove that all elements of X have property P is to let A be the set of all elements of X having property P.

In a similar vein, we see:

THEOREM 2.2. *(Definition by Recursion) Let $(X, \leq)$ be a well-ordering. Let Y be any non-empty set and let g be a function from $\mathcal{P}(Y)$ into Y. Then there is a unique function f from X into Y such that for every $a \in X$,*

$$f(a) = g(\{f(x) |\ x \in X, x < a\}).$$

[Less formally, this just says that $f(a)$ is defined in terms of the $f(x)$'s for $x < a$.]

As in the previous section, we wish to pick out particular well-orderings, called *ordinal numbers*, such that each well-ordering is isomorphic to *exactly one* ordinal number. We do this so that the well order $\leq$ of the ordinal is as natural as possible–that is, is give by $\in$. The precise definition we obtain is as follows:

DEFINITION 2.2. A set X is an ordinal number iff (i) $x \in y \in X \Rightarrow x \in X$ (equivalently, $y \in X \Rightarrow y \subseteq X$), and (ii) X is well-ordered by the relation $\leq$ defined by: $a \leq b$ iff $a \in b$ or $a = b$.

Condition (i) is frequently expressed by saying "X is transitive" and condition (ii) is loosely expressed by saying "$\in$ well-orders X." Note that technically X is not a well-ordering, but $(X, \leq)$ is–however condition (ii) determines $\leq$ completely from X. Notice, of course, that most sets aren't even *linearly* ordered by $\in$–in fact, one of the usual (but somewhat technical) axioms of set theory implies that if X is linearly ordered by $\in$, then in fact it is well-ordered by $\in$. Thus the conditions in (ii) could be expanded to read: (ii)*: $x \in y, y \in z, z \in X \Rightarrow x \in z$, $x, y \in X \Rightarrow x = y \vee x \in y \vee y \in x$. $(x \notin x)$ follows by the usual axioms.

Notice that the finite cardinal numbers and ω, as defined in the previous section, are also ordinal numbers. The following lemma gives some of the basic properties of ordinals. By convention, we normally use Greek letters $\alpha, \beta, \ldots$ to stand for ordinals.

LEMMA 2.3. *(1) If α is an ordinal and $x \in \alpha$ then x is an ordinal. (2) If α, β are ordinals then either $\alpha \in \beta$ or $\alpha = \beta$ or $\beta \in \alpha$. (3) If α is an ordinal then $\alpha + 1 = \alpha \cup \{\alpha\}$ is an ordinal. (4) If X is a set of ordinals then $\bigcup X$ is an ordinal.*

NOTATION 4. If α, β are ordinals we write $\alpha < \beta$ for $\alpha \in \beta$. Part (1) of the lemma states that if α is an ordinal then $\alpha = \{\beta |\ \beta \text{ is an ordinal, } \beta < \alpha\}$. The ordinal $\alpha + 1$ is the *immediate successor* of α–that is, $\alpha < \alpha + 1$ and there is no ordinal β such that $\alpha < \beta < \alpha + 1$. Similarly, $\bigcup X$ is the least upper bound of the set X of ordinals.

The class of all ordinals is not a set, but we can still think of it as well-ordered by $\leq$. Further, we can prove things about the class of all ordinals by induction, and define functions on ordinals by recursion.

Finally we note that ordinals do have the property for which we introduced them.

THEOREM 2.4. *Let $(X, \leq)$ be a well-ordering. Then there is exactly one ordinal α such that $(X, \leq) \cong (\alpha, \leq)$.*

We distinguish between two types of non-zero ordinals as follows:

DEFINITION 2.3. α is a successor ordinal iff $\alpha = \beta + 1$ for some ordinal β; α is a limit ordinal iff $\alpha \neq 0$ and α is not a successor ordinal.

Note that α is a limit ordinal iff $\alpha \neq 0$ and $\bigcup \alpha = \alpha$. If X is any non-empty set of ordinals not containing a largest ordinal, then $\bigcup X$ is a limit ordinal.

