

Lecture Notes for
Mathematics 571
Fall 2016

Model Theory

written by

Lou van den Dries and
C. Ward Henson

Mathematics Department
University of Illinois
1409 West Green Street
Urbana, Illinois 61801

`vddries@math.uiuc.edu`
`henson@math.uiuc.edu`

`www.math.uiuc.edu/~henson/`
`www.math.uiuc.edu/~vddries/`
`www.math.uiuc.edu/People/vddries.html`

CONTENTS

Introduction	1
1. Compactness	3
Appendix 1.A: Ultrafilters	8
Appendix 1.B: From prestructures to structures	9
2. Boolean Algebras	13
3. Theories and Types	21
4. Elementary Maps	28
5. Saturated Models	36
6. Löwenheim-Skolem Theorems	42
7. Quantifier Elimination	45
8. Algebraically Closed Fields	53
9. \mathbb{Z} -groups	59
10. Model Theoretic Algebraic Closure	64
11. Algebraic Closure in Minimal Structures	67
12. Real Closed Fields	73
13. Homogeneity	78
14. Omitting Types	83
15. Skolem Hulls	95
16. Indiscernibles	99
17. Cantor-Bendixson Rank	104
18. Morley Rank	105
19. Universal Sentences and Substructures	122
20. Existential Formulas and Model Completeness	127
21. Characterizing Definability	133
22. Systems of Definable Sets and Functions	142
23. Selected Problems	149
24. Some Problem Settings	153

INTRODUCTION

This course gives an introduction to the methods of model theory for first order logic. Model theory is the branch of logic that deals with mathematical structures and the formal languages they interpret. First order logic is the most important formal language and its model theory has many connections to the main body of mathematics.

The central object of study in model theory is the collection of *definable* sets and functions on a structure \mathcal{M} . Let M be the underlying set of \mathcal{M} . A set $A \subseteq M^n$ is *definable* (without parameters) in \mathcal{M} if there is a formula $\varphi(x_1, \dots, x_n)$ in the first order language associated to \mathcal{M} such that

$$A = \{(a_1, \dots, a_n) \in M^n \mid \varphi \text{ is true of } (a_1, \dots, a_n) \text{ in } \mathcal{M}\}.$$

A function is definable if its graph is a definable set. We develop a number of tools for analyzing the definable sets and functions on a structure \mathcal{M} , including the method of quantifier elimination. A fundamental idea involves passing to elementary extensions of \mathcal{M} (for example, to ultrapowers of \mathcal{M}). This enriches the underlying set M without changing the structure of the category of definable sets and functions.

As a focus for developing “pure” model theory, we prove Morley’s Theorem: *if T is a complete theory in a countable language, and T is κ -categorical for some uncountable κ , then T is categorical for all uncountable κ .* In developing the tools needed to prove this theorem, we introduce *stability*, one of the key concepts of modern model theory.

We present many applications and examples in order to show how model theory can be useful in mathematics. For example, we treat the model theory of the field of real numbers (real closed fields) and show how this can be used to obtain the solution to Hilbert’s 17th Problem: a rational function over \mathbb{R} is positive semi-definite iff it is a sum of squares. Our treatment of real closed fields allows us to show that the definable sets in the field \mathbb{R} (equipped with names for its elements) are exactly the semi-algebraic sets.

A prerequisite is exposure to the syntax and semantics of first order logic, and experience with expressing mathematical properties via first order formulas. A good undergraduate course in logic will usually provide the necessary background. The canonical prerequisite at UIUC is the first half of Math 570. The lecture notes for Math 570 (written by Prof. van den Dries) are available at <http://www.math.uiuc.edu/~vddries/>

1. COMPACTNESS

The main purpose of this chapter is to give a proof of the Compactness Theorem for arbitrary first order languages. We do this using ultraproducts. The ultraproduct construction has the virtue of being explicit and algebraic in character, so it is accessible to mathematicians who know little about formal logic.

Fix a first order language L . Let I be a nonempty set and let U be an ultrafilter¹ on I . Consider a family of L -structures $(\mathcal{M}_i \mid i \in I)$. For each $i \in I$ let M_i denote the underlying set of the structure \mathcal{M}_i and take $M = \prod(M_i \mid i \in I)$ to be the cartesian product of the sets M_i .

We define an interpretation² \mathcal{M} of L as follows:

- (i) the underlying set of \mathcal{M} is the cartesian product $M = \prod(M_i \mid i \in I)$;
- (ii) for each constant symbol c of L we set

$$c^{\mathcal{M}} = (c^{\mathcal{M}_i} \mid i \in I);$$

- (iii) for each n and each n -ary function symbol F of L we let $F^{\mathcal{M}}$ be the function defined on M^n by

$$F^{\mathcal{M}}(f_1, \dots, f_n) = (F^{\mathcal{M}_i}(f_1(i), \dots, f_n(i)) \mid i \in I);$$

- (iv) for each n and each n -ary predicate symbol P of L we let $P^{\mathcal{M}}$ be the n -ary relation on M defined by

$$P^{\mathcal{M}}(f_1, \dots, f_n) \iff \{i \in I \mid P^{\mathcal{M}_i}(f_1(i), \dots, f_n(i))\} \in U;$$

- (v) $=^{\mathcal{M}}$ is the binary relation on M defined by

$$f =^{\mathcal{M}} g \iff \{i \in I \mid f(i) = g(i)\} \in U.$$

Note that constants and function symbols are treated in this construction in a “coordinatewise” way, exactly as we would do in forming the cartesian product of algebraic structures. Only in defining the interpretations of predicate symbols and $=$ (clauses (iv) and (v)) do we do something novel, and only there does the ultrafilter enter into the definition.

For the algebraic part of \mathcal{M} we have the following easy fact, proved by a straightforward argument using induction on terms:

1.1. Lemma. *For any L -term $t(x_1, \dots, x_n)$ and any $f_1, \dots, f_n \in M$,*

$$t^{\mathcal{M}}(f_1, \dots, f_n) = (t^{\mathcal{M}_i}(f_1(i), \dots, f_n(i)) \mid i \in I).$$

The following result gives the most important model theoretic property of this construction:

¹See Appendix A of this chapter for some basic facts about ultrafilters.

²See Appendix B of this chapter for an explanation of the words “interpretation”, “prestructure”, and “structure” and for some basic relations among them.

1.2. Proposition. For any L -formula $\varphi(x_1, \dots, x_n)$ and any $f_1, \dots, f_n \in M$ we have

$$\mathcal{M} \models \varphi[f_1, \dots, f_n] \iff \{i \in I \mid \mathcal{M}_i \models \varphi[f_1(i), \dots, f_n(i)]\} \in U.$$

Proof. The proof is by induction on formulas $\varphi(x_1, \dots, x_n)$, where x_1, \dots, x_n is an arbitrary list of distinct variables.

Basis step: Here φ is an atomic formula of the form $P(t_1, \dots, t_m)$, where P is an m -place predicate symbol or the equality symbol $=$. Our assumptions ensure that any variable occurring in a term t_j , $j = 1, \dots, m$, is among x_1, \dots, x_n ; thus we may write each such t_j as $t_j(x_1, \dots, x_n)$.

Let (f_1, \dots, f_n) range over M^n ; let $g_j(i) = t_j^{\mathcal{M}_i}(f_1(i), \dots, f_n(i))$ for each $j = 1, \dots, m$ and $i \in I$. Note that $g_j \in M$ for each $j = 1, \dots, m$. Then we have:

$$\begin{aligned} \mathcal{M} \models \varphi[f_1, \dots, f_n] &\iff P^{\mathcal{M}}(t_1^{\mathcal{M}}(f_1, \dots, f_n), \dots, t_m^{\mathcal{M}}(f_1, \dots, f_n)) \\ &\iff P^{\mathcal{M}}(g_1, \dots, g_m) \\ &\iff \{i \mid P^{\mathcal{M}_i}(g_1(i), \dots, g_m(i))\} \in U \\ &\iff \left\{i \mid P^{\mathcal{M}_i}\left(t_1^{\mathcal{M}_i}(f_1(i), \dots, f_n(i)), \dots, t_m^{\mathcal{M}_i}(f_1(i), \dots, f_n(i))\right)\right\} \in U \\ &\iff \{i \mid \mathcal{M}_i \models \varphi[f_1(i), \dots, f_n(i)]\} \in U. \end{aligned}$$

(Lemma 1.1 is used in the second equivalence.)

Induction step: Here we consider three cases: (1) φ is $\neg\varphi_1$ for some formula φ_1 ; (2) φ is $(\varphi_1 \wedge \varphi_2)$ for some formulas φ_1, φ_2 (3); φ is $\exists y\varphi_1$ for some formula φ_1 and some variable y .

Case (1) φ is $\neg\varphi_1$:

$$\begin{aligned} \mathcal{M} \models \varphi[f_1, \dots, f_n] &\iff \mathcal{M} \not\models \varphi_1[f_1, \dots, f_n] \\ &\iff \{i \mid \mathcal{M}_i \models \varphi_1[f_1(i), \dots, f_n(i)]\} \notin U \\ &\iff^* \{i \mid \mathcal{M}_i \not\models \varphi_1[f_1(i), \dots, f_n(i)]\} \in U \\ &\iff \{i \mid \mathcal{M}_i \models \neg\varphi_1[f_1(i), \dots, f_n(i)]\} \in U \\ &\iff \{i \mid \mathcal{M}_i \models \varphi[f_1(i), \dots, f_n(i)]\} \in U \end{aligned}$$

In the equivalence \star we use the fact that for any subset A of I , A is not in U if and only if $I \setminus A$ is in U .

Case (2) φ is $(\varphi_1 \wedge \varphi_2)$:

$$\begin{aligned} \mathcal{M} \models \varphi[f_1, \dots, f_n] &\iff \mathcal{M} \models (\varphi_1 \wedge \varphi_2)[f_1, \dots, f_n] \\ &\iff \mathcal{M} \models \varphi_1[f_1(i), \dots, f_n(i)] \text{ and } \\ &\quad \mathcal{M} \models \varphi_2[f_1(i), \dots, f_n(i)] \\ &\iff \{i \mid \mathcal{M}_i \models \varphi_1[f_1(i), \dots, f_n(i)]\} \in U \text{ and } \\ &\quad \{i \mid \mathcal{M}_i \models \varphi_2[f_1(i), \dots, f_n(i)]\} \in U \\ &\iff^* \{i \mid \mathcal{M}_i \models \varphi_1[f_1(i), \dots, f_n(i)]\} \cap \\ &\quad \{i \mid \mathcal{M}_i \models \varphi_2[f_1(i), \dots, f_n(i)]\} \in U \\ &\iff \{i \mid \mathcal{M}_i \models \varphi_1[f_1(i), \dots, f_n(i)] \text{ and } \\ &\quad \mathcal{M}_i \models \varphi_2[f_1(i), \dots, f_n(i)]\} \in U \end{aligned}$$

(continuing)

$$\begin{aligned} \mathcal{M} \models \varphi[f_1, \dots, f_n] &\Leftrightarrow \{i \mid \mathcal{M}_i \models (\varphi_1 \wedge \varphi_2)[f_1(i), \dots, f_n(i)]\} \in U \\ &\Leftrightarrow \{i \mid \mathcal{M}_i \models \varphi[f_1(i), \dots, f_n(i)]\} \in U \end{aligned}$$

In the equivalence \star we use the fact that for any subsets A and B of I , A and B are in U if and only if $A \cap B$ is in U .

Case (3) φ is the formula $\exists y\varphi_1$: By passing to a logically equivalent formula, we may assume that y is not among x_1, \dots, x_n .

$$\begin{aligned} &\mathcal{M} \models \varphi[f_1, \dots, f_n] \\ \Leftrightarrow &\mathcal{M} \models \exists y\varphi_1[f_1, \dots, f_n] \\ \Leftrightarrow &\text{for some } g \in M : \mathcal{M} \models \varphi_1[f_1, \dots, f_n, g] \\ \Leftrightarrow &\text{for some } g \in M : \{i \mid \mathcal{M}_i \models \varphi_1[f_1(i), \dots, f_n(i), g(i)]\} \in U \\ \Leftrightarrow^\dagger &\{i \mid \text{for some } a \in M_i : \mathcal{M}_i \models \varphi_1[f_1(i), \dots, f_n(i), a]\} \in U \\ \Leftrightarrow &\{i \mid \mathcal{M} \models \exists y\varphi_1[f_1(i), \dots, f_n(i)]\} \in U \end{aligned}$$

To see the “ \Leftarrow ”-part of the equivalence \dagger , use the Axiom of Choice to obtain a function g such that for any $i \in \{i \mid \text{for some } a \in M_i : \mathcal{M}_i \models \varphi_1[f_1(i), \dots, f_n(i), a]\}$ we have $\mathcal{M}_i \models \varphi_1[f_1(i), \dots, f_n(i), g(i)]$. For all other values of i the value of $g(i)$ can be arbitrary. \square

1.3. Corollary. *The interpretation \mathcal{M} defined above is a prestructure.*

Proof. Applying Proposition 1.2 to the equality axioms, we see that they are valid in \mathcal{M} . \square

1.4. Definition (Ultraproduct of a family of L -structures). Let $(\mathcal{M}_i \mid i \in I)$ be a family of L -structures and U an ultrafilter on I . Let \mathcal{M} be the prestructure for L that is defined above. The *ultraproduct* $\prod_U \mathcal{M}_i$ of the given family of L -structures $(\mathcal{M}_i \mid i \in I)$ with respect to U is defined to be the L -structure \mathcal{N} obtained by taking the quotient of \mathcal{M} by the congruence $=^{\mathcal{M}}$ as described in Appendix 2 of this chapter.

1.5. Notation. Let $(\mathcal{M}_i \mid i \in I)$ and \mathcal{M} be as above. For each $f \in \prod M_i$ we let f/U denote the equivalence class of f under the equivalence relation $=^{\mathcal{M}}$. As f varies, f/U gives an arbitrary element of the underlying set of the ultraproduct $\prod_U \mathcal{M}_i$.

The *ultrapower of the L -structure \mathcal{M}_0 with respect to U* is the ultraproduct $\prod_U \mathcal{M}_i$ with \mathcal{M}_i equal to \mathcal{M}_0 for every $i \in I$. We denote this structure by \mathcal{M}_0^I/U .

1.6. Fact. Let I be a nonempty set and let U be the principal ultrafilter on I that is generated by the singleton set $\{j\}$, where j is a fixed element of I . For every family $\{\mathcal{M}_i \mid i \in I\}$ of L -structures, the ultraproduct $\prod_U \mathcal{M}_i$ is isomorphic to \mathcal{M}_j .

The next theorem is the main result of this chapter; it is basic to any use of the ultraproduct construction in model theory. This result was originally proved by the Polish logician Jerzy Łoś.

1.7. Theorem (Fundamental Theorem of Ultraproducts). *Let an indexed family of L -structures and an ultrafilter U be given as described above. For any L -formula $\varphi(x_1, \dots, x_n)$ and any sequence $f/U = (f_1/U, \dots, f_n/U)$,*

$$\prod_U \mathcal{M}_i \models \varphi[f/U] \text{ if and only if } \{i \in I \mid \mathcal{M}_i \models \varphi[f_1(i), \dots, f_n(i)]\} \in U.$$

Proof. This is an immediate consequence of Propositions 1.2 and 1.25. \square

1.8. Corollary. *If σ is an L -sentence, then*

$$\prod_U \mathcal{M}_i \models \sigma \text{ if and only if } \{i \in I \mid \mathcal{M}_i \models \sigma\} \in U.$$

Proof. This is a special case of Theorem 1.7. \square

Now we use the ultraproduct construction to prove the Compactness Theorem, which is one of the most important tools in model theory. First we need a basic definition:

1.9. Definition. Let T be a set of sentences in L and let \mathcal{M} be an L -structure. We say that \mathcal{M} is a *model* of T , and write $\mathcal{M} \models T$, if every sentence in T is true in \mathcal{M} .

1.10. Theorem (Compactness Theorem). *Let T be any set of sentences in L . If every finite subset of T has a model, then T has a model.*

Proof. Assume that every finite subset of T has a model. Let I be the set of all finite subsets of T . For each $i \in I$ let \mathcal{M}_i be any model of i , which exists by assumption. We will obtain the desired model of T as an ultraproduct $\prod_U \mathcal{M}_i$ for a suitably chosen ultrafilter U on I .

Let S be the family of all subsets of I of the form $I_\sigma = \{i \in I \mid \sigma \in i\}$, where $\sigma \in T$. Note that S has the finite intersection property; indeed, each finite intersection $I_{\sigma_1} \cap \dots \cap I_{\sigma_n}$ has $\{\sigma_1, \dots, \sigma_n\}$ as an element. So there exists an ultrafilter U on I that contains S , by Corollary 1.21.

We complete the proof by showing that the ultraproduct $\prod_U \mathcal{M}_i$ is a model of T . Given $\sigma \in T$, we see that $\mathcal{M}_i \models \sigma$ whenever $\sigma \in i$, because of the way we chose \mathcal{M}_i . Hence $\{i \mid \mathcal{M}_i \models \sigma\} \supseteq I_\sigma \in U$. It follows from Theorem 1.7 that each such σ is true in $\prod_U \mathcal{M}_i$. \square

1.11. Remark. Note that the preceding proof yields the following result: Let T be a set of sentences in L and let \mathcal{C} be a class of L -structures. Suppose each finite subset of T has a model in \mathcal{C} . Then T has a model that is an ultraproduct of structures from \mathcal{C} .

The Compactness Theorem is a very useful tool for building models of a given set of sentences, and nearly everything we do in this course depends on it in one way or another. We give a number of examples of this in the rest of this chapter.

1.12. **Corollary.** *Let L be a first order language and let κ be an infinite cardinal number. If T is a set of sentences in L such that for each positive integer n there is a model of T with at least n elements, then T has a model with at least κ many elements. (In particular, this holds if T has at least one infinite model.)*

Proof. Expand L by adding a set of κ many new constant symbols; let L' be the new language. Let T' be T together with all sentences $\neg(c_1 = c_2)$, where c_1 and c_2 are distinct new constants. Our hypothesis implies that every finite subset of T' has a model. Therefore T' itself has a model \mathcal{N} , by the Compactness Theorem. Let \mathcal{M} be the reduct of \mathcal{N} to the original language L ; \mathcal{M} is a model of T and has at least κ many elements. \square

1.13. **Fact.** Let T be a set of sentences in a first order language L and let $\varphi(x)$ be a formula in L . For each L -structure \mathcal{M} let $\varphi^{\mathcal{M}}$ denote the set of tuples a from M such that $\mathcal{M} \models \varphi(a)$. Suppose that the set $\varphi^{\mathcal{M}}$ is finite whenever \mathcal{M} is a model of T . Then there is a positive integer N such that $\varphi^{\mathcal{M}}$ has at most N elements for every model \mathcal{M} of T . This can be proved using the Compactness Theorem in a manner similar to the proof of the previous result.

1.14. **Remark.** The preceding results demonstrate a fundamental limitation on the expressive power of first order logic: only *finite* cardinalities can be “expressed” by first order formulas. There is no way to express any bound on the sizes of definable sets other than a uniform finite upper bound. We will see later on how to control more precisely the cardinality of models like the one constructed above. In particular, it turns out to be possible to make the model have *precisely* κ many elements, as long as the number of symbols in the language L is less than or equal to κ .

1.15. **Definition.** Let Γ be a set of L -formulas and let the family $(x_j \mid j \in J)$ include all variables that occur free in some member of Γ . Let \mathcal{M} be an L -structure. We say that Γ is *satisfiable in \mathcal{M}* if there exist elements $(a_j \mid j \in J)$ of M such that $\mathcal{M} \models \Gamma[a_j \mid j \in J]$.

1.16. **Definition.** Let T be a set of sentences in L and Γ a set of L -formulas. We say that Γ is *consistent with T* if for every finite subsets F of T and G of Γ there exists a model \mathcal{M} of F such that G is satisfiable in \mathcal{M} .

The next result is a version of the Compactness Theorem for formulas.

1.17. **Corollary.** *Let T be a set of sentences in L and Γ a set of L -formulas, and assume that Γ is consistent with T . Then Γ is satisfiable in some model of T .*

Proof. Let $(x_j \mid j \in J)$ include all variables that occur free in some member of Γ . Let $(c_j \mid j \in J)$ be new constants and consider the language $L(c_j \mid j \in J)$. Apply the Compactness Theorem to the set $T \cup \Gamma(c_j \mid j \in J)$ of $L(c_j \mid j \in J)$ -sentences. \square

APPENDIX 1.A: ULTRAFILTERS

Here we present some prerequisites about ultrafilters:

1.18. **Definition.** Let I be a nonempty set. A *ultrafilter on I* is a collection U of subsets of I that satisfies:

- (a) $\emptyset \notin U$ and $I \in U$;
- (b) for all $A, B \in U$, we have $A \cap B \in U$;
- (c) for all $A \in U$ and $B \subseteq I$, if $A \subseteq B$ then $B \in U$.
- (d) for all $A \subseteq I$, either $A \in U$ or $I \setminus A \in U$.

Note. A *proper filter* on I is a collection U of subsets of I satisfying (a), (b), (c) in 1.18.

1.19. **Definition.** Let I be a set and let S be a collection of subsets of I ; S has the *finite intersection property (FIP)* if for every integer n and every finite subcollection $\{A_1, \dots, A_n\}$ of S , the intersection $A_1 \cap \dots \cap A_n$ is nonempty.

1.20. **Lemma.** Let I be a nonempty set and U a collection of subsets of I . Then U is an ultrafilter on I iff U has the FIP and is maximal among collections of subsets of I having the FIP.

Proof. (\Rightarrow) It is immediate (by induction) from 1.18(b) that U is closed under finite intersections; hence no such intersection can be empty, by 1.18(a). For any $A \subseteq I$, if $A \notin U$ then by 1.18(d) we have $I \setminus A \in U$ and hence $U \cup \{A\}$ fails FIP. This shows that U is maximal among collections with FIP.

(\Leftarrow) Suppose U is maximal among collections of subsets of I that have the FIP. Conditions (a), (b), and (c) obviously hold for U . For (d), suppose $A \subseteq I$ and $A \notin U$. Then $U \cup \{A\}$ must fail the FIP. That is, there must exist $A_1, \dots, A_n \in U$ such that $A_1 \cap \dots \cap A_n \cap A = \emptyset$. But then $I \setminus A \supseteq A_1 \cap \dots \cap A_n \in U$ and hence $I \setminus A \in U$. Therefore (d) also holds for U . \square

In the next proof we use the Axiom of Choice in the form of Zorn's Lemma, which we formulate as follows:

ZORN'S LEMMA: If (Λ, \leq) is a nonempty partially ordered set with the property that every linearly ordered subset of (Λ, \leq) has an upper bound in (Λ, \leq) , then (Λ, \leq) has a maximal element.

1.21. **Lemma.** Let I be a nonempty set and let S be a collection of subsets of I . There exists an ultrafilter on I that contains S if and only if S has the finite intersection property.

Proof. (\Rightarrow) Since each ultrafilter has the FIP, the same is true of any sub-collection of an ultrafilter.

(\Leftarrow) Let Λ_S be the family of all collections R of subsets of I such that R has FIP and $R \supseteq S$. It is immediate that the union of every nonempty linearly ordered subset of (Λ_S, \subseteq) has FIP and contains S . Therefore, Zorn's Lemma implies that (Λ_S, \subseteq) has a maximal element, which must be an ultrafilter by 1.20. \square

APPENDIX 1.B: FROM PRESTRUCTURES TO STRUCTURES

Let L be any first order language.

1.22. Definition. An *interpretation* \mathcal{M} of L consists of

- (i) a nonempty set M , the *underlying set* of \mathcal{M} ;
- (ii) for each constant symbol c of L an element $c^{\mathcal{M}}$ of M , the *interpretation of c in \mathcal{M}* ;
- (iii) for each n and each n -ary function symbol F of L a function $F^{\mathcal{M}}$ from M^n to M , the *interpretation of F in \mathcal{M}* ;
- (iv) for each n and each n -ary predicate symbol P of L a subset $P^{\mathcal{M}}$ of M^n , the *interpretation of P in \mathcal{M}* ;
- (v) a subset $=^{\mathcal{M}}$ of M^2 , the *interpretation of $=$ in \mathcal{M}* .

Suppose \mathcal{M} is an interpretation of L . For each L -term $t(x_1, \dots, x_n)$ we define the interpretation of t in \mathcal{M} by induction on t ; it is a function from M^n to M and it is denoted by $t^{\mathcal{M}}$. By induction on formulas we likewise define the satisfaction relation

$$\mathcal{M} \models \varphi[a_1, \dots, a_n]$$

where $\varphi(x_1, \dots, x_n)$ is an L -formula and $a_1, \dots, a_n \in M$. Formally this is identical to what is done for L -structures, with which we assume the reader is familiar. The only difference here is that we are allowing an arbitrary binary relation to be used as the interpretation of $=$; that is, we are temporarily treating $=$ as if it were a non-logical symbol.

1.23. Definition. A *prestructure* \mathcal{M} for L is an interpretation of L in which the logical equality axioms are valid; that is,

- (i) $=^{\mathcal{M}}$ is an equivalence relation on M ;
- (ii) for any n , any n -ary function symbol F of L , and any elements $a_1, b_1, \dots, a_n, b_n$ of M such that $a_1 =^{\mathcal{M}} b_1, \dots, a_n =^{\mathcal{M}} b_n$ one has

$$F^{\mathcal{M}}(a_1, \dots, a_n) =^{\mathcal{M}} F^{\mathcal{M}}(b_1, \dots, b_n);$$

- (iii) for any n , any n -ary predicate symbol P of L , and any elements $a_1, b_1, \dots, a_n, b_n$ of M such that $a_1 =^{\mathcal{M}} b_1, \dots, a_n =^{\mathcal{M}} b_n$ one has

$$P^{\mathcal{M}}(a_1, \dots, a_n) \iff P^{\mathcal{M}}(b_1, \dots, b_n).$$

When $=^{\mathcal{M}}$ is an equivalence relation on M , universal algebraists express conditions (ii) and (iii) by saying that $=^{\mathcal{M}}$ is a *congruence* with respect to the functions $F^{\mathcal{M}}$ mentioned in (ii) and the relations $P^{\mathcal{M}}$ mentioned in (iii).

Note that \mathcal{M} is a *structure* for L if it is an interpretation of L and $=^{\mathcal{M}}$ is the identity relation on M , that is $a =^{\mathcal{M}} b \Leftrightarrow a = b$ for any $a, b \in M$. (In that case, \mathcal{M} trivially satisfies the equality axioms and hence it is a prestructure.)

When \mathcal{M} is a prestructure for L , we define the *quotient of \mathcal{M} by $=^{\mathcal{M}}$* as follows; it is a structure for L . We will denote it here by \mathcal{N} and by $\mathcal{M}/=^{\mathcal{M}}$.

(i) The underlying set N of \mathcal{N} is the set of all equivalence classes of $=^{\mathcal{M}}$. We denote the equivalence class of $a \in M$ with respect to $=^{\mathcal{M}}$ by $[a]$, and we let $\pi: M \rightarrow N$ denote the quotient map that takes each $a \in M$ to its equivalence class ($\pi(a) = [a]$).

(ii) For each constant symbol c of L the interpretation of c in \mathcal{N} is $[c^{\mathcal{M}}]$.

(iii) For each n and each n -ary function symbol F of L the interpretation of F in \mathcal{N} is the function $F^{\mathcal{N}}: N^n \rightarrow N$ defined by

$$F^{\mathcal{N}}([a_1], \dots, [a_n]) = [F^{\mathcal{M}}(a_1, \dots, a_n)]$$

for every $a_1, \dots, a_n \in M$. The fact that $=^{\mathcal{M}}$ is a congruence for $F^{\mathcal{M}}$ ensures that the right hand side of this definition depends only on the equivalence classes $[a_1], \dots, [a_n]$ and not on their representatives a_1, \dots, a_n .

(iv) For each n and each n -ary predicate symbol P of L the interpretation of P in \mathcal{N} is the n -ary relation $P^{\mathcal{N}}$ on N defined by

$$P^{\mathcal{N}}([a_1], \dots, [a_n]) \iff P^{\mathcal{M}}(a_1, \dots, a_n)$$

for every $a_1, \dots, a_n \in M$. The fact that $=^{\mathcal{M}}$ is a congruence for $P^{\mathcal{M}}$ ensures that the right hand side of this definition depends only on the equivalence classes $[a_1], \dots, [a_n]$ and not on their representatives a_1, \dots, a_n .

Since \mathcal{N} is to be a structure, the interpretation $=^{\mathcal{N}}$ of $=$ in \mathcal{N} must be the identity relation on N . Note that we have

$$[a] =^{\mathcal{N}} [b] \iff a =^{\mathcal{M}} b$$

for all $a, b \in M$. Hence the identity interpretation of $=$ in \mathcal{N} is the same as the one we would get if we treated $=$ as another predicate symbol of L and applied clause (iv) of this construction.

Our definition of the quotient structure \mathcal{N} can be summarized by saying that the quotient map π from M onto N is a strong homomorphism of \mathcal{M} onto \mathcal{N} .

The following Lemma is easily proved by induction on terms.

1.24. Lemma. *For any L -term $t(x_1, \dots, x_n)$ and any $a_1, \dots, a_n \in M$,*

$$t^{\mathcal{N}}([a_1], \dots, [a_n]) = [t^{\mathcal{M}}(a_1, \dots, a_n)].$$

The following result gives the main content of this quotient construction from a model theoretic point of view. It says that no difference between a prestructure \mathcal{M} and its quotient structure \mathcal{N} can be expressed in first order logic. It justifies the usual practice of only considering structures in

model theory. (However, prestructures are often used, at least implicitly, in the construction of structures; this happens in the usual proof of the completeness theorem for first order logic, for example.)

1.25. Proposition. *Let \mathcal{M} be a prestructure for L and \mathcal{N} its quotient structure as described above. For any L -formula $\varphi(x_1, \dots, x_n)$ and any $a_1, \dots, a_n \in M$*

$$\mathcal{N} \models \varphi[[a_1], \dots, [a_n]] \iff \mathcal{M} \models \varphi[a_1, \dots, a_n].$$

Proof. By induction on the formula φ . When φ is an atomic formula, this equivalence follows from the preceding Lemma and the fact that π is a strong homomorphism. The induction step is an immediate consequence of the definition of \models and (for quantifiers) the fact that π is surjective. \square

EXERCISES

1.26. Let I be a nonempty set and let U be an ultrafilter on I .

(a) If A_1, \dots, A_n are subsets of I and if the set $A_1 \cup \dots \cup A_n$ is in U , then for some $j = 1, \dots, n$ the set A_j is in U .

(b) If A_1, \dots, A_n are subsets of I and if the set $A_1 \cap \dots \cap A_n$ is in U , then for all $j = 1, \dots, n$ the set A_j is in U .

(c) The ultrafilter U is called *principal* if there exists $A \in U$ such that $A \subseteq B$ holds for every $B \in U$. Show that U is principal iff some element of U is finite iff some element of U is a singleton set (a set of the form $\{i\}$ for some $i \in I$). Further, if $\{i\} \in U$, show that U consists of all sets $A \subseteq I$ that contain i as an element.

1.27. Let I be a nonempty set, U an ultrafilter on I , and J an element of U . Define V to be the set of $X \subseteq J$ such that $X \in U$.

- Show that V is an ultrafilter on J .
- Show that if $(\mathcal{M}_i \mid i \in I)$ is a family of L -structures, then $\Pi_U(\mathcal{M}_i \mid i \in I)$ is isomorphic to $\Pi_V(\mathcal{M}_j \mid j \in J)$

1.28. Let I be an index set and U an ultrafilter on I . Let $(\mathcal{M}_i \mid i \in I)$ and $(\mathcal{N}_i \mid i \in I)$ be families of L -structures. If \mathcal{M}_i can be embedded in \mathcal{N}_i for all $i \in I$, show that $\Pi_U \mathcal{M}_i$ can be embedded in $\Pi_U \mathcal{N}_i$.

1.29. Let \mathcal{M} be any L -structure. Show that \mathcal{M} can be embedded in some ultraproduct of a family of finitely generated substructures of \mathcal{M} .

1.30. Let L be the first order language whose only nonlogical symbol is the binary predicate symbol $<$. Let $\mathcal{M} = (\mathbb{N}, <)$ and let $\mathcal{N} = \mathcal{M}^I/U$ be an ultrapower of \mathcal{M} where I is countably infinite and U is a nonprincipal ultrafilter on I .

- Show that \mathcal{N} is a linear ordering.
- Show that the range of the diagonal embedding of \mathcal{M} into \mathcal{N} is a proper initial segment of \mathcal{N} . Give an explicit description of an element of \mathcal{N} that is not in the range of this embedding.

- Show that \mathcal{N} is not a well ordering; that is, describe an infinite descending sequence in \mathcal{N} .

1.31. Let L be the first order language whose nonlogical symbols consist of a binary predicate symbol $<$, a binary function symbol $+$ and a constant symbol 0 . Let \mathbb{Z} be the ordered abelian group of all the integers, considered as an L -structure. Let I be any countable infinite set and let U be a nonprincipal ultrafilter on I . Consider the ultrapower \mathbb{Z}^I/U .

- Show that \mathbb{Z}^I/U is an ordered abelian group.
- Find a natural embedding of \mathbb{Z} into this group so that the image of the embedding is a convex subgroup.
- Show that \mathbb{Z}^I/U contains a nonzero element b that is divisible in \mathbb{Z}^I/U by every positive integer n . (This means that for each $n \geq 1$ there exists a in \mathbb{Z}^I/U that satisfies $b = a + \cdots + a$ (n times).) Such an element can be produced explicitly.

2. BOOLEAN ALGEBRAS

Revision Note: this chapter will be extended to add a small amount of material about boolean algebras that is needed for developing CB rank (chapter 17).

Propositional logic and boolean algebras share a historical origin in the work of Boole. This line was continued by Peirce and Schröder, and gradually the algebraic aspects of propositional logic got purified in the notion of *boolean algebra*. This has turned out to be useful in areas such as set theory, measure theory, and model theory. Skolem, Tarski, and Stone are among those who further developed the subject of boolean algebras in the 20th century.

A *boolean algebra* B is a set B equipped with distinguished elements 0 and 1 (the zero and unit of B), a unary operation $'$ on B (complementation), and binary operations \vee (join) and \wedge (meet) on B such that for all $x, y, z \in B$,

$$\begin{array}{ll}
 (1) & x \vee 0 = x, & x \wedge 1 = x, \\
 (2) & x \vee x' = 1, & x \wedge x' = 0, \\
 (3) & x \vee y = y \vee x, & x \wedge y = y \wedge x, \\
 (4) & (x \vee y) \vee z = x \vee (y \vee z), & (x \wedge y) \wedge z = x \wedge (y \wedge z), \\
 (5) & x \vee (y \wedge z) = (x \vee y) \wedge (x \vee z), & x \wedge (y \vee z) = (x \wedge y) \vee (x \wedge z).
 \end{array}$$

In the rest of this chapter B denotes a boolean algebra, and we let a, b, c, x, y range over its underlying set B . We leave it as an exercise to prove the following identities in B :

- $0' = 1, 1' = 0,$
- $x \vee x = x, x \wedge x = x,$
- $x'' = x,$
- $(x \vee y)' = x' \wedge y', (x \wedge y)' = x' \vee y',$
- $x \vee y = y \iff x \wedge y = x.$

The reader should also verify that we have a partial ordering on B defined by $x \leq y :\iff x \vee y = y$, and that the element $x \vee y$ is the least upper bound in B of the subset $\{x, y\}$, and that $x \wedge y$ is the largest lower bound of $\{x, y\}$. Also, $0 \leq x \leq 1$ for all x, y ; and for all x , we see x' is the unique element y of B that satisfies $x \wedge y = 0$ and $x \vee y = 1$. In this way all ingredients of a boolean algebra can be defined in terms of its partial ordering.

The laws that define boolean algebras have an obvious symmetry, which suggests the following construction. To B we associate the boolean algebra B^{op} , the *opposite of B* : it has the same underlying set as B , with zero $0^{\text{op}} := 1$, unit $1^{\text{op}} := 0$, the same complementation operation as B , and

join and meet operations \vee^{op} and \wedge^{op} defined as follows:

$$x \vee^{\text{op}} y := x \wedge y, \quad x \wedge^{\text{op}} y = x \vee y.$$

Note that $(B^{\text{op}})^{\text{op}} = B$.

A *subalgebra* of B is a subset A of B such that $0, 1 \in A$ and for all $x, y \in A$ we have $x', x \vee y, x \wedge y \in A$. Such A will be considered as a boolean algebra in its own right, by taking 0 as the zero of A , 1 as the unit of A , and by taking the complementation, join and meet operations on A to be restrictions of the corresponding operations on B .

Let P be a subset of B . There is a smallest subalgebra of B (with respect to inclusion) that contains P , namely the intersection of all subalgebras of B that contain P . The smallest subalgebra of B containing P is called the *subalgebra of B generated by P* ; if it happens to be B itself, then we say that P *generates B* . We shall prove below that if B is finitely generated (that is, B has a finite generating set), then B itself is finite, and we shall completely describe the finite boolean algebras.

Example. Let S be a set. Then the powerset $\mathcal{P}(S)$ becomes a boolean algebra by setting $0 := \emptyset$, $1 := S$, and for $X, Y \subseteq S$,

$$X' := S \setminus X, \quad X \vee Y := X \cup Y, \quad X \wedge Y := X \cap Y.$$

We call this boolean algebra the *powerset algebra on S* , to be denoted just by $\mathcal{P}(S)$. Note that the partial ordering \leq of this boolean algebra is given by inclusion: for $X, Y \subseteq S$,

$$X \leq Y \iff X \subseteq Y.$$

Note also that a subalgebra of $\mathcal{P}(S)$ is just a collection A of subsets of S such that $S \in A$, and $S \setminus X, X \cup Y \in A$ for all $X, Y \in A$, because it follows that then also $\emptyset \in A$ and $X \cap Y \in A$ for all $X, Y \in A$.

A subalgebra of $\mathcal{P}(S)$ is also called an *algebra on S* . For example, the collection of subsets of S that are finite or cofinite is an algebra on S . (A subset X of S is *cofinite* if $S \setminus X$ is finite.) It is the subalgebra of $\mathcal{P}(S)$ generated by the set $\{\{s\} \mid s \in S\}$ of singletons.

The main goal in this chapter is to prove the Tarski-Stone theorem that every boolean algebra is isomorphic to an algebra on a set. We actually obtain a strong version of this, the Stone representation theorem. But we first need to define the notion of isomorphism of boolean algebras.

Another example. Let A be a set of propositional atoms and $\Sigma \subseteq \text{Prop}(A)$ a set of propositions. Let p, q, r range over $\text{Prop}(A)$.

We say that p is Σ -equivalent to q if $\Sigma \models p \leftrightarrow q$. It is clear that Σ -equivalence is an equivalence relation on $\text{Prop}(A)$. Let p/Σ be the Σ -equivalence class of p , and let $\text{Prop}(A)/\Sigma$ be the set of equivalence classes.

It is easy to check that $\text{Prop}(A)/\Sigma$ becomes a boolean algebra by setting $0 := \perp/\Sigma$, $1 := \top/\Sigma$, and

$$(p/\Sigma)' := (\neg p)/\Sigma, \quad p/\Sigma \vee q/\Sigma := (p \vee q)/\Sigma, \quad p/\Sigma \wedge q/\Sigma := (p \wedge q)/\Sigma.$$

We call this boolean algebra the **Lindenbaum algebra** of Σ . Note that it is generated by $\{a/\Sigma : a \in A\}$.

Morphisms and ideals. In this subsection C and D also denote boolean algebras (besides B). We say that a map $\varphi: B \rightarrow C$ is a (*homo*)*morphism* of boolean algebras if for all x, y ,

$$\varphi(0_B) = 0_C, \quad \varphi(1_B) = 1_C, \quad \varphi(x') = \varphi(x)',$$

$$\varphi(x \vee y) = \varphi(x) \vee \varphi(y), \quad \varphi(x \wedge y) = \varphi(x) \wedge \varphi(y).$$

Note that then $\varphi(B)$ is a subalgebra of C . Note also that if $\varphi: B \rightarrow C$ and $\psi: C \rightarrow D$ are boolean algebra morphisms, then the composition $\psi \circ \varphi: B \rightarrow D$ is a boolean algebra morphism.

For example, any map $f: S \rightarrow T$ induces a boolean algebra morphism

$$Y \mapsto f^{-1}(Y): \mathcal{P}(T) \rightarrow \mathcal{P}(S)$$

between the powerset algebras. As a special case, if $S \subseteq T$ and $f: S \rightarrow T$ is the inclusion map, this gives the boolean algebra morphism

$$Y \mapsto Y \cap S: \mathcal{P}(T) \rightarrow \mathcal{P}(S).$$

If $\varphi: B \rightarrow C$ is a bijective boolean algebra morphism, then $\varphi^{-1}: C \rightarrow B$ is also a boolean algebra morphism, and such φ is called a *boolean algebra isomorphism*. For example, the identity map $\text{id}_B: B \rightarrow B$ is a boolean algebra isomorphism. We also have the boolean algebra isomorphism

$$x \mapsto x': B \rightarrow B^{\text{op}}.$$

We say that B and C are *isomorphic* if there exists a boolean algebra isomorphism $B \rightarrow C$. Note that if $\varphi: B \rightarrow C$ is an injective boolean algebra morphism, then $\varphi: B \rightarrow \varphi(B)$ is a boolean algebra isomorphism. We leave it as an exercise to check that a boolean algebra morphism $B \rightarrow C$ is injective iff the only element in B with image 0_C is 0_B .

2.1. Lemma. *Let $X_1, \dots, X_n \subseteq S$, and let A be the algebra on S generated by $\{X_1, \dots, X_n\}$. Then A is finite and isomorphic to $\mathcal{P}(\{1, \dots, m\})$ for some $m \leq 2^n$. In particular, $|A| \leq 2^{2^n}$.*

Proof. For $\varepsilon: \{1, \dots, n\} \rightarrow \{-1, 1\}$, put

$$X^\varepsilon := X_1^{\varepsilon(1)} \cap \dots \cap X_n^{\varepsilon(n)},$$

where for any $Y \subseteq S$ we put $Y^1 := Y$ and $Y^{-1} := S \setminus Y$. Note that if $\varepsilon, \varepsilon': \{1, \dots, n\} \rightarrow \{-1, 1\}$ are distinct, then X^ε and $X^{\varepsilon'}$ are disjoint. Let S_1, \dots, S_m be the distinct *nonempty* sets among the sets X^ε . Then $m \leq 2^n$,

S_1, \dots, S_m are disjoint, and $S_1 \cup \dots \cup S_m = S$. It follows easily that the map

$$I \mapsto S(I) := \bigcup_{i \in I} S_i : \mathcal{P}(\{1, \dots, m\}) \rightarrow \mathcal{P}(S)$$

is an injective boolean algebra morphism. Its image must be A , which is therefore isomorphic to $\mathcal{P}(\{1, \dots, m\})$. \square

We call B *trivial* if $|B| = 1$, equivalently, $0 = 1$ in B . The power set algebra $\mathcal{P}(\emptyset)$ is trivial, and there is up-to-unique-isomorphism exactly one trivial boolean algebra. If $|S| = 1$, then the power set algebra $\mathcal{P}(S)$ has size 2, and there is up-to-unique isomorphism exactly one boolean algebra of size 2; we shall denote it by $\mathbb{2}$.

An *ideal* of B is a set $I \subseteq B$ such that $0 \in I$ and for all a, b , if $a \leq b \in I$, then $a \in I$, and if $a, b \in I$, then $a \vee b \in I$. For example, the collection of finite subsets of a set S is an ideal in the powerset algebra $\mathcal{P}(S)$. It may help to think of the elements in an ideal as “small” in a sense determined by the ideal. If I is an ideal of B , then

$$I \vee b := \{x \in B \mid x \leq a \vee b \text{ for some } a \in I\}$$

is easily checked to be the smallest ideal of B , with respect to inclusion, that contains $I \cup \{b\}$.

If $\varphi: B \rightarrow C$ is a morphism of boolean algebras, then $\{b: \varphi(b) = 0\}$ is an ideal of B , called the *kernel* of φ . Conversely, let I be an ideal of B . Then we have a equivalence relation $=_I$ on B given by

$$a =_I b :\iff a \vee x = b \vee x \text{ for some } x \in I.$$

Let a/I denote the $=_I$ -equivalence class of a . We make the set $B/I := \{a/I \mid a \in B\}$ into a boolean algebra by letting $0/I$ be its zero, $1/I$ be its unit, and setting

$$(a/I)' = a'/I, \quad a/I \vee b/I = (a \vee b)/I, \quad a/I \wedge b/I = (a \wedge b)/I.$$

Then the map $\varphi: B \rightarrow B/I$ given by $\varphi(b) = b/I$ is a boolean algebra morphism with kernel I .

Among the ideals of B are $\{0\}$ (the trivial ideal) and B itself. An ideal I of B is said to be *proper* if $1 \notin I$, equivalently, $I \neq B$. A *maximal ideal* of B is a proper ideal \mathfrak{m} of B such that the only ideals of B containing \mathfrak{m} are \mathfrak{m} and B . We now have the following key lemma.

2.2. Lemma. *Let I be a proper ideal of B . Then*

- (1) *if $b \notin I$, then $I \vee b'$ is a proper ideal of B ;*
- (2) *there is a maximal ideal $\mathfrak{m} \supseteq I$ of B ;*
- (3) *I is the intersection of the maximal ideals $\mathfrak{m} \supseteq I$ of B .*

Proof. Item (1) is because if $a \vee b' = 1$, then $b \leq a$. Item (2) follows from Zorn’s lemma applied to the collection of proper ideals of B containing I ,

ordered by inclusion. For (3), note that if $b \notin I$, then by (2) there is a maximal ideal $\mathfrak{m} \supseteq I \vee b'$, and we have $b \notin \mathfrak{m}$ for such \mathfrak{m} . \square

Here are some easy consequences of this lemma:

2.3. Corollary. *A proper ideal I of B is maximal iff for all b we have $b \in I$ or $b' \in I$. Assigning to each boolean algebra morphism $\varphi: B \rightarrow \mathbb{2}$ its kernel yields a bijection from the set of boolean algebra morphisms $B \rightarrow \mathbb{2}$ onto the set of maximal ideals of B .*

Ultrafilters, and the Stone representation. Sometimes we prefer a notion opposite to that of *ideal*, namely *filter*, because filters will correspond in some logical settings to *theories*, and this correspondence would be less direct in terms of ideals. A *filter of B* is a set $F \subseteq B$ such that $1 \in F$ and for all a, b , if $a \geq b \in F$, then $a \in F$, and if $a, b \in F$, then $a \wedge b \in F$; equivalently, it is an ideal of B^{op} . Note also that for any $F \subseteq B$, F is a filter of B iff $\{a' \mid a \in F\}$ is an ideal of B . A filter F of B is said to be *proper* if $0 \notin F$, equivalently, $F \neq B$. A *maximal filter of B* (also called: *ultrafilter of B*) is a proper filter \mathfrak{u} of B such that the only filters of B containing \mathfrak{u} are \mathfrak{u} and B . Here is the filter version of some facts about ideals.

2.4. Lemma. *Let F be a proper filter of B .*

- (1) *F is the intersection of the ultrafilters $\mathfrak{u} \supseteq F$ of B .*
- (2) *F is an ultrafilter iff for all b we have $b \in F$ or $b' \in F$.*

The set of ultrafilters of B is called the *Stone space* of B , and denoted by $\text{St}(B)$. (We use the word “space” since we shall give a topology to $\text{St}(B)$.)

Put $[a] := \{\mathfrak{u} \in \text{St}(B) \mid a \in \mathfrak{u}\}$. So for all a and for all $\mathfrak{u} \in \text{St}(B)$,

$$\mathfrak{u} \in [a] \iff a \in \mathfrak{u}.$$

2.5. Theorem. *The map $a \mapsto [a]: B \mapsto \mathcal{P}(\text{St}(B))$ is an injective morphism of boolean algebras.*

Proof. First, $[0] = \emptyset$ because $0 \notin \mathfrak{u}$ for all $\mathfrak{u} \in \text{St}(B)$, and $[1] = \text{St}(B)$. It is also trivial to check that

$$[a'] = \text{St}(B) \setminus [a], \quad [a \vee b] = [a] \cup [b], \quad [a \wedge b] = [a] \cap [b].$$

So the above map is indeed a morphism of boolean algebras. Its kernel is

$$\{a \mid a \notin \mathfrak{u} \text{ for all } \mathfrak{u} \in \text{St}(B)\},$$

which is $\{0\}$, because if $a \neq 0$, then $\{b \mid b \geq a\}$ is a proper filter and thus $a \in \mathfrak{u}$ for some $\mathfrak{u} \in \text{St}(B)$. \square

In particular, every boolean algebra is isomorphic to an algebra on a set. In view of Lemma 2.1 this has the following consequence:

2.6. Corollary. *Every finitely generated boolean algebra is finite, and isomorphic to the powerset algebra $\mathcal{P}(\{1, \dots, m\})$ for exactly one m .*

Boolean rings. Let S be a set, and associate to each set $A \subseteq S$ its characteristic function $\chi_A: S \rightarrow \{0, 1\}$ defined by $\chi_A(s) = 1$ for $s \in A$, and $\chi_A(s) = 0$ for $s \in S \setminus A$. Let us add and multiply such functions modulo 2: $1 + 1 = 0$, so $x = -x$ for $x \in \{0, 1\}$. Then for $A, B \subseteq S$ we get

$$\chi_{S \setminus A} = 1 - \chi_A, \quad \chi_{A \cap B} = \chi_A \cdot \chi_B, \quad \chi_{A \cup B} = \chi_A + \chi_B + \chi_A \cdot \chi_B,$$

which expresses in some sense the boolean operations in terms of addition and multiplication modulo 2. We can also express conversely addition and multiplication modulo 2 of such functions in terms of the basic boolean operations: with $A \Delta B := (A \cup B) \setminus (A \cap B)$, we have

$$\chi_A + \chi_B = \chi_{A \Delta B}, \quad \chi_A \cdot \chi_B = \chi_{A \cap B}.$$

The set $A \Delta B$ is called the *symmetric difference* of A and B . In fact, we can turn any boolean algebra B into a ring B_{ring} having the same underlying set as B , and with addition and multiplication given by

$$a + b := a \Delta b := (a \vee b) \wedge (a \wedge b)', \quad ab = a \wedge b.$$

Then B_{ring} is a commutative ring with 0 as its zero element and 1 as its multiplicative unit. (Verifying directly the associative law for addition and the distributive law takes a little effort; another way to see that these laws hold in B_{ring} is to note that B is isomorphic to an algebra on some set S , so that we can use the bijective correspondence between subsets of S and their characteristic functions, and apply the above observations on adding and multiplying such functions.) If we identify the elements $0, 1 \in \mathbb{2}$ with the corresponding integers 0 and 1 modulo 2, then $\mathbb{2}_{\text{ring}}$ is exactly the ring $\{0, 1\}$ of integers modulo 2.

In the ring B_{ring} we have $x^2 = x$ for all x . A ring in which this identity holds is called a *boolean ring*. (For us, a ring has by definition a unit 1.) A boolean ring also satisfies the identities $x + x = 0$ and $xy = yx$; to see why, consider $(x + 1)^2$ and $(x + y)^2$.

2.7. Lemma. *Every boolean ring equals B_{ring} for some boolean algebra B .*

We leave the proof as an exercise. It follows that a boolean ring equals B_{ring} for a unique boolean algebra B . Note that $\mathbb{2}_{\text{ring}} = \{0, 1\}$

Let B and C be boolean algebras, and $\varphi: B \rightarrow C$ a map. It is easy to check that φ is a boolean algebra morphism iff $\varphi: B_{\text{ring}} \rightarrow C_{\text{ring}}$ is a ring morphism (where the latter includes the requirement $\varphi(1) = 1$).

Let $I \subseteq B$. It is easy to check that I is an ideal of B iff I is an ideal of B_{ring} in the sense of (commutative) rings, that is, $0 \in I$, $a + b \in I$ for all $a, b \in I$, and $ax \in I$ for all $a \in I$ and all x in the ring. If I is an ideal of B , then $a/I = a + I$ for all $a \in B$, and thus $(B/I)_{\text{ring}} = B_{\text{ring}}/I$.

The Stone space. Returning to the Stone space $\text{St}(B)$ of a boolean algebra B , let us call the sets $[a] \subseteq \text{St}(B)$ *basic*. Since $[a] \cap [b] = [a \wedge b]$ for all $a, b \in B$, the collection of basic sets is a base for a topology on $\text{St}(B)$, called the **Stone topology**, which makes $\text{St}(B)$ a topological space.

So a set $U \subseteq \text{St}(B)$ is open iff U is a union of basic sets. Since the complement in $\text{St}(B)$ of any basic set is also basic, the basic sets are closed as well as open (clopen), and a set $F \subseteq \text{St}(B)$ is closed iff F is an intersection of basic sets. The key fact about the space $\text{St}(B)$ is the following:

2.8. Theorem. *$\text{St}(B)$ is a compact hausdorff space, and its clopen sets are exactly the basic sets.*

Proof. Let \mathbf{u} and \mathbf{v} be distinct points of $\text{St}(B)$. Viewed as ultrafilters on B , we have $a \in B$ such that $a \in \mathbf{u}$ and $a \notin \mathbf{v}$. Then $\mathbf{u} \in [a]$ and $\mathbf{v} \in [-a]$, and $[a] \cap [-a] = \emptyset$. Thus $\text{St}(B)$ is hausdorff.

To show that $\text{St}(B)$ is compact, let $(F_i)_{i \in I}$ be a family of closed subsets of $\text{St}(B)$ with $I \neq \emptyset$ and $\bigcap F_i = \emptyset$; it is enough to get $i_1, \dots, i_n \in I$ such that $F_{i_1} \cap \dots \cap F_{i_n} = \emptyset$. Representing each F_i as an intersection of basic sets, we reduce to the case that each F_i is basic, say $F_i = [a_i]$, $a_i \in B$. If $a_{i_1} \wedge \dots \wedge a_{i_n} \neq 0$ for all $i_1, \dots, i_n \in I$, then we have an ultrafilter $\mathbf{u} \in \text{St}(B)$ with $a_i \in \mathbf{u}$ for all $i \in I$, that is, $\mathbf{u} \in F_i$ for all $i \in I$, contradicting $\bigcap_i F_i = \emptyset$. So there exist $i_1, \dots, i_n \in I$ such that $a_{i_1} \wedge \dots \wedge a_{i_n} = 0$, and so $F_{i_1} \cap \dots \cap F_{i_n} = \emptyset$, as desired.

Let $X \subseteq \text{St}(B)$ be clopen. Then X is compact, so as a union of basic sets, it is union of finitely many basic sets, and thus a basic set. \square

The way we showed that $\text{St}(B)$ is hausdorff actually shows that $\text{St}(B)$ is totally disconnected: for any two distinct points of $\text{St}(B)$ there is a clopen set containing one but not the other.

For any topological space S , let $\text{Clopen}(S)$ be the boolean subalgebra of $\mathcal{P}(S)$ consisting of the clopen subsets of S . Thus by Theorem 2.5, we have an isomorphism

$$b \mapsto [b]: B \rightarrow \text{Clopen}(\text{St}(B))$$

from the boolean algebra B onto the boolean algebra of clopen sets in $\text{St}(B)$. It is important that assigning to a boolean algebra its Stone space extends to a contravariant functor from the category of boolean algebras and boolean algebra morphisms into the category of topological spaces with continuous maps as morphisms: any morphism $f: A \rightarrow B$ between boolean algebras A and B induces a continuous map

$$\text{St}(f): \text{St}(B) \rightarrow \text{St}(A), \quad \text{St}(f)(\mathbf{u}) := f^{-1}(\mathbf{u})$$

between their Stone spaces. Here is a key fact about such f :

2.9. Proposition. *f is injective iff $\text{St}(f)$ is surjective.*

Proof. First, suppose f is injective. Consider an ultrafilter \mathfrak{u} of A . The set $S = \{f(a) \mid a \in \mathfrak{u}\}$ has the finite intersection property in B ; that is, for any $f(a_1), \dots, f(a_n) \in S$ we have $f(a_1) \wedge \dots \wedge f(a_n) = f(a_1 \wedge \dots \wedge a_n) \neq f(0) = 0$. Hence there exists an ultrafilter \mathfrak{v} on B with $S \subseteq \mathfrak{v}$. Note that for every $a \in A$, we have $f(a) \in \mathfrak{v} \iff a \in \mathfrak{u}$ because \mathfrak{u} is an ultrafilter on A . Then $\text{St}(f)(\mathfrak{v}) = f^{-1}(\mathfrak{v}) = \mathfrak{u}$, and hence $\text{St}(f)$ is surjective.

Next assume f is not injective. Hence there exists $a \in A$ with $a \neq 0$ and $f(a) = 0$. Take $\mathfrak{u} \in \text{St}(A)$ with $a \in \mathfrak{u}$. Then clearly $\mathfrak{u} \neq f^{-1}(\mathfrak{v})$ for all $\mathfrak{v} \in \text{St}(B)$, so $\text{St}(f)$ is not surjective. \square

EXERCISE

2.10. Suppose $f: A \rightarrow B$ is a morphism between boolean algebras. Show that f is surjective iff $\text{St}(f)$ is injective.

3. THEORIES AND TYPES

In this chapter we discuss a few basic topics in model theory; they are closely tied to the Compactness Theorem.

3.1. Definition. An L -theory is a set T of L -sentences that is closed under logical consequence; *i.e.*, for any L -sentence σ , if $T \models \sigma$, then $\sigma \in T$.

Note. If Σ is a set of L -sentences and T is the set of L -sentences σ such that $\Sigma \models \sigma$, then T is an L -theory. In this case, we will say that T is *axiomatized* by Σ .

3.2. Definition. Let L be a first order language, let Σ be a set of L -sentences, and let \mathcal{K} be a class of L -structures.

- We let $\text{Mod}(\Sigma)$ denote the class of all L -structures that are models of Σ . (The language L is to be understood from the context.)
- The set Σ is *complete* if it is satisfiable and for every L -sentence σ , either $\Sigma \models \sigma$ or $\Sigma \models \neg\sigma$.
- An L -theory T is a *completion* of Σ if T is complete and $T \supseteq \Sigma$.
- The *theory* of \mathcal{K} is defined by

$$\text{Th}(\mathcal{K}) = \{\sigma \mid \sigma \text{ is an } L\text{-sentence and } \mathcal{M} \models \sigma \text{ for all } \mathcal{M} \in \mathcal{K}\}.$$
- For $\mathcal{K} = \{\mathcal{M}\}$ we write $\text{Th}(\mathcal{M})$ instead of $\text{Th}(\{\mathcal{M}\})$. Note that $\text{Th}(\mathcal{M})$ is a complete theory.
- We say \mathcal{K} is *axiomatizable* if $\mathcal{K} = \text{Mod}(\Sigma)$ for some Σ .

3.3. Definition. If \mathcal{M} and \mathcal{N} are two structures for the same language L , we say that \mathcal{M} and \mathcal{N} are *elementarily equivalent*, and write $\mathcal{M} \equiv \mathcal{N}$, if $\text{Th}(\mathcal{M}) = \text{Th}(\mathcal{N})$.

3.4. Definition. Let $L_1 \subseteq L_2$ be two first order languages and let Σ_i be a set of L_i -sentences for $i = 1, 2$. We say that Σ_2 is an *extension* of Σ_1 , if Σ_1 is contained in the set of logical consequences of Σ_2 . Further, Σ_2 is said to be a *conservative extension* of Σ_1 , if, in addition, $\Sigma_2 \models \sigma \Rightarrow \Sigma_1 \models \sigma$ for every sentence σ of L_1 .

3.5. Fact. If Σ_2 is an extension of Σ_1 and every model of Σ_1 has an expansion to a model of Σ_2 , then Σ_2 is a conservative extension of Σ_1 .

The following result, which is just a restatement of the Compactness Theorem, expresses a fundamental property of logical consequence in first order logic.

3.6. Corollary. *If $\Sigma \models \sigma$, then there is a finite set $\Sigma_0 \subseteq \Sigma$ with $\Sigma_0 \models \sigma$.*

Proof. Assume $\Sigma \models \sigma$, so $\Sigma \cup \{\neg\sigma\}$ has no model. Hence there exists a finite $\Sigma' \subseteq \Sigma \cup \{\neg\sigma\}$ such that Σ' has no model, by the Compactness Theorem. There exists a finite $\Sigma_0 \subseteq \Sigma$ with $\Sigma' \subseteq \Sigma_0 \cup \{\neg\sigma\}$. Evidently $\Sigma_0 \cup \{\neg\sigma\}$ cannot have a model, and therefore $\Sigma_0 \models \sigma$. \square

The following result is a variation on the same theme:

3.7. Corollary. *Let L be a first order language and let Σ, Σ' be sets of sentences in L . Suppose that for every model \mathcal{M} of Σ there exists $\gamma \in \Sigma'$ such that $\mathcal{M} \models \gamma$. Then there exists a finite subset $\{\gamma_1, \dots, \gamma_m\}$ of Σ' such that $\Sigma \models \gamma_1 \vee \dots \vee \gamma_m$.*

Proof. Apply the Compactness Theorem to $\Sigma \cup \{\neg\gamma \mid \gamma \in \Sigma'\}$. □

3.8. Definition. Let Σ be a set of L -sentences. We denote by $S_0(\Sigma)$ the set of all complete L -theories that are extensions of Σ .

