Class Notes for
Mathematics 571
Spring 2010

Model Theory

written by
C. Ward Henson

Mathematics Department
University of Illinois
1409 West Green Street
Urbana, Illinois 61801
email: henson@math.uiuc.edu
www: http://www.math.uiuc.edu/~henson/

Introduction

The purpose of Math 571 is to give a thorough 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 is a rich and interesting subject with significant applications to the main body of mathematics. Model theory began as a serious subject in the 1950s with the work of Abraham Robinson and Alfred Tarski, and since then it has been an active and successful area of research.

Beyond the core techniques and results of model theory, Math 571 places a lot of emphasis on examples and applications, in order to show clearly the variety of ways in which model theory can be useful in mathematics. For example, we give a thorough treatment of the model theory of the field of real numbers (real closed fields) and show how this can be used to obtain the characterization of positive semi-definite rational functions that gives a solution to Hilbert's 17th Problem.

A highlight of Math 571 is a proof of 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.* The machinery needed for this proof includes the concepts of *Morley rank and degree* for formulas in ω-stable theories. The methods needed for this proof illustrate ideas that have become central to modern research in model theory.

To succeed in Math 571, it is necessary to have 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 course at UIUC is Math 570, but this covers many things that are not needed as background for Math 571.

In the lecture notes for Math 570 (written by Prof. van den Dries) the material necessary for Math 571 is presented in sections 2.3 through 2.6 (pages 24–37 in the 2009 version). These lecture notes are available at http://www.math.uiuc.edu/ vddries/410notes/main.dvi.

A standard undergraduate text in logic is *A Mathematical Introduction to Logic* by Herbert B. Enderton (Academic Press; second edition, 2001). Here the material needed for Math 571 is covered in sections 2.0 through 2.2 (pages 67–104).

This material is also discussed in *Model Theory* by David Marker (see sections 1.1 and 1.2, and the first half of 1.3, as well as many of the exercises at the end of chapter 1) and in many other textbooks in model theory.

For Math 571 it is not necessary to have any exposure to a proof system for first order logic, nor to Gödel's completeness theorem. Math 571 begins with a proof of the compactness theorem for first order languages, and this is all one needs for model theory.

We close this introduction by discussing a number of books of possible interest to anyone studying model theory.

The first two books listed are now the standard graduate texts in model theory; they can be used as background references for most of what is done in Math 571.

David Marker, *Model Theory: an Introduction.*

Bruno Poizat, *A Course in Model Theory.*

The next book listed was the standard graduate text in model theory from its first publication in the 1960s until recently. It is somewhat out of date and incomplete from a modern viewpoint, but for much of the content of Math 571 it is a suitable reference.

C. C. Chang and H. J. Keisler, *Model Theory.*

Another recent monograph on model theory is *Model Theory* by Wilfrid Hodges. This book contains many results and examples that are otherwise only available in journal articles, and gives a very comprehensive treatment of basic model theory. However it is very long and it is organized in a complicated way that makes things hard to find. The author extracted a shorter and more straightforward text entitled *A Shorter Model Theory*, which is published in an inexpensive paperback edition.

In the early days of the subject (i.e., 1950s and 1960s), Abraham Robinson was the person who did the most to make model theory a useful tool in the main body of mathematics. Along with Alfred Tarski, he created much of modern model theory and gave it its current style and emphasis. He published three books in model theory, and they are still interesting to read:
(a) *Intro. to Model Theory and the Metamathematics of Algebra*, 1963;
(b) *Complete Theories*, 1956; new edition 1976;
(c) *On the Metamathematics of Algebra*, 1951.

The final reference listed here is *Handbook of Mathematical Logic*, Jon Barwise, editor; this contains expository articles on most parts of logic. Of particular interest to students in model theory are the following chapters:
A.1. An introduction to first-order logic, Jon Barwise.
A.2. Fundamentals of model theory, H. Jerome Keisler.
A.3. Ultraproducts for algebraists, Paul C. Eklof.
A.4. Model completeness, Angus Macintyre.

Contents

1. Ultraproducts and the Compactness Theorem

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 ultrafilter[1] on I. Consider a family of L-structures $(\mathcal{A}_i \mid i \in I)$. For each $i \in I$ let A_i denote the underlying set of the structure $\mathcal{A}_i$ and take $A = \prod(A_i \mid i \in I)$ to be the cartesian product of the sets A_i.

We define an interpretation[2] $\mathcal{A}$ of L as follows:

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

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

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

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

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

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

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

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

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 $\mathcal{A}$ we have the following easy fact, proved by a straightforward argument using induction on terms:

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

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

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

1.2. **Proposition.** *For any L-formula $\varphi(x_1, \ldots, x_n)$ and any $f_1, \ldots, f_n \in A$*

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

[1]See Appendix 1 of this chapter for some basic facts about filters and ultrafilters.

[2]See Appendix 2 of this chapter for an explanation of the words "interpretation", "prestructure", and "structure" and for some basic relations among them.

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. In the basis step of the induction 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 A^n; let $g_j(i) = t_j^{\mathcal{A}_i}(f_1(i), \dots, f_n(i))$ for each $j = 1, \dots, m$ and $i \in I$. Note that $g_j \in A$ for each $j = 1, \dots, m$. Then we have:

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

(Lemma 1.1 is used in the second equivalence.)

In the induction step of the proof 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 $\neg\varphi_1$:

$$\begin{array}{lll} \mathcal{A} \models \varphi[f_1, ...f_n] & \Leftrightarrow & \mathcal{A} \not\models \varphi_1[f_1, ...f_n] \\ & \Leftrightarrow & \{i \mid \mathcal{A}_i \models \varphi_1[f_1(i), ..., f_n(i)]\} \notin U \\ & \Leftrightarrow^\star & \{i \mid \mathcal{A}_i \not\models \varphi_1[f_1(i), ..., f_n(i)]\} \in U \\ & \Leftrightarrow & \{i \mid \mathcal{A}_i \models \neg\varphi_1[f_1(i), ..., f_n(i)]\} \in U \\ & \Leftrightarrow & \{i \mid \mathcal{A}_i \models \varphi[f_1(i), ..., f_n(i)]\} \in U \end{array}$$

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 $(\varphi_1 \wedge \varphi_2)$:

$$\begin{array}{llrl} \mathcal{A} \models \varphi[f_1, ...f_n] & \Leftrightarrow & \mathcal{A} & \models (\varphi_1 \wedge \varphi_2)[f_1, ...f_n] \\ & \Leftrightarrow & \mathcal{A} & \models \varphi_1[f_1(i), ..., f_n(i)] \text{ and} \\ & & \mathcal{A} & \models \varphi_2[f_1(i), ..., f_n(i)] \\ & \Leftrightarrow & \{i \mid \mathcal{A}_i & \models \varphi_1[f_1(i), ..., f_n(i)]\} \in U \text{ and} \\ & & \{i \mid \mathcal{A}_i & \models \varphi_2[f_1(i), ..., f_n(i)]\} \in U \\ & \Leftrightarrow^\star & \{i \mid \mathcal{A}_i & \models \varphi_1[f_1(i), ..., f_n(i)]\} \cap \\ & & \{i \mid \mathcal{A}_i & \models \varphi_2[f_1(i), ..., f_n(i)]\} \in U \\ & \Leftrightarrow & \{i \mid \mathcal{A}_i & \models \varphi_1[f_1(i), ..., f_n(i)] \text{ and} \\ & & \mathcal{A}_i & \models \varphi_2[f_1(i), ..., f_n(i)]\} \in U \\ & \Leftrightarrow & \{i \mid \mathcal{A}_i & \models (\varphi_1 \wedge \varphi_2)[f_1(i), ..., f_n(i)]\} \in U \\ & \Leftrightarrow & \{i \mid \mathcal{A}_i & \models \varphi[f_1(i), ..., f_n(i)]\} \in U \end{array}$$

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 $\exists y\varphi_1$: We may assume that y is not among the variables $x_1, ..., x_n$.

$$
\begin{array}{ll}
 & \mathcal{A} \models \varphi[f_1, ..., f_n] \\
\Leftrightarrow & \mathcal{A} \models \exists y\varphi_1[f_1, ..., f_n] \\
\Leftrightarrow & \text{for some } g \in A : \mathcal{A} \models \varphi_1[f_1, ..., f_n, g] \\
\Leftrightarrow & \text{for some } g \in A : \{i \mid \mathcal{A}_i \models \varphi_1[f_1(i), ..., f_n(i), g(i)]\} \in U \\
\Leftrightarrow^\dagger & \{i \mid \text{for some } a \in A_i : \mathcal{A}_i \models \varphi_1[f_1(i), ..., f_n(i), a]\} \in U \\
\Leftrightarrow & \{i \mid \mathcal{A} \models \exists y\varphi_1[f_1(i), ..., f_n(i)]\} \in U
\end{array}
$$

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

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

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

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

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

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

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 $\{\mathcal{A}_i \mid i \in I\}$ of L-structures, the ultraproduct $\prod_U \mathcal{A}_i$ is isomorphic to $\mathcal{A}_j$.

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 Łos.

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 $\varphi(x_1, \ldots, x_n)$ and any sequence $f/U = (f_1/U, \ldots, f_n/U)$,*

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

Proof. This is an immediate consequence of Propositions 1.2 and 1.29. □

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

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

Proof. This is a special case of Theorem 1.7. □

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 $\mathcal{A}$ be an L-structure. We say that $\mathcal{A}$ is a *model* of T, and write $\mathcal{A} \models T$, if every sentence in T is true in $\mathcal{A}$.

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 $\mathcal{A}_i$ be any model of i, which exists by assumption. We will obtain the desired model of T as an ultraproduct $\prod_U \mathcal{A}_i$ for a suitably chosen ultrafilter U on I.

Let S be the family of all the subsets of I of the form $I_\sigma = \{i \in I : \sigma \in i\}$, where $\sigma \in T$. Note that S has the finite intersection property; indeed, each finite intersection $I_{\sigma_1} \cap ... \cap I_{\sigma_n}$ has $\{\sigma_1, ..., \sigma_n\}$ as an element. So there exists an ultrafilter U on I that contains S, by Corollary 1.25.

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

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 $\mathcal{A}$, by the Compactness Theorem. Let $\mathcal{B}$ be the reduct of $\mathcal{A}$ to the original language L; $\mathcal{A}$ 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 $\mathcal{A}$ let $\varphi^{\mathcal{A}}$ denote the set of tuples a from A such that $\mathcal{A} \models \varphi(a)$. Suppose that the set $\varphi^{\mathcal{A}}$ is finite whenever $\mathcal{A}$ is a model of T. Then there is a positive integer N such that $\varphi^{\mathcal{A}}$ has at most N elements for every model $\mathcal{A}$ 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 $\mathcal{A}$ be an L-structure. We say that Γ is *satisfiable in* $\mathcal{A}$ if there exist elements $(a_j \mid j \in J)$ of A such that $\mathcal{A} \models \Gamma[a_j \mid j \in J]$.

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 $\mathcal{A}$ of F such that G is satisfiable in $\mathcal{A}$.

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 filters and ultrafilters:

1.18. **Definition.** (1) Let I be a nonempty set. A *filter on* I is a collection F of subsets of I that satisfies:

(a) $\emptyset \notin F$ and $I \in F$;

(b) for all $A, B \in F, A \cap B \in F$;

(c) for all $A \in F$ and $B \subseteq I$, if $A \subseteq B$ then $B \in F$.

(2) Let F be a filter on I; F is an *ultrafilter* if it is maximal under $\subseteq$ among filters on I. F is *principal* if there exists a subset A of I such that F is exactly the collection of all sets B that satisfy $A \subseteq B \subseteq I$.

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 $\{A_1, \ldots, A_n\}$ of S, the intersection $A_1 \cap \ldots \cap A_n$ is nonempty.

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

Proof. ($\Rightarrow$) It is immediate (by induction) from the definition of filter that any filter is closed under finite intersections. Since no filter contains the empty set, this shows that each filter has the FIP; hence the same is true of any subcollection of a filter.

($\Leftarrow$) Let $S \subseteq \mathcal{P}(I)$ have the FIP; we want to find a filter $F \supseteq S$. We know that all supersets of finite intersections of elements of S must be elements of F. Thus we are led to define

$$F = \{A \mid A \subseteq I \text{ and there exist } A_1, \ldots, A_n \in S \text{ with } A \supseteq A_1 \cap \ldots \cap A_n\}$$

Since S has the FIP, we see that F does not contain the empty set. It is easy to check that conditions (b) and (c) in the definition of filter are satisfied. $\square$

Remark: The filter defined in the preceding proof is evidently the smallest filter on I containing S. Thus it is called *the filter generated by S*.

1.21. **Facts.** Let I be a nonempty set.
(a) If S is a collection of subsets of I and S has the FIP, then for any $A \subseteq I$, either $S \cup \{A\}$ or $S \cup \{I \setminus A\}$ has the FIP.
(b) Suppose J is an index set and for each $j \in J$, S_j is a collection of subsets of I that has the FIP. Suppose that the family $\{S_j \mid j \in J\}$ is *directed*, in the sense that for any $j_1, j_2 \in J$ there is $j_3 \in J$ such that $S_{j_1} \cup S_{j_2} \subseteq S_{j_3}$. Let S be the union of the family $\{S_j \mid j \in J\}$. Then S has the FIP.

1.22. **Lemma.** *Let I be a nonempty set and let F be a filter on I. F is an ultrafilter if and only if for each $A \subseteq I$, either $A \in F$ or $I \setminus A \in F$.*

Proof. ($\Rightarrow$) If F is an ultrafilter on I and $A \subseteq I$, then by Fact 1.21(a), either $F \cup \{A\}$ or $F \cup \{I \setminus A\}$ has the FIP. Therefore by Lemma 1.20, $F \cup \{A\}$ or $F \cup \{I \setminus A\}$ is contained in a filter. But F is maximal, so the only filter it can be contained in must be F itself. Hence $F \cup \{A\} \subseteq F$ or $F \cup \{I \setminus A\} \subseteq F$.

($\Leftarrow$) Suppose F is a filter on I with the property that for any $A \subseteq I$ either $A \in F$ or $I \setminus A \in F$. We have to show that F is maximal among the filters on I under set-theoretic inclusion. If F is not maximal, then there is a filter G on I with $F \subseteq G$ and $G \neq F$. Take any set $A \in G \setminus F$. Since $A \notin F$ we must have $I \setminus A \in F \subseteq G$. But then A and $I \setminus A$ are in G, and this implies $\emptyset \in G$, which is a contradiction. $\square$

1.23. **Facts.** Let I be a nonempty set and let U be an ultrafilter on I.
(a) If $A_1, \ldots, A_n$ are subsets of I and if the set $A_1 \cup \cdots \cup A_n$ is in U, then for some $j = 1, \ldots, n$ the set A_j is in U.
(b) If $A_1, \ldots, A_n$ are subsets of I and if the set $A_1 \cap \cdots \cap A_n$ is in U, then for all $j = 1, \ldots, n$ the set A_j is in U.
(c) The ultrafilter U is principal iff some element of U is a finite set iff some element of U is a singleton set (a set of the form $\{i\}$ for some $i \in I$).

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

Zorn's Lemma: If $(\Lambda, \leq)$ is a nonempty partially ordered set with the property that every linearly ordered subset of $(\Lambda, \leq)$ has an upper bound in $(\Lambda, \leq)$, then $(\Lambda, \leq)$ has a maximal element.