It is frequently more convenient to break proofs by induction, or definitions by recursion, into cases according to whether an ordinal is a successor or a limit ordinal. For example, the recursive definition of ordinal addition is as follows:

if $\beta = 0$ then $\alpha + \beta = \alpha$,
if $\beta = \gamma + 1$ then $\alpha + \beta = (\alpha + \gamma) + 1$,
if β is a limit then $\alpha + \beta = \bigcup \{\alpha + \gamma |\ \gamma < \beta\}$.

While most linear orderings $(X, \leq)$ are not well-orderings, there is no restriction on the sets X in well-orderings, by the next theorem. This means that proof by induction can (in principle) be applied to any set.

THEOREM 2.5. *(AC) For every set X there is some $\leq$ which well-orders X.*

As an immediate consequence of the two preceeding theorems we have:

COROLLARY 2.6. *(AC) For every set X there is some ordinal α such that $X \sim \alpha$.*

The ordinal α is not unique unless $\alpha < \omega$, since if $\omega \leq \alpha$ then $\alpha \sim \alpha + 1$, but the *least* such ordinal will be the *cardinality* of X.

DEFINITION 2.4. κ is a cardinal number iff κ is an ordinal number and for every $\alpha < \kappa$ we have $\alpha \prec \kappa$ (equivalently, for every ordinal α such that $\alpha \sim \kappa$ we have $\kappa \leq \alpha$).

This fills the lacuna in the preceding section. Note that the cardinal numbers are well-ordered by $\leq$, and $<$ is $\in$ on cardinal numbers.

The way we will customarily use proof by induction on an arbitrary set X is as follows: let $|X| = \kappa$ so there is some one-to-one function h mapping κ onto X. Write x_α for $h(\alpha)$. Then $X = \{x_\alpha | \ \alpha < \kappa\}$ and we prove what we want about x_α by induction on $\alpha < \kappa$. Note that for each $\alpha < \kappa$ we have $|\{x_\beta| \ \beta < \alpha\}| = \alpha < \kappa$.

The class of infinite cardinals can be indexed by the class of ordinals by using the following definition by recursion:

$\omega(0) = \omega$,
$\omega(\gamma + 1) = (\omega(\gamma))^+$,
β a limit $\Rightarrow \omega(\beta) = \bigcup\{\omega(\gamma)| \ \gamma < \beta\}$.

We normally write ω_γ instead of $\omega(\gamma)$.

We finally need to introduce the concept of cofinality in order to make the important distinction between regular and singular cardinals.

Definition 2.5. Let α, β be limit ordinals. Then α is cofinal in β iff there is a strictly increasing function $f \in \alpha_\beta$ such that $\bigcup\{f(\gamma)| \ \gamma < \alpha\} = \beta$.

Certainly β is confinal in β. ω is cofinal in every countable limit ordinal, but ω is not cofinal in ω_1.

Definition 2.6. Let β be a limit ordinal. Then the cofinality of β is $cf(\beta)$ equals the least α such that α is confinal in β.

Lemma 2.7. *For any limit ordinal β, $cf(\beta) \leq \beta$ and $cf(\beta)$ is a cardinal.*

Definition 2.7. Let κ be an infinite cardinal. Then κ is regular iff $\kappa = cf(\kappa)$. κ is singular iff $cf(\kappa) < \kappa$.

Definition 2.8. κ is a successor cardinal iff $\kappa = \lambda^+$ for some cardinal λ, i.e. $\kappa = \omega_{\beta+1}$ for some β.

Definition 2.9. κ is a limit cardinal iff $\kappa \geq \omega$ and κ is not a successor cardinal, i.e., $\kappa = \omega_\alpha$ for some limit ordinal α.

The division of infinite cardinals into regular and singular is almost the same as the division into successor and limit.