Note that $S_0(\Sigma) \neq \emptyset$ iff Σ is satisfiable.

We put a natural topology on $S_0(\Sigma)$ (and regard it as the *space* of completions of Σ) as follows: for each L -sentence σ , let

$$[\sigma] = \{T \in S_0(\Sigma) \mid \sigma \in T\} = \{T \in S_0(\Sigma) \mid T \models \sigma\}.$$

Note that the family

$$\mathcal{B} = \{[\sigma] \mid \sigma \text{ is an } L\text{-sentence}\}$$

is closed under finite intersections and unions; indeed, for any L -sentences σ and τ , we see that

$$[\sigma] \cap [\tau] = [\sigma \wedge \tau] \text{ and } [\sigma] \cup [\tau] = [\sigma \vee \tau].$$

The *logic topology* on $S_0(\Sigma)$ is the topology for which \mathcal{B} is the family of basic open sets. That is, for each $T \in S_0(\Sigma)$, the basic open neighborhoods of T are the sets $[\sigma]$ where $\sigma \in T$.

Evidently this is a Hausdorff topology. Moreover, each set of the form $[\sigma]$ is closed as well as open (clopen), since

$$S_0(\Sigma) \setminus [\sigma] = [\neg\sigma]$$

for all L -sentences σ . Furthermore, the logic topology on $S_0(\Sigma)$ is *compact*; this is an immediate consequence of Corollary 3.7.

There is a close relation between closed sets in $S_0(\Sigma)$ and (not necessarily complete) L -theories T' that extend Σ . For such a theory T' , define

$$K(T') = \{T \in S_0(\Sigma) \mid T' \subseteq T\} = \bigcap \{[\sigma] \mid \sigma \in T'\}.$$

Then $K(T')$ is closed, because it is the intersection of a family of clopen sets. Conversely, if K is a closed set in $S_0(\Sigma)$, then there is a set Σ' of L -sentences such that the open set $S_0(\Sigma) \setminus K$ is equal to the union of the basic open sets $[\sigma]$ with $\sigma \in \Sigma'$. Taking T' to be the set of all logical consequences of $\Sigma \cup \{\neg\sigma \mid \sigma \in \Sigma'\}$ in L , we have that T' extends Σ and $K(T') = K$. The correspondence between T' and $K(T')$ is a bijection between L -theories T' that extend Σ and closed subsets in $S_0(\Sigma)$. Note that $K(T')$ is nonempty iff T' is satisfiable.

3.9. Proposition. *Let Σ be a satisfiable set of L -sentences. The space $S_0(\Sigma)$ with the logic topology is a totally disconnected, compact Hausdorff space. Its closed sets are the sets of the form*

$$K(T') = \{T \in S_0(\Sigma) \mid T' \subseteq T\}$$

where T' is an L -theory extending Σ . Moreover, the clopen subsets of $S_0(\Sigma)$ in this topology are exactly the sets of the form $[\sigma]$, where σ is an L -sentence.

Proof. It remains only to prove that each clopen set $C \subseteq S_0(\Sigma)$ is of the form $[\sigma]$ for some sentence σ . Since C is open, it is the union of a family of basic open sets. Since C is closed, hence compact, this family can be taken to be finite. In $S_0(\Sigma)$, a union of finitely many basic open sets is itself a basic open set. \square

3.10. Fact. Let Σ be a satisfiable L -theory and let σ, τ be L -sentences. Then $[\sigma] = [\tau]$ in $S_0(\Sigma)$ iff σ and τ are equivalent over Σ (i.e., $\Sigma \models (\sigma \leftrightarrow \tau)$).

Next we discuss how the results above show how to identify $S_0(\Sigma)$ with the Stone space of a boolean algebra of obvious logical significance.

Let $B_0(\Sigma)$ be the set of equivalence classes of L -sentences under the equivalence relation \equiv_Σ of logical equivalence over Σ ($\sigma \equiv_\Sigma \tau$ iff $\Sigma \models (\sigma \leftrightarrow \tau)$). For each L -sentence σ , let σ/Σ denote the equivalence class of σ under this relation. We make $B_0(\Sigma)$ into a Boolean algebra by taking 0 to be the class of an inconsistent sentence and 1 to be the class of a valid sentence, and by defining $\sigma/\Sigma \vee \tau/\Sigma = (\sigma \vee \tau)/\Sigma$ and $\sigma/\Sigma \wedge \tau/\Sigma = (\sigma \wedge \tau)/\Sigma$, and $(\sigma/\Sigma)' = (\neg\sigma)/\Sigma$. Note that $\sigma/\Sigma \leq \tau/\Sigma$ in the ordering on $B_0(\Sigma)$ iff $\Sigma \models (\sigma \rightarrow \tau)$. If Σ is satisfiable, then $0 \neq 1$ in $B_0(\Sigma)$; otherwise, $0 = 1$ and so $B_0(\Sigma)$ is the trivial (1-element) boolean algebra.

The discussion above proves the following result:

3.11. Proposition. *Let Σ be a satisfiable set of L -sentences.*

(1) *The topological space $S_0(\Sigma)$ is homeomorphic to the Stone space of the boolean algebra $B_0(\Sigma)$, under the correspondence taking a complete theory T extending Σ to the ultrafilter $\{\sigma/\Sigma \mid \sigma \in T\}$ on $B_0(\Sigma)$.*

(2) *In particular, this correspondence maps an arbitrary closed subset $K = K(T')$ of $S_0(\Sigma)$ to the filter $\{\sigma/\Sigma \mid T' \models \sigma\}$ on $B_0(\Sigma)$. Every proper filter on $B_0(\Sigma)$ is of this form for some satisfiable L -theory $T' \supseteq \Sigma$.*

Proof. Exercise for the reader. \square

Types

Next we introduce *types*; they provide a way of describing the first order expressible properties of elements in a structure. To do this, we simply extend what is done earlier in this chapter from sentences to formulas.

Let $x = (x_i \mid i \in I)$ be a fixed family of distinct variables. Let $\Phi(x)$ be a set of L -formulas $\varphi(x)$ whose free variables all occur among $(x_i \mid i \in I)$. We call $\Phi(x)$ a *partial x -type (in L)*.

Let Σ be a set of L -sentences and $\Phi(x)$ a partial x -type in L . We say $\Phi(x)$ is Σ -*realizable* if there exists a model \mathcal{M} and a family $a \in M^I$ such that $\mathcal{M} \models \varphi[a]$ for every $\varphi(x) \in \Phi(x)$. In this situation we also say that $\Phi(x)$ is *realizable* in \mathcal{M} and that $\Phi(x)$ is *realized by a in \mathcal{M}* .

3.12. Lemma. *Let Σ be a set of L -sentences and $\Phi(x)$ a partial x -type in L . The following are equivalent:*

- (1) $\Phi(x)$ is a maximal Σ -realizable partial x -type in L .
- (2) $\Phi(x)$ is Σ -realizable and for each L -formula $\varphi(x)$, either $\varphi(x) \in \Phi(x)$ or $\neg\varphi(x) \in \Phi(x)$.
- (3) There is a model \mathcal{M} of Σ and a family $a \in M^I$ such that

$$\Phi(x) = \{\varphi(x) \mid \varphi(x) \text{ is an } L\text{-formula and } \mathcal{M} \models \varphi[a]\}.$$

Proof. Exercise. □

3.13. Definition. Let Σ be a set of L -sentences.

- Any $\Phi(x)$ satisfying the equivalent conditions in 3.12 is called a Σ -*realizable x -type in L* .
- We denote the set of all Σ -realizable x -types in L by $S_x(\Sigma)$.
- If \mathcal{M} is a model of Σ and $a \in M^I$, we denote the set

$$\{\varphi(x) \mid \varphi(x) \text{ is an } L\text{-formula and } \mathcal{M} \models \varphi[a]\}$$

by $\text{tp}_{\mathcal{M},x}(a)$ or simply by variants such as $\text{tp}_{\mathcal{M}}(a)$, $\text{tp}_x(a)$ or even $\text{tp}(a)$ if the variables or the structure or both are understood. We refer to this set as *the type of a in \mathcal{M}* .

3.14. Remark. By Lemma 3.12 we see that every $\Phi(x) \in S_x(\Sigma)$ is of the form $\text{tp}_{\mathcal{M},x}(a)$ for some model \mathcal{M} of Σ and some $a \in M^I$.

3.15. Notation. If x consists of n distinct variables, it is common to write $S_n(\Sigma)$ instead of $S_x(\Sigma)$, in situations where the precise variables used are not important or are understood. In such a situation, the elements of $S_n(\Sigma)$ may be referred to as *n -types*.

Next we discuss the relation between types and complete theories. As above, we have a fixed language L and family $x = (x_i \mid i \in I)$ of distinct variables, and a set Σ of L -sentences.

Let $c = (c_i \mid i \in I)$ be a family of distinct new constant symbols and let L_c be the language $L(c_i \mid i \in I)$ extending L . To avoid confusion, we write Σ_c for Σ considered as a set of L_c -sentences. Based on the simple observation given in the next result, we may identify $S_x(\Sigma)$ with $S_0(\Sigma_c)$

using the bijection that takes an x -type $\Phi(x)$ to the set of L_c -sentences $\Phi(c)$.

3.16. Lemma. *Let \mathcal{M}, \mathcal{N} be L -structures, $a \in M^I$, and $b \in N^I$. The x -type $\text{tp}_{\mathcal{M}}(a)$ can be identified with the complete L_c -theory $\text{Th}(\mathcal{M}, a_i)_{i \in I}$. In particular, $\text{tp}_{\mathcal{M}}(a) = \text{tp}_{\mathcal{N}}(b)$ if and only if $(\mathcal{M}, a_i)_{i \in I} \equiv_{L(c)} (\mathcal{N}, b_i)_{i \in I}$*

Proof. Set $\Phi(x) = \text{tp}_{\mathcal{M}}(a)$, and consider the set of formulas $\Phi(c) = \{\varphi(c) \mid \varphi(x) \in \Phi(x)\}$ in the language L_c . Evidently $\Phi(c) \subseteq \text{Th}(\mathcal{M}, a_i)_{i \in I}$. Moreover, it is an easy exercise in changing bound variables to show that every sentence in $\text{Th}(\mathcal{M}, a_i)_{i \in I}$ is logically equivalent to a sentence in $\Phi(c)$. \square

3.17. Remark. Let $\Phi(x)$ be a partial x -type in L and $\varphi(x)$ an L -formula; let σ be a set of L -sentences.

- $\Sigma \models_L \varphi(x) \iff \Sigma \models_{L(c)} \varphi(c)$;
- $\Phi(x)$ is Σ -realizable $\iff \Sigma \cup \Phi(c)$ has a model;
- $\Phi(x)$ is a Σ -realizable x -type $\iff \Sigma \subseteq \Phi(c)$ and $\Phi(c)$ is a complete $L(c)$ -theory.

We define the *logic topology* on the space of x -types $S_x(\Sigma)$ so that the bijection by which we identify $S_x(\Sigma)$ with $S_0(\Sigma_c)$ is a homeomorphism, when we put the logic topology defined above on $S_0(\Sigma_c)$. That is, the basic open sets for the logic topology on $S_x(\Sigma)$ are the sets of the form

$$[\varphi(x)] = \{\Phi(x) \in S_x(\Sigma) \mid \varphi \in \Phi\}$$

where $\varphi(x)$ varying over the L -formulas whose free variables are among $(x_i \mid i \in I)$. As in our discussion of $S_0(\Sigma)$, this yields that the clopen subsets in $S_x(\Sigma)$ are exactly the basic open sets defined above. It also yields that the closed sets in $S_x(\Sigma)$ are the sets of the form

$$K(\Phi'(x)) := \{\Phi(x) \in S_x(\Sigma) \mid \Phi'(x) \subseteq \Phi(x)\}$$

where $\Phi'(x)$ ranges over partial x -types in L .

We also introduce the boolean algebra $B_x(\Sigma)$ of equivalence classes of L -formulas $\varphi(x)$, relative to the relation of logical equivalence over Σ . As done above for sentences, we denote the equivalence class of $\varphi(x)$ by $\varphi(x)/\Sigma$ and define the boolean algebra operations on these classes by:

$\varphi(x)/\Sigma \vee \psi(x)/\Sigma = (\varphi(x) \vee \psi(x))/\Sigma$ and $\varphi(x)/\Sigma \wedge \psi(x)/\Sigma = (\varphi(x) \wedge \psi(x))/\Sigma$, and $(\varphi(x)/\Sigma)' = (\neg\varphi(x))/\Sigma$.

The following result is immediate from Corollary 3.9 and the preceding discussion.

3.18. Proposition. *Let Σ be a satisfiable set of L -sentences and let $x = (x_i \mid i \in I)$ be a family of distinct variables.*

(1) *The topological space $S_x(\Sigma)$ is homeomorphic to the Stone space of the boolean algebra $B_x(\Sigma)$, under the correspondence taking a Σ -realizable x -type $\Phi(x)$ to the ultrafilter $\{\varphi(x)/\Sigma \mid \varphi(x) \in \Phi(x)\}$ on $B_x(\Sigma)$.*

(2) *In particular, this correspondence maps an arbitrary closed subset $K(\Phi'(x))$ of $S_x(\Sigma)$ to the filter $\{\varphi(x)/\Sigma \mid \Phi'(x) \models \varphi(x)\}$ on $B_x(\Sigma)$. Every*

proper filter on $B_x(\Sigma)$ is of this form for some Σ -realizable partial x -type $\Phi'(x)$ in L .

Types over a set of parameters

Later we will need the formalism of x -types over A , where A is a subset of a model \mathcal{M} of an L -theory T . In such a situation, we denote $(\mathcal{M}, a)_{a \in A}$ by \mathcal{M}_A and take T_A to be $\text{Th}(\mathcal{M}_A)$; thus T_A is a complete L_A -theory. It specifies the elementary properties of elements of A within a model \mathcal{M} of T . (The model \mathcal{M} is arbitrary except that $A \subseteq M$ and $\mathcal{M}_A \models T_A$. Note that any model of T_A is isomorphic to an L_A -structure of the form \mathcal{M}_A , where $\mathcal{M} \models T$ and $A \subseteq M$.)

3.19. Definition. An x -type over A for the theory T is a T_A -realizable x -type in L_A . The space of all x -types over A for the theory T , namely the space $S_x(T_A)$, will be denoted by $S_x(A)$ if the theory T and model \mathcal{M} containing A are understood. If the model \mathcal{M} needs to be specified, to avoid confusion, we will denote $S_x(T_A)$ by $S_x^{\mathcal{M}}(A)$.

Further, if $\mathcal{M} \models T$, $A \subseteq M$, and $a \in M^I$, then we will denote the type of a in \mathcal{M}_A by $tp_{\mathcal{M}}(a/A)$.

3.20. Fact. Let T be an L -theory, \mathcal{M} a model of T , and A a subset of M . Let $\Phi(x)$ be an x -type in L_A . Then $\Phi(x) \in S_x(A)$ iff $\Phi(x)$ is finitely satisfiable in the given structure \mathcal{M}_A .

An application of type spaces

To close this chapter we give an application of the topology of type spaces that will be used later (for example, when we consider quantifier elimination and model completeness).

Let T be a satisfiable L -theory. Let $\Delta(x)$ be a nonempty set of L -formulas whose free variables are among $x = (x_1, \dots, x_n)$, with $n \geq 0$.

3.21. Proposition. Assume that $\Delta(x)$ is closed under disjunction and conjunction (up to equivalence over T). Let $\varphi(x)$ be an L -formula. The following are equivalent:

- (1) $T \models \varphi$ or $T \models \neg\varphi$ or $T \models \varphi \leftrightarrow \psi$ for some formula $\psi \in \Delta(x)$.
- (2) For every $p(x), q(x) \in S_x(T)$, if $\varphi \in p(x)$ and $\neg\varphi \in q(x)$, then there exists $\psi \in \Delta(x)$ such that $\psi \in p(x)$ and $\neg\psi \in q(x)$.

Proof. (1 \Rightarrow 2): If $\varphi \in p(x)$ and $\neg\varphi \in q(x)$, then neither $T \models \varphi$ nor $T \models \neg\varphi$ hold. Thus there exists $\psi \in \Delta(x)$ such that $T \models \varphi \leftrightarrow \psi$. It follows that $\psi \in p(x)$ and $\neg\psi \in q(x)$.

(2 \Rightarrow 1): Assume that condition (2) holds and that neither $T \models \varphi$ nor $T \models \neg\varphi$. Let K denote the clopen set $[\varphi]$ in $S_x(T)$. Note that both K and its complement K^c are nonempty. Let $\mathcal{S}_{\Delta}(x)$ be the family of basic open sets of the form $[\psi]$ where $\psi \in \Delta(x)$.

We will first show that K is the union of a family of basic open sets from $\mathcal{S}_{\Delta}(x)$. Fix $p(x) \in K$; condition (2) implies that there is a subset $\Delta(x)'$ of

$\Delta(x)$ such that $\bigcup\{[\neg\psi'] \mid \psi' \in \Delta(x)'\}$ contains K^c as a subset and does not have $p(x)$ as an element. Since K^c is compact, the set $\Delta(x)'$ can be taken to be finite. Since $\Delta(x)$ is closed under conjunction, there is a single formula ψ' from $\Delta(x)$ such that $K^c \subseteq [\neg\psi']$ and $p(x) \in [\psi']$. That is, $p(x) \in [\psi'] \subseteq K$. Therefore K is the union of a family of basic open sets from $\mathfrak{S}_\Delta(x)$.

Since K is compact, it is a *finite* union of such basic open sets. Since $\Delta(x)$ is closed under disjunction, there must be a single formula $\psi \in \Delta(x)$ such that $K = [\psi]$, and therefore $T \models \varphi \leftrightarrow \psi$, as desired. \square

3.22. Corollary. *Consider the setting of Proposition 3.21. Assume that $\Delta(x)$ is closed under negation, disjunction and conjunction (up to equivalence over T). The following are equivalent:*

- (1) *Every L -formula $\varphi(x)$ is T -equivalent to a formula in $\Delta(x)$.*
- (2) *For every distinct $p(x), q(x) \in S_x(T)$, there is a formula $\psi(x) \in \Delta(x)$ such that $\psi(x) \in p(x)$ and $\neg\psi(x) \in q(x)$.*

Proof. This follows quickly from Proposition 3.21. Note that the set $\Delta(x)$ considered here contains a valid formula and an inconsistent formula.

Because of the usefulness of this result, and to illustrate boolean algebra ideas, we sketch an alternate proof of $(2 \Rightarrow 1)$ that does not depend on 3.21. Assuming (2), let A be the boolean subalgebra of $B_x(T)$ that is generated by $\{\varphi(x)/T \mid \varphi(x) \in \Delta(x)\}$. To prove (1), we need to show $B_x(T) = A$. The inclusion map $i: A \rightarrow B_x(T)$ induces a continuous map $\text{St}(i): \text{St}(B_x(T)) \rightarrow \text{St}(A)$, under which an ultrafilter \mathfrak{u} on $B_x(T)$ is mapped to the ultrafilter $\mathfrak{u} \cap A$ on A . Because $\text{St}(B_x(T)) = S_x(T)$, condition (2) implies that $\text{St}(i)$ is injective. Exercise 2.10 yields that i is surjective, and hence $B_x(T) = A$. \square

3.23. Remark. An important case of Corollary 3.22 is when $\Delta(x)$ consists of the quantifier free L -formulas with all free variables in x .

EXERCISES

3.24. Show that the Compactness Theorem (Theorem 1.10) can be derived from Corollary 3.6 by a trivial argument.

3.25. Let T be an L -theory and let \mathcal{K} be the set of all L -structures that are *not* models of T . Show that T is axiomatizable by a finite set of L -sentences iff \mathcal{K} is axiomatizable.

4. ELEMENTARY MAPS

4.1. Definition. Let \mathcal{M}, \mathcal{N} be L -structures and let f be a function from a subset A of M into N . We say f is *elementary* (with respect to \mathcal{M}, \mathcal{N}) if for every L -formula $\varphi(x_1, \dots, x_m)$ and every $a_1, \dots, a_m \in A$

$$\mathcal{M} \models \varphi[a_1, \dots, a_m] \Leftrightarrow \mathcal{N} \models \varphi[f(a_1), \dots, f(a_m)].$$

If the domain of the function f is all of M and f is elementary with respect to \mathcal{M}, \mathcal{N} , then f is called an *elementary embedding from \mathcal{M} into \mathcal{N}* .

4.2. Fact. Let $\mathcal{M}, \mathcal{N}, A, f$ be as in 4.1. The function f is elementary with respect to \mathcal{M}, \mathcal{N} if and only if $(\mathcal{M}, a)_{a \in A} \equiv (\mathcal{N}, f(a))_{a \in A}$. In particular, if there exists a function $f: A \rightarrow N$ that is elementary with respect to \mathcal{M}, \mathcal{N} for some subset A of M (including the empty set), then $\mathcal{M} \equiv \mathcal{N}$.

4.3. Fact. If f is an elementary function then f must be 1-1. Moreover, if f is an elementary embedding of \mathcal{M} into \mathcal{N} , then f is an embedding of \mathcal{M} into \mathcal{N} .

4.4. Definition. Let \mathcal{M}, \mathcal{N} be L -structures and suppose $M \subseteq N$. We say \mathcal{M} is an *elementary substructure* of \mathcal{N} and write $\mathcal{M} \preceq \mathcal{N}$ if the inclusion map is an elementary embedding from \mathcal{M} into \mathcal{N} . In this case we also refer to \mathcal{N} as an *elementary extension* of \mathcal{M} and write $\mathcal{N} \succeq \mathcal{M}$.

The importance of elementary extensions for model theoretic arguments is indicated by the following remark.

4.5. Remark. Let L be a first order language and let \mathcal{M}, \mathcal{N} be L -structures that satisfy $\mathcal{M} \preceq \mathcal{N}$. An important property of elementary extensions is that each relation R on M that is definable in \mathcal{M} has a canonical extension R' to a relation on N that is definable in \mathcal{N} .

To obtain this extension, take any L -formula $\varphi(x_1, \dots, x_m, y_1, \dots, y_n)$ and any $b_1, \dots, b_n \in M$ such that

$$R = \{(a_1, \dots, a_m) \in M^m \mid \mathcal{M} \models \varphi[a_1, \dots, a_m, b_1, \dots, b_n]\}$$

and set

$$R' = \{(a_1, \dots, a_m) \in N^m \mid \mathcal{N} \models \varphi[a_1, \dots, a_m, b_1, \dots, b_n]\}$$

It is an easy exercise to show that R' does not depend on the specific $L(M)$ -formula $\varphi(x_1, \dots, x_m, b_1, \dots, b_n)$ used in defining R . Note that the parameters needed to define R' in \mathcal{N} are exactly the same as the parameters used to define R in \mathcal{M} .

This correspondence between (certain) relations R on M and R' on N preserves all structural properties that can be expressed in first order logic. For example: it is an isomorphism with respect to Boolean operations and projections; also, if R is the graph of a function, then so is R' .

4.6. Facts. (a) Let g be an elementary embedding from \mathcal{M} into \mathcal{N} , and let f be any restriction of g to a subset A of M . Then f is elementary with

respect to \mathcal{M}, \mathcal{N} . (In particular, this holds if g is an isomorphism from \mathcal{M} onto \mathcal{N} .)

(b) If f_1 is elementary (with respect to $\mathcal{M}_1, \mathcal{M}_2$) and f_2 is elementary (with respect to $\mathcal{M}_2, \mathcal{M}_3$), and if the range of f_1 is contained in the domain of f_2 , then the composition $f_2 \circ f_1$ is elementary (with respect to $\mathcal{M}_1, \mathcal{M}_3$).

(c) If f is elementary (with respect to \mathcal{M}, \mathcal{N}), then f^{-1} is also elementary (with respect to \mathcal{N}, \mathcal{M}).

4.7. Fact. Let I be an index set and U be an ultrafilter on I . Fix a first order language L and an L -structure \mathcal{M} . Consider the ultrapower \mathcal{M}^I/U of \mathcal{M} . Define a function δ on M by setting $\delta(a) = g_a/U$, where g_a is the constant function with $g_a(i) = a$ for all $i \in I$. Then δ is an elementary embedding from \mathcal{M} into \mathcal{M}^I/U . (This is called the *diagonal embedding*; often one identifies a with $\delta(a)$ for each $a \in M$ and thereby regards \mathcal{M} as an elementary substructure of \mathcal{M}^I/U .)

The following result gives a useful tool for showing that \mathcal{M} is an elementary substructure of \mathcal{N} . Note that the condition in this Theorem refers to truth of formulas *only* in the structure \mathcal{N} .

4.8. Theorem (Tarski-Vaught Test for \preceq). *Let \mathcal{N} be an L -structure and suppose $M \subseteq N$. Then M is the underlying set of an elementary substructure of \mathcal{N} if and only if for every formula $\psi(x_1, \dots, x_m, y)$ in L and every sequence a_1, \dots, a_m in M , if $\mathcal{N} \models \exists y \psi[a_1, \dots, a_m]$, then there exists $b \in M$ such that $\mathcal{N} \models \psi[a_1, \dots, a_m, b]$.*

Proof. (\Rightarrow) This follows immediately from the definition of elementary substructure.

(\Leftarrow): Suppose M and \mathcal{N} satisfy the given conditions. We first need to show that M is the underlying set of a substructure of \mathcal{N} . If c is a constant symbol in L , apply the given conditions on M, \mathcal{N} to the formula $\psi(y)$ equal to $y = c$; this shows that $c^{\mathcal{N}} \in M$. If F is an m -ary function symbol in L , apply the given conditions on M, \mathcal{N} to the formula $\psi(x_1, \dots, x_m, y)$ equal to $F(x_1, \dots, x_m) = y$; this shows that M is closed under the function $F^{\mathcal{N}}$. Hence there exists $\mathcal{M} \subseteq \mathcal{N}$ whose underlying set is M .

We need to show that for any formula $\varphi(x_1, \dots, x_m)$ and any $a_1, \dots, a_m \in M$,

$$\mathcal{M} \models \varphi[a_1, \dots, a_m] \Leftrightarrow \mathcal{N} \models \varphi[a_1, \dots, a_m].$$

The proof is by induction on the formula φ . By changing bound variables if necessary, we may restrict attention to formulas $\varphi(x_1, \dots, x_m)$ that have no bound occurrences of any x_j , $j = 1, \dots, m$.

In the basis step φ is an atomic formula; the displayed equivalence follows from the fact that \mathcal{M} is a substructure of \mathcal{N} .

In the induction step, the cases of propositional connectives are trivial. In the remaining case φ is of the form $\exists y \psi(x_1, \dots, x_m, y)$, where the statement

to be proved is assumed to be true for ψ and y is not among x_1, \dots, x_m . Then we have:

$$\begin{aligned}
\mathcal{M} \models \varphi &\Leftrightarrow \mathcal{M} \models \exists y \psi[a_1, \dots, a_m] \\
&\Leftrightarrow \mathcal{M} \models \psi[a_1, \dots, a_m, b] \text{ for some } b \in M \\
&\Leftrightarrow \mathcal{N} \models \psi[a_1, \dots, a_m, b] \text{ for some } b \in M \\
&\Leftrightarrow \mathcal{N} \models \exists y \psi[a_1, \dots, a_m] \\
&\Leftrightarrow \mathcal{N} \models \varphi[a_1, \dots, a_m]
\end{aligned}$$

In the third equivalence we used the induction hypothesis and in the fourth we used the hypothesis of the implication we are proving as well as the fact that y is distinct from all of x_1, \dots, x_m . \square

4.9. Facts (Unions of Chains). Let (I, \leq) be a linearly ordered set. For each $i \in I$ let \mathcal{M}_i be an L -structure, and suppose this indexed family of structures is a chain. That is, for each $i, j \in I$, we suppose $i \leq j \Rightarrow \mathcal{M}_i \subseteq \mathcal{M}_j$.

(1) There is a well defined structure whose universe is the union of the sets M_i and which is an extension of each \mathcal{M}_i ; moreover, such a structure is unique. (For obvious reasons, this structure is called the *union* of the given chain of structures.)

(2) If, in addition, $\mathcal{M}_i \preceq \mathcal{M}_j$ holds whenever $i, j \in I$ and $i \leq j$, then the union of this chain of structures is an elementary extension of each \mathcal{M}_i . (In this situation we refer to $(\mathcal{M}_i \mid i \in I)$ as an *elementary chain* of L -structures.)

A useful way of proving that functions are elementary is the *back-and-forth* method, a version of which we now describe.

4.10. Definition. Let \mathcal{M} and \mathcal{N} be L -structures. A *partial isomorphism* from \mathcal{M} to \mathcal{N} is a bijection $f: A \rightarrow B$ with $A \subseteq M$ and $B \subseteq N$ such that:

(i) for each (m -ary) relation symbol R in L and all $a_1, \dots, a_m \in A$,

$$R^{\mathcal{M}}(a_1, \dots, a_m) \iff R^{\mathcal{N}}(f(a_1), \dots, f(a_m));$$

(ii) For each (n -ary) function symbol F in L (including constant symbols, with $n = 0$) and all $a_1, \dots, a_n, a_{n+1} \in B$,

$$F^{\mathcal{M}}(a_1, \dots, a_n) = a_{n+1} \iff F^{\mathcal{N}}(f(a_1), \dots, f(a_n)) = f(a_{n+1}).$$

4.11. Definition. Let \mathcal{M}, \mathcal{N} be L -structures and let \mathcal{F} be a collection of partial isomorphisms from \mathcal{M} to \mathcal{N} . We say \mathcal{F} is a *back-and-forth system* from \mathcal{M} to \mathcal{N} if it has the following properties:

- (“back”) For each $f \in \mathcal{F}$ and each $b \in N$ there is some $g \in \mathcal{F}$ such that g extends f and b is in the range of g .
- (“forth”) For each $f \in \mathcal{F}$ and each $a \in M$ there is some $g \in \mathcal{F}$ such that g extends f and a is in the domain of g .

4.12. Lemma. Let \mathcal{F} be a back-and-forth system from \mathcal{M} to \mathcal{N} and let $t(x_1, \dots, x_n)$ be an L -term. Every $f \in \mathcal{F}$ has the following property: for all $a_1, \dots, a_n, a_{n+1} \in \text{dom}(f)$,

$$t^{\mathcal{M}}(a_1, \dots, a_n) = a_{n+1} \iff t^{\mathcal{N}}(f(a_1), \dots, f(a_n)) = f(a_{n+1}).$$

Proof. We argue by induction of the number of function symbols that occur in t . The equivalence is trivial if t is x_i for some $i = 1, \dots, n$. If t contains a single function symbol F , then t is of the form $F(y_1, \dots, y_m)$ where y_1, \dots, y_m are among x_1, \dots, x_n , and the desired equivalence is given in the definition of partial isomorphism. For the induction step, t is of the form $F(t_1, \dots, t_m)$ where each t_j is a term with strictly fewer occurrences of function symbols than t . For $j = 1, \dots, m$, let $a'_j = t_j^{\mathcal{M}}(a_1, \dots, a_n)$, and note that $t^{\mathcal{M}}(a_1, \dots, a_n) = F^{\mathcal{M}}(a'_1, \dots, a'_m)$. By assumption, there exists $g \in \mathcal{F}$ that extends f and is defined on a'_1, \dots, a'_m . By the definition of a'_j and the induction assumption applied to g , we have $g(a'_j) = t_j^{\mathcal{N}}(g(a_1), \dots, g(a_n))$ for $j = 1, \dots, m$, and hence $t^{\mathcal{N}}(f(a_1), \dots, f(a_n)) = t^{\mathcal{N}}(g(a_1), \dots, g(a_n)) = F^{\mathcal{N}}(g(a'_1), \dots, g(a'_m))$. Since g is a partial isomorphism, we have

$$F^{\mathcal{M}}(a'_1, \dots, a'_m) = a_{n+1} \iff F^{\mathcal{N}}(g(a'_1), \dots, g(a'_m)) = g(a_{n+1}).$$

From the equations above and the fact that $g(a_{n+1}) = f(a_{n+1})$ we conclude

$$t^{\mathcal{M}}(a_1, \dots, a_n) = a_{n+1} \iff t^{\mathcal{N}}(f(a_1), \dots, f(a_n)) = f(a_{n+1}).$$

as desired. \square

4.13. Proposition. *Let \mathcal{M}, \mathcal{N} be L -structures and let \mathcal{F} be a nonempty back-and-forth system from \mathcal{M} to \mathcal{N} . Then each function in \mathcal{F} is elementary with respect to \mathcal{M}, \mathcal{N} . In particular, $\mathcal{M} \equiv \mathcal{N}$.*

Proof. Let $\varphi(x_1, \dots, x_n)$ be an L -formula, f a function in \mathcal{F} , and a_1, \dots, a_n elements of the domain of f . We must prove

$$\mathcal{M} \models \varphi[a_1, \dots, a_n] \iff \mathcal{N} \models \varphi[f(a_1), \dots, f(a_n)].$$

This is done by induction on $\varphi(x_1, \dots, x_n)$.

In the base case of the induction φ is an atomic formula, of the form $R(t_1, \dots, t_m)$ where t_1, \dots, t_m are L -terms whose variables are among x_1, \dots, x_n . For $j = 1, \dots, m$, let $a'_j = t_j^{\mathcal{M}}(a_1, \dots, a_n)$. Note that $\varphi^{\mathcal{M}}(a_1, \dots, a_n)$ and $R^{\mathcal{M}}(a'_1, \dots, a'_m)$ have the same truth value.

Let $g \in \mathcal{F}$ extend f and be defined on a'_1, \dots, a'_m . By Lemma 4.12 we have $g(a'_j) = t_j^{\mathcal{N}}(g(a_1), \dots, g(a_n))$ which is the same as $t_j^{\mathcal{N}}(f(a_1), \dots, f(a_n))$. Hence $\varphi^{\mathcal{N}}(f(a_1), \dots, f(a_n))$ and $R^{\mathcal{N}}(g(a'_1), \dots, g(a'_m))$ have the same truth value.

Finally, since g is a partial isomorphism, we have

$$R^{\mathcal{M}}(a'_1, \dots, a'_m) \iff R^{\mathcal{N}}(g(a'_1), \dots, g(a'_m))$$

from which we conclude

$$\varphi^{\mathcal{M}}(a_1, \dots, a_n) \iff \varphi^{\mathcal{N}}(f(a_1), \dots, f(a_n))$$

as desired.

The induction steps for propositional connectives are trivial. The induction steps for quantifiers follow from the “back-and-forth” properties satisfied by \mathcal{F} .

The final statement follows because \mathcal{F} is nonempty. \square

Theory of dense linear orderings without endpoints

We illustrate the use of these ideas by treating the theory of dense linear orderings without endpoints. Let L denote the language whose only nonlogical symbol is a binary predicate symbol $<$. Let DLO denote the theory of dense linear orderings without maximum or minimum element, axiomatized by a set of L -sentences.

4.14. Example. Each L -formula is equivalent in DLO to a quantifier-free L -formula.

Proof. We apply Corollary 3.22. Fix an L -formula $\varphi(x_1, \dots, x_m)$. Let Σ be the set of quantifier-free L -formulas whose free variables are among x_1, \dots, x_m . We will verify condition (2) of Corollary 3.22. To that end, consider two dense linear orderings without endpoints $(M, <)$ and $(N, <)$ and elements $a_1, \dots, a_m \in M, b_1, \dots, b_m \in N$. We assume that $(M, <) \models \varphi[a_1, \dots, a_m]$ and that every quantifier-free L -formula satisfied in $(M, <)$ by a_1, \dots, a_m is satisfied in $(N, <)$ by b_1, \dots, b_m . We need to show $(N, <) \models \varphi[b_1, \dots, b_m]$.

Let \mathcal{F} be the set of all order preserving functions from a finite subset of M into N . An easy argument shows that \mathcal{M} is a local isomorphism from $(M, <)$ onto $(N, <)$. Our assumptions ensure that there exists $f \in \mathcal{F}$ such that f is defined on $\{a_1, \dots, a_m\}$ and satisfies $f(a_i) = b_i$ for all $i = 1, \dots, m$. By Proposition 4.13, the function f is elementary with respect to $(M, <)$ and $(N, <)$. Since $(M, <) \models \varphi[a_1, \dots, a_m]$, we conclude $(N, <) \models \varphi[b_1, \dots, b_m]$, as needed. \square

Note that we have proved in passing that every two models of DLO are elementarily equivalent, since there is a local isomorphism from one onto the other. Hence DLO is complete.

Theory of equality

We complete this chapter by analyzing the simplest logical theory, which is the theory of equality. Let L denote the language of $=$, without any nonlogical symbols. Note that an L -structure is simply a nonempty set. For each $n \geq 0$ let σ_n be a sentence in L that expresses the statement that the universe has at most n elements (so $\neg\sigma_0$ is logically valid). For each $n \geq 1$ let T_n be the theory in L axiomatized by $\sigma_n \wedge \neg\sigma_{n-1}$ and let T_∞ be the theory axiomatized by the set $\{\neg\sigma_n \mid n \geq 1\}$. Thus T_n is the theory of sets of size n ($n \geq 1$) and T_∞ is the theory of infinite sets.

4.15. Example (Theories in the language of equality).

- (i) Each formula in the pure language of $=$ is logically equivalent to a Boolean combination of quantifier free formulas and the sentences σ_n for $n \geq 1$.
- (ii) The complete theories in the language of $=$ are equivalent to T_∞ and T_n for $n \geq 1$. For each such theory T , every formula in the language of $=$ is equivalent in T to a quantifier free formula.

Proof. (i) We apply Corollary 3.22. Fix a formula $\varphi(x_1, \dots, x_m)$ in the language of $=$. Let Δ be the set of all Boolean combinations of quantifier free formulas whose variables are among x_1, \dots, x_m and the sentences σ_n for $n \geq 1$. We want to verify condition (2) of Corollary 3.22. To that end, consider two sets M, N and elements $a_1, \dots, a_m \in M, b_1, \dots, b_m \in N$. We assume that $M \models \varphi[a_1, \dots, a_m]$ and that every formula in Δ that is satisfied by a_1, \dots, a_m in M is satisfied by b_1, \dots, b_m in N . We need to show $N \models \varphi[b_1, \dots, b_m]$.

Our hypotheses ensure that for all $n \geq 1$ we have $M \models \sigma_n \iff N \models \sigma_n$. Therefore, either M and N have the same finite cardinality or both M, N are infinite. Moreover, for each $1 \leq i < j \leq m$, we also have that $a_i = a_j \iff b_i = b_j$. Therefore there is a bijection f from $\{a_1, \dots, a_m\}$ onto $\{b_1, \dots, b_m\}$ such that $f(a_i) = b_i$ for all $i = 1, \dots, m$. Let \mathcal{F} be the set of all 1-1 functions g that extend f and map a finite subset of M into N . It is easy to check that \mathcal{F} is a local isomorphism from (M, a_1, \dots, a_m) onto (N, b_1, \dots, b_m) . By Proposition 4.13, we have $(M, a_1, \dots, a_m) \equiv (N, b_1, \dots, b_m)$, and hence $N \models \varphi[b_1, \dots, b_m]$.

(ii) For finite n , any two models of T_n are isomorphic, hence elementarily equivalent, so T_n is complete in these cases. On the other hand, T_∞ has only infinite models; the back-and-forth argument used to prove (i) shows that any two infinite sets are elementarily equivalent, which proves that T_∞ is also complete. If T is any complete theory in the language of equality, and M is one of its models, then M is a model of T_∞ or of T_n for some $n \geq 1$, depending on the cardinality of M . Therefore T is equal to $\text{Th}(M)$, which contains one of these theories, say T_j where $j \geq 1$ or $j = \infty$. But we showed that each such T_j is complete, from which it follows easily that T and T_j are equal. \square

EXERCISES

4.16. Let L be the language whose only nonlogical symbol is a binary predicate symbol $<$. Let \mathcal{M} be an L -structure that is a dense linear ordering without endpoints. Let $\varphi(x, y_1, \dots, y_n)$ be any L -formula (with x a single variable) and let $a_1, \dots, a_n \in M$. Show that the definable set

$$\{a \in M \mid \mathcal{M} \models \varphi[a, a_1, \dots, a_n]\}$$

is the union of a finite number of open intervals (whose endpoints are in $M \cup \{-\infty, +\infty\}$) and a finite subset of M . (Remark: this property of the ordered structures $\langle M, < \rangle \models \text{DLO}$ is expressed by saying that $\langle M, < \rangle$ is *o-minimal*. The study of o-minimal ordered structures more generally has been an important theme in model theory for several decades.)

4.17. Show that if T is the empty theory in the language L of equality, then the space $S_0(T)$ consists of a sequence of points $(T_n \mid n \geq 1)$ that are isolated, together with a point T_∞ to which this sequence converges.

4.18. Let L be the pure language of $=$, so L has no nonlogical symbols, and let σ be any L -sentence. Show that if σ is satisfiable, then σ is true in some finite set.

4.19. Let $\mathcal{M} \subseteq \mathcal{N}$ be L -structures. Suppose that for every finite sequence $a_1, \dots, a_m \in M$ and every $b \in N$ there is an automorphism of \mathcal{N} that fixes each element of a_1, \dots, a_m and moves b into A . Show that $\mathcal{M} \preceq \mathcal{N}$.

4.20. Let K be a field and let L be the first order language of vector spaces over K ; the nonlogical symbols of L are a constant 0 , a binary function symbol $+$, and a unary function symbol F_a for each $a \in K$. Given a K -vector space V , we regard V as an L -structure in the obvious way: 0 is interpreted by the identity element of V , $+$ is interpreted by the addition of V , and each F_a is interpreted by the operation of scalar multiplication by a . Suppose $W \subseteq V$ are infinite dimensional K -vector spaces. Use the previous exercise to prove that $W \preceq V$. Use this result to show that any two infinite K -vector spaces are elementarily equivalent.

4.21. Let \mathcal{M} be an L -structure and A a nonempty subset of M . The *diagram of A in \mathcal{M}* , denoted by $\text{Diag}_A(\mathcal{M})$, is the set of all quantifier-free L_A -sentences that are true in $(\mathcal{M}, a)_{a \in A}$. Suppose A is a set of generators for \mathcal{M} and \mathcal{N} is another L -structure. Show that there is a 1-1 correspondence between embeddings of \mathcal{M} into \mathcal{N} and expansions of \mathcal{N} that are models of $\text{Diag}_A(\mathcal{M})$.

4.22. Let \mathcal{M} be an L -structure and A a nonempty subset of M . The *elementary diagram of A in \mathcal{M}* , denoted by $\text{EDiag}_A(\mathcal{M})$, is the set of all L_A -sentences that are true in $(\mathcal{M}, a)_{a \in A}$. Suppose A is a set of generators for \mathcal{M} and \mathcal{N} is another L -structure. Show that there is a 1-1 correspondence between elementary embeddings of \mathcal{M} into \mathcal{N} and expansions of \mathcal{N} that are models of $\text{EDiag}_A(\mathcal{M})$.

4.23. Let I be an index set and U an ultrafilter on I . Let $(\mathcal{M}_i \mid i \in I)$ and $(\mathcal{N}_i \mid i \in I)$ be families of L -structures. If \mathcal{M}_i can be elementarily embedded in \mathcal{N}_i for all $i \in I$, show that $\prod_U \mathcal{M}_i$ can be elementarily embedded in $\prod_U \mathcal{N}_i$.

4.24. Let \mathcal{M} be an infinite L -structure and κ an infinite cardinal. Show that there exists an ultrapower of \mathcal{M} that has cardinality at least κ . (Compare Corollary 1.12.) It follows that every infinite L -structure has an elementary extension of cardinality at least κ .

5. SATURATED MODELS

In this chapter we prove that every satisfiable theory T has models that are rich, in a certain sense. This is the first of several such notions that turn out to be useful in model theory. (See Chapter 13.)

5.1. Definition. Let \mathcal{M} be an L -structure and let κ be an infinite cardinal. We say that \mathcal{M} is κ -saturated if the following condition holds for every $A \subseteq M$ with $\text{card}(A) < \kappa$: whenever \mathcal{C} is a collection of A -definable subsets of M , if \mathcal{C} has FIP, then $\bigcap \mathcal{C} \neq \emptyset$.

Recall that a subset X of M is A -definable in \mathcal{M} iff there is an L_A -formula $\varphi(x)$ such that $X = \{b \in M \mid \mathcal{M}_A \models \varphi[b]\}$.

It is easy to translate Definition 5.1 into equivalent forms expressed in terms of realizing types:

5.2. Fact. Let \mathcal{M} be an L -structure and let κ be an infinite cardinal. Let x be a single variable in L . The following are equivalent:

- (1) \mathcal{M} is κ -saturated.
- (2) the following condition holds for every $A \subseteq M$ with $\text{card}(A) < \kappa$: if $\Phi(x)$ is a set of L_A -formulas that is finitely satisfiable in \mathcal{M}_A , then $\Phi(x)$ is realizable in \mathcal{M}_A .
- (3) For every $A \subseteq M$ with $\text{card}(A) < \kappa$, every type $p(x) \in S_x^{\mathcal{M}}(A)$ is realizable in \mathcal{M}_A .

5.3. Facts. Let \mathcal{M} be an L -structure and κ an infinite cardinal.

- (a) If \mathcal{M} is infinite and κ -saturated, then M has cardinality at least κ .
- (b) If \mathcal{M} is finite, then \mathcal{M} is τ -saturated for every infinite cardinal τ .
- (c) If \mathcal{M} is κ -saturated and A is a subset of M having cardinality $< \kappa$, then the expansion \mathcal{M}_A is also κ -saturated.

Definition 5.1 and its equivalent formulations in 5.2 refer only to formulas with a single free variable. The following result shows that κ -saturated structures realize partial types in many variables in a very rich way.

5.4. Theorem. *Let κ be an infinite cardinal and suppose \mathcal{M} is a κ -saturated L -structure. Suppose $A \subset M$ has cardinality $< \kappa$. Let $\Phi(x)$ be a set of L_A -formulas with free variables among $x = (x_i \mid i \in I)$, where I has cardinality $\leq \kappa$. If $\Phi(x)$ is finitely satisfiable in \mathcal{M}_A , then $\Phi(x)$ is satisfiable in \mathcal{M}_A .*

Proof. Let \mathcal{M} , A , and $\Phi(x)$ be as in the statement of the Theorem. Extend $\Phi(x)$ so that it is maximal among sets of L_A -formulas with free variables among $(x_i \mid i \in I)$ that are finitely satisfiable in \mathcal{M}_A .

Let $<$ be a well ordering of I such that the order type of $(I, <)$ is the cardinal of I . As a consequence, each proper initial segment of $(I, <)$ has cardinality $< \kappa$. For each $j \in I$ let $\Phi_{\leq j}$ be the set of formulas in Φ whose free variables are among $(x_i \mid i \leq j)$. Note that the maximality of Φ ensures that if φ is any L_A -formula whose free variables are among $(x_i \mid i \leq j)$, then

either $\varphi \in \Phi_{\leq j}$ or $\neg\varphi \in \Phi_{\leq j}$. Moreover, $\Phi_{\leq j}$ is closed under conjunction and under application of the existential quantifier $\exists x_j$.

We need to obtain a family $(b_i \mid i \in I)$ of elements of M that satisfies Φ in \mathcal{M}_A ; we do this by induction over the well ordering $(I, <)$.

Fix $j \in I$ and suppose we have already obtained $(b_i \mid i < j)$ that satisfy all the formulas from Φ that have free variables among $(x_i \mid i < j)$. Let Φ' be the result of substituting b_i for all free occurrences of x_i in $\Phi_{\leq j}$, for all $i < j$. We see that Φ' is an x_j -type in $L_{A \cup \{b_i \mid i < j\}}$ that is finitely satisfiable in $\mathcal{M}_{A \cup \{b_i \mid i < j\}}$. Since $A \cup \{b_i \mid i < j\}$ has cardinality $< \kappa$ and \mathcal{M} is κ -saturated, by 5.2 there exists b_j in M that satisfies Φ' in $\mathcal{M}_{A \cup \{b_i \mid i < j\}}$. It follows that the family $(b_i \mid i \leq j)$ satisfies $\Phi_{\leq j}$ in \mathcal{M}_A .

The result of this construction is a family $(b_i \mid i \in I)$ of elements of M such that for each $j \in I$, the family $(b_i \mid i \leq j)$ satisfies $\Phi_{\leq j}$ in \mathcal{M}_A . Hence $(b_i \mid i \in I)$ satisfies Φ in \mathcal{M}_A , as desired. \square

5.5. Corollary. *Let \mathcal{M} be a κ -saturated L -structure. If $\mathcal{N} \equiv \mathcal{M}$ and $\text{card}(\mathcal{N}) \leq \kappa$, then there is an elementary embedding of \mathcal{N} into \mathcal{M} .*

Proof. Let $(c_i \mid i \in I)$ be an enumeration of \mathcal{N} , so $\text{card}(I) \leq \kappa$. Let $\Phi(x_i \mid i \in I)$ be the set of all L -formulas with the indicated free variables that are satisfied by $(c_i \mid i \in I)$ in \mathcal{N} . Since $\mathcal{N} \equiv \mathcal{M}$, the set Φ is finitely satisfiable in \mathcal{M} . Apply Theorem 5.4 to obtain a family $(b_i \mid i \in I)$ of elements of M that satisfies Φ in \mathcal{M} . The function $f: \mathcal{N} \rightarrow M$ that satisfies $f(c_i) = b_i$ for all $i \in I$ is an elementary embedding of \mathcal{N} into \mathcal{M} . \square

5.6. Remark. Suppose $\mathcal{M} \preceq \mathcal{N}$ and $A \subseteq M$, and consider the complete L_A -theory $T_A = \text{Th}(\mathcal{M}_A)$. Then \mathcal{N}_A is a model of T_A . In particular, it makes sense to speak of a type over A being realized in \mathcal{N}_A .

Next we prove the existence of κ -saturated models. We construct such a model by taking the union of a suitable elementary chain. The following result is the main tool needed for building this chain.

5.7. Lemma. *Let \mathcal{M} be an L -structure and let x be a single variable in L . There exists an elementary extension \mathcal{N} of \mathcal{M} such that every $p(x) \in S_x^{\mathcal{M}}(M)$ is realized in \mathcal{N}_M .*

Proof. For every $p(x) \in S_x^{\mathcal{M}}(M)$ let c_p be a new constant symbol, and let L' be the language obtained by adding all these constants to L_M . Let Σ be $\text{Th}(\mathcal{M}_M)$ together with all sentences $\varphi(c_p)$ where $\varphi(x) \in p(x) \in S_x^{\mathcal{M}}(M)$. Consider distinct p_1, \dots, p_n from $S_x^{\mathcal{M}}(M)$ and formulas $\varphi_i(x) \in p_i(x)$ for $i = 1, \dots, n$. Since each type $p_i(x)$ is finitely satisfiable in \mathcal{M}_M , we may take $b_i \in M$ that satisfies $\varphi_i(x)$ in \mathcal{M}_M . Hence $(\mathcal{M}_M, b_1, \dots, b_n)$ is a model of $\text{Th}(\mathcal{M}_M) \cup \{\varphi_1(c_{p_1}), \dots, \varphi_n(c_{p_n})\}$, where we interpret c_{p_i} by b_i for each $i = 1, \dots, n$. Since each $p(x) \in S_x^{\mathcal{M}}(M)$ is closed under finite conjunctions, it follows that every finite subset of Σ has a model. By the Compactness

Theorem, we get a model \mathcal{N}' of Σ . Note that for each distinct $a, b \in M$, the sentence $\neg a = b$ is in Σ , so a, b have distinct interpretations in \mathcal{N}' . By passing to a model isomorphic to \mathcal{N}' , we may assume that each $a \in M$ has itself as its interpretation in \mathcal{N}' . Let \mathcal{N} be the L -reduct of \mathcal{N}' ; so we have $M \subseteq N$ and, because $\mathcal{N}' \models \text{Th}(\mathcal{M}_M)$, we also have $\mathcal{M} \preceq \mathcal{N}$. Finally, for each $p(x) \in S_x^{\mathcal{M}}(M)$ we have that the interpretation of c_p realizes $p(x)$ in \mathcal{N}' and therefore it realizes $p(x)$ in \mathcal{N} . This shows that \mathcal{N} has the desired properties. \square

5.8. Theorem (Existence of Saturated Models). *For every infinite cardinal number κ , every structure has a κ -saturated elementary extension.*

Proof. Let κ^+ denote the smallest cardinal number $> \kappa$ and let $\Lambda = \{\alpha \mid \alpha \text{ is an ordinal } < \kappa^+\}$, ordered by $<$. We obtain the desired structure as the union of an elementary chain of structures, indexed by the well-ordered set $(\Lambda, <)$. The chain of structures is defined by induction, as follows: to begin, we let $\mathcal{M}_0 = \mathcal{M}$. Given $\alpha \in \Lambda$, we define \mathcal{M}_α assuming that \mathcal{M}_β is defined for all $\beta < \alpha$. If $\alpha = \beta + 1$ for some β , let \mathcal{M}_α be one of the elementary extensions of \mathcal{M}_β that are described in Lemma 5.7. Otherwise α is a limit ordinal and we define $\mathcal{M}_\alpha = \bigcup_{\beta < \alpha} \mathcal{M}_\beta$.

The chain of structures defined by this procedure is an elementary chain; one proves by induction on $\beta \in \Lambda$ that $\mathcal{M}_\alpha \preceq \mathcal{M}_\beta$ holds for all $\alpha < \beta$, using Fact 4.9 at limit ordinals.

Finally, let $\mathcal{N} = \bigcup_{\alpha \in \Lambda} \mathcal{M}_\alpha$. We show that this is the required structure.

Note that $\mathcal{M}_\alpha \preceq \mathcal{N}$ for every $\alpha \in \Lambda$, again using Fact 4.9. In particular, \mathcal{N} is an elementary extension of \mathcal{M} .

We will complete the proof by showing that \mathcal{N} is κ^+ -saturated (which is more than we need to prove). Let $A \subseteq N$ satisfy $\text{card}(A) \leq \kappa$. Since the cofinality of the ordered set Λ is $\kappa^+ > \kappa$ there exists $\eta \in \Lambda$ such that $A \subseteq M_\eta$.

Let $\Phi(x)$ be any x -type in L_A that is finitely satisfiable in \mathcal{N}_A . Since $\mathcal{M}_\eta \preceq \mathcal{N}$ and $A \subseteq M_\eta$, we see that $\Phi(x)$ is finitely satisfiable in $\mathcal{M}_{\eta, A}$. Hence there exists $p(x) \in S_x^{\mathcal{M}_\eta}(M_\eta)$ satisfying $\Phi(x) \subseteq p(x)$. By construction, this implies that $\Phi(x)$ is satisfied by some b in $\mathcal{M}_{\eta+1, M_\eta}$. Since $\mathcal{M}_{\eta+1} \preceq \mathcal{N}$, it follows that b satisfies $\Phi(x)$ in \mathcal{N}_A , as desired. \square

5.9. Proposition. *Let \mathcal{M}, \mathcal{N} be κ -saturated L -structures with $\mathcal{M} \equiv \mathcal{N}$. Let \mathcal{F} be the set of all functions $f: A \rightarrow B$ that are elementary with respect to \mathcal{M}, \mathcal{N} (so $A \subseteq M$ and $B \subseteq N$) with A, B of cardinality $< \kappa$. Then \mathcal{F} is a nonempty back-and-forth system from \mathcal{M} to \mathcal{N} . If, in addition, M, N have cardinality $= \kappa$, then \mathcal{M} and \mathcal{N} are isomorphic.*

Proof. Note that $\mathcal{F} \neq \emptyset$ follows from $\mathcal{M} \equiv \mathcal{N}$, because the empty function is in \mathcal{F} .

Consider $f: A \rightarrow B$ in \mathcal{F} and take $a \in M$. Define $p(x) = \text{tp}_x^{\mathcal{M}^A}(a) \in S_x^{\mathcal{M}}(A)$. Let $q(x) \in S_x^{\mathcal{N}}(B)$ be the image of $p(x)$ under f . That is, $q(x)$ consists of all L_B -formulas of the form $\varphi(x, f(a_1), \dots, f(a_n))$ where $\varphi(x, y_1, \dots, y_n)$ is an L -formula and $\varphi(x, a_1, \dots, a_n) \in p(x)$. Our saturation assumption ensures that $q(x)$ is realized in \mathcal{N}_B by some $b \in N$. It follows that f extended by taking a to b is in \mathcal{F} . This proves the “forth” property for \mathcal{F} , and the “back” property is proved similarly by exchanging the roles of \mathcal{M} and \mathcal{N} .

Finally, assume that $\text{card}(M) = \text{card}(N) = \kappa$. By a usual back-and-forth construction of length κ , build a family $(f_\alpha \mid \alpha < \kappa)$ of members of \mathcal{F} with the properties: (i) $f_\alpha \subseteq f_\beta$ whenever $\alpha < \beta < \kappa$; (ii) $M = \cup\{\text{dom}(f_\alpha) \mid \alpha < \kappa\}$; and $N = \cup\{\text{range}(f_\alpha) \mid \alpha < \kappa\}$. Then $\cup_{\alpha < \kappa} f_\alpha$ is an isomorphism from \mathcal{M} onto \mathcal{N} . \square

5.10. Corollary. *Let T be a complete theory and let κ be an infinite cardinal. Then, up to isomorphism T has at most one κ -saturated model of cardinality κ .*

Proof. Immediate from 5.9. \square

5.11. Remark. An L -structure \mathcal{M} is called *saturated* (without mention of any cardinal) if it is κ -saturated for $\kappa = \text{card}(M)$.

The existence of κ -saturated models can also be proved directly using ultraproducts. However, when $\kappa > \omega_1$ it is technically rather difficult to prove the existence of an ultrafilter U for which the ultrapower \mathcal{M}^I/U is κ -saturated, and this is why we used a different method. On the other hand, when the language is countable and $\kappa = \omega_1$, it is relatively easy to obtain ω_1 -saturated ultraproducts, as we now show.

5.12. Theorem. *Let U be a nonprincipal ultrafilter on a countable (infinite) set I . Let L be a countable language and $(\mathcal{M}_i \mid i \in I)$ a family of L -structures. Then the ultraproduct $\prod_U \mathcal{M}_i$ is ω_1 -saturated.*

Proof. We may assume $I = \mathbb{N}$. Since U is nonprincipal it contains every cofinite subset of \mathbb{N} . Let M be the cartesian product $\prod_{i \in \mathbb{N}} M_i$.

We denote the ultraproduct $\prod_U \mathcal{M}_i$ by \mathcal{N} and its underlying set by N . Let A be any countable subset of N . Let $\Phi(x)$ be a set of L_A -formulas such that every finite subset of $\Phi(x)$ is satisfiable in \mathcal{N}_A . We must show that the entire set $\Phi(x)$ is satisfiable in \mathcal{N}_A .

Let $(\varphi_k(x, y_k) \mid k \in \mathbb{N})$ be a family of L -formulas and $(b_k \mid k \in \mathbb{N})$ a family of finite tuples from A such that $\Phi(x)$ is $\{\varphi_k(x, b_k) \mid k \in \mathbb{N}\}$. This is possible because the language L_A is countable. For convenience of notation we will take each tuple b_k to be of length 1 (i.e., to be an element of A). For each $k \in \mathbb{N}$, let f_k be an element of the cartesian product M for which b_k is the equivalence class f_k/U . (See the notation in 1.5.)

For each $k \in \mathbb{N}$ and $i \in \mathbb{N}$, let

$$C_k(i) = \{u \in M_i \mid \mathcal{M}_i \models \varphi_k[u, f_k(i)]\}.$$

Using the Fundamental Theorem of Ultraproducts and the hypothesis that each finite subset of $\Phi(x)$ is satisfiable in (\mathcal{N}_A) we have that

$$\{i \in \mathbb{N} \mid C_0(i) \cap \dots \cap C_k(i) \neq \emptyset\} \in U$$

for each $k \in \mathbb{N}$.

Define G_k for $k \in \mathbb{N}$ by setting $G_0 = \mathbb{N}$ and for $k \geq 1$

$$G_k := \{i \in \mathbb{N} \mid i \geq k \text{ and } C_0(i) \cap \dots \cap C_k(i) \neq \emptyset\}.$$

Note that $\mathbb{N} = G_0 \supseteq G_1 \supseteq \dots \supseteq G_k$ and that $G_k \in U$, for all $k \in \mathbb{N}$. Moreover, $\bigcap \{G_k \mid k \in \mathbb{N}\} = \emptyset$; therefore we may define $d(i)$ for each $i \in \mathbb{N}$ to be the largest $k \in \mathbb{N}$ such that $i \in G_k$.

Now we construct an element g of M whose equivalence class g/U will satisfy every formula from $\Phi(x)$ in \mathcal{N}_A . Fix $i \in \mathbb{N}$ and define $g(i)$ as follows. If $d(i) = 0$ let $g(i)$ be an arbitrary element of M_i . If $d(i) \geq 1$, choose $g(i)$ to be an element of $C_0(i) \cap \dots \cap C_{d(i)}(i)$, which is guaranteed to be nonempty by the definition of $d(i)$.