Remark: An element of $(\Lambda, \leq)$ is maximal if no other element is strictly larger than it. There may be many maximal elements. We will often use Zorn's Lemma where Λ is a collection of sets and $\leq$ is the set containment relation $\subseteq$. In that situation the hypothesis of Zorn's Lemma states that whenever C is a subcollection of Λ and C is a chain under $\subseteq$, the union of C is a subset of some element of Λ. (This restricted formulation is easily seen to be equivalent to Zorn's Lemma; it is known as *Hausdorff's Maximum Principle.*)

1.24. **Theorem.** *Let I be a nonempty set. Every filter on I is contained in an ultrafilter on I.*

Proof. Let F be a filter on I. Let $\Lambda = \{G \mid G \text{ is a filter on } I \text{ and } F \subseteq G\}$. Partially order Λ by set inclusion $\subseteq$. We want to apply Zorn's Lemma to $(\Lambda, \leq)$. Suppose C is a chain in Λ. It is easy to show that the union of C is a filter, and hence it is in Λ. (Compare Fact 1.21(b).) Zorn's Lemma yields the existence of a maximal element G in Λ. That is, G is a filter that contains F and G is maximal among all filters on I that contain F. In particular, G is maximal as a filter on I; by definition G is an ultrafilter. $\square$

1.25. **Corollary.** *Let I be a nonempty set and let S be a collection of subsets of I. If S has the FIP, then there is an ultrafilter on I that contains S.*

Proof. Immediate from Lemma 1.20 and Theorem 1.24. $\square$

Appendix 1.B: From prestructures to structures

Let L be any first order language.

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

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

$$\mathcal{A} \models \varphi[a_1, \ldots, a_n]$$

where $\varphi(x_1, \ldots, x_n)$ is an L-formula and $a_1, \ldots, a_n \in A$. 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.

1.27. **Definition.** A *prestructure* $\mathcal{A}$ for L is an interpretation of L in which the logical equality axioms are valid; that is,
(i) $=^{\mathcal{A}}$ is an equivalence relation on A;
(ii) for any n, any n-ary function symbol F of L, and any elements $a_1, b_1, \ldots, a_n, b_n$ of A such that $a_1 =^{\mathcal{A}} b_1, \ldots, a_n =^{\mathcal{A}} b_n$ one has

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

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

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

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

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

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

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

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

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

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

for every $a_1, \dots, a_n \in A$. The fact that $=^{\mathcal{A}}$ is a congruence for $F^{\mathcal{A}}$ 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 $\mathcal{B}$ is the n-ary relation $P^{\mathcal{B}}$ on B defined by

$$P^{\mathcal{B}}([a_1], \dots, [a_n]) \quad \Longleftrightarrow \quad P^{\mathcal{A}}(a_1, \dots, a_n)$$

for every $a_1, \dots, a_n \in A$. The fact that $=^{\mathcal{A}}$ is a congruence for $P^{\mathcal{A}}$ 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 $\mathcal{B}$ is to be a structure, the interpretation $=^{\mathcal{B}}$ of $=$ in $\mathcal{B}$ must be the identity relation on B. Note that we have

$$[a] =^{\mathcal{B}} [b] \quad \Longleftrightarrow \quad a =^{\mathcal{A}} b$$

for all $a, b \in A$. Hence the identity interpretation of $=$ in $\mathcal{B}$ 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 $\mathcal{B}$ can be summarized by saying that the quotient map π from A onto B is a strong homomorphism of $\mathcal{A}$ onto $\mathcal{B}$.

The following Lemma is easily proved by induction on terms.

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

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

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 $\mathcal{A}$ and its quotient structure $\mathcal{B}$ 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.)

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

$$\mathcal{B} \models \varphi[[a_1],\ldots,[a_n]] \quad\Longleftrightarrow\quad \mathcal{A} \models \varphi[a_1,\ldots,a_n].$$

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. □

Exercises

1.30. 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 $(\mathcal{A}_i \mid i \in I)$ is a family of L-structures, then $\Pi_U(\mathcal{A}_i \mid i \in I)$ is isomorphic to $\Pi_V(\mathcal{A}_j \mid j \in J)$

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

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

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

• Show that $\mathcal{B}$ is a linear ordering.

• Show that the range of the diagonal embedding of $\mathcal{A}$ into $\mathcal{B}$ is a proper initial segment of $\mathcal{B}$. Give an explicit description of an element of B that is not in the range of this embedding.

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

1.34. 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. Theories and Types

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

A *theory* consists of a first order language L together with a set T of sentences in L; often the language is determined by the context. We may refer to T as an *L-theory* to indicate which language is intended.

2.1. **Definition.** Let L be a first order language, let T, T' be L-theories, let σ be an L-sentence, and let $\mathcal{K}$ be a class of L-structures.

• T is *satisfiable* if it has at least one model; $\text{Mod}(T)$ denotes the class of all models of T.

• σ is a *logical consequence* of T (and we write $T \models \sigma$) if σ is true in every model of T.

• T is *complete* if it is satisfiable and for every L-sentence σ, either $T \models \sigma$ or $T \models \neg\sigma$.

• T and T' are *equivalent* if they have the same logical consequences in L; this is the same as saying that each sentence in T is a logical consequence of T' and each sentence in T' is a logical consequence of T. When T and T' are equivalent we will also say that T is *axiomatized* by T'.

• The *theory* of $\mathcal{K}$ is defined by

$$\text{Th}(\mathcal{K}) = \{\sigma \mid \sigma \text{ is an } L\text{-sentence and } \mathcal{A} \models \sigma \text{ for all } \mathcal{A} \in \mathcal{K}\}.$$

If $\mathcal{K} = \{\mathcal{A}\}$ we write $\text{Th}(\mathcal{A})$ instead of $\text{Th}(\{\mathcal{A}\})$. We say $\mathcal{K}$ is *axiomatizable* if $\mathcal{K} = \text{Mod}(T)$ for some theory T.

The following results are easy consequences of the definitions.

2.2. **Facts.** Let $\mathcal{K}$, $\mathcal{K}_1$ and $\mathcal{K}_2$ be classes of L-structures, and let T, T_1, and T_2 be L-theories.

(1) $\mathcal{K}_1 \subseteq \mathcal{K}_2 \Rightarrow \text{Th}(\mathcal{K}_1) \supseteq \text{Th}(\mathcal{K}_2)$;
(2) $T_1 \subseteq T_2 \Rightarrow \text{Mod}(T_1) \supseteq \text{Mod}(T_2)$;
(3) T_1 and T_2 are equivalent iff $\text{Mod}(T_1) = \text{Mod}(T_2)$;
(4) $\text{Mod}(\text{Th}(\mathcal{K})) \supseteq \mathcal{K}$, with equality if $\mathcal{K}$ is axiomatizable;
(5) $\text{Th}(\mathcal{K})$ contains its logical consequences;
(6) $\text{Th}(\text{Mod}(T)) \supseteq T$ and T axiomatizes $\text{Th}(\text{Mod}(T))$.
(7) T is of the form $\text{Th}(\mathcal{A})$, where $\mathcal{A}$ is an L-structure, iff T is finitely satisfiable and it is $\subseteq$-maximal among finitely satisfiable sets of L-sentences.
(8) Among L-theories containing T, the complete theories are those equivalent to theories of the form $\text{Th}(\mathcal{A})$, where $\mathcal{A}$ is a model of T.

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

2.4. **Fact.** Let T be a theory. Then T is complete iff T has a model and any two models of T are elementarily equivalent.

2.5. **Definition.** Let $L_1 \subseteq L_2$ be two first order languages and let T_i be a theory in L_i for $i = 1, 2$. We say that T_2 is an *extension* of T_1 (and, equivalently, that T_1 is a *subtheory* of T_2), if T_1 is contained in the set of logical consequences of T_2. Further, T_2 is said to be a *conservative* extension of T_1, if, in addition, $T_2 \models \sigma \Rightarrow T_1 \models \sigma$ for every sentence σ of L_1.

2.6. **Fact.** If T_2 is an extension of T_1 and every model of T_1 has an expansion to a model of T_2, then T_2 is a conservative extension of T_1.

The following result, which is just a restatement of the Compactness Theorem, expresses a fundamental property of logical consequence in first order logic. It shows that the relation $T \models \sigma$ has finitary character, and thus (in principle) can be analyzed by some sort of (possibly abstract) "proof system" with the property that only finitely many sentences appear in each "proof." Each presentation of Gödel's Completeness Theorem (which we do not need for model theory) gives such a proof system.

2.7. **Corollary.** *If $T \models \sigma$, then there is a finite set $T_0 \subseteq T$ with $T_0 \models \sigma$.*

Proof. Assume $T \models \sigma$, so $T \cup \{\neg\sigma\}$ has no model. Hence there exists a finite $T' \subseteq T \cup \{\neg\sigma\}$ such that T' has no model, by the Compactness Theorem. There exists a finite $T_0 \subseteq T$ with $T' \subseteq T_0 \cup \{\neg\sigma\}$. Evidently $T_0 \cup \{\neg\sigma\}$ cannot have a model, and therefore $T_0 \models \sigma$. □

The following result is a variation on the same theme:

2.8. **Corollary.** *Let L be a first order language and let S and T be sets of sentences in L. Suppose that for every model $\mathcal{A}$ of T there exists $\gamma \in S$ such that $\mathcal{A} \models \gamma$. Then there exists a finite subset $\{\gamma_1, \ldots, \gamma_m\}$ of S such that $T \models \gamma_1 \vee \ldots \vee \gamma_m$.*

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

2.9. **Definition.** Let T be a satisfiable L-theory. We denote by $S_0(T)$ the set of all theories of the form $\mathrm{Th}(\mathcal{A})$, where $\mathcal{A}$ is a model of T.

Note that we may regard $S_0(T)$ as the set of complete L-theories that extend T, up to equivalence of theories. We think of it as the *space of completions of T*.

We put a natural topology on $S_0(T)$ as follows: for each L-sentence σ, let

$$[\sigma] = \{T' \in S_0(T) \mid \sigma \in T'\} = \{T' \in S_0(T) \mid T' \models \sigma\}.$$

Note that the family

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

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 $S_0(T)$ is the topology for which $\mathcal{F}$ is the family of basic open sets. That is, for each $T' \in S_0(T)$, 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, since

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

for all L-sentences σ.

Furthermore, the logic topology on $S_0(T)$ is *compact*; this is an immediate consequence of Corollary 2.8.

Note also that there is a close relation between closed sets in $S_0(T)$ and L-theories T_1 that extend T. For such a theory T_1, define

$$K(T_1) = \{T' \in S_0(T) \mid T_1 \subseteq T'\} = \bigcap\{[\sigma] \mid \sigma \in T_1\}.$$

Then $K(T_1)$ is closed, because it is the intersection of a family of clopen sets. Conversely, if K is a closed set in $S_0(T)$, then there is a set Σ of L-sentences such that the open set $S_0(T) \setminus K$ is equal to the union of the basic open sets $[\sigma]$ with $\sigma \in \Sigma$. Taking $T_1 = T \cup \{\neg\sigma \mid \sigma \in \Sigma\}$ we have that T_1 extends T and $K(T_1) = K$. Note that $K(T_1)$ is nonempty iff T_1 is satisfiable.

2.10. **Proposition.** *Let T be a satisfiable L-theory. The space $S_0(T)$ with the logic topology is a totally disconnected, compact Hausdorff space. Its closed sets are the sets of the form*

$$\{T' \in S_0(T) \mid T_1 \subseteq T'\}$$

where T_1 is a set of L-sentences containing T. Moreover, the clopen subsets of $S_0(T)$ 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 $C \subseteq S_0(T)$ 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 $S_0(T)$, a union of finitely many basic open sets is itself a basic open set. □

2.11. **Fact.** Let T be a satisfiable L-theory and let σ, τ be L-sentences. Then $[\sigma] = [\tau]$ iff σ and τ are equivalent over T (*i.e.*, $\sigma \leftrightarrow \tau$ is a logical consequence of T).

These results show that the topological space $S_0(T)$ by itself characterizes the relation of equivalence of L-sentences over the theory T.

Types

Next we introduce *types*; they provide a way of describing the first order expressible properties of elements of a structure.

Fix $n \geq 1$ and let $x_1, \ldots, x_n$ be a fixed sequence of distinct variables.

2.12. **Definition.** Let $\mathcal{A}$ be an L-structure and consider $a_1, \ldots, a_n \in A$.

• The *type of* $(a_1, \ldots, a_n)$ *in* $\mathcal{A}$ is the set of L-formulas

$$\{\varphi(x_1, \ldots, x_n) \mid \mathcal{A} \models \varphi[a_1, \ldots, a_n]\};$$

we denote this set by $\mathrm{tp}_{\mathcal{A}}(a_1, \ldots, a_n)$ or simply by $\mathrm{tp}(a_1, \ldots, a_n)$ if the structure $\mathcal{A}$ is understood.

• An *n-type in L* is a set of formulas of the form $\mathrm{tp}_{\mathcal{A}}(a_1, \ldots, a_n)$ for some L-structure $\mathcal{A}$ and some $a_1, \ldots, a_n \in A$. A *partial n-type in L* is a subset of an n-type in L.

• If $\Gamma(x_1, \ldots, x_n)$ is a partial n-type in L, we say $(a_1, \ldots, a_n)$ *realizes* Γ *in* $\mathcal{A}$ if every formula in Γ is true of $a_1, \ldots, a_n$ in $\mathcal{A}$.

• If $\Gamma(x_1, \ldots, x_n)$ is a partial n-type in L and $\mathcal{A}$ is an L-structure, we say that Γ is *realized in* $\mathcal{A}$ if there is some n-tuple in A that realizes Γ in $\mathcal{A}$. If no such n-tuple exists, then we say that $\mathcal{A}$ *omits* Γ.

2.13. **Facts.** Let $\Gamma(x_1, \ldots, x_n)$ be a set of formulas in L, all of whose free variables are among $x_1, \ldots, x_n$.

• $\Gamma(x_1, \ldots, x_n)$ is an n-type in L if and only if it is a maximal (finitely) satisfiable subset of the set of all formulas in L whose free variables are among $x_1, \ldots, x_n$.

• $\Gamma(x_1, \ldots, x_n)$ is a partial n-type in L if and only if it is (finitely) satisfiable.

2.14. **Definition.** Let T be a theory in L and let $\Gamma = \Gamma(x_1, \ldots, x_n)$ be a partial n-type in L.

• Γ is *consistent with* T if $T \cup \Gamma$ is finitely satisfiable. This is equivalent to saying that Γ is realized in some model of T.

• The set of all n-types that contain T is denoted by $S_n(T)$. These are exactly the n-types in L that are consistent with T.

Let $c_1, \ldots, c_n$ be distinct new constant symbols and let L_n be the language $L(c_1, \ldots, c_n)$ extending T. Let T_n denote the theory whose set of sentences is T but whose language is L_n. The simple observation we give next allows us to identify $S_n(T)$ with $S_0(T_n)$ using the bijection that takes an n-type $\Gamma(x_1, \ldots, x_n)$ to the set of L_n-sentences $\Gamma(c_1, \ldots, c_n)$.

2.15. **Lemma.** *Let $\mathcal{A}, \mathcal{B}$ be L-structures, let a be an n-tuple in $\mathcal{A}$ and let b be an n-tuple in $\mathcal{B}$. The n-type $\mathrm{tp}_{\mathcal{A}}(a)$ can be identified with the complete theory $\mathrm{Th}(\mathcal{A}, a_1, \ldots, a_n)$. In particular, $\mathrm{tp}_{\mathcal{A}}(a) = \mathrm{tp}_{\mathcal{B}}(b)$ if and only if $(\mathcal{A}, a_1, \ldots, a_n) \equiv (\mathcal{B}, b_1, \ldots, b_n)$.*

Proof. Set $\Gamma(x_1,\ldots,x_n) = \mathrm{tp}_{\mathcal{A}}(a)$, and consider the set of formulas

$$\Gamma(c_1,\ldots,c_n) = \{\varphi(c_1,\ldots,c_n) \mid \varphi(x_1,\ldots,x_n) \in \Gamma(x_1,\ldots,x_n)\}$$

in the language L_n. Evidently $\Gamma(c_1,\ldots,c_n) \subseteq \mathrm{Th}(\mathcal{A},a_1,\ldots,a_n)$. Moreover, it is an easy exercise in changing bound variables to show that every sentence in $\mathrm{Th}(\mathcal{A},a_1,\ldots,a_n)$ is logically equivalent to a sentence in $\Gamma(c_1,\ldots,c_n)$. □

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

$$[\varphi(x_1,\ldots,x_n)] = \{\Gamma(x_1,\ldots,x_n) \in S_n(T) \mid \varphi \in \Gamma\}$$

where $\varphi(x_1,\ldots,x_n)$ is any L-formula whose free variables are among $x_1,\ldots,x_n$. The following result is immediate from Corollary 2.10.