Theorem 2.8. *(1) Every successor cardinal is regular. (2) if $\kappa = \omega_\alpha$ is a limit cardinal, then $cf(\kappa) = cf(\alpha)$–hence if κ is regular then $\kappa = \omega_\kappa$.*

Regular limit cardinals are called *inaccessible cardinals*–their existence cannot be proved from the usual axioms of set theory.

With cofinalities we can state a few more laws of cardinal computation, continuing the list from the previous section.

(9) $\kappa^{cf(\kappa)} > \kappa$ for every cardinal $\kappa \geq \omega$.
(10) Assume that $|I| < cf(\kappa)$ and for every $i \in I$, $|X_i| < \kappa$.
Then $|\bigcup_{i \in I} X_i| < \kappa$.

It is frequently tempting to assume GCH because it simplifies many computations, e.g.: Assuming GCH we have, for any cardinals $\kappa, \lambda \geq \omega$, $\kappa^\lambda = \kappa$ if $\lambda < cf(\kappa)$, $\kappa^\lambda = \kappa^+$ if $cf(\kappa) \leq \lambda \leq \kappa$, $\kappa^\lambda = \lambda^+$ if $\kappa \leq \lambda$.

Appendix B: Notes on Validities and Logical Consequence

1. Some Useful Validities of Sentential Logic

1) Excluded Middle
$\models \phi \vee \neg\phi$
$\models \neg(\phi \wedge \neg\phi)$

2) Modus Ponens
$\phi, \phi \to \psi \models \psi$

3) Conjunction
$\phi, \psi \models \phi \wedge \psi$

4) Transitivity of Implication
$\phi \to \psi, \psi \to \theta \models \phi \to \theta$

5) Plain Ol' True as Day
$\phi \wedge \psi \models \phi$
$\phi \models \phi \vee \psi$
$\phi \to (\psi \to \theta), \phi \to \psi \models \phi \to \theta$
$\phi \models \psi \to \phi$
$\neg\psi \models \psi \to \phi$
$\phi \vdash\dashv \neg\neg\phi$

6) Proof by Contradiction
$\neg\phi \to (\psi \wedge \neg\psi) \models \phi$
$\neg\phi \to \phi \models \phi$

7) Proof by Cases
$\phi \to \psi, \theta \to \psi, \models (\phi \vee \theta) \to \psi$
$\phi \to \psi, \neg\phi \to \psi \models \psi$

8) De Morgan's Laws
$\neg(\phi \vee \psi) \vdash\dashv \neg\phi \wedge \neg\psi$
$\neg(\phi \wedge \psi) \vdash\dashv \neg\phi \vee \neg\psi$

9) Distributive Laws
$\phi \wedge (\psi \vee \theta) \vdash\dashv (\phi \wedge \psi) \vee (\phi \wedge \theta)$

$\phi \vee (\psi \wedge \theta) \vdash\dashv (\phi \vee \psi) \wedge (\phi \vee \theta)$

10) Contraposition
$\phi \to \psi \vdash\dashv \neg\psi \to \neg\phi$

11) The connectives $\wedge$ and $\vee$ are both commutative and associative.

2. Some Facts About Logical Consequence

1) $\Sigma \cup \{\phi\} \models \psi$ iff $\Sigma \models \phi \to \psi$

2) If $\Sigma \models \phi_i$ for each $i = 1, \ldots, n$ and $\{\phi_1, \ldots, \phi_n\} \models \psi$ then $\Sigma \models \psi$.

3) $\Sigma \models \phi$ iff $\Sigma \cup \{\neg\phi\}$ is not satisfiable.