It is obvious that for each $k \in \mathbb{N}$, we have $g(i) \in C_k(i)$ whenever $d(i) \geq k$ and $d(i) \geq 1$. Therefore $\{i \in \mathbb{N} \mid g(i) \in C_0(i)\} \supseteq G_1$ and for $k \geq 1$, $\{i \in \mathbb{N} \mid g(i) \in C_k(i)\} \supseteq G_k$. Recalling the definition of $C_k(i)$ and that the sets G_k are all in U , it follows that for each $k \in \mathbb{N}$

$$\{i \in \mathbb{N} \mid \mathcal{M}_i \models \varphi_k[g(i), f_k(i)]\} \in U.$$

The Fundamental Theorem of Ultraproducts implies that g/U satisfies $\varphi_k(x, b_k)$ in \mathcal{N}_A for all $k \in \mathbb{N}$. That is, g/U satisfies $\Phi(x)$ in \mathcal{N}_A , as desired. \square

5.13. Remark. Let I be any index set and let U be an ultrafilter on I . We say that U is *countably incomplete* if there exist sets $(F_k \mid k \in \mathbb{N})$ from U whose intersection $\bigcap \{F_k \mid k \in \mathbb{N}\}$ is not in U . The proof of the preceding result can be slightly modified to show that if U is a countably incomplete ultrafilter on I and $(\mathcal{M}_i)_{i \in I}$ is any family of L structures indexed by I , where L is a countable language, then the ultraproduct $\prod_U \mathcal{M}_i$ is ω_1 -saturated.

EXERCISES

5.14. Show that the linear ordering $(\mathbb{R}, <)$ is ω -saturated but not ω_1 -saturated. (Note that $(\mathbb{R}, <) \models DLO$, so you can use Example 4.14.)

5.15. Show that no infinite well ordering is ω -saturated.

5.16. Let I be a countable infinite set and U a nonprincipal ultrafilter on I .

- Let \mathcal{M} be the linear ordering $(\mathbb{Q}, <)$. Show that the cardinality of the ultrapower \mathcal{M}^I/U is exactly 2^ω . (Note that it is not enough to prove that the ultrapower is uncountable; it is possible that $\omega_1 < 2^\omega$.)

- More generally, let L be any first order language and let \mathcal{M}_i be a countable infinite L -structure for each $i \in I$. Show that the cardinality of the ultraproduct $\prod_U(\mathcal{M}_i \mid i \in I)$ is exactly 2^ω .

6. LÖWENHEIM-SKOLEM THEOREMS

6.1. Theorem (Downward Löwenheim-Skolem Theorem). *Let \mathcal{N} be an infinite L -structure and let A be a subset of N . Let κ be an infinite cardinal that satisfies $\text{card}(L) \leq \kappa$ and $\text{card}(A) \leq \kappa \leq \text{card}(N)$. There exists an elementary substructure \mathcal{M} of \mathcal{N} such that $\text{card}(M) = \kappa$ and $A \subseteq M$.*

Proof. By enlarging A within N if necessary, we may assume $\text{card}(A) = \kappa$.

We expand the language L by a procedure known as “Skolemization”. For each L -formula $\varphi(x, y)$, where x is a finite list of variables of length $n \geq 0$ and y is a single variable, we add a new n -ary function symbol f_φ . Let L' be the language whose signature consists of L together with all the new function symbols f_φ . (Note that if x is empty, then $n = 0$ and f_φ is a constant symbol.) Since we have $\text{card}(L) \leq \kappa$, the number of L -formulas is also $\leq \kappa$ and hence $\text{card}(L') \leq \kappa$ as well.

Using the Axiom of Choice, we expand \mathcal{N} to an L' -structure \mathcal{N}' by interpreting each f_φ in such a way that

$$\mathcal{N}' \models \forall x_1 \dots \forall x_n (\exists y \varphi(x, y) \rightarrow \varphi(x, f_\varphi(x))).$$

Finally, let \mathcal{M}' be the substructure of \mathcal{N}' generated by A . Our construction yields $\mathcal{M}' \preceq \mathcal{N}'$ by Theorem 4.8 (the Tarski-Vaught Test). Since $A \subseteq M'$ and $\text{card}(L') \leq \kappa$ we have $\kappa = \text{card}(A) \leq \text{card}(M') \leq \kappa$ and therefore $\text{card}(M') = \kappa$.

Finally, taking \mathcal{M} to be the L -reduct of \mathcal{M}' yields the desired elementary substructure of \mathcal{N} . □

We illustrate the use of Theorem 6.1 by proving the existence of *countable* ω -saturated models, under suitable hypotheses.

6.2. Theorem (Countable ω -saturated Models). *Assume that L is a countable language and let T be a complete theory in L with only infinite models. The following are equivalent:*

- (1) *The theory T has a countable ω -saturated model.*
- (2) *For each $n \geq 1$, the type space $S_n(T)$ is countable.*

Proof. (1 \Rightarrow 2) Let \mathcal{M} be a countable, ω -saturated model of T . By Theorem 5.4, every n -type consistent with T is realized in \mathcal{M} . Hence $S_n(T)$ must be countable.

(2 \Rightarrow 1) This proof is patterned after the proofs of Lemma 5.7 and Theorem 5.8, with appropriate modifications to keep structures countable.

Assume $S_n(T)$ is countable for each $n \geq 1$. It follows that for every model \mathcal{M} of T and every finite subset F of M , the set $S_1(T_F)$ is countable. Indeed, there is an obvious embedding of $S_1(T_F)$ into $S_{k+1}(T)$, where k is the cardinality of F ; namely, if $F = \{a_1, \dots, a_k\}$ and $\mathcal{N} \succeq \mathcal{M}$, map the type of b in $(\mathcal{N}, a_1, \dots, a_k)$ to the type of (b, a_1, \dots, a_k) in \mathcal{N} .

Using Lemma 5.7 followed by the use of Theorem 6.1 (for the cardinal $\kappa = \omega$) we may prove the following version of Lemma 5.7 for the current situation: *Let \mathcal{M} be a countable model of T . There exists a countable elementary extension \mathcal{N} of \mathcal{M} such that for any finite subset F of M , every 1-type over F is realized in \mathcal{N}_F .*

Now let \mathcal{M} be any countable model of T . Build an elementary chain $(\mathcal{M}_n \mid n \in \mathbb{N})$ by setting $\mathcal{M}_0 = \mathcal{M}$ and by applying the statement in the previous paragraph to obtain \mathcal{M}_{n+1} from \mathcal{M}_n for each $n \in \mathbb{N}$. The union of this elementary chain is a countable ω -saturated elementary extension of \mathcal{M} . \square

6.3. Remark. A complete theory T in a countable language is called *small* if $S_n(T)$ is countable for every $n \geq 1$.

6.4. Corollary. *If T is a complete theory in a countable language and T has only countably many countable models, up to isomorphism, then T has a countable ω -saturated model.*

Proof. Each type consistent with T is realized in a countable model. Under the hypotheses of this Corollary, this implies there are only countably many n -types consistent with T , for each $n \geq 1$. Hence the previous result applies and yields the existence of a countable ω -saturated model. \square

6.5. Theorem (Upward Löwenheim-Skolem Theorem). *Let \mathcal{M} be an infinite L -structure and let κ be an infinite cardinal that satisfies $\text{card}(L) \leq \kappa$ and $\text{card}(M) \leq \kappa$. There exists an elementary extension \mathcal{N} of \mathcal{M} such that $\text{card}(N) = \kappa$.*

Proof. Since \mathcal{M} is infinite, it has an elementary extension \mathcal{N}' whose cardinality is $\geq \kappa$ (for example, a κ -saturated elementary extension). By Theorem 6.1 there exists an elementary substructure \mathcal{N} of \mathcal{N}' such that $M \subseteq N$ and $\text{card}(N) = \kappa$. It follows easily that $\mathcal{M} \preceq \mathcal{N}$. Indeed, if $\varphi(x)$ is any L -formula and a is any tuple from M , then we have $\mathcal{M} \models \varphi[a] \Leftrightarrow \mathcal{N}' \models \varphi[a] \Leftrightarrow \mathcal{N} \models \varphi[a]$. \square

6.6. Fact. If \mathcal{M} is finite, then $\text{Th}(\mathcal{M})$ is absolutely categorical, in the sense that any model \mathcal{N} of $\text{Th}(\mathcal{M})$ must be isomorphic to \mathcal{M} . In particular, a finite structure cannot have any proper elementary extension or any proper elementary substructure.

It is a consequence of Theorem 6.5 together with some elementary reasoning that a first order theory can be absolutely categorical *only* when it is the theory of a fixed finite structure. For a theory with at least one infinite model, the only categoricity we can expect is that of the following Definition.

6.7. Definition. Let T be a theory in L and let κ be any cardinal. We say T is κ -categorical if T has a model of cardinality equal to κ , and any two models of T that are both of cardinality κ are isomorphic.

6.8. Theorem (Categoricity Test for Completeness). *Let T be a satisfiable theory that has only infinite models. If T is κ -categorical for some cardinal $\kappa \geq \text{card}(L)$, then T is complete.*

Proof. Suppose T is a theory that is κ -categorical, where $\kappa \geq \text{card}(L)$ and T has only infinite models. We need to show that if $\mathcal{M}, \mathcal{N} \models T$, then $\mathcal{M} \equiv \mathcal{N}$. By use of Theorems 6.1 and 6.5 (as needed), we find models $\mathcal{M}', \mathcal{N}'$ with $\mathcal{M}' \equiv \mathcal{M}$, $\mathcal{N}' \equiv \mathcal{N}$, and $\text{card}(\mathcal{M}') = \kappa = \text{card}(\mathcal{N}')$. (If $\kappa < \text{card}(M)$, use the Downward Löwenheim-Skolem Theorem to find \mathcal{M}' with $\mathcal{M}' \equiv \mathcal{M}$ and $\text{card}(\mathcal{M}') = \kappa$. If $\kappa > \text{card}(M)$, use the Upward Löwenheim-Skolem Theorem.)

Since T is κ -categorical we have $\mathcal{M}' \cong \mathcal{N}'$, and hence $\mathcal{M}' \equiv \mathcal{N}'$ by Proposition 4.13. Therefore $\mathcal{M} \equiv \mathcal{M}' \equiv \mathcal{N}' \equiv \mathcal{N}$, and hence $\mathcal{M} \equiv \mathcal{N}$. \square

EXERCISES

6.9. Show that the “no finite models” assumption in Theorem 6.8 is necessary. That is, give an example of an infinite cardinal κ and a κ -categorical theory T in a language whose cardinality is at most κ , such that T is not complete.

6.10. Let L be the language whose only nonlogical symbol is the unary predicate symbol P . Let T be the theory of all L -structures \mathcal{M} such that $P^{\mathcal{M}}$ is infinite. Give a clear mathematical description of the space $S_0(T)$ of all complete extensions of T , including its topology.

6.11. Let L be the first order language with two binary function symbols \cap and \cup , a unary function symbol c , and two constant symbols 0 and 1 . For each set S let $\mathcal{P}(S)$ denote the L -structure based on the *power set* of S . That is, the underlying set of $\mathcal{P}(S)$ is the collection of all subsets of S , we interpret \cap, \cup, c as intersection, union, and complement, respectively, and we interpret $0, 1$ as \emptyset, S , respectively. Let \mathcal{K} be the class of all L -structures that are isomorphic to $\mathcal{P}(S)$ for some set S . Show that \mathcal{K} is not axiomatizable.

6.12. Let κ be an infinite cardinal and let G be a simple group of cardinality equal to κ . If τ is any infinite cardinal $\leq \kappa$, show that G has a subgroup H such that $\text{card}(H) = \tau$ and H is simple. (Note that a group is simple iff whenever a, b are elements not equal to the identity element, then a is a finite product of some conjugates of b and some conjugates of b^{-1} .)

7. QUANTIFIER ELIMINATION

The method of quantifier elimination, which we introduce in this chapter, is the most important practical technique in applications of model theory. At the simplest level, it can be useful in showing that a theory is complete, or, more generally, in classifying the models of a set of sentences Σ up to elementary equivalence (*i.e.*, determining all the completions of Σ). Similarly, it can be a significant tool for understanding the structure of the type spaces $S_x(\Sigma)$. Its most important use is in analyzing the properties of definable sets in models of Σ .

7.1. Definition. Let Σ be a set of L -sentences. We say Σ has *Quantifier Elimination (QE)* if for every $n \geq 1$ and every L -formula $\varphi(x_1, \dots, x_n)$ there exists a quantifier free L -formula $\psi(x_1, \dots, x_n)$ such that $T \models \varphi \leftrightarrow \psi$.

7.2. Remark. Suppose Σ is a set of L -sentences with QE, and assume that L contains at least one constant symbol, say c . Then every L -sentence is Σ -equivalent to a quantifier-free L -sentence, although this is not required directly by the defining condition. (Proof: Let σ be a sentence and regard it as a formula $\varphi(x)$. Since Σ has QE, there is a quantifier-free formula $\psi(x)$ such that $\Sigma \models \forall x(\sigma \leftrightarrow \psi(x))$. Since x does not occur free in σ , this implies $\Sigma \models \sigma \leftrightarrow \psi(c)$.)

Here are some simple examples of useful facts about models of Σ that follow from QE.

7.3. Facts. Let Σ be a set of sentences that has QE.

- (1) If \mathcal{M}, \mathcal{N} are models of Σ and they have a common substructure \mathcal{A} , then $\mathcal{M}_A \equiv \mathcal{N}_A$. (That is, the identity map on A is elementary with respect to \mathcal{M}, \mathcal{N} .)
- (2) If \mathcal{M}, \mathcal{N} are models of Σ and $\mathcal{M} \subseteq \mathcal{N}$, then $\mathcal{M} \preceq \mathcal{N}$.
- (3) If Σ is satisfiable, and there exists an L -structure \mathcal{A} such that \mathcal{A} embeds into every model of Σ , then Σ is complete.

Proof. We prove (1). Let σ be an L_A -sentence. Then σ is of the form $\varphi(a_1, \dots, a_n)$ where $\varphi(x_1, \dots, x_n)$ is an L -formula and $a_1, \dots, a_n \in A$. Because Σ has QE, there is a quantifier-free L -formula $\psi(x_1, \dots, x_n)$ that is equivalent to $\varphi(x_1, \dots, x_n)$ in all models of Σ . Hence $\sigma \leftrightarrow \psi(a_1, \dots, a_n)$ is true in \mathcal{M}_A and in \mathcal{N}_A . Since $\psi(a_1, \dots, a_n)$ is quantifier free, we have

$$\mathcal{M}_A \models \psi(a_1, \dots, a_n) \iff \mathcal{N}_A \models \psi(a_1, \dots, a_n)$$

and therefore

$$\mathcal{M}_A \models \sigma \iff \mathcal{N}_A \models \sigma.$$

Since σ was arbitrary, this shows $\mathcal{M} \equiv \mathcal{N}$. □

7.4. Remark. If Σ is a set of L -sentences that has QE, and if L' is an extension of L by adding only constant symbols, then Σ continues to have QE when considered as a set of L' -sentences.

Our first criterion for QE comes directly from Corollary 3.22.

7.5. Theorem. *Let Σ be a set of L -sentences. The following conditions are equivalent:*

- (1) Σ has quantifier elimination.
- (2) For each $n \geq 1$, every type in $S_n(\Sigma)$ is determined by the quantifier-free formulas it contains.

Proof. (1 \Rightarrow 2): Obvious.

(2 \Rightarrow 1): Let $\varphi(x_1, \dots, x_n)$ be any L -formula, $n \geq 1$. Apply Corollary 3.22 with $\Delta(x)$ taken to be the set of all quantifier free L -formulas whose free variables are among x_1, \dots, x_n . \square

Next we make the preceding result more useful for applications by relating it to extensions of embeddings.

7.6. Notation. Let \mathcal{M}, \mathcal{N} be L -structures.

- We let $\text{Sub}(\mathcal{M}, \mathcal{N})$ denote the set of all functions f such that f is an embedding of a substructure of \mathcal{M} into \mathcal{N} .
- Similarly, we let $\text{Sub}_0(\mathcal{M}, \mathcal{N})$ denote the set of all $f \in \text{Sub}(\mathcal{M}, \mathcal{N})$ such that the domain of f is a finitely generated substructure of \mathcal{M} .

Note that $\text{Sub}_0(\mathcal{M}, \mathcal{N}) \subseteq \text{Sub}(\mathcal{M}, \mathcal{N})$ and they could be empty.

In order to deal efficiently with substructure embeddings, we need some lemmas and notation.

7.7. Notation. Let \mathcal{M} be an L -structure and A a nonempty subset of M . We denote by $\langle A \rangle_{\mathcal{M}}$ the substructure of \mathcal{M} that is generated by A .

7.8. Fact. Let \mathcal{M} be an L -structure and A a nonempty subset of M . The underlying set of $\langle A \rangle_{\mathcal{M}}$ consists of all elements of M of the form $t^{\mathcal{M}}(a_1, \dots, a_n)$ where $t(x_1, \dots, x_n)$ is an L -term and $a_1, \dots, a_n \in A$.

7.9. Lemma. *Let \mathcal{M}, \mathcal{N} be L -structures and $f \in \text{Sub}(\mathcal{M}, \mathcal{N})$. Then*

- (1) *The range of f is a substructure of \mathcal{N} .*
- (2) *For each L -term $t(x_1, \dots, x_n)$ and each a_1, \dots, a_n in the domain of f ,*

$$t^{\mathcal{N}}(f(a_1), \dots, f(a_n)) = f(t^{\mathcal{M}}(a_1, \dots, a_n)).$$

- (3) *For each quantifier-free L -formula $\varphi(x_1, \dots, x_n)$ and each a_1, \dots, a_n in the domain of f ,*

$$\mathcal{N} \models \varphi[f(a_1), \dots, f(a_n)] \Leftrightarrow \mathcal{M} \models \varphi[a_1, \dots, a_n].$$

Proof. (1) We need to show that $c^{\mathcal{N}}$ is in the range of f for any constant symbol c of L and that the range of f is closed under the application of $F^{\mathcal{N}}$ for any function symbol F of L . If c is a constant symbol of L , then $c^{\mathcal{M}}$ is in the domain of f and we have $c^{\mathcal{N}} = f(c^{\mathcal{M}})$. If F is an n -ary function symbol of L and a_1, \dots, a_n are in the domain of f (so $f(a_1), \dots, f(a_n)$

are arbitrary elements of the range of f), we have $F^{\mathcal{N}}(f(a_1), \dots, f(a_n)) = f(F^{\mathcal{M}}(a_1, \dots, a_n))$, which is in the range of f .

(2) This is proved by induction on terms.

(3) This is proved by induction on formulas. Part (2) yields the base case, in which atomic formulas are treated. The induction steps for propositional connectives are trivial. \square

7.10. Lemma. *Let \mathcal{M}, \mathcal{N} be L -structures. Let J be a nonempty set and consider two functions $\alpha: J \rightarrow M$, $\beta: J \rightarrow N$. Let $(x_j \mid j \in J)$ be a family of distinct variables. Suppose that for any quantifier-free formula $\varphi(x_j \mid j \in J)$ whose variables are among $(x_j \mid j \in J)$ we have*

$$\mathcal{M} \models \varphi[\alpha(j) \mid j \in J] \iff \mathcal{N} \models \varphi[\beta(j) \mid j \in J].$$

Then there exists an embedding f from $\langle \{\alpha(j) \mid j \in J\} \rangle_{\mathcal{M}}$ into \mathcal{N} such that $f(\alpha(j)) = \beta(j)$ for all $j \in J$. Moreover, f is unique with these properties and its range is $\langle \{\beta(j) \mid j \in J\} \rangle_{\mathcal{N}}$.

Proof. The underlying set of $\langle \{\alpha(j) \mid j \in J\} \rangle_{\mathcal{A}}$ consists exactly of those elements of A that can be written in the form $t^A(\alpha(j) \mid j \in J)$ where $t(x_j \mid j \in J)$ is any L -term whose variables are among $(x_j \mid j \in J)$. If t_1, t_2 are two such terms and $t_1^A(\alpha(j) \mid j \in J) = t_2^A(\alpha(j) \mid j \in J)$, then our assumptions yield that $t_1^{\mathcal{B}}(\beta(j) \mid j \in J) = t_2^{\mathcal{B}}(\beta(j) \mid j \in J)$. (Consider the quantifier-free formula $t_1 = t_2$.) Thus we may define a function f on $\langle \{\alpha(j) \mid j \in J\} \rangle_{\mathcal{A}}$ by

$$f(t^A(\alpha(j) \mid j \in J)) = t^{\mathcal{B}}(\beta(j) \mid j \in J)$$

where t ranges over the L -terms whose variables are among $(x_j \mid j \in J)$. It is routine to show that this f has the desired properties. \square

Next we present the fundamental test for Quantifier Elimination.

7.11. Theorem. *Let Σ be a set of L -sentences. The following conditions are equivalent:*

- (1) Σ has QE.
- (2) Whenever \mathcal{M}, \mathcal{N} are models of Σ , f is in $\text{Sub}(\mathcal{M}, \mathcal{N})$, and $a \in M$, there exists an elementary extension \mathcal{N}' of \mathcal{N} and a function g in $\text{Sub}(\mathcal{M}, \mathcal{N}')$ such that g extends f and $a \in \text{dom}(g)$.
- (3) Whenever \mathcal{M}, \mathcal{N} are ω -saturated models of Σ , either $\text{Sub}_0(\mathcal{M}, \mathcal{N})$ is empty or it is a back-and-forth system from \mathcal{M} to \mathcal{N} .

Proof. (1 \Rightarrow 2): Assume (1) and the hypotheses of (2). Let $A = \text{dom}(f)$ and $B = \text{range}(f)$. For each L_A -formula φ , we let φ^f denote the L_B -formula obtained by replacing each occurrence of a by $f(a)$, for every $a \in A$.

Let $p(x) = tp_{\mathcal{M}}(a/A)$ and let $p^f(x) = \{\varphi^f(x) \mid \varphi(x) \in p(x)\}$. Consider finitely many formulas $\varphi_1(x), \dots, \varphi_m(x)$ from $p(x)$. Since $A \neq \emptyset$, the language L_A has some constant symbols, and we know that Σ has QE as a set

of L_A -sentences. Hence there exists a quantifier free L_A -sentence σ such that

$$\Sigma \models_{L_A} \exists x(\varphi_1(x) \wedge \cdots \wedge \varphi_m(x)) \leftrightarrow \sigma.$$

and hence also

$$\Sigma \models_{L_B} \exists x(\varphi_1^f(x) \wedge \cdots \wedge \varphi_m^f(x)) \leftrightarrow \sigma^f.$$

Since $p(x)$ is realized in \mathcal{M}_A , we have that $\mathcal{M}_A \models \sigma$, and since σ is quantifier free, this implies $\mathcal{N}_B \models \sigma^f$, and hence that

$$\mathcal{N}_B \models \exists x(\varphi_1^f(x) \wedge \cdots \wedge \varphi_m^f(x)).$$

It follows from this argument that $p^f(x)$ is finitely satisfied in \mathcal{N}_B . Therefore we may take \mathcal{N}' to be an elementary extension of \mathcal{N} such that $p^f(x)$ is realized, say by b , in \mathcal{N}'_B . We see that $p^f(x) = \text{tp}_{\mathcal{N}'}(b/B)$. Indeed, if $\psi(x)$ is any L_B -formula, there exists an L_A -formula $\varphi(x)$ such that $\varphi^f(x) = \psi(x)$; we know either $\varphi(x)$ or $\neg\varphi(x)$ is in $p(x)$ and therefore one of $\varphi^f(x)$, $\neg\varphi^f(x)$ (which equals $(\neg\varphi)^f(x)$) must be in $p^f(x)$.

Finally, we may define the desired function g extending f by letting $g(a) = b$ and extending g to be defined on the substructure of \mathcal{M} generated by $A \cup \{a\}$. As noted in Fact 7.8, any element of this substructure is of the form $t^{\mathcal{M}}(a_1, \dots, a_m, a)$ for some L -term t and some $a_1, \dots, a_m \in A$, and we define g on this element to be $t^{\mathcal{N}'}(f(a_1), \dots, f(a_m), b)$. Then $g \in \text{Sub}(\mathcal{M}, \mathcal{N})$, and g extends f and satisfies $g(a) = b$. The information needed to show that g is well defined and to check these final details is contained in $p(x)$ and $p^f(x)$, as shown in Lemma 7.10.

(2 \Rightarrow 3): Assume \mathcal{M}, \mathcal{N} are ω -saturated models of Σ . When we apply statement (2) to \mathcal{M}, \mathcal{N} , we may take \mathcal{N}' to be \mathcal{N} itself, since the type realized by $g(a)$ in \mathcal{N}' over a finite set of generators for the range of f can be realized in \mathcal{N} . (Then we argue as in the previous paragraph.) This shows that $\text{Sub}_0(\mathcal{M}, \mathcal{N})$ has the “forth” property in Definition 4.11. Applying the same argument to the opposite pair \mathcal{N}, \mathcal{M} shows that $\text{Sub}_0(\mathcal{M}, \mathcal{N})$ also has the “back” property in that Definition. That is, $\text{Sub}_0(\mathcal{M}, \mathcal{N})$ is indeed a back-and-forth system from \mathcal{M} to \mathcal{N} , as desired.

(3 \Rightarrow 1): We verify condition (2) in Theorem 7.5. Fix $n \geq 1$ and let p, q be any two types in $\mathcal{S}_n(\Sigma)$. Suppose a_1, \dots, a_n realizes p in $\mathcal{M} \models \Sigma$ and b_1, \dots, b_n realizes q in $\mathcal{N} \models \Sigma$. By Theorem 5.8 we may assume that \mathcal{M} and \mathcal{N} are ω -saturated. Suppose p and q contain exactly the same quantifier-free formulas. Using Lemma 7.10 we get an isomorphism f from $\langle a_1, \dots, a_n \rangle_{\mathcal{M}}$ onto $\langle b_1, \dots, b_n \rangle_{\mathcal{N}}$ with $f(a_i) = b_i$ for all $i = 1, \dots, n$. Then $f \in \text{Sub}_0(\mathcal{M}, \mathcal{N})$, so by statement (3) and Proposition 4.13 we conclude that f is elementary, and thus $p = q$, as desired. \square

The following observation gives a variant of the criterion for QE in the previous Theorem.

7.12. Corollary. *Let Σ be a set of L -sentences. The following conditions are equivalent:*

(1) Σ has QE.

(2') Whenever \mathcal{M}, \mathcal{N} are models of Σ , f is in $\text{Sub}_0(\mathcal{M}, \mathcal{N})$, and $a \in M$, there exists an elementary extension \mathcal{N}' of \mathcal{N} and a function g in $\text{Sub}_0(\mathcal{M}, \mathcal{N}')$ such that g extends f and $a \in \text{dom}(g)$.

Proof. It is clear that 7.11(2) implies (2'). Further, it is clear that the proof of $(2 \Rightarrow 3)$ in 7.11 only needs what is provided by (2'). Alternatively, it is clear by a Zorn's Lemma argument (using facts about elementary embeddings) that (2') implies 7.11(2). \square

Example: discrete linear orderings without endpoints

Consider the language L_0 whose only nonlogical symbol is a binary predicate $<$. Let T_{dis} be the theory of discrete linear orderings without minimum or maximum element, formulated in L_0 . (A linear ordering without endpoints is *discrete* if each element has a unique successor and a unique predecessor.) The theory T_{dis} does not admit QE, as can be seen by considering the formula $\exists z(x < z \wedge z < y)$. However, T_{dis} can be analyzed by applying the *method* of quantifier elimination. This means that we formulate a carefully chosen extension, show that the extension has QE, and then use this fact to draw conclusions about T_{dis} .

To obtain the extension of T_{dis} that we will use, let L be the extension of L_0 obtained by adding unary function symbols p and s . T is the theory in L of all linear orderings with functions p and s such that for each element x , $p(x)$ is the predecessor of x in the ordering and $s(x)$ is the successor of x . If \mathcal{A} is any model of T , it is obvious that the reduct of \mathcal{A} to L_0 is a model of T_{dis} . Moreover, each model \mathcal{M}_0 of T_{dis} expands in a unique way to a model of T , because the predecessor function and the successor function are definable in \mathcal{M}_0 .

7.13. Example. The theory T of discrete linear orderings without minimum or maximum element, equipped with the predecessor and successor functions, has quantifier elimination and is complete. Therefore T_{dis} is also complete.

Proof. We verify condition (2') in Corollary 7.12. Let \mathcal{M}, \mathcal{N} be models of T and let \mathcal{M}_0 be the substructure of \mathcal{M} generated by the elements a_1, \dots, a_m . We may assume $a_1 < \dots < a_m$ in \mathcal{M} . Further, let f be an embedding of \mathcal{M}_0 into \mathcal{N} . We may suppose that no subset of $\{a_1, \dots, a_m\}$ generates \mathcal{M}_0 . It follows that for each $k \in \mathbb{N}$ and each $i = 2, \dots, m-1$, the k -th successor of a_i is less than a_{i+1} and the k -th predecessor of a_i is greater than a_{i-1} in \mathcal{M} .

Now let a be any element of M that is not in M_0 . We must extend the embedding f so that it is defined on a as well as its predecessors and successors, and gives an embedding into an elementary extension \mathcal{N}' of \mathcal{N} . To accomplish this, we take \mathcal{N}' to be any ω -saturated elementary extension of \mathcal{N} , which exists by Theorem 5.8.

For each $i = 1, \dots, m$ let C_i be the set of all predecessors and successors of a_i in \mathcal{M} , including a_i itself. Then each C_i is a convex set in \mathcal{M} and M_0 is the disjoint union of the sets C_1, \dots, C_m . Moreover, a either lies between C_i and C_{i+1} for some $i = 1, \dots, m-1$, or it lies below C_1 , or it lies above C_m , in the ordering of \mathcal{M} .

Since f is an embedding with respect to ordering and also to the predecessor and successor functions, each set $f(C_i)$ is a convex set in \mathcal{N} that consists of all the predecessors and successors of $f(a_i)$. This remains true when we move up to \mathcal{N}' . Moreover, the convex sets $f(C_1), \dots, f(C_m)$ are disjoint and are arranged in order from left to right in the ordering of \mathcal{N}' . A simple saturation argument shows that there exist elements d_1, \dots, d_{m+1} of \mathcal{N}' such that

$$d_1 < f(C_1) < d_2 < f(C_2) \dots < f(C_{m-1}) < d_m < f(C_m) < d_{m+1}.$$

Note that the same system of inequalities will hold if we replace any d_j by any one of its predecessors or successors. We now extend f to be an embedding defined on the substructure of \mathcal{M} generated by M_0 and a by defining $f(a) = d_j$ for a suitable value of j . An easy argument shows that this extends to an embedding of the entire substructure.

This completes the proof that T has QE. To conclude that T is complete, we apply Fact 7.3(3), using the fact that the structure $(\mathbb{Z}, <, p, s)$, in which $p(n) = n - 1$ and $s(n) = n + 1$ for all $n \in \mathbb{Z}$, can be embedded into every model of T .

Finally, it follows that T_{dis} is complete, since T is a conservative extension of T_{dis} . Indeed, let \mathcal{M}, \mathcal{N} be models of T_{dis} and let $\mathcal{M}', \mathcal{N}'$ be their unique expansions to models of T . Since T is complete we have $\mathcal{M}' \equiv \mathcal{N}'$. Taking reducts to L_0 we have $\mathcal{M} \equiv \mathcal{N}$. Since \mathcal{M}, \mathcal{N} were arbitrary models of T_{dis} this shows that T_{dis} is complete. \square

Another criterion for QE

When we are trying to prove that a theory has QE using Theorem 7.11(2) or Corollary 7.12, it is sometimes inconvenient that we must extend a given embedding f in $\text{Sub}(\mathcal{M}, \mathcal{N})$ to *every* element a of the model \mathcal{M} . The next result gives a criterion for QE in which we get to choose which element a to treat.

7.14. Theorem. *Let Σ be a set of L -sentences such that $\text{card}(L) = \kappa$. The following conditions are equivalent:*

- (1) Σ has QE.
- (2) Suppose \mathcal{M}, \mathcal{N} are models of Σ , $\text{card}(M) \leq \kappa$, and \mathcal{N} is κ^+ -saturated. Then for every $f \in \text{Sub}(\mathcal{M}, \mathcal{N})$, either $\text{dom}(f) = M$ or f has a proper extension to a function $g \in \text{Sub}(\mathcal{M}, \mathcal{N})$.

Proof. (1 \Rightarrow 2): Use Theorems 5.4 and 7.11(2).

(2 \Rightarrow 1): We assume condition (2) of this Theorem and use it to verify condition (2) of Theorem 7.11. Suppose \mathcal{M}, \mathcal{N} are arbitrary models of T ,

\mathcal{A}_0 is a substructure of \mathcal{M} , and f is an embedding of \mathcal{A}_0 into \mathcal{N} . Fix $a \in M$. We must show that f can be extended to an embedding of $\mathcal{B} = \langle \mathcal{A}_0 \cup \{a\} \rangle_{\mathcal{M}}$ into an elementary extension of \mathcal{N} .

Since B has cardinality $\leq \kappa$, we may consider an elementary substructure \mathcal{M}' of \mathcal{M} such that \mathcal{M}' has cardinality $\leq \kappa$ and contains B , by Theorem 6.1. Further, we may consider an elementary extension \mathcal{N}' of \mathcal{N} such that \mathcal{N}' is κ^+ -saturated, by Theorem 5.8. We will show that f can be extended to an embedding of \mathcal{M}' into \mathcal{N}' , which implies the desired condition.

Let Ω be the set of all extensions of f to embeddings whose domain is a substructure of \mathcal{M}' and whose range is a substructure of \mathcal{N}' . We consider Ω as a partially ordered set with $g \leq h$ defined to mean that h is an extension of g . If C is a chain in (Ω, \leq) , then one checks easily that the union of C is an element of Ω . Therefore, by Zorn's Lemma there is a maximal element g of (Ω, \leq) . Applying condition (2) of the Theorem to the embedding g and the models \mathcal{M}' and \mathcal{N}' , we see that g can be maximal only if it is defined on all of \mathcal{M}' . In particular g is defined on B , and thus it gives an extension of f as needed to verify condition (2) of Theorem 7.11. \square

EXERCISES

7.15. Let L be any first order language and let L' be any first order language that extends L by the addition of some set of new constant symbols. Let Σ be a set of L -sentences. Show that Σ has QE over L if and only if Σ has QE over L' . (Therefore, in showing that Σ has QE, it does no harm to assume that its language contains at least one constant symbol.)

7.16. Let L be a first order language and let T be an L -theory that has QE and is complete.

- If L contains at least one constant symbol, show that there exists a single L -structure that embeds into every model of T .
- Even when L has no constant symbol, show that there exists a single L -structure that embeds into every ω -saturated model of T . (That is, the converse to Corollary 5.5(2) is true.)

7.17. Let L be the language whose only nonlogical symbol is a binary predicate symbol $<$. Let \mathcal{M} be an L -structure that is a discrete linear ordering without endpoints. Let $\varphi(x, y_1, \dots, y_n)$ be any L -formula (with x a single variable) and let $a_1, \dots, a_n \in A$. Show that the definable set

$$\{a \in A \mid \mathcal{M} \models \varphi[a, a_1, \dots, a_n]\}$$

is the union of a finite number of open intervals (whose endpoints are in $A \cup \{-\infty, +\infty\}$) and a finite subset of A .

7.18. **Remark.** The result in Exercise 7.17 is expressed by saying that the theory T_{dis} is *o-minimal*.

7.19. Let K be a field and let L be the first order language of vector spaces over K , as described in Exercise 4.20. Let T be the theory of infinite K -vector spaces.

- Show that T has quantifier elimination and use this to prove that T is complete. (Compare Exercise 4.20.)
- Let $\mathcal{M} \models T$ and $X \subseteq A$. Give a clear mathematical description of the space of 1-types over X . That is, describe the space $S_1(X)$, including its topology.
- Let κ be any infinite cardinal $\geq \text{card}(K)$. Which models of T are κ -saturated?
- Show that there exist models \mathcal{M}, \mathcal{N} of T such that there does not exist any back-and-forth system from \mathcal{M} to \mathcal{N} . (Yet $\mathcal{M} \equiv \mathcal{N}$ since T is complete.)

7.20. Let \mathcal{Q} be the ordered field of rational numbers, considered as a structure for the first order language whose nonlogical symbols are the constant symbols $0, 1$, the binary predicate symbol $<$, and the binary function symbols $+, -, \times$, all with the obvious interpretations in \mathcal{Q} .

- Show that if $X \subseteq \mathbb{Q}$ is definable in \mathcal{Q} by a quantifier-free formula (in which some elements of \mathbb{Q} may be used as parameters), then there exists $q \in \mathbb{Q}$ such that the interval (q, ∞) in \mathbb{Q} is either contained in X or disjoint from X .
- Use the preceding result to show that $\text{Th}(\mathcal{Q})$ does not have QE.

8. ALGEBRAICALLY CLOSED FIELDS

We illustrate the use of Theorem 7.11 by using it to show that the theory of algebraically closed fields has QE. After proving this result we will show how it can be used to obtain some interesting consequences for algebraically closed fields. (These are well-known fundamental facts, whose model theoretic proofs give insight into the resonance between model theory and algebraic geometry.)

We formulate this theory in the language L_r of rings; this language has binary function symbols $+$, $-$, \times and constants $0, 1$. We let ACF denote a set of L_r -sentences that axiomatizes the class of algebraically closed fields considered as L_r -structures: ACF consists of the first order axioms for fields together with axioms asserting, for each $n \geq 1$, that every nontrivial polynomial of degree n has a root:

$$\forall y_0 \dots \forall y_n (y_0 \neq 0 \rightarrow \exists x (y_0 x^n + \dots + x_{n-1} x + x_n = 0)).$$

The best known algebraically closed field is the field \mathbb{C} of complex numbers (Fundamental Theorem of Algebra). Also, for each prime number p , the union of the canonical chain of finite fields of characteristic p , namely

$$\tilde{\mathbb{F}}_p = \bigcup_{n=1}^{\infty} \mathbb{F}_{p^n}$$

is algebraically closed. (Here \mathbb{F}_{p^n} denotes the field with p^n elements.) Indeed, $\tilde{\mathbb{F}}_p$ is the algebraic closure of $\mathbb{F}_p \cong \mathbb{Z}/p\mathbb{Z}$; that is, it is algebraically closed and algebraic over \mathbb{F}_p .

In our proof of QE for ACF we use a small amount of the basic theory of fields, mainly concerning simple properties of polynomials in one variable over a given field. These concern the process of extending a field by adjoining a root of a given polynomial. Iterating this procedure, one shows that every field is contained in an algebraically closed field. Most graduate texts in algebra contain this basic material.

More advanced properties of algebraically closed fields, such as the uniqueness of the algebraic closure of a field and the properties of transcendence bases for algebraically closed fields, are not needed for this proof of QE for ACF . Indeed, they can be proved efficiently using the model theoretic ideas discussed here, as we show in Chapter 11.

We regard any field K as an L_r -structure in the obvious way, and note that the substructures of K are exactly the subrings of K . Obviously the rings that occur as substructures of some field are exactly the domains. If F, K are fields, then the elements of $\text{Sub}(F, K)$ are exactly the ring isomorphisms $f: R \rightarrow S$, where R is a subring of F and S is a subring of K .

8.1. Theorem. *The set of axioms ACF has QE.*

Proof. We verify condition (2) of Theorem 7.11. We need to consider algebraically closed fields F, K , as well as a ring isomorphism $f: R \rightarrow S$ from $\text{Sub}(F, K)$ and an element a of F . We must prove that f can be extended to an embedding of a subring R' of F into K , where $R' \supseteq R \cup \{a\}$.

First we consider the field of fractions \bar{R} of R inside F . It is easy to see that f can be extended (in a unique way) to an embedding of \bar{R} into K , which we also denote by f ; for each $b, c \in R$ with $c \neq 0$ we define $f(b/c) = f(b)/f(c)$. The range of this extended f is obviously the field of fractions of S inside K . From this argument we see that we may assume that R is a subfield of F and S is a subfield of K .

Next we consider the algebraic closure \tilde{R} of R in F . Given $a \in \tilde{R}$, let $p(x)$ be the minimal polynomial of a over R , so that $p(x)$ is an irreducible polynomial in $R[x]$ and $p(a) = 0$. Let $q(x)$ be the corresponding polynomial in $S[x]$, obtained by applying f to the coefficients of $p(x)$. Since f is an isomorphism of fields, $q(x)$ is irreducible in $S[x]$. The field K is algebraically closed, so $q(x)$ has a root in this field. Let $b \in K$ be such a root. It is an elementary exercise to show that f can be extended (in a unique way) to an isomorphism from $R(a)$ onto $S(b) \subseteq K$ such that $f(a) = b$. Applying this type of extension inductively to all elements of \tilde{R} , we extend f to an embedding of \tilde{R} into K . Hence we may assume that R is algebraically closed in F .

Finally, consider $a \in F \setminus R$; we have that a is transcendental over R . Let K' be a κ -saturated elementary extension of K , where $\kappa > \text{card}(R)$. Since $\text{card}(S) = \text{card}(R) < \kappa$, and hence $\text{card}(S[x]) < \kappa$, there must be an element b of K' that is transcendental over S . We may extend f (in a unique way) to an isomorphism of $R[a]$ onto $S[b] \subseteq K'$ such that $f(a) = b$; namely, for each polynomial $p(x)$ in $R[x]$ we define $f(p(a)) = p(b)$.

In all cases we have extended the original $f \in \text{Sub}(F, K)$ to an embedding whose domain is a subring of F that contains a given element a and whose range is a substructure of an elementary extension of K . Consequently, ACF satisfies condition 7.11(2), completing the proof. \square

For any integer $n \geq 2$, let σ_n denote the L_r -sentence $1 + \dots + 1 = 0$ in which there are n occurrences of 1 in the summation. For each prime p , we let $ACF_p = ACF \cup \{\sigma_p\}$, and $ACF_0 = ACF \cup \{-\sigma_p \mid p \text{ is a prime}\}$.

For each prime number p , the models of ACF_p are the algebraically closed fields that have characteristic p . Similarly, the models of ACF_0 are the algebraically closed fields of characteristic 0. Note that $\tilde{\mathbb{F}}_p$ is a model of ACF_p for each prime p , and \mathbb{C} is a model of ACF_0 .

8.2. Corollary. (i) For each prime p , ACF_p is complete, and ACF_0 is also complete. These sets axiomatize all the completions of ACF .

(ii) For each sentence σ in the language of rings, $ACF_0 \models \sigma$ iff $ACF_p \models \sigma$ for all sufficiently large primes p iff $ACF_p \models \sigma$ for infinitely many primes p .

Proof. (i) We apply Fact 7.3(3). For each prime p , the field \mathbb{F}_p with p elements embeds in every field of characteristic p , and thus into every model of ACF_p . A similar argument, replacing \mathbb{F}_p by the ring \mathbb{Z} , shows that ACF_0 is complete. Every model of ACF , namely every algebraically closed field K , is a model of one of these sets of axioms (the one determined by the characteristic of K), so there cannot be any other completions of ACF .

(ii) Let σ be a sentence in L_r . If $ACF_0 \models \sigma$ then by Corollary 3.6 there is an integer n such that $ACF \cup \{\neg\sigma_p \mid p \text{ is a prime } \leq n\} \models \sigma$. This proves the other two conditions. Conversely, suppose $ACF_0 \not\models \sigma$. Because ACF_0 is complete we have $ACF_0 \models \neg\sigma$ so there exists a positive integer n such that $ACF \cup \{\neg\sigma_p \mid p \text{ is a prime } \leq n\} \models \neg\sigma$. It follows that there can only exist finitely many primes p such that $ACF_p \models \sigma$. \square

8.3. Application. *Suppose K is any algebraically closed field and $f: K^n \rightarrow K^n$ is a polynomial map. If f is 1-1, then f is onto.*

Proof. This striking model theoretic proof was discovered by James Ax. Let $f = (f_1, \dots, f_n)$ be a polynomial map from K^n to itself, so each f_j is defined by a polynomial in n variables with coefficients from K . Let d be a positive integer larger than the degrees of all the polynomials that are involved in defining f_1, \dots, f_n .

It is easy to construct a sentence τ_d in the language of rings such that for any field k , we have $k \models \tau_d$ if and only if for every polynomial map $f: k^n \rightarrow k^n$ defined by polynomials over k having degree at most d , if f is 1-1 then f is onto. We are trying to show $K \models \tau_d$ for each algebraically closed field K . By Corollary 8.2(ii) it suffices to prove $ACF_p \models \tau_d$ for every prime p . Moreover, because ACF_p is complete, it suffices to find for each prime p an algebraically closed field K_p of characteristic p such that $K_p \models \tau_d$. We will prove for every prime p that $\tilde{\mathbb{F}}_p$, the algebraic closure of the prime field \mathbb{F}_p of characteristic p , satisfies the sentence τ_d .

Fix a prime p and let $f: \tilde{\mathbb{F}}_p^n \rightarrow \tilde{\mathbb{F}}_p^n$ be a polynomial map that is 1-1. Fix any element (a_1, \dots, a_n) in $\tilde{\mathbb{F}}_p^n$. There is a finite subfield k of $\tilde{\mathbb{F}}_p$ that contains a_1, \dots, a_n and all coefficients of the polynomials that define the coordinate functions of f . Therefore f restricted to k^n is a 1-1 map into k^n . Since k^n is finite this implies that the restriction of f to k^n maps onto k^n . In particular (a_1, \dots, a_n) is in the range of f , proving that f is onto. \square

The fact that ACF has QE implies that ACF is *model complete*: that is, whenever F and K are algebraically closed fields, and F is a subfield of K , then F is an elementary substructure of K . The following result, Hilbert's Nullstellensatz, is an easy consequence of this fact.

8.4. Application. *Let K be a field and suppose f_1, \dots, f_m are polynomials in the variables x_1, \dots, x_n with coefficients in K . If the system of equations $f_1(x_1, \dots, x_n) = \dots = f_m(x_1, \dots, x_n) = 0$ has a solution in some extension*

field of K , then this system has a solution in some finite degree (and hence algebraic) extension of K .

Proof. Let K and f_1, \dots, f_m be as stated, and suppose the system of equations $f_1 = \dots = f_m = 0$ has a solution in the extension field $K' \supseteq K$. Without loss of generality we may take K' to be algebraically closed, since every field is contained in an algebraically closed field. Let \tilde{K} be the algebraic closure of K in K' ; then \tilde{K} is itself an algebraically closed field. It follows from Theorem 8.1 that $\tilde{K} \preceq K'$. Note that the existence of a solution of the system of equations $f_1(x_1, \dots, x_n) = \dots = f_m(x_1, \dots, x_n) = 0$ can be expressed by an existential L_K -sentence. (Constants are needed to name the elements of K that appear as coefficients in the polynomials.) This sentence is true in K' , and therefore it is true in \tilde{K} , and the coordinates of a solution to $f_1 = \dots = f_m = 0$ in \tilde{K} generate a finite degree extension of K . \square

8.5. Definition. Let K be an algebraically closed field and X a subset of K^n . We say that X is *constructible* if it is a finite Boolean combination of zero sets of polynomials in $K[x_1, \dots, x_n]$.

8.6. Remark. It follows from Theorem 8.1 that *all* definable sets in K^n are constructible, and the converse is obviously true. In particular, the collection of all constructible sets is closed under projections.

8.7. Application (Chevalley). *Let K be an algebraically closed field. If X is a constructible subset of K^n and if $h = (h_1, \dots, h_m)$ is a polynomial map over K from K^n to K^m , then $h(X)$ is a constructible subset of K^m .*

Proof. To say that h is “over K ” means that h_1, \dots, h_m are polynomials with coefficients from K . Let $\varphi(y)$ be an $L_{r,K}$ -formula that defines X in K_K ; here we write y for y_1, \dots, y_n . Then $h(X)$ is defined in K_K by the $L_{r,K}$ -formula $\psi(x)$ given by

$$\exists y_1 \dots \exists y_n (h_1(y) = x_1 \wedge \dots \wedge h_m(y) = x_m \wedge \varphi(y)).$$

Therefore $h(X)$ is also constructible, by Remark 8.6. \square

8.8. Definition. An infinite L -structure \mathcal{M} is *minimal* if every definable subset of M is either finite or cofinite. That is, given any L -formula $\varphi(x, y_1, \dots, y_n)$ (in which x is a single variable) and any parameters $a_1, \dots, a_n \in M$, the set

$$\{a \in M \mid \mathcal{M} \models \varphi[a, a_1, \dots, a_n]\}$$

is either finite or cofinite as a subset of M .

A set of L -sentences Σ is *strongly minimal* if every infinite model of Σ is minimal.

An L -structure \mathcal{M} is *strongly minimal* if $\text{Th } \mathcal{M}$ is strongly minimal.

8.9. Proposition. *Let Σ be a strongly minimal set of L -sentences and let $\varphi(x, y_1, \dots, y_n)$ be any L -formula (in which x is a single variable). Then there exists an integer m with the property that for any model \mathcal{M} of Σ and any $b_1, \dots, b_n \in M$, either the set $\{a \in M \mid \mathcal{M} \models \varphi[a, b_1, \dots, b_n]\}$ or its complement in M has at most m elements.*

Proof. The proof is a straightforward compactness argument. Suppose Σ is strongly minimal and $\varphi(x, y_1, \dots, y_n)$ is a formula for which the conclusion fails. For each m there must be a model \mathcal{M} of T and parameters $b_1, \dots, b_n \in M$, such that both the set $\{a \in M \mid \mathcal{M} \models \varphi[a, b_1, \dots, b_n]\}$ and its complement have at least $m + 1$ elements. This fact can be expressed by a formula $\psi_{>m}(b_1, \dots, b_n)$. From the compactness theorem applied to $\Sigma \cup \{\psi_{>m}(y_1, \dots, y_n) \mid m \geq 1\}$ it follows that Σ has a model \mathcal{M} with parameters $b_1, \dots, b_n \in M$, such that both the set $\{a \in M \mid \mathcal{M} \models \varphi[a, b_1, \dots, b_n]\}$ and its complement are infinite. Therefore \mathcal{M} is a nonminimal model of Σ , which is a contradiction. \square

8.10. Corollary. *The set of axioms ACF is strongly minimal.*

Proof. Let K be an algebraically closed field and let X be a definable subset of K . By Theorem 8.1 we can define X by a quantifier free L_K -formula $\varphi(x)$, in which x is a single variable. The formula φ is equivalent to a Boolean combination of finitely many equations of the form $p(x) = 0$ where $p(x)$ is a polynomial with coefficients in K . We may assume that all the polynomials that appear in φ are nonconstant. Therefore, either X or $K \setminus X$ is contained in the union of finitely many zero sets of nonconstant polynomials. It follows that X or $K \setminus X$ must be finite. \square

8.11. Remark. In Chapter 11 we prove some results about the structure of strongly minimal models that can be used to show easily that every algebraically closed field is uniquely determined by its characteristic and its transcendence degree (Steinitz's Theorem).

EXERCISES

8.12. Let K be an algebraically closed field, considered as an L_r -structure; let A be any subset of K and let k be the subfield of K generated by A . Let \mathcal{M} denote the $L_{r,A}$ -structure $(K, a)_{a \in A}$.

- For any $a, b \in K$, show that $\text{tp}_{\mathcal{M}}(a) = \text{tp}_{\mathcal{M}}(b)$ iff either a, b are both transcendental over k or both a, b are algebraic over k and have the same minimal polynomial over k .

8.13. If Σ is a set of L -sentences, a model \mathcal{M} of Σ is called *existentially closed* in $\text{Mod}(\Sigma)$ if it satisfies the following condition: whenever $\mathcal{M} \subseteq \mathcal{N} \models \Sigma$, $\varphi(x_1, \dots, x_m, y_1, \dots, y_n)$ is a quantifier-free formula, and $a_1, \dots, a_m \in M$, then $\mathcal{N} \models \exists y_1 \dots \exists y_n \varphi[a_1, \dots, a_m]$ implies $\mathcal{M} \models \exists y_1 \dots \exists y_n \varphi[a_1, \dots, a_m]$.

- Consider a set of axioms Σ for the class of fields (in the language L_r). Show that a field K is existentially closed in the class of all fields iff K is algebraically closed.

9. \mathbb{Z} -GROUPS

In this chapter we will apply the *method of quantifier elimination* to analyze the first order theory and definable sets of the ordered abelian group of the integers $(\mathbb{Z}, +, -, <, 0)$.

It is easy to see that $\text{Th}(\mathbb{Z}, +, -, <, 0)$ does not have QE. The number 1 is definable (as the smallest positive element of \mathbb{Z}) as are the divisibility predicates D_n defined for $n \geq 2$ by

$$D_n(x) \iff \exists y(x = ny).$$

Here we are using ny to represent the term $y + \dots + y$ in which there are n occurrences of y . Neither 1 nor D_n can be defined by quantifier free formulas in $(\mathbb{Z}, +, -, <, 0)$. It turns out that if we add symbols for the element 1 and the predicates D_n to the language, and thus take the structure $(\mathbb{Z}, +, -, <, 0, 1, D_n)_{n \geq 2}$ as the basic object of study, then the resulting theory does have QE and we do achieve a useful analysis of the definable sets. Further, we are able to axiomatize this theory using a clear and simple set of sentences.

Let L be the language of this structure. It has binary function symbols $+$, $-$, a binary relation symbol $<$, constant symbols $0, 1$, and an infinite family of unary relation symbols D_n for $n \geq 2$. In L we formulate the theory T of \mathbb{Z} -groups, which has the following axioms: (a) the axioms of ordered abelian groups; (b) the axiom that 1 is the smallest positive element; (c) the divisibility axioms (given above in the displayed formula) that define each D_n in terms of the group structure; and (d) the congruence axioms:

$$\forall x(D_n(x+1) \vee D_n(x+2) \vee \dots \vee D_n(x+n)).$$

for each $n \geq 2$. (Here we write k in place of the term $k1$ for each positive integer k .) These congruence axioms express the property of division by n with remainder.

9.1. Lemma. *For each $n \geq 2$ and $1 \leq i < j \leq n$*

- (1) $T \models \forall x \forall y ((D_n(x) \wedge D_n(y)) \rightarrow D_n(x+y));$
- (2) $T \models \forall x (D_n(x) \rightarrow D_n(-x));$
- (2) $T \models \forall x (D_n(x+i) \rightarrow \neg D_n(x+j)).$

Proof. We argue informally in T . (1) If $x = nu$ and $y = nv$ then $x + y = n(u + v)$. (2) If $x = nu$ then $-x = n(-u)$. (3) Argue by contradiction; suppose $1 \leq i < j \leq n$, $x + i = nu$, and $x + j = nv$. Then $j - i = n(v - u)$. It follows that $0 < v - u < 1$, contradicting one of the axioms of T . \square

9.2. Theorem. *The theory T of \mathbb{Z} -groups has quantifier elimination. Moreover, T is complete and therefore $T = \text{Th}(\mathbb{Z}, +, -, <, 0, 1, D_n)_{n \geq 2}$.*

In proving this Theorem we use explicit methods for eliminating quantifiers, rather than the model theoretic methods presented in Chapter 7. To do this we need to introduce some definitions and a Lemma.

9.3. Definition. (a) An *existential* formula is a formula in prenex normal form that has only \exists quantifier symbols in its prefix. (b) An existential formula is *primitive* if it is of the form

$$\exists x_1 \dots \exists x_n \varphi(x_1, \dots, x_n)$$

where φ is a conjunction of literals; a literal is either an atomic formula or the negation of an atomic formula. (c) A *universal* formula is a formula in prenex normal form that has only \forall quantifier symbols in its prefix.

9.4. Lemma. *Let T be an L -theory. If every primitive existential formula with a single existential quantifier is equivalent in T to a quantifier free formula, then T has quantifier elimination.*

Proof. It suffices to prove that every prenex formula is equivalent in T to a quantifier free formula. We do this by induction on the number of quantifiers in the prefix of the prenex formula.

We show first that every existential formula with just one existential quantifier is equivalent in T to a quantifier free formula. Each such formula is logically equivalent to a disjunction of primitive existential formulas, each of which also has just a single existential quantifier. Each of these disjuncts is equivalent in T to a quantifier free formula, by hypothesis. Hence the original existential formula is equivalent in T to a quantifier free formula.

By taking negations, it follows that every universal formula with a single quantifier in its prefix is also equivalent in T to a quantifier free formula.

The induction step is carried out by using the above results to eliminate the innermost quantifier in the prefix, and then using the induction hypothesis to eliminate the remaining quantifiers. \square

Proof of Theorem 9.2. We will give an explicit proof of quantifier elimination. The completeness of T follows using Fact 7.3(3), using the fact that the structure $(\mathbb{Z}, +, -, <, 0, 1, D_n)_{n \geq 2}$ can be embedded in every model of T .

Let φ be any existential L -formula with a single existential quantifier, of the form $\exists x \psi$ with ψ quantifier free. We first observe that we may assume ψ is a positive Boolean combination of atomic formulas (i.e. using only the connectives \wedge, \vee). This is because each negation of an atomic formula is equivalent in T to a positive combination of atomic formulas. Namely: $\neg t = s$ is equivalent to $t < s \vee s < t$; $\neg t < s$ is equivalent to $s < t \vee s = t$; and $\neg D_n(t)$ is equivalent to $D_n(t+1) \vee \dots \vee D_n(t+(n-1))$ by Lemma 9.1 and the congruence axioms of T . By putting ψ in disjunctive normal form and distributing the existential quantifier $\exists x$ over the connective \vee , we see that φ is equivalent in T to a disjunction of existential formulas $\exists x \theta$ where each θ is a conjunction of atomic formulas. Arguing as in the proof of Lemma 9.4 it suffices to prove that every such formula is equivalent in T to a quantifier free formula.

We next observe that every atomic formula in L is equivalent in T either to an atomic formula in which x does not occur or to one of the following: $nx = t$, $nx < t$, $t < nx$, or $D_m(nx + t)$, where n is an integer > 0 and t is a term not containing x . In such atomic formulas we will call n a “coefficient of x ”, and m a “divisor”.

Let $\theta(x, y_1, \dots, y_k)$ be any conjunction of atomic formulas as in the previous paragraph. We may assume that x actually occurs in θ , since otherwise $\exists x\theta(x, y_1, \dots, y_k)$ is equivalent to the quantifier free formula $\theta(0, y_1, \dots, y_k)$. We show next that θ is equivalent to an L -formula of the same form in which the only coefficient of x that occurs is 1. Let N be the least common multiple of all coefficients of x that occur in θ . Multiplying each term in θ by a suitable positive integer, we may assume that every coefficient of x in θ is equal to N . (If n is a coefficient of x in θ and $N = dn$, then we replace $nx = t$ by $Nx = dt$, $nx < t$ by $Nx < dt$, $t < nx$ by $dt < Nx$, and $D_m(nx + t)$ by $D_{dm}(Nx + dt)$.) Let $\theta'(z, y_1, \dots, y_k)$ be the result of replacing each occurrence of Nx in θ by z . Evidently $\exists x\theta(x, y_1, \dots, y_k)$ is equivalent in T to $\exists z(D_N(z) \wedge \theta'(z, y_1, \dots, y_k))$.

Therefore we need only consider $\theta(x, y_1, \dots, y_k)$ that are conjunctions of atomic formulas of the form $x = t$, $x < t$, $t < x$, or $D_m(x + t)$, where t is a term not containing x , and in which at least one atomic formula of the form $D_m(x + t)$ occurs. We will now show that $\varphi = \exists x\theta(x, y_1, \dots, y_k)$ is equivalent in T to a quantifier free formula, by treating a series of cases.

Let M be the least common multiple of all divisors occurring in θ .

Case (1): θ contains at least one conjunct of the form $x = t$. Then φ is equivalent to $\theta(t, y_1, \dots, y_k)$.

Case (2): θ contains no conjuncts of the form $x = t$ but does contain at least one conjunct of the form $x < t$. Let t_1, \dots, t_p be all terms t such that $x < t$ occurs in θ . Then φ is equivalent to the disjunction of all formulas $\theta(t_i - j, y_1, \dots, y_k)$ where $1 \leq i \leq p$ and $1 \leq j \leq M$. Arguing informally in T we can see this as follows: suppose x witnesses the truth of θ , and t represents the minimum of t_1, \dots, t_p ; choose $j \in \{1, \dots, M\}$ such that $D_M(x - (t - j))$ holds. The axioms of T guarantee that such a choice exists and (using Lemma 9.1) is unique. It is now easy to see that $t - j$ also witnesses the truth of θ .

Case (3): θ contains no conjuncts of the form $x = t$ but does contain at least one conjunct of the form $t < x$. Let t_1, \dots, t_p be all terms t such that $t < x$ occurs in θ . Then φ is equivalent to the disjunction of all formulas $\theta(t_i + j, y_1, \dots, y_k)$ where $1 \leq i \leq p$ and $1 \leq j \leq M$.

Case (4): θ contains only atomic formulas of the form $D_m(x + t)$. In this case φ is equivalent to the disjunction of all formulas of the form $\theta(j, y_1, \dots, y_k)$ where $1 \leq j \leq M$.