2.16. **Proposition.** *Let T be a satisfiable L-theory and $n \geq 0$. The space $S_n(T)$ with the logic topology is a totally disconnected, compact Hausdorff space. Its closed sets are the sets of the form*

$$\{\Gamma' \in S_n(T) \mid \Gamma \subseteq \Gamma'\}$$

where Γ is a set of L-formulas whose free variables are among $x_1,\ldots,x_n$ such that $\Gamma \supseteq T$.
Moreover, the clopen subsets of $S_n(T)$ in this topology are exactly the sets of the form $[\varphi(x_1,\ldots,x_n)]$, where $\varphi(x_1,\ldots,x_n)$ is an L-formula whose free variables are among $x_1,\ldots,x_n$.
Furthermore, two L-formulas $\varphi(x_1,\ldots,x_n)$ and $\psi(x_1,\ldots,x_n)$ are equivalent over T iff the basic open sets $[\varphi(x_1,\ldots,x_n)]$ and $[\psi(x_1,\ldots,x_n)]$ are equal.

Types over a set of parameters

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

2.17. **Definition.** An n-type over X for the theory T is an n-type in $L(X)$ that is consistent with T_X. The space of all n-types over X for the theory T, namely the space $S_n(T_X)$, will be denoted by $S_n(X)$ if the theory T and model $\mathcal{A}$ containing X are understood.

2.18. **Fact.** Let T be an L-theory, $\mathcal{A}$ a model of T, and X a subset of A. Let $\Gamma(x_1,\ldots,x_n)$ be an n-type in $L(X)$. Then $\Gamma \in S_n(X)$ iff Γ is finitely satisfiable in the given structure $(\mathcal{A},a)_{a\in X}$.

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).

Let T be a satisfiable L-theory and $n \geq 0$. Let Σ be a nonempty set of L-formulas whose free variables are among $x_1, \dots, x_n$. Assume that Σ is closed under disjunction and conjunction (up to equivalence over T).

2.19. **Proposition.** *Let $\varphi(x_1, \dots, x_n)$ be an L-formula. The following are equivalent:*
(1) $T \models \varphi$ or $T \models \neg\varphi$ or $T \models \varphi \leftrightarrow \sigma$ for some formula $\sigma \in \Sigma$.
(2) For every $T_1, T_2 \in S_n(T)$, if $\varphi \in T_1$ and $\neg\varphi \in T_2$, then there exists $\sigma \in \Sigma$ such that $\sigma \in T_1$ and $\neg\sigma \in T_2$.

Proof. ($1 \Rightarrow 2$): If $\varphi \in T_1$ and $\neg\varphi \in T_2$, then neither $T \models \varphi$ nor $T \models \neg\varphi$ hold. Thus there exists $\sigma \in \Sigma$ such that $T \models \varphi \leftrightarrow \sigma$. It follows that $\sigma \in T_1$ and $\neg\sigma \in T_2$.

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

We will first show that K is the union of a family of basic open sets from $\mathcal{S}$. Fix $T_1 \in K$; condition (1) implies that there is a subset Σ' of Σ such that $\bigcup\{[\neg\sigma'] \mid \sigma' \in \Sigma'\}$ contains K^c as a subset and does not have T_1 as an element. Since K^c is compact, the set Σ' can be taken to be finite. Since Σ is closed under conjunction, there is a single formula σ' from Σ such that $K^c \subseteq [\neg\sigma']$ and $T_1 \in [\sigma']$. That is, $T_1 \in [\sigma'] \subseteq K$. Therefore K is the union of a family of basic open sets from $\mathcal{S}$.

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

□

Exercises

2.20. Show that the Compactness Theorem (Theorem 1.10) can be derived from Corollary 2.7 by a trivial argument.

2.21. 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 equivalent to a finite L-theory iff $\mathcal{K}$ is axiomatizable.

3. Elementary Maps

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

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

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

3.2. **Fact.** Let $\mathcal{A}, \mathcal{B}, X, f$ be as in the Definition. The function f is elementary with respect to $\mathcal{A}, \mathcal{B}$ if and only if $(\mathcal{A}, a)_{a \in X} \equiv (\mathcal{B}, f(a))_{a \in X}$. In particular, if there exists a function $f\colon X \to B$ that is elementary with respect to $\mathcal{A}, \mathcal{B}$ for some subset X of A (including the empty set), then $\mathcal{A} \equiv \mathcal{B}$.

3.3. **Fact.** If f is an elementary function then f must be 1-1. Moreover, if f is an elementary embedding of $\mathcal{A}$ into $\mathcal{B}$, then f is an embedding of $\mathcal{A}$ into $\mathcal{B}$.

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

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

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

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

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

and set

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

It is an easy exercise to show that R' does not depend on the specific $L(A)$-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 $\mathcal{B}$ are exactly the same as the parameters used to define R in $\mathcal{A}$.

This correspondence between (certain) relations R on A and R' on B 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'.

3.6. **Facts.** (a) Let g be an isomorphism from $\mathcal{A}$ onto $\mathcal{B}$, and let f be any restriction of g to a subset X of A. Then f is elementary with respect to $\mathcal{A}, \mathcal{B}$.
(b) If g is elementary (with respect to $\mathcal{A}, \mathcal{B}$) and f is elementary (with respect to $\mathcal{B}, \mathcal{C}$), and if the range of g is contained in the domain of f, then the composition $f \circ g$ is elementary (with respect to $\mathcal{A}, \mathcal{C}$).
(c) If f is elementary (with respect to $\mathcal{A}, \mathcal{B}$), then f^{-1} is elementary (with respect to $\mathcal{B}, \mathcal{A}$).

3.7. **Fact.** Let I be an index set and U be an ultrafilter on I. Fix a first order language L and an L-structure $\mathcal{A}$. Consider the ultrapower $\mathcal{A}^I/_U$ of $\mathcal{A}$. Define a function δ on A 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 $\mathcal{A}$ into $\mathcal{A}^I/_U$. (This is called the *diagonal embedding*; often one identifies a with $\delta(a)$ for each $a \in A$ and thereby regards $\mathcal{A}$ as an elementary substructure of $\mathcal{A}^I/_U$.)

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

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

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

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

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

$$\mathcal{A} \models \varphi[a_1, \ldots, a_m] \Leftrightarrow \mathcal{B} \models \varphi[a_1, \ldots, a_m].$$

The proof is by induction on the formula φ. By changing bound variables if necessary, we may restrict attention to formulas $\varphi(a_1, \ldots, x_m)$ that have no bound occurrences of any x_j, $j = 1, \ldots, m$.

In the basis step φ is an atomic formula; the displayed equivalence follows from the assumption that $\mathcal{A}$ is a substructure of $\mathcal{B}$.

In the induction step, the cases of propositional connectives are trivial. In the remaining case φ is of the form $\exists y \psi(x_1, \ldots, x_m, y)$, where the statement to be proved is assumed to be true for ψ and y is not among $x_1, \ldots, x_m$. Then we have:

$$\begin{aligned}
\mathcal{A} \models \varphi \quad &\Leftrightarrow \quad \mathcal{A} \models \exists y \psi[a_1, \ldots, a_m] \\
&\Leftrightarrow \quad \mathcal{A} \models \psi[a_1, \ldots, a_m, b] \text{ for some } b \in A \\
&\Leftrightarrow \quad \mathcal{B} \models \psi[a_1, \ldots, a_m, b] \text{ for some } b \in A \\
&\Leftrightarrow \quad \mathcal{B} \models \exists y \psi[a_1, \ldots, a_m] \\
&\Leftrightarrow \quad \mathcal{B} \models \varphi[a_1, \ldots, a_m]
\end{aligned}$$

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 $x_1, \ldots, x_m$. □

3.9. **Facts** (Unions of Chains)**.** Let $(I, \leq)$ be a linearly ordered set. For each $i \in I$ let $\mathcal{A}_i$ be an L-structure, and suppose this indexed family of structures is a chain. That is, for each $i, j \in I$, we suppose $i \leq j \Rightarrow \mathcal{A}_i \subseteq \mathcal{A}_j$.
(1) There is a well defined structure whose universe is the union of the sets A_i and which is an extension of each $\mathcal{A}_i$; 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, $\mathcal{A}_i \preceq \mathcal{A}_j$ holds whenever $i, j \in I$ and $i \leq j$, then the union of this chain of structures is an elementary extension of each $\mathcal{A}_i$. (In this situation we refer to $(\mathcal{A}_i \mid i \in I)$ as an *elementary chain* of L-structures.

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

3.10. **Definition.** Let $\mathcal{A}, \mathcal{B}$ be L-structures and let $\mathcal{F}$ be a nonempty family of functions. We say $\mathcal{F}$ *is a local isomorphism from* $\mathcal{A}$ *onto* $\mathcal{B}$ if it has the following properties:
• Each function in $\mathcal{F}$ is an embedding from a substructure of $\mathcal{A}$ into $\mathcal{B}$.
• (*"back"*) For each $f \in \mathcal{F}$ and each $b \in B$ 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 $a \in A$ there is some $g \in \mathcal{F}$ such that g extends f and a is in the domain of g.
We say $\mathcal{A}$ is locally isomorphic to $\mathcal{B}$ if there is a local isomorphism from $\mathcal{A}$ onto $\mathcal{B}$.

To work effectively with local isomorphisms, we need some facts about maps between substructures.

3.11. **Lemma.** *Let $\mathcal{A}, \mathcal{B}$ be L-structures and let f be an embedding of a substructure of $\mathcal{A}$ into $\mathcal{B}$. Then*
(1) The range of f is a substructure of $\mathcal{B}$.
(2) For each L-term $t(x_1, \ldots, x_n)$ and each $a_1, \ldots, a_n$ in the domain of f,

$$t^{\mathcal{B}}(f(a_1), \ldots, f(a_n)) = f(t^{\mathcal{A}}(a_1, \ldots, a_n)).$$

(3) For each quantifier-free L-formula $\varphi(x_1, \ldots, x_n)$ and each $a_1, \ldots, a_n$ in

the domain of f*,*

$$\mathcal{B} \models \varphi[f(a_1), \ldots, f(a_n)] \quad \Leftrightarrow \quad \mathcal{A} \models \varphi[a_1, \ldots, a_n].$$

Proof. (1) We need to show that $c^{\mathcal{B}}$ 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 $F^{\mathcal{B}}$ for any function symbol F of L. If c is a constant symbol of L, then $c^{\mathcal{A}}$ is in the domain of f and we have $c^{\mathcal{B}} = f(c^{\mathcal{A}})$. If F is an n-ary function symbol of L and $a_1, \ldots, a_n$ are in the domain of f (so $f(a_1), \ldots, f(a_n)$ are arbitrary elements of the range of f), we have $F^{\mathcal{B}}(f(a_1), \ldots, f(a_n)) = f(F^{\mathcal{A}}(a_1, \ldots, a_n))$, 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. □

3.12. **Proposition.** *Let* $\mathcal{A}, \mathcal{B}$ *be* L*-structures and let* $\mathcal{F}$ *be a local isomorphism from* $\mathcal{A}$ *onto* $\mathcal{B}$*. Then each function in* $\mathcal{F}$ *is elementary with respect to* $\mathcal{A}, \mathcal{B}$*. In particular,* $\mathcal{A} \equiv \mathcal{B}$*.*

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

$$\mathcal{A} \models \varphi[a_1, \ldots, a_n] \quad \Longleftrightarrow \quad \mathcal{B} \models \varphi[f(a_1), \ldots, f(a_n)].$$

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

In the base case of the induction φ is an atomic formula, and the desired equivalence is contained in Lemma 3.11(3). 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. □

When constructing local isomorphisms, the following notation and result are often useful.

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

3.14. **Lemma.** *Let* $\mathcal{A}, \mathcal{B}$ *be* L*-structures. Let* J *be a nonempty set and consider two functions* $\alpha \colon J \to A$*,* $\beta \colon J \to B$*. 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*

$$\mathcal{A} \models \varphi[\alpha(j) \mid j \in J] \quad \Leftrightarrow \quad \mathcal{B} \models \varphi[\beta(j) \mid j \in J].$$

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

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 $t^{\mathcal{A}}(\alpha(j) \mid j \in J)$ 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 $t_1^{\mathcal{A}}(\alpha(j) \mid j \in J) = t_2^{\mathcal{A}}(\alpha(j) \mid j \in J)$, 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

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

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. $\square$

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, formulated as a (finite) set of L-sentences.

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

Proof. We apply Proposition 2.19. Fix an L-formula $\varphi(x_1, \ldots, x_m)$. Let Σ be the set of quantifier-free L-formulas whose free variables are among $x_1, \ldots, x_m$. We will verify condition (2) of Proposition 2.19. To that end, consider two dense linear orderings without endpoints $(A, <)$ and $(B, <)$ and elements $a_1, \ldots, a_m \in A, b_1, \ldots, b_m \in B$. We assume that $(A, <) \models \varphi[a_1, \ldots, a_m]$ and that every quantifier-free L-formula satisfied in $(A, <)$ by $a_1, \ldots, a_m$ is satisfied in $(B, <)$ by $b_1, \ldots, b_m$. We need to show $(B, <) \models \varphi[b_1, \ldots, b_m]$.

Let $\mathcal{F}$ be the set of all order preserving functions from a finite subset of A into B. An easy argument shows that $\mathcal{A}$ is a local isomorphism from $(A, <)$ onto $(B, <)$. Our assumptions ensure that there exists $f \in \mathcal{F}$ such that f is defined on $\{a_1, \ldots, a_m\}$ and satisfies $f(a_i) = b_i$ for all $i = 1, \ldots, m$. By Proposition 3.12, the function f is elementary with respect to $(A, <)$ and $(B, <)$. $\square$

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. 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.

3.16. **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 Proposition 2.19. Fix a formula $\varphi(x_1, \dots, x_m)$ in the language of $=$. Let Σ 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 Proposition 2.19. To that end, consider two sets A, B and elements $a_1, \dots, a_m \in A, b_1, \dots, b_m \in B$. We assume that $A \models \varphi[a_1, \dots, a_m]$ and that every formula in Σ that is satisfied by $a_1, \dots, a_m$ in A is satisfied by $b_1, \dots, b_m$ in B. We need to show $B \models \varphi[b_1, \dots, b_m]$.

Our hypotheses ensure that for all $n \geq 1$ we have $A \models \sigma_n \iff B \models \sigma_n$. Therefore, either A and B have the same finite cardinality or both A, B are infinite. Moreover, we also have that for each $1 \leq i < j \leq m$, $a_i = a_j \iff b_i = b_j$. Therefore there is a bijection f from $\{a_1, \dots, a_n\}$ onto $\{b_1, \dots, b_n\}$ such that $f(a_i) = b_i$ for all $i = 1, \dots, n$. Let $\mathcal{F}$ be the set of all 1-1 functions g that extend f and map a finite subset of A into B. It is easy to check that $\mathcal{F}$ is a local isomorphism from $(A, a_1, \dots, a_n)$ onto $(B, b_1, \dots, b_n)$. By Proposition 3.12, $(A, a_1, \dots, a_n) \equiv (B, b_1, \dots, b_n)$, and hence $B \models \varphi[b_1, \dots, b_m]$.

(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 A is one of its models, then A is a model of T_∞ or of T_n for some $n \geq 1$, depending on the cardinality of A. Therefore T is equivalent to $\mathrm{Th}(\mathcal{A})$, 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 equivalent. $\square$

The previous result allows one to 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.

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

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

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

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

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

3.22. 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.

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

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

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

3.24. 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. 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 Section 12.)

4.1. **Definition.** Let $\mathcal{A}$ be an L-structure and let κ be an infinite cardinal. We say that $\mathcal{A}$ is *κ-saturated* if the following condition holds: if X is any subset of A having cardinality $< \kappa$ and $\Gamma(x)$ is any 1-type in $L(X)$ that is finitely satisfiable in $(\mathcal{A}, a)_{a \in X}$, then $\Gamma(x)$ itself is satisfiable in $(\mathcal{A}, a)_{a \in X}$.