Appendix C: Gothic Alphabet

a	$\mathfrak{A}\mathfrak{a}$	b	$\mathfrak{B}\mathfrak{b}$	c	$\mathfrak{C}\mathfrak{c}$	d	$\mathfrak{D}\mathfrak{d}$
e	$\mathfrak{E}\mathfrak{e}$	f	$\mathfrak{F}\mathfrak{f}$	g	$\mathfrak{G}\mathfrak{g}$	h	$\mathfrak{H}\mathfrak{h}$
i	$\mathfrak{I}\mathfrak{i}$	j	$\mathfrak{J}\mathfrak{j}$	k	$\mathfrak{K}\mathfrak{k}$	l	$\mathfrak{L}\mathfrak{l}$
m	$\mathfrak{M}\mathfrak{m}$	n	$\mathfrak{N}\mathfrak{n}$	o	$\mathfrak{O}\mathfrak{o}$	p	$\mathfrak{P}\mathfrak{p}$
q	$\mathfrak{Q}\mathfrak{q}$	r	$\mathfrak{R}\mathfrak{r}$	s	$\mathfrak{S}\mathfrak{s}$	t	$\mathfrak{T}\mathfrak{t}$
u	$\mathfrak{U}\mathfrak{u}$	v	$\mathfrak{V}\mathfrak{v}$	w	$\mathfrak{W}\mathfrak{w}$	x	$\mathfrak{X}\mathfrak{x}$
y	$\mathfrak{Y}\mathfrak{y}$	z	$\mathfrak{Z}\mathfrak{z}$				

Bibliography

Index

ω-homogeneous, 77, 78
ω-saturated, 76, 77
Łoš-Vaught, 70

atom, 74
atomic, 74, 75
axiom, 19, 37, 43
Axiom of Choice, 101

basic
  formula, 67

cardinal numbers, 97
cardinality, 97
categoricity, 57
categorictiy, 70
chain, 69
  union of, 69
  elementary, 69
cofinality, 102
compactness, 17, 54
complete, 17, 50
  formula, 74
  type, 74
completeness, 63
connective, 7
connectives, 24
consequence, 43
consistent, 19, 40, 62
Continuum Hypothesis, 99
countable, 98
countably universal, 75, 77

De Morgan's Laws, 103
deducible, 19
  logically, 38
deduction, 37
Deduction Theorem, 20
Definition by Recurion, 100
diagram
  basic, 67
  elementary, 67
Distributive Laws, 103

elementarily equivalent, 43
embeddable, 67
  elementarily, 67
equality, 24
Excluded Middle, 103
expression, 8
extension, 66
  elementary, 66

finite, 98
finitely satisfiable, 18
finiteness, 41, 54
formula, 67
  atomic, 25, 67, 70

generalization, 37, 39
Generalized Continuum Hypothesis, 99

Henkin set, 51
Henkin Sets, 63

inaccessible cardinals, 102
inconsistent, 40
independent, 20
induction, 100
  proof by, 100
interpretation, 11
isomorphism, 44

Löwenheim-Skolem, 54
length, 8, 51
locally realize, 74
logical consequence, 14, 35
logical equivalence, 15, 36

model, 30, 34, 43, 49
  canonical, 63
Modus Ponens, 103
modus ponens, 19, 38

normal form
  conjunctive, 16
  disjunctive, 16

omit, 61
  locally, 63
omitting types, 74

ordinals, 100

power set, 98
prime, 74
principal, 74

quantifiers, 24

realize, 61
  locally, 63
recursion, 70
  definition by, 12, 18, 64, 100
reduct, 29
regular cardinal, 102
rule, 37

satisfiability, 13, 31, 33, 34, 67
  finite, 62
sentence, 8, 24
singular cardinal, 102
soundness, 38
structure, 28
  canonical, 50
submodel, 66
  elementary, 66
substitution, 32
successor, 26, 101
symbol, 8, 23
  constant, 24
  function, 24
  logical, 24
  non-logical, 24
  relation, 24
syntax, 8

tautology, 14, 35, 38
terms, 25
theory, 43
  complete, 43
truth assignment, 11
truth table, 15
type, 61
  complete, 65, 74
types
  omitting, 66

unique readability, 9, 25
universe, 44

valid, 38
validity, 14, 19, 33
variable, 24
  free, 27

well-ordering, 100
witness, 50