This completes the proof that $\varphi = \exists x\theta(x, y_1, \dots, y_k)$ is equivalent in T to a quantifier free formula. \square

Our objective was to analyze the ordered abelian group $(\mathbb{Z}, +, -, <, 0)$. Let L_0 be the language of this structure and let T_0 be the theory in L_0 whose axioms are (i) the axioms of ordered abelian groups; (ii) the existence of a smallest positive element; (iii) the congruence axioms

$$\forall x \exists y (x + 1 = ny \vee x + 2 = ny \vee \dots \vee x + n = ny).$$

for each $n \geq 2$. It is clear that each model of T_0 can be expanded in a unique way to a model of T . Indeed, one simply lets 1 be interpreted by the smallest positive element of the model and takes D_n to be interpreted as “divisibility by n ” for each $n \geq 2$. Therefore T is a conservative extension of T_0 , from which it follows that T_0 is complete and therefore $T_0 = \text{Th}(\mathbb{Z}, +, -, <, 0)$.

We obtain a deeper result if we expand L_0 to add the constant symbol 1 and extend T_0 by adding the axiom stating that 1 is the smallest positive element. Let L_1 be the resulting language and T_1 the resulting theory. Evidently each model of T_1 expands uniquely to a model of T ; therefore T_1 is complete and $T_1 = \text{Th}(\mathbb{Z}, +, -, <, 0, 1)$. By looking closer at the relation between T and T_1 we obtain the following result:

9.5. Corollary. *T_1 is model complete; that is if \mathcal{M}, \mathcal{N} are models of T_1 and $\mathcal{M} \subseteq \mathcal{N}$, then $\mathcal{M} \preceq \mathcal{N}$.*

Proof. Let \mathcal{M}, \mathcal{N} be models of T_1 with $\mathcal{M} \subseteq \mathcal{N}$. Let \mathcal{N}' be the unique expansion of \mathcal{N} to a model of T . The set M is the universe of a substructure of \mathcal{N}' , which we denote by \mathcal{M}' . We will show that \mathcal{M}' is a model of T . Therefore, since it is an expansion of \mathcal{M} , it is the unique expansion of this structure to a model of T .

To show that \mathcal{M}' is a model of T we need only consider the divisibility axioms, which define D_n in terms of the abelian group structure. The congruence axioms are implied by the divisibility axioms over T_1 , which we know is satisfied by \mathcal{M}' (since it is satisfied by \mathcal{M}). Fix an element a of M . If $a = nb$ for some b in M , then this equation also holds in \mathcal{N} , which implies that a satisfies $D_n(x)$ in \mathcal{N}' since it is a model of T . Therefore a satisfies $D_n(x)$ in \mathcal{M}' , since it is a substructure of \mathcal{N}' . Conversely, suppose a is not of the form nb in \mathcal{M} . There must exist a unique $k = 1, \dots, n - 1$ and some $b \in M$ satisfying $a + i = nb$ in \mathcal{M} . This equation also holds in \mathcal{N}' , which implies that $D_n(x)$ must be false of a in that structure. Hence $D_n(x)$ is also false of a in \mathcal{M}' by the substructure condition.

Thus we have proved \mathcal{M}' is a model of T . Since T has QE and is therefore model complete itself, we conclude $\mathcal{M}' \preceq \mathcal{N}'$. It follows by restricting to L_1 that $\mathcal{M} \preceq \mathcal{N}$, and the proof is complete. \square

The key point in the preceding proof is that both $D_n(x)$ and $\neg D_n(x)$ are equivalent in T to universal formulas of L_1 .

Note that T_0 is not model complete. Indeed, the function $f: \mathbb{Z} \rightarrow \mathbb{Z}$ defined by setting $f(n) = 2n$ for all n is clearly an embedding of $(\mathbb{Z}, +, -, <, 0)$ into itself but it is not an elementary embedding.

We finish this chapter by using Theorem 9.2 to characterize the subsets of \mathbb{Z} that are definable in $(\mathbb{Z}, +, -, <, 0)$. It turns out to be necessary to distinguish the positive part of a definable set from the negative part, as these can be defined independently of each other.

9.6. Definition. Let $A \subseteq \mathbb{N}$. We call A *eventually periodic* if there are $n \geq 0$ and $p > 0$ such that for all $m \in \mathbb{N}$,
if $m \geq n$, then $m \in A \iff m + p \in A$.

Evidently $A \subseteq \mathbb{N}$ is eventually periodic if and only if it is the union of a finite number of arithmetic progressions and a finite set. Moreover, the collection of all eventually periodic sets is a Boolean algebra of subsets of \mathbb{N} . Note that each eventually periodic set A is definable in $(\mathbb{Z}, +, -, <, 0)$ as is $-A = \{-n \mid n \in A\}$.

9.7. Corollary. *The subsets of \mathbb{Z} that are definable in $(\mathbb{Z}, +, -, <, 0)$ are exactly the sets of the form $(-A) \cup B$, where A and B are eventually periodic subsets of \mathbb{N} .*

Proof. Let \mathcal{P} be the collection of all subsets of \mathbb{Z} of the form $-A \cup B$, where A and B are eventually periodic subsets of \mathbb{N} . Clearly every set in \mathcal{P} is definable in $(\mathbb{Z}, +, -, <, 0)$. It is routine to show that \mathcal{P} is closed under union, intersection, and complement in \mathbb{Z} . By this remark and Theorem 9.2 it suffices to show that each set defined by an atomic L -formula in the structure $(\mathbb{Z}, +, -, <, 0, 1, D_n)_{n \geq 2}$ belongs to \mathcal{P} . Arguing as in the proof of Theorem 9.2 we see it suffices to consider atomic formulas $\varphi(x)$ of the following forms: $nx = t$, $nx < t$, $t < nx$, and $D_m(nx + t)$, where n is a positive integer, $m \geq 2$, and t is a term without variables. In each case it is easy to see that the set defined in $(\mathbb{Z}, +, -, <, 0, 1, D_n)_{n \geq 2}$ by $\varphi(x)$ belongs to \mathcal{P} . \square

10. MODEL THEORETIC ALGEBRAIC CLOSURE

10.1. Definition. Let \mathcal{M} be an L -structure and $A \subseteq M$. An element a of M is *algebraic over A in \mathcal{M}* if there is an L -formula $\varphi(x, y_1, \dots, y_n)$ and elements e_1, \dots, e_n of A such that

- (i) $\mathcal{M} \models \varphi[a, e_1, \dots, e_n]$, and
- (ii) $\{c \in M \mid \mathcal{M} \models \varphi[c, e_1, \dots, e_n]\}$ is finite.

The set of elements of M that are algebraic over A in \mathcal{M} is denoted by $\text{acl}_{\mathcal{M}}(A)$, or simply by $\text{acl}(A)$ when the structure \mathcal{M} is understood.

A is *algebraically closed in \mathcal{M}* if $\text{acl}_{\mathcal{M}}(A) = A$.

Note. It is immediate from Definition 10.1 that the operation $\text{acl}_{\mathcal{M}}$ has *finite character*; that is, for every subset A of M we have that $\text{acl}_{\mathcal{M}}(A)$ is the union of the sets $\text{acl}_{\mathcal{M}}(F)$ where F ranges over the finite subsets of A .

10.2. Fact. If $A \subseteq M$ and $\text{card}(A), \text{card}(L) \leq \kappa$, then $\text{card}(\text{acl}_{\mathcal{M}}(A)) \leq \kappa$.

10.3. Proposition. *Let \mathcal{M} be an L -structure. The operation $\text{acl}_{\mathcal{M}}(A)$ defined on all subsets A of M is a closure operation. That is, it satisfies the following two properties for $A, B \subseteq M$:*

- (1) $A \subseteq \text{acl}_{\mathcal{M}}(A)$; and
- (2) if $B \subseteq \text{acl}_{\mathcal{M}}(A)$ then $\text{acl}_{\mathcal{M}}(B) \subseteq \text{acl}_{\mathcal{M}}(A)$.

Proof. (1) If $a \in A$, then $a \in \text{acl}_{\mathcal{M}}(A)$ is witnessed by the formula $x = y_1$ with parameter a .

(2) Assume $B \subseteq \text{acl}_{\mathcal{M}}(A)$ and $a \in \text{acl}_{\mathcal{M}}(B)$. Let $\varphi(x, y_1, \dots, y_n)$ and $e_1, \dots, e_n \in B$ witness the fact that $a \in \text{acl}_{\mathcal{M}}(B)$ as in Definition 10.1. Let m be the cardinality of the set $\{c \in M \mid \mathcal{M} \models \varphi[c, e_1, \dots, e_n]\}$. By changing the formula φ if necessary we may assume that for every $b_1, \dots, b_n \in M$ the set $\{c \in M \mid \mathcal{M} \models \varphi[c, b_1, \dots, b_n]\}$ has cardinality at most m , while we continue to have $\mathcal{M} \models \varphi[a, e_1, \dots, e_n]$. Similarly, let $\psi_j(y_j, z_1, \dots, z_p)$ and $f_1, \dots, f_p \in A$ witness the fact that $e_j \in \text{acl}_{\mathcal{M}}(A)$ for each $j = 1, \dots, p$. (We have unified the lists of parameters and added extra variables in the formulas to ensure that the parameters are the same for each j . There is no loss of generality in doing so.) Then the formula $\sigma(x, z_1, \dots, z_p)$ given by

$$\exists y_1 \dots \exists y_n (\varphi(x, y) \wedge \psi_1(y_1, z) \wedge \dots \wedge \psi_n(y_n, z))$$

with parameters f_1, \dots, f_p witnesses the fact that $a \in \text{acl}_{\mathcal{M}}(A)$. (Here we write y for y_1, \dots, y_n and z for z_1, \dots, z_p .) \square

10.4. Fact. For each $A, B \subseteq M$, the following properties follow quickly from the ones given in Proposition 10.3:

- (3) If $A \subseteq B$, then $\text{acl}_{\mathcal{M}}(A) \subseteq \text{acl}_{\mathcal{M}}(B)$.
- (4) The set $\text{acl}_{\mathcal{M}}(A)$ is algebraically closed: $\text{acl}_{\mathcal{M}}(\text{acl}_{\mathcal{M}}(A)) = \text{acl}_{\mathcal{M}}(A)$.

Proof. (3) Applying 10.3(1) to B yields $A \subseteq \text{acl}_{\mathcal{M}}(B)$; then apply 10.3(2) to this containment.

(4) Applying 10.3(2) to $B = \text{acl}_{\mathcal{M}}(A)$ gives $\text{acl}_{\mathcal{M}}(\text{acl}_{\mathcal{M}}(A)) \subseteq \text{acl}_{\mathcal{M}}(A)$; 10.3(1) gives the reverse containment. \square

The following result shows that model theoretic algebraic closure is well behaved with respect to elementary maps.

10.5. Proposition. *Let \mathcal{M}, \mathcal{N} be L -structures, $A \subseteq M$, and $B \subseteq N$. If the function $f: A \rightarrow B$ is elementary with respect to \mathcal{M}, \mathcal{N} , then f can be extended to a function $g: \text{acl}_{\mathcal{M}}(A) \rightarrow \text{acl}_{\mathcal{N}}(B)$ that is elementary with respect to \mathcal{M}, \mathcal{N} . Moreover, if f is surjective, then any such g must also be surjective.*

Proof. Let $\mathcal{M}, \mathcal{N}, A, B, f$ be as given in the Proposition. Let Ω be the set of all functions $g: A' \rightarrow B'$ such that $A \subseteq A' \subseteq \text{acl}_{\mathcal{M}}(A)$, $B \subseteq B' \subseteq \text{acl}_{\mathcal{N}}(B)$, g is elementary with respect to \mathcal{M}, \mathcal{N} , and g extends f . It is easy to show that (Ω, \subseteq) is closed under unions of linearly ordered chains, so it satisfies the hypothesis of Zorn's Lemma. Therefore there exists $g \in \Omega$ that is maximal under \subseteq . We must show that the domain of g is $\text{acl}_{\mathcal{M}}(A)$. If not, let $a \in \text{acl}_{\mathcal{M}}(A) \setminus A'$. By Proposition 10.3 we have $a \in \text{acl}_{\mathcal{M}}(A')$. Let $\varphi(x, y_1, \dots, y_n)$ be an L -formula and $e_1, \dots, e_n \in A'$ be parameters that witness the fact that $a \in \text{acl}_{\mathcal{M}}(A')$. Moreover, we may suppose that φ and e have been chosen so that the finite set $U = \{c \in M \mid \mathcal{M} \models \varphi[c, e_1, \dots, e_n]\}$ has the smallest possible cardinality. Let this cardinality be m .

Since g is an elementary map, the set

$$V = \{c \in N \mid \mathcal{N} \models \varphi[c, g(e_1), \dots, g(e_n)]\}$$

also has cardinality m . Moreover, g maps $A' \cap U$ bijectively onto $B' \cap V$. Since $A' \cap U$ has cardinality $< m$ (as it does not contain a) there must exist $b \in V \setminus B'$. Extend g to the map g' defined on $A' \cup \{a\}$ by setting $g'(a) = b$. We will show that g' is elementary with respect to \mathcal{M}, \mathcal{N} , contradicting the maximality of g .

To that end, suppose $\psi(x, z_1, \dots, z_p)$ is any L -formula and $f_1, \dots, f_p \in A'$ are such that $\mathcal{M} \models \psi[a, f_1, \dots, f_p]$. The formula $\varphi(x, y) \wedge \psi(x, z)$ and the parameters e, f witness the fact that $a \in \text{acl}_{\mathcal{M}}(A')$. Therefore the choice of φ and e ensure that

$$\mathcal{M} \models \forall x(\varphi \rightarrow \psi)[e_1, \dots, e_n, f_1, \dots, f_p].$$

Since g is elementary, we have

$$\mathcal{N} \models \forall x(\varphi \rightarrow \psi)[g(e_1), \dots, g(e_n), g(f_1), \dots, g(f_p)].$$

Therefore our choice of b implies

$$\mathcal{N} \models \psi[b, g(f_1), \dots, g(f_p)].$$

This completes the proof that g' is elementary and therefore we may conclude that the domain of the maximal g in Ω is all of $\text{acl}_{\mathcal{M}}(A)$.

Finally, suppose the given function f has range B . Let $g: \text{acl}_{\mathcal{M}}(A) \rightarrow \text{acl}_{\mathcal{N}}(B)$ be any extension of f that is elementary with respect to \mathcal{M}, \mathcal{N} . Let C be the range of g and suppose, by way of getting a contradiction, that C is a proper subset of $\text{acl}_{\mathcal{N}}(B)$. Since $B \subseteq C$ we have $\text{acl}_{\mathcal{N}}(C) = \text{acl}_{\mathcal{N}}(B)$.

Applying the first part of this Proposition to g^{-1} we see that g^{-1} should have an extension that maps $\text{acl}_{\mathcal{N}}(B)$ into $\text{acl}_{\mathcal{M}}(A)$ and is elementary with respect to \mathcal{N}, \mathcal{M} . But since the range of g^{-1} is all of $\text{acl}_{\mathcal{M}}(A)$ and since the extension, being an elementary function, must be 1-1, this is clearly impossible. This contradiction proves that the range of g must be $\text{acl}_{\mathcal{N}}(B)$, as claimed. \square

From the previous result we can derive the fact that model theoretic algebraic closure is to a large extent independent of the model within which it is computed.

10.6. Corollary. *Suppose \mathcal{M}, \mathcal{N} are L -structures with $\mathcal{M} \preceq \mathcal{N}$. Then $\text{acl}_{\mathcal{M}}(A) = \text{acl}_{\mathcal{N}}(A)$ for every $A \subseteq M$.*

Proof. Apply Proposition 10.5 to the identity map on A . \square

EXERCISES

10.7. Let M be an infinite set, considered as a structure for the language of pure equality. For each $A \subseteq M$, show that $\text{acl}_M(A) = A$.

10.8. Let $\mathcal{M} \models DLO$. For each $A \subseteq M$, show that $\text{acl}_{\mathcal{M}}(A) = A$.

10.9. Let K be a field and let L be the language of vector spaces over K . (See Exercises 4.20 and 7.19.) For each infinite K -vector space V (considered as an L -structure) and each $A \subseteq V$, show that $\text{acl}_V(A)$ is the K -linear subspace of V spanned by A .

10.10. Consider the theory T_{dis} of discrete linear orderings without endpoints. (See Example 7.13.) For $\mathcal{M} \models T_{dis}$ and $A \subseteq M$, describe $\text{acl}_{\mathcal{M}}(A)$.

11. ALGEBRAIC CLOSURE IN MINIMAL STRUCTURES

Throughout this chapter let \mathcal{M} denote an infinite minimal L -structure. We will write $\text{cl}(A)$ in place of $\text{acl}_{\mathcal{M}}(A)$ for $A \subseteq M$.

From Proposition 10.3 we know that cl is a closure operation of finite character on the subsets of M . When \mathcal{M} is minimal, cl is actually a *pregeometry*; this means that cl also satisfies the *Exchange Property*:

11.1. Proposition. *Let \mathcal{M} be an infinite minimal structure. Let $A \subseteq M$ and $a, b \in M$. If $a \notin \text{cl}(A)$ and $b \notin \text{cl}(A)$, then*

$$a \in \text{cl}(A \cup \{b\}) \iff b \in \text{cl}(A \cup \{a\}).$$

Proof. We argue by contradiction. Suppose $a, b \notin \text{cl}(A)$, $a \in \text{cl}(A \cup \{b\})$, and $b \notin \text{cl}(A \cup \{a\})$. Let the formula $\varphi(x, y, z_1, \dots, z_p)$ and the parameters $e_1, \dots, e_p \in X$ witness the fact that $a \in \text{cl}(A \cup \{b\})$ (where b is included as a parameter to be substituted for the variable y). Let r be the cardinality of the finite set $\{c \in M \mid \mathcal{M} \models \varphi[c, b, e_1, \dots, e_p]\}$, which contains a as an element. Let $\psi(y, z_1, \dots, z_p)$ be a formula expressing that there are at most r values of x for which $\varphi(x, y, z_1, \dots, z_p)$ is true. Note that $\mathcal{M} \models \psi[b, e_1, \dots, e_p]$. Since $b \notin \text{cl}(A \cup \{a\})$, the set

$$\{b' \in M \mid \mathcal{M} \models \varphi[a, b', e_1, \dots, e_p] \text{ and } \mathcal{M} \models \psi[b', e_1, \dots, e_p]\}$$

must be infinite; since \mathcal{M} is minimal this set must be cofinite in M . Let s be the number of elements of M that are not in this set.

Now consider a formula $\sigma(x, z_1, \dots, z_p)$ that expresses the statement that $\varphi(x, y, z_1, \dots, z_p) \wedge \psi(y, z_1, \dots, z_p)$ holds for all but s many values of y . The set

$$\{c \in M \mid \mathcal{M} \models \sigma[c, e_1, \dots, e_p]\}$$

has a as an element; since $a \notin \text{cl}(A)$ and \mathcal{M} is minimal, this set must be cofinite. Let a_0, \dots, a_r be distinct elements of this set. For each $j = 0, \dots, r$ we have that the set

$$\{b' \in M \mid \mathcal{M} \models \varphi[a_j, b', e_1, \dots, e_p] \text{ and } \mathcal{M} \models \psi[b', e_1, \dots, e_p]\}$$

must be cofinite in M , which is infinite. Therefore the intersection of these sets is also cofinite, hence nonempty. That is, there must exist a single $b' \in M$ such that for each $j = 0, \dots, r$ we have

$$\mathcal{M} \models \varphi[a_j, b', e_1, \dots, e_p] \text{ and } \mathcal{M} \models \psi[b', e_1, \dots, e_p]$$

which is a contradiction. □

11.2. Definition. Let cl be a pregeometry on the set M ; let $A, B \subseteq M$.

- (1) A is *closed* if $\text{cl}(A) = A$.
- (2) $\text{cl}(A)$ is the *closure* of A .
- (3) (B closed) A *spans* B if $\text{cl}(A) = B$.
- (4) A is *independent* if $a \notin \text{cl}(A \setminus \{a\})$ for all $a \in A$. Otherwise A is *dependent*.
- (5) (B closed) A is a *basis* for B if A is independent and A spans B .

The next result gives the most important facts about independent sets and spanning sets in a pregeometry.

11.3. Theorem. *Let cl be a pregeometry on the set M ; let $A, B \subseteq M$ with B closed.*

- (1) *A is independent if and only if each finite subset of A is independent.*
- (2) *A is a basis for B if and only if A is maximal among independent subsets of B . Consequently every closed set has a basis. Indeed, every independent subset of B is contained in a basis for B .*
- (3) *If A spans B , then there exists $C \subseteq A$ such that C is a basis of B .*
- (4) *A is a basis for B if and only if A is minimal among subsets of B that span B .*
- (5) *Suppose A is a basis for B and $a \in B$. Then there is a smallest finite set $F \subseteq A$ such that $a \in \text{cl}(F)$. We will call F the support of a in A .*
- (6) *Any two bases for B have the same cardinality.*

Proof. (1) Suppose A is independent and let C be any subset of A . For each $a \in C$ we have $C \setminus \{a\} \subseteq A \setminus \{a\}$ and therefore $\text{cl}(C \setminus \{a\}) \subseteq \text{cl}(A \setminus \{a\})$. Since A is independent this implies $a \notin \text{cl}(C \setminus \{a\})$. Therefore C is independent. In particular every finite subset of A is independent. Conversely, suppose A is dependent, so there exists $a \in A$ such that $a \in \text{cl}(A \setminus \{a\})$. Therefore there is a finite subset C of $A \setminus \{a\}$ such that $a \in \text{cl}(C)$. It follows that $C \cup \{a\}$ is a dependent finite subset of A .

(2) Suppose A is a basis for B . For each $a \in B \setminus A$ we have $a \in \text{cl}(A)$, from which it follows that $A \cup \{a\}$ is dependent. It follows that A is maximal among independent subsets of B . Conversely, suppose A is maximal among independent subsets of B . Then for each $a \in B \setminus A$ the set $A \cup \{a\}$ is dependent. If $a \notin \text{cl}(A)$ then there exists $b \in A$ with $b \in \text{cl}((A \cup \{a\}) \setminus \{b\})$. Since A is independent we have $b \notin \text{cl}(A \setminus \{b\})$. The Exchange Property implies $a \in \text{cl}((A \setminus \{b\}) \cup \{b\}) = \text{cl}(A)$. This contradiction proves $a \in \text{cl}(A)$. Since $a \in B$ was arbitrary, this proves that A spans B and therefore A is a basis for B .

Suppose A is any independent subset of B . Let Ω be the collection of all independent subsets of B that contain A . Part (1) of this Theorem implies that if \mathcal{C} is any subset of Ω that is a chain under \subseteq , then $\cup \mathcal{C}$ is independent and thus is a member of Ω . Zorn's Lemma implies the existence of maximal elements of Ω under \subseteq . Any such set is a basis of B , by what was proved in the preceding paragraph.

(3) The proof is similar to the second part of (2). Given A spanning B , let Ω be the collection of all independent subsets of A . By Zorn's Lemma and (1) there exists $C \in \Omega$ that is maximal with respect to \subseteq . By the argument in the previous paragraph, $\text{cl}(C) = \text{cl}(A)$ and therefore C is an independent set spanning B . By (2), C is a basis for B .

(4) Suppose A is a basis for B and C is a proper subset of A . For each $a \in A \setminus C$ we have $a \notin \text{cl}(A \setminus \{a\}) \supseteq \text{cl}(C)$, which shows that C does

not span B . Conversely suppose A is minimal among sets that span B . We must show A is independent. Otherwise there exists $a \in A$ such that $a \in \text{cl}(A \setminus \{a\})$. It follows that $\text{cl}(A \setminus \{a\}) = \text{cl}(A)$, contradicting the assumption that A is a minimal spanning set.

(5) Since cl has finite character we know there exists a finite $F \subseteq A$ with $a \in \text{cl}(F)$. Let F be such a set of smallest cardinality. We will show $a \notin \text{cl}(A \setminus \{b\})$ for each $b \in F$. It follows that F must be contained in any subset A of A that satisfies $a \in \text{cl}(A)$. If $b \in F$ then we have $a \notin \text{cl}(F \setminus \{b\})$ by the minimality of F . The Exchange Property implies $b \in \text{cl}((F \setminus \{b\}) \cup \{a\})$. Since $b \notin \text{cl}(A \setminus \{b\})$ we see it is impossible for a to be in $\text{cl}(A \setminus \{b\})$.

(6) Let U and V be bases for B . The case where one of the bases is infinite can be proved using a simple counting argument based on the finite character of cl . Suppose V is infinite and $\text{card}(U) \leq \text{card}(V)$. For each $a \in U$ there exists a finite set $F(a) \subseteq V$ such that $a \in \text{cl}(F(a))$. Let $F = \cup\{F(a) \mid a \in U\}$. Evidently F spans B , and since V is a basis for B it follows from (3) that $F = V$. Since V is infinite it follows that U is also infinite and indeed that $\text{card}(V) = \text{card}(F) \leq \text{card}(U)$. Hence $\text{card}(U) = \text{card}(V)$.

Now we handle the finite case. Let U be a finite basis for B and let V be any independent subset of B . By the last sentence of (2) and what is proved in the previous paragraph, V must be finite. We will show $\text{card}(V) \leq \text{card}(U)$. To do this we prove the following by induction on the cardinality of V :
there exists $W \subseteq U$ such that $W \cup V$ is a basis for B , $W \cap V = \emptyset$, and $\text{card}(W \cup V) = \text{card}(U)$.

As basis step we consider the case $\text{card}(V) = 0$. Evidently we may take $W = U$ when $V = \emptyset$.

For the induction step, consider an independent set $V \subseteq B$ and suppose the statement is true for all independent sets that are strictly smaller than V . Fix $a \in V$ and let $Z = V \setminus \{a\}$. By the induction hypothesis applied to Z , there exists $W \subseteq U$ such that $W \cup Z$ is a basis for B , $W \cap Z = \emptyset$, and $\text{card}(W \cup Z) = \text{card}(U)$. Let A be the support of a in $W \cup Z$ (see (5)). Since $Z \cup \{a\} = V$ is independent, so $a \notin \text{cl}(Z)$, we cannot have $A \subseteq Z$, and thus A meets W . Let b be any element of $A \cap W$.

We complete the induction step by showing that $W \setminus \{b\}$ is the desired subset of U for $V = Z \cup \{a\}$. This requires us to prove:

- (i) $(W \setminus \{b\}) \cup (Z \cup \{a\})$ is independent;
- (ii) $(W \setminus \{b\}) \cup (Z \cup \{a\})$ spans B ;
- (iii) $(W \setminus \{b\}) \cap (Z \cup \{a\}) = \emptyset$; and
- (iv) $\text{card}((W \setminus \{b\}) \cup (Z \cup \{a\})) = \text{card}(W \cup Z)$, which equals $\text{card}(U)$.

To begin proving these items, we note

$$(\#) a \notin \text{cl}((W \setminus \{b\}) \cup Z);$$

otherwise, the support of a in $W \cup Z$ would be contained in $(W \setminus \{b\}) \cup Z$ (see (5)), but this set does not contain b , whereas the support does contain

it. Since $(W \setminus \{b\}) \cup Z \subseteq W \cup Z$ is independent, $(\#)$ implies (i). Moreover, we know $a \in B = \text{cl}(W \cup Z)$, so $(\#)$ and the Exchange Principle yield $b \in \text{cl}((W \setminus \{b\}) \cup (Z \cup \{a\}))$, from which we get (ii).

The only way (iii) could fail is for $a \in (W \setminus \{b\}) \cup Z$ to hold; this would imply $A = \{a\}$, from which it follows that $b = a$ and hence also $a \in Z$, a contradiction.

Finally, we note that the map taking b to a and being the identity on the rest of $W \cup Z$ is a bijection from $W \cup Z$ onto $(W \setminus \{b\}) \cup (Z \cup \{a\})$. This verifies (iv) and completes the induction proof.

To finish the proof of (6), simply note that the induction argument above shows that if U, V are two finite bases of B , then $\text{card}(V) \leq \text{card}(U)$ and $\text{card}(U) \leq \text{card}(V)$. \square

11.4. Definition. Let cl be a pregeometry on the set M and $A \subseteq M$. The *rank of A with respect to cl in M* , denoted by $\text{rank}(A)$, is the unique cardinality of a basis for the closed set $\text{cl}(A)$.

11.5. Definition. Let \mathcal{M} be a minimal L -structure. The *rank of \mathcal{M}* , denoted $\text{rank}(\mathcal{M})$, is the rank of the set M with respect to the pregeometry $\text{acl}_{\mathcal{M}}$.

11.6. Proposition. *Let \mathcal{M} and \mathcal{N} be L -structures with $\mathcal{M} \equiv \mathcal{N}$ and \mathcal{M} minimal. Suppose $A \subseteq M$ and $B \subseteq N$, and let $f: A \rightarrow B$ be a function that is elementary with respect to \mathcal{M}, \mathcal{N} . For each $a \in M \setminus \text{acl}_{\mathcal{M}}(A)$ and each $b \in N \setminus \text{acl}_{\mathcal{N}}(B)$ the extension of f that takes a to b is also elementary with respect to \mathcal{M}, \mathcal{N} .*

Proof. Otherwise there exists an L -formula $\varphi(x, y_1, \dots, y_n)$ and parameters $e_1, \dots, e_n \in A$ such that $\mathcal{M} \models \varphi[a, e_1, \dots, e_n]$ and $\mathcal{N} \models \neg\varphi[b, f(e_1), \dots, f(e_n)]$. Since a is not algebraic over A and b is not algebraic over B , the sets $\{c \in M \mid \mathcal{M} \models \varphi[c, e_1, \dots, e_n]\}$ and $\{d \in N \mid \mathcal{N} \models \neg\varphi[d, f(e_1), \dots, f(e_n)]\}$ are both infinite. Since f is elementary it follows that $\{c \in M \mid \mathcal{M} \models \neg\varphi[c, e_1, \dots, e_n]\}$ is infinite. This contradicts the assumption that \mathcal{M} is minimal. \square

11.7. Corollary. *Let \mathcal{M} and \mathcal{N} be L -structures with $\mathcal{M} \equiv \mathcal{N}$ and \mathcal{M} minimal. Suppose we have independent sets $A \subseteq M$ (with respect to $\text{acl}_{\mathcal{M}}$) and $B \subseteq N$ (with respect to $\text{acl}_{\mathcal{N}}$), and let $f: A \rightarrow B$ be a bijection. Then f is elementary with respect to \mathcal{M}, \mathcal{N} .*

Proof. Let Ω be the collection of subsets $C \subseteq A$ such that the restriction of f to C is elementary with respect to \mathcal{M}, \mathcal{N} . We regard \emptyset as an element of Ω (justified since $\mathcal{M} \equiv \mathcal{N}$). The partially ordered set (Ω, \subseteq) satisfies the hypothesis of Zorn's Lemma, so there exists $C \in \Omega$ that is maximal with respect to \subseteq . We need to show $C = A$. If not, let a be any element of $A \setminus C$ and let $b = f(a) \in B \setminus f(C)$. By Proposition 11.6, f restricted to $C \cup \{a\}$ is elementary with respect to \mathcal{M}, \mathcal{N} . This contradicts the maximality of C and proves $C = A$. \square

11.8. Theorem. *Let T be a complete strongly minimal theory, and let \mathcal{M}, \mathcal{N} be infinite models of T .*

(1) There is an elementary embedding of \mathcal{M} into \mathcal{N} if and only if $\text{rank}(\mathcal{M}) \leq \text{rank}(\mathcal{N})$.

(2) \mathcal{M} and \mathcal{N} are isomorphic if and only if $\text{rank}(\mathcal{M}) = \text{rank}(\mathcal{N})$.

(3) T is κ -categorical for every cardinal number $\kappa > \text{card}(L)$.

Proof. (1 \Leftarrow) Since acl defines a pregeometry in each of these structures, there exist bases A for M (with respect to $\text{acl}_{\mathcal{M}}$) and B for N (with respect to $\text{acl}_{\mathcal{N}}$). By hypothesis $\text{card}(A) \leq \text{card}(B)$ so there is a 1-1 function f from A into B . The preceding Corollary yields that f is elementary with respect to \mathcal{M}, \mathcal{N} . By Proposition 10.5, f can be extended to a function $g: M \rightarrow N$ that is elementary with respect to \mathcal{M}, \mathcal{N} . It follows easily that g is an elementary embedding from \mathcal{M} into \mathcal{N} .

(1 \Rightarrow) As above, there is a basis A for M (with respect to $\text{acl}_{\mathcal{M}}$). If f is an elementary embedding of \mathcal{M} into \mathcal{N} , then $f(A)$ is independent with respect to $\text{acl}_{\mathcal{N}}$. By Theorem 11.3(2) there is a basis B for \mathcal{N} that contains $f(A)$. It follows that $\text{card}(A) \leq \text{card}(B)$ and hence $\text{rank}(\mathcal{M}) \leq \text{rank}(\mathcal{N})$.

(2) The argument is similar to (1).

(3) By the Löwenheim-Skolem Theorems (Theorems 6.1 and 6.5) there exist models of T having cardinality κ . Let \mathcal{M} and \mathcal{N} be two such models of T . As in the proof of (1), let A be a basis for M (with respect to $\text{acl}_{\mathcal{M}}$) and B for N (with respect to $\text{acl}_{\mathcal{N}}$). Because acl is of finite character in each model, and the number of L -formulas is $< \kappa$, a counting argument shows that A and B must each be of cardinality equal to κ . Now use part (2). \square

11.9. Example. Let K be a field and let L be the language of vector spaces over K . Let Σ be the set of L -sentences whose models are the infinite vector spaces over K . (See Exercises 4.20 and 7.19.)

- The set of axioms Σ is strongly minimal.

It follows that Chapter 10 applies to infinite K -vector spaces. Exercise 10.9 shows that algebraic closure in the sense of model theory and linear span in the sense of linear algebra are identical, when applied to subsets of a fixed infinite vector space over K .

- Let V, W be infinite K -vector spaces and let $X \subseteq V, Y \subseteq W$ be K -linear subspaces. Suppose $F: X \rightarrow Y$ is a K -linear isomorphism. Then F is an elementary map in the sense of the L -structures V, W .

- If V is an infinite K -vector space and $X \subseteq V$ is a K -linear subspace, then the model theoretic rank of X in the sense of algebraic closure in V does not depend on V . Moreover, this rank is the same as the dimension of X in the sense of linear algebra.

- Theorem 11.3 implies all of the standard facts about linearly independent sets, spanning sets, and bases, for arbitrary vector spaces over K .

11.10. Example. Let K be an algebraically closed field, considered as an L_r -structure, and $A \subseteq K$; let k be the subfield of K generated by A . For

each $a \in K$ we have that $a \in \text{acl}_K(A)$ iff there is a nonconstant polynomial $p(x)$ with coefficients in k such that $p(a) = 0$ in K . (This follows from the fact that ACF has QE; see Theorem 8.1.) In other words, the concept *algebraic closure* has the same meaning whether we interpret it model theoretically or algebraically, when we are working in an algebraically closed field. In particular, A is independent with respect to the closure operator acl_K if and only if A is algebraically independent in the field K ; further, A is a basis for K in the sense of acl_K if and only if A is a transcendence basis for K . Hence the results in this Chapter yield an immediate proof of Steinitz's Theorem: Every algebraically closed field K has a transcendence basis B , and K is determined up to isomorphism by $\text{card}(B)$. In particular, ACF_0 and each ACF_p are κ -categorical for each uncountable cardinal κ .

EXERCISES

11.11. Prove the statements made in Examples 11.9 and 11.10.

11.12. Let L be the language of pure equality and let T be the theory in L of all infinite sets. From Example 4.15 we know that T admits QE and is complete.

- Show that T is strongly minimal.
- Explain the meaning of the rank of a given model of T , in the sense of Definition 11.5.

11.13. Let L be the language whose nonlogical symbols consist of a unary function symbol F . Let T be the theory in L of the class of all L -structures (A, f) in which f is a bijection from A onto itself and f has no finite cycles. Note that (\mathbb{Z}, S) is a model of T , where $S(a) = a + 1$ for all $a \in \mathbb{Z}$.

- Show that T admits QE and is complete; therefore $T = \text{Th}(\mathbb{Z}, S)$.
- Show that T is strongly minimal.
- Explain the meaning of the rank of a given model of T , in the sense of Definition 11.5.

11.14. Let \mathcal{M} be an infinite strongly minimal L -structure and let κ be an infinite cardinal.

- Show that \mathcal{M} is κ -saturated iff the rank of \mathcal{M} is $\geq \kappa$. (See 11.5.)

11.15. Let L be the language whose only nonlogical symbol is a binary predicate symbol $<$. Let \mathcal{M} be any infinite linear ordering, considered as an L -structure.

- Show that \mathcal{M} is not strongly minimal.

11.16. Let L be the language whose only nonlogical symbol is a binary predicate symbol $<$. Proceeding as in Example 7.13, show that the L -structure $(\mathbb{N}, <)$ is minimal. (By the preceding exercise, it is not strongly minimal.)

12. REAL CLOSED FIELDS

Note: the final version of this chapter will include some development of the basic theory of ordered fields, including treatment of real closed fields. In the most recent versions of Math 571, Lou used a note on this material that he had written awhile ago. Lou will be rewriting this chapter

We now consider the class of *real closed ordered fields*; these are ordered fields that satisfy an appropriate notion of being algebraically closed. We regard them as structures for the language L_{or} of ordered rings, which is the expansion of L_r by adding a binary relation symbol \leq . This class is axiomatizable, by a set of L_{or} -sentences that we denote $RCOF$, which consists of the axioms for ordered fields together with the sentences:

(1) $\forall y(y \geq 0 \rightarrow \exists x(y = x^2))$;

and for each odd $n \geq 1$

(2_n) $\forall y_1 \dots \forall y_n \exists x(x^n + y_1 x^{n-1} + \dots + y_{n-1} x + y_n = 0)$.

These sentences assert that every non-negative element has a square root, and that every monic polynomial of odd degree has a root.

Note that a real closed field has a unique ordering that is compatible with the field structure (since squares must be ≥ 0).

The best known real closed ordered field is the ordered field \mathbb{R} of real numbers. (This is easy to prove using facts from calculus applied to polynomials.) Evidently, the subfield of all real algebraic numbers is also a model of $RCOF$.

In 1926-27 Artin and Schreier developed the theory of ordered fields and proved that every ordered field K has a *real closure* (by which we mean a real closed ordered field that is an algebraic extension of K .) Moreover, any two real closures of an ordered field K are isomorphic over K . The uniqueness of the real closure will play a role in our analysis of $RCOF$.

Further, it can be proved that if K is a real closed ordered field, then the irreducible monic polynomials in $K[x]$ are exactly the polynomials $x - a$ for $a \in K$ together with the polynomials $x^2 + ax + b$ such that $a, b \in K$ satisfy $4b - a^2 > 0$. We will also use this fact without including a proof.

A full discussion of real closed ordered fields may be found in Serge Lang's book *Algebra*.

12.1. Fact. The L_{or} -substructures of real closed ordered fields are the ordered domains. In particular, the ordering on an ordered domain D can be extended to its field of fractions F by making the fraction $a/b > 0$ iff $a, b > 0$ or $a, b < 0$. Then the real closure of F is a real closed ordered field that contains D as an ordered subring.

We are going to show $RCOF$ has QE using Theorem 7.11, and to do that we first need the following Lemma:

12.2. Lemma. *Let $F \subseteq K$ be real closed ordered fields and suppose b is an element of $K \setminus F$. The isomorphism type of b over F (in the language L_{or}) is determined by the set of elements $f \in F$ such that $f < b$.*

Proof. The uniqueness of the real closure implies that F is algebraically closed in K .

Suppose K' is another real closed ordered field extension of F and b' is in $K' \setminus F$. Suppose further that for all $f \in F$ we have $f < b \iff f < b'$. Consider the map g defined on $F[b]$ by taking $g(b) = b'$ and $g(f) = f$ for all $f \in F$. Since b and b' are both transcendental over F , this is a ring isomorphism from $F[b]$ onto $F[b']$. We need to show that g is order preserving. That is, for any polynomial $p(x)$ with coefficients in F we have to prove

$$(\star) \quad p(b) > 0 \iff p(b') > 0.$$

This condition is trivial when $p(x)$ is constant, so we may assume that $p(x)$ is nonconstant. Without loss of generality we may assume that $p(x)$ is monic since the equivalence (\star) is preserved under multiplication by nonzero elements of F . Moreover, we may assume that $p(x)$ is irreducible in $F[x]$, since the product of any finitely many polynomials that satisfy (\star) will again satisfy (\star) . If $p(x)$ has degree 1, so it is of the form $x - f$ for some $f \in F$, then condition (\star) just says $b > f \iff b' > f$, which we were given to be true. Thus we may assume that $p(x)$ is of degree > 1 . Since it is irreducible, it has no roots in F . As noted above, this implies that $p(x) = x^2 + fx + g$ for some $f, g \in F$ that satisfy $4g - f^2 > 0$. Completing the square shows that $p(b) > 0$ in K and $p(b') > 0$ in K' , so condition (\star) holds for $p(x)$, as needed. \square

12.3. Theorem (Tarski). *RCOF has QE.*

Proof. We apply Theorem 7.11, verifying condition (2) of that result. Therefore we need to consider real closed ordered fields F, K as well as a substructure F_0 of F and an ordered ring embedding h of F_0 into K . Let K_0 be the range of h . Given any element a of F , we must show that h can be extended to an embedding of $F_0[a]$ into some elementary extension of K .

Let κ be a cardinal that satisfies $\kappa > \text{card}(F_0)$ and $\kappa > \omega = \text{card}(L_{or})$. Let K' be a κ -saturated elementary extension of K .

We know that F_0 is an ordered subring of F . Since the field of fractions of F_0 is uniquely determined as an ordered field over F_0 , we can extend the embedding h to be defined on the field generated in F by F_0 . Therefore we may assume that F_0 is already an ordered subfield of F . A similar argument using the uniqueness of the real closure of an ordered field shows that we may also assume that F_0 is itself a real closed ordered field; in particular, we may assume that F_0 is algebraically closed in F .

Now we have an element a of $F \setminus F_0$. Note that if $b_1, \dots, b_m, c_1, \dots, c_n$ are finitely many elements of F_0 and they satisfy $b_i < a < c_j$ for all i, j , then we can find $a' \in K \subseteq K'$ such that $f(b_i) < a' < f(c_j)$ for all i, j . Indeed, we may take $a' = \frac{1}{2}(f(u) + f(v))$ where $u = \max(b_1, \dots, b_m)$ and $v = \min(c_1, \dots, c_n)$. Since K' is κ -saturated and $\text{card}(F_0) < \kappa$, it is therefore possible to find $a' \in K'$ that satisfies $f < a \Leftrightarrow h(f) < a'$ for all $f \in F_0$. Using Lemma 12.2 we conclude that h can be extended to an embedding of ordered domains from $F_0[a]$ onto $K_0[a']$ by setting $h(p(a)) = p(a')$ for all polynomials $p(x) \in F_0[x]$. This completes the proof. \square

12.4. Corollary. *RCOF is complete; hence RCOF axiomatizes the L_{or} -theory of the ordered field \mathbb{R} of real numbers.*

Proof. Every ordered field has characteristic 0 and therefore contains an isomorphic copy of the ordered ring \mathbb{Z} . The completeness of *RCOF* follows from Fact 7.3(3). \square

12.5. Remark. The theory of the ordered field \mathbb{R} is decidable; this follows immediately from the fact that it is axiomatized by *RCOF*, so it is a complete theory for which one has a computable set of axioms. This decidability result was a large part of Tarski's original motivation for proving that *RCOF* admits quantifier elimination. It is of some interest to computer scientists, because instances of certain problems in areas such as robotics can be formulated as sentences in the language of ordered rings, and the "feasibility" of a given problem instance corresponds to the truth of the sentence in \mathbb{R} . For this reason some computer scientists have tried to find efficient algorithms for deciding $\text{Th}(\mathbb{R})$ and have implemented these algorithms in software systems. However, the systems do not perform very well, and it has been shown that the computational complexity of $\text{Th}(\mathbb{R})$ is sufficiently high that no feasible algorithm for deciding it can exist. Current interest emphasizes subproblems that are defined by syntactic restrictions.

12.6. Fact. Let K be a real closed field, considered as an L_{or} -structure. Let X be any subset of K^n that is definable in K (allowing parameters from K).

- (a) If $n = 1$, then X is the finite union of points from K and open intervals whose endpoints are in $K \cup \{-\infty, +\infty\}$.
- (b) The closure of X and the interior of X are also definable subsets of K^n , where K is given the topology defined using its ordering.

12.7. Remark. Statement (a) of the preceding Fact is expressed by saying that *RCOF* is *o-minimal*. The study of *o-minimal* structures is an active area of research today.

Sets $X \subseteq \mathbb{R}^n$ definable in the ordered field \mathbb{R} by quantifier-free formulas (allowing parameters from \mathbb{R}) are called *semi-algebraic*. Tarski's Theorem (that *RCOF* has QE) yields that the system of semi-algebraic sets is closed under projections, and, more generally, under polynomial maps. Statement

(b) in the preceding Fact shows that it is closed under the operations of forming the closure and the interior.

Artin and Schreier developed the theory of real closed ordered fields, in part toward solving Hilbert's 17th Problem. This problem asked for a characterization of positive definite rational functions with coefficients in the real numbers or, more generally, in a given ordered field. As our last result we give a model theoretic proof of the solution to this problem in the case where the ordered field is a real closed ordered field. For a more general discussion see Abraham Robinson's book *Model Theory*, for example, or the article by Angus Macintyre in *The Handbook of Mathematical Logic*.

12.8. Corollary. *Let F be a real closed ordered field, and let p, q be polynomials in the variables x_1, \dots, x_n with coefficients in F . Suppose that the rational function $f = p/q$ is positive semi-definite, in the sense that for any $a \in F^n$ with $q(a) \neq 0$, one has $f(a) = p(a)/q(a) \geq 0$. Then f is equal to a sum of finitely many squares of rational functions in the field of rational functions $F(x_1, \dots, x_n)$.*

Proof. Assume that $f = p/q$ is positive semi-definite. If f is not a sum of squares in the field $F(x_1, \dots, x_n)$ then this field has an ordering in which the element f is negative. To show this, use Zorn's Lemma and take P to be a maximal subset of $F(x_1, \dots, x_n)$ that contains $-f$ and all nonzero squares, does not contain 0, and is closed under $+$ and \times . The desired linear ordering on the field $F(x_1, \dots, x_n)$ is defined by taking

$$g < h \iff (h - g) \in P.$$

This ordering on $F(x_1, \dots, x_n)$ obviously extends the original ordering on F . Let K be a real closed ordered field that extends $F(x_1, \dots, x_n)$ with this ordering. Now consider the polynomials $p(x_1, \dots, x_n)$ and $q(x_1, \dots, x_n)$ as terms in the language $L_{or, F}$. We see that the sentence $\exists x_1 \dots \exists x_n (p(x_1, \dots, x_n)q(x_1, \dots, x_n) < 0)$ is true in K_F . (Note that $q \neq 0 \wedge p/q < 0$ is equivalent to $pq < 0$ in ordered fields.) By Tarski's Theorem, this sentence is equivalent in $RCOF$ to a quantifier free sentence, so that it is also true in F_F . But this sentence is false in F_F by hypothesis, contradicting the assumption that f is not a sum of squares. \square

EXERCISES

12.9. Let K be a countable ordered field, considered as an L_{or} -structure, and let $T = \text{Th}(K)$. Show that there exists a 1-type $p \in S_1(T)$ that is not realized in K . Therefore, no countable ordered field is ω -saturated.

12.10. Let R be an ordered field. Let x be a transcendental element over R and consider the field $R(x)$ of rational functions in x with coefficients in R .

- Show that there is linear ordering $<$ on $R(x)$ that makes $R(x)$ into an ordered field, such that $r < x$ for all $r \in R$.

- Show that this ordering is unique.
- Show how to embed the ordered field $R(x)$ with this ordering into a suitable ultrapower of R .
- Describe all the embeddings of the field $R(x)$ into an ultrapower of R . (Each one induces a field ordering on $R(x)$.)

12.11. Use the preceding Exercise and results in this Chapter to show that $RCOF$ is not κ -categorical for any infinite cardinal κ . (For example, construct models of $RCOF$ of cardinality κ , such that one has an ordering of cofinality ω and the other has an ordering of uncountable cofinality.)

13. HOMOGENEITY

In this chapter we construct models that are not only highly saturated but are also homogeneous in a strong sense. Such models play the same role in the setting of general mathematical structures that uncountable algebraically closed fields play in algebra and number theory. That is, they contain “ideal elements” with which one can directly calculate and they support useful functions and relations; thus they provide a convenient framework for certain mathematical arguments. Use of such rich models is a key feature of modern model theory.

13.1. Definition. Let L be a first order language, \mathcal{M} an L -structure, and κ an infinite cardinal number. We say \mathcal{M} is *strongly κ -homogeneous* if it has the following property for every subset A of M of cardinality $< \kappa$: any map from A into M that is elementary with respect to \mathcal{M} can be extended to an automorphism of \mathcal{M} .

We construct a strongly homogeneous model as the union of a well ordered elementary chain. The next result is needed at the successor stage when we are defining this elementary chain by induction.

13.2. Lemma. *Suppose \mathcal{M} is κ -saturated and $\mathcal{N} \preceq \mathcal{M}$ satisfies $\text{card}(\mathcal{N}) < \kappa$. Then any elementary map f between subsets of \mathcal{N} can be extended to an elementary embedding of \mathcal{N} into \mathcal{M} .*

Proof. Suppose the domain of the elementary mapping f is A . Then we have $(\mathcal{N}, a)_{a \in A} \equiv (\mathcal{M}, f(a))_{a \in A}$. Moreover, it is easy to see that $(\mathcal{M}, f(a))_{a \in A}$ must also be κ -saturated, since $\text{card}(A) < \kappa$. By Corollary 5.5 there is an elementary embedding of $(\mathcal{N}, a)_{a \in A}$ into $(\mathcal{M}, f(a))_{a \in A}$. This yields an elementary embedding of \mathcal{N} into \mathcal{M} that extends f . \square

13.3. Theorem (Existence of Strongly Homogeneous Models). *For every infinite cardinal κ , every structure has a κ -saturated elementary extension \mathcal{M} such that every reduct of \mathcal{M} is strongly κ -homogeneous.*

Proof. Let \mathcal{M}_0 be any structure and κ an infinite cardinal. Let $\tau = \kappa^+$. Using induction over the well ordered set $\{\alpha \mid \alpha < \tau\}$ we construct an elementary chain of structures $(\mathcal{M}_\alpha \mid \alpha < \tau)$ such that $\mathcal{M}_{\alpha+1}$ is $\text{card}(M_\alpha)^+$ -saturated for every $\alpha < \tau$. To construct this sequence, at each successor stage (α to $\alpha + 1$) we apply Theorem 5.8 to \mathcal{M}_α ; at limit stages we take the union of the previously defined structures. Finally, the desired elementary extension \mathcal{M} of \mathcal{M}_0 is obtained by setting $\mathcal{M} = \cup\{\mathcal{M}_\alpha \mid \alpha < \tau\}$.

Note that any subset A of M that has cardinality $< \tau$ must be a subset of M_α for some $\alpha < \tau$. (Here we use the fact that $\tau = \kappa^+$ is a regular cardinal.) From this it is immediate that \mathcal{M} is τ -saturated (as in the proof of Theorem 5.8).

It remains to show that every reduct of \mathcal{M} is strongly κ -homogeneous. Let L be any sublanguage of the language of \mathcal{M} . Note that the chain of reducts

$(\mathcal{M}_\alpha|L \mid \alpha < \tau)$ is an elementary chain such that $\mathcal{M}_{\alpha+1}|L$ is $\text{card}(M_\alpha)^+$ -saturated for every $\alpha < \tau$. (See Theorem 5.4.) Moreover, $\mathcal{M}|L$ is the union of this chain.

Let f be any mapping between subsets of M that is elementary with respect to $\mathcal{M}|L$, such that the domain and range of f have cardinality $< \kappa$. As noted above, the domain and range of f are both contained in M_α for some $\alpha < \tau$. Moreover, such a mapping f is elementary with respect to $\mathcal{M}_\alpha|L$, since $\mathcal{M}_\alpha|L \preceq \mathcal{M}|L$. Without loss of generality we may assume that α is a limit ordinal.

An ordinal β can be written in a unique way as $\beta = \lambda + n$ for some limit ordinal λ and some integer $n \in \mathbb{N}$. We call β *odd* or *even* according to whether the integer n is odd or even. Note that each limit ordinal is even.

Applying Lemma 13.2 to $\mathcal{M}_\alpha|L$ and f we obtain an elementary embedding f_α from $\mathcal{M}_\alpha|L$ into $\mathcal{M}_{\alpha+1}|L$ that extends f . We proceed by induction to obtain a sequence of elementary embeddings f_β from $\mathcal{M}_\beta|L$ into $\mathcal{M}_{\beta+1}|L$, for β in the interval $\alpha \leq \beta < \tau$, such that $f_{\beta+1}$ is always an extension of f_β^{-1} . It follows that $f_{\beta+2}$ is an extension of f_β for all $\alpha \leq \beta < \tau$. At successor ordinals the mapping $f_{\beta+1}$ is obtained by applying Lemma 13.2 to $\mathcal{M}_{\beta+1}|L$ and f_β^{-1} . At limit ordinals λ the induction construction is continued by first taking g to be the union of all the elementary mappings f_β such that $\beta < \lambda$ and β is even, and then applying Lemma 13.2 to extend g to an elementary embedding f_λ of $\mathcal{M}_\lambda|L$ into $\mathcal{M}_{\lambda+1}|L$. Finally, let h be the union of the mappings f_β such that $\alpha \leq \beta < \tau$ and β is even. It is easy to show that h is an automorphism of $\mathcal{M}|L$ and that it extends the original elementary mapping f . \square

The strongly homogeneous models constructed in the proof of Theorem 13.3 are very large. In some situations it is useful to control the cardinality of strongly homogeneous models, as we do in the next result.

13.4. Theorem (Countable strongly ω -homogeneous Models). *Assume that L is a countable language, and let Σ be a complete set of L -sentences. For each $n \in \mathbb{N}$ let \mathcal{T}_n be a countable collection of partial n -types in L , with each partial type in each \mathcal{T}_n being Σ -satisfiable. Then there is a countable strongly ω -homogeneous model of Σ that realizes every partial n -type in \mathcal{T}_n for each n .*

Proof. Since Σ is complete and L and the sets \mathcal{T}_n are countable, there is a countable model \mathcal{M}_0 of T in which all the given partial types are realized. We inductively construct an elementary chain $(\mathcal{M}_n \mid n \in \mathbb{N})$ of countable structures and for each $n \geq 1$ a countable set \mathcal{F}_n of automorphisms of \mathcal{M}_n , such that the following conditions are satisfied: (1) for all $n \geq 0$, every elementary map between finite subsets of \mathcal{M}_n extends to an automorphism of \mathcal{M}_{n+1} that is a member of \mathcal{F}_{n+1} ; (2) for all $n \geq 1$ each automorphism of \mathcal{M}_n in \mathcal{F}_n extends to an automorphism of \mathcal{M}_{n+1} in \mathcal{F}_{n+1} . We also take \mathcal{F}_0 to

be empty. To see that this can be done, consider a countable model \mathcal{M}_n of Σ together with a countable set \mathcal{F}_n of automorphisms of \mathcal{M}_n . Using Theorem 13.3 there is a strongly ω_1 -homogeneous elementary extension \mathcal{N} of \mathcal{M}_n . (Of course \mathcal{N} need not be countable.) Since \mathcal{F}_n is countable and there are only countably many maps between finite subsets of \mathcal{M}_n , there is a countable set \mathcal{F} of automorphisms of \mathcal{N} with the property that each automorphism in \mathcal{F}_n and each elementary map between finite subsets of \mathcal{M}_n extends to an automorphism of \mathcal{N} that is in \mathcal{F} . By the Downward Löwenheim-Skolem Theorem there is a countable structure $\mathcal{M}_{n+1} \preceq \mathcal{N}$ such that $M_n \subseteq M_{n+1}$ and such that M_{n+1} is closed under f and f^{-1} for each $f \in \mathcal{F}$. In particular, we have $\mathcal{M}_n \subseteq \mathcal{M}_{n+1}$. Let \mathcal{F}_{n+1} be the set of restrictions of members of \mathcal{F} to \mathcal{M}_{n+1} . Then \mathcal{M}_{n+1} and \mathcal{F}_{n+1} have the desired properties.

The desired model \mathcal{M} of Σ is the union of the chain $(\mathcal{M}_n \mid n \in \mathbb{N})$. Note that by construction every automorphism of \mathcal{M}_n that is a member of \mathcal{F}_n extends to an automorphism of \mathcal{M} . \square

13.5. Corollary. *If the structure \mathcal{M} is κ -saturated and has cardinality κ , then \mathcal{M} is strongly κ -homogeneous.*

Proof. Suppose \mathcal{M} is κ -saturated and has cardinality κ . Let $A \subseteq M$ have cardinality $< \kappa$ and suppose $f: A \rightarrow M$ is an elementary map with respect to \mathcal{M} . Then the structures $(\mathcal{M}, a)_{a \in A}$ and $(\mathcal{M}, f(a))_{a \in A}$ are elementarily equivalent, and both of them are κ -saturated and have cardinality κ . Therefore, these two structures are isomorphic by Proposition 5.9; any isomorphism between them is an automorphism of \mathcal{M} that extends f . Hence \mathcal{M} is strongly κ -isomorphic. \square

In the rest of this chapter we explore the relations among several notions of “richness” for L -structures.

13.6. Definition. Let L be a first order language, \mathcal{M} an L -structure, and κ an infinite cardinal number.

(1) \mathcal{M} is κ -homogeneous if it has the following property for every subset A of M of cardinality $< \kappa$: any elementary mapping of A into \mathcal{M} can be extended to an elementary mapping of $A \cup \{a\}$ into \mathcal{M} , for each $a \in A$.

(2) \mathcal{M} is κ -universal if every structure \mathcal{N} that satisfies $\text{card}(N) < \kappa$ and $\mathcal{N} \equiv \mathcal{M}$ can be elementarily embedded into \mathcal{M} .

13.7. Theorem. *Let κ be an infinite cardinal number.*

(a) *Any strongly κ -homogeneous structure is κ -homogeneous.*

(b) *Any κ -saturated structure is κ -homogeneous and κ^+ -universal.*

(c) *Assume $\text{card}(L) < \kappa$. Any structure that is κ -homogeneous and κ -universal is κ -saturated.*

(d) *Any κ -homogeneous structure that is of cardinality κ is strongly κ -homogeneous.*

Proof. (a) Any restriction of an automorphism is an elementary map.

(b) Let \mathcal{M} be κ -saturated. Corollary 5.5 shows that \mathcal{M} is κ^+ -universal. To show that \mathcal{M} is κ -homogeneous, consider a subset A of M whose cardinality is less than κ and let $f: A \rightarrow M$ be an elementary map with respect to the structure \mathcal{M} . Then $(\mathcal{M}, a)_{a \in A} \equiv (\mathcal{M}, f(a))_{a \in A}$ and both of these structures are κ -saturated. Therefore, for any $b \in A$ there exists $c \in A$ such that $(\mathcal{M}, b, a)_{a \in A} \equiv (\mathcal{M}, c, f(a))_{a \in A}$. The desired extension of f can be obtained by setting $f(b) = c$.

(c) Let \mathcal{M} be κ -homogeneous and κ -universal, and suppose $\text{card}(L) < \kappa$. Let $A \subseteq M$ with $\text{card}(A) < \kappa$, and consider a 1-type $\Phi(x)$ in L_A that is finitely satisfiable in $\text{Th}((\mathcal{M}, a)_{a \in A})$. There is an L_A structure $(\mathcal{N}, f(a))_{a \in A}$ and an element $b \in N$ such that b realizes $\Phi(x)$ in $(\mathcal{N}, f(a))_{a \in A}$. Since the cardinality of L_A is $< \kappa$, the Downward Löwenheim-Skolem Theorem implies that we may assume $\text{card}(N) < \kappa$. Since \mathcal{M} is κ -universal, there exists an elementary embedding g of \mathcal{N} into \mathcal{M} . The composition $g \circ f$ maps A into M and is an elementary map with respect to \mathcal{M} . Since $\text{card}(A) < \kappa$ and \mathcal{M} is κ -homogeneous, there is an elementary map h that extends $g \circ f$ and such that $g(b)$ is in the range of h . If $c \in \text{dom}(h)$ satisfies $h(c) = g(b)$, then c must realize $\Phi(x)$ in $(\mathcal{M}, a)_{a \in A}$.

(d) Let \mathcal{M} be a κ -homogeneous structure, and let $f: A \rightarrow M$ be an elementary map with $\text{card}(A) < \kappa$ and $A \subseteq M$. Then we have $(\mathcal{M}, a)_{a \in A} \equiv (\mathcal{M}, f(a))_{a \in A}$. We can inductively extend f to an increasing chain of elementary mappings whose union is an automorphism of \mathcal{M} . Thus \mathcal{M} is strongly κ -homogeneous. \square

EXERCISES

13.8. Let \mathcal{M}, \mathcal{N} be L -structures that are elementarily equivalent. Show that there exist elementary extensions \mathcal{M}' of \mathcal{M} and \mathcal{N}' of \mathcal{N} such that $\mathcal{M}' \cong \mathcal{N}'$.