4.2. **Facts.** Let $\mathcal{A}$ be an L-structure and κ an infinite cardinal.
(a) If $\mathcal{A}$ is infinite and κ-saturated, then the underlying set of $\mathcal{A}$ has cardinality at least κ.
(b) If $\mathcal{A}$ is finite, then $\mathcal{A}$ is κ-saturated for every κ.
(c) If $\mathcal{A}$ is κ-saturated and X is a subset of A having cardinality $< \kappa$, then the expansion $(\mathcal{A}, a)_{a \in X}$ is also κ-saturated.

Definition 4.1 refers only to realizing 1-types. The following result shows that κ-saturated structures realize partial types in a very rich way.

4.3. **Theorem.** *Let κ be an infinite cardinal and suppose $\mathcal{A}$ is a κ-saturated L-structure. Suppose $X \subset A$ has cardinality $< \kappa$. Let $\Gamma(x_j \mid j \in J)$ be a set of $L(X)$-formulas, where J has cardinality $\leq \kappa$. If Γ is finitely satisfiable in $(\mathcal{A}, a)_{a \in X}$, then Γ is satisfiable in $(\mathcal{A}, a)_{a \in X}$.*

Proof. Let X, J, and $\Gamma(x_j \mid j \in J)$ be as in the statement of the Theorem. Extend Γ so that it is maximal among sets of $L(X)$-formulas with free variables among $(x_j \mid j \in J)$ that are finitely satisfiable in $(\mathcal{A}, a)_{a \in X}$.

Let $<$ be a well ordering of J such that the order type of $(J, <)$ is the cardinal of J. As a consequence, each proper initial segment of $(J, <)$ has cardinality $< \kappa$. For each $k \in J$ let $\Gamma_{\leq k}$ be the set of formulas in Γ whose free variables are among $(x_j \mid j \leq k)$. Note that the maximality of Γ ensures that if φ is any $L(X)$-formula whose free variables are among $(x_j \mid j \leq k)$, then either $\varphi \in \Gamma_{\leq k}$ or $\neg\varphi \in \Gamma_{\leq k}$. Moreover, $\Gamma_{\leq k}$ is closed under conjunction and under application of the existential quantifier $\exists x_k$.

We need to obtain a family $(a_j \mid j \in J)$ of elements of A that satisfies Γ in $(\mathcal{A}, a)_{a \in X}$; we do this by induction over the well ordering $(J, <)$.

Fix $k \in J$ and suppose we have already obtained $(a_j \mid j < k)$ that satisfy all the formulas from Γ that have free variables among $(x_j \mid j < k)$. Let Γ' be the result of substituting a_j for all free occurrences of x_j in $\Gamma_{\leq k}$, for all $j < k$. We see that Γ' is a 1-type (with x_k the allowed free variable) in $L(X \cup \{a_j \mid j < k\})$ that is finitely satisfiable in $(\mathcal{A}, a)_{a \in X \cup \{a_j | j < k\}}$. Since $X \cup \{a_j \mid j < k\}$ has cardinality $< \kappa$ and $\mathcal{A}$ is κ-saturated, there exists a_k in A that satisfies Γ' in $(\mathcal{A}, a)_{a \in X \cup \{a_j | j < k\}}$. It follows that the family $(a_j \mid j \leq k)$ satisfies $\Gamma_{\leq k}$ in $(\mathcal{A}, a)_{a \in X}$.

The result of this construction is a family $(a_j \mid j \in J)$ of elements of A such that for each $k \in J$, the family $(a_j \mid j \le k)$ satisfies $\Gamma_{\le k}$ in $(\mathcal{A}, a)_{a\in X}$. Hence $(a_j \mid j \in J)$ satisfies Γ in $(\mathcal{A}, a)_{a\in X}$, as desired. $\square$

4.4. **Corollary.** *Let $\mathcal{A}$ be a κ-saturated L-structure. If $\mathcal{B} \equiv \mathcal{A}$ and* $\mathrm{card}(B) \le \kappa$*, then there is an elementary embedding of $\mathcal{B}$ into $\mathcal{A}$.*

Proof. Let $(b_j \mid j \in J)$ be an enumeration of B, so $\mathrm{card}(J) \le \kappa$. Let $\Gamma(x_j \mid j \in J)$ be the set of all L-formulas with the indicated free variables that are satisfied by $(b_j \mid j \in J)$ in $\mathcal{B}$. Since $\mathcal{B} \equiv \mathcal{A}$, the set Γ is finitely satisfiable in $\mathcal{A}$. Apply Theorem 4.3 to obtain a family $(a_j \mid j \in J)$ of elements of A that satisfies Γ in $\mathcal{A}$. The function $f\colon B \to A$ that satisfies $f(b_j) = a_j$ for all $j \in J$ is an elementary embedding of $\mathcal{B}$ into $\mathcal{A}$. $\square$

4.5. **Remark.** Suppose $\mathcal{A} \preceq \mathcal{B}$ and $X \subseteq A$, and consider the complete $L(X)$-theory $T_X = \mathrm{Th}((\mathcal{A}, a)_{a\in X})$. Then $(\mathcal{B}, a)_{a\in X}$ is a model of T_X. In particular, it makes sense to speak of a type over X being realized in $(\mathcal{B}, a)_{a\in X}$.

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.

4.6. **Lemma.** *Let $\mathcal{A}$ be an L-structure. There exists an elementary extension $\mathcal{B}$ of $\mathcal{A}$ such that for any subset X of A, every 1-type over X is realized in $(\mathcal{B}, a)_{a\in X}$.*

Proof. Take J to be a set whose cardinality is the number of $L(A)$-formulas, and let I be the collection of all finite subsets of J. Let U be an ultrafilter on I such that for every $j \in J$, U has the set $\{i \in I \mid j \in i\}$ as an element. This is possible because the family of all such sets has the FIP. Let $\mathcal{B}$ be the ultrapower $\mathcal{A}^I/U$ of $\mathcal{A}$ and let δ be the diagonal embedding of $\mathcal{A}$ into $\mathcal{B}$. We know that δ is an elementary embedding, so we may regard $\mathcal{B}$ as an elementary extension of $\mathcal{A}$. (Replace $\mathcal{B}$ by a structure that is isomorphic to $\mathcal{B}$ such that the composition of this isomorphism with δ is the identity on A.)

Let X be any subset of A and let $\Gamma(x)$ be a 1-type in $L(X)$ that is finitely satisfiable in $(\mathcal{A}, a)_{a\in X}$. We need to show that Γ is realized in $(\mathcal{B}, \delta(a))_{a\in X}$. Let $\alpha\colon J \to \Gamma$ be any surjective function. For each $i \in I$ choose $a_i \in A$ that satisfies the formula $\alpha(j)$ in $(\mathcal{A}, a)_{a\in X}$ for every j in the finite set i. Let f be the element $(a_i \mid i \in I)$ of A^I. It follows from Theorem 1.7 that f/U satisfies $\Gamma(x)$ in $(\mathcal{B}, \delta(a))_{a\in X}$. $\square$

4.7. **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 $\Lambda = \{\alpha \mid \alpha$ 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 $\mathcal{A}_0 = \mathcal{A}$. Given $\alpha \in \Lambda$, we define $\mathcal{A}_\alpha$ assuming that $\mathcal{A}_\beta$ is defined for all $\beta < \alpha$. If $\alpha = \beta + 1$ for some β, let $\mathcal{A}_\alpha$ be one of the elementary extensions of $\mathcal{A}_\beta$ that are described in Lemma 4.6. Otherwise α is a limit ordinal and we define $\mathcal{A}_\alpha = \bigcup_{\beta<\alpha} \mathcal{A}_\beta$.

The chain of structures defined by this procedure is an elementary chain; one proves by induction on $\beta \in \Lambda$ that $\mathcal{A}_\alpha \preceq \mathcal{A}_\beta$ holds for all $\alpha < \beta$, using Fact 3.9 at limit ordinals.

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

Note that $\mathcal{A}_\alpha \preceq \mathcal{B}$ for every $\alpha \in \Lambda$, again using Fact 3.9. In particular, $\mathcal{B}$ is an elementary extension of $\mathcal{A}$.

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

Let $\Gamma(x)$ be any 1-type in $L(X)$ that is finitely satisfiable in $(\mathcal{B}, a)_{a\in X}$. Since $\mathcal{A}_\eta \preceq \mathcal{B}$ and $X \subseteq A_\eta$, we see that Γ is finitely satisfiable in $(\mathcal{A}_\eta, a)_{a\in X}$. By construction, this implies that $\Gamma(x)$ is satisfied by some b in $(\mathcal{A}_{\eta+1}, a)_{a\in X}$. Since $\mathcal{A}_{\eta+1} \preceq \mathcal{B}$, it follows that b satisfies $\Gamma(x)$ in $(\mathcal{B}, a)_{a\in X}$, as desired. □

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 $\mathcal{A}^I/U$ 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.

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

Proof. We may assume $I = \mathbb{N}$. Since U is nonprincipal it contains every cofinite subset of $\mathbb{N}$. For each $i \in \mathbb{N}$, let A_i be the underlying set of $\mathcal{A}_i$ and let A be the cartesian product $\prod_{i\in\mathbb{N}} A_i$.

We denote the ultraproduct $\prod_U \mathcal{A}_i$ by $\mathcal{B}$ and its underlying set by B. Let C be any countable subset of B. Let $\Gamma(x_1, \dots, x_n)$ be a set of $L(C)$-formulas such that every finite subset of Γ is satisfiable in $(\mathcal{B}, a)_{a\in C}$. We must show that the entire set Γ is satisfiable in $(\mathcal{B}, a)_{a\in C}$.

For convenience of notation, we will take $n = 1$ and write x for x_1. 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 C such that $\Gamma(x)$ is $\{\varphi_k(x, b_k) \mid k \in \mathbb{N}\}$. This is possible because the language $L(C)$ 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 C). For each $k \in \mathbb{N}$, let f_k be an element of the cartesian product A for which b_k is the $=^{\mathcal{A}}$ equivalence class of f_k.

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

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

Using the Fundamental Theorem of Ultraproducts and the hypothesis that each finite subset of $\Gamma(x)$ is satisfiable in $(\mathcal{B}, a)_{a \in C}$ we have that

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

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

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

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

Note that $\mathbb{N} = G_0 \supseteq G_1 \supseteq \ldots G_k$ and that $G_k \in U$, for all $k \in \mathbb{N}$. Moreover, $\cap\{G_k \mid k \in \mathbb{N}\} = \emptyset$; 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 A whose $=^{\mathcal{A}}$ equivalence class $[g]$ will satisfy every formula from $\Gamma(x)$ in $(\mathcal{B}, a)_{a \in C}$. Fix $i \in \mathbb{N}$ and define $g(i)$ as follows. If $d(i) = 0$ let $g(i)$ be an arbitrary element of A_i. If $d(i) \geq 1$, choose $g(i)$ to be an element of $C_0(i) \cap \ldots \cap C_{d(i)}(i)$, which is guaranteed to be nonempty by the definition of $d(i)$.

It is obvious that for each $k \in \mathbb{N}$, $g(i) \in C_k(i)$ holds 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 $k \in \mathbb{N}$

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

The Fundamental Theorem of Ultraproducts implies that $[g]$ satisfies $\varphi_k(x, b_k)$ in $(\mathcal{B}, a)_{a \in C}$ for all $k \in \mathbb{N}$. $\square$

4.9. **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 $\cap\{F_k \mid k \in \mathbb{N}\}$ 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 $(\mathcal{A}_i)_{i \in I}$ is any family of L structures indexed by I, where L is a countable language, then the ultraproduct $\prod_U \mathcal{A}_i$ is ω_1-saturated.

4.10. 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 3.15.)

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

4.12. Let I be a countable infinite set and U a nonprincipal ultrafilter on I.
• Let $\mathcal{A}$ be the linear ordering $(\mathbb{Q}, <)$. Show that the cardinality of the ultrapower $\mathcal{A}^I/U$ 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 $\mathcal{A}_i$ be a countable infinite L-structure for each $i \in I$. Show that the cardinality of the ultraproduct $\Pi_U(\mathcal{A}_i \mid i \in I)$ is exactly 2^ω.

5. Quantifier Elimination

The method of quantifier elimination, which we introduce in this chapter, is often useful in applications for analyzing definable sets.

5.1. **Definition.** *T admits 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 $T \models \varphi \leftrightarrow \psi$.

Our first criterion for QE comes directly from Proposition 2.19.

5.2. **Theorem.** *Let T be a satisfiable L-theory. The following conditions are equivalent:*
(1) T admits quantifier elimination.
(2) For each $n \geq 1$, every type in $S_n(T)$ 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 Proposition 2.19 with Σ taken to be the set of all quantifier-free formulas whose 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.

5.3. **Notation.** For L-structures $\mathcal{A}, \mathcal{B}$ we let $\mathcal{F}_0(\mathcal{A}, \mathcal{B})$ denote the set of all functions f such that f is an embedding of a finitely generated substructure of $\mathcal{A}$ into $\mathcal{B}$.

Note that $\mathcal{F}_0(\mathcal{A}, \mathcal{B})$ could be empty.

5.4. **Theorem.** *Let T be a satisfiable L-theory. The following conditions are equivalent:*
(1) T admits quantifier elimination;
(2) Whenever $\mathcal{A}, \mathcal{B}$ are models of T, f is in $\mathcal{F}_0(\mathcal{A}, \mathcal{B})$, and $a \in A$, there exists an elementary extension $\mathcal{B}'$ of $\mathcal{B}$ and a function g in $\mathcal{F}_0(\mathcal{A}, \mathcal{B}')$ such that g extends f and a is in the domain of g.
(3) Whenever $\mathcal{A}, \mathcal{B}$ are ω-saturated models of T, either $\mathcal{F}_0(\mathcal{A}, \mathcal{B})$ is empty or it is a local isomorphism from $\mathcal{A}$ onto $\mathcal{B}$.

Proof. $(1 \Rightarrow 2)$: Let $a_1, \dots, a_n$ be generators for the domain of f and let $\Gamma(x_1, \dots, x_n, y) = \mathrm{tp}_{\mathcal{A}}(a_1, \dots, a_n, a)$. Since T admits QE and f preserves the truth of quantifier-free formulas, we see that $\mathrm{tp}_{\mathcal{B}}(f(a_1), \dots, f(a_n)) = \mathrm{tp}_{\mathcal{A}}(a_1, \dots, a_n)$. Note that if $\varphi_1, \dots, \varphi_m \in \Gamma$, then $\exists y(\varphi_1 \wedge \cdots \wedge \varphi_m)$ is satisfied by $a_1, \dots, a_n$ in $\mathcal{A}$. Therefore each such formula is also satisfied by $f(a_1), \dots, f(a_n)$ in $\mathcal{B}$. It follows that the 1-type $\Gamma(f(a_1), \dots, f(a_n), y)$ over $\{f(a_1), \dots, f(a_n)\}$ is finitely satisfiable in $(\mathcal{B}, f(a_1), \dots, f(a_n))$. By Lemma 4.6 there exists $\mathcal{B}' \succeq \mathcal{B}$ and $b \in B'$ such that b realizes $\Gamma(f(a_1), \dots, f(a_n), y)$

in $(\mathcal{B}, f(a_1), \ldots, f(a_n))$. Using Lemma 3.14 we may extend f to an embedding g of $\langle a_1, \ldots, a_n, a\rangle_{\mathcal{A}}$ into $\mathcal{B}'$ with $g(a) = b$.

$(2 \Rightarrow 3)$: Assume $\mathcal{A}, \mathcal{B}$ are ω-saturated models of T. When we apply statement (2) to $\mathcal{A}, \mathcal{B}$, we may take $\mathcal{B}'$ to be $\mathcal{B}$ itself, since the type realized by $g(a)$ in $\mathcal{B}'$ over a finite set of generators for the range of f can be realized in $\mathcal{B}$. (Then we argue as in the previous paragraph.) This shows that $\mathcal{F}_0(\mathcal{A}, \mathcal{B})$ has the "forth" property in Definition 3.10. Applying the same argument to the opposite pair $\mathcal{B}, \mathcal{A}$ shows that $\mathcal{F}_0(\mathcal{A}, \mathcal{B})$ also has the "back" property in that Definition. That is, $\mathcal{F}_0(\mathcal{A}, \mathcal{B})$ is indeed a local isomorphism from $\mathcal{A}$ onto $\mathcal{B}$, as desired.

$(3 \Rightarrow 1)$: We verify condition (2) in Theorem 5.2. Fix $n \geq 1$ and p, q be any two types in $S_n(T)$. Suppose $a_1, \ldots, a_n$ realizes p in $\mathcal{A} \models T$ and $b_1, \ldots, b_n$ realizes q in $\mathcal{B} \models T$. By Theorem 4.7 we may assume that $\mathcal{A}$ and $\mathcal{B}$ are ω-saturated. Suppose p and q contain exactly the same quantifier-free formulas. Using Lemma 3.14 we get an isomorphism f from $\langle a_1, \ldots, a_n\rangle_{\mathcal{A}}$ onto $\langle b_1, \ldots, b_n\rangle_{\mathcal{B}}$ with $f(a_i) = b_i$ for all $i = 1, \ldots, n$. Then $f \in \mathcal{F}_0(\mathcal{A}, \mathcal{B})$, so by statement (3) and Proposition 3.12 we conclude that f is elementary, and thus $p = q$, as desired. $\square$

5.5. **Corollary.** *Let T be a satisfiable L-theory that admits quantifier elimination.*
(1) Suppose $\mathcal{A}, \mathcal{B}$ are models of T. If some substructure of $\mathcal{A}$ can be embedded in $\mathcal{B}$, then $\mathcal{A} \equiv \mathcal{B}$.
(2) If some single L-structure can be embedded into every ω-saturated model of T, then T is complete.

Proof. (1) Let $\mathcal{A}, \mathcal{B}$ be models of T such that some substructure of $\mathcal{A}$ can be embedded in $\mathcal{B}$. By Theorem 4.7 there exist ω-saturated models $\mathcal{A}', \mathcal{B}'$ such that $\mathcal{A}' \succeq \mathcal{A}$ and $\mathcal{B}' \succeq \mathcal{B}$. Then some substructure of $\mathcal{A}'$ embeds in $\mathcal{B}'$, and hence $\mathcal{F}_0(\mathcal{A}', \mathcal{B}')$ is nonempty. It follows from Theorem 5.4(3) that $\mathcal{A}' \equiv \mathcal{B}'$ and hence $\mathcal{A} \equiv \mathcal{B}$.

(2) By part (1) and Theorem 4.7, the assumptions in (2) imply that any two models of T are elementarily equivalent. $\square$

Theory of 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, but it can be analyzed by applying the *method* of quantifier elimination. This means that we formulate a carefully chosen extension, show that the extension admits 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 $\mathcal{A}_0$ of T_{dis} expands in a unique way to a model of T, because the predecessor function and the successor function are definable in $\mathcal{A}_0$.

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

Proof. We verify condition (2) in Theorem 5.4. Let $\mathcal{A}, \mathcal{B}$ be models of T and let $\mathcal{A}_0$ be the substructure of $\mathcal{A}$ generated by the elements $a_1, \ldots, a_m$. We may assume $a_1 < \ldots < a_m$ in $\mathcal{A}$. Further, let f be an embedding of $\mathcal{A}_0$ into $\mathcal{B}$. We may suppose that no subset of $\{a_1, \ldots, a_m\}$ generates $\mathcal{A}_0$. It follows that for each $k \in \mathbb{N}$ and each $i = 2, \ldots, m-1$, 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 $\mathcal{A}$.

Now let a be any element of A that is not in $\mathcal{A}_0$. 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 $\mathcal{B}'$ of $\mathcal{B}$. To accomplish this, we take $\mathcal{B}'$ to be any ω-saturated elementary extension of $\mathcal{B}$, which exists by Theorem 4.7.

For each $i = 1, \ldots, m$ let C_i be the set of all predecessors and successors of a_i in $\mathcal{A}$, including a_i itself. Then each C_i is a convex set in $\mathcal{A}$ and A_0 is the disjoint union of the sets $C_1, \ldots, C_m$. Moreover, a either lies between C_i and C_{i+1} for some $i = 1, \ldots, m-1$, or it lies below C_1, or it lies above C_m, in the ordering of $\mathcal{A}$.

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 $\mathcal{B}$ that consists of all the predecessors and successors of $f(a_i)$. This remains true when we move up to $\mathcal{B}'$. Moreover, the convex sets $f(C_1), \ldots, f(C_m)$ are disjoint and are arranged in order from left to right in the ordering of $\mathcal{B}'$. A simple saturation argument shows that there exist elements $d_1, \ldots, d_{m+1}$ of $\mathcal{B}'$ such that

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

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 $\mathcal{A}$ generated by $\mathcal{A}_0$ 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 admits QE. To conclude that T is complete, we apply Corollary 5.5, 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 $\mathcal{A}, \mathcal{B}$ be models of T_{dis} and let $\mathcal{A}', \mathcal{B}'$ be their unique expansions to models of T. Since T is complete we have $\mathcal{A}' \equiv \mathcal{B}'$. Taking reducts to L_0 we have $\mathcal{A} \equiv \mathcal{B}$. Since $\mathcal{A}, \mathcal{B}$ were arbitrary models of T_{dis} this shows that T_{dis} is complete. □

Another criterion for QE

When we are trying to prove that a theory admits QE using Theorem 5.4 it is sometimes inconvenient that we must extend a given embedding f to *every* element a of the model $\mathcal{A}$. The next result gives a criterion for QE in which we get to choose which element a to treat; the cost is that we must consider embeddings f defined on an arbitrary substructure of $\mathcal{A}$. (In Theorem 5.4 the domain of f is finitely generated.)

5.7. **Theorem.** *Let T be a satisfiable theory in a language L with $\kappa = \mathrm{card}(L)$. The following conditions are equivalent:*
(1) T admits quantifier elimination;
(2) Suppose $\mathcal{A}, \mathcal{B}$ are models of T, $\mathrm{card}(A) \leq \kappa$, and $\mathcal{B}$ is κ^+-saturated; suppose further that $\mathcal{A}_0$ is a proper substructure of $\mathcal{A}$ and that f is an embedding of $\mathcal{A}_0$ into $\mathcal{B}$; then f can be extended properly to an embedding of some substructure $\mathcal{C}$ of $\mathcal{A}$ into $\mathcal{B}$.

Proof. ($1 \Rightarrow 2$): Use Theorems 4.3 and 5.4, statement (2).

($2 \Rightarrow 1$): We assume condition (2) of this Theorem and use it to verify condition (2) of Theorem 5.4. Suppose $\mathcal{A}, \mathcal{B}$ are models of T, $\mathcal{A}_0$ is a finitely generated substructure of $\mathcal{A}$, and f is an embedding of $\mathcal{A}_0$ into $\mathcal{B}$. Fix $a \in A$ and let $\mathcal{C}$ be the substructure of $\mathcal{A}$ generated by $\mathcal{A}_0$ and a. We must show that f can be extended to an embedding of $\mathcal{C}$ into an elementary extension of $\mathcal{B}$.

Since $\mathcal{C}$ has cardinality at most κ, we may replace $\mathcal{A}$ by one of its elementary substructures that contains $\mathcal{C}$ and has cardinality at most κ. Moreover, we replace $\mathcal{B}$ by one of its elementary extensions that is κ^+-saturated. (See Theorem 4.7.)

Let Ω be the set of all extensions of f to embeddings whose domain is a substructure of $\mathcal{A}$ and whose range is a substructure of $\mathcal{B}$. 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 $(\Omega, \leq)$, 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 $(\Omega, \leq)$. Applying condition (2) of the Theorem to the embedding g and the models $\mathcal{A}$ and $\mathcal{B}$, we see that g can be maximal only if it is defined on all of $\mathcal{A}$. In particular g is defined on $\mathcal{C}$, and thus it gives an extension of f as needed to verify condition (2) of Theorem 5.4. □

5.8. 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 T be an L-theory and let T' be the L'-theory whose set of sentences is identical to T. Show that T admits QE if and only if T' admits QE. (Therefore, in showing that T admits QE, it does no harm to assume that its language contains at least one constant symbol.)

5.9. Let L be a first order language and let T be an L-theory that admits 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.)

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

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

is the union of a finite number of open intervals (whose endpoints are in A) and a finite subset of A.

5.11. Let K be a field and let L be the first order language of vector spaces over K, as described in Exercise 3.6. Let T be the theory of infinite K-vector spaces.
• Show that T admits quantifier elimination and use this to prove that T is complete. (Compare Exercise 3.6.)
• Let $\mathcal{A} \models T$ 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 $\mathcal{A}, \mathcal{B}$ of T such that there does not exist any local isomorphism from $\mathcal{A}$ onto $\mathcal{B}$. (Yet $\mathcal{A} \equiv \mathcal{B}$ since T is complete.)

5.12. 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 $S \subseteq \mathbb{Q}$ 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 S or disjoint from S.
• Use the preceding result to show that $\text{Th}(\mathcal{Q})$ does not admit quantifier elimination.

6. Löwenheim-Skolem Theorems

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

Proof. Without loss of generality we may assume $\operatorname{card}(X) = \kappa$; if $\operatorname{card}(X) < \kappa$ then we may replace X by X' such that $X \subseteq X' \subseteq B$ and $\operatorname{card}(X') = \kappa$. Furthermore we may also assume without loss of generality that $c^{\mathcal{B}} \in X$ for all constants in L, since L has at most κ constant symbols.

Now we construct a substructure $\mathcal{A}$ of $\mathcal{B}$ whose universe contains X and such that $\mathcal{A}$ and $\mathcal{B}$ satisfy the condition in Theorem 3.8. This amounts to a family of closure conditions on the universe of $\mathcal{A}$; namely, this set should contain each element of X (including $c^{\mathcal{B}}$ for every constant symbol c of L), it should be closed under all functions $F^{\mathcal{B}}$ where F is a function symbol of L, and it should also be closed under the existential conditions in Theorem 3.8. We take A to be the smallest subset of B that satisfies these conditions. The process of constructing such a set is familiar, but we spell it out carefully in order to show that A has the required cardinality.

Construct a sequence $A_0 \subseteq A_1 \subseteq \ldots$ of subsets of B inductively as follows: $A_0 = X$; for each $n \geq 0$, A_{n+1} is A_n together with an element $b \in B$ from each of the following situations:

(A) for each function symbol F of L and each tuple $\bar{a} \in A_n$, put $b = F^{\mathcal{B}}(\bar{a})$ into A_{n+1}.

(B) for each existential formula $\exists y \varphi[\bar{x}, y] \in L$ and each tuple $\bar{a} \in A_n$, if $\mathcal{B} \models \exists y \varphi[\bar{a}]$, then put into A_{n+1} some $b \in B$ for which $\mathcal{B} \models \varphi[\bar{a}, b]$.

By construction, $A_\infty = \bigcup_n A_n$ contains X and hence contains $c^{\mathcal{B}}$ for every constant c of L; moreover, A_∞ is closed under every function $F^{\mathcal{B}}$ where F is a function symbol of L. Therefore A_∞ is the universe of a substructure $\mathcal{A}$ of $\mathcal{B}$ and A contains the set X. Moreover, $\mathcal{A}$ and $\mathcal{B}$ satisfy the condition in Theorem 3.8, so $\mathcal{A}$ is an elementary substructure of $\mathcal{B}$. Therefore the only remaining condition to prove is $\operatorname{card}(A) = \kappa$. For this, it suffices to show $\operatorname{card}(A_n) \leq \kappa$ for all n. We prove this by induction on n. For $n = 0$ this was given as part of the hypotheses. Suppose A_n has cardinality $\leq \kappa$. Observe that an upper bound on the number of elements b that we add to A_n to obtain A_{n+1} is the product of (a) the number of formulas and function symbols in L and (b) the number of finite tuples from A_n. Each of these numbers is bounded above by κ and therefore the set A_{n+1} also has at most κ elements. (For the elementary facts about cardinal numbers that we are using here, see the book *Naive Set Theory* by P. Halmos.) $\square$

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. Then T has a countable ω-saturated model if and only if for each $n \geq 1$ there are only countably many n-types in L that are consistent with T.*

Proof. ($\Rightarrow$) Let $\mathcal{A}$ be a countable, ω-saturated model of T. By Theorem 4.3, every n-type consistent with T is realized in $\mathcal{A}$. Hence $S_n(T)$ must be countable.

($\Leftarrow$) This proof is patterned after the proofs of Lemma 4.6 and Theorem 4.7, with appropriate modifications to keep structures countable.

Assume $S_n(T)$ is countable for each $n \geq 1$. It follows that for every model $\mathcal{A}$ of T and every finite subset F of A, 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 $\mathcal{B} \succeq \mathcal{A}$, map the type of b in $(\mathcal{B}, a_1, \dots, a_k)$ to the type of $(b, a_1, \dots, a_k)$ in $\mathcal{B}$.

Using Lemma 4.6 followed by the use of Theorem 6.1 (for the cardinal $\kappa = \omega$) we may prove the following version of Lemma 4.6 for the current situation: *Let $\mathcal{A}$ be a countable model of T. There exists a countable elementary extension $\mathcal{B}$ of $\mathcal{A}$ such that for any finite subset F of A, every 1-type over F is realized in $(\mathcal{B}, a)_{a \in F}$.*

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

6.3. **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. $\square$

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

Proof. Since $\mathcal{A}$ is infinite, it has an elementary extension $\mathcal{B}'$ whose cardinality is $\geq \kappa$. By Theorem 6.1 there exists an elementary substructure $\mathcal{B}$ of $\mathcal{B}'$ such that $A \subseteq B$ and $\mathrm{card}(B) = \kappa$. It follows easily that $\mathcal{A} \preceq \mathcal{B}$. Indeed, if $\varphi(\bar{x})$ is any L-formula and $\bar{a}$ is any tuple from A (so $\bar{a}$ is also in B), then we have

$$\mathcal{A} \models \varphi[\bar{a}] \Leftrightarrow \mathcal{B}' \models \varphi[\bar{a}] \Leftrightarrow \mathcal{B} \models \varphi[\bar{a}].$$

$\square$

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

It is a consequence of Theorem 6.4 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.

6.6. **Definition.** Let T be a theory in L and let κ be any cardinal number. 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.

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

Proof. Suppose T is a theory that is κ-categorical, where $\kappa \geq \mathrm{card}(L)$ and T has only infinite models. We need to show that if $\mathcal{A}, \mathcal{B} \models T$, then $\mathcal{A} \equiv \mathcal{B}$. By use of Theorems 6.1 and 6.4 (as needed), we find models $\mathcal{A}', \mathcal{B}'$ with $\mathcal{A}' \equiv \mathcal{A}$, $\mathcal{B}' \equiv \mathcal{B}$, and $\mathrm{card}(A') = \kappa = \mathrm{card}(B')$. (If $\kappa < \mathrm{card}(A)$, use the Downward Löwenheim-Skolem Theorem to find $\mathcal{A}'$ with $\mathcal{A}' \equiv \mathcal{A}$ and $\mathrm{card}(A') = \kappa$. If $\kappa > \mathrm{card}(A)$, use the Upward Löwenheim-Skolem Theorem.)

Since T is κ-categorical we have $\mathcal{A}' \cong \mathcal{B}'$, and hence $\mathcal{A}' \equiv \mathcal{B}'$ by Proposition 3.12. Therefore $\mathcal{A} \equiv \mathcal{A}' \equiv \mathcal{B}' \equiv \mathcal{B}$, and hence $\mathcal{A} \equiv \mathcal{B}$ since $\equiv$ is an equivalence relation. $\square$

Exercises

6.8. Theorem 6.6 states that if T is an L-theory with no finite models and if T is κ-categorical for some cardinal with $\kappa \geq \mathrm{card}(L)$, then T is complete. Show that the "no finite models" assumption 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.

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

6.10. 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 $\mathbb{K}$ be the class of all L-structures that are isomorphic to $\mathcal{P}(S)$ for some set S. Show that $\mathbb{K}$ is not axiomatizable.