13.9. Let \mathcal{M}, \mathcal{N} be L -structures and let f be a map from a subset of M into N . Assume f is elementary with respect to \mathcal{M}, \mathcal{N} . Show that there exist elementary extensions \mathcal{M}' of \mathcal{M} and \mathcal{N}' of \mathcal{N} and an isomorphism g of \mathcal{M}' onto \mathcal{N}' such that g is an extension of f .

13.10. Let L be the language whose nonlogical symbols consist of infinitely many constant symbols $\{c_n \mid n \in \mathbb{N}\}$. Let Σ be the set of L -sentences $c_m \neq c_n$ for all distinct $m, n \in \mathbb{N}$. It follows from Example 4.15 that T admits QE. Every model of T has a substructure isomorphic to $(\mathbb{N}, n)_{n \in \mathbb{N}}$, so T is complete by Corollary 7.3.

- Which countable model of T is ω -saturated?
- Which countable models of T are strongly ω -homogeneous.

13.11. Let L be the language whose nonlogical symbols consist of a unary function symbol F . Let T be the theory in L of the class of all L -structures (M, f) in which f is a bijection from M onto itself and f has no finite cycles. From Exercise 11.13 we know that T admits QE and is complete, and we know that T is strongly minimal and we understand the meaning of the dimension of a model of T .

- Which countable models of T are strongly ω -homogeneous?

13.12. Let \mathcal{M} be an L -structure and $A \subseteq M$. Recall that $R \subseteq M^m$ is called *A-definable in \mathcal{M}* if there is an L -formula $\varphi(x_1, \dots, x_m, y_1, \dots, y_n)$ and parameters a_1, \dots, a_n from A such that

$$R = \{(u_1, \dots, u_m) \in M^m \mid \mathcal{M} \models \varphi[u_1, \dots, u_m, a_1, \dots, a_n]\}.$$

Now suppose \mathcal{M} is κ -saturated and strongly κ -homogeneous and $A \subseteq M$ has $\text{card}(A) < \kappa$. Suppose further that $R \subseteq M^m$ is M -definable in \mathcal{M} .

- Show that R is A -definable in \mathcal{M} iff R is fixed setwise by every automorphism of \mathcal{M} that fixes A pointwise.

14. OMITTING TYPES

In the first section of this chapter we investigate when a set Σ of L -sentences has a model \mathcal{M} in which a given set of L -formulas Φ *fails* to be realized. In that situation we say \mathcal{M} *omits* Φ .

The main result of this section (Theorem 14.3) gives a sufficient condition for a given countable family of partial types to be omitted from some model of a set of sentences Σ in a countable language. If Σ is complete, this result yields a condition on a type $p \in S_n(\Sigma)$ that is *necessary and sufficient* for the existence of a model of Σ that omits p . (Theorem 14.9).

In the second section we apply these results to the study of models of Σ that realize the smallest possible collection of types (*atomic* and *prime* models).

In the last section we use results from the earlier sections to characterize and study ω -categoricity of Σ , where Σ is complete and its language is countable. The results in this section show that ω -categoricity is a rich and robust concept.

Most results in this chapter require that the language be countable.

OMITTING A PARTIAL TYPE

14.1. Definition. Let Σ be a satisfiable set of L -sentences and let $\Phi(x)$ be a partial x -type in L , where $x = x_1, \dots, x_n$. We say that Σ *locally omits* $\Phi(x)$ if for any Σ -realizable formula $\psi(x)$, there is a formula $\varphi(x) \in \Phi(x)$ such that $\psi(x) \wedge \neg\varphi(x)$ is Σ -realizable.

14.2. Remark. The notion of “local omitting” may seem more natural when rephrased topologically: Σ locally omits $\Phi(x)$ iff the closed subset $\{p(x) \in S_x(\Sigma) \mid \Phi(x) \subseteq p(x)\}$ has empty interior in the space $S_x(\Sigma)$.

14.3. Theorem (Omitting Types Theorem). *Suppose L is a countable language, and let Σ be a satisfiable set of L -sentences. For each $k \geq 1$ let Φ_k be a partial n_k -type in L that is locally omitted by Σ . Then there is a countable model \mathcal{M} of Σ such that for all k , the partial type Φ_k is omitted in \mathcal{M} .*

Proof. In order to keep the notation simpler, we first consider the case of a single partial 1-type $\Phi(x)$. Let L' be the language obtained from L by adding a countable set of new constants $\{c_n \mid n \geq 1\}$. Let $\sigma_1, \sigma_2, \dots$ list the sentences of L' . Starting with Σ we construct a chain $\Sigma = \Sigma_0 \subseteq \Sigma_1 \subseteq \dots$ of satisfiable sets of L' -sentences such that each Σ_{m+1} is a finite extension of Σ_m and the following conditions are satisfied for all $m \geq 1$:

- (a) Σ_m contains σ_m or $\neg\sigma_m$;
- (b) If $\sigma_m = \exists y\psi(y) \in \Sigma_m$, then $\psi(c_p) \in \Sigma_m$ for some $p \geq 1$;
- (c) There is some $\varphi(x) \in \Phi(x)$ such that $\neg\varphi(c_m) \in \Sigma_m$.

First we show that it is sufficient to construct such a chain. Let $\Sigma' = \bigcup_{m=1}^{\infty} \Sigma_m$. Then $\Sigma \subseteq \Sigma'$ and Σ' is a maximal satisfiable set of L' -sentences. Also, for any L' -formula $\psi(y)$, if $\exists y\psi(y)$ is in Σ' , then $\psi(c_p) \in \Sigma'$ for some p . Note that these conditions also appear in the usual proof of the Gödel Completeness Theorem. As in that proof, we define an L' -prestructure \mathcal{M} with $M = \{c_n \mid n \geq 1\}$ by interpreting the nonlogical symbols of L' as follows:

(1) If P is a k -ary predicate symbol in L , let

$$P^{\mathcal{M}}(c_{i_1}, \dots, c_{i_k}) \Leftrightarrow P(c_{i_1}, \dots, c_{i_k}) \in \Sigma'.$$

(2) If F is a k -ary function symbol in L , take $F^{\mathcal{M}}(c_{i_1}, \dots, c_{i_k})$ to be the earliest c_p for which $F(c_{i_1}, \dots, c_{i_k}) = c_p$ is in Σ' . (Note that the sentence $\exists y(F(c_{i_1}, \dots, c_{i_k}) = y)$ is in Σ' since it is a valid sentence and Σ' is maximal satisfiable; therefore by (b) above there exists c_p such that $F(c_{i_1}, \dots, c_{i_k}) = c_p$ is in Σ' .)

(3) If c is a constant in L' , take $c^{\mathcal{M}}$ to be the earliest c_p for which $c = c_p$ is in Σ' . (As in (2), condition (b) above ensures that some such c_p exists.)

It is routine to show (by induction on the length of sentences) that for any L' -sentence σ , one has

$$\mathcal{M} \models \sigma \Leftrightarrow \sigma \in \Sigma'.$$

(In this induction, condition (b) is used for the handling of existential quantifiers.) In particular, this shows that \mathcal{M} is a prestructure, since any instance of an equality axiom (indeed, any valid sentence) is a member of Σ' .

Furthermore, no element of M satisfies all formulas in $\Phi(x)$; this is ensured by condition (c) above. Thus the reduct of \mathcal{M} to L nearly satisfies the conclusion of the Theorem; the only problem is that \mathcal{M} will generally be a prestructure rather than a structure. ($\mathcal{M} \models c_m = c_n$ if and only if the sentence $c_m = c_n$ is in Σ' , and this may happen even when $m \neq n$.)

By applying the construction discussed in Appendix B of Chapter 1, we obtain an L' -structure \mathcal{N} as a quotient of the prestructure \mathcal{M} , and we have $\mathcal{N} \models \Sigma'$ and \mathcal{N} omits $\Phi(x)$. The reduct of \mathcal{N} to L is the desired model of Σ that omits $\Phi(x)$.

This argument shows that it suffices to construct $\Sigma = \Sigma_0 \subseteq \Sigma_1 \subseteq \dots$ satisfying the conditions above, including (a),(b),(c). We set $\Sigma_0 = \Sigma$ and define Σ_m for $m \geq 1$ by induction. Given Σ_{m-1} , with $m \geq 1$, we construct Σ_m in three steps as follows:

(i) Let

$$\Sigma'_{m-1} = \begin{cases} \Sigma_{m-1} \cup \{\sigma_m\}, & \text{if this is satisfiable} \\ \Sigma_{m-1} \cup \{\neg\sigma_m\}, & \text{otherwise.} \end{cases}$$

Note that Σ'_{m-1} is satisfiable.

(ii) Suppose $\sigma_m \in \Sigma'_{m-1}$ and $\sigma_m = \exists y \psi(y)$. Choose c_p to be the first new constant not occurring in ψ or Σ'_{m-1} . Note that $\Sigma'_{m-1} \cup \{\psi(c_p)\}$ is satisfiable. (Otherwise $\Sigma'_{m-1} \models \neg \psi(c_p)$ so $\Sigma'_{m-1} \models \forall y \neg \psi(y)$; i.e., $\Sigma'_{m-1} \models \neg \sigma_m$, which is a contradiction.) Let $\Sigma''_{m-1} = \Sigma'_{m-1} \cup \{\psi(c_p)\}$.

(iii) We show that $\Sigma''_{m-1} \not\models \varphi(c_m)$ for some $\varphi \in \Phi$. Suppose otherwise. There are L -formulas $\psi_1(z), \dots, \psi_k(z)$, with $z = z_1, \dots, z_N$, such that $\Sigma''_{m-1} = \Sigma \cup \{\psi_j(c_1, \dots, c_N) \mid j = 1, \dots, k\}$. (Choose $N \geq m$ so the new constants in Σ''_{m-1} are among c_1, \dots, c_N . Choose variables z_1, \dots, z_N not occurring in $\Sigma''_{m-1} \setminus \Sigma$. Let $\psi_j(z)$ be the result of replacing c_i by z_i in the j th sentence of $\Sigma''_{m-1} \setminus \Sigma$, for $i = 1, \dots, N$ and $j = 1, \dots, k$.) So $\Sigma \cup \{\psi_j(c_1, \dots, c_N) \mid j = 1, \dots, k\} \models \sigma(c_m)$ for all $\varphi \in \Phi$. Consider

$$\tau(z_m) = \exists z_1 \dots \exists z_{m-1} \exists z_{m+1} \dots \exists z_N \bigwedge_{j=1}^k \psi_j(z_1, \dots, z_N).$$

We see that $\tau(z_m)$ is Σ -realizable, and $\Sigma \cup \{\tau(z_m)\} \models \varphi(z_m)$ for all $\varphi \in \Phi$. Hence $\tau(x)$ is Σ -realizable and $\Sigma \cup \{\tau(x)\} \models \varphi(x)$ for all $\varphi \in \Phi$, contradicting the hypothesis that Φ is locally omitted. So, we can finally define Σ_m to be equal to the set $\Sigma''_{m-1} \cup \{\neg \varphi(c_m)\}$, where $\varphi \in \Phi$ is chosen so that this set is satisfiable. It is clear that Σ_m satisfies the conditions (a), (b), and (c). (Note that if $m \geq 1$ and σ_m is in Σ_m , then σ_m must be in Σ'_{m-1} , since otherwise $\neg \sigma_m \in \Sigma'_m \subseteq \Sigma_m$. This ensures that condition (b) remains true of Σ_m .)

This completes the proof of the Theorem for a single partial 1-type. The proof can easily be modified to cover countably many partial types $\Phi_k(x_1, \dots, x_{n_k})$, $k \geq 1$ (each locally omitted by Σ). Requirements (a) and (b) of the construction remain unchanged. For (c), enumerate all finite sequences $\alpha = (k, c_{i_1}, \dots, c_{i_{n_k}})$ where $k \geq 1$, and $c_{i_1}, \dots, c_{i_{n_k}}$ are new constants. Then condition (c) becomes:

For all $m \geq 1$, the set Σ_m contains a sentence $\neg \varphi(c_{i_1}, \dots, c_{i_{n_k}})$ where $\varphi(x_1, \dots, x_{n_k}) \in \Phi_k$ and $\alpha = (k, c_{i_1}, \dots, c_{i_{n_k}})$ is the m^{th} sequence in the enumeration of all such sequences. \square

14.4. Remark. The Omitting Types Theorem can be rephrased more topologically as follows: *Suppose L is a countable language, and let Σ be a satisfiable set of L -sentences. For each positive integer m , let X_m be a meager subset of $S_m(\Sigma)$. Then there is a countable model \mathcal{M} of Σ that omits every type in the union $\bigcup\{X_m \mid m \geq 1\}$.*

First, we show how this follows from Theorem 14.3. Because each X_m is meager, we may let $(X_{m,n} \mid m, n \geq 1)$ be closed subsets of $S_m(\Sigma)$, such that each $X_{m,n}$ has empty interior and $X_m \subseteq \bigcup\{X_{m,n} \mid n \geq 1\}$ for all $m \geq 1$. Because $X_{m,n}$ is closed, it can be written as

$$X_{m,n} = \{p(x_1, \dots, x_m) \mid \Phi_{m,n}(x_1, \dots, x_m) \subseteq p(x_1, \dots, x_m)\}$$

for some set $\Phi_{m,n}(x_1, \dots, x_m)$ of L -formulas. Because $X_{m,n}$ has empty interior in $S_m(\Sigma)$, it follows that $\Phi_{m,n}$ is locally omitted by Σ . Theorem 14.3 therefore implies that there is a countable model \mathcal{M} of Σ that omits every type in

$$\bigcup_{m=1}^{\infty} \{p(x_1, \dots, x_m) \mid \Phi_{m,n}(x_1, \dots, x_m) \subseteq p(x_1, \dots, x_m) \text{ for some } n \geq 1\};$$

hence \mathcal{M} omits every type in the union $\bigcup\{X_m \mid m \geq 1\}$.

For the reverse direction, suppose Σ and $\Phi_k(x_1, \dots, x_{n_k})$ are as in the hypotheses of the Omitting Types Theorem. For each $m \geq 1$ let X_m be the set of types $p \in S_m(\Sigma)$ such that $p \supseteq \Phi_{n_k}$ for some k with $n_k = m$. Then each X_m is a meager subset of $S_m(\Sigma)$ and hence, by the statement above, there is a countable model of Σ that omits every type in $\bigcup\{X_m \mid m \geq 1\}$. In particular, this model omits Φ_k for every $k \geq 1$.

14.5. Remark. The Omitting Types Theorem (as stated here) is false for uncountable languages. An example of a partial 1-type that is locally omitted by a set of sentences Σ , but not omitted in any model of Σ is the following: Let I be an uncountable set and let L be the language whose nonlogical symbols are the distinct constants $\{c_i \mid i \in I\} \cup \{d_n \mid n \in \mathbb{N}\}$. Take Σ to be the set of sentences $\neg c_i = c_j$ for all distinct $i, j \in I$. Then $\Phi(x) = \{x = d_n \mid n \in \mathbb{N}\}$ is locally omitted by Σ , but not omitted by Σ . Indeed, every model of Σ is uncountable, while any structure that omits $\Phi(x)$ must be countable.

PRINCIPAL TYPES, ATOMIC MODELS

Fix a satisfiable set of L -sentences Σ and variables $x = x_1, \dots, x_n$.

14.6. Definition. (1) An L -formula $\varphi(x)$ is Σ -complete if $\varphi(x)$ is Σ -realizable and for every L -formula $\psi(x)$, either $\Sigma \models \varphi(x) \rightarrow \psi(x)$ or $\Sigma \models \varphi(x) \rightarrow \neg\psi(x)$.

(2) A type $p(x) \in S_x(\Sigma)$ is *principal* if it contains a Σ -complete formula.

14.7. Remark. If $\varphi(x)$ is Σ -complete, then there is a unique type $p(x) \in S_x(\Sigma)$ that has $\varphi(x) \in p(x)$ (and this type is principal by definition). Indeed, $p(x) = \{\psi(x) \mid \Sigma \models \varphi(x) \rightarrow \psi(x)\}$.

14.8. Remark. Note that $p(x) \in S_x(\Sigma)$ is principal iff the singleton $\{p(x)\}$ is an open set in $S_x(\Sigma)$. That is, principal x -types are exactly the isolated points of $S_x(\Sigma)$.

14.9. Theorem. *Let Σ be a complete set of sentences in a countable language L , and let $p(x) \in S_x(\Sigma)$. Then $p(x)$ is principal if and only if it is realized in every model of Σ (if and only if it is realized in every countable model of Σ).*

Proof. Consider $p(x) \in S_x(\Sigma)$. If $p(x)$ is not principal, then $\{p(x)\}$ has empty interior, and therefore $p(x)$ is locally omitted (see 14.2). The Omitting Types Theorem yields that $p(x)$ is omitted by some model of Σ . For the converse, suppose $p(x)$ is principal and let $\varphi(x) \in p(x)$ be a Σ -complete formula. Since Σ is complete, it follows that $\Sigma \models \exists x_1 \dots \exists x_n \varphi(x)$ and hence $p(x)$ is realized in every model of Σ .

The second equivalence is a consequence of the Downward Löwenheim-Skolem Theorem, since L is countable. \square

14.10. Definition. A structure \mathcal{M} is *atomic* if every n -type realized in \mathcal{M} is principal relative to $\text{Th}(\mathcal{M})$. (Equivalently: \mathcal{M} is atomic if every n -tuple in M satisfies a $\text{Th}(\mathcal{M})$ -complete formula.)

14.11. Fact. (a) If $\varphi(x_1, \dots, x_n, y)$ is a Σ -complete formula, then so is $\exists y \varphi(x_1, \dots, x_n, y)$;

(b) if \mathcal{M} is atomic and $a_1, \dots, a_n \in M$, then $(\mathcal{M}, a_1, \dots, a_n)$ is also atomic.

14.12. Theorem. *Suppose \mathcal{M} and \mathcal{N} are countable atomic structures for the same language. If $\mathcal{M} \equiv \mathcal{N}$, then $\mathcal{M} \cong \mathcal{N}$.*

Proof. Let $\mathcal{M} \equiv \mathcal{N}$ be countable and atomic. Let \mathcal{F} be the set of all functions $f: A \rightarrow B$, where $A \subseteq M$ and $B \subseteq N$ are finite and f is elementary with respect to \mathcal{M}, \mathcal{N} . We will show that \mathcal{F} is a back-and-forth system from \mathcal{M} to \mathcal{N} . Since M and N are countable, it is then easy to build an isomorphism from \mathcal{M} onto \mathcal{N} as the union of an increasing chain of maps from \mathcal{F} .

Consider $f: A \rightarrow B$ from \mathcal{F} , and write $A = \{a_1, \dots, a_n\}$. Let a be any element of M . We must extend f to $g \in \mathcal{F}$ with $a \in \text{dom}(g)$. Let $p(x, y) = \text{tp}_{\mathcal{M}}(a_1, \dots, a_n, a)$ with $x = x_1, \dots, x_n$ and y a single variable. Since \mathcal{M} is atomic, there exists a $\text{Th}(\mathcal{M})$ -complete formula $\varphi(x, y)$ in $p(x, y)$. Since f is elementary with respect to \mathcal{M}, \mathcal{N} , we have that $\mathcal{N} \models \exists y \varphi(x, y)[f(a_1), \dots, f(a_n)]$. Thus we may choose $b \in N$ such that $\mathcal{N} \models \varphi[f(a_1), \dots, f(a_n), b]$. Since $\varphi(x, y)$ is a complete formula for $\text{Th}(\mathcal{M}) = \text{Th}(\mathcal{N})$, we conclude that $\text{tp}_{\mathcal{N}}(f(a_1), \dots, f(a_n), b) = \text{tp}_{\mathcal{M}}(a_1, \dots, a_n, a)$. Hence the extension of f defined on $A \cup \{a\}$ by sending a to b is elementary; it is the desired element of \mathcal{F} .

This proves the “forth” property for \mathcal{F} ; the “back” property is proved by exchanging the roles of \mathcal{M} and \mathcal{N} . \square

14.13. Corollary. *A countable atomic model is strongly ω -homogeneous.*

Proof. This is immediate from the proof of the previous result. \square

14.14. Definition. Let Σ be a set of L -sentences and \mathcal{M} a model of Σ . We say \mathcal{M} is a *prime model* of Σ if for every model \mathcal{N} of Σ , there is an elementary embedding of \mathcal{M} into \mathcal{N} .

14.15. Remark. If Σ has a prime model, then Σ must be complete. If L is countable, then any prime model of Σ must be countable.

14.16. Theorem. *Let Σ be a complete set of sentences in a countable language and let \mathcal{M} be a model of Σ . Then \mathcal{M} is a prime model of Σ if and only if \mathcal{M} is countable and atomic.*

Proof. (\Rightarrow) Suppose \mathcal{M} is a prime model of Σ . Since the language is countable, Σ has a countable model by the Downward Löwenheim-Skolem Theorem. Since \mathcal{M} can be embedded in this model, \mathcal{M} must also be countable. To show \mathcal{M} is atomic, consider an arbitrary n -type $p(x) = \text{tp}_{\mathcal{M}}(a)$ realized in \mathcal{M} . We must show $p(x)$ is principal. Given any model \mathcal{N} of Σ , let f be an elementary embedding of \mathcal{M} into \mathcal{N} . Evidently $p(x)$ is realized in \mathcal{N} by $(f(a_1), \dots, f(a_n))$. That is, $p(x)$ is realized in every model of Σ , so $p(x)$ is a principal type by Theorem 14.9. This shows that \mathcal{M} is atomic.

(\Leftarrow) Suppose the model \mathcal{M} of Σ is countable and atomic. Let \mathcal{N} be any model of Σ . Since Σ is complete, we have $\mathcal{N} \equiv \mathcal{M}$. Let \mathcal{F} be the set of all functions $f: A \rightarrow B$, where $A \subseteq \mathcal{M}$ and $B \subseteq \mathcal{N}$ are finite and f is elementary with respect to \mathcal{M}, \mathcal{N} . Arguing as in the proof of Theorem 14.12, we have that \mathcal{F} has the “forth” property from \mathcal{M} to \mathcal{N} . Since \mathcal{M} is countable, it is then easy to build an elementary embedding from \mathcal{M} into \mathcal{N} as the union of an increasing chain of maps from \mathcal{F} . \square

For the rest of this section we consider the *existence* of atomic models.

Note that if \mathcal{M} is atomic, then so is every elementary substructure of \mathcal{M} . Therefore, if Σ is a complete set of sentences in a countable language and Σ has an atomic model, then it has a countable atomic model, by the Downward Löwenheim-Skolem Theorem.

14.17. Theorem. *Let Σ be a complete set of sentences in a countable language L . Then Σ has an atomic model if and only if the set of principal types is dense in $S_n(\Sigma)$ for every $n \geq 1$. Equivalently, this is true if and only if every Σ -realizable L -formula $\varphi(x)$ is an element of some principal type $p(x) \in S_x(\Sigma)$.*

Proof. (\Rightarrow) Let \mathcal{M} be an atomic model of Σ . If $\varphi(x_1, \dots, x_n)$ is Σ -realizable, then it is realized in \mathcal{M} , say by (a_1, \dots, a_n) , since Σ is complete. That is, $\varphi(x_1, \dots, x_n)$ is a member of $\text{tp}_{\mathcal{M}}(a_1, \dots, a_n)$, which is a principal type, since \mathcal{M} is atomic.

(\Leftarrow) For each $n \geq 1$ let

$$\Delta_n(x) = \{\neg\varphi(x) \mid \varphi(x) \text{ is a } \Sigma\text{-complete formula}\}$$

where $x = x_1, \dots, x_n$. A model of Σ is atomic iff it omits every Δ_n . We use the Omitting Types Theorem to prove that there exists a model of Σ that omits all of these partial types. Hence we must show that Δ_n is locally omitted by Σ for each $n \geq 1$. That is, we must show that the closed subset $\{p(x) \in S_x(\Sigma) \mid \Delta_n(x) \subseteq p(x)\}$ has empty interior in the space $S_x(\Sigma)$. If not, the density of principal types in $S_x(\Sigma)$ ensures that there is a principal

type $p(x)$ that contains $\Delta_n(x)$. That is, $p(x)$ contains the negation of every Σ -complete formula $\varphi(x)$. This is a contradiction. \square

The next result gives a useful sufficient condition for the existence of an atomic model of a theory T in a countable language. Its proof is based on the following topological fact.

14.18. Fact. Let X be a compact hausdorff space with a basis of clopen subsets. Suppose the set of isolated points in X is not dense. Then $\text{card}(X) \geq 2^\omega$.

Proof. By assumption, there exists a nonempty clopen set U that contains no isolated points. We may take distinct $p, q \in U$ and then find a clopen set $U_0 \subseteq U$ with $p \in U_0$ and $q \notin U_0$. Then let $U_1 = U \setminus U_0$. We have partitioned U into two clopen sets, each of which is nonempty. In particular, U_0, U_1 are disjoint and neither one contains an isolated point. Applying this construction inductively, we can construct an infinite binary tree of nonempty clopen subsets of U . Each branch of this tree is a decreasing chain of nonempty closed sets, so by the compactness of X its intersection is nonempty. Moreover, the sets in two distinct branches of this tree are eventually disjoint, so their intersections are disjoint. Hence distinct branches of the tree give rise to distinct elements of X . Since there are 2^ω many branches on this tree, we conclude $\text{card}(X) \geq 2^\omega$. \square

14.19. Corollary. Let Σ be a complete set of sentences in a countable language, and suppose that for all $n \geq 1$ we have $\text{card}(S_n(\Sigma)) < 2^\omega$. Then Σ has a countable atomic model.

Proof. For each $n \geq 1$, the space $S_n(\Sigma)$ is compact hausdorff with a basis of clopen sets, so the topological fact 14.18 applies. Therefore, our cardinality assumption implies that the isolated points are dense in $S_n(\Sigma)$ for all $n \geq 1$. By Theorem 14.17, Σ has a countable atomic model. \square

14.20. Corollary. Let Σ be a complete set of sentences in a countable language. If Σ has a countable ω -saturated model, then T also has a countable atomic model.

Proof. The assumptions yield that $S_n(\Sigma)$ is countable for all $n \geq 1$. Hence the preceding Corollary implies that Σ has a countable atomic model. \square

Fix a countable language L and a complete set of L -sentences Σ .

14.21. **Theorem** (Engeler, Ryll-Nardzewski, Svenonius). *Suppose Σ has only infinite models. The following conditions are equivalent:*

- (1) Σ is ω -categorical;
- (2) For each $n \geq 1$, every Σ -realizable n -type is principal;
- (3) For each $n \geq 1$, there are only finitely many Σ -realizable n -types;
- (4) For each $n \geq 1$ and $x = x_1, \dots, x_n$, there are finitely many formulas $\varphi_1(x), \dots, \varphi_{k_n}(x)$ such that each formula $\sigma(x)$ is Σ -equivalent to $\varphi_j(x)$ for some $j = 1, \dots, k_n$.
- (5) Every model of Σ is atomic.

Proof. Statements (2), (3), and (4) are easily seen to be equivalent using topological arguments, after recalling (A) that the Σ -equivalence classes of L -formulas $\varphi(x)$ correspond exactly to the clopen subsets of $S_x(\Sigma)$, and (B) that the principal types are exactly the isolated points of $S_x(\Sigma)$. So (2) implies that $S_x(\Sigma)$ is a compact hausdorff space in which every point is isolated; this yields that $S_x(\Sigma)$ is finite. If (3) holds, then $S_x(\Sigma)$ has only finitely many clopen subsets, which is equivalent to (4). Finally, if $S_x(\Sigma)$ has only finitely many clopen subsets, then it has only finitely many closed subsets; in particular it has only finitely many elements.

(1) \Rightarrow (2): If Σ is ω -categorical, then any two countable models of Σ realize the same types. Therefore, applying the Downward Löwenheim-Skolem Theorem, no Σ -realizable type is omitted from any model of Σ . This implies that all types are principal.

(2) \Rightarrow (5): Immediate from the definitions.

(5) \Rightarrow (1): Immediate by Theorem 14.12. □

14.22. **Remark.** Let \mathcal{M} be a countable structure for a countable first order language and suppose $T = \text{Th}(\mathcal{M})$ is ω -categorical. Let G be the automorphism group of \mathcal{M} , acting coordinatewise on M^n for each $n \geq 1$. Then G has only finitely many distinct orbits on M^n for each n . This is an immediate consequence of condition (3) in Theorem 14.21 and the fact that the unique countable model of an ω -categorical theory is strongly ω -homogeneous. (See Theorems 6.2 and 13.4 or, alternatively, Theorem 14.12 and the existence results for atomic models in the previous section.)

Infinite permutation groups arising as $\text{Aut}(\mathcal{M})$ for an ω -categorical structure \mathcal{M} for a countable language have turned out to be very interesting. They are treated in the book *Oligomorphic Permutation Groups* by Peter Cameron, and remain an object of active research.

The next result gives a sufficient condition for $\text{Th}(\mathcal{M})$ to be ω -categorical that is based on automorphism group considerations of the kind just discussed.

14.23. Theorem. *Let L be a countable language, let \mathcal{M} be any L -structure, and $T = \text{Th}(\mathcal{M})$. If $G = \text{Aut}(\mathcal{M})$ has only finitely many orbits on M^n for each $n \geq 1$, then T is ω -categorical.*

Proof. Let \mathcal{M} satisfy the given hypotheses. We will show that \mathcal{M} realizes every T -realizable type. The automorphism condition on \mathcal{M} implies that \mathcal{M} can only realize finitely many n -types for each n . Therefore T is ω -categorical since it satisfies condition (3) of Theorem 14.21.

So, let $p(x_1, \dots, x_n)$ be any T -realizable n -type. Given $\varphi(x_1, \dots, x_n) \in p(x_1, \dots, x_n)$, we have $T \models \exists x_1 \dots \exists x_n \varphi(x_1, \dots, x_n)$ (since T is complete), and hence there is $a \in M^n$ with $\mathcal{M} \models \varphi[a]$. Let $F \subseteq M^n$ be a finite set that selects one n -tuple from each orbit under the action of G . The a realizing φ in \mathcal{M} can be taken from F . Suppose $p(x_1, \dots, x_n)$ is not realized in \mathcal{M} . For each $a \in F$ we get $\psi_a \in p(x_1, \dots, x_n)$ such that $\mathcal{M} \models \neg \psi_a[a]$. Consider $\varphi(x_1, \dots, x_n) = \bigwedge_{a \in F} \psi_a \in p$. As argued above, there is some $a \in F$ satisfying $\varphi(x_1, \dots, x_n)$ in \mathcal{M} . However, this implies that a satisfies ψ_a in \mathcal{M} , since ψ_a is a conjunct of φ . This is a contradiction. \square

14.24. Example. Let $\mathcal{M} = (\mathbb{Q}, <)$. We will show directly (using additional structure on \mathbb{Q}) that $G = \text{Aut}(\mathcal{M})$ has only finitely many orbits on \mathbb{Q}^n for each n . The main idea is the following: given $q_0 < q_1 < \dots < q_n$ in \mathbb{Q} , there is $g \in G$ such that $g(k) = q_k$ for $k = 0, \dots, n$. Indeed, we may define g by:

$$g(r) = \begin{cases} r + q_0, & \text{for } r < 0 \\ rq_1 - (r - 1)q_0, & \text{for } 0 \leq r < 1 \\ \vdots & \vdots \\ (r - k)q_{k+1} - (r - k - 1)q_k, & \text{for } k \leq r < k + 1 \\ \vdots & \vdots \\ r + q_n - n, & \text{for } n \leq r. \end{cases}$$

This construction together with Theorem 14.23 proves that $\text{Th}(\mathbb{Q}, <)$ is ω -categorical. (However, this proof does not show, as does Cantor's back and forth argument, that the theory of $(\mathbb{Q}, <)$ is axiomatized by the sentence asserting that it is a dense linear ordering without endpoints. Indeed, it gives no axiomatization of $\text{Th}(\mathbb{Q}, <)$ at all.)

The next result illustrates robustness of the class of ω -categorical structures for a countable language.

14.25. Theorem. *Let Σ be an ω -categorical complete set of sentences in a countable language.*

(a) *If \mathcal{M} is any model of Σ and $a_1, \dots, a_n \in M$, then $\text{Th}(\mathcal{M}, a_1, \dots, a_n)$ is also ω -categorical.*

(b) *If \mathcal{M} is any model of Σ and if \mathcal{N} is any structure for any countable language such that the universe of \mathcal{N} and all of its interpretations of predicate symbols and function symbols (including constant symbols) are 0-definable*

in \mathcal{M} , then $\text{Th}(\mathcal{N})$ is also ω -categorical. (The universe of \mathcal{N} is allowed to be a set of n -tuples from the universe of \mathcal{M} .) In particular, the theory of any reduct of \mathcal{M} to a smaller language is ω -categorical.

Proof. (a) It suffices to consider the case where $n = 1$. If $a \in M$, the automorphism group of (\mathcal{M}, a) is the stabilizer of a in $G = \text{Aut}(\mathcal{M})$, which is the subgroup $G_a = \{\sigma \in G \mid \sigma(a) = a\}$. For $b, c \in M^n$, the orbits of b and c under G_a are equal if and only if the orbits of (b, a) and (c, a) are equal under G . Since G has only finitely many orbits acting on M^{n+1} , it follows that G_a has only finitely many orbits acting on M^n .

(b) If $N \subseteq M^d$, then we may regard N^n as a subset of M^{dn} . The definability assumptions ensure that $\text{Aut}(\mathcal{M})$ acts by automorphisms on N , and N^n is invariant under the action of $\text{Aut}(\mathcal{M})$ on M^{dn} . It follows that each $\text{Aut}(\mathcal{N})$ orbit in N^n is a union of some $\text{Aut}(\mathcal{M})$ orbits in M^{dn} . Hence the number of such orbits (in both cases) must be finite. By Theorem 14.23, $\text{Th}(\mathcal{N})$ is ω -categorical. \square

14.26. Remark. Using part (a) of the previous result, part (b) can be strengthened to allow the predicate and function symbol interpretations given by \mathcal{N} to be A -definable, where A is a fixed finite subset of M . However, simple examples show that part (b) does not necessarily hold if the interpretations given by \mathcal{N} are A -definable for some *infinite* set $A \subseteq M$. (For example take \mathcal{M} to be the set \mathbb{N} and take \mathcal{N} to be the structure $(\mathbb{N}, 0, 1, 2, 3, \dots)$ in which every element of \mathbb{N} is named.)

We close this chapter with a result that is interesting but somewhat curious; it's main significance is that it leads naturally to an open question that has been actively worked on, with only partial success, for more than 50 years.

14.27. Theorem (Vaught). *Let T be a complete theory in a countable language. If T is not ω -categorical, then T has at least 3 nonisomorphic countable models.*

Proof. Suppose that the number of nonisomorphic countable models of T is countable. In that case, there are only countably many T -realizable n -types for all $n \geq 1$ (as countably many countable models realize all the types that can be realized). Hence by Corollary 14.19, T has a countable atomic model \mathcal{M} and by Theorem 6.2, T has a countable ω -saturated model \mathcal{N} . If $\mathcal{M} \cong \mathcal{N}$ then every T -realizable type is realized in \mathcal{M} and hence is principal, so T is ω -categorical, by Theorem 14.21.

So we assume that $\mathcal{M} \not\cong \mathcal{N}$. For some $n \geq 1$ there exists at least one n -type $p(x_1, \dots, x_n)$ that is T -realizable and not principal. Then \mathcal{N} realizes $p(x_1, \dots, x_n)$ and \mathcal{M} doesn't. Let b_1, \dots, b_n realize $p(x_1, \dots, x_n)$ in \mathcal{N} and let $T' = \text{Th}(\mathcal{N}, b_1, \dots, b_n)$ in $L(b_1, \dots, b_n)$.

Note that T' cannot satisfy condition (4) of Theorem 14.21, since T does not satisfy this condition. Indeed, any L -formulas that are inequivalent in

T will remain inequivalent in T' . So T' has at least two nonisomorphic countable models. Let $(\mathcal{N}', c_1, \dots, c_n)$ be a countable model of T' that is not isomorphic to $(\mathcal{N}, b_1, \dots, b_n)$. Then $\mathcal{M} \not\cong \mathcal{N}'$, since \mathcal{N}' realizes p and \mathcal{M} doesn't. Moreover, $\mathcal{N} \not\cong \mathcal{C}$, since \mathcal{N} is ω -saturated and \mathcal{N}' isn't; otherwise, by Proposition 5.9, $(\mathcal{N}, b_1, \dots, b_n)$ and $(\mathcal{N}', c_1, \dots, c_n)$ would be isomorphic. \square

Quite a lot is known about the countable models of a complete theory in a countable language, and this topic has been an active one in research in model theory up to the present day. However, the following difficult problem is still open:

VAUGHT'S CONJECTURE: Let T be a complete theory in a countable language. If T has an uncountable number of nonisomorphic countable models then T has 2^ω many nonisomorphic countable models.

EXERCISES

14.28. Let T be a complete L -theory and let Φ be a partial n -type in L . If T has a model that omits Φ , show that Φ is locally omitted by T .

14.29. Let T be a complete theory in a countable language. For each positive integer k let $\Phi_k(x_1, \dots, x_{n_k})$ be a partial n_k -type in L that is omitted in some model \mathcal{M}_k of T . Show that there is a single countable model \mathcal{M} of T that omits $\Phi_k(x_1, \dots, x_{n_k})$ for all k .

14.30. Let L be a countable language and let L' be the result of adding countably many new predicate symbols $\{P_1, P_2, \dots\}$ to L . Let T be a complete theory in the language L' and let $\Phi(x_1, \dots, x_n)$ be a set of formulas in L . Let T_m be the set of sentences in T that contain P_j only for $j = 1, \dots, m$. Assume that for each m , the theory T_m has a model that omits $\Phi(x_1, \dots, x_n)$. Show that T has a model that omits $\Phi(x_1, \dots, x_n)$.

14.31. Here is an example of a complete theory T in a countable language such that no T -realizable type is principal. This is *the theory of infinitely many independent unary predicates*. Let L be the language with unary predicate symbols $(U_n \mid n \geq 1)$. Define T to be the theory $\text{Th}(\mathcal{M})$, where \mathcal{M} is the following structure: take $M = \mathcal{P}(\mathbb{N})$ and define the interpretation of each U_n by taking $U_n^{\mathcal{M}}(\alpha) \Leftrightarrow n \in \alpha$ for each $\alpha \subseteq \mathbb{N}$.

(a) Show that T is the theory axiomatized by the sentences $\sigma_{F,G}$ (for F, G disjoint finite subsets of \mathbb{N}) given by $\exists x \left(\bigwedge_{j \in F} U_j(x) \wedge \bigwedge_{j \in G} \neg U_j(x) \right)$.

(b) Show that no formula in L is T -complete (and hence T has no principal types). Indeed, if $\varphi(x)$ is a formula (with x a single variable) that does not contain U_n , show that both $\varphi(x) \wedge U_n(x)$ and $\varphi(x) \wedge \neg U_n(x)$ are T -realizable.

14.32. Let T be one of the following theories. (Each is a complete theory in a countable language, with no finite models.)

+ Equality on an infinite set with infinitely many named elements.

- + Infinite vector spaces over a countable field K .
- + ACF_p for a fixed characteristic p .
- + Bijections without a finite cycle.
- + Discrete linear orderings without endpoints.
- + Discrete linear orderings with minimum but no maximum.
- + Descending equivalence relations with infinite splitting of classes.
- + Dense linear orderings with increasing sequence of named elements.

For each of these theories, do the following:

- Show that T has a countable atomic model.
- Try to describe the countable atomic model of T as a clear, specific mathematical structure. (According to Theorem 14.12, the countable infinite atomic model of a complete theory is unique up to isomorphism, if such a model exists.)
- For each principal T -realizable n -type p , try to give explicitly a complete formula contained in p .

15. SKOLEM HULLS

Consider a language L and a set Σ of L -sentences. By the Downward Löwenheim-Skolem Theorem, we know that if $\mathcal{M} \models \Sigma$ and $A \subset M$, then there is an elementary substructure \mathcal{N} of \mathcal{M} with $A \subseteq N$, and that the cardinality of N can be taken to be $\leq \max(\text{card}(A), \text{card}(L))$. However, in general there are many such elementary substructures and we have no clear way to choose coherently among them. For later purposes it will be useful to remedy this problem.

One way around this issue, which we develop in this chapter, is to enlarge the language L to L' , and to extend Σ to a set Σ' of L' -sentences such that

(a) Every substructure of a model of Σ' is an elementary substructure.

If Σ' satisfies (a) and we have $\mathcal{M}' \models \Sigma'$ and a nonempty $A \subset M'$, then $\langle A \rangle_{\mathcal{M}'}$ is a natural elementary substructure of \mathcal{M}' which contains A .

We also want the models of Σ and of Σ' to be closely related to each other in a useful way; the connection we choose to emphasize here is

(b) Every model of Σ is the reduct to L of some model of Σ' .

In particular, (b) implies that Σ' is a conservative extension of Σ ; that is: for every L -sentence σ , if $\Sigma' \models \sigma$, then $\Sigma \models \sigma$. (Proof: otherwise we would have $\Sigma' \models \sigma$ while $\Sigma \cup \{\neg\sigma\}$ would be satisfiable. But then we would have a model \mathcal{M} of $\Sigma \cup \{\neg\sigma\}$ which could not be expanded to a model of Σ' , contradicting (b).) Therefore, moving from Σ to Σ' does not change the L -theories that can be considered.

Finally, in order to control the size of L' -substructures, we want

(c) $\text{card}(L') = \text{card}(L)$.

Now we turn to investigating the meaning of conditions (a),(b),(c) above, and to showing how they can be satisfied.

First we note that condition (a) implies QE.

15.1. Proposition. *Let Σ be a set of L -sentences. If every substructure of a model of Σ is an elementary substructure, then Σ has QE.*

Proof. This is an immediate consequence of our criteria for QE in Chapter 7. For example, we verify 7.5(2). Let $x = x_1, \dots, x_n$ and consider $p(x) \in S_x(\Sigma)$. Working in an ω -saturated model \mathcal{M} of Σ , let $a \in M^n$ realize $p(x)$ in \mathcal{M} . Let $\Phi(x)$ be the set of all quantifier-free formulas $\varphi(x)$ that are members of $p(x)$. Suppose $b \in M^n$ realizes $\Phi(x)$ in \mathcal{M} . Since \mathcal{M} is ω -saturated, every type $q(x) \in S_x(\Sigma)$ that contains $\Phi(x)$ is realized in \mathcal{M} and thus is of the form $\text{tp}_{\mathcal{M}}(b)$ for some such b . So to verify condition 7.5(2) it suffices to show that every such b realizes $p(x)$ in \mathcal{M} .

There is an isomorphism f from $\langle a \rangle_{\mathcal{M}}$ onto $\langle b \rangle_{\mathcal{M}}$ such that $f(a_i) = (b_i)$ for all $i = 1, \dots, n$. Indeed, f is given by setting $f(t^{\mathcal{M}}(a)) = t^{\mathcal{M}}(b)$ for every L -term $t(x)$. By assumption, $\langle a \rangle_{\mathcal{M}} \preceq \mathcal{M}$ and $\langle b \rangle_{\mathcal{M}} \preceq \mathcal{M}$. Therefore, f is an elementary map, and hence $\text{tp}_{\mathcal{M}}(b) = \text{tp}_{\mathcal{M}}(a) = p(x)$, as desired. \square

Next we determine the exact meaning of condition (a).

15.2. Proposition. *Let Σ be a set of L -sentences. The following are equivalent:*

- (1) *Every substructure of a model of Σ is an elementary substructure.*
- (2) *For any L -formula $\varphi(x, y)$, where $x = x_1, \dots, x_m$ with $m \geq 1$ and y is a single variable, there are L -terms $t_1(x), \dots, t_n(x)$ such that*

$$\Sigma \models \exists y \varphi(x, y) \rightarrow (\varphi(x, t_1(x)) \vee \dots \vee \varphi(x, t_n(x))).$$

- (3) *The same as (2), but only for quantifier-free formulas $\varphi(x, y)$.*

Proof. (2 \Rightarrow 1) Assume (2) and suppose $\mathcal{A} \subseteq \mathcal{M} \models \Sigma$. We apply the Tarski-Vaught Test 4.8 to prove $\mathcal{A} \preceq \mathcal{M}$. So suppose $\varphi(x, y)$ is an L -formula and $a \in A^m$ is such that $\mathcal{M} \models \exists y \varphi[a_1, \dots, a_m]$. Then we get L -terms $t_1(x), \dots, t_n(x)$ satisfying the condition in (2). Letting $b_j = t_j^{\mathcal{M}}(a)$ for $j = 1, \dots, n$, we see that $\mathcal{M} \models \varphi[a_1, \dots, a_m, b_j]$ for some j . But $b_j \in A$ for all j , so we have verified the condition in Theorem 4.8.

(1 \Rightarrow 3) We argue by contradiction, and begin by assuming that (3) fails. So we have a quantifier-free L -formula $\varphi(x, y)$, where $x = x_1, \dots, x_m$ with $m \geq 1$ and y is a single variable, such that for every finite sequence $t_1(x), \dots, t_n(x)$ of L -terms, the set of L -formulas

$$\{\neg \varphi(x, t_j(x)) \mid j = 1, \dots, n\} \cup \{\exists y \varphi(x, y)\}$$

is Σ -realizable. By the compactness theorem we get $\mathcal{M} \models \Sigma$ and $a \in M^m$ such that $\mathcal{M} \models \neg \varphi[a, t^{\mathcal{M}}(a)]$ holds for every L -term $t(x)$ while $\mathcal{M} \models \exists y \varphi[a]$. Let $\mathcal{A} = \langle a \rangle_{\mathcal{M}}$, which is nonempty (and is thus a substructure) because $m \geq 1$. Note that the elements of A are all of the form $t^{\mathcal{M}}(a)$ for some L -term $t(x)$. Since $\varphi(x, y)$ is quantifier-free, it follows that $\mathcal{A} \models \neg \varphi[a, b]$ for every $b \in A$. Hence $\mathcal{A} \models \neg \exists y \varphi(x, y)[a]$, and so \mathcal{A} is not an elementary substructure of \mathcal{M} .

To complete the proof we note that (3) implies that whenever $\varphi(x, y)$ is a quantifier-free L -formula, then $\exists y \varphi(x, y)$ is Σ -equivalent to the quantifier-free L -formula $(\varphi(x, t_1(x)) \vee \dots \vee \varphi(x, t_n(x)))$. From this it follows that Σ has QE. Hence (2) and (3) are equivalent. \square

Next we show how to obtain Σ' satisfying conditions (a),(b),(c) above; indeed, we satisfy (a) by ensuring that condition (3) in Proposition 15.2 holds in a very strong way that depends very simply on Σ .

Let L be a first order language and Σ a set of L -sentences. We say that Σ *has Skolem functions* if for every quantifier-free L -formula $\varphi(x_1, \dots, x_m, y)$ there is an m -ary function symbol f in L such that

$$\Sigma \models \exists y \varphi(x_1, \dots, x_m, y) \rightarrow \varphi(x_1, \dots, x_m, f(x_1, \dots, x_m)).$$

Note that if Σ has Skolem functions, then so does any extension of Σ in the same language. Note also that L must contain a constant symbol (apply the definition to the formula $\exists y(y = y)$).

15.3. Proposition (Skolemization). *Let L be a first order language. There exists a first order language $L^{sk} \supseteq L$ and a set $\text{sk}(L)$ of L^{sk} -sentences with the following properties:*

- (1) $\text{sk}(L)$ has Skolem functions;
- (2) Every L -structure has an expansion to a model of $\text{sk}(L)$;
- (3) L^{sk} has the same cardinality as L .

Proof. Inductively build an increasing sequence of first order languages $(L_k \mid k \in \mathbb{N})$ with $L_0 = L$ and an increasing sequence $(\Sigma_k \mid k \in \mathbb{N})$ with $\Sigma_0 = \emptyset$, such that Σ_k is a set of L_k -sentences for each $k \geq 1$. To obtain L_{k+1} from L_k we add a new function symbol f_φ for each quantifier-free L_k -formula $\varphi(x_1, \dots, x_m, y)$; to obtain Σ_{k+1} from Σ_k add all sentences of the form

$$\forall x_1 \dots \forall x_m (\exists y \varphi(x_1, \dots, x_m, y) \rightarrow \varphi(x_1, \dots, x_m, f_\varphi(x_1, \dots, x_m)))$$

where $\varphi(x_1, \dots, x_m, y)$ is a quantifier-free L_k -formula. Finally, let L^{sk} be the union of all the languages L_k and let $\text{sk}(L)$ be the union $\cup \{\Sigma_k \mid k \in \mathbb{N}\}$.

To prove (1), we note that each quantifier-free formula of L^{sk} is an L_k -formula for some k . Indeed, for each quantifier-free L_k -formula $\varphi(x, y)$, the needed sentence is in Σ_{k+1} .

To prove (2), we see that each L_k -structure that is a model of Σ_k can be expanded to an L_{k+1} -structure that is a model of Σ_{k+1} , using the axiom of choice to interpret each new function symbol appropriately.

To prove (3), we note that in constructing L_{k+1} from L_k , we added one new symbol for each quantifier-free L_k -formula; hence $\text{card}(L_{k+1}) = \text{card}(L_k)$ for each $k \in \mathbb{N}$. It follows that $\text{card}(L^{sk}) = \text{card}(L_0) = \text{card}(L)$. \square

15.4. Definition. Let Σ be a set of L -sentences, and suppose L^{sk} and $\text{sk}(L)$ satisfy the conditions in Proposition 15.3. Then the set $\Sigma' = \Sigma \cup \text{sk}(L)$ of L^{sk} -sentences will be called a *Skolemization* of Σ .

It is immediate from Proposition 15.3 that for any set Σ of L -sentences, the Skolemization $\Sigma' = \Sigma \cup \text{sk}(L)$ satisfies conditions (a),(b),(c) above.

15.5. Definition. Let Σ be a set of L -sentences and let Σ' be a Skolemization of Σ . If $\mathcal{M}' \models \Sigma'$ and $A \subseteq M'$, then $\langle A \rangle_{\mathcal{M}'}$ will be called the *Skolem hull* of A in \mathcal{M}' .

Note that even when $A = \emptyset$, the Skolem hull of A is nonempty, and hence is an elementary substructure of \mathcal{M}' , because L^{sk} always contains constant symbols.

We close this chapter with a summary of the key facts about Skolemizations that we will use later.

15.6. Corollary. *Let Σ be a set of L -sentences and let Σ' be a Skolemization of Σ . Then*

(a) Every substructure of a model of Σ' is an elementary substructure. As a consequence, Σ' has QE.

(b) Every model of Σ is the reduct to L of some model of Σ' .

(c) For every $\mathcal{M}' \models \Sigma'$ and every $A \subseteq M'$, the cardinality of the Skolem hull $\langle A \rangle_{\mathcal{M}'}$ is at most $\max(\text{card}(A), \text{card}(L))$.

(d) If $\mathcal{M}', \mathcal{N}' \models \Sigma'$ and $f: A \rightarrow B$ is a (not necessarily surjective) map that is elementary with respect to $\mathcal{M}', \mathcal{N}'$ (so $A \subseteq M'$ and $B \subseteq N'$), then f extends to an elementary embedding of the Skolem hull $\langle A \rangle_{\mathcal{M}'}$ into $\langle B \rangle_{\mathcal{N}'}$, and such an extension is unique.

16. INDISCERNIBLES

16.1. Definition. Let \mathcal{M} be an L -structure, $(I, <)$ a linear ordering, and $(a_i \mid i \in I)$ a family of elements of M . We say $(a_i \mid i \in I)$ is an *indiscernible family* (with respect to the ordering $<$ on I and the structure \mathcal{M}) if it has the following property:

$$\mathcal{M} \models \varphi[a_{i_1}, \dots, a_{i_n}] \Leftrightarrow \mathcal{M} \models \varphi[a_{j_1}, \dots, a_{j_n}]$$

for all L -formulas $\varphi(x_1, \dots, x_n)$ and all $i_1 < \dots < i_n$ and $j_1 < \dots < j_n$ from $(I, <)$.

If $(I, <) = (\mathbb{N}, <)$, then we call $(a_i \mid i \in \mathbb{N})$ with this property an *indiscernible sequence*.

Note that if $(a_i \mid i \in I)$ is an indiscernible family in \mathcal{M} and there exist distinct $i, j \in I$ for which $a_i = a_j$, then $a_i = a_j$ holds for all $i, j \in I$. (Apply the definition to the formula $x = y$.) Therefore, if $(a_i \mid i \in I)$ is non constant, then the function $i \mapsto a_i$ is 1-1 on I .

In the next proof we apply the classical Ramsey's Theorem from combinatorics. For each $n \geq 1$ and $A \subseteq \mathbb{N}$ we write $[A]^n$ for the set of all subsets α of A of cardinality $= n$. Given such α , we write it in the normal form $\alpha = \{i_1, \dots, i_n\}$ with $i_1 < \dots < i_n$.

Ramsey's Theorem. Let $C: [\mathbb{N}]^n \rightarrow F$ be a function, with $n \geq 1$ and F a finite set. There exists an infinite $H \subseteq \mathbb{N}$ such that C is constant on $[H]^n$.

16.2. Proposition. *Let Σ be a set of L -sentences with an infinite model. There exists a model \mathcal{M} of Σ and a non constant indiscernible sequence $(a_k \mid k \in \mathbb{N})$ in \mathcal{M} .*

Proof. Let L' be the language obtained from L by adding distinct new constants $(c_k \mid k \in \mathbb{N})$. Let Σ' be the set consisting of Σ together with all the sentences $c_k \neq c_l$ for distinct $k, l \in \mathbb{N}$, and all the sentences

$$\varphi(c_{i_1}, \dots, c_{i_n}) \leftrightarrow \varphi(c_{j_1}, \dots, c_{j_n})$$

where $\varphi(x_1, \dots, x_n)$ is any L -formula and $i_1 < \dots < i_n$ and $j_1 < \dots < j_n$ are in \mathbb{N} . (We will refer to these last sentences as the "indiscernibility axioms" in Σ' .)

If \mathcal{N} is a model of Σ' and \mathcal{M} is the reduct of \mathcal{N} to L , then $(c_k^{\mathcal{N}} \mid k \in \mathbb{N})$ is evidently a non constant indiscernible sequence in \mathcal{M} . Hence it suffices to show that Σ' has a model, which we do using the Compactness Theorem.

Let \mathcal{M} be any infinite model of Σ and let $\alpha: \mathbb{N} \rightarrow M$ be any 1-1 function. Let Δ be any finite subset of Σ' . Let K be the set of $k \in \mathbb{N}$ such that c_k occurs in some member of Δ . Let ψ_1, \dots, ψ_m be all the indiscernibility axioms that occur in Σ . We may assume that there exist L -formulas $\varphi_j(x_1, \dots, x_n)$ such that for each $j = 1, \dots, m$ the sentence ψ_j is logically equivalent to

$$\varphi_j(c_{i_1}, \dots, c_{i_n}) \leftrightarrow \varphi_j(c_{j_1}, \dots, c_{j_n}).$$

for some sequences $i_1 < \dots < i_n$ and $j_1 < \dots < j_n$ from K .

We now define a coloring function $C: [\mathbb{N}]^n \rightarrow \mathcal{P}(\{1, \dots, m\})$, to which we will apply Ramsey's Theorem. Namely, for each $i_1 < \dots < i_n$ from \mathbb{N} we take $C(i_1, \dots, i_n)$ to be the set of all $j \in \{1, \dots, m\}$ such that

$$\mathcal{M} \models \varphi_j[\alpha(i_1), \dots, \alpha(i_n)].$$

By Ramsey's Theorem there is an infinite set $H \subseteq \mathbb{N}$ such that C is constant on $[H]^n$; that is, $C(i_1, \dots, i_n) = C(j_1, \dots, j_n)$ whenever $i_1 < \dots < i_n$ and $j_1 < \dots < j_n$ are sequences from H .

Let $g: K \rightarrow H$ be any increasing function. We obtain a model of Δ by using \mathcal{M} to interpret the symbols of L and by interpreting c_k as $\alpha(g(k))$ for each $k \in K$. (The other c_k do not occur in Δ .) This shows that every finite subset of Σ' has a model. \square

16.3. Remark. It is sometimes useful to extend the previous result in the following way. Suppose $\varphi(x, y)$ is an L -formula and \mathcal{M} is a model of Σ with an infinite subset $A \subset M$ such that $\{(a, b) \in A^2 \mid \mathcal{M} \models \varphi[a, b]\}$ is a linear ordering on A . By taking an elementary extension if necessary, we may assume that there is a function $\alpha: \mathbb{N} \rightarrow A$ such that $\mathcal{M} \models \varphi[\alpha(k), \alpha(l)]$ for all $k < l$ in \mathbb{N} . Using this function in the above proof yields a non constant indiscernible sequence $(a_k \mid k \in \mathbb{N})$ in a model of Σ such that $\varphi(a_k, a_l)$ holds for $k, l \in \mathbb{N}$ if and only if $k < l$.

Likewise, if $\varphi(x)$ is an L -formula for which there is a model \mathcal{N} of Σ such that $\{a \in N \mid \mathcal{N} \models \varphi[a]\}$ is infinite, then the above proof can be modified to produce an indiscernible sequence $(a_k \mid k \in \mathbb{N})$ in a model \mathcal{M} of Σ such that $\mathcal{M} \models \varphi[a_k]$ for all $k \in \mathbb{N}$.

16.4. Definition. Let $(a_i \mid i \in I)$ be an indiscernible family in \mathcal{M} , with $(I, <)$ an infinite linear ordering, and let $(x_k \mid k \in \mathbb{N})$ be a fixed sequence of distinct variables. The *type* of $(a_i \mid i \in I)$ in \mathcal{M} is the set of all L -formulas $\varphi(x_1, \dots, x_n)$ such that $\mathcal{M} \models \varphi[a_{i_1}, \dots, a_{i_n}]$ for some (equivalently, every) $i_1 < \dots < i_n$ from $(I, <)$.

16.5. Proposition. *Let $(a_i \mid i \in I)$ be an indiscernible family in \mathcal{M} , with $(I, <)$ infinite, and let $(J, <)$ be another infinite linear ordering. There exists $\mathcal{N} \equiv \mathcal{M}$ and an indiscernible family $(b_j \mid j \in J)$ in \mathcal{N} having the same type as $(a_i \mid i \in I)$.*

Proof. An easy application of the Compactness Theorem. \square

EHRENFEUCHT-MOSTOWSKI MODELS

We now combine the construction of indiscernible families with the Skolem hulls discussed in the previous chapter to produce models that have a large group of automorphisms and models that realize few types, even over an arbitrary countable set of parameters from the model.

The starting point is a complete theory T in a first order language L . We suppose T has infinite models. Let T' be a Skolemization of T , in the language L' (Definition 15.4). Using Proposition 16.2, take any infinite

model \mathcal{M} of T' with a non constant indiscernible sequence $(a_k \mid k \in \mathbb{N})$ and let Φ be the type in \mathcal{M} of this sequence. We will refer to such a Φ as *the type of a non constant indiscernible sequence in a model of a Skolemization T' of T* .

Given such a Φ , we construct a model of T' for each infinite ordered set $(I, <)$, which we will denote as $\Phi(I, <)$. To do this, let \mathcal{N} be a model of T' and $(b_i \mid i \in I)$ an indiscernible family in \mathcal{N} that has type Φ . Take $\Phi(I, <)$ to be the Skolem hull of $\{b_i \mid i \in I\}$ in \mathcal{N} (Definition 15.5). Since Φ contains the formula $x_1 \neq x_2$, we may take $b_i = i$ with no loss of generality; that is, we may take $\Phi(I, <)$ to be generated by I as an L' -structure.

Note that this construction is canonical. Let $(b'_i \mid i \in I)$ be an indiscernible family in another model \mathcal{N}' of T' (with the same ordered index set) and suppose that $(b'_i \mid i \in I)$ has type Φ in \mathcal{N}' . Let $X = \{b_i \mid i \in I\}$ and $Y = \{b'_i \mid i \in I\}$ and consider the map $f: X \rightarrow Y$ defined by $f(b_i) = b'_i$ for all $i \in I$. Since these indiscernible families have the same type, f is an elementary map with respect to $\mathcal{N}, \mathcal{N}'$. Therefore, by Corollary 15.6(d), f extends to an isomorphism from $\langle X \rangle_{\mathcal{N}}$ onto $\langle Y \rangle_{\mathcal{N}'}$.

Using a similar idea we can make this construction functorial. Suppose $(I, <)$ and $(J, <)$ are infinite linear orderings, that $(b_i \mid i \in I)$ has type Φ in \mathcal{N} and that $(b'_j \mid j \in J)$ has type Φ in \mathcal{N}' . For each (strictly) order preserving function $F: I \rightarrow J$ we consider F as a map from $\{b_i \mid i \in I\}$ to $\{b'_j \mid j \in J\}$ by taking each b_i to $b'_{F(i)}$. Since these indiscernible sequences have the same type, this defines an elementary map with respect to $\mathcal{N}, \mathcal{N}'$. Therefore it extends in a unique way to an elementary embedding of the Skolem hulls, by Corollary 15.6(d). We denote this extension by $\Phi(F)$.