6.11. 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 $\operatorname{card}(H) = \tau$ 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. Algebraically Closed Fields

We illustrate the use of Theorem 5.4 by using it to show that the theory ACF of algebraically closed fields admits quantifier elimination. After proving this result we will show how it can be used to obtain some interesting consequences for algebraically closed fields.

We formulate ACF in the language L_r of rings; this language has binary function symbols $+, -, \times$ and constants $0, 1$. In addition to the first order axioms for fields, ACF contains axioms asserting, for each $n \geq 1$, that every nontrivial polynomial of degree n has a root:

$$\forall x_0 \dots \forall x_n \big(x_n \neq 0 \to \exists y(x_n y^n + \cdots + x_1 y + x_0 = 0)\big).$$

Of course the best known algebraically closed field is the field of complex numbers. (This fact is known as the Fundamental Theorem of Algebra.)

In our proof of quantifier elimination 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 (for example, Serge Lang's *Algebra*) contain this 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 quantifier elimination for ACF. Indeed, they can be proved efficiently using the model theoretic ideas discussed here, as we show in Section 10.

7.1. **Theorem.** *The theory ACF admits quantifier elimination.*

Proof. We will verify condition (2) of Theorem 5.4. We need to consider algebraically closed fields $\mathcal{A}, \mathcal{B}$, as well as a finitely generated substructure $\mathcal{A}_0$ of $\mathcal{A}$ and an embedding f of $\mathcal{A}_0$ into $\mathcal{B}$. Given any element a of A, we must prove that f can be extended to an embedding of the substructure of $\mathcal{A}$ generated by A_0 and a into $\mathcal{B}$. Let $\mathcal{B}_0$ be the range of f; then $\mathcal{B}_0$ is a substructure of $\mathcal{B}$ and f is an isomorphism of $\mathcal{A}_0$ onto $\mathcal{B}_0$.

Since $\mathcal{A}_0$ is a substructure of $\mathcal{A}$, it is a subring. Let $\mathcal{A}_0'$ be the field of fractions of $\mathcal{A}_0$ inside $\mathcal{A}$. It is easy to see that f can be extended (in a unique way) to an embedding of $\mathcal{A}_0'$ into $\mathcal{B}$, which we also denote by f; for each $b, c \in A_0$ with $c \neq 0$ we define $f(b/c) = f(b)/f(c)$. Note that $\mathcal{A}_0'$ is not necessarily finitely generated as an L_r-structure, but that is not a problem for our construction; clearly $\mathcal{A}_0'$ is countable, and that is all we will need later. Let $\mathcal{B}_0'$ be the range of this extended f; $\mathcal{B}_0'$ is obviously the field of fractions of $\mathcal{B}_0$ inside $\mathcal{B}$.

Suppose a is algebraic over $\mathcal{A}_0'$. Let $p(x)$ be the minimal polynomial of a over $\mathcal{A}_0'$, so that $p(x)$ is an irreducible polynomial in $\mathcal{A}_0'[x]$ and $p(a) = 0$.

Let $q(x)$ be the corresponding polynomial in $\mathcal{B}'_0[x]$, obtained by applying f to the coefficients of $p(x)$. Since f is an isomorphism of fields, $q(x)$ is irreducible in $\mathcal{B}'_0[x]$. The field $\mathcal{B}$ is algebraically closed, so $q(x)$ has a root in this field. Let b be such a root. It is an elementary exercise to show that f can be extended (in a unique way) to an isomorphism from $\mathcal{A}'_0(a)$ onto $\mathcal{B}'_0(b) \subseteq \mathcal{B}$ such that $f(a) = b$.

Otherwise a must be transcendental over $\mathcal{A}'_0$. Let $\mathcal{B}'$ be an uncountable elementary extension of $\mathcal{B}$. Since $\mathcal{B}'_0$ is a countable substructure of $\mathcal{B}$, there is an element b of $\mathcal{B}'$ that is transcendental over $\mathcal{B}'_0$. We may extend f (in a unique way) to an isomorphism of $\mathcal{A}'_0(a)$ onto $\mathcal{B}'_0(b) \subseteq \mathcal{B}$ such that $f(a) = b$.

In all cases we have extended the original f to an embedding g whose domain is a substructure of $\mathcal{A}$ that contains A_0 and a and whose range is a substructure of an elementary extension of $\mathcal{B}$. By restricting g to $\langle A_0 \cup \{a\}\rangle_{\mathcal{A}}$ we see that condition (2) in Theorem 5.4 holds for any two models of ACF. Consequently, ACF admits QE. $\square$

For each p that is a prime number or 0, we let ACF_p denote the theory of algebraically closed fields that have characteristic p. For any integer $n \geq 2$ let σ_n denote the L_r-sentence $1 + \cdots + 1 = 0$ in which there are n occurrences of 1 in the summation. For each prime p, the theory ACF_p is axiomatized over ACF by the single sentence σ_p. Moreover, ACF_0 is axiomatized over ACF by $\{\neg\sigma_p \mid p \text{ is a prime }\}$.

7.2. **Corollary.** *(i) For each p (0 or a prime), the theory ACF_p is complete. (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 Corollary 5.5. For each p (0 or a prime) let F_p be the prime field of characteristic p. (So F_p is $\mathbb{Q}$ if $p = 0$ and $\mathbb{Z}/p\mathbb{Z}$ if p is a prime.) Evidently F_p embeds in every field of characteristic p, and thus into every model of ACF_p.

(ii) Let σ be a sentence in L_r. If $ACF_0 \models \sigma$ then by Corollary 2.7 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 that there exists a positive integer n such that $ACF \cup \{\neg\sigma_p \mid p \text{ is a prime } \leq n\} \models \neg\sigma$. If follows that there can only exist finitely many primes p such that $ACF_p \models \sigma$. $\square$

7.3. **Example.** Suppose K is any algebraically closed field and $f\colon K^n \to K^n$ 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 $f = (f_1, \ldots, f_n)$ 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 degree 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, $k \models \tau_d$ if and only if for every polynomial map $f\colon k^n \to k^n$ 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 7.2(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 if K_p is the algebraic closure of the prime field of characteristic p, then $K_p \models \tau_d$.

Fix a prime p and let K_p be the algebraic closure of the prime field of characteristic p. Let $f\colon K_p^n \to K_p^n$ be a polynomial map that is 1-1. Fix any element $(a_1, \dots, a_n)$ in K_p^n. There is a finite subfield k of K_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. □

The fact that ACF admits 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, which is a weak form of Hilbert's Nullstellensatz, is an easy consequence of this fact.

7.4. **Corollary.** *Let $k \subseteq K$ be fields 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 K, then this system has a solution in some algebraic extension of k.*

Proof. Without loss of generality we may take K to be algebraically closed, since every field is contained in an algebraically closed field. Let $\widetilde{k}$ be the algebraic closure of k in K; then $\widetilde{k}$ is itself an algebraically closed field. It follows from Theorem 7.1 that $\widetilde{k} \preceq K$. 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 $\widetilde{k}$. □

7.5. **Definition.** Let K be an algebraically closed field and S a subset of K^n. We say that S is *constructible* if it is a finite Boolean combination of zero sets of polynomials that have coefficients in K.

7.6. **Remark.** It follows from Theorem 7.1 that *all* definable sets in K^n are constructible, and conversely. In particular, the collection of all constructible sets is closed under projections.

7.7. **Corollary** (Chevalley). *Let K be an algebraically closed field. If S 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 $h(S)$ 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 $\varphi(\bar{x})$ be a formula in $L(K)$ that defines S in K. Then $h(S)$ is defined in K by the $L(K)$-formula

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

so that $h(S)$ is also constructible, by Remark 7.6. □

7.8. **Definition.** An infinite L-structure $\mathcal{A}$ is *minimal* if every definable subset of A 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 $a_1, \dots, a_n \in A$, the set

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

is either finite or cofinite as a subset of A.

A theory T is *strongly minimal* if every infinite model of T is minimal.

7.9. **Proposition.** *Let T be a strongly minimal theory and let $\varphi(x, y_1, \dots, y_n)$ be any formula in the language of T (in which x is a single variable). Then there exists an integer N with the property that for any model $\mathcal{A}$ of T and any $b_1 \dots, b_n \in A^n$, either the set $\{a \in A \mid \mathcal{A} \models \varphi[a, b_1 \dots, b_n]\}$ or its complement in A has at most N elements.*

Proof. The proof is a straightforward compactness argument. Suppose T is strongly minimal and $\varphi(x, y_1, \dots, y_n)$ is a formula for which the conclusion fails. For each N there must be a model $\mathcal{A}$ of T and parameters $b_1 \dots, b_n \in A^n$, such that both the set $\{a \in A \mid \mathcal{A} \models \varphi[a, b_1 \dots, b_n]\}$ and its complement have more than N elements. This fact can be expressed by a formula $\psi_N(b_1 \dots, b_n)$. From the compactness theorem it follows that T has a model $\mathcal{A}$ with parameters $b_1 \dots, b_n \in A^n$, such that both the set $\{a \in A \mid \mathcal{A} \models \varphi[a, b_1 \dots, b_n]\}$ and its complement are infinite. Therefore $\mathcal{A}$ is a nonminimal model of T, which is a contradiction. □

7.10. **Corollary.** *The theory ACF is strongly minimal.*

Proof. Let K be an algebraically closed field and let S be a definable subset of K. By Theorem 7.1 we can define S 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 S or $K \setminus S$ is contained in the union of finitely many zero sets nonconstant polynomials $p(x)$. It follows that S or $K \setminus S$ must be finite. □

Exercises

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

• For any $a, b \in K$, show that $\text{tp}_{\mathcal{A}}(a) = \text{tp}_{\mathcal{A}}(b)$ 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.

7.12. If T is an L-theory, a model $\mathcal{A}$ of T is called *existentially closed* in $\text{Mod}(T)$ if it satisfies the following condition: whenever $\mathcal{A} \subseteq \mathcal{B} \models T$, $\varphi(x_1, \ldots, x_m, y_1, \ldots, y_n)$ is a quantifier-free formula, and $a_1, \ldots, a_m \in A$, then $\mathcal{B} \models \exists y_1 \ldots \exists y_n \varphi[a_1, \ldots, a_m]$ implies $\mathcal{A} \models \exists y_1 \ldots \exists y_n \varphi[a_1, \ldots, a_m]$.

• Let T be the theory 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.

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 $\mathrm{Th}\,(\mathbb{Z}, +, -, <, 0)$ does not admit 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

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

Here we are using ny to represent the term $y + \cdots + y$ in which there are n copies 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 admit 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:

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

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.

8.1. **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. □

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

In proving this Theorem we use explicit methods for eliminating quantifiers, rather than the model theoretic methods presented in Section 5. To do this we need to introduce some definitions and a Lemma.

8.3. **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. (a) A *universal* formula is a formula in prenex normal form that has only $\forall$ quantifier symbols in its prefix.

8.4. **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 admits 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 existential 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 8.2. We will give an explicit proof of quantifier elimination. The completeness of T follows using Corollary 5.5, 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 8.1 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 disjuction of existential formulas $\exists x \theta$ where each θ is a conjunction of atomic formulas. Arguing as in the proof of Lemma 8.4 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 conjucts of the form $x = t$ but does contain at least one conjuct 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 8.1) is unique. It is now easy to see that $t - j$ also witnesses the truth of θ.

Case (3): θ contains no conjucts of the form $x = t$ but does contain at least one conjuct 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

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

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 $T_0 = \mathrm{Th}(\mathbb{Z}, +, -, <, 0)$.

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 $T_1 = \mathrm{Th}(\mathbb{Z}, +, -, <, 0, 1)$. By looking closer at the relation between T and T_1 we obtain the following result:

8.5. **Corollary.** *T_1 is model complete; that is if $\mathcal{A}, \mathcal{B}$ are models of T_1 and $\mathcal{A} \subseteq \mathcal{B}$, then $\mathcal{A} \preceq \mathcal{B}$.*

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

To show that $\mathcal{A}'$ 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 $\mathcal{A}'$ (since it is satisfied by $\mathcal{A}$). Fix an element a of A. If $a = nb$ for some b in A, then this equation also holds in $\mathcal{B}$, which implies that a satisfies $D_n(x)$ in $\mathcal{B}'$ since it is a model of T. Therefore a satisfies $D_n(x)$ in $\mathcal{A}'$, since it is a substructure of $\mathcal{B}'$. Conversely, suppose a is not of the form nb in $\mathcal{A}$. There must exist a unique $k = 1, \ldots, n-1$ and some $b \in A$ satisfying $a + i = nb$ in $\mathcal{A}$. This equation also holds in $\mathcal{B}'$, which implies that $D_n(x)$ must be false of a in that structure. Hence $D_n(x)$ is also false of a in $\mathcal{A}'$ by the substructure condition.

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

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 $f \colon \mathbb{Z} \to \mathbb{Z}$ 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 8.2 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.

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

$$\text{if } m \geq n, \text{ then } m \in A \Longleftrightarrow m + p \in A.$$

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 $-A = \{-n \mid n \in A\}$.

8.7. **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 8.2 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 8.2 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 $\mathcal{P}$. $\square$

9. Model Theoretic Algebraic Closure

9.1. **Definition.** Let $\mathcal{A}$ be an L-structure and $X \subseteq A$. An element a of A is *algebraic over* X *in* $\mathcal{A}$ if there is an L-formula $\varphi(x, y_1, \ldots, y_n)$ and elements $e_1, \ldots, e_n$ of X such that
(i) $\mathcal{A} \models \varphi[a, e_1, \ldots, e_n]$, and
(ii) $\{c \in A \mid \mathcal{A} \models \varphi[c, e_1, \ldots, e_n]\}$ is finite.

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

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

9.2. **Fact.** Let $\mathcal{A}$ be an algebraically closed field and $X \subseteq A$; let k be the subfield of $\mathcal{A}$ generated by X. For each $a \in A$ we have that $a \in \operatorname{acl}_{\mathcal{A}}(X)$ iff there is a nonconstant polynomial $p(x)$ with coefficients in k such that $p(a) = 0$ in $\mathcal{A}$. (This follows from the fact that ACF admits QE; see Theorem 7.1.) 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.

9.3. **Proposition.** *Let $\mathcal{A}$ be an L-structure. The operation $\operatorname{acl}_{\mathcal{A}}(X)$ defined on all subsets X of A is a closure operation. That is, it satisfies the following two properties for $X, Y \subseteq A$:*
(1) $X \subseteq \operatorname{acl}_{\mathcal{A}}(X)$; and
(2) if $Y \subseteq \operatorname{acl}_{\mathcal{A}}(X)$ then $\operatorname{acl}_{\mathcal{A}}(Y) \subseteq \operatorname{acl}_{\mathcal{A}}(X)$.
Moreover, $\operatorname{acl}_{\mathcal{A}}$ has finite character; that is,
(3) $\operatorname{acl}_{\mathcal{A}}(X)$ is the union of the sets $\operatorname{acl}_{\mathcal{A}}(F)$ where F ranges over the finite subsets of X.

Proof. (1) If $a \in X$, then $a \in \operatorname{acl}_{\mathcal{A}}(X)$ is witnessed by the formula $x = y_1$ with parameter a.
(2) Assume $Y \subseteq \operatorname{acl}(X)$ and $a \in \operatorname{acl}(Y)$. Let $\varphi(x, y_1, \ldots, y_n)$ and $e_1, \ldots, e_n \in Y$ witness the fact that $a \in \operatorname{acl}(Y)$ as in Definition 9.1. Let N be the cardinality of the set $\{c \in A \mid \mathcal{A} \models \varphi[c, e_1, \ldots, e_n]\}$. By changing the formula φ if necessary we may assume that for every $b_1, \ldots, b_n \in A$ the set $\{c \in A \mid \mathcal{A} \models \varphi[c, b_1, \ldots, b_n]\}$ has cardinality at most N, while we continue to have $\mathcal{A} \models \varphi[a, e_1, \ldots, e_n]$ Similarly, let $\psi_j(y_j, z_1, \ldots, z_p)$ and $f_1, \ldots, f_p \in X$ witness the fact that $e_j \in \operatorname{acl}(X)$ for each $j = 1, \ldots, p$. (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 $\sigma(x, z_1, \ldots, z_p)$ given by

$$\exists y_1 \ldots \exists y_n (\varphi(x, \bar{y}) \wedge \psi_1(y_1, \bar{z}) \wedge \ldots \wedge \psi_n(y_n, \bar{z}))$$

with parameters $f_1, \ldots, f_p$ witnesses the fact that $a \in \operatorname{acl}(X)$.
(3) Definition 9.1 implies that $\operatorname{acl}(X)$ is contained in the union of the sets $\operatorname{acl}(F)$ where F ranges over the finite subsets of X. Part (2) of this Proposition implies the reverse containment. □

9.4. **Remark.** It follows from Proposition 9.3 that for every $X \subseteq A$, the set $\operatorname{acl}_{\mathcal{A}}(X)$ is algebraically closed in $\mathcal{A}$. (Just apply part (2) to $Y = \operatorname{acl}(X)$ and then use part (1).)