To summarize, our construction begins with the following data: (1) A complete L -theory T with infinite models and a Skolemization T' of T ; L' is the language of T' . (2) The type Φ of a non constant indiscernible sequence in a model of T' . The construction yields the following: for each infinite linear ordering $(I, <)$ it gives a model $\Phi(I, <)$ of T' ; this model is generated as an L' -structure by the set I , and $(I, <)$ itself is an indiscernible family in $\Phi(I, <)$. Moreover, for each order preserving map $F: (I, <) \rightarrow (J, <)$ of infinite linear orderings, we have an elementary embedding $\Phi(F)$ from $\Phi(I, <)$ into $\Phi(J, <)$ that extends F . Finally, this is functorial; that is, Φ maps the identity function on $(I, <)$ to the identity on $\Phi(I, <)$, and satisfies $\Phi(F) \circ \Phi(G) = \Phi(F \circ G)$ whenever F, G are order preserving maps that can be composed.

The models $\Phi(I, <)$ obtained via this construction are known as *Ehrenfeucht-Mostowski* models, in recognition of the logicians who first introduced these techniques.

We give two applications of this construction.

16.6. Corollary. *Let T be a complete L -theory with infinite models. For any cardinal κ such that $\text{card}(L) \leq \kappa$ there is a model \mathcal{M} of T such that M has cardinality κ and \mathcal{M} has 2^κ automorphisms (which is the maximum possible number).*

Proof. Let Φ be the type of a non constant indiscernible sequence in a model of a Skolemization T' of T . Recall that the language of T' has the same cardinality as the language of T . Let $I = \kappa \times \mathbb{Z}$ with the lexicographic ordering $(\alpha, m) < (\beta, n)$ iff $(\alpha < \beta$ or $(\alpha = \beta$ and $m < n))$. Note that $\Phi(I, <)$ has cardinality κ , since it is generated by a set of cardinality κ in a language of cardinality at most κ . Note also that $(I, <)$ has 2^κ many automorphisms. (For each function $\varphi: \kappa \rightarrow \{0, 1\}$, the map taking (α, n) to $(\alpha, n + \varphi(\alpha))$ is an automorphism of $(I, <)$.) Also, we know that each automorphism of $(I, <)$ extends to an automorphism of $\Phi(I, <)$. Therefore the reduct of $\Phi(I, <)$ to L is a model of T of cardinality κ that has 2^κ many automorphisms. \square

16.7. Corollary. *Let L be a countable first order language and let T be a complete L -theory with infinite models. For every uncountable cardinal κ , there is a model \mathcal{M} of T such that M has cardinality κ , but for every countable subset $A \subseteq M$ and every $n \geq 1$, only countably many n -types are realized in \mathcal{M}_A .*

Proof. Let Φ be the type of a non constant indiscernible sequence in a model of a Skolemization T' of T (so the language L' of T' is countable). Let $\mathcal{M} = \Phi(\kappa, <)$. Then the reduct of \mathcal{M} to L is a model of T of cardinality κ . We will show it satisfies the condition in this Corollary. Indeed, we will show that for every countable subset $A \subseteq M$ and every $n \geq 1$, only countably many n -types are realized in the L'_A -structure \mathcal{M}_A .

Let A be a countable subset of M . For each $a \in A$ there is an L' -term t_a and a finite sequence s_a from κ such that a is the value of $t_a(s_a)$ in \mathcal{M} . Let S be the subset of κ consisting of all ordinals that occur in s_a for some $a \in A$. Since A is countable and each s_a is finite, we see that S is countable.

Suppose X, Y are subsets of κ that contain S , and that $f: X \rightarrow Y$ is order preserving and is the identity on S . By Corollary 15.6(d), f has a unique extension to an elementary embedding from $\langle X \rangle_{\mathcal{M}}$ to $\langle Y \rangle_{\mathcal{M}}$, which we denote by \tilde{f} . Both of these Skolem hulls are elementary substructures of \mathcal{M} . Since f is the identity on S , its extension is the identity on $\langle S \rangle_{\mathcal{M}}$, which contains A . Therefore, for any tuple a_1, \dots, a_n in $\langle X \rangle_{\mathcal{M}}$, the types realized by (a_1, \dots, a_n) and by $(\tilde{f}(a_1), \dots, \tilde{f}(a_n))$ in \mathcal{M}_A are the same.

Suppose $\alpha_1, \dots, \alpha_n$ and β_1, \dots, β_n are finite sequences of the same length from κ ; we will say that these sequences are S -equivalent if there is an order preserving map that is the identity on S and takes α_j to β_j for each $j = 1, \dots, n$.

Since S is countable, there exists a countable subset X of κ such that any finite sequence in κ is S -equivalent to some sequence in X . (To S we need to add at most ω many ordinals from each cut in κ that is determined by S .) Note that $\langle X \rangle_{\mathcal{M}}$ is countable.

Let (a_1, \dots, a_n) be any n -tuple from M . For each $j = 1, \dots, n$, let t_j be an L' -term and s_j a finite sequence from κ such that a_j is the value of $t_j(s_j)$ in \mathcal{M} . Let $\alpha_1, \dots, \alpha_p$ be the ordinals that occur in the sequences s_1, \dots, s_n . Let β_1, \dots, β_p be S -equivalent to $\alpha_1, \dots, \alpha_p$ with $\beta_i \in X$ for all $i = 1, \dots, p$, and let f be an order preserving map that is the identity on S and takes α_i to β_i for each i . Then $\tilde{f}(a_j)$ is an element of $\langle X \rangle_{\mathcal{M}}$ for each $j = 1, \dots, n$. Moreover, as noted above, (a_1, \dots, a_n) and $(\tilde{f}(a_1), \dots, \tilde{f}(a_n))$ realize the same type in \mathcal{M}_A .

Since $\langle X \rangle_{\mathcal{M}}$ is countable, this shows that for any countable $A \subseteq M$, only countably many n -types are realized in the L'_A -structure \mathcal{M}_A . Hence the same is true if we replace \mathcal{M} by its reduct to L . \square

16.8. Corollary. *Let T be a complete theory in a countable language L , and let κ be an uncountable cardinal. Suppose T is κ -categorical. Then for every $\mathcal{M} \models T$, every countable $A \subseteq M$, and every $n \geq 1$, the space $S_n(A)$ of all n -types over A realized in models of T is countable.*

Proof. We begin with a countable set A in a model \mathcal{M} of T , which we take to be countable by the Downward Löwenheim-Skolem Theorem, and we fix $n \geq 1$. Arguing by contradiction, suppose $S_n(A)$ is uncountable. Since κ is uncountable, we may use the Compactness Theorem and the Löwenheim-Skolem Theorems to create a model \mathcal{N}_1 of T that is an elementary extension of \mathcal{M} , has cardinality κ , and is such that uncountably many types over A from $S_n(A)$ are realized in $\mathcal{N}_{1,A}$. On the other hand, Corollary 16.7 yields a model \mathcal{N}_2 of T having cardinality κ such that for all countable subsets B of \mathcal{N}_2 , only countably many types from $S_n(B)$ are realized in $\mathcal{N}_{2,B}$. Since T is κ -categorical, the models \mathcal{N}_1 and \mathcal{N}_2 are isomorphic, which is impossible. This contradiction completes the proof. \square

17. CANTOR-BENDIXSON RANK

Revision Note: This section still needs to be typed, following notes that Lou used in teaching Math 571 the last two times.

18. MORLEY RANK

Revision Note: this treatment of Morley Rank will be substantially revised to use what is planned for the previous chapter (CB rank) and to provide a basic foundation of the concept of stability. The ultimate goal of the chapter is still the development of Morley Rank and the proof of Morley's Theorem, following approximately the same line as is given here.

In this chapter T is a complete L -theory. Also, x and y denote finite tuples of variables, $x = x_1, \dots, x_m$ and $y = y_1, \dots, y_n$; we will write $\forall x$ instead of $\forall x_1 \dots \forall x_m$, and similarly for $\exists x$ and other strings of variables. If \mathcal{A} is an L -structure and $\varphi(x)$ is an $L(A)$ -formula, we will use the canonical interpretation of $\varphi(x)$ in \mathcal{A} (which corresponds to interpreting $\varphi(x)$ in $(\mathcal{A}, a)_{a \in A}$).

18.1. Definition. We define a relation " $RM_x(\mathcal{A}, \varphi(x)) \geq \alpha$ ", where $\mathcal{A} \models T$, $\varphi(x)$ is an $L(A)$ -formula, and α is an ordinal; the definition is by induction on α .

- (1) $RM_x(\mathcal{A}, \varphi(x)) \geq 0$ iff $\mathcal{A} \models \exists x \varphi(x)$;
- (2) $RM_x(\mathcal{A}, \varphi(x)) \geq \alpha + 1$ iff there is an elementary extension \mathcal{B} of \mathcal{A} and a sequence $(\varphi_k(x) \mid k \in \mathbb{N})$ of $L(B)$ -formulas such that
 - (a) $\mathcal{B} \models \forall x (\varphi_k(x) \rightarrow \varphi(x))$ for all $k \in \mathbb{N}$;
 - (b) $\mathcal{B} \models \forall x \neg(\varphi_k(x) \wedge \varphi_l(x))$ for all distinct $k, l \in \mathbb{N}$; and
 - (c) $RM_x(\mathcal{B}, \varphi_k(x)) \geq \alpha$ for all $k \in \mathbb{N}$;
- (3) for λ a limit ordinal, $RM_x(\mathcal{A}, \varphi(x)) \geq \lambda$ iff $RM_x(\mathcal{A}, \varphi(x)) \geq \alpha$ for all $\alpha < \lambda$.

18.2. Lemma. *Suppose $\mathcal{A} \models T$ and $\varphi(x)$ is an $L(A)$ -formula. Let S be the set of ordinals α such that $RM_x(\mathcal{A}, \varphi(x)) \geq \alpha$ holds. Then exactly one of the following alternatives holds:*

- (1) S is empty;
- (2) S is the class of all ordinals;
- (3) $S = \{\alpha \mid \alpha \text{ is an ordinal and } \alpha \leq \gamma\}$ for some ordinal γ .

Proof. The main point is to show, by induction on the ordinal α , that if $RM_x(\mathcal{A}, \varphi(x)) \geq \alpha$ and $\alpha > \beta \geq 0$, then $RM_x(\mathcal{A}, \varphi(x)) \geq \beta$. The key step in this induction is when α is a successor ordinal. Assume $RM_x(\mathcal{A}, \varphi(x)) \geq \alpha + 1$. Then we get an elementary extension \mathcal{B} of \mathcal{A} and a sequence $(\varphi_k(x) \mid k \in \mathbb{N})$ of $L(B)$ -formulas as in clause (2) of Definition 18.1. By the induction assumption we have that $RM_x(\mathcal{B}, \varphi_k(x)) \geq 0$ for each k , so each $\varphi_k(x)$ is satisfiable in \mathcal{B} . It follows that $\varphi(x)$ is also satisfiable in \mathcal{B} , and hence also in \mathcal{A} , since $\varphi(x)$ is an $L(A)$ -formula and $\mathcal{A} \preceq \mathcal{B}$. Therefore $RM_x(\mathcal{A}, \varphi(x)) \geq 0$. If $\beta + 1 < \alpha + 1$, then $\beta < \alpha$, so by the induction hypothesis we have that $RM_x(\mathcal{B}, \varphi_k(x)) \geq \beta$ for each k . Therefore $RM_x(\mathcal{A}, \varphi(x)) \geq \beta + 1$. Finally, if β is a limit ordinal $< \alpha$, we have (from the induction hypothesis)

$RM_x(\mathcal{A}, \varphi(x)) \geq \delta$ for every $\delta + 1 < \beta$. But this implies $RM_x(\mathcal{A}, \varphi(x)) \geq \beta$, since β is a limit ordinal.

From the first part of this proof, we have that S is an initial segment of the ordinals. If it is not the class of all ordinals, then it equals $\{\alpha \mid \alpha < \beta\}$, where β is the least ordinal not in S . If S is nonempty, β must be a successor ordinal, since S is closed upwards under limits by Definition 18.1(3). Condition (3) holds when γ is the predecessor of β . \square

The previous result allows us to define a *value* for “ $RM_x(\mathcal{A}, \varphi(x))$ ” in the following natural way; this is called the *Morley rank of $\varphi(x)$* .

18.3. Definition (Morley rank). Let \mathcal{A} be a model of T and let $\varphi(x)$ be an $L(\mathcal{A})$ -formula. If $RM_x(\mathcal{A}, \varphi(x)) \geq \alpha$ is false for all ordinals α , then we write $RM_x(\mathcal{A}, \varphi(x)) = -\infty$. If $RM_x(\mathcal{A}, \varphi(x)) \geq \alpha$ holds for all ordinals α , then we write $RM_x(\mathcal{A}, \varphi(x)) = +\infty$. Otherwise we define $RM_x(\mathcal{A}, \varphi(x))$ to be the greatest ordinal α for which $RM_x(\mathcal{A}, \varphi(x)) \geq \alpha$ holds. To indicate that $RM_x(\mathcal{A}, \varphi(x))$ is an ordinal we write $0 \leq RM_x(\mathcal{A}, \varphi(x)) < +\infty$ or we say that the $L(\mathcal{A})$ -formula $\varphi(x)$ is *ranked*.

18.4. Lemma. *Let \mathcal{A} is a model of T and $\varphi(x, y)$ an L -formula. If a is a finite tuple of elements of \mathcal{A} , then the value of $RM_x(\mathcal{A}, \varphi(x, a))$ depends only on $\text{tp}_{\mathcal{A}}(a)$.*

Proof. Let $\varphi(x, y)$ be an L -formula. It suffices to prove for each ordinal α that the truth of the relation “ $RM_x(\mathcal{A}, \varphi(x, a)) \geq \alpha$ ” only depends on $\text{tp}_{\mathcal{A}}(a)$. We do this by induction on α . The initial step $\alpha = 0$ and the induction step when α is a limit ordinal are trivial.

So, suppose the statement of the Lemma holds for all ordinals $\alpha < \beta + 1$. For $j = 1, 2$, let \mathcal{A}_j be a model of T and a_j a finite tuple from \mathcal{A}_j , and assume that $\text{tp}_{\mathcal{A}_1}(a_1) = \text{tp}_{\mathcal{A}_2}(a_2)$. We assume $RM_x(\mathcal{A}_1, \varphi(x, a_1)) \geq \beta + 1$ and need to prove $RM_x(\mathcal{A}_2, \varphi(x, a_2)) \geq \beta + 1$.

The assumption yields an elementary extension \mathcal{B}_1 of \mathcal{A}_1 , a sequence $(\varphi_k(x, z_k) \mid k \in \mathbb{N})$ of L -formulas and, for each $k \in \mathbb{N}$, a finite tuple b_k from \mathcal{B}_1 such that the formulas $(\varphi_k(x, b_k) \mid k \in \mathbb{N})$ witness that $RM_x(\mathcal{A}_1, \varphi(x, a_1)) \geq \beta + 1$. That is:

- (a) $\mathcal{B}_1 \models \forall x (\varphi_k(x, b_k) \rightarrow \varphi(x, a_1))$ for all $k \in \mathbb{N}$;
- (b) $\mathcal{B}_1 \models \forall x \neg (\varphi_k(x, b_k) \wedge \varphi_l(x, b_l))$ for all distinct $k, l \in \mathbb{N}$; and
- (c) $RM_x(\mathcal{B}_1, \varphi_k(x, b_k)) \geq \beta$ for all $k \in \mathbb{N}$.

Now let \mathcal{B}_2 be any ω -saturated elementary extension of \mathcal{A}_2 . We know that $\text{tp}_{\mathcal{B}_1}(a_1) = \text{tp}_{\mathcal{B}_2}(a_2)$. Since \mathcal{B}_2 is ω -saturated, we may construct inductively a sequence $(c_k \mid k \in \mathbb{N})$ of finite tuples from \mathcal{B}_2 such that for all $k \in \mathbb{N}$

$$\text{tp}_{\mathcal{B}_2}(a_2 c_0 \dots c_k) = \text{tp}_{\mathcal{B}_1}(a_1 b_0 \dots b_k).$$

It follows that

- (a) $\mathcal{B}_2 \models \forall x (\varphi_k(x, c_k) \rightarrow \varphi(x, a_2))$ for all $k \in \mathbb{N}$;
- (b) $\mathcal{B}_2 \models \forall x \neg (\varphi_k(x, c_k) \wedge \varphi_l(x, c_l))$ for all distinct $k, l \in \mathbb{N}$; and

(c) $RM_x(\mathcal{B}_2, \varphi_k(x, c_k)) \geq \beta$ for all $k \in \mathbb{N}$.

(Statements (a) and (b) are immediate; for (c) we use the induction hypothesis.) That is, the formulas $(\varphi_k(x, c_k) \mid k \in \mathbb{N})$ and the model \mathcal{B}_2 witness that $RM_x(\mathcal{A}_2, \varphi(x, a_2)) \geq \beta + 1$. \square

18.5. Notation. Let $\varphi(x, y)$ be an L -formula and a a tuple of elements of a model \mathcal{A} of T . We will write $RM(\varphi(x, a))$ in place of $RM_x(\mathcal{A}, \varphi(x, a))$, as long as the type $\text{tp}_{\mathcal{A}}(a)$ and the tuple of variables x are understood.

18.6. Lemma. *Let \mathcal{A} be an ω -saturated model of T and let $\varphi(x)$ be an $L(\mathcal{A})$ -formula. In applying Definition 18.1, in the clause defining $RM_x(\mathcal{A}, \varphi(x)) \geq \alpha + 1$ one may take the elementary extension \mathcal{B} to be \mathcal{A} itself.*

Proof. Exactly like the argument for the successor ordinal induction step in the proof of Lemma 18.4. \square

18.7. Lemma (Properties of Morley rank). *Let \mathcal{A} be a model of T and let $\varphi(x), \psi(x)$ be $L(\mathcal{A})$ -formulas.*

- (1) $RM(\varphi(x)) = 0$ iff the number of tuples $u \in \mathcal{A}$ for which $\mathcal{A} \models \varphi(u)$ is finite and > 0 ;
- (2) if $\mathcal{A} \models \forall x(\varphi(x) \rightarrow \psi(x))$, then $RM(\varphi(x)) \leq RM(\psi(x))$;
- (3) $RM(\varphi(x) \vee \psi(x)) = \max(RM(\varphi(x)), RM(\psi(x)))$;
- (4) if $0 \leq \beta < RM(\varphi(x)) < +\infty$, then there exists an elementary extension \mathcal{B} of \mathcal{A} and an $L(\mathcal{B})$ -formula $\chi(x)$ such that $\mathcal{B} \models \chi(x) \rightarrow \varphi(x)$ and $RM(\chi(x)) = \beta$.

Proof. (1) Note that if $\varphi(x)$ is satisfied in \mathcal{A} by infinitely many distinct values of x , say by $(u_k \mid k \in \mathbb{N})$, then the formulas $\varphi_k(x, u_k)$ that express $x = u_k$ have Morley rank ≥ 0 and thus witness that $\varphi(x)$ has Morley rank ≥ 1 .

(2) One proves by induction on the ordinal α that if $\mathcal{A} \models \forall x(\varphi(x) \rightarrow \psi(x))$ and $RM(\varphi(x)) \geq \alpha$, then $RM(\psi(x)) \geq \alpha$.

(3) Since $\mathcal{A} \models \forall x(\varphi(x) \rightarrow (\varphi(x) \vee \psi(x)))$, part (2) yields that $RM(\varphi(x)) \leq RM(\varphi(x) \vee \psi(x))$. Likewise, $RM(\psi(x)) \leq RM(\varphi(x) \vee \psi(x))$, so we have $\max(RM(\varphi(x)), RM(\psi(x))) \leq RM(\varphi(x) \vee \psi(x))$. To get the reverse inequality, one proves by induction on the ordinal α that $RM(\varphi(x) \vee \psi(x)) \geq \alpha$ implies $RM(\varphi(x)) \geq \alpha$ or $RM(\psi(x)) \geq \alpha$.

(4) Let \mathcal{F} be all formulas $\psi(x)$ with parameters from an elementary extension \mathcal{B} of \mathcal{A} such that $\psi(x) \rightarrow \varphi(x)$ is valid in \mathcal{B} . Suppose β is an ordinal that is not the Morley rank of any formula in \mathcal{F} . Therefore, if $\psi(x)$ is any such formula and $RM(\psi(x)) \geq \beta$, one has $RM(\psi(x)) \geq \beta + 1$. Now prove by induction on the ordinal α that if $\psi(x) \in \mathcal{F}$ and $RM(\psi(x)) \geq \beta$, then $RM(\psi(x)) \geq \alpha$. From this it follows that no ordinal $\geq \alpha$ is the Morley rank of a formula in \mathcal{F} . Statement (4) follows immediately from this result. \square

18.8. **Remark.** Part (4) of the previous result shows that the ordinals that occur as Morley ranks of formulas $\varphi(x, a)$ form an initial segment of the class of all ordinals. Moreover, the number of such ordinal ranks is $\leq \kappa$, where κ is the maximum of the number of types of finite tuples (over the empty set) in models of T and the cardinality of L . Since every type is a set of L -formulas, $\kappa \leq 2^{\text{card}(L)}$. Therefore, there exists an ordinal $\alpha_T < (2^{\text{card}(L)})^+$ such that the set of ordinal Morley ranks is exactly the set of ordinals $< \alpha_T$.

18.9. **Lemma** (Morley degree). *Let \mathcal{A} be a model of T and $\varphi(x)$ a ranked $L(\mathcal{A})$ -formula. There exists a finite bound on the integers k such that there exists an elementary extension \mathcal{B} of \mathcal{A} and $L(\mathcal{B})$ -formulas $(\varphi_j(x) \mid 0 \leq j < k)$ that satisfy the conditions*

- (a) $RM(\varphi_j(x)) = RM(\varphi(x))$ for all $j < k$;
- (b) $\mathcal{B} \models \forall x(\varphi_j(x) \rightarrow \varphi(x))$ for all $j < k$;
- (c) $\mathcal{B} \models \forall x \neg(\varphi_i(x) \wedge \varphi_j(x))$ for all distinct $i, j < k$.

Moreover, the maximum value of k depends only on $\text{tp}_{\mathcal{A}}(a)$. If \mathcal{A} is ω -saturated, a sequence of such formulas with maximal k can be found for \mathcal{B} equal to \mathcal{A} itself.

Proof. Let \mathcal{A} be a model of T , $\varphi(x, y)$ an L -formula, and a a tuple from \mathcal{A} ; assume $RM(\varphi(x, a)) = \alpha$ is an ordinal.

Suppose \mathcal{B} is an elementary extension of \mathcal{A} and $(\varphi_j(x) \mid 0 \leq j < k)$ is a sequence of $L(\mathcal{B})$ -formulas that satisfy conditions (a),(b),(c) in the statement of the Lemma. For each $0 \leq j < k$, let $\psi_j(x, y_j)$ be an L -formula and b_j a finite tuple from \mathcal{B} such that φ_j is $\psi(x, b_j)$. The fact that conditions (a),(b),(c) hold is equivalent to a property of the type realized by a, b_0, \dots, b_{k-1} in \mathcal{B} . (For clause (a) we apply Lemma 18.4.) Therefore, the existence of \mathcal{B} and k such formulas $(\varphi_j(x) \mid 0 \leq j < k)$ depends only on the type realized by a in \mathcal{A} . Moreover, if such a sequence of k formulas exists for *some* elementary extension of \mathcal{A} , and if \mathcal{B} is any *specific* ω -saturated elementary extension of \mathcal{A} , then we can find such a sequence of k formulas for \mathcal{B} . (Just realize the type of the parameter sequence b_0, \dots, b_{k-1} over a in \mathcal{B} .)

Therefore, in proving that the maximum value of k exists, we may assume that \mathcal{A} is ω -saturated and restrict ourselves to considering sequences of $L(\mathcal{A})$ -formulas $(\varphi_j(x) \mid 0 \leq j < k)$.

Let Λ be the set of finite sequences from $\{0, 1\}$; for $\sigma, \tau \in \Lambda$ we write $\sigma \subseteq \tau$ to mean that τ is an extension of σ . If $\sigma \subseteq \tau$ and the length of τ is exactly one more than the length of σ , then we call τ an *immediate* extension of σ and write τ as $\sigma 0$ or $\sigma 1$ to indicate which is the last entry in the sequence τ . In our construction we use the fact that (Λ, \subseteq) is a well-founded partial ordering whose least element is the empty sequence (denoted \emptyset).

We build a nonempty subset S of Λ that is closed under restriction ($\sigma \subseteq \tau \in S$ implies $\sigma \in S$); further, for each $\sigma \in S$ we define an $L(\mathcal{A})$ -formula

φ_σ of Morley rank α . This is done by induction on the binary tree (Λ, \subseteq) . For the basis step, we put $\emptyset \in S$ and define $\varphi_\emptyset = \varphi(x, a)$. For the induction step, consider $\sigma \in \Lambda$ and suppose we have dealt with all $\tau \in \Lambda$ that are shorter than σ . If $\sigma \notin S$, then neither immediate extension of σ gets put into S . If $\sigma \in S$, there are two cases. First, suppose there is an $L(A)$ -formula $\psi(x)$ such that both $\varphi_\sigma \wedge \psi$ and $\varphi_\sigma \wedge \neg\psi$ have Morley rank equal to α . In that case we choose such a formula ψ , put both immediate extensions of σ into S , and set $\varphi_{\sigma 0} = \varphi_\sigma \wedge \psi$ and $\varphi_{\sigma 1} = \varphi_\sigma \wedge \neg\psi$. Second, if no such ψ exists, then neither immediate extension of σ gets put into S . (Note that in this latter case, for every $L(A)$ -formula $\psi(x)$, one of the formulas $\varphi_\sigma \wedge \psi$ and $\varphi_\sigma \wedge \neg\psi$ has Morley rank $= \alpha$ and the other one has Morley rank $< \alpha$. (See Lemma 18.7(3).)

Next we prove that S is finite. Otherwise, by König's Lemma, there is an infinite branch in S . That is, there exists a function $f: \mathbb{N} \rightarrow \{0, 1\}$ such that for all $k \in \mathbb{N}$ the sequence $f|k = f(0), \dots, f(k-1)$ is in S . For all $n \geq 1$, let $\chi_n(x)$ be the $L(A)$ -formula $\varphi_{f|n} \wedge \neg\varphi_{f|n+1}$. It is easy to check that the sequence $(\chi_n \mid n \geq 1)$ witnesses that $RM(\varphi(x, a)) \geq \alpha + 1$, which is a contradiction.

Let S_0 denote the set of leaves of the finite binary tree S ; that is, S_0 contains those $\sigma \in S$ such that no proper extension of σ is in S . Then S is exactly the set of $\sigma \in \Lambda$ such that some extension of σ is in S_0 . Note that if σ, τ are distinct elements of S_0 then there is a sequence η such that one of σ, τ is an extension of $\eta 0$ and the other one is an extension of $\eta 1$. Hence φ_σ and φ_τ are contradictory in \mathcal{A} . Our construction of S ensures that if $\sigma \in S$ is not in S_0 , then both $\sigma 0$ and $\sigma 1$ are in S and, moreover, φ_σ is logically equivalent to $\varphi_{\sigma 0} \vee \varphi_{\sigma 1}$. A simple argument shows that $\varphi(x, a) = \varphi_\emptyset$ is logically equivalent to the disjunction of all formulas φ_σ with σ ranging over S_0 .

Let $d = \text{card}(S_0)$ and let $\chi_0, \dots, \chi_{d-1}$ enumerate the formulas φ_σ with $\sigma \in S_0$. Our construction has ensured that $(\chi_i \mid 0 \leq i < d)$ satisfies conditions (a),(b),(c) in the statement of the Lemma. Moreover, in \mathcal{A} the formula $\varphi(x, a)$ is equivalent to the disjunction of χ_i for $0 \leq i < d$. Suppose now that $(\varphi_j(x) \mid 0 \leq j < e)$ is any sequence of $L(A)$ -formulas that satisfy conditions (a),(b),(c) and that $e > d$. Consider any i with $0 \leq i < d$ and distinct r, s with $0 \leq r, s < e$. By our construction, χ_i is φ_σ for some σ that is a leaf in S . Using Lemma 18.7 and the fact that φ_r and φ_s are contradictory in \mathcal{A} , it follows that at most one of $\chi_i \wedge \varphi_r$ and $\chi_i \wedge \varphi_s$ can have Morley rank $= \alpha$. Since $d < e$, the pigeonhole principle implies that there must exist at least one value of r with $0 \leq r < k$ such that $\chi_j \wedge \varphi_r$ has Morley rank $< \alpha$ for all $0 \leq j < d$. As noted above, $\varphi(x, a)$ is equivalent to the disjunction of χ_i for $0 \leq i < d$. Therefore, φ_r is equivalent to the disjunction of the formulas $\chi_i \wedge \varphi_r$ with $0 \leq i < d$. Using Lemma 18.7(3) it follows that φ_r itself has Morley rank $< \alpha$. This contradiction proves the Lemma. \square

18.10. **Definition.** Given a ranked $L(A)$ -formula $\varphi(x)$, the greatest integer whose existence is proved in Lemma 18.9 is called the *Morley degree* of $\varphi(x)$ and it is denoted $dM(\varphi(x))$.

18.11. **Lemma** (Properties of Morley degree). *Let \mathcal{A} be an ω -saturated model of T and let $\varphi(x), \psi(x)$ be $L(A)$ -formulas.*

(1) *If $\varphi(x)$ is ranked and $dM(\varphi(x)) = d$, with the latter statement witnessed by the sequence $(\varphi_j \mid 0 \leq j < d)$ of $L(A)$ -formulas, then each $\varphi_j(x)$ has Morley degree 1.*

(2) *if $0 \leq RM(\varphi(x)) = RM(\psi(x)) < +\infty$ and $\mathcal{A} \models \varphi(x) \rightarrow \psi(x)$, then $dM(\varphi(x)) \leq dM(\psi(x))$;*

(3) *if $0 \leq RM(\varphi(x)) = RM(\psi(x)) < +\infty$ then $dM(\varphi(x) \vee \psi(x)) \leq dM(\varphi(x)) + dM(\psi(x))$, with equality if $\mathcal{A} \models \neg(\varphi(x) \wedge \psi(x))$;*

(4) *if $0 \leq RM(\psi(x)) < RM(\varphi(x)) < +\infty$, then $dM(\varphi(x) \vee \psi(x)) = dM(\varphi(x))$.*

Proof. (1) Lemmas 18.7 and 18.9 ensure that each $\varphi_j(x)$ has a Morley degree. If for some j the formula $\varphi_j(x)$ has Morley degree > 1 , then there exist two $L(A)$ -formulas witnessing that fact. Replacing $\varphi_j(x)$ by them in the sequence $(\varphi_j \mid 0 \leq j < d)$ witnesses that $\varphi(x)$ has Morley degree $\geq d + 1$, a contradiction.

(2) Any sequence of $L(A)$ -formulas of length d , witnessing that d is the Morley degree of $\varphi(x)$, will witness that the Morley degree of $\psi(x)$ is $\geq d$.

(3) Let $\alpha = RM(\varphi(x)) = RM(\psi(x))$; Lemma 18.7(3) yields $RM(\varphi(x) \vee \psi(x)) = \alpha$. Suppose $d = dM(\varphi(x) \vee \psi(x))$, witnessed by the sequence $(\chi_j(x) \mid 0 \leq j < d)$ of $L(A)$ -formulas. For each j , at least one of the formulas $\chi_j(x) \wedge \varphi(x)$ and $\chi_j(x) \wedge \psi(x)$ has Morley rank $= \alpha$ by Lemma 18.7(3). Let k be the number of values of j for which $RM(\chi_j(x) \wedge \varphi(x)) = \alpha$, and arrange the formulas so that this occurs for $0 \leq j < k$. Therefore $RM(\chi_j(x) \wedge \psi(x)) = \alpha$ for $k \leq j < d$. These sequences witness that $k \leq dM(\varphi(x))$ and $d - k \leq dM(\psi(x))$ and hence $d \leq dM(\varphi(x)) + dM(\psi(x))$.

Now suppose $\mathcal{A} \models \neg(\varphi(x) \wedge \psi(x))$. Let $(\varphi_j(x) \mid 0 \leq j < k)$ witness that $dM(\varphi(x)) = k$ and $(\psi_j(x) \mid 0 \leq j < l)$ witness that $dM(\psi(x)) = l$. Then $(\varphi_0(x), \dots, \varphi_{k-1}(x), \psi_0(x), \dots, \psi_{l-1}(x))$ witnesses that $dM(\varphi(x) \vee \psi(x)) \geq k + l$. Combined with the first part of the proof, this shows

$$dM(\varphi(x)) + dM(\psi(x)) \geq k + l \geq dM(\varphi(x)) + dM(\psi(x)).$$

(4) Argue as in the first part of the proof of (3); note that since $RM(\psi(x)) < \alpha$, one has $k = d$. \square

18.12. **Lemma.** *Let $\mathcal{A} \models T$ and $C \subseteq A$. Let $p(x)$ be a type (in $L(C)$) of a finite tuple that is consistent with $\text{Th}((\mathcal{A}, a)_{a \in C})$. Assume that some formula in $p(x)$ is ranked. Then there exists a formula $\varphi_p(x)$ in $p(x)$ that determines $p(x)$ in the following sense:*

$$p(x) \text{ consists exactly of the } L(C)\text{-formulas } \psi(x) \text{ such that } RM(\psi(x) \wedge \varphi_p(x)) = RM(\varphi_p(x)) \text{ and } dM(\psi(x) \wedge \varphi_p(x)) = dM(\varphi_p(x)).$$

Indeed, such a formula can be obtained by taking $\varphi_p(x)$ to be a formula $\varphi(x)$ in $p(x)$ with least possible Morley rank and degree, in lexicographic order.

Proof. Choose $\varphi_p(x) \in p(x)$ as specified in the last sentence of the Lemma. That is, $\varphi_p(x)$ is a formula in $p(x)$ of least possible Morley rank and, among members of $p(x)$ having that rank, $dM(\varphi_p(x))$ is least possible.

If $\psi(x)$ is any formula in $p(x)$, then also $\psi(x) \wedge \varphi_p(x) \in p(x)$ and hence $RM(\psi(x) \wedge \varphi_p(x)) \geq RM(\varphi_p(x))$ by our choice of $\varphi_p(x)$. Hence $RM(\psi(x) \wedge \varphi_p(x)) = RM(\varphi_p(x))$ by Lemma 18.7. A similar argument using Lemma 18.11 proves $dM(\psi(x) \wedge \varphi_p(x)) = dM(\varphi_p(x))$.

Conversely, suppose $\psi(x)$ is any $L(C)$ -formula with $RM(\psi(x) \wedge \varphi_p(x)) = RM(\varphi_p(x))$ and $dM(\psi(x) \wedge \varphi_p(x)) = dM(\varphi_p(x))$. By way of contradiction, suppose $\psi(x) \notin p(x)$, in which case $\neg\psi(x) \in p(x)$. But then $RM(\neg\psi(x) \wedge \varphi_p(x)) = RM(\varphi_p(x))$. In that case Lemma 18.11 yields $dM(\varphi_p(x)) \geq dM(\psi(x) \wedge \varphi_p(x)) + dM(\neg\psi(x) \wedge \varphi_p(x)) > dM(\psi(x) \wedge \varphi_p(x))$, which is a contradiction. \square

18.13. Definition. Let $\mathcal{A} \models T$ and $C \subseteq A$. Let $p(x)$ be a type (in $L(C)$) of a finite tuple that is consistent with $\text{Th}((\mathcal{A}, a)_{a \in C})$. We define $RM(p(x))$ to be the least Morley rank of a formula in $p(x)$. If some formula in $p(x)$ is ranked, we define $dM(p(x))$ to be the least Morley degree of a formula $\varphi(x)$ in $p(x)$ that satisfies $RM(\varphi(x)) = RM(p(x))$.

18.14. Definition. Let λ be an infinite cardinal. We say T is λ -stable if for every model \mathcal{A} of T and every $C \subseteq A$ of cardinality $\leq \lambda$, at most λ many types (in $L(C)$) of finite tuples are consistent with $\text{Th}((\mathcal{A}, a)_{a \in C})$.

18.15. Theorem. Let L be countable. The following conditions are equivalent:

- (1) T is ω -stable;
- (2) for any $\mathcal{A} \models T$ and any $L(\mathcal{A})$ -formula $\varphi(x)$, $RM(\varphi(x)) < +\infty$;
- (3) T is λ -stable for every $\lambda \geq \omega$.

Proof. (1 \Rightarrow 2): We prove the contrapositive. Let \mathcal{A} be an ω -saturated model of T . Every Morley rank of a formula with parameters from some model of T is the Morley rank of some $L(\mathcal{A})$ -formula. Hence there exists an ordinal α_T such that for any formula $\varphi(x)$ with parameters from a model of T , if $RM(\varphi(x)) \geq \alpha_T$ then $RM(\varphi(x)) = +\infty$. (In fact, by 18.7(4), α_T can be chosen so that these Morley ranks are exactly the ordinals $< \alpha_T$, but we do not need that here.)

Suppose $\varphi(x)$ is any $L(\mathcal{A})$ -formula whose Morley rank is $+\infty$. Then $RM(\varphi(x)) \geq \alpha_T + 1$, so there exist two $L(\mathcal{A})$ -formulas $\psi_1(x), \psi_2(x)$ that are contradictory in \mathcal{A} and have Morley rank $\geq \alpha_T$, and such that $\psi_j(x) \rightarrow \varphi(x)$ is valid in \mathcal{A} for $j = 1, 2$. (Indeed, there is a whole infinite sequence of such formulas.) By choice of α_T , this ensures that $\psi_1(x), \psi_2(x)$ both have Morley rank $+\infty$. Using Lemma 18.7 we see that $\varphi(x) \wedge \psi_1(x)$ and $\varphi(x) \wedge \neg\psi_1(x)$ both have Morley rank $+\infty$.

As in the proof of Lemma 18.9 we let Λ be the set of finite sequences from $\{0, 1\}$ partially ordered by extension, and we use the other notation established in that proof. Suppose there exists a formula $\varphi(x)$ with parameters from a model \mathcal{A} of T whose Morley rank equals $+\infty$. Without loss of generality we may take \mathcal{A} to be ω_1 -saturated. Using the argument in the previous paragraph inductively, we may construct a family $(\varphi_\sigma(x) \mid \sigma \in \Lambda)$ of $L(\mathcal{A})$ -formulas such that each $\varphi_\sigma(x)$ has Morley rank $+\infty$, $\varphi_\emptyset(x)$ is $\varphi(x)$, and for each $\sigma \in \Lambda$ there is a formula $\psi(x)$ such that $\varphi_{\sigma 0} = \varphi_\sigma(x) \wedge \psi(x)$ and $\varphi_{\sigma 1} = \varphi_\sigma(x) \wedge \neg\psi(x)$. Let C be the set of all parameters from \mathcal{A} that occur in $\varphi_\sigma(x)$ for some $\sigma \in \Lambda$; note that C is countable. For each function $f: \mathbb{N} \rightarrow \{0, 1\}$, let $\Sigma_f(x)$ be the set of formulas $\{\varphi_{f|n} \mid n \in \mathbb{N}\}$. Our construction ensures that each $\Sigma_f(x)$ is satisfiable in $(\mathcal{A}, a)_{a \in C}$. Moreover, if f, g are distinct functions, then $\Sigma_f(x)$ and $\Sigma_g(x)$ are contradictory in \mathcal{A} ; indeed, if $n \in \mathbb{N}$ is the least integer with $f(n) \neq g(n)$ and $\sigma = f|n = g|n$, then one of these sets contains $\varphi_{\sigma 0}(x)$ and the other one contains $\varphi_{\sigma 1}(x)$ and these two formulas are contradictory in $(\mathcal{A}, a)_{a \in C}$. For each function f , let $p_f(x)$ be the type realized in $(\mathcal{A}, a)_{a \in C}$ by some specific realization of $\Sigma_f(x)$. Then $\{p_f(x) \mid f: \mathbb{N} \rightarrow \{0, 1\}\}$ is a family of uncountably many types of finite tuples consistent with $\text{Th}((\mathcal{A}, a)_{a \in C})$. This contradicts (1).

(2 \Rightarrow 3): Let $\mathcal{A} \models T$ and $C \subseteq A$ with $\text{card}(C) \leq \lambda$. We need to show that there are at most λ many types $p(x)$ (in $L(C)$) of a finite tuple that are consistent with $\text{Th}(\mathcal{A}, a)_{a \in C}$. Given such $p(x)$, condition (2) ensures that Lemma 18.12 applies, so that $p(x)$ is determined by a formula $\varphi_p(x)$ in the way described there. Since there are at most λ many $L(C)$ -formulas (here we use the assumption that L is countable), there are at most λ many such types $p(x)$. \square

We complete this chapter by showing that every uncountable model of an ω -stable theory in a countable language contains nonconstant sequences of ordered indiscernibles, even when names for moderately large sets of parameters are added to the language. First we need some notation and a technical lemma.

18.16. Notation. Let \mathcal{A} be an L -structure. If b is a tuple in A and B is any subset of A , we will write $\text{tp}_{\mathcal{A}}(b/B)$ for the type (in $L(B)$) realized by b in $(\mathcal{A}, a)_{a \in B}$.

18.17. Lemma. *Assume T is ω -stable. Suppose $\mathcal{A} \models T$ and $C \subseteq A$. Let $\varphi(x)$ be a ranked $L(C)$ -formula, and set $(\alpha, d) = (RM(\varphi(x)), dM(\varphi(x)))$. Suppose $(a_k \mid k \in \mathbb{N})$ is a sequence of finite tuples (of the same length) from A and for each $k \in \mathbb{N}$ define $p_k(x) = \text{tp}_{\mathcal{A}}(a_k/C \cup \{a_0, \dots, a_{k-1}\})$. Assume that $\mathcal{A} \models \varphi(a_k)$ and $(RM(p_k(x)), dM(p_k(x))) = (\alpha, d)$, for all $k \in \mathbb{N}$. Then $(a_k \mid k \in \mathbb{N})$ is an indiscernible sequence in $(\mathcal{A}, a)_{a \in C}$.*

Proof. We prove by induction on $n \in \mathbb{N}$ that whenever $i_0 < \dots < i_n$ are in \mathbb{N} , $\text{tp}_{\mathcal{A}}(a_{i_0} \dots a_{i_n}/C) = \text{tp}_{\mathcal{A}}(a_0 \dots a_n/C)$.

In the basis case, $n = 0$. Take any $i \in \mathbb{N}$. Since $\varphi(x) \in \text{tp}_{\mathcal{A}}(a_i/C)$ we have $(RM(\text{tp}_{\mathcal{A}}(a_i/C)), dM(\text{tp}_{\mathcal{A}}(a_i/C))) \leq (\alpha, d)$ lexicographically. On the other hand, $\text{tp}_{\mathcal{A}}(a_i/C) \subseteq p_i(x)$; thus our assumptions yield $(RM(\text{tp}_{\mathcal{A}}(a_i/C)), dM(\text{tp}_{\mathcal{A}}(a_i/C))) \geq (\alpha, d)$. It follows that $(RM(\text{tp}_{\mathcal{A}}(a_i/C)), dM(\text{tp}_{\mathcal{A}}(a_i/C))) = (\alpha, d)$. Lemma 18.12 implies $\text{tp}_{\mathcal{A}}(a_i/C) = \text{tp}_{\mathcal{A}}(a_0/C)$.

For the induction step $n > 0$. Consider any $i_0 < \dots < i_n$ from \mathbb{N} . As argued in the previous paragraph, $\text{tp}_{\mathcal{A}}(a_{i_n}/C \cup \{a_{i_0}, \dots, a_{i_{n-1}}\})$ has Morley rank α and degree d , as does $p_n(x)$. Both of these types contain the formula $\varphi(x)$. Applying Lemma 18.12 we conclude that for any $L(C)$ -formula $\psi(x, y_0, \dots, y_{n-1})$ we have:

- (i) $\mathcal{A} \models \psi[a_n, a_0, \dots, a_{n-1}]$ if and only if the formula $\varphi(x) \wedge \psi(x, a_0, \dots, a_{n-1})$ has Morley rank α and degree d ;
- (ii) $\mathcal{A} \models \psi[a_{i_n}, a_{i_0}, \dots, a_{i_{n-1}}]$ if and only if the formula $\varphi(x) \wedge \psi(x, a_{i_0}, \dots, a_{i_{n-1}})$ has Morley rank α and degree d .

The induction hypothesis states that $\text{tp}_{\mathcal{A}}(a_{i_0} \dots a_{i_{n-1}}/C) = \text{tp}_{\mathcal{A}}(a_0 \dots a_{n-1}/C)$. This implies that the right hand sides of statements (i) and (ii) are equivalent to each other. Therefore we conclude $\text{tp}_{\mathcal{A}}(a_{i_0} \dots a_{i_n}/C) = \text{tp}_{\mathcal{A}}(a_0 \dots a_n/C)$ as claimed. \square

18.18. Proposition. *Assume T is an ω -stable L -theory with L countable. Suppose $\mathcal{A} \models T$ and $C \subseteq A$. Assume that A is uncountable and $\text{card}(C) < \text{card}(A)$. Then there exists a nonconstant sequence of ordered indiscernibles in $(\mathcal{A}, a)_{a \in C}$.*

Proof. We may assume C is infinite. Let $\lambda = \text{card}(C)$. We begin an inductive construction by noting that the formula $x = x$ is satisfied by $> \lambda$ many elements of A in \mathcal{A} . Choose an $L(A)$ formula $\varphi(x)$ that is satisfied by $> \lambda$ many elements of A in \mathcal{A} and has the minimum possible Morley rank and degree; say these are (α, d) . Note that $\alpha > 0$ since $\varphi(x)$ is satisfied by infinitely many elements. By adding finitely many elements to C we may assume that $\varphi(x)$ is an $L(C)$ -formula.

We will construct a sequence $(a_k \mid k \in \mathbb{N})$ of elements of A that satisfy $\varphi(x)$ in \mathcal{A} such that for all $k \in \mathbb{N}$, the Morley rank and degree of $\text{tp}_{\mathcal{A}}(a_k/C \cup \{a_0, \dots, a_{k-1}\})$ is exactly (α, d) .

First we obtain a_0 with this property. If no such element of A exists, we have $(RM(\text{tp}_{\mathcal{A}}(a/C)), dM(\text{tp}_{\mathcal{A}}(a/C))) < (\alpha, d)$ for all $a \in A$. For each $a \in A$ we therefore have an $L(C)$ -formula $\psi_a(x)$ that is satisfied by a and has $(RM(\psi_a(x)), dM(\psi_a(x))) < (\alpha, d)$. There are at most λ such formulas while there are $> \lambda$ many values of a . Therefore there is a set of $> \lambda$ many values of a for which $\psi_a(x)$ is the same formula $\psi(x)$. But this contradicts the minimum choice of (α, d) . This proves a_0 exists.

For the induction step we have a_0, \dots, a_{k-1} and seek a_k . This is handled by the same argument as in the previous paragraph, replacing C by $C \cup \{a_0, \dots, a_{k-1}\}$.

Finally, by Lemma 18.17 the resulting sequence $(a_k \mid k \in \mathbb{N})$ is indiscernible over C in \mathcal{A} . \square

EXERCISES

18.19. Let \mathcal{A} be an ω -saturated L -structure and let $X \subseteq A^m$ be A -definable in \mathcal{A} . Assume that $0 \leq \beta < \alpha = RM(X) < +\infty$. (Since $L(A)$ -formulas that are equivalent in $(\mathcal{A}, a)_{a \in A}$ have the same Morley rank by Lemma 17.7(2), we may refer without ambiguity to the Morley rank of a definable set.) Show that there is an infinite family $(Y_n \mid n \in \mathbb{N})$ of pairwise disjoint A -definable subsets of X such that $RM(Y_n) = \beta$ for all $n \in \mathbb{N}$.

18.20. Let \mathcal{A} be an L -structure, let $\varphi(x)$ be an $L(A)$ -formula and let $t(x)$ an L -term, with $x = x_1, \dots, x_m$. Show that the formulas $(\varphi(x) \wedge y = t(x))$ and $\varphi(x)$ have the same Morley rank. (Here y is a single, new variable. The Morley rank of $\varphi(x)$ is taken with respect to the variables x and the Morley rank of $(\varphi(x) \wedge y = t(x))$ is taken with respect to the variables x, y .)

18.21. Let L be the language whose nonlogical symbols consist of a constant symbol e , a unary function symbol i , and a binary function symbol p . Let G be a group, considered as an L -structure by interpreting e as the identity element, $i(g)$ as the inverse of g , and $p(g, h)$ as the product of g and h in G . Assume that the theory of G is ω -stable. Show that G satisfies the descending chain condition on G -definable subgroups. That is, if $G \supseteq H_0 \supseteq H_1 \supseteq \dots$ are subgroups of G and each H_n is G -definable, show that the sequence $(H_n \mid n \in \mathbb{N})$ is eventually constant. (Hint: use cosets to show that the Morley ranks of the sets H_n would otherwise yield an infinite, strictly decreasing sequence of ordinals.)

18.22. Let L be the language whose nonlogical symbols are the unary predicate symbols P_1, \dots, P_n . Let T be the L -theory whose axioms express that the sets P_1, \dots, P_n are infinite and that they form a partition of the underlying set of the L -structure being considered. Show that T admits QE and is complete. Show that the formula $x = x$ has Morley rank 1 and Morley degree n in models of T .

18.23. Let L be a countable language and let T be a complete L -theory with infinite models. Suppose that for every model \mathcal{A} of T and every countable $C \subseteq A$, the space of 1-types $S_1(C)$ is countable. Show that T is ω -stable.

18.24. Let L be a countable language and let T be a complete L -theory with infinite models. Suppose that for every model \mathcal{A} of T and every $L(A)$ -formula $\psi(x)$ in which x is a single variable, one has $RM(\psi(x)) < +\infty$. Show that for every model \mathcal{A} of T , every $n \geq 1$, and every $L(A)$ -formula $\varphi(x_1, \dots, x_n)$, one has $RM(\varphi(x_1, \dots, x_n)) < +\infty$. (Hint: use the preceding exercise together with a careful reading of the proof of Theorem 17.15.)

The goal of the rest of this chapter is to prove the following important theorem. The ideas developed by Morley for its proof had a strong influence on the development of pure model theory during the last decades of the 20th century.

For the rest of this chapter L is a countable first order language and T is a complete L -theory with infinite models.

18.25. Theorem (Morley's Theorem). *If T is κ -categorical for one uncountable cardinal κ , then T is κ -categorical for all uncountable κ .*

To prove this theorem, we make use of all the tools that were developed in the last few chapters. In particular, Morley rank plays a key role in the proof. Its use is justified by the following result.

18.26. Proposition. *If T is κ -categorical for some uncountable κ , then T is ω -stable. Therefore, for every satisfiable formula $\varphi(x)$ with parameters from some model of T , $RM(\varphi(x))$ is an ordinal and so $dM(\varphi(x))$ is defined.*

Proof. Let κ be an uncountable cardinal and suppose that T is κ -categorical. Let \mathcal{A} be the unique (up to isomorphism) model of T with $\text{card}(\mathcal{A}) = \kappa$. By the Löwenheim-Skolem Theorems and the uniqueness of \mathcal{A} , every model of T of cardinality $\leq \omega_1$ is isomorphic to an elementary substructure of \mathcal{A} .

Arguing by contradiction, suppose T is not ω -stable; that is, there is a model \mathcal{B} of T and a countable subset $C \subseteq B$ such that uncountably many types of finite tuples are consistent with $\text{Th}((\mathcal{B}, b)_{b \in C})$. By passing to a different model, we may assume that uncountably many such types are actually realized in $(\mathcal{B}, b)_{b \in C}$ and that $\text{card}(B) = \omega_1$.

Putting the two previous paragraphs together, we may assume that we have a countable $C \subseteq A$ such that uncountably many types of finite tuples are realized in $(\mathcal{A}, a)_{a \in C}$.

However, Corollary 16.7 yields a model \mathcal{A}' of T having cardinality κ and satisfying the property that for every countable subset $C \subseteq A'$ and every $n \geq 1$, only countably many n -types are realized in $(\mathcal{A}', a)_{a \in C}$. Obviously \mathcal{A} and \mathcal{A}' cannot be isomorphic, which contradicts the assumption that T is κ -categorical.

Finally, the second sentence of the Proposition follows from the first using Theorem 18.15. \square

18.27. Lemma. *If T is ω -stable, then for every infinite cardinal κ and every regular cardinal $\lambda \leq \kappa$, T has a λ -saturated model of cardinality κ .*

Proof. Let $\lambda \leq \kappa$ be infinite cardinals with λ regular. By Theorem 18.15, T is κ -stable. Therefore we may build an elementary chain $(\mathcal{A}_\alpha \mid \alpha < \lambda)$ of models of T , all having cardinality equal to κ , such that for all $\alpha < \lambda$, every type of a finite tuple that is consistent with $\text{Th}((\mathcal{A}_\alpha, a)_{a \in A_\alpha})$ is realized in

$\mathcal{A}_{\alpha+1}$. Let \mathcal{A} be the union of this chain, so $\mathcal{A} \models T$ and $\text{card}(A) = \kappa$. Let C be any subset of A of cardinality $< \lambda$. Since λ is regular, there exists an $\alpha < \lambda$ such that $C \subseteq A_\alpha$. Any type in $L(C)$ of a finite tuple that is consistent with $\text{Th}((\mathcal{A}, a)_{a \in C})$ is realized in $(\mathcal{A}_\alpha, a)_{a \in C}$ and hence in $(\mathcal{A}, a)_{a \in C}$. Therefore \mathcal{A} is λ -saturated. \square

18.28. Corollary. *If κ is an uncountable cardinal and T is κ -categorical, then the unique model of T of cardinality κ is κ -saturated.*

Proof. Let \mathcal{A} be the unique model of T with cardinality κ . By Proposition 18.26, T is ω -stable. If κ is regular, then by taking $\lambda = \kappa$ in the previous result there is a κ -saturated model of cardinality κ ; this model is isomorphic to \mathcal{A} . If κ is not regular, then it is a limit cardinal. For any cardinal $\tau < \kappa$, we may apply the previous result with $\lambda = \tau^+$ to show that \mathcal{A} is τ^+ -saturated. Since κ is a limit of such cardinals, we conclude that \mathcal{A} is κ -saturated in this case too. \square

18.29. Remark. If T is ω -stable, then it can be proved that T has a κ -saturated model of cardinality κ for every infinite cardinal κ , without assuming categoricity. However, the proof of this result uses properties of Morley rank beyond the ones we developed.

18.30. Definition. Suppose \mathcal{A} is a model of T and $C \subseteq A$. Let a be a finite tuple from A and take $p(x) = \text{tp}_A(a/C)$ to be the type realized by a in $(\mathcal{A}, c)_{c \in C}$. When we say that $p(x)$ is *principal* we mean that it is principal relative to the $L(C)$ -theory $\text{Th}((\mathcal{A}, c)_{c \in C})$. (Here $p(x)$ is a complete type in $L(C)$ and $\text{Th}((\mathcal{A}, c)_{c \in C})$ is the set of sentences in $p(x)$, so there is no possible ambiguity.) That is, there exists an L -formula $\varphi(x, y)$ and a tuple d from C such that for any $L(C)$ -formula $\psi(x)$ in $p(x)$, the formula $\varphi(x, d) \rightarrow \psi(x)$ is valid in $(\mathcal{A}, c)_{c \in C}$. When this condition holds, we will say $\varphi(x, d)$ is a *complete formula in $p(x)$* . If $D \subseteq C$ and d is a tuple from D , then we say $p(x)$ is *principal over D* .

18.31. Lemma. *Let $\mathcal{A} \models T$ and $C \subset A$, and suppose a, b, a_0, \dots, a_n are finite tuples from A .*

- (1) *if $\text{tp}_A(a/C)$ is principal and every coordinate of b is either a coordinate of a or a member of C , then $\text{tp}_A(b/C)$ is principal;*
- (2) *$\text{tp}_A(ab/C)$ is principal if and only if $\text{tp}_A(a/C)$ and $\text{tp}_A(b/C \cup \{a\})$ are principal;*
- (3) *$\text{tp}_A(a_0 \dots a_n/C)$ is principal if and only if $\text{tp}_A(a_j/C \cup \{a_0, \dots, a_{j-1}\})$ is principal for each $0 \leq j \leq n$.*

Proof. (1) Suppose $\text{tp}_A(a/C)$ is principal and write a as a_1, \dots, a_m where $a_j \in A$ for each j . Let $\varphi(x_1, \dots, x_m)$ be a complete formula in $\text{tp}_A(a/C)$.

First, we treat the case where every coordinate of b is a coordinate of a ; say b is b_1, \dots, b_n and for each $j = 1, \dots, n$ let $\pi(j)$ be an element of $\{1, \dots, m\}$

for which $b_j = a_{\pi(j)}$. Then the formula

$$\exists x_1 \dots \exists x_m (\varphi(x_1, \dots, x_m) \wedge y_1 = x_{\pi(1)} \wedge \dots \wedge y_n = x_{\pi(n)})$$

is a complete formula in $\text{tp}_{\mathcal{A}}(b/C)$.

Note that the argument in the previous paragraph covers the case where b is a permutation of a . Therefore, to complete the proof of part (1) it suffices to show that $\text{tp}_{\mathcal{A}}(ac/C)$ is principal for each tuple $c = c_1, \dots, c_k$ from C . This type contains the complete formula

$$\varphi(x_1, \dots, x_m) \wedge x_{m+1} = c_1 \wedge \dots \wedge x_{m+k} = c_k.$$

(2) First suppose that $\text{tp}_{\mathcal{A}}(ab/C)$ contains the complete formula $\varphi(x, y)$. Then $\exists y \varphi(x, y)$ is a complete formula in $\text{tp}_{\mathcal{A}}(a/C)$ and $\varphi(a, y)$ is a complete formula in $\text{tp}_{\mathcal{A}}(b/C \cup \{a\})$. Conversely, suppose $\varphi(x)$ is a complete formula in $\text{tp}_{\mathcal{A}}(a/C)$ and $\psi(x, y)$ is an $L(C)$ -formula such that $\psi(a, y)$ is a complete formula in $\text{tp}_{\mathcal{A}}(b/C \cup \{a\})$. Then $\psi(x, y) \wedge \varphi(y)$ is a complete formula in $\text{tp}_{\mathcal{A}}(ab/C)$.

(3) This is proved by induction on n using part (2). \square

18.32. Definition. Let \mathcal{A} be an L -structure and $C \subseteq A$. We say that \mathcal{A} is *constructible over C* if there is an ordinal γ and a family $(a_\alpha \mid \alpha < \gamma)$ such that $A = C \cup \{a_\alpha \mid \alpha < \gamma\}$ and $\text{tp}_{\mathcal{A}}(a_\beta/C \cup \{a_\alpha \mid \alpha < \beta\})$ is principal for all $\beta < \gamma$.

18.33. Remark. Let \mathcal{A} be an L -structure and $C \subseteq A$, and assume \mathcal{A} is constructible over C . Then there exists an ordinal γ and a family $(a_\alpha \mid \alpha < \gamma)$ as in Definition 18.32 that also satisfies: $a_\alpha \notin C$ for all $\alpha < \gamma$ and $a_\alpha \neq a_\beta$ for all $\alpha < \beta < \gamma$. (From the original family remove all a_β that are members of C or equal some a_α with $\alpha < \beta$; it is easy to verify that the thinned family still witnesses that \mathcal{A} is constructible over C .)

18.34. Lemma. Let \mathcal{A} be an L -structure and $C \subseteq A$, and suppose that \mathcal{A} is constructible over C . Then $(\mathcal{A}, c)_{c \in C}$ is atomic; that is, $\text{tp}_{\mathcal{A}}(a/C)$ is principal for each finite tuple a from A .

Proof. Let $(a_\alpha \mid \alpha < \gamma)$ satisfy the conditions in Definition 18.32 and the preceding remark. That is,

$$A = C \cup \{a_\alpha \mid \alpha < \gamma\};$$

$\text{tp}_{\mathcal{A}}(a_\beta/C \cup \{a_\alpha \mid \alpha < \beta\})$ is principal for all $\beta < \gamma$;

$a_\alpha \notin C$ for all $\alpha < \gamma$; and

$a_\alpha \neq a_\beta$ for all $\alpha < \beta < \gamma$.

For convenience, set $C_\beta = \{a_\alpha \mid \alpha < \beta\}$ for each $\beta \leq \gamma$.

Let b be a finite tuple from C_γ . We say b is *good* if it is a permutation of a tuple $a_{\beta_1}, \dots, a_{\beta_n}$ such that $\beta_1 < \dots < \beta_n < \gamma$ and $\text{tp}_{\mathcal{A}}(a_{\beta_j}/C \cup C_{\beta_j})$ is principal over $C \cup \{\alpha_{\beta_1}, \dots, \alpha_{\beta_{j-1}}\}$ for each $j = 1, \dots, n$. Lemma 18.31 implies that $\text{tp}_{\mathcal{A}}(b/C)$ is principal whenever b is a good tuple from C_γ .

Now we prove, by induction on $\beta \leq \gamma$, that each finite tuple $b = b_1, \dots, b_n$ of distinct elements of C_β can be extended to a good tuple from C_β . So, let b be a finite tuple from $C_{\beta+1}$; we may assume that a_β occurs in b (or the desired result follows immediately from the induction hypothesis) and without loss of generality $b_n = a_\beta$. There are distinct $b'_1, \dots, b'_p \in C_\beta$ such that $\text{tp}_A(a_\beta/C \cup C_\beta)$ is principal over $C \cup \{b'_1, \dots, b'_p\}$. Let b' be the tuple obtained from $b_1, \dots, b_{n-1}, b'_1, \dots, b'_p$ by eliminating any b'_i that also occurs among b_1, \dots, b_{n-1} . Then b' is contained in C_β ; by the induction hypothesis it can be extended to a good tuple d from C_β . This argument is completed by noting that d, a_β is a good tuple from $C_{\beta+1}$ that extends b .

Now we prove the Lemma. Let b be any finite tuple from A ; we want to show that $\text{tp}_A(b/C)$ is principal. By Lemma 18.31 we may assume that no coordinate of b is in C and that the coordinates of b are distinct. By what was proved in the previous paragraph, there is a good tuple b' from C_γ that extends b . As noted above, $\text{tp}_A(b'/C)$ is principal. Hence Lemma 18.31 yields that $\text{tp}_A(b/C)$ is also principal, as desired. \square

18.35. Proposition. *Suppose T is ω -stable. Let $\mathcal{A} \models T$ and $C \subseteq A$. There exists $\mathcal{B} \preceq \mathcal{A}$ such that $C \subseteq B$ and \mathcal{B} is constructible over C .*

Proof. If C is the underlying set of an elementary substructure of \mathcal{A} , then take \mathcal{B} to be that structure. Otherwise there is an $L(C)$ -formula $\varphi(x)$ that is satisfied in \mathcal{A} but not by any element of C (by the Tarski-Vaught criterion). Choose such a formula with least possible Morley rank and degree. Let $(\alpha, d) = (RM(\varphi(x)), dM(\varphi(x)))$.

We claim that $\varphi(x)$ is a complete formula for a type $p(x)$ over C that is consistent with $\text{Th}((\mathcal{A}, a)_{a \in C})$. Otherwise there is an $L(C)$ -formula $\psi(x)$ such that $\varphi(x) \wedge \psi(x)$ and $\varphi(x) \wedge \neg\psi(x)$ are both consistent with $\text{Th}((\mathcal{A}, a)_{a \in C})$. But one of these formulas must have $(RM, dM) < (\alpha, d)$, which is impossible.

Let a_0 be an element of A that satisfies $\varphi(x)$ in \mathcal{A} . As shown above, $\text{tp}_A(a_0/C)$ is principal and $a_0 \notin C$. Continue inductively as long as possible to construct a sequence of distinct elements a_α in $A \setminus C$ for an initial segment of ordinals α such that whenever a_α is defined, we have that $\text{tp}_A(a_\alpha/C \cup \{a_\delta \mid \delta < \alpha\})$ is principal. Since A is a set, this construction must stop. If γ is the first ordinal at which the construction cannot be continued, then $C \cup \{a_\alpha \mid \alpha < \gamma\}$ is the underlying set of an elementary substructure of \mathcal{A} that is constructible over C . \square

Next we prove the main technical result of this chapter, from which Morley's Theorem is an easy consequence.

18.36. Theorem. *Suppose T is ω -stable. Assume κ is an uncountable cardinal and that every model of T of cardinality κ is κ -saturated. Then every uncountable model of T is saturated; that is, if $\mathcal{A} \models T$ and $\lambda = \text{card}(A)$ is uncountable, then \mathcal{A} is λ -saturated.*

Proof. Assume T is ω -stable and that κ, λ are uncountable cardinals. We will prove the contrapositive of the statement in the Theorem. That is, we assume T has a model \mathcal{A} of cardinality λ that is not λ -saturated and we obtain the same kind of model of cardinality κ .

So, there is a subset C of A of cardinality $< \lambda$ and a type $p(x)$ of a finite tuple over C such that $p(x)$ is consistent with $\text{Th}((\mathcal{A}, a)_{a \in C})$ but is not realized in $(\mathcal{A}, a)_{a \in C}$.

By Proposition 18.18 there is a nonconstant sequence $(a_k \mid k \in \mathbb{N})$ of ordered indiscernibles in $(\mathcal{A}, a)_{a \in C}$. Let $I = \{a_k \mid k \in \mathbb{N}\}$. Note that

(A) for each $L(C \cup I)$ -formula $\varphi(x)$ that is satisfiable in $(\mathcal{A}, a)_{a \in C \cup I}$ there exists $\psi(x) \in p(x)$ such that $\varphi(x) \wedge \neg\psi(x)$ is satisfiable in $(\mathcal{A}, a)_{a \in C \cup I}$ since otherwise $p(x)$ would be realized in $(\mathcal{A}, a)_{a \in C}$.

Let C_0 be any countable subset of C . For each $L(C_0 \cup I)$ -formula $\varphi(x)$ that is satisfiable in $(\mathcal{A}, a)_{a \in C_0 \cup I}$ let ψ_φ be one of the formulas ψ satisfying (A) for φ . Since $C_0 \cup I$ is countable, there is a countable set C_1 such that $C_0 \subseteq C_1 \subseteq C$ and such that the parameters of ψ_φ are in C_1 for all $L(C_0 \cup I)$ -formulas $\varphi(x)$ that are satisfiable in $(\mathcal{A}, a)_{a \in C_0 \cup I}$. Continue this inductively to define C_k for all $k \in \mathbb{N}$ and let $C' = \bigcup \{C_k \mid k \in \mathbb{N}\}$. This countable set satisfies $C' \subseteq C$ and the parameters of ψ_φ are in C' for all $L(C' \cup I)$ -formulas $\varphi(x)$ that are satisfiable in $(\mathcal{A}, a)_{a \in C' \cup I}$. Let $p'(x)$ be the restriction of $p(x)$ to C' . We have:

(B) for each $L(C' \cup I)$ -formula $\varphi(x)$ that is satisfiable in $(\mathcal{A}, a)_{a \in C' \cup I}$ there exists $\psi(x) \in p'(x)$ such that $\varphi(x) \wedge \neg\psi(x)$ is satisfiable in $(\mathcal{A}, a)_{a \in C' \cup I}$.

Note also that $(a_k \mid k \in \mathbb{N})$ is a sequence of ordered indiscernibles in $(\mathcal{A}, a)_{a \in C'}$.

By Proposition 16.5 there is a model of $\text{Th}((\mathcal{A}, a)_{a \in C'})$ that contains a family $(b_\alpha \mid \alpha < \kappa)$ of ordered indiscernibles having the same type as $(a_k \mid k \in \mathbb{N})$. We may assume this model is of the form $(\mathcal{B}, a)_{a \in C'}$. Using Proposition 18.35 there is $\mathcal{B}' \preceq \mathcal{B}$ such that $C' \cup \{b_\alpha \mid \alpha < \kappa\} \subseteq \mathcal{B}'$ and \mathcal{B}' is constructible over $C' \cup \{b_\alpha \mid \alpha < \kappa\}$.

We show that $p'(x)$ is not realized in $(\mathcal{B}', a)_{a \in C'}$. Suppose otherwise, that $p'(x)$ is realized by the finite tuple b in $(\mathcal{B}', a)_{a \in C'}$. By Lemma 18.34, we have that $\text{tp}_{\mathcal{B}'}(b/C' \cup \{b_\alpha \mid \alpha < \kappa\})$ is principal; it contains a complete formula that we may write as $\varphi(x, b_{\alpha_0}, \dots, b_{\alpha_n})$ where $\varphi(x, y_0, \dots, y_n)$ is an $L(C')$ -formula and $\alpha_0 < \dots < \alpha_n < \kappa$. So, for each formula $\psi(x)$ in $p'(x)$ we have that $\varphi(x, b_{\alpha_0}, \dots, b_{\alpha_n}) \rightarrow \psi(x)$ is valid in $(\mathcal{B}', b_{\alpha_0}, \dots, b_{\alpha_n}, a)_{a \in C'}$. But $b_{\alpha_0}, \dots, b_{\alpha_n}$ and a_0, \dots, a_n realize the same type over C' . Hence $\varphi(x, a_0, \dots, a_n) \rightarrow \psi(x)$ is valid in $(\mathcal{A}, a_0, \dots, a_n, a)_{a \in C'}$ for each formula $\psi(x)$ in $p'(x)$. This contradicts (B) and confirms the claim that $p'(x)$ is not realized in $(\mathcal{B}', a)_{a \in C'}$.

We finish the proof by using the downward Löwenheim-Skolem Theorem to get $\mathcal{B}'' \preceq \mathcal{B}'$ such that $C' \subseteq \mathcal{B}''$ and $\text{card}(\mathcal{B}'') = \kappa$. Then \mathcal{B}'' is a model

of T that has cardinality κ but is not κ -saturated. (Indeed, it is not even ω_1 -saturated.) \square

Now we put all the pieces together to give a proof of the main result:

Proof of Theorem 18.25 (Morley's Theorem). Suppose κ is an uncountable cardinal and T is κ -categorical. By Proposition 18.26, T is ω -stable. By Corollary 18.28, every model of T of cardinality κ is κ -saturated. Let λ be any uncountable cardinal. By Theorem 18.36, every model of T of cardinality λ is λ -saturated. Using Corollary 5.10 and the fact that T is complete, we conclude that any two models of T of cardinality λ are isomorphic. That is, T is λ -categorical. \square

19. UNIVERSAL SENTENCES AND SUBSTRUCTURES

Revision Note: this chapter will be merged with the next one.

A main point of this chapter is to axiomatize the class of all structures that can be embedded in models of a given set of sentences. This is a relatively easy application of the compactness theorem, and could be treated as soon as that theorem has been proved. It is used in the next chapter.

19.1. Definition. (1) A *universal* sentence is a sentence in prenex normal form that has only \forall quantifier symbols in its prefix.

(2) If Σ is a set of L -sentences, Σ_{\forall} denotes the set of all universal L -sentences σ such that $\Sigma \models \sigma$.

(3) If \mathcal{M} is an L -structure, $\text{Th}_{\forall}(\mathcal{M})$ denotes the set of all universal sentences that are true in \mathcal{M} ; it is called the *universal theory* of \mathcal{M} .

19.2. Theorem. *Let Σ be a set of L -sentences, and let \mathcal{M} be an L -structure. Then \mathcal{M} can be embedded in some model of Σ if and only if $\mathcal{M} \models \Sigma_{\forall}$.*

Proof. Suppose \mathcal{M} is embedded in an L -structure \mathcal{N} . Lemma 7.9(3) implies that any universal sentence true in \mathcal{N} is true in \mathcal{M} . Therefore, if \mathcal{N} is a model of Σ , then \mathcal{M} is a model of Σ_{\forall} .

For the converse direction, suppose $\mathcal{M} \models \Sigma_{\forall}$.

Let $\text{Diag}(\mathcal{M}) = \text{Diag}_M(\mathcal{M})$ denote the set of quantifier-free L_M -sentences that are true in $(\mathcal{M}, a)_{a \in M}$. (See Exercise 4.21.) To embed \mathcal{M} in a model of Σ it suffices to show that $\Sigma \cup \text{Diag}(\mathcal{M})$ has a model. To see this, say \mathcal{N}' is a model of this theory. The reduct \mathcal{N} of \mathcal{N}' to L is a model of Σ . Furthermore, the function $f: M \rightarrow N$ defined by taking each $f(a)$ to be the interpretation of a in \mathcal{N}' is an embedding of \mathcal{M} into \mathcal{N} .

By the Compactness Theorem it thus suffices to show that each finite subset of $\Sigma \cup \text{Diag}(\mathcal{M})$ has a model. Suppose otherwise. Since $\text{Diag}(\mathcal{M})$ is closed under conjunction, there would exist a sentence $\sigma \in \text{Diag}(\mathcal{M})$ such that $\Sigma \models \neg\sigma$. Let a_1, \dots, a_m be all the constants from M that occur in σ ; let z_1, \dots, z_m be distinct variables that do not occur in σ . Let $\varphi(z_1, \dots, z_m)$ be the L -formula that results from σ by replacing each occurrence of a_i by z_i for each $i = 1, \dots, m$. Note that $\varphi(a_1, \dots, a_m)$ is equal to σ . Since none of the constants a_1, \dots, a_m occurs in Σ we have that $\Sigma \models \forall z_1 \dots \forall z_m \neg\varphi(z_1, \dots, z_m)$. Since \mathcal{M} is a model of Σ_{\forall} and φ is quantifier free, we have $\mathcal{M} \models \neg\varphi[a_1, \dots, a_m]$, which implies that $\neg\varphi(a_1, \dots, a_m)$ is in $\text{Diag}(\mathcal{M})$. That is, we have $\neg\sigma \in \text{Diag}(\mathcal{M})$, which is a contradiction. \square

19.3. Corollary. *Let Σ be a set of L -sentences, and let \mathcal{M} be an L -structure. If each finitely generated substructure of \mathcal{M} can be embedded in some model of Σ , then \mathcal{M} can be embedded in some model of Σ .*

Proof. Suppose \mathcal{M} cannot be embedded in any model of Σ . By Theorem 19.2, there is a quantifier free L -formula $\varphi(x_1, \dots, x_n)$ whose universal closure $\forall x_1 \dots \forall x_n \varphi(x_1, \dots, x_n)$ is a consequence of Σ but is false in \mathcal{M} . Let a_1, \dots, a_n be elements of M such that $\mathcal{M} \models \neg\varphi[a_1, \dots, a_n]$. If \mathcal{M}' is taken to be the substructure of \mathcal{M} generated by the finite set $\{a_1, \dots, a_n\}$, then we see that $\forall x_1 \dots \forall x_n \varphi(x_1, \dots, x_n)$ is false in \mathcal{M}' . Hence \mathcal{M}' is a finitely generated substructure of \mathcal{M} that is not a model of Σ_{\forall} and thus cannot be embedded in any model of Σ . \square

19.4. Remark. It is easy to show that any structure can be embedded into a suitably chosen ultraproduct of its finitely generated substructures. This can be used to give another proof of the preceding Corollary.

19.5. Corollary. *Let \mathcal{M} and \mathcal{N} be L -structures. The structure \mathcal{M} can be embedded in some elementary extension of \mathcal{N} if and only if $\mathcal{M} \models \text{Th}_{\forall}(\mathcal{N})$.*

Proof. First note that to obtain an elementary extension of \mathcal{N} into which \mathcal{M} can be embedded, it suffices to simply have $\mathcal{N}' \equiv \mathcal{N}$ into which \mathcal{M} can be embedded. Indeed, given such an \mathcal{N}' , we may assume it is sufficiently saturated to ensure that there is an elementary embedding $f: \mathcal{N} \rightarrow \mathcal{N}'$. Then, by passing to a suitable structure isomorphic to \mathcal{N}' , we may obtain such a situation with f the identity map. That is, after these changes \mathcal{M} can still be embedded into \mathcal{N}' and \mathcal{N}' is an elementary extension of \mathcal{N} .

With this modification, the equivalence to be proved is exactly the conclusion of Theorem 19.2 in the case where $\Sigma = \text{Th}(\mathcal{N})$ is complete. \square

19.6. Corollary. *Let Σ be a set of L -sentences. The following conditions are equivalent:*

- (1) Σ is axiomatized by a set of universal sentences;
- (2) Σ is axiomatized by Σ_{\forall} ;
- (3) Every substructure of a model of Σ is a model of Σ .

Proof. (1) \Leftrightarrow (2): Evidently Σ_{\forall} contains any set of universal sentences that axiomatizes Σ , if any such set exists.

(2) \Rightarrow (3): Suppose $\mathcal{M} \subseteq \mathcal{N} \models \Sigma$. Lemma 7.9 implies that $\mathcal{M} \models \Sigma_{\forall}$; therefore $\mathcal{M} \models \Sigma$ by (2).

(3) \Rightarrow (2): We see that condition (2) is equivalent to $\text{Mod}(\Sigma_{\forall}) \subseteq \text{Mod}(\Sigma)$. Suppose $\mathcal{M} \models \Sigma_{\forall}$. By Theorem 19.2 there is $\mathcal{N} \in \text{Mod}(\Sigma)$ with $\mathcal{M} \subseteq \mathcal{N}$. Then (3) implies that \mathcal{M} is a model of Σ . \square

The following result is not difficult to prove, but it turns out to be useful in many application settings.

19.7. Remark. Suppose Σ is a set of universal L -sentences and that Σ has QE. Then Σ has Skolem terms; that is, for any L -formula $\varphi(x, y)$, where

$x = x_1, \dots, x_m$ with $m \geq 1$ and y is a single variable, there are L -terms $t_1(x), \dots, t_n(x)$ such that

$$\Sigma \models \exists y \varphi(x, y) \rightarrow (\varphi(x, t_1(x)) \vee \dots \vee \varphi(x, t_n(x))).$$

Proof. This is immediate using Proposition 15.2. It suffices to verify condition 15.2(1). Given $\mathcal{M} \subseteq \mathcal{N} \models \Sigma$, we must show $\mathcal{M} \preceq \mathcal{N}$. Since Σ consists of universal sentences, we have $\mathcal{M} \models \Sigma$, and since Σ has QE, we therefore also have $\mathcal{M} \preceq \mathcal{N}$. \square

19.8. Example. An instructive example of the preceding remark comes from the theory of real closed ordered fields. Let L be the language obtained from L_{or} by adding, for each even $n \in \mathbb{N}$, an $(n+1)$ -ary new function symbol g_n . Let $RCOF^*$ be the set of L -sentences obtained from the axioms $RCOF$ for real closed ordered fields by adding sentences asserting:

- (1) for all a in the given ordered field, $g_0(a)^2 = |a|$ and $g_0(a) \geq 0$;
- (2_n ; n even ≥ 2) for all a_0, \dots, a_n in the given ordered field, $g_n(a_0, \dots, a_n)$ is the least root of the polynomial

$$p(a, x) = a_0 + a_1x + \dots + a_nx^n + x^{n+1}.$$

We want to show that $\Sigma = RCOF^*$ satisfies the assumptions in Remark 19.7.

Note that the functions (g_n) that we are adding to a real closed ordered field are 0-definable in $RCOF$ and algebraic. In particular, each real closed ordered field expands in a unique way to a model of $RCOF^*$. (As discussed later, in Chapter 21, in this situation we say that $RCOF^*$ is an *extension by definitions of RCOF*.)

First we need to see that $RCOF^*$ can be axiomatized by universal L -sentences. The sentences described in (1) and (2) can obviously be expressed by universal L -sentences. (If $\mathcal{M} \models RCOF$ and $a \in M^{n+1}$, then $b \in M$ is the least root of $p(a, x)$ in \mathcal{M} iff $p(a, b) = 0$ and for all $c < b$ in \mathcal{M} , $p(a, c) \neq 0$.) Further, the functions interpreting $g_0, g_2, g_4 \dots$ give Skolem functions for the axioms in $RCOF$, allowing them to be expressed as universal sentences in the language L .

Second, we must show that $RCOF^*$ has QE. Since the functions we add to L_{or} are all definable in $RCOF^*$ by formulas from L_{or} , it follows that every L -formula is $RCOF^*$ -equivalent to an L_{or} -formula. Since $RCOF$ has QE, as shown in Theorem 12.3, the L_{or} -formula can be taken to be quantifier free.

Thus Remark 19.7 applies to $RCOF^*$. Hence we have expanded $RCOF$ by 0-definable algebraic functions, achieving a mathematically equivalent theory $RCOF^*$ in which *all* formulas have definable Skolem functions; indeed, they are given piecewise by functions obtained by interpreting L -terms, hence by algebraic functions that are expressed explicitly as compositions of the basic functions in L .

We finish this chapter by giving two examples around the problem of understanding substructures. In each case we consider an axiomatizable class $\mathcal{C} = \text{Mod}(\Sigma)$, and we are interested in the class \mathcal{SC} of structures that can be embedded in members of \mathcal{C} . Of course $\mathcal{SC} = \text{Mod}(\Sigma_{\forall})$, by Theorem 19.2. But what if we want an explicit and comprehensible set of axioms for \mathcal{SC} ?

19.9. Example. Any substructure of a field is a ring with identity. Moreover, a ring with identity can be embedded in a field iff the ring is an integral domain. Therefore, if Σ is a set of axioms for the class of fields, then Σ_{\forall} is axiomatized by the axioms for commutative rings with identity together with the cancellation law for multiplication. Here we consider rings with identity as structures for the language with binary function symbols $+$, $-$, \times for addition, subtraction, and multiplication (resp.), and with constants $0, 1$.

19.10. Example. Now consider the class of groups, considered as structures for the language with a binary function symbol for the group operation and a constant symbol for the identity element. In this setting, a substructure of a group is a monoid.

The question “which monoids can be embedded in a group” has a more complex answer than the corresponding question for rings and fields. The cancellation rule must, of course, hold in any monoid that can be embedded in a group, but this is not a sufficient condition in general. (It is sufficient for finite monoids and for abelian monoids.)

Let Σ be a (finite) set of axioms for the class of groups, in the language specified above. By Theorem 19.2, we know that a monoid M embeds in a group iff $M \models \Sigma_{\forall}$. But Σ_{\forall} is rather complicated; for example, it is not even a decidable set of sentences. (This is due to the fact that the word problem for groups, which was shown to be undecidable by Boone and Novikov, is effectively reducible to Σ_{\forall} . Indeed, any universal sentence in this language is logically equivalent to a conjunction of sentences of the form

$$\forall x_1 \dots \forall x_k [(w_1 = v_1 \wedge \dots \wedge w_m = v_m) \rightarrow (w_{m+1} = v_{m+1} \vee \dots \vee w_{m+n} = v_{m+n})]$$

for some terms $w_1, v_1, \dots, w_{m+n}, v_{m+n}$ in the language of groups. Although our language does not include the inverse operation, we can express that x and y are inverse by the equation $x \cdot y = e$, where e is the constant symbol representing the identity element. It is therefore easy to show that each specific instance of the word problem for groups can be formulated as the question whether a specific sentence of the displayed form is in Σ_{\forall} .)

The theory Σ_{\forall} is in any case known to be computably enumerable, since Σ is finite, and thus Σ_{\forall} can be axiomatized by a decidable set of universal sentences. (If Δ is a computably enumerable set of L -sentences, with Δ enumerated effectively as $(\sigma_0, \sigma_1, \sigma_2, \dots)$, then $\Delta' = \{\sigma_0 \wedge \dots \wedge \sigma_n \mid n \in \mathbb{N}\}$ axiomatizes Δ and is decidable, because the members of Δ' are effectively enumerated in order of increasing length.) Indeed, with a bit of work, when

Σ axiomatizes the class of all groups, one can describe a comprehensible (but still rather complicated) set of universal axioms for Σ_{\forall} .

20. EXISTENTIAL FORMULAS AND MODEL COMPLETENESS

Revision Note: this chapter will be revised to give a more complete treatment of model companions, especially for inductive theories. The previous chapter (on universal axiomatizations and embeddings) will be moved into this one.

Let L be a fixed language. It is useful to have some notation that refers to the quantifier complexity of L -formulas. For $n \geq 0$, we let \exists_n denote the class of all formulas that are in prenex normal form with a prenex part consisting of n alternating blocks of quantifiers, beginning on the left with \exists . Also, \forall_n denotes the dual class of all formulas that are in prenex normal form with a prenex part consisting of n alternating blocks of quantifiers, beginning on the left with \forall . Thus \exists_0 and \forall_0 each consist of the quantifier free formulas; \exists_1 is the set of *existential* formulas; and \forall_1 is the set of *universal* formulas (which already appeared in Chapter 19). It is customary to refer to the formulas in \forall_2 as $\forall\exists$ formulas; they will play a special role in this chapter.

Now let Σ be a set of L -sentences. For each $n \geq 0$ we let $\exists_n(\Sigma)$ denote the set of all L -formulas that are Σ -equivalent to some L -formula in \exists_n . Similarly, $\forall_n(\Sigma)$ denotes the set of all L -formulas that are Σ -equivalent to some L -formula in \forall_n .

20.1. Theorem. *Let Σ be a set of L -sentences, and let $\varphi(x)$ be an L -formula, with $x = x_1, \dots, x_m$. The following conditions are equivalent:*

- (1) $\varphi(x)$ is Σ -equivalent to an existential L -formula.
- (2) $\varphi(x)$ persists upwards in models of Σ ; that is, if \mathcal{M}, \mathcal{N} are models of Σ and $\mathcal{M} \subseteq \mathcal{N}$, then for any $a \in M^n$,

$$\mathcal{M} \models \varphi[a] \quad \Rightarrow \quad \mathcal{N} \models \varphi[a].$$

Proof. (1 \Rightarrow 2) Let $\psi(x)$ be an existential L -formula that is Σ -equivalent to $\varphi(x)$. Suppose $\psi(x)$ is the formula $\exists y_1 \dots \exists y_n \theta(x, y)$, where $y = y_1, \dots, y_n$ and θ is quantifier free. Suppose also that $\mathcal{M} \subseteq \mathcal{N}$ are models of Σ and $\mathcal{M} \models \varphi(x)[a]$, where $a \in M^m$. Then also $\mathcal{M} \models \psi(x)[a]$, so there exists $b \in M^n$ such that $\mathcal{M} \models \theta(x, y)[a, b]$. Since θ is quantifier free, we conclude that $\mathcal{N} \models \theta(x, y)[a, b]$ and hence $\mathcal{N} \models \psi(x)[a]$. Therefore we have $\mathcal{N} \models \varphi(x)[a]$.

(2 \Rightarrow 1) Suppose (2) holds of the L -formula $\varphi(x)$. We apply Proposition 3.21 to Σ and $\varphi(x)$, with $\Delta(x)$ equal to the set of L -formulas $\psi(x)$ that are Σ -equivalent to an existential formula. Note that Δ is closed under disjunction and conjunction, as needed to apply 3.21.

We verify condition 3.21(2) in this setting. Then condition 3.21(1) follows, which yields that either $\Sigma \models \varphi(x)$ or $\Sigma \models \neg\varphi(x)$ or there is an existential L -formula that is Σ -equivalent to $\varphi(x)$. In the first two cases we see that

$\varphi(x)$ is Σ -equivalent either to $\exists y(y = x)$ or to $\exists y(y \neq x)$, both of which are existential. Hence to prove (1) it suffices to verify condition 3.21(2).

Thus we consider types $p(x), q(x) \in S_x(\Sigma)$ and assume $\varphi(x) \in p(x)$ and $\neg\varphi(x) \in q(x)$. We need to show there exists some $\psi(x) \in \exists_1(\Sigma)$ such that $\psi(x) \in p(x)$ and $\neg\psi(x) \in q(x)$. We argue by contradiction; thus we assume for the rest of the proof that no such $\psi(x) \in \exists_1(\Sigma)$ exists; the proof is finished when we derive a contradiction.

Consider $\mathcal{M}, \mathcal{N} \models \Sigma$ as well as $a \in M^m$ and $b \in N^m$ that realize $p(x)$ in \mathcal{M} and $q(x)$ in \mathcal{N} respectively. Our assumptions yield that $\mathcal{M} \models \varphi(x)[a]$ and $\mathcal{N} \models \neg\varphi(x)[b]$. Further, we have that for every L -formula $\psi(x) \in \exists_1(\Sigma)$, if $\mathcal{M} \models \psi(x)[a]$ then $\mathcal{N} \models \psi(x)[b]$.

Let $c = c_1, \dots, c_m$ be distinct constant symbols not in L , and consider the language L_c obtained from L by adding c . Under the hypotheses described in the previous paragraph, we see that the L_c -structure $(\mathcal{M}, a_1, \dots, a_m)$ satisfies every universal L_c -sentence that is true in $(\mathcal{N}, b_1, \dots, b_m)$. Therefore, by Corollary 19.5 we conclude that there is an elementary extension $(\mathcal{N}', b_1, \dots, b_m)$ of $(\mathcal{N}, b_1, \dots, b_m)$ and an embedding f of $(\mathcal{M}, a_1, \dots, a_m)$ into $(\mathcal{N}', b_1, \dots, b_m)$. Condition (2) in the Lemma implies that $\mathcal{N}' \models \varphi[b_1, \dots, b_m]$, which in turn implies $\mathcal{N} \models \varphi[b_1, \dots, b_m]$. This contradiction completes the proof. \square

For completeness we give the analogous characterization of $\forall_1(\Sigma)$.

20.2. Corollary. *Let Σ be a set of L -sentences, and let $\psi(x)$ be an L -formula, with $x = x_1, \dots, x_m$. The following conditions are equivalent:*

- (1) $\psi(x)$ is Σ -equivalent to a universal L -formula.
- (2) $\psi(x)$ persists downwards in models of Σ ; that is, if \mathcal{M}, \mathcal{N} are models of Σ and $\mathcal{M} \subseteq \mathcal{N}$, then for any $b \in N^m$,

$$\mathcal{N} \models \psi[b] \Rightarrow \mathcal{M} \models \psi[b].$$

Proof. Apply 20.1 to $\varphi(x) = \neg\psi(x)$. \square

20.3. Definition. Let Σ be a set of L -sentences. We say Σ is *model complete* if whenever \mathcal{M}, \mathcal{N} are models of Σ and $\mathcal{M} \subseteq \mathcal{N}$, one has $\mathcal{M} \preceq \mathcal{N}$.

20.4. Remark. Any theory that admits Quantifier Elimination is model complete.

20.5. Proposition. *Let Σ be a set of L -sentences. The following conditions are equivalent:*

- (1) Σ is model complete;
- (2) Every universal L -formula is in $\exists_1(\Sigma)$;
- (3) Every existential L -formula is in $\forall_1(\Sigma)$;
- (4) Every L -formula is in $\exists_1(\Sigma)$;
- (5) Every L -formula is in $\forall_1(\Sigma)$.

Proof. The implications (2) \Rightarrow (4) and (3) \Rightarrow (5) are proved by straightforward inductions on formulas, after putting a given formula into an equivalent prenex form. Their converses are trivial. The equivalence (1) \Leftrightarrow (4) is immediate using Theorem 20.1. Likewise, (1) \Leftrightarrow (5) follows from Corollary 20.2. \square

20.6. Definition. Let \mathcal{M}, \mathcal{N} be L -structures with $\mathcal{M} \subseteq \mathcal{N}$. We say that \mathcal{M} is *existentially closed* in \mathcal{N} (and write $\mathcal{M} \subseteq_{ec} \mathcal{N}$) if for each existential L -formula $\varphi(x)$ and each $a \in M^n$, if $\mathcal{N} \models \varphi(x)[a]$, then $\mathcal{M} \models \varphi(x)[a]$.

20.7. Remark. An existential formula is called *primitive* if it is of the form $\exists y_1 \dots \exists y_n \psi(x_1, \dots, x_m, y_1, \dots, y_n)$ where $\psi(x, y)$ is a conjunction of atomic formulas and negations of atomic formulas. It is easy to see that every existential formula is logically equivalent to a disjunction of primitive existential formulas; just write the quantifier free part of the formula in disjunctive normal form and then use prenex rules to distribute the \exists quantifiers over the disjunctions. It follows that (in the preceding definition) \mathcal{M} is existentially closed in \mathcal{N} if and only if the condition in the definition holds for every *primitive* existential formula $\varphi(x)$. In many situations this is a useful observation, because primitive existential formulas often express solvability conditions that have a clear mathematical meaning in the context of the structures being studied.

Note. In the rest of this chapter we use several times the *elementary diagram* $\text{EDiag}(\mathcal{M})$ and the *diagram* $\text{Diag}(\mathcal{M})$ of a structure \mathcal{M} . We take $\text{EDiag}(\mathcal{M}) = \text{EDiag}_M(\mathcal{M})$ to be the L_M -theory of the expanded structure $\mathcal{M}_M = (\mathcal{M}, a)_{a \in M}$; further, $\text{Diag}(\mathcal{M}) = \text{Diag}_M(\mathcal{M})$ is the set of all quantifier free L_M -sentences in $\text{EDiag}(\mathcal{M})$. See Exercises 4.21 and 4.22 for precise statements of their main properties.

20.8. Proposition. *Let \mathcal{M}, \mathcal{N} be L -structures with $\mathcal{M} \subseteq \mathcal{N}$. Then \mathcal{M} is existentially closed in \mathcal{N} if and only if there is an elementary extension \mathcal{M}' of \mathcal{M} and an embedding f of \mathcal{N} into \mathcal{M}' such that $f(a) = a$ for all $a \in M$.*

Proof. (\Leftarrow) Suppose $\mathcal{M} \subseteq \mathcal{N}$, $\mathcal{M}' \succeq \mathcal{M}$, and $f : \mathcal{N} \rightarrow \mathcal{M}'$ is an embedding that restricts to the identity on M . Consider an existential L -formula $\exists y_1 \dots \exists y_n \psi(x, y)$, where $\psi(x, y)$ is quantifier free, with $x = x_1, \dots, x_m$ and $y = y_1, \dots, y_n$. Consider $a = a_1, \dots, a_m \in M^m$. Suppose that $\mathcal{N} \models \exists y_1 \dots \exists y_n \psi(x, y)[a]$. Because f is an embedding and $\psi(x, y)$ is quantifier free, we have $\mathcal{M}' \models \exists y_1 \dots \exists y_n \psi(x, y)[a]$. (Recall $f(a_i) = a_i$ for all $i = 1, \dots, m$.) Since $\mathcal{M} \preceq \mathcal{M}'$, we conclude $\mathcal{M} \models \exists y_1 \dots \exists y_n \psi(x, y)[a]$. Therefore \mathcal{M} is existentially closed in \mathcal{N} .

(\Rightarrow) Consider $\Phi = \text{EDiag}(\mathcal{M}) \cup \text{Diag}(\mathcal{N})$ in the language L_N . To get \mathcal{M}' as desired, it suffices to show that Φ is satisfiable. To see this, suppose $(\mathcal{M}', f(b))_{b \in N}$ is a model of Δ . By changing to an isomorphic model we may assume that $f(a) = a$ for all $a \in M$. Then \mathcal{M}' and f have the desired properties.

Arguing by contradiction, we suppose Φ is unsatisfiable. Since $\text{Diag}(\mathcal{N})$ is closed under conjunction there exists a sentence $\sigma \in \text{Diag}(\mathcal{N})$ such that $\text{EDiag}(\mathcal{M}) \models \neg\sigma$. We may write σ as $\varphi(a_1, \dots, a_m, b_1, \dots, b_n)$ where $\varphi(x_1, \dots, x_m, y_1, \dots, y_n)$ is a quantifier free L -formula, $a_1, \dots, a_m \in M$, and $b_1, \dots, b_n \in N \setminus M$. Since the constants b_1, \dots, b_n do not occur in $\text{EDiag}(\mathcal{M})$ it follows that

$$\text{EDiag}(\mathcal{M}) \models \forall y_1 \dots \forall y_n \neg\varphi(a_1, \dots, a_m, y_1, \dots, y_n).$$

Equivalently we have

$$\text{EDiag}(\mathcal{M}) \models \neg\exists y_1 \dots \exists y_n \varphi(a_1, \dots, a_m, y_1, \dots, y_n).$$

But then $\exists y_1 \dots \exists y_n \varphi(x_1, \dots, x_m, y_1, \dots, y_n)$ is an existential formula that is true of a_1, \dots, a_m in \mathcal{N} and false of a_1, \dots, a_m in \mathcal{M} . Thus \mathcal{M} is not existentially closed in \mathcal{N} , a contradiction. \square

20.9. Definition. Let Σ be a set of L -sentences. We denote by $\Sigma_{\forall\exists}$ the set of all $\forall\exists$ L -sentences σ such that $\Sigma \models \sigma$.

20.10. Proposition. Let Σ be a set of L -sentences and let \mathcal{M} be an L -structure. Then $\mathcal{M} \models \Sigma_{\forall\exists}$ if and only if there is a model \mathcal{N} of Σ such that $\mathcal{M} \subseteq \mathcal{N}$ and \mathcal{M} is existentially closed in \mathcal{N} .

Proof. (\Leftarrow) Exercise.

(\Rightarrow) Suppose $\mathcal{M} \models \Sigma_{\forall\exists}$. We construct a model \mathcal{N} of Σ so that \mathcal{M} is an existentially closed substructure of \mathcal{N} . Let Φ be the set of universal sentences in $\text{EDiag}(\mathcal{M})$ and consider $\Sigma' = \Sigma \cup \text{Diag}(\mathcal{M}) \cup \Phi$ in the language L_M . Suppose Σ' has a model \mathcal{N}' . Without loss of generality we may assume $a^{\mathcal{N}'} = a$ for all $a \in M$. Let \mathcal{N} be the reduct of \mathcal{N}' to L . Since $\mathcal{N}' \models \text{Diag}(\mathcal{M})$, we have $\mathcal{M} \subseteq \mathcal{N}$. Evidently $\mathcal{N} \models \Sigma$. For any universal L -formula $\varphi(x_1, \dots, x_n)$ and any $a_1, \dots, a_n \in M$, we have that $\mathcal{M} \models \varphi[a_1, \dots, a_n]$ iff $\mathcal{N} \models \varphi[a_1, \dots, a_n]$; in one direction this is because \mathcal{M} is a substructure of \mathcal{N} and in the other it is because $\mathcal{N}' \models \Sigma'$. By taking negations, we see that \mathcal{M} is an existentially closed substructure of \mathcal{N} .

Therefore it suffices to show Σ' is satisfiable. Arguing by contradiction, we suppose otherwise. In that case we would have sentences $\varphi(a) \in \text{Diag}(\mathcal{M})$, where $a = a_1, \dots, a_n$ is in M^n , and $\forall y_1 \dots \forall y_n \psi(a, y) \in (\text{EDiag}(\mathcal{M}))_{\forall}$ so that $\Sigma \models \neg\varphi(a) \vee \exists y_1 \dots \exists y_n \neg\psi(a, y)$. Here $y = y_1, \dots, y_n$. Note that there is no loss of generality in assuming that we have only a single sentence from each set, because $\text{Diag}(\mathcal{M})$ is closed under conjunction, and the conjunction of finitely many universal sentences is logically equivalent to a universal sentence. Because the constants a do not appear in Σ , we have

$$\Sigma \models \forall x_1 \dots \forall x_m (\neg\varphi(x) \vee \exists y_1 \dots \exists y_n \neg\psi(x, y))$$

where $x = x_1, \dots, x_m$.

Since $\mathcal{M} \models \Sigma_{\forall\exists}$, this $\forall\exists$ sentence is true in \mathcal{M} . In particular $\mathcal{M} \models \neg(\varphi(x) \vee \neg\forall y \psi(x, y))[a]$. This contradiction to the choice of φ and ψ completes the proof. \square

20.11. **Definition.** Let Σ be a set of L -sentences. We call Σ *inductive* if $\text{Mod}(\Sigma)$ is closed under unions of arbitrary chains; *i.e.*, whenever $(I, <)$ is a linearly ordered set and $(\mathcal{M}_i \mid i \in I)$ is a chain of models of Σ , so that $\mathcal{M}_i \subseteq \mathcal{M}_j$ holds for all $i < j$ in I , the union $\bigcup\{\mathcal{M}_i \mid i \in I\}$ is a model of Σ .

20.12. **Theorem.** Let Σ be a set of L -sentences. The following conditions are equivalent:

- (1) Σ is axiomatized by $\Sigma_{\forall\exists}$.
- (2) Σ is axiomatized by some set of $\forall\exists$ sentences.
- (3) Σ is inductive.
- (4) Whenever $(\mathcal{M}_n \mid n \in \mathbb{N})$ is an ω -chain of models of Σ , so that $\mathcal{M}_n \subseteq \mathcal{M}_{n+1}$ holds for all n , the union $\bigcup_1^\infty \mathcal{M}_n$ is a model of Σ .

Proof. (1) \Leftrightarrow (2) and (3) \Rightarrow (4) are obvious.

(2) \Rightarrow (3) Consider a linearly ordered set $(I, <)$ and a chain $(\mathcal{M}_i \mid i \in I)$ of models of Σ . Let \mathcal{N} be the union of $(\mathcal{M}_i \mid i \in I)$. It suffices to show that any $\forall\exists$ -sentence σ that is true in \mathcal{M}_i for all $i \in I$ must also be true in \mathcal{N} . Suppose σ is $\forall x_1 \dots \forall x_m \exists y_1 \dots \exists y_n \varphi(x, y)$, where $\varphi(x, y)$ is quantifier-free, and that $\mathcal{M}_i \models \sigma$ for all $i \in I$. Given any $a \in N^m$, there exists $i \in I$ such that $a \in M_i^m$. Since $\mathcal{M}_i \models \sigma$, there exists $b \in M_i^n \subseteq N^n$ such that $\mathcal{M}_i \models \varphi(x, y)[a, b]$. Since $\varphi(x, y)$ is quantifier free and $\mathcal{M}_i \subseteq \mathcal{N}$, we have $\mathcal{N} \models \varphi(x, y)[a, b]$. Since $a \in N^m$ was arbitrary, this shows σ is true in \mathcal{N} .

(4) \Rightarrow (1) It suffices to show $\text{Mod}(\Sigma_{\forall\exists}) \subseteq \text{Mod}(\Sigma)$. We use an elementary chains argument.

Suppose we are given that $\mathcal{M}_0 \models \Sigma_{\forall\exists}$. By Proposition 20.10 we know there exists $\mathcal{M}_1 \models \Sigma$, such that $\mathcal{M}_0 \subseteq_{ec} \mathcal{M}_1$. By Proposition 20.8 there exists $\mathcal{M}_2 \succeq \mathcal{M}_0$ and an embedding f of \mathcal{M}_1 into \mathcal{M}_2 that is the identity on M_0 . By changing \mathcal{M}_2 to an isomorphic model if necessary, we may assume $M_1 \subseteq M_2$ and $f(a) = a$ for all $a \in M_1$. Therefore $\mathcal{M}_0 \subseteq \mathcal{M}_1 \subseteq \mathcal{M}_2$. Because $\mathcal{M}_0 \preceq \mathcal{M}_2$ we have $\mathcal{M}_2 \models \Sigma_{\forall\exists}$.

Since $\mathcal{M}_2 \models \Sigma_{\forall\exists}$ we are in position to continue this construction inductively. This yields an ω -chain of structures $(\mathcal{M}_n \mid n \in \mathbb{N})$ such that $\mathcal{M}_{2k} \preceq \mathcal{M}_{2k+2}$ for all $k \geq 0$ and $\mathcal{M}_{2k+1} \models \Sigma$ for all $k \geq 0$.

Now let $\mathcal{N} = \bigcup \mathcal{M}_n = \bigcup \mathcal{M}_{2k} = \bigcup \mathcal{M}_{2k+1}$. $(\mathcal{M}_{2k})_{k \in \mathbb{N}}$ is an elementary chain, and hence $\mathcal{N} \succeq \mathcal{M}_0$; $(\mathcal{M}_{2k+1})_{k \in \mathbb{N}}$ is an ω -chain of models of Σ , and hence $\mathcal{N} \models \Sigma$ by condition (4). Therefore \mathcal{M}_0 is also a model of Σ . \square

20.13. **Corollary.** If Σ is a model complete set of L -sentences, then Σ is inductive, so Σ is axiomatized by a set of $\forall\exists$ sentences; in particular, Σ is axiomatized by $\Sigma_{\forall\exists}$.

Proof. If Σ is a model complete theory then any chain of models is an elementary chain. Hence the union of the chain is an elementary extension of each of them, and thus it is a model of Σ . Theorem 20.12 yields that Σ is axiomatized by $\forall\exists$ -sentences. \square

20.14. **Corollary** (Robinson's Test). *Let Σ be a set of L -sentences. Then Σ is model complete if and only if for any two models \mathcal{M}, \mathcal{N} of Σ , if $\mathcal{M} \subseteq \mathcal{N}$, then \mathcal{M} is existentially closed in \mathcal{N} .*

Proof. (\Rightarrow) This is an immediate consequence of the definitions.

(\Leftarrow) Let $\mathcal{M}_0 \subseteq \mathcal{N}_0$ be models of Σ . We need to show that $\mathcal{M}_0 \preceq \mathcal{N}_0$. We build a chain as in the proof of Theorem 20.12, starting with \mathcal{M}_0 . Using Proposition 20.8 repeatedly we get a chain $\mathcal{M}_0 \subseteq_{ec} \mathcal{N}_0 \subseteq_{ec} \mathcal{M}_1 \subseteq_{ec} \mathcal{N}_1 \subseteq_{ec} \dots$, all being models of Σ with $\mathcal{M}_0 \preceq \mathcal{M}_1 \preceq \dots$ and $\mathcal{N}_0 \preceq \mathcal{N}_1 \preceq \mathcal{N}_2 \preceq \dots$. Let \mathcal{N}' be the union of the whole chain; then $\mathcal{N}' = \bigcup \mathcal{M}_n = \bigcup \mathcal{N}_n$. Thus \mathcal{N}' is simultaneously an elementary extension of both \mathcal{M}_0 and \mathcal{N}_0 . It follows that $\mathcal{M}_0 \preceq \mathcal{N}_0$. \square

20.15. **Remark.** Suppose Σ is a set of L -sentences and that L has at most κ symbols. Using the Löwenheim-Skolem Theorem we may show that it suffices to consider only models of cardinality $\leq \kappa$ in the definition of model completeness of Σ or in Robinson's Test for model completeness of Σ . For example, suppose Σ has a pair of models $\mathcal{M}' \subseteq \mathcal{N}'$ and that \mathcal{M}' is not existentially closed in \mathcal{N}' . Let \mathcal{M} be an elementary submodel of \mathcal{M}' whose cardinality is $\leq \kappa$; then let \mathcal{N} be an elementary submodel of \mathcal{N}' that contains \mathcal{M} as a substructure and has cardinality $\leq \kappa$. Then $\mathcal{M} \subseteq \mathcal{N}$ are models of Σ , both have cardinality $\leq \kappa$, and \mathcal{M} is not existentially closed in \mathcal{N} .

21. CHARACTERIZING DEFINABILITY

In this chapter we present a number of basic results, all of which concern definability in one way or another. The results discussed here include the characterizations of definability due to Svenonius and Beth. We also present Robinson's Joint Consistency Lemma and Craig's Interpolation Theorem, and discuss the relations among these fundamental properties of first order logic.

Our first results give necessary and sufficient conditions for a relation to be 0-definable in a given structure. As background, we note that if \mathcal{M} is an L -structure and τ is an automorphism of \mathcal{M} , then we consider τ as acting on each finite power M^n by taking $\tau(a_1, \dots, a_n) = (\tau(a_1), \dots, \tau(a_n))$. Since the interpretation of every formula in \mathcal{M} is invariant under these actions (*i.e.*, automorphisms are elementary maps), a necessary condition for a set $R \subset M^n$ to be 0-definable in \mathcal{M} is that it be setwise invariant under every automorphism of \mathcal{M} . These first results examine the extent to which some weakening of the converse might hold. (The converse itself cannot hold in full generality; indeed, there are many structures, such as $(\mathbb{N}, <)$, that have no automorphisms except for the identity. Evidently a structure for a countable language can have at most countably many 0-definable sets, since there are only countably many formulas. If the structure is infinite, there must be uncountably many sets that are not 0-definable; for a rigid structure like $(\mathbb{N}, <)$, every set will be invariant under its automorphisms, while most of its sets will not be 0-definable.)

If L is a language and S is a family of new symbols that are being added to the signature, we denote the extension language by $L(S)$.

21.1. Theorem. *Let L be any first order language and let P be an n -ary predicate symbol that is not in L . Suppose (\mathcal{M}, R) is an ω -saturated $L(P)$ -structure and that \mathcal{M} is strongly ω -homogeneous. The following conditions are equivalent:*

- (1) R is 0-definable in \mathcal{M} ;
- (2) every automorphism of \mathcal{M} leaves R setwise invariant.

Proof. (1) \Rightarrow (2) is immediate, since automorphisms are elementary maps.

(2) \Rightarrow (1): Assume every automorphism of \mathcal{M} leaves R setwise invariant. By way of contradiction, assume that R is not 0-definable in \mathcal{M} . To show the existence of the desired automorphism of \mathcal{M} , it suffices to find n -tuples a and b from M^n that realize the same n -type in \mathcal{M} but such that $R(a)$ is true and $R(b)$ is false. We would then have an automorphism of \mathcal{M} taking a to b , since \mathcal{M} is strongly ω -homogeneous; this automorphism would not leave R invariant, contradicting our hypothesis.

Let $x = x_1, \dots, x_n$ be a list of distinct variables. For any $a \in M^n$ we regard $\text{tp}_{\mathcal{M}}(a)$ as the set of L -formulas $\varphi(x)$ such that $\mathcal{M} \models \varphi(x)[a]$.

Since (\mathcal{M}, R) is ω -saturated, it suffices to show there exists a in R so that $\text{tp}_{\mathcal{M}}(a) \cup \{\neg P(x)\}$ is $\text{Th}(\mathcal{M}, R)$ -realizable. If b realizes this partial type in (\mathcal{M}, R) , then a and b have the desired properties.

Let $\Sigma(x)$ be the set of all L -formulas $\varphi(x)$ such that $(\mathcal{M}, R) \models \forall x(\neg P(x) \rightarrow \varphi(x))$. If $a \in M^n$ realizes $\Sigma(x) \cup \{P(x)\}$ in (\mathcal{M}, R) , then obviously $\text{tp}_{\mathcal{M}}(a) \cup \{\neg P(x)\}$ is $\text{Th}(\mathcal{M}, R)$ -realizable. Therefore, since (\mathcal{M}, R) is ω -saturated it suffices to prove that $\Sigma(x) \cup \{P(x)\}$ is $\text{Th}(\mathcal{M}, R)$ -realizable.

Arguing by contradiction, suppose $\Sigma(x) \cup \{P(x)\}$ is not $\text{Th}(\mathcal{M}, R)$ -realizable. Then there is a formula $\varphi(x)$ in $\Sigma(x)$ such that $(\mathcal{M}, R) \models \varphi(x) \rightarrow \neg P(x)$. Thus $(\mathcal{M}, R) \models \forall x(\neg P(x) \leftrightarrow \varphi(x))$ and hence $\neg\varphi(x)$ 0-defines R in \mathcal{M} . This contradicts the assumption that R is not 0-definable in \mathcal{M} . \square

21.2. Corollary (Svenonius's Theorem). *Let \mathcal{M} be an L -structure and let R be an n -ary relation on M . The following conditions are equivalent:*

- (1) R is 0-definable in \mathcal{M} .
- (2) For every elementary extension (\mathcal{N}, S) of (\mathcal{M}, R) , every automorphism of \mathcal{N} leaves S setwise invariant.

Proof. (1) \Rightarrow (2): If R is 0-definable in \mathcal{M} , then S is 0-definable in \mathcal{N} , by the same L -formula, and thus S would be invariant under every automorphism of \mathcal{N} .

(1) \Rightarrow (2): By Theorem 13.3 we may take (\mathcal{N}, S) to be an elementary extension of (\mathcal{M}, R) such that (\mathcal{N}, S) is ω -saturated and \mathcal{N} is strongly ω -homogeneous. Note that S is 0-definable in \mathcal{N} if and only if R is 0-definable in \mathcal{M} . Now apply Theorem 21.1 to (\mathcal{N}, S) . \square

21.3. Remark. For applications, it is useful to note that $R \subseteq M^n$ is A -definable in \mathcal{M} iff R is 0-definable in $\mathcal{M}_A = (\mathcal{M}, a)_{a \in A}$. Further, if \mathcal{M} is κ -saturated (resp. strongly κ -homogeneous) and $A \subseteq M$ has cardinality $< \kappa$, then \mathcal{M}_A is also κ -saturated (resp. strongly κ -homogeneous). Hence we get a suitable version of the previous results for A -definability in place of 0-definability.

We give another application of Theorem 21.1:

21.4. Corollary. *Let \mathcal{M} be a κ -saturated and strongly κ -homogeneous L -structure and let $A \subseteq M$ have cardinality $< \kappa$. Let $\text{Aut}_A(\mathcal{M})$ denote the group of all automorphisms of \mathcal{M} that fix A pointwise. Fix $u \in M^n$. Then:*

- (1) u is A -definable in \mathcal{M} iff u is fixed by every $\tau \in \text{Aut}_A(\mathcal{M})$.
- (2) u is algebraic over A in \mathcal{M} iff the set $\{\tau(u) \mid \tau \in \text{Aut}_A(\mathcal{M})\}$ is finite.

Proof. The direction \Rightarrow in both statements is trivial. We prove the direction \Leftarrow in (2); the corresponding proof for (1) is essentially the same.

Assume that $R = \{\tau(u) \mid \tau \in \text{Aut}_A(\mathcal{M})\} \subseteq M^n$ is finite. Note that the automorphism group of \mathcal{M}_A is exactly $\text{Aut}_A(\mathcal{M})$. Our assumptions on \mathcal{M} , A , and R ensure that \mathcal{M}_A is strongly κ -homogeneous and that (\mathcal{M}_A, R) is

κ -saturated. Therefore we may apply Theorem 21.1. Since R is given to be setwise invariant under every automorphism of \mathcal{M}_A , this yields that R is 0-definable in \mathcal{M}_A . That is, there is an L_A -formula $\varphi(x)$ whose interpretation in \mathcal{M}_A is exactly the finite set R . Since $u \in R$, this shows that u is algebraic over A . \square

21.5. Definition. Let Σ be a satisfiable set of sentences in a language that contains $L(P)$, where P is an n -ary predicate symbol not in L .

(a) We say Σ 0-defines P explicitly over L if there is an L -formula $\varphi(x_1, \dots, x_n)$ such that $\Sigma \models \forall x_1 \dots \forall x_n (P(x) \leftrightarrow \varphi(x))$.

(b) We say Σ 0-defines P implicitly over L if for any models \mathcal{M} and \mathcal{N} of Σ that have the same reduct to L , we have that $P^{\mathcal{M}}$ and $P^{\mathcal{N}}$ are identical.

21.6. Theorem (Beth's Definability Theorem). *Let T be a satisfiable theory in a language L' that contains $L(P)$, where P is an n -ary predicate symbol. Then T 0-defines P explicitly over L if and only if T 0-defines P implicitly over L .*

Proof. (\Rightarrow) Obvious.

(\Leftarrow) Assume that T 0-defines P implicitly over L but that T does not 0-define P explicitly. Let x be the sequence x_1, \dots, x_n of distinct variables.

Consider the L' -theory T' consisting of T together with all sentences of the form $\neg \forall x (P(x) \leftrightarrow \varphi(x))$ where $\varphi(x)$ is an L -formula. We claim that T' is unsatisfiable. Note that if (\mathcal{M}, R) is the reduct to $L(P)$ of a model of T' , then R cannot be 0-definable in \mathcal{M} .

If T' is satisfiable, use Theorem 13.3 to get an ω -saturated model \mathcal{M}' of T' such that every reduct of \mathcal{M}' to a sublanguage of $L(P)$ is strongly ω -homogeneous. Let \mathcal{M} be the reduct of \mathcal{M}' to L and $R = P^{\mathcal{M}'}$. Since R is not 0-definable in \mathcal{M} , by Theorem 21.1 there is an automorphism σ of \mathcal{M} such that $\sigma R \neq R$. Let \mathcal{N}' be the unique L' -structure with underlying set A that is determined by requiring that the function σ is an isomorphism from \mathcal{M}' onto \mathcal{N}' . Then \mathcal{N}' is a model of T , the reduct of \mathcal{N}' to L is \mathcal{M} (since σ is an automorphism of \mathcal{M}), and $P^{\mathcal{N}'} = \sigma(R) \neq R = P^{\mathcal{M}'}$. This contradicts the assumption that T 0-defines P implicitly over L .

Therefore T' is unsatisfiable. Consequently, there exist L -formulas $\varphi_1(x), \dots, \varphi_k(x)$ such that T together with the sentences $\neg \forall x (P(x) \leftrightarrow \varphi_j(x))$ ($j = 1, \dots, k$) is unsatisfiable. Therefore

$$T \models \forall x (P(x) \leftrightarrow \varphi_1(x)) \vee \dots \vee \forall x (P(x) \leftrightarrow \varphi_k(x)).$$

Let T_0 be the $L(P)$ -theory consisting of all $L(P)$ -sentences in T . In particular the sentence

$$\forall x (P(x) \leftrightarrow \varphi_1(x)) \vee \dots \vee \forall x (P(x) \leftrightarrow \varphi_k(x))$$

is in T_0 . Therefore we see that T_0 0-defines P implicitly over L .

For each $j = 1, \dots, k$ and each $L(P)$ -formula ψ let $\psi^{(j)}$ denote the L -formula that results from ψ by replacing every occurrence of $P(u)$ by $\varphi_j(u)$ with suitable change of bound variables in φ_j . Finally, let $T_j = \{\psi^{(j)} \mid \psi \text{ is an } L(P)\text{-sentence and } T_0 \models \psi\}$. Note that if \mathcal{M} is any model of T_j and we set $R = \{a \in A^n \mid \mathcal{M} \models \varphi_j[a]\}$, then $(\mathcal{M}, R) \models T_0$.

Fix $j \in \{1, \dots, k\}$. Note that the $L(P)$ -theories $T_0 \cup T_j$ and $T_0 \cup \{\forall x(P(x) \leftrightarrow \varphi_j(x))\}$ are equivalent; indeed, their models are exactly the $L(P)$ -structures (\mathcal{M}, R) where $\mathcal{M} \models T_j$ and $R = \{a \in A^n \mid \mathcal{M} \models \varphi_j[a]\}$. Therefore there is an $L(P)$ -sentence ψ_j such that $T_0 \models \psi_j$ and

$$T_0 \models \psi_j^{(j)} \leftrightarrow \forall x(P(x) \leftrightarrow \varphi_j(x)).$$

Therefore

$$T_0 \models \forall x(P(x) \leftrightarrow \bigwedge_{j=1}^k (\psi_j^{(j)} \rightarrow \varphi_j(x))).$$

But this sentence is in $L(P)$ and hence we have

$$T \models \forall x(P(x) \leftrightarrow \bigwedge_{j=1}^k (\psi_j^{(j)} \rightarrow \varphi_j(x))).$$

showing that T explicitly 0-defines P over L . □

21.7. Fact. Let T be a satisfiable theory in a language that contains $L(F)$, where F is an n -ary function symbol (a constant symbol if $n = 0$). We say that T 0-defines F explicitly over L if there is an L -formula $\varphi(x_1, \dots, x_n, y)$ such that

$$T \models \forall x_1 \dots \forall x_n \forall y (F(x_1, \dots, x_n) = y \leftrightarrow \varphi(x_1, \dots, x_n, y)).$$

Further, T 0-defines F implicitly over L if for any models \mathcal{M}, \mathcal{N} of T that have the same reduct to L , $F^{\mathcal{M}}$ and $F^{\mathcal{N}}$ are identical.

Beth's Definability Theorem then holds also for functions: T 0-defines F explicitly over L if and only if T 0-defines F implicitly over L . This follows immediately from what is proved above, by applying 21.6 to a new $(n + 1)$ -ary relation symbol P and adding the axiom $\forall x_1 \dots \forall x_n \forall y (P(x_1, \dots, x_n, y) \iff F(x_1, \dots, x_n) = y)$.

21.8. Definition. Let $L \subseteq L'$ be first order languages and let $T \subseteq T'$ be theories in these languages. T' is an *extension by definitions* of T if T' is a conservative extension of T and if every formula in L' is equivalent in T' to a formula in L .

21.9. Fact. Suppose $L \subseteq L'$, $T \subseteq T'$ and assume that T' is a conservative extension of T . Assume further that for every simple atomic formula φ' in L' there exists a formula φ in L s.t. $T' \models \varphi' \leftrightarrow \varphi$. Then T' is an extension by definitions of T . (By a "simple atomic formula" we mean one of the form $P(v_1, \dots, v_n)$ or $f(v_1, \dots, v_n) = w$ or $c = w$, where v_1, \dots, v_n and w are distinct variables.)

21.10. Corollary. *Let $L \subseteq L'$ be first order languages and let $T \subseteq T'$ be theories in these languages. Then T' is an extension by definitions of T if and only if every model of T has a unique expansion that is a model of T' .*

Proof. Suppose T' is an extension by definitions of T . Let \mathcal{M} be any model of T . If P is any predicate symbol of L' , then there exists a formula φ in L such that $T' \models \forall x_1 \dots \forall x_n (P(x_1, \dots, x_n) \leftrightarrow \varphi(x_1, \dots, x_n))$. This gives an interpretation of P on A . This interpretation is well defined because, if φ_1 and φ_2 are two formulas in L such that $T' \models \forall x_1 \dots \forall x_n (P(x_1, \dots, x_n) \leftrightarrow \varphi_i(x_1, \dots, x_n)), i = 1, 2$, then $T' \models \forall x_1 \dots \forall x_n (\varphi_1(x_1, \dots, x_n) \leftrightarrow \varphi_2(x_1, \dots, x_n))$, so $T \models \forall x_1 \dots \forall x_n (\varphi_1(x_1, \dots, x_n) \leftrightarrow \varphi_2(x_1, \dots, x_n))$, since T' is a conservative extension of T .