9.5. **Fact.** Suppose $\mathcal{A}$ is an algebraically closed field and k is a subfield of $\mathcal{A}$. Let K be the set of all $a \in A$ for which there is a nonconstant polynomial $p(x)$ with coefficients in k such that $p(a) = 0$ in $\mathcal{A}$. Using Fact 9.2 and the preceding remark, we have that K is a subfield of $\mathcal{A}$ and that K is algebraically closed in $\mathcal{A}$ (in either of the two equivalent senses of this term). This illustrates the power of quantifier elmination in an algebraic setting.

The following result shows that the model theoretic algebraic closure is to a large extent independent of the structure within which it is computed. In particular, it implies that if $\mathcal{A}, \mathcal{B}$ are L-structures that satisfy $\mathcal{A} \preceq \mathcal{B}$, then $\operatorname{acl}_{\mathcal{A}}(X) = \operatorname{acl}_{\mathcal{B}}(X)$ for any $X \subseteq A$. (Take f to be the identity map on $\operatorname{acl}_{\mathcal{A}}(X)$ in the following Proposition.)

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

Proof. Let $\mathcal{A}, \mathcal{B}, X, Y, f$ be as given in the Proposition. Let Ω be the set of all functions $g \colon X' \to Y'$ such that $X \subseteq X' \subseteq \operatorname{acl}_{\mathcal{A}}(X)$, $Y \subseteq Y' \subseteq \operatorname{acl}_{\mathcal{B}}(Y)$, g is elementary with respect to $\mathcal{A}, \mathcal{B}$, and g extends f. It is easy to show that $(\Omega, \subseteq)$ 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 $\operatorname{acl}_{\mathcal{A}}(X)$. If not, let $a \in \operatorname{acl}_{\mathcal{A}}(X) \setminus X'$. By Proposition 9.3 we have $a \in \operatorname{acl}_{\mathcal{A}}(X')$. Let $\varphi(x, y_1, \dots, y_n)$ be an L-formula and $e_1, \dots, e_n \in X'$ parameters that witness the fact that $a \in \operatorname{acl}(X')$. Moreover, we may suppose that φ and $\bar{e}$ have been chosen so that the finite set $U = \{c \in A \mid \mathcal{A} \models \varphi[c, e_1, \dots, e_n]\}$ has the smallest possible cardinality. Let this cardinality be N.

Since g is an elementary map, the set $V = \{c \in B \mid \mathcal{B} \models \varphi[c, g(e_1), \dots, g(e_n)]\}$ also has cardinality N. Moreover, g maps $X' \cap U$ bijectively onto $Y' \cap V$. Since $X' \cap U$ has cardinality $< N$ (as it does not contain a) there must exist $b \in V \setminus Y'$. Extend g to the map g' defined on $X' \cup \{a\}$ by setting $g'(a) = b$. We will show that g' is elementary with respect to $\mathcal{A}, \mathcal{B}$, contradicting the maximality of g.

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

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

Since g is elementary, we have

$$\mathcal{B} \models \forall x(\varphi \to \psi)[g(e_1), \ldots, g(e_n), g(f_1), \ldots, g(f_p)].$$

Therefore our choice of b implies

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

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 $\mathrm{acl}_{\mathcal{A}}(X)$.

Finally, suppose the given function f has range Y. Let $g\colon \mathrm{acl}_{\mathcal{A}}(X) \to \mathrm{acl}_{\mathcal{B}}(Y)$ be any extension of f that is elementary with respect to $\mathcal{A}, \mathcal{B}$. Let Z be the range of g and suppose, by way of getting a contradiction, that Z is a proper subset of $\mathrm{acl}_{\mathcal{B}}(Y)$. Since $Y \subseteq Z$ we have $\mathrm{acl}_{\mathcal{B}}(Z) = \mathrm{acl}_{\mathcal{B}}(Y)$. Applying the first part of this Proposition to g^{-1} we see that g^{-1} should have an extension that maps $\mathrm{acl}_{\mathcal{B}}(Y)$ into $\mathrm{acl}_{\mathcal{A}}(X)$ and is elementary with respect to $\mathcal{B}, A$. But since the range of g^{-1} is all of $\mathrm{acl}_{\mathcal{A}}(X)$ 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 $\mathrm{acl}_{\mathcal{B}}(Y)$, as claimed. $\square$

Exercises

9.7. Let A be an infinite set, considered as a structure for the language of pure equality. For each $X \subseteq A$, show that $\mathrm{acl}_{\mathcal{A}}(X) = X$.

9.8. Let $\mathcal{A} \models DLO$. For each $X \subseteq A$, show that $\mathrm{acl}_{\mathcal{A}}(X) = X$.

9.9. Let K be a field and let L be the language of vector spaces over K. (See Exercises 3.6 and 5.4.) For each infinite K-vector space V (considered as an L-structure) and each $X \subseteq V$, show that $\mathrm{acl}_V(X)$ is the K-linear subspace of V spanned by X.

9.10. Consider the theory T_{dis} of discrete linear orderings without endpoints. (See Example 5.6.) For $\mathcal{A} \models T_{dis}$ and $X \subseteq A$, describe $\mathrm{acl}_{\mathcal{A}}(X)$.

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

From Proposition 9.3 we know that cl is a closure operation of finite character on the subsets of A. When $\mathcal{A}$ is minimal, cl is actually a *(combinatorial) pregeometry*; this means that cl also satisfies the *Exchange Property*:

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

$$a \in \mathrm{cl}(X \cup \{b\}) \iff b \in \mathrm{cl}(X \cup \{a\}).$$

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

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

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

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

$$\{c \in A \mid \mathcal{A} \models \sigma[c, e_1, \ldots, e_p]\}$$

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

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

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

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

which is a contradiction. □

10.2. **Definition.** Let cl be a pregeometry on the set A; let $X, Y \subseteq A$.
(1) X is *closed* if $\mathrm{cl}(X) = X$.
(2) $\mathrm{cl}(X)$ is the *closure* of X.
(3) (Y closed) X *spans* Y if $\mathrm{cl}(X) = Y$.
(4) X is *independent* if $a \notin \mathrm{cl}(X \setminus \{a\})$ for all $a \in X$.
(5) (Y closed) X is a *basis* for Y if X is independent and X spans Y.

10.3. **Theorem.** *Let* cl *be a pregeometry on the set A; let $X, Y \subseteq A$ with Y closed.*
(1) X is independent if and only if each finite subset of X is independent.
(2) X is a basis for Y if and only if X is maximal among independent subsets of Y. Consequently every closed set has a basis. Indeed, every independent subset of Y is contained in a basis for Y.
(3) If X spans Y, then there exists $Z \subseteq X$ such that Z is a basis of Y.
(4) X is a basis for Y if and only if X is minimal among subsets of Y that span Y.
(5) Suppose X is a basis for Y and $a \in Y$. Then there is a smallest finite set $F \subseteq X$ such that $a \in \mathrm{cl}(F)$. We will call F the support of a in X.
(6) Any two bases for Y have the same cardinality.

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

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

Suppose X is any independent subset of Y. Let Ω be the collection of all independent subsets of Y that contain X. 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 Y, by what was proved in the preceding paragraph.

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

(4) Suppose X is a basis for Y and Z is a proper subset of X. For each $a \in X \setminus Z$ we have $a \notin \mathrm{cl}(X \setminus \{a\}) \supseteq \mathrm{cl}(Z)$, which shows that Z does not span Y. Conversely suppose X is minimal among sets that span Y. We must show X is independent. Otherwise there exists $a \in X$ such that

$a \in \mathrm{cl}(X \setminus \{a\})$. It follows that $\mathrm{cl}(X \setminus \{a\}) = \mathrm{cl}(X)$, contradicting the assumption that X is a minimal spanning set.

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

(6) Let U and V be bases for Y. 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 $\mathrm{card}(U) \leq \mathrm{card}(V)$. For each $a \in U$ there exists a finite set $F(a) \subseteq V$ such that $a \in \mathrm{cl}(F(a))$. Let $F = \cup\{F(a) \mid a \in U\}$. Evidently F spans Y, and since V is a basis for Y it follows from (3) that $F = V$. Since V is infinite it follows that U is also infinite and indeed that $\mathrm{card}(V) = \mathrm{card}(F) \leq \mathrm{card}(U)$. Hence $\mathrm{card}(U) = \mathrm{card}(V)$.

Now we handle the finite case. Let U be a finite basis for Y and let V be any independent subset of Y. By what we proved above, V must be finite. We will show that $\mathrm{card}(V) \leq \mathrm{card}(U)$, which suffices to complete the proof of (4). To do this we prove the following statement by induction on the cardinality of V:
there exists $W \subseteq U$ such that $W \cap V = \emptyset$, $W \cup V$ is a basis for Y, and $\mathrm{card}(W \cup V) = \mathrm{card}(U)$.

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

For the induction step consider an independent set V and suppose the statement is true for all independent sets that are smaller than V. Fix $a \in V$ and let $Z = V \setminus \{a\}$. By the induction hypothesis there exists $W \subseteq U$ such that $W \cap Z = \emptyset$, $W \cup Z$ is a basis for Y, and $\mathrm{card}(W \cup Z) = \mathrm{card}(U)$. Let A be the support of a in $W \cup Z$. Since V is independent, A must meet W. Let b be any element of $A \cap W$. By (5) we have $a \notin \mathrm{cl}((W \setminus \{b\}) \cup Z)$.

We complete the proof by showing that $W \setminus \{b\}$ is the desired subset of U for $V = Z \cup \{a\}$. The Exchange Property yields $b \in \mathrm{cl}((W \setminus \{b\}) \cup (Z \cup \{a\}))$. Therefore $(W \setminus \{b\}) \cup (Z \cup \{a\})$ spans Y. Since b is in the support of a in $W \cup Z$ we have $a \notin \mathrm{cl}((W \setminus \{b\}) \cup Z)$. It follows that $(W \setminus \{b\}) \cup (Z \cup \{a\})$ is independent, that $(W \setminus \{b\}) \cap (Z \cup \{a\}) = \emptyset$, and that $\mathrm{card}((W \setminus \{b\}) \cup (Z \cup \{a\})) = \mathrm{card}(W \cup Z) = \mathrm{card}(U)$. This completes the proof. □

10.4. **Definition.** Let cl be a pregeometry on the set A and $X \subseteq A$. The *dimension* of X, which is denoted $\dim(X)$, is the unique cardinality of a basis for the closed set $\mathrm{cl}(X)$.

10.5. **Definition.** Let $\mathcal{A}$ be a minimal L-structure. The *dimension* of $\mathcal{A}$, denoted $\dim(\mathcal{A})$, is the dimension of the set A with respect to the pregeometry $\mathrm{acl}_{\mathcal{A}}$.

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

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

10.7. **Corollary.** *Let $\mathcal{A}$ and $\mathcal{B}$ be L-structures with $\mathcal{A} \equiv \mathcal{B}$ and $\mathcal{A}$ minimal. Suppose we have independent sets $X \subseteq A$ (with respect to $\mathrm{acl}_{\mathcal{A}}$) and $Y \subseteq B$ (with respect to $\mathrm{acl}_{\mathcal{B}}$), and let $f \colon X \to Y$ be a 1-1 function. Then f is elementary with respect to $\mathcal{A}, \mathcal{B}$.*

Proof. Let Ω be the collection of subsets $S \subseteq X$ such that the restriction of f to S is elementary with respect to $\mathcal{A}, \mathcal{B}$. We regard $\emptyset$ as an element of Ω (justified since $\mathcal{A} \equiv \mathcal{B}$). The partially ordered set $(\Omega, \subseteq)$ satisfies the hypothesis of Zorn's Lemma, so there exists $S \in \Omega$ that is maximal with respect to $\subseteq$. We need to show $S = X$. If not, let a be any element of $X \setminus S$ and let $b = f(a) \in Y \setminus f(S)$. By Proposition 10.6, f restricted to $S \cup \{a\}$ is elementary with respect to $\mathcal{A}, \mathcal{B}$. This contradicts the maximality of S and proves $S = X$. $\square$

10.8. **Theorem.** *Let T be a complete strongly minimal theory with infinite models, and let $\mathcal{A}, \mathcal{B}$ be models of T.*
(1) There is an elementary embedding of $\mathcal{A}$ into $\mathcal{B}$ if and only if $\dim(\mathcal{A}) \leq \dim(\mathcal{B})$.
(2) $\mathcal{A}$ and $\mathcal{B}$ are isomorphic if and only if $\dim(\mathcal{A}) = \dim(\mathcal{B})$.
(3) T is κ-categorical for every cardinal number $\kappa >$ $\mathrm{card}(L)$.

Proof. ($1\Leftarrow$) Since acl defines a pregeometry in each of these structures, there exist bases X for A (with respect to $\mathrm{acl}_{\mathcal{A}}$) and Y for B (with respect to $\mathrm{acl}_{\mathcal{B}}$). By hypothesis $\mathrm{card}(X) \leq \mathrm{card}(Y)$ so there is a 1-1 function f from X into Y. The preceding Corollary yields that f is elementary with respect to $\mathcal{A}, \mathcal{B}$. By Proposition 9.6, f can be extended to a function $g \colon A \to B$ that is elementary with respect to $\mathcal{A}, \mathcal{B}$. It follows easily that g is an elementary embedding from $\mathcal{A}$ into $\mathcal{B}$.

(1⇒) As above, there is a basis X for A (with respect to $\mathrm{acl}_{\mathcal{A}}$). If f is an elementary embedding of $\mathcal{A}$ into $\mathcal{B}$, then $f(X)$ is independent with respect to $\mathrm{acl}_{\mathcal{B}}$. By Theorem 10.3(2) there is a basis Y for $\mathcal{B}$ that contains $f(X)$. It follows that $\mathrm{card}(X) \leq \mathrm{card}(Y)$ and hence $\dim(\mathcal{A}) \leq \dim(\mathcal{B})$.

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

(3) By the Löwenheim-Skolem Theorems (Theorems 6.1 and 6.4) there exist models of T having cardinality κ. Let $\mathcal{A}$ and $\mathcal{B}$ be two such models of T. As in the proof of (1), let X be a basis for A (with respect to $\mathrm{acl}_{\mathcal{A}}$) and Y for B (with respect to $\mathrm{acl}_{\mathcal{B}}$). Because acl is of finite character in each model, and there are fewer than κ formulas in L, a counting argument shows that X and Y must each be of cardinality equal to κ. Now use part (2). □

Exercises

10.9. Let L be the language of pure equality and let T be the theory in L of all infinite sets. From Example 3.16 we know that T admits QE and is complete.
• Show that T is strongly minimal.
• Explain the meaning of the dimension of a given model of T, in the sense of Section 10.

10.10. 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. From Problem 2.2 we know that T admits QE and is complete. Note that $(\mathbb{Z}, S)$ is a model of T, where $S(a) = a + 1$ for all $a \in \mathbb{Z}$; therefore $T = \mathrm{Th}(\mathbb{Z}, S)$.
• Show that T is strongly minimal.
• Explain the meaning of the dimension of a given model of T, in the sense of Section 10.