Similarly, if F is any function symbol of L' , then there exists a formula φ in L such that $T' \models \forall x_1 \dots \forall x_n \forall y (F(x_1, \dots, x_n) = y \leftrightarrow \varphi(x_1, \dots, x_n, y))$. Note that this implies that $T \models \forall x_1 \dots \forall x_n \exists! y \varphi(x_1, \dots, x_n, y)$. Therefore the formula $\varphi(x_1, \dots, x_n, y)$ defines on every model of T the graph of a totally defined function. This gives a well-defined interpretation of F on A . Similarly for constant symbols c in L' , using the formula $c = y$ in the same way.

An easy induction argument on formulas shows that this expansion of \mathcal{M} is a model of T' . Furthermore, it is the only such model, because any model of T' has to interpret the relation, function and constant symbols of L' according to the L -formulas by which they are explicitly 0-defined in T' .

For the converse, let T' be an L' theory such that every model of T has a unique expansion that is a model of T' .

First we show that T' is a conservative expansion of T . Let σ be any L -sentence proved by T' . If \mathcal{M} is any model of T , then it has a (unique) expansion \mathcal{M}' that is a model of T' . Since σ is true in \mathcal{M}' , and σ is an L -sentence, it must be true in \mathcal{M} . Therefore T proves σ .

To complete the proof we need to show that every formula in L' is equivalent in T' to a formula in L . This is an immediate consequence of Fact 21.9 and Beth's Definability Theorem (using Fact 21.7 in the case of function symbols and constant symbols). \square

21.11. Definition. Let T_i be a theory in L_i for each $i = 1, 2$, where L_1 and L_2 do not have any nonlogical symbols in common. Call T_1 and T_2 *equivalent by definitions* if there is a theory T in some language L that contains the union of L_1 and L_2 such that T is simultaneously an extension by definitions of T_1 and of T_2 .

When T_1 and T_2 are equivalent by definitions, we may regard them as being interchangeable. In particular, there is a bijective correspondence between models of T_1 and models of T_2 that preserves all properties of mathematical

significance. Namely, expand any model of T_1 to a model of T and then take the reduct of this model to L_2 .

21.12. Example. Let T_1 be the theory of Boolean rings in the language L_1 with nonlogical symbols $\{+, -, \times, 0, 1\}$. Let T_2 be the theory of Boolean algebras in the language L_2 with nonlogical symbols $\{\wedge, \vee, (\cdot)^c, 0, 1\}$. As is well known, every Boolean ring can be regarded as a Boolean algebra, and vice versa. This is because of the fact that in a Boolean ring $+, -, \times, 0, 1$ can be defined in terms of $\wedge, \vee, (\cdot)^c, 0, 1$ and vice versa. If T is the theory axiomatized by the sentences that express these definitions, then T is easily seen to be an extension by definitions of both T_1 and T_2 . Therefore T_1 and T_2 are equivalent by definitions. This model theoretic fact expresses in a complete way the relation between Boolean rings and Boolean algebras.

In the rest of this chapter we discuss Craig's Interpolation Theorem, which gives another important characteristic property of first order logic. At the end of this chapter we use Craig's Theorem to give another proof of Beth's Theorem.

21.13. Theorem (Robinson's Joint Consistency Lemma). *Let L be a first order language and let L_1 and L_2 be extensions of L whose intersection is L . For each $i = 1, 2$ let T_i be a satisfiable theory in L_i . If there is a complete theory T in L such that $T \subseteq T_1 \cap T_2$, then $T_1 \cup T_2$ is satisfiable.*

For the proof of this theorem we need the following preliminary result:

21.14. Lemma. *Let L be a first order language and let L_1 and L_2 be extensions of L whose intersection is L . Suppose \mathcal{M}_j is an L_j structure for $j = 1, 2$, and $\mathcal{M}_1 \upharpoonright L \equiv \mathcal{M}_2 \upharpoonright L$. Then there exists an elementary extension \mathcal{M}'_1 of \mathcal{M}_1 and $f: \mathcal{M}_2 \rightarrow \mathcal{M}'_1$ such that f is an elementary embedding of $\mathcal{M}_2 \upharpoonright L$ into $\mathcal{M}'_1 \upharpoonright L$.*

Proof. For any L -structure \mathcal{M} , recall that $\text{EDiag}(\mathcal{M})$ denote the first order theory of the L_M -structure $(\mathcal{M}, a)_{a \in M}$. It is a simple but key fact that \mathcal{M} can be elementarily embedded into an L -structure \mathcal{N} iff \mathcal{N} has an expansion that is a model of $\text{EDiag}(\mathcal{M})$. (See Exercise 4.22.)

We first show that to prove the Lemma it suffices to prove that $\Sigma = \text{EDiag}(\mathcal{M}_1) \cup \text{EDiag}(\mathcal{M}_2 \upharpoonright L)$ is satisfiable. If so, let \mathcal{M}'_1 be the reduct of a model of Σ to L_1 . Without loss of generality we may assume that for each $a \in M_1$, the interpretation of a in \mathcal{M}'_1 is a itself. This implies that \mathcal{M}'_1 is an elementary extension of \mathcal{M}_1 . Moreover, there is an elementary embedding of $\mathcal{M}_2 \upharpoonright L$ into $\mathcal{M}'_1 \upharpoonright L$ because \mathcal{M}'_1 has an expansion that is a model of $\text{EDiag}(\mathcal{M}_2 \upharpoonright L)$.

Arguing by contradiction, suppose $\text{EDiag}(\mathcal{M}_1) \cup \text{EDiag}(\mathcal{M}_2 \upharpoonright L)$ is not satisfiable. By the Compactness Theorem and the fact that $\text{EDiag}(\mathcal{M}_1)$ and $\text{EDiag}(\mathcal{M}_2 \upharpoonright L)$ are closed under conjunction, there exist $\sigma_1 \in \text{EDiag}(\mathcal{M}_1)$ and $\sigma_2 \in \text{EDiag}(\mathcal{M}_2 \upharpoonright L)$ such that $\{\sigma_1, \sigma_2\}$ has no model.

There exist $a_1, \dots, a_m \in M_1, b_1, \dots, b_n \in M_2$, and L -formulas τ_1, τ_2 , such that $\sigma_1 = \tau_1(a_1, \dots, a_m)$ and $\sigma_2 = \tau_2(b_1, \dots, b_n)$. We may assume that $M_1 \cap M_2$ is empty, and also that τ_1 and τ_2 have no free variables in common. Let the free variables in τ_1 be $y = y_1, \dots, y_m$ and those of τ_2 be $z = z_1, \dots, z_n$. Since $\{\sigma_1, \sigma_2\}$ has no model, it follows that $\exists y \tau_1 \wedge \exists z \tau_2$ has no model. But $\exists z \tau_2$ is true in \mathcal{M}_2 , because $\sigma_2 \in \text{EDiag}(\mathcal{M}_2|L)$. Moreover, $\mathcal{M}_1|L \equiv \mathcal{M}_2|L$, so $\exists z \tau_2$ is true in \mathcal{M}_1 as well. Further, $\exists y \tau_1$ is true in \mathcal{M}_1 , because $\sigma_1 \in \text{EDiag}(\mathcal{M}_1)$. Therefore \mathcal{M}_1 is a model of $\exists y \tau_1 \wedge \exists z \tau_2$. This is a contradiction. \square

Proof of Theorem 21.13. We are given, for $i = 1, 2$, a satisfiable L_i -theory T_i ; $T = T_1 \cap T_2$ is assumed to be a complete theory in $L = L_1 \cap L_2$. We wish to show that $T_1 \cup T_2$ is satisfiable.

Let \mathcal{M}_1 be a model of T_1 , and let \mathcal{N}_1 be a model of T_2 . Now $T = \text{Th}(\mathcal{M}_1|L) = \text{Th}(\mathcal{N}_1|L)$ since T is complete; therefore $\mathcal{M}_1|L \equiv \mathcal{N}_1|L$. By the preceding Lemma there is a model $\mathcal{N}_2 \succeq \mathcal{N}_1$ and a map $f_1 : M_1 \rightarrow N_2$ that elementarily embeds $\mathcal{M}_1|L$ into $\mathcal{N}_2|L$. Next we apply the Lemma to $(\mathcal{M}_1, a)_{a \in M_1}$ in the language L_{1, M_1} and $(\mathcal{N}_2, f_1(a))_{a \in M_1}$ in the language L_{2, M_1} . We then have

$$(\mathcal{M}_1|L, a)_{a \in M_1} \equiv (\mathcal{N}_2|L, f_1(a))_{a \in M_1},$$

and these two structures are the reducts to L_{M_1} of $(\mathcal{M}_1, a)_{a \in M_1}$ and $(\mathcal{N}_2, f_1(a))_{a \in M_1}$ (respectively). Using the lemma again, we see that there exists an elementary extension $(\mathcal{M}_2, a)_{a \in M_1}$ of $(\mathcal{M}_1, a)_{a \in M_1}$ and an elementary embedding g_1 of $(\mathcal{N}_2|L, f_1(a))_{a \in M_1}$ into $(\mathcal{M}_2|L, a)_{a \in M_1}$. Note that these last two structures are reducts to the language L_{M_1} , which is the intersection of L_{1, M_1} and L_{2, M_1} . So we have $\mathcal{N}_1 \preceq \mathcal{N}_2, \mathcal{M}_1 \preceq \mathcal{M}_2$ and maps $f_1 : \mathcal{M}_1 \rightarrow \mathcal{N}_2$ and $g_1 : \mathcal{N}_2 \rightarrow \mathcal{M}_2$ that are elementary embeddings with respect to formulas in the language L . In addition, we have that $g_1(f_1(a)) = a$ for each $a \in M_1$.

We continue inductively in this way. The result is a pair of elementary chains $\mathcal{M}_1 \preceq \mathcal{M}_2 \preceq \mathcal{M}_3 \preceq \dots$ and $\mathcal{N}_1 \preceq \mathcal{N}_2 \preceq \mathcal{N}_3 \preceq \dots$ and mappings $f_n : \mathcal{M}_n \rightarrow \mathcal{N}_{n+1}$ and $g_n : \mathcal{N}_{n+1} \rightarrow \mathcal{M}_{n+1}$ that are elementary embeddings with respect to formulas of L and that satisfy $g_n(f_n(a)) = a$ for all $a \in M_n$ and $f_{n+1}(g_n(b)) = b$ for all $b \in N_{n+1}$. Note that for all $n \geq 1, f_{n+1} = f_n$ on M_n and $g_{n+1} = g_n$ on N_{n+1} .

Now let $\mathcal{M} = \cup \mathcal{M}_n, \mathcal{N} = \cup \mathcal{N}_n, f = \cup f_n$, and $g = \cup g_n$. We have $\mathcal{M} \models T_1$ since $\mathcal{M}_1 \models T_1$ and $\mathcal{M}_1 \preceq \mathcal{M}$; similarly $\mathcal{N} \models T_2$. Moreover we see that f is an isomorphism of $\mathcal{M}|L$ onto $\mathcal{N}|L$ whose inverse is g . We can replace \mathcal{N} by an isomorphic copy \mathcal{N}' such that $\mathcal{N}'|L = \mathcal{M}|L$, using the mapping f to rename all elements. It then follows that we can define a structure \mathcal{C} for $L_1 \cup L_2$ so that $\mathcal{C}|L_1 = \mathcal{M}$ and $\mathcal{C}|L_2 = \mathcal{N}'$. The fact that $\mathcal{N}'|L = \mathcal{M}|L$ guarantees that the interpretations of symbols of $L_1 \cap L_2$ are well-defined. We see that \mathcal{C} is necessarily a model for $T = T_1 \cup T_2$, because $\mathcal{C}|L_1 = \mathcal{M} \models T_1$ and $\mathcal{C}|L_2 = \mathcal{N}' \models T_2$. This completes the proof. \square

21.15. Theorem (Craig's Interpolation Theorem). *Let L be a first order language and let φ and ψ be L -sentences such that $\varphi \models \psi$. Then there is a sentence θ such that*

(i) $\varphi \models \theta$ and $\theta \models \psi$, and

(ii) every predicate, function, or constant symbol (excluding equality) that occurs in θ occurs also in both φ and ψ .

Proof. Assume $\varphi \models \psi$; let L_1 be the language of φ and L_2 that of ψ ; take L to be the common language, containing $=$ at least. It suffices to show that $T_0 \models \psi$ where $T_0 = \{\sigma \in L : \varphi \models \sigma\}$; if this holds, then there is a finite subset F of T_0 such that $F \models \psi$. Taking θ to be the conjunction of the formulas in F will give the desired sentence. If $T_0 \not\models \psi$ then $T_0 \cup \{\neg\psi\}$ is satisfiable. Let T_1 be a complete extension in L_2 of $T_0 \cup \{\neg\psi\}$ and set $T = T_1 \cap L$, so T is a complete theory in L . We claim that $T \cup \{\varphi\}$ is satisfiable in L_1 . If not, there is a sentence $\sigma \in T$ such that $\varphi \models \neg\sigma$. But then $\neg\sigma \in T_0 \subseteq T_1$, which implies $\neg\sigma \in T$, a contradiction. We apply Theorem 21.13 to T_1 and $T \cup \{\varphi\}$. Since both sets are satisfiable and since T is complete, $T \cup T_1 \cup \{\varphi\}$ is satisfiable. In particular $\{\varphi, \neg\psi\}$ is satisfiable, which is a contradiction. \square

21.16. Remark. It is possible to have sentences φ and ψ that have no predicate, function, or constant symbol in common, yet satisfy $\varphi \models \psi$. For example, φ might be unsatisfiable or ψ might be valid. If logic with identity is considered (as we do here), then there are more interesting examples, such as the following:

$$\forall x \forall y [x = y] \models \forall x \forall y [P(x) \leftrightarrow P(y)].$$

Examples like this explain why the conclusion of Craig's Theorem allows the equality symbol to occur in the interpolating sentence θ .

If in Craig's Theorem one only considers sentences φ and ψ without equality, then it can be shown that there is an interpolating sentence that contains only symbols that occur in both φ and ψ . If there are no such symbols and neither formula contains equality, then it can be shown that either φ is unsatisfiable or ψ is valid.

Next we note that Craig's Theorem yields easily a joint consistency result that is formally a strengthening of Robinson's Lemma:

21.17. Corollary (Joint Consistency Theorem). *Let L be a first order language and let L_1 and L_2 be extensions of L whose intersection is L . For each $i = 1, 2$ let T_i be a satisfiable theory in L_i . If the L -theory $T_1 \cap T_2$ is satisfiable, then the $L_1 \cup L_2$ -theory $T_1 \cup T_2$ is also satisfiable.*

Proof. If $T_1 \cup T_2$ is not satisfiable, by the Compactness Theorem and the fact that both T_1 and T_2 are closed under conjunction, there are sentences $\sigma_1 \in T_1$ and $\sigma_2 \in T_2$ such that $\{\sigma_1, \sigma_2\}$ has no model. Therefore $\sigma_1 \models \neg\sigma_2$. Applying Craig's Theorem, there is an L -sentence σ such that $\sigma_1 \models \sigma$ and

$\sigma \models \neg\sigma_2$. But this implies that $T_1 \models \sigma$ and $T_2 \models \neg\sigma$, contradicting the assumption that $T_1 \cap T_2$ is satisfiable. \square

Finally, we give an alternate proof of Beth's Definability Theorem that uses Craig's Theorem:

Assume that T 0-defines P implicitly over L . For each symbol α of L_1 that is not in L , let α' denote a symbol of the same type and arity as α , which does not occur in L_1 . Let L_2 denote the language that contains L and that contains α' for each symbol α of L_1 that is not in L . Let T' be the theory in L_2 that results from T by leaving every symbol of L unchanged and by replacing every occurrence of any other symbol α of L_1 by the corresponding symbol α' . We observe that $T \cup T' \models \forall \bar{x}[P(\bar{x}) \leftrightarrow P'(\bar{x})]$. Indeed, consider any model of $T \cup T'$ and let (\mathcal{M}, R, R') denote its reduct to $L(P, P')$. Then (\mathcal{M}, R) and (\mathcal{M}, R') are both reducts of models of T . It follows from our hypothesis (that T implicitly 0-defines P over L) that $R = R'$. Therefore there exist finite subsets $\Sigma \subseteq T$ and $\Sigma' \subseteq T'$ such that $\Sigma \cup \Sigma' \models \forall x(P(\bar{x}) \leftrightarrow P'(\bar{x}))$. By adding finitely many sentences from $T \cup T'$ to each of these finite sets, we can ensure that Σ' is precisely the result of replacing every occurrence of a symbol α of L_1 that is not in L by the corresponding α' . In particular, Σ' will contain P' in exactly the same places that Σ contains P .

We now add new constants c_1, \dots, c_n to the language of $T \cup T'$. Evidently $\Sigma \cup \Sigma' \models P(c_1, \dots, c_n) \rightarrow P'(c_1, \dots, c_n)$. Let σ be the conjunction of all the sentences in Σ and let σ' be the conjunction of the sentences in Σ' . Then $\sigma \wedge P(c_1, \dots, c_n) \models (\sigma' \rightarrow P'(c_1, \dots, c_n))$. Obviously the common language of the sentences $\sigma \wedge P(c_1, \dots, c_n)$ and $(\sigma' \rightarrow P'(c_1, \dots, c_n))$ is $L(c_1, \dots, c_n)$. Now apply Craig's Theorem to the above: there is an L -formula $\theta(x_1, \dots, x_n)$ such that $\sigma \wedge P(c_1, \dots, c_n) \models \theta(c_1, \dots, c_n)$ and $\theta(c_1, \dots, c_n) \models (\sigma' \rightarrow P'(c_1, \dots, c_n))$. That is to say:

- (a) $\Sigma \models (P(c_1, \dots, c_n) \rightarrow \theta(c_1, \dots, c_n))$ and
- (b) $\Sigma' \models (\theta(c_1, \dots, c_n) \rightarrow P'(c_1, \dots, c_n))$.

From (b) we conclude that $\Sigma \models (\theta(c_1, \dots, c_n) \rightarrow P(c_1, \dots, c_n))$. (To see this, consider a formal derivation of $(\theta(c_1, \dots, c_n) \rightarrow P'(c_1, \dots, c_n))$ from Σ' . For each symbol α of L_1 that is not in L , replace every occurrence of α' in this derivation by α . The result is a formal derivation of $(\theta(c_1, \dots, c_n) \rightarrow P(c_1, \dots, c_n))$ from Σ .) Therefore $\Sigma \models (\theta(c_1, \dots, c_n) \leftrightarrow P(c_1, \dots, c_n))$. But Σ does not in fact contain the new constants c_i , so we conclude $\Sigma \models \forall x(P(\bar{x}) \leftrightarrow \theta(\bar{x}))$. This shows that T explicitly 0-defines P , since Σ is a subset of T .

22. SYSTEMS OF DEFINABLE SETS AND FUNCTIONS

We take the point of view that Model Theory is the study of sets and functions that are definable in a given mathematical structure using formulas of first order logic with equality. In this chapter we explore the collections of definable sets and characterize them using simple “geometric” properties. This discussion can be read at any time, as it depends only on simple facts about the interpretation of first order formulas.

Let L be a first order language and \mathcal{M} an L -structure; let M be the underlying set of \mathcal{M} .

22.1. Definition. A set $X \subseteq M^m$ is *definable in \mathcal{M}* if there is an L -formula $\varphi(x_1, \dots, x_m, y_1, \dots, y_n)$ and elements b_1, \dots, b_n of M such that

$$X = \{(a_1, \dots, a_m) \in M^m \mid \mathcal{M} \models \varphi[a_1, \dots, a_m, b_1, \dots, b_n]\}.$$

If $A \subseteq M$ and the above equation holds for some φ and some $b_1, \dots, b_n \in A$, then we say that X is *A-definable in \mathcal{M}* . Let $X \subseteq M^m$ and $Y \subseteq M^n$; a function $f: X \rightarrow Y$ is *A-definable in \mathcal{M}* if the graph of f is A -definable in \mathcal{M} . (We regard the graph of f as a subset of M^{m+n} .) Note that this implies that the domain X and the range $f(X)$ of f are also A -definable in \mathcal{M} . Indeed, if $R \subseteq M^{m+n}$ is the graph of f , and we consider distinct variables $x = x_1, \dots, x_m, y = y_1, \dots, y_n$, then the domain of f is defined by the formula

$$\varphi(x) = \exists y_1 \dots \exists y_n R(x, y)$$

and the range of f is defined by

$$\psi(y) = \exists x_1 \dots \exists x_m R(x, y)$$

Now we begin to analyze the nature of the sets and functions that are definable in a given structure \mathcal{M} . We want to explain them in a way that is intelligible to any mathematician, so we move the syntax of first order logic far into the background. The most basic logical operations used to build up first order formulas are the propositional connectives \neg, \vee, \wedge and the existential quantifier $\exists x$ where x is any variable ranging over the set M . They have the following meanings:

- \neg stands for the negation, “not”,
- \vee stands for the disjunction, “or”
- \wedge stands for the conjunction, “and”,
- $\exists x$ stands for the existential quantifier, “there exists x ”.

These logical operations correspond to familiar elementary mathematical operations on sets; namely, the basic propositional connectives correspond to Boolean operations on sets and the existential quantifiers correspond to projection operations on Cartesian products. We illustrate this now in a simple setting: let x, y be variables ranging over nonempty sets M, N (respectively), and let $\varphi(x, y)$ and $\psi(x, y)$ denote conditions on (x, y) defining subsets Φ and Ψ (respectively) of $M \times N$. We consider the conditions that

can be built up from $\varphi(x, y)$ and $\psi(x, y)$ using the basic logical operations (on the left below), and the sets that are defined by them (on the right):

$\neg\varphi(x, y)$	defines	the complement of Φ in $M \times N$,
$\varphi(x, y) \vee \psi(x, y)$	defines	the union $\Phi \cup \Psi$,
$\varphi(x, y) \wedge \psi(x, y)$	defines	the intersection $\Phi \cap \Psi$,
$\exists x\varphi(x, y)$	defines	the projection $\pi(\Phi)$

where $\pi(x, y) = y$ is the projection onto the second coordinate.

To illustrate the usefulness of these simple ideas, consider a given function $f: M \rightarrow N$. The image $f(M)$ of M under f can be defined by the equivalence

$$y \in f(M) \iff \exists x[f(x) = y].$$

Let Γ be the graph of f , which is defined as a subset of $M \times N$ by the condition $f(x) = y$. The displayed equivalence exhibits the fact that $f(M)$ is the projection of Γ under the projection map π onto the second coordinate.

There are three other logical operations that are often used in mathematics:

\rightarrow	stands for	the implication, “if . . . , then”,
\leftrightarrow	stands for	the equivalence, “if and only if”,
$\forall x$	stands for	the universal quantifier, “for all x .”

As is familiar, these operations can be defined in terms of the basic ones. Indeed, $\varphi \rightarrow \psi$ is equivalent to $\neg\varphi \vee \psi$, $\varphi \leftrightarrow \psi$ is equivalent to $(\varphi \wedge \psi) \vee (\neg\varphi \wedge \neg\psi)$ and $\forall x\varphi$ is equivalent to $\neg\exists x\neg\varphi$. Therefore we see that these three logical operations correspond to elementary set operations that can be constructed by applying the basic ones several times.

Simple and familiar logical equivalences often capture mathematical facts that seem complicated when viewed without the use of logical notation. For example, the familiar equivalence

$$\forall y\varphi(x, y) \iff \neg\exists y\neg\varphi(x, y)$$

shows that the set defined by $\forall y\varphi(x, y)$ can be obtained from Φ by first taking the complement in $M \times N$, then projecting onto the first coordinate, and then taking the complement of that set in M . This technique is particularly useful when dealing with logically complicated notions, such as continuity or differentiability, which we express in the usual way with ϵ 's and δ 's and quantifiers over them. In such cases we often deal with conditions having more than two variables and with repeated quantifiers.

We use several additional notational conventions. A condition $\varphi(x, y)$ defining a subset of $M \times N$ is sometimes viewed as defining a condition on triples (x, y, z) , where z ranges over a nonempty set P ; in that case $\varphi(x, y)$ defines a subset of $M \times N \times P$. In such a situation we indicate the condition also as $\varphi(x, y, z)$. This is similar to the situation in algebra where one routinely regards a polynomial $p(x, y)$ as a polynomial in three variables x, y, z in which all monomials containing z are taken to have coefficient 0.

It is also useful to consider conditions obtained by substitution. As above, let $f: M \rightarrow N$ be a function and let Γ be the graph of f . The condition $\varphi(x, f(x))$ defines a subset X of M . This condition is equivalent to

$$\exists y[f(x) = y \wedge \varphi(x, y)].$$

Therefore X can be obtained by applying the projection π' to $\Gamma \cap \Phi$, where $\pi'(x, y) = x$ is the projection of $M \times N$ onto the first coordinate.

We will show that the essential features of the collection of A -definable sets in \mathcal{M} are captured by the following definition:

22.2. Definition. Let M be a nonempty set. A *definability system on M* is a sequence $\mathcal{S} = (\mathcal{S}_m)_{m \in \mathbb{N}}$ such that for each $m \geq 0$:

- (1) \mathcal{S}_m is a Boolean algebra of subsets of M^m that contains \emptyset and M^m as elements;
- (2) if $X \in \mathcal{S}_m$, then $M \times X$ and $X \times M$ belong to \mathcal{S}_{m+1} ;
- (3) $\{(a_1, \dots, a_m) \in M^m \mid a_1 = a_m\} \in \mathcal{S}_m$;
- (4) if $X \in \mathcal{S}_{m+1}$, then $\pi(X) \in \mathcal{S}_m$, where $\pi: M^{m+1} \rightarrow M^m$ is the projection map on the first m coordinates.

If $X \subseteq M^m$ we say X *belongs to \mathcal{S}* if $X \in \mathcal{S}_m$. If $X \subseteq M^m$ and $Y \subseteq M^n$ and if $f: X \rightarrow Y$ is a function, then we say f *belongs to \mathcal{S}* if the graph of f belongs to \mathcal{S} .

22.3. Notation. Let L be a first order language and \mathcal{M} an L -structure with underlying set M ; let A be any subset of M . We write $\mathcal{D}(\mathcal{M}, A)$ for the system $(\mathcal{S}_m)_{m \in \mathbb{N}}$, where for each $m \geq 0$, \mathcal{S}_m is the collection of all subsets of M^m that are A -definable in M .

22.4. Proposition. *Let L be a first order language and \mathcal{M} an L -structure with underlying set M ; let A be any subset of M . Then $\mathcal{D}(\mathcal{M}, A)$ is a definability system on M .*

Proof. Exercise. The informal remarks above make it easy to prove this result. □

Our next result is a converse to Proposition 22.4. It states that each definability system is *closed under definability*. If M is a nonempty set and $X \subseteq M^k$, then for each $m \geq 0$ we regard X^m as a subset of M^{km} .

22.5. Theorem. *Let M be a nonempty set and let \mathcal{S} be a definability system on M . Let L be a first order language and \mathcal{N} an L -structure whose underlying set is N ; let A be a subset of N . Suppose all of the following sets belong to \mathcal{S} :*

- (a) N ;
- (b) $\{c^N\}$ for each constant symbol c in L ;
- (c) $\{s\}$ for each $s \in A$;

- (d) $R^{\mathcal{N}}$ for each relation symbol R in L ;
(e) the graph of $f^{\mathcal{N}}$ for each function symbol f in L .
Then every set that is A -definable in \mathcal{N} belongs to \mathcal{S} .

We first give a series of basic results about definability systems that will be used in the proof of Theorem 22.5. For these lemmas we fix a nonempty set M and a definability system \mathcal{S} on M .

22.6. Lemma. *If X and Y belong to \mathcal{S} , then $X \times Y$ belongs to \mathcal{S} .*

Proof. Suppose $X \subseteq M^m$ and $Y \subseteq M^n$. Then

$$X \times Y = (X \times M^n) \cap (M^m \times Y).$$

Condition (2) of Definition 22.2 (used repeatedly) followed by condition (1) yields that this set belongs to \mathcal{S} . \square

22.7. Lemma. *For all $1 \leq i < j \leq m$, the diagonal set*

$$\Delta^m(i, j) := \{(a_1, \dots, a_m) \in X^m \mid a_i = a_j\}$$

belongs to \mathcal{S} .

Proof. Let $k := j - i + 1$. Condition (3) in Definition 22.2 gives that the diagonal set $\Delta^k(1, j - i + 1)$ belongs to \mathcal{S} , and

$$\Delta^m(i, j) = X^{i-1} \times \Delta^k(1, j - i + 1) \times X^{m-j}.$$

This set belongs to \mathcal{S} by repeated use of condition (2) of Definition 22.2. (See also Lemma 22.6.) \square

22.8. Lemma. *Let $X \in \mathcal{S}_n$ and let $i(1), \dots, i(n) \in \{1, \dots, m\}$. Then the set $Y \subseteq M^m$ defined by*

$$Y := \{(a_1, \dots, a_m) \in X^m \mid (a_{i(1)}, \dots, a_{i(n)}) \in X\}$$

belongs to \mathcal{S} .

Proof. Note that for any $a_1, \dots, a_m \in M$, the tuple (a_1, \dots, a_m) is in Y iff

$$\exists y_1 \dots \exists y_n (x_{i(1)} = y_1 \wedge \dots \wedge x_{i(n)} = y_n \wedge (y_1, \dots, y_n) \in X).$$

Let D_j denote the diagonal set $\Delta^{m+n}(i(j), m + j)$ and let $\pi_j: M^{m+j} \rightarrow M^{m+j-1}$ denote the projection map onto the first coordinates, for each $j = 1, \dots, n$. The displayed condition shows that

$$X = \pi_1(\dots \pi_n(D_1 \cap \dots \cap D_n \cap (X^m \times Y)) \dots).$$

Using Definition 22.2 and Lemma 22.7 we see that $X \in \mathcal{S}_m$. \square

22.9. Lemma. *Suppose $X \subseteq M^m$, $Y \subseteq M^n$, and $Z \subseteq M^p$ all belong to \mathcal{S} . Let $f: X \rightarrow Y$ and $g: Y \rightarrow Z$ be functions that belong to \mathcal{S} . Then their composition $g \circ f: X \rightarrow Z$ also belongs to \mathcal{S} .*

Proof. Let x be a tuple of variables that ranges over M^m and let y range over M^n and z range over M^p similarly. Use Lemma 22.8 and the equivalence

$$(x, z) \in \Gamma(g \circ f) \Leftrightarrow \exists y((x, y) \in \Gamma(f) \wedge (y, z) \in \Gamma(g)).$$

□

22.10. Lemma. *Suppose $X \subseteq M^m$ belongs to \mathcal{S} . Let $f = (f_1, \dots, f_n): X \rightarrow M^n$ be a function with coordinate functions $f_j: X \rightarrow M$. The function f belongs to \mathcal{S} if and only if all of the coordinate functions f_j belong to \mathcal{S} .*

Proof. (\Rightarrow) Fix j ($1 \leq j \leq n$) and let $\pi_j: M^n \rightarrow M$ be the projection map defined by $\pi_j(x_1, \dots, x_n) = x_j$. Using Lemma 22.7 we see that π_j belongs to \mathcal{S} , since its graph is a diagonal set. Noting that $f_j = \pi_j \circ f$, Lemma 22.9 completes the proof of this direction.

(\Leftarrow) Let x range over M^m and $y = (y_1, \dots, y_n)$ range over M^n . The graph of f is defined by the equivalence

$$(x, y) \in \Gamma(f) \Leftrightarrow ((x, y_1) \in \Gamma(f_1) \wedge \dots \wedge (x, y_n) \in \Gamma(f_n)).$$

If the functions f_1, \dots, f_n all belong to \mathcal{S} , this equivalence together with Lemma 22.8 and condition (1) of Definition 22.2 show that f belongs to \mathcal{S} . □

PROOF OF THEOREM 22.5.

Let L , \mathcal{N} , and A be as in the statement of the Theorem, and let \mathcal{S} be any definability system to which all sets listed in conditions (a)–(e) of the Theorem belong. We must show that every A -definable set in \mathcal{N} belongs to \mathcal{S} .

Let k be such that $N \subseteq M^k$. As noted above, we consider N^m as a subset of M^{km} for each $m \geq 0$.

First we prove the following statement by induction on the complexity of terms:

Let t be an L_A -term, and let x_1, \dots, x_m be a sequence of distinct variables that includes all variables of t ; the function $t^{\mathcal{N}}: N^m \rightarrow N$ defined by interpreting t in \mathcal{N} belongs to the definability system \mathcal{S} .

In the basic step of this induction t is either a constant symbol c or an element of A , or one of the variables x_i . In the first case the graph of the function $t^{\mathcal{N}}$ is $N^m \times \{c^{\mathcal{N}}\}$, and the second case is similar; in the third case it is the intersection of N^{m+1} with k diagonal sets. In each case this shows the graph belongs to \mathcal{S} .

For the induction step, we consider the case where t is of the form $f(t_1, \dots, t_n)$ where f is an n -ary function symbol of L and t_1, \dots, t_n are L_A -terms of which the statement being proved is true. Let $G: N^m \rightarrow N^n$ be the function with coordinate functions $t_j^{\mathcal{N}}$, $j = 1, \dots, n$. Lemma 22.10

shows that G belongs to \mathcal{S} ; Lemma 22.9 shows that $t^{\mathcal{N}} = f^{\mathcal{N}} \circ G$ belongs to \mathcal{S} . This completes the inductive proof of this statement about terms.

Now we prove the following statement about formulas from which Theorem 22.5 follows immediately; the proof is by induction on formulas:

Let φ be an L_A -formula and let x_1, \dots, x_m be a sequence of distinct variables that includes all free variables of φ ; the set

$$\varphi^{\mathcal{N}} := \{(a_1, \dots, a_m) \in N^m \mid \mathcal{N} \models \varphi[a_1, \dots, a_m]\}$$

belongs to the definability system \mathcal{S} .

In the basic step of this induction φ is an atomic formula of the form $R(t_1, \dots, t_n)$ where R is an n -ary relation symbol of L and t_1, \dots, t_n are $L(A)$ -terms. Let $G: N^m \rightarrow N^n$ be the function defined above using the terms t_1, \dots, t_n . As shown there, G belongs to \mathcal{S} . We see that $\varphi^{\mathcal{N}}$ is defined by the equivalence

$$(a_1, \dots, a_m) \in \varphi^{\mathcal{N}} \Leftrightarrow$$

$$\exists y_1 \dots \exists y_n ((a_1, \dots, a_m, y_1, \dots, y_n) \in \Gamma(G) \wedge (y_1, \dots, y_n) \in R^{\mathcal{N}}).$$

(Strictly speaking note that each $\exists y_j$ stands for a sequence of k existential quantifiers over M .) This shows that $\varphi^{\mathcal{N}}$ is the result of applying a sequence of kn projections to the set

$$\Gamma(G) \cap (N^m \times R^{\mathcal{N}})$$

which shows that $\varphi^{\mathcal{N}}$ belongs to \mathcal{S} .

Now we consider the cases of the induction step where φ is constructed from formulas α and β using propositional connectives. We have a list x_1, \dots, x_m of distinct variables that include all free variables of φ , and thus they also include all free variables of α and β . We apply the induction hypothesis to α and β and this list of variables, obtaining that the sets $\alpha^{\mathcal{N}}$ and $\beta^{\mathcal{N}}$, which are both subsets of N^m , belong to \mathcal{S} . It follows immediately from condition (1) of Definition 22.2 that $\varphi^{\mathcal{N}}$ also belongs to \mathcal{S} .

The other case of the induction step concerns the situation where φ is of the form $\exists y\psi$. We may assume that y is not in the list of variables x_1, \dots, x_m . (Otherwise perform a change of bound variables that replaces y by some completely new variable. Since this does not increase the complexity of ψ , we may still apply the induction hypothesis to the new situation.) If y is not in the list x_1, \dots, x_m , then we apply the induction hypothesis to the formula ψ and the list of variables x_1, \dots, x_m, y . Evidently $\varphi^{\mathcal{N}} = \pi(\psi^{\mathcal{N}})$, where $\pi: N^{m+1} \rightarrow N^m$ is the projection on the first m coordinates. Condition (4) of Definition 22.2 yields that $\varphi^{\mathcal{N}}$ belongs to \mathcal{S} . This completes the proof of Theorem 22.5. \square

Let M be a nonempty set. Given two definability systems $\mathcal{S}(1)$ and $\mathcal{S}(2)$ on M , we say that $\mathcal{S}(2)$ contains $\mathcal{S}(1)$, and we write $\mathcal{S}(1) \subseteq \mathcal{S}(2)$, if $\mathcal{S}(1)_m \subseteq \mathcal{S}(2)_m$ for all $m \geq 0$. This defines a partial ordering on the collection of all definability systems on X . Any family $(\mathcal{S}(i))_{i \in I}$ of definability systems on

X has a greatest lower bound \mathcal{S} in the collection of all definability systems on M ; namely, we just take

$$\mathcal{S}_m = \bigcap \{\mathcal{S}(i)_m \mid i \in I\} \text{ for each } m.$$

Suppose $\mathcal{F} = (\mathcal{F}_m)_{m \in \mathbb{N}}$ where \mathcal{F}_m is a collection of subsets of M^m for each $m \geq 0$. Obviously there is at least one definability system \mathcal{S} that contains \mathcal{F} ; just let \mathcal{S}_m be the collection of all subsets of M^m for all $m \geq 0$. It follows that there exists a smallest definability system \mathcal{S} on M that contains \mathcal{F} . Namely, let \mathcal{S} be the greatest lower bound (intersection) of all definability systems on M that contain \mathcal{F} . We call this *the definability system on M generated by \mathcal{F}* .

22.11. Corollary. *Let L be a first order language and let \mathcal{N} be an L -structure with underlying set N ; let A be a subset of N . Then $\mathcal{D}(\mathcal{N}, A)$ is the definability system on M generated by the sets listed in (b)–(e) of Theorem 22.5.*

Proof. Exercise. □

EXERCISES

In the following Exercises, let \mathcal{S} be a definability system on the nonempty set M .

22.12. Show that there is a language L and an L -structure \mathcal{M} based on the set M such that $\mathcal{S} = \mathcal{D}(\mathcal{M}, \emptyset)$.

22.13. Let I be a finite index set and let $X \in \mathcal{S}_m$ be the union of the sets $X_i \in \mathcal{S}_m$, i ranging over I . Show that a function $f: X \rightarrow M^n$ belongs to \mathcal{S} if and only if all of its restrictions $f|X_i$ belong to \mathcal{S} .

22.14. Let $X \subseteq M^{m+n}$ and $x \in M^m$, and put $X_x := \{y \in M^n \mid (x, y) \in X\}$. Show that if X belongs to \mathcal{S} and $k \in \mathbb{N}$, then the set $\{x \in M^m \mid \text{card}(X_x) \leq k\}$ also belongs to \mathcal{S} .

22.15. Let the sets X, Y, Z and the function $f: X \times Y \rightarrow Z$ belong to \mathcal{S} . Show that the sets $\{a \in X \mid f(a, \cdot): Y \rightarrow Z \text{ is injective}\}$ and $\{a \in X \mid f(a, \cdot): Y \rightarrow Z \text{ is surjective}\}$ also belong to \mathcal{S} .

22.16. Suppose $M = \mathbb{R}$ and the order relation $\{(x, y) \in \mathbb{R}^2 \mid x < y\}$ belongs to \mathcal{S} . Suppose $X \subseteq \mathbb{R}^m$ belongs to \mathcal{S} . Show that the topological closure $\text{cl}(X)$ of X and the interior $\text{int}(X)$ of X in \mathbb{R}^m also belong to \mathcal{S} .

22.17. Suppose $M = \mathbb{R}$ and the order relation $\{(x, y) \in \mathbb{R}^2 \mid x < y\}$ belongs to \mathcal{S} . Suppose that the function $f: \mathbb{R}^{m+1} \rightarrow \mathbb{R}$ belongs to \mathcal{S} . Show that the set

$$X := \{a \in \mathbb{R}^m \mid f(a, t) \text{ tends to a limit } \ell(a) \text{ as } t \rightarrow +\infty\}$$

belongs to \mathcal{S} , and the limit function $\ell: A \rightarrow \mathbb{R}$ so defined also belongs to \mathcal{S} .

23. SELECTED PROBLEMS

23.1. Problem. Define a *simple* atomic formula to be one of the following four kinds: $x_1 = x_2$ or $R(x_1, \dots, x_m)$ or $y = c$ or $y = F(x_1, \dots, x_m)$, where x_1, \dots, x_m, y are distinct variables. (Here R is any predicate symbol, c is any constant symbol, and F is any function symbol.) Show that every L -formula is logically equivalent to an L -formula whose atomic formulas are all simple.

23.2. Problem. Let L be a first order language and let T be a theory in L that can be axiomatized by a set of universal sentences. Let $\varphi(u, v)$ be a quantifier-free L -formula such that $T \models \forall u \exists v \varphi(u, v)$. Show that there exists a finite sequence of terms $t_1(u), \dots, t_k(u)$ in L for which

$$T \models \forall u (\varphi(u, t_1(u)) \vee \dots \vee \varphi(u, t_k(u))).$$

23.3. Problem. Let L be an arbitrary first order language, and let \mathcal{M} and \mathcal{N} be finite L -structures. Show that $\mathcal{M} \equiv \mathcal{N}$ if and only if \mathcal{M} and \mathcal{N} are isomorphic.

23.4. Problem. Let I be a countably infinite set, and let U be any non-principal ultrafilter on I ; for each $i \in I$ let A_i be a countable infinite set. Show that the ultraproduct $\prod_U A_i$ has the same cardinality as the set of real numbers (which is the same as the cardinality of the set of all subsets of \mathbb{N}). Note that it is rather easy to show that $\prod_U A_i$ is uncountable, and this problem asks you to do more than that.

23.5. Problem. Let I be a countable (infinite) set and let U be a non-principal ultrafilter on I . Show that there is no countable collection S of subsets of I such that U is the filter generated by S .

23.6. Problem. Let K be any ultraproduct of fields. Show that there is a ring R that is a direct product of fields and a maximal ideal M in R such that K is isomorphic to the quotient of R by M . Conversely, show that each such quotient is isomorphic to an ultraproduct of fields.

23.7. Problem. Let L be the first order language with a binary relation symbol $<$ and a binary function symbol $+$. Let \mathcal{K} be the collection of infinite L -structures that are linearly ordered abelian groups. Show that any two members of \mathcal{K} satisfy exactly the same existential sentences of L . (Consider the two linearly ordered abelian groups \mathbb{Z} and \mathbb{Q} . Show that \mathbb{Z} embeds in every member of \mathcal{K} . Show that every member of \mathcal{K} embeds into an elementary extension of \mathbb{Q} . Show that \mathbb{Q} embeds into an elementary extension of \mathbb{Z} . Show that the desired result follows from these three statements.)

23.8. Problem. Show that the theory of the ordered field $(\mathbb{Q}, +, -, \times, 0, 1, <)$ does not admit quantifier elimination. (Hint: analyze the subsets of \mathbb{Q} that are definable by quantifier free formulas in this structure; show that if X is such a set, then there is a rational number

q such that the interval (q, ∞) is either contained in X or is disjoint from X . Find a definable subset of \mathbb{Q} that does not have this property.)

23.9. Problem. Let K be a finite field and let V be a vector space over K . A *symplectic form* on V is a bilinear function $(\cdot, \cdot): V \times V \rightarrow K$ such that $(x, x) = 0$ for all $x \in V$; the form is *nondegenerate* if for all nonzero $x \in V$ there exists $y \in V$ such that $(x, y) \neq 0$. We consider V with a symplectic form as a structure by formalizing the vector space structure as in Exercise 4.20 and by adding binary predicates P_a for each $a \in K$ and taking $P_a(x, y)$ to mean $(x, y) = a$. Let \mathcal{K} be the class of all infinite K -vector spaces with nondegenerate symplectic forms, and $T = \text{Th}(\mathcal{K})$. Show T admits quantifier elimination and use this to show T is complete.

23.10. Problem. Let F be an algebraically closed field and let K be a subfield of F . Given any $a \in F$, the 1-type of a over K (denoted $tp(a/K)$) is the 1-type realized by a in the expanded structure $F_K = (F, +, -, \times, 0, 1, \{b \mid b \in K\})$. (a) Show that a_1 and a_2 have the same 1-type over K if and only if they have the same minimal polynomial over K . (This includes the case where both a_1 and a_2 are transcendental over K .) (b) Show that the 1-type of a over K is principal (relative to the theory $\text{Th}(F_K)$) if and only if a is algebraic over K .

23.11. Problem. Let \mathcal{M} be an L -structure and X a subset of M that is M -definable in \mathcal{M} . Suppose that X is minimal, in the sense that every set $Y \subseteq X$ that is M -definable in \mathcal{M} is finite or cofinite. For each such Y define $\text{cl}(Y) = \text{acl}_{\mathcal{M}}(Y) \cap X$. Show that (X, cl) is a pregeometry.

23.12. Problem. Let p be a prime number or 0 and consider algebraically closed fields of characteristic p . Use properties of model theoretic algebraic closure on algebraically closed fields to prove the following facts:

- (1) Let k be a field of characteristic p , and let K_1 and K_2 be algebraically closed extension fields of k such that K_i is algebraic over k (in the sense of fields) for both $i = 1, 2$. Show that there is an isomorphism f from K_1 onto K_2 such that f is the identity map on k . (Uniqueness of algebraic closure.)
- (2) Let K be an algebraically closed field of characteristic p . Call $X \subset K$ **algebraically independent** if whenever $q(x_1, \dots, x_m)$ is a polynomial with coefficients in \mathbb{Z} and a_1, \dots, a_m are distinct elements of X such that $q(a_1, \dots, a_m) = 0$ in K , then q must be the 0 polynomial in K . Call X a **transcendence base** for K if X is algebraically independent in K and K is algebraic over the subfield generated by X . (a) Show that there exists a transcendence base for K ; (b) Show that X is a transcendence base for K iff X is a minimal subset of K with the property that K is algebraic over the field generated by X ; (c) Show that K is determined up to isomorphism by the cardinality of a transcendence base for K .

23.13. Problem. Show there exist vector spaces \mathcal{M}, \mathcal{N} over \mathbb{Q} such that \mathcal{M} and \mathcal{N} are elementarily equivalent but there does not exist any back and forth system for them.

23.14. **Problem.** Let \mathcal{M} be a κ -saturated L -structure, where $\kappa > \text{card}(L)$. Let $A \subseteq M$, with $\text{card}(A) < \kappa$. Show that $\text{acl}_{\mathcal{M}}(A)$ is equal to the intersection of all elementary substructures of \mathcal{M} that contain A .

23.15. **Problem.** Consider the structure $\mathcal{M} = (\mathcal{F}, \subseteq)$, where \mathcal{F} is the collection of all finite subsets of \mathbb{N} . Let τ be an infinite cardinal and τ^+ be the next cardinal larger than τ . Show that every τ^+ -saturated model that is elementarily equivalent to \mathcal{M} has cardinality at least 2^τ .

23.16. **Problem.** Let T be a complete strongly minimal theory in the first order language L . Suppose T has an infinite model. Show that for each cardinal number $\kappa \geq \text{card}(L)$ there is a κ -saturated model of T that has cardinality equal to κ . (Think about the dimension of such a model \mathcal{M} with respect to the pregeometry $\text{acl}_{\mathcal{M}}$ on M .)

23.17. **Problem.** Consider the language in which there are countably many unary relation symbols $\{P_n \mid n \in \mathbf{N}\}$ and a doubly indexed countable set of constant symbols $\{c_{ij} \mid i, j \in \mathbf{N}\}$. Let T be the theory in this language that asserts that the interpretations of the symbols $\{P_n \mid n \in \mathbf{N}\}$ are mutually disjoint sets, that the elements interpreting the various constants $\{c_{ij} \mid i, j \in \mathbf{N}\}$ are pairwise distinct, together with all sentences of the form $P_i(c_{ij})$. Show that T is a complete theory. Show that T has 2^ω mutually nonisomorphic countable models. Show that T has (nonetheless) a countable ω -saturated model and describe it.

23.18. **Problem.** Let F be an ordered field. Suppose that for any formula $\varphi(x)$ in the language of ordered rings for which $F \models \exists x \varphi(x)$, there is a finite element a of F such that $F \models \varphi[a]$. Show that there is an Archimedean ordered field K that is elementarily equivalent to F . (An element of an ordered field is *finite* if it is bounded above in absolute value by some element of the prime field. An ordered field is *Archimedean* if all of its elements are finite. Note that an ordered field is Archimedean if and only if it is isomorphic to an ordered subfield of \mathbb{R} .)

23.19. **Problem.** Let T be a theory in a countable language L and let $\Sigma(x)$ and $\Gamma(y)$ be partial 1-types that are consistent with T . Suppose that for every L -formula $\varphi(x, y)$ there is $\sigma(x) \in \Sigma(x)$ with the property that for all finite sets $\{\gamma_1(y), \dots, \gamma_n(y)\}$ of formulas from $\Gamma(y)$: if $\{\varphi, \gamma_1, \dots, \gamma_n\}$ is consistent with T then $\{\varphi, \gamma_1, \dots, \gamma_n, \neg \sigma\}$ is consistent with T . Show that T has a model realizing $\Gamma(y)$ and omitting $\Sigma(x)$.

23.20. **Problem.** Let T be a complete theory in a countable language. Assume that T has infinite models and is not ω -categorical. Show that T has at least two non-isomorphic countable models that are strongly ω -homogeneous.

23.21. **Problem.** Let L be the language with a unary function symbol as its only nonlogical symbol. Suppose that (A, f) is an L -structure whose theory is ω -categorical. Show that (A, f) is *uniformly locally finite*. That is, show that there exists a function $\alpha: \mathbb{N} \rightarrow \mathbb{N}$ such that for any nonempty

finite subset F of A , the substructure of (A, f) that is generated by F has at most $\alpha(\text{card}(F))$ elements.

23.22. Problem. Consider the language L_r of rings with 1. Let R be any ring with 1 and let $M_2(R)$ be the ring of 2×2 matrices over R with the identity matrix as its 1. Considering both rings as L_r -structures, let $T = \text{Th}(R)$ and $T_2 = \text{Th}(M_2(R))$. In each part below, is the given implication true for all rings R ?

- (a) If T is ω -categorical, then T_2 is ω -categorical.
- (b) If T is uncountably categorical, then T_2 is uncountably categorical.
- (c) If T is ω -stable, then T_2 is ω -stable.

23.23. Problem. Let T be a complete theory in a countable language, and assume T has infinite models. Show that there exists a countable model \mathcal{M} of T such that \mathcal{M} is isomorphic to a proper elementary substructure of itself. (Hint: take \mathcal{M} to be a Skolem hull generated by a suitable sequence of indiscernibles.)

23.24. Problem. Let L be countable and let \mathcal{M} be any countable L -structure. Show that there exists a sequence $(\mathcal{N}_n \mid n \in \mathbb{N})$ of L -structures such that (i) each \mathcal{N}_n is a proper elementary extension of \mathcal{M} , (ii) $\mathcal{N}_{n+1} \preceq \mathcal{N}_n$ for each $n \in \mathbb{N}$, and $M = \bigcap \{N_n \mid n \in \mathbb{N}\}$. (Hint: take \mathcal{N}_1 to be a Skolem hull generated by M and a suitable sequence of indiscernibles.)

23.25. Problem. Let T_0 and T be theories in L , with $T_0 \subseteq T$. We say that T is *preserved under substructures relative to T_0* if for any models \mathcal{M}, \mathcal{N} of T_0 , if $\mathcal{N} \models T$ and $\mathcal{M} \subseteq \mathcal{N}$, then $\mathcal{M} \models T$. Show that T is preserved under substructures relative to T_0 if and only if there is a set of universal sentences Σ such that T is axiomatized by $T_0 \cup \Sigma$.

23.26. Problem. Let \mathcal{M}, \mathcal{N} be L -structures and let T be a theory in L . Show that there exists a model \mathcal{C} of T such that \mathcal{M} and \mathcal{N} can be embedded in \mathcal{C} if and only if for all universal sentences σ and τ of L , if $T \models \sigma \vee \tau$ then both \mathcal{M}, \mathcal{N} satisfy σ or both \mathcal{M}, \mathcal{N} satisfy τ . (Note: any finite disjunction of universal sentences is logically equivalent to a universal sentence.)

23.27. Problem. Let L be any first order language, let κ be a cardinal number at least as large as the cardinality of L , and let T be a theory in L that has a set of $\forall\exists$ axioms. Show that if T is κ -categorical, then T is model complete.

23.28. Problem. Let $L \subseteq L'$ be first order languages. (a) Let T' be a theory in L' and suppose that each L -structure has at most one expansion that is a model of T' . Show that there is a theory T in L such that $\text{Mod}(T)$ is exactly the class of L -structures that have an expansion that is a model of T' . (b) Give an example of L, L', T' such that the class of reducts to L of models of T' is not $\text{Mod}(T)$ for any theory T in L .

24. SOME PROBLEM SETTINGS

These problems are stated in a very general way, analogous to how model theoretic questions actually arise in applications. They are meant to give realistic practice in the use of model theoretic methods and concepts in real mathematical settings. They can be answered using the concepts and methods presented in this book.

We begin with a series of model theoretic questions that can be applied to any first order theory T ; here L denotes the language in which T is expressed. Inevitably some of these questions will be very difficult to answer in some settings, while others will be easier and can be answered in several different ways. In each setting, the objective should be to obtain a system of answers that gives the clearest possible picture of the model theoretic properties of the class of structures being considered.

- Suppose T is introduced as $T = \text{Mod}(\mathcal{K})$ for some interesting class \mathcal{K} of L -structures. Check whether or not \mathcal{K} is indeed equal to $\text{Mod}(T)$; in other words, is \mathcal{K} an axiomatizable class. If not, characterize the models of T .
- Describe some familiar structures that are models of T .
- Characterize the models of T up to isomorphism, as much as possible. In which cardinals is T categorical? How many countable models does T have, up to isomorphism?
- Does T admit quantifier elimination? If not, can you identify a convenient extension by definitions that does admit QE?
- Is T complete? If not, describe the completions of T as simply and directly as possible. That is, characterize the models of T up to elementary equivalence.
- Is T model complete? If not, what are the existentially closed models of T ? Is the class of existentially closed models of T axiomatizable?
- Is T strongly minimal? More generally, characterize the families of definable sets in models of T . (Here we mean sets of *elements* of the underlying set of the model, not tuples, and parameters are allowed in the definitions.)
- Does T have a countable model that is ω -saturated and strongly ω -homogeneous. If so, can you describe it in an explicit and concrete way using familiar mathematical objects.
- Does T have a countable atomic model? If so, can you describe it in an explicit and concrete way using familiar mathematical objects. For each n -tuple a from this model, try to give explicitly a complete formula satisfied by a .
- Characterize the definable closure of a set A in a model of T . That is, given a model \mathcal{M} of T and a set $A \subseteq M$, for which elements $a \in M$ is the set $\{a\}$ definable over A in \mathcal{M} ?
- Characterize the algebraic closure of a set A in a model of T .

- Does T have any saturated models? (A model is saturated if it is κ -saturated, where κ is the cardinality of its underlying set.)
- Let \mathcal{M} be any model of T . What totally defined functions $f: M \rightarrow M$ are definable in \mathcal{M} ? (Here we allow parameters in the definitions.)
- Is T ω -stable? If so, what is the Morley rank and Morley degree of the formula $x = x$?

Next we turn to a list of basic theories to which the questions given above should be addressed. All of the structures considered are mathematically very simple, so treating them does not require much background knowledge about them. The first eight settings are very combinatorial in nature; they involve bijections, equivalence relations, families of sets, and graphs. The last five settings involve linear orderings, vector spaces, and Boolean algebras, all of which are basic objects from undergraduate mathematics.

If T is an L -theory and L_0 is a sublanguage of L , then the *restriction of T to L_0* is the L_0 -theory T_0 consisting of all L_0 -sentences σ such that $T \models \sigma$.

A. *Bijections(1)*:

- L is the language with a unary function symbol F ;
- T is the L -theory of the class of all L -structures (M, f) where f is a bijection from M onto itself and f has no finite cycles.

B. *Bijections(2)*:

- L is the language with a unary function symbol F and a constant symbol 0 ;
- T is the L -theory of the class of all L -structures (M, f, a) where f is a bijection from M onto $M \setminus \{a\}$ and f has no finite cycles.

C. *Equivalence relations(1)*:

- L is the language with a single, binary predicate symbol E ;
- T is the L -theory of equivalence relations that have infinitely many classes, and each class is infinite.

D. *Equivalence relations(2)*:

- L is the language with a single, binary predicate symbol E ;
- T is the L -theory of equivalence relations that have exactly one class of size n , for each integer $n \geq 1$.

Also consider the language L' obtained from L by adding unary predicates P_n for each $n \geq 1$ and the extension T' of T obtained by adding for each $n \geq 1$ the sentence expressing the condition “ $P_n(x)$ if and only if the equivalence class of x has cardinality n ”.

E. *Descending chain of equivalence relations*:

- L is the language with infinitely many binary predicate symbols E_n for $n \in \mathbb{N}$;
- T is the theory whose axioms assert that each E_n is an equivalence relation, that E_0 has infinitely many equivalence classes, and that each

equivalence class of each E_n is the union of infinitely many equivalence classes of E_{n+1} .

Also consider the language L_n whose nonlogical symbols are E_0, \dots, E_n and the L_n -theory T_n that is the restriction of T to L_n .

F. Descending chain of subsets:

- L is the language whose nonlogical symbols are countably many unary predicates $(P_n \mid n \in \mathbb{N})$;
- T is the L -theory axiomatized by the sentences asserting that $P_0 \supseteq P_1 \supseteq P_2 \supseteq \dots$, and that the complement of P_0 and each $P_n \setminus P_{n+1}$ are infinite.

Also consider the language L_n whose nonlogical symbols are P_0, \dots, P_n and the L_n -theory T_n that is the restriction of T to L_n .

G. Independent subsets:

- L is the language whose nonlogical symbols are countably many unary predicates $(P_n \mid n \in \mathbb{N})$;
- T is the L -theory axiomatized by the sentences

$$\exists x \left(\bigwedge_{j \in F} U_j(x) \wedge \bigwedge_{j \in G} \neg U_j(x) \right)$$

where F, G are arbitrary finite subsets of \mathbb{N} .

Also consider the language L_n whose nonlogical symbols are P_0, \dots, P_n and the L_n -theory T_n that is the restriction of T to L_n .

H. Random graph:

- L is the language with one binary predicate symbol E ;
- T is the L -theory of all graphs having the following extension properties (k ranges over integers ≥ 1): if A and B are any two disjoint sets of k vertices, then there exists a vertex v such that for each $a \in A$ and each $b \in B$ there is an edge between v and a but no edge between v and b .

A graph is a structure $G = (V, E)$ where V is a nonempty set and E is a symmetric, irreflexive binary relation on V ; $E(x, y)$ means that there is an edge between x and y .

I. Dense linear orderings with an increasing sequence:

- L is the language with a binary relation symbol $<$ and infinitely many constant symbols $\{c_n \mid n \in \mathbb{N}\}$;
- T is the L -theory whose axioms are the axiom for dense linear orderings without endpoints and the sentences $c_i < c_j$ for each $i < j$ in \mathbb{N} .

J. Discrete linear orderings:

- L is the language with a binary relation symbol $<$;
- T is the L -theory whose models are exactly the linear orderings in which each element has an immediate successor and an immediate predecessor.

Also consider the language L' obtained from L by adding unary function symbols p and s and the extension T' of T obtained by adding axioms ensuring that $s(x)$ is the immediate successor of x and $p(x)$ is the immediate predecessor of x , for any x in the underlying set.

K. Divisible abelian ordered groups:

- L is the language with binary function symbols $+$ and $-$, a binary predicate symbol $<$, and a constant symbol 0 ;
- T is the L -theory of nontrivial divisible ordered abelian groups.

L. Atomless Boolean algebras:

- L is the language with two binary function symbols \cap and \cup , a unary function symbol $(\cdot)^c$, and two constant symbols 0 and 1 ;
- T is the L -theory of atomless Boolean algebras. (We take the functions symbols to be interpreted as intersection, union, and complement, and the constants to be interpreted as the bottom and top elements of the Boolean algebra.)

An *atom* in a Boolean algebra is a nonzero element that is minimal among nonzero elements, in the natural partial ordering on the algebra. A Boolean algebra is *atomless* if it contains no atoms; note that this is equivalent to the condition that for any nonzero element a there exist nonzero elements b and c that are disjoint and whose union is a .

M. Atomic Boolean algebras:

- L is the language with two binary function symbols \cap and \cup , a unary function symbol, and two constant symbols 0 and 1 ; they are interpreted as in the previous item.
- T is the L -theory of atomic Boolean algebras.

Also consider the language L' obtained from L by adding unary predicate symbols A_n for each $n \geq 1$; take T' to be the L' -theory obtained from T by adding axioms stating that each $A_n(x)$ means “ x contains at least n mutually disjoint atoms”.

A Boolean algebra is *atomic* if each nonzero element contains at least one atom; note that this is equivalent to the condition that any nonzero element is the supremum of a collection of atoms.