10.11. Let K be a field and let L be the language of vector spaces over K. Let T be the theory in L of all infinite vector spaces over K. (See Exercises 3.6, 5.4, and 9.3.)
• Show that T is strongly minimal.
It follows that Section 10 applies to infinite K-vector spaces. Exercise 9.3 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 $F\colon X \to Y$ is a K-linear isomorphism. Show that F is an elementary map in the sense of the L-structures V, W. (Note that if K is a finite field, and X, Y are finitely generated, then they are not models of T.)
• If V is an infinite K-vector space and $X \subseteq V$ is a K-linear subspace, show that the model theoretic dimension of X in the sense of algebraic closure in V does not depend on V. Show that this dimension is the same as the

dimension of X in the sense of linear algebra.
• Check that Theorem 10.3 implies all of the standard facts about linearly independent sets, spanning sets, and bases, for arbitrary vector spaces over K.

10.12. Let T be a strongly minimal L-theory and let κ be an infinite cardinal. Let $\mathcal{A}$ be an infinite model of T.
• Show that $\mathcal{A}$ is κ-saturated iff the dimension of $\mathcal{A}$ in the sense of Section 10 is $\geq \kappa$.

10.13. Let L be the language whose only nonlogical symbol is a binary predicate symbol $<$. Let $\mathcal{A}$ be any infinite linear ordering, considered as an L-structure.
• Show that $\mathrm{Th}(\mathcal{A})$ is not strongly minimal.

11. Real Closed Ordered Fields

We now consider the theory $RCOF$ of *real closed ordered fields*, formulated in the language L_{or} of ordered rings. The axioms of this theory are the axioms for ordered rings and the sign change property (intermediate value property) for polynomials with coefficients in the field. This last set of axioms can be formulated as follows: for each $n > 0$ let $t_n(x, y_0, \ldots, y_n)$ be the term $y_n \cdot x^n + \ldots + y_1 \cdot x + y_0$. The sign change property for polynomials of degree at most n is expressed by the following sentence:

$$\forall a \forall b \forall \bar{y}((a < b \wedge t_n(a, \bar{y}) < 0 \wedge t_n(b, \bar{y}) > 0) \rightarrow \exists x(t_n(x, \bar{y}) = 0 \wedge a \leq x \leq b)).$$

Note that the sign change property implies that every positive element is a square and that all polynomials of odd degree have roots. (If $a > 0$ then $x^2 - a$ changes sign over the interval $[0, a+1]$; for large enough $b > 0$, a given odd degree polynomial will change sign on the interval $[-b, b]$.) These properties give an equivalent set of axioms for $RCOF$ as can be shown by an algebraic argument. Two other algebraic characterizations of real closed ordered fields among ordered fields are: (1) no proper algebraic extension can be ordered; (2) the extension formed by adding $\sqrt{-1}$ is algebraically closed. A full discussion of real closed ordered fields may be found in Serge Lang's book *Algebra*.

Obvious models for the theory $RCOF$ are the ordered fields of all the real numbers and all algebraic real numbers. In 1926-27 Artin and Schreier developed the theory of ordered fields and proved that every ordered field $\mathcal{A}$ has a *real closure* (by which we mean a real closed ordered field that is an algebraic extension of $\mathcal{A}$.) Moreover, any two real closures of an ordered field $\mathcal{A}$ are isomorphic over $\mathcal{A}$.

We are going to show that the theory $RCOF$ admits quantifier elimination, using Theorem 5.7. In order to verify condition (2) of that Theorem, we need the following Lemma.

11.1. **Lemma.** *Let $F \subseteq K$ be real closed ordered fields and suppose $b \in K \setminus F$. 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 $b' \in K' \setminus F$. 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 $p(x) \in F[x]$ with coefficients in F we have to prove

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

We are given that this condition holds when $p(x)$ is constant, so we may assume that $p(x)$ is nonconstant. Without loss of generality we may assume that the polynomial $p(x)$ is monic since the equivalence $(\star)$ is preserved under multiplication by elements of F. Moreover, we may assume that $p(x)$ is irreducible in $F[x]$, since the product of any two 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 know is true. Thus we may assume that $p(x)$ is of degree > 1. Since it is irreducible, it has no roots in F. Since F is algebraically closed in K and in K', this means that $p(x)$ also has no roots in either of these fields. Since $p(x)$ is positive for large enough values of x (in K and in K'), it follows from the sign change property for polynomials over real closed ordered fields (the axioms of type (iii)) that $p(x)$ must be everywhere positive in K and in K'. In particular, both $p(b)$ and $p(b')$ must be positive, proving that condition $(\star)$ is true. □

11.2. **Theorem** (Tarski)**.** *The theory RCOF of real closed ordered fields admits quantifier elimination.*

Proof. We apply Theorem 5.7, verifying condition (2) of that result. We are using a countable language so $\kappa = \omega$. Therefore we need to consider real closed ordered fields $\mathcal{A}, \mathcal{B}$ with $\mathcal{A}$ countable and $\mathcal{B}$ ω_1-saturated. We also consider a proper substructure $\mathcal{A}_0$ of $\mathcal{A}$ and an embedding f of $\mathcal{A}_0$ into $\mathcal{B}$. We must show that f can be extended properly to an embedding of some substructure of $\mathcal{A}$ into $\mathcal{B}$.

Let $\mathcal{B}_0$ be the range of f; then $\mathcal{B}_0$ is a substructure of $\mathcal{B}$ and f is an isomorphism of $\mathcal{A}_0$ onto $\mathcal{B}_0$.

We know that $\mathcal{A}_0$ is an ordered subring of $\mathcal{A}$. Since the field of fractions of $\mathcal{A}_0$ is uniquely determined as an ordered field over $\mathcal{A}_0$, we can extend the embedding f to be defined on the field generated in $\mathcal{A}$ by $\mathcal{A}_0$. Therefore we may assume that $\mathcal{A}_0$ is already an ordered subfield of $\mathcal{A}$. A similar argument using the uniqueness of the real closure of an ordered field shows that we may also assume that $\mathcal{A}_0$ is itself a real closed ordered field; in particular, we may assume that $\mathcal{A}_0$ is algebraically closed in $\mathcal{A}$.

Let b be any element of $A \setminus A_0$. Since $\mathcal{B}$ is ω_1-saturated and $\mathcal{A}_0$ is countable, it is possible to find $b' \in B$ that satisfies $a < b \Leftrightarrow f(a) < b'$ for all $a \in A_0$. Using Lemma 11.1 we conclude that f can be extended to an embedding of ordered fields from $\mathcal{A}_0(b)$ onto $\mathcal{B}_0(b')$ by setting $f(b) = b'$. This completes the proof. □

11.3. **Corollary.** *The theory RCOF is complete; hence RCOF is equal to the 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 field $\mathbb{Q}$. The completeness of *RCOF* follows from Corollary 5.5. □

11.4. **Remark.** The theory of the ordered field $\mathbb{R}$ is decidable; this follows immediately from the fact that it equals $RCOF$, 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 $RCOF$ 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 $RCOF$ is sufficiently high that no feasible algorithm for deciding it can exist. Current interest emphasizes subproblems that are defined by syntactic restrictions.

11.5. **Fact.** Let S be any subset of K^n that is definable in the real closed ordered field $(K, +, -, \times, <, 0, 1)$.

(a) If n = 1, show that S must be a finite union of points from K and open intervals whose endpoints are in K.

(b) Show that the closure of S and the interior of S are also definable subsets of K^n, where K is given the topology defined using its ordering.

11.6. **Remark.** Statement (a) of the preceding Fact is expressed by saying that the theory $RCOF$ is *o-minimal*. The study of o-minimal structures is an active area of research today. Combined with Tarski's Theorem, to the effect that definable sets in real closed ordered fields can be defined using quantifier free formulas, statement (b) shows that for each $n \geq 1$, the collection of subsets of $\mathbb{R}^n$ that are quantifier free definable in the ordered field $\mathbb{R}$ 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.*

11.7. **Corollary.** *Let F be a real closed ordered field, and let p, q be polynomials in the variables $x_1, \ldots, x_n$ 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 $F(x_1, \ldots, x_n)$.*

Proof. If f is not a sum of squares in the field $F(x_1, \ldots, x_n)$ 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 $F(x_1, \ldots, x_n)$ that contains $-f$ and all nonzero squares, does not contain 0, and is closed under $+$ and $\times$. The desired linear ordering on the field $F(x_1, \ldots, x_n)$ is defined by taking

$$g < h \iff (h - g) \in P.$$

This ordering on $F(x_1, \ldots, x_n)$ obviously extends the original ordering on F. Let K be a real closed ordered field that extends $F(x_1, \ldots, x_n)$ with this ordering. Now consider the polynomials $p(x_1, \ldots, x_n)$ and $q(x_1, \ldots, x_n)$ as terms in the language $L(F)$. We see that the sentence $\exists x_1 \ldots \exists x_n(p(x_1, \ldots, x_n)q(x_1, \ldots, x_n) < 0)$ is true in K. (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. But this sentence is false in F by hypothesis, contradicting the assumption that f is not a sum of squares. □

Exercises

11.8. Let K be a countable ordered field, considered as an L_{or}-structure, and let $T = \mathrm{Th}(K)$. 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.

11.9. 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 $R(x)$.)

11.10. Use the preceding Exercise and results in Section 11 to show that the theory *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.)

12. Homogeneous Models

In this chapter we prove that every satisfiable theory T has models that are not only highly saturated but are also very homogeneous. Such models play the same role in the setting of very general mathematical structures that the field of complex numbers plays in algebra and number theory. That is, they contain many "ideal elements" with which one can directly calculate and they support many 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.

12.1. **Definition.** Let L be a first order language, $\mathcal{A}$ an L-structure, and κ an infinite cardinal number. We say $\mathcal{A}$ is *strongly κ-homogeneous* if it has the following property for every subset S of A of cardinality $< \kappa$: any elementary map of S into $\mathcal{A}$ can be extended to an automorphism of $\mathcal{A}$.

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.

12.2. **Lemma.** *Suppose $\mathcal{A}$ is κ-saturated and $\mathcal{B} \preceq \mathcal{A}$ satisfies* $\mathrm{card}(B) < \kappa$*. Then any elementary map f between subsets of $\mathcal{B}$ can be extended to an elementary embedding of $\mathcal{B}$ into $\mathcal{A}$.*

Proof. Suppose the domain of the elementary mapping f is S. Then we have $(\mathcal{B}, a)_{a\in S} \equiv (\mathcal{A}, f(a))_{a\in S}$. Moreover, it is easy to see that $(\mathcal{A}, f(a))_{a\in S}$ must also be κ-saturated, since $\mathrm{card}(S) < \kappa$. By Corollary 4.4 there is an elementary embedding of $(\mathcal{B}, a)_{a\in S}$ into $(\mathcal{A}, f(a))_{a\in S}$. This yields an elementary embedding of $\mathcal{B}$ into $\mathcal{A}$ that extends f. $\square$

12.3. **Theorem** (Existence of Strongly Homogeneous Models)**.** *For every infinite cardinal number κ, every structure has a κ-saturated elementary extension $\mathcal{A}$ such that every reduct of $\mathcal{A}$ is strongly κ-homogeneous.*

Proof. Let $\mathcal{A}_0$ be any structure and κ an infinite cardinal number. Let $\tau = \kappa^+$. Using induction over the well ordered set $\{\alpha \mid \alpha < \tau\}$ we construct an elementary chain of structures $\mathcal{A}_\alpha$ $(\alpha < \tau)$ such that $\mathcal{A}_{\alpha+1}$ is $\mathrm{card}(A_\alpha)^+$-saturated for every $\alpha < \tau$. To construct this sequence, at each successor stage (α to $\alpha + 1$) we apply Theorem 4.7 to $\mathcal{A}_\alpha$; at limit stages we take the union of the previously defined structures. Finally, the desired elementary extension $\mathcal{A}$ of $\mathcal{A}_0$ is obtained by setting $\mathcal{A} = \cup\{\mathcal{A}_\alpha \mid \alpha < \tau\}$.

Note that any subset S of A that has cardinality $< \tau$ must be a subset of A_α for some $\alpha < \tau$. (Here we use the fact that $\tau = \kappa^+$ is a regular cardinal.) From this it is immediate that $\mathcal{A}$ is τ-saturated (as in the proof of Theorem 4.7).

It remains to show that every reduct of $\mathcal{A}$ is strongly κ-homogeneous. Let L be any sublanguage of the language of $\mathcal{A}$. Note that the chain of reducts

$(\mathcal{A}_\alpha|L \mid \alpha < \tau)$ is an elementary chain such that $\mathcal{A}_{\alpha+1}|L$ is card$(A_\alpha)^+$-saturated for every $\alpha < \tau$. Moreover, $\mathcal{A}|L$ is the union of this chain.

Let f be any mapping between subsets of A that is elementary with respect to $\mathcal{A}|L$, such that the domain and range of f have cardinality $< \kappa$. As noted above, the domain and range of f are both contained in A_α for some $\alpha < \tau$. Moreover, such a mapping f is elementary with respect to $\mathcal{A}_\alpha|L$, since $\mathcal{A}_\alpha|L \preceq \mathcal{A}|L$. 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 12.2 to $\mathcal{A}_\alpha|L$ and f we obtain an elementary embedding f_α from $\mathcal{A}_\alpha|L$ into $\mathcal{A}_{\alpha+1}|L$ that extends f. We proceed by induction to obtain a sequence of elementary embeddings f_β from $\mathcal{A}_\beta|L$ into $\mathcal{A}_{\beta+1}|L$, 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 12.2 to $\mathcal{A}_{\beta+1}|L$ 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 12.2 to extend g to an elementary embedding f_λ of $\mathcal{A}_\lambda|L$ into $\mathcal{A}_{\lambda+1}|L$. 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 $\mathcal{A}|L$ and that it extends the original elementary mapping f. □

12.4. **Fact.** If $\mathcal{A}$ is κ-saturated, then every reduct of $\mathcal{A}$ is also κ-saturated, by Theorem 4.3. (This was used in the preceding proof.)

The strongly homogeneous models constructed in the proof of Theorem 12.3 are very large. In some situations it is useful to control the cardinality of strongly homogeneous models, as we do in the next result.

12.5. **Theorem** (Countable strongly ω-homogeneous Models)**.** *Assume that L is a countable language, and let T be a complete theory in L. 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 consistent with T. Then there is a countable strongly ω-homogeneous model of T that realizes every partial n-type in $\mathcal{T}_n$ for each n.*

Proof. Since T is complete, there is a countable model $\mathcal{A}_0$ of T in which all the given partial types are realized. We inductively construct an elementary chain $(\mathcal{A}_n \mid n \in \mathbb{N})$ of countable structures and for each $n \geq 1$ a countable set $\mathcal{F}_n$ of automorphisms of $\mathcal{A}_n$, such that the following conditions are satisfied: (1) for all $n \geq 0$, every elementary map between finite subsets of $\mathcal{A}_n$ extends to an automorphism of $\mathcal{A}_{n+1}$ that is a member of $\mathcal{F}_{n+1}$; (2) for all $n \geq 1$ each automorphism of $\mathcal{A}_n$ in $\mathcal{F}_n$ extends to an automorphism of $\mathcal{A}_{n+1}$ 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 $\mathcal{A}_n$ of T together with a countable set $\mathcal{F}_n$ of automorphisms of $\mathcal{A}_n$. Using Theorem 12.3 there is a strongly ω_1-homogeneous elementary extension $\mathcal{B}$ of $\mathcal{A}_n$. (Of course $\mathcal{B}$ need not be countable.) Since $\mathcal{F}_n$ is countable and there are only countably many maps between finite subsets of $\mathcal{A}_n$, there is a countable set $\mathcal{F}$ of automorphisms of $\mathcal{B}$ with the property that each automorphism in $\mathcal{F}_n$ and each elementary map between finite subsets of $\mathcal{A}_n$ extends to an automorphism of $\mathcal{B}$ that is in $\mathcal{F}$. By the Downward Löwenheim-Skolem Theorem there is a countable structure $\mathcal{A}_{n+1} \preceq \mathcal{B}$ such that $A_n \subseteq A_{n+1}$ and such that A_{n+1} is closed under f and f^{-1} for each $f \in \mathcal{F}$. In particular, we have $\mathcal{A}_n \subseteq \mathcal{A}_{n+1}$. Let $\mathcal{F}_{n+1}$ be the set of restrictions of members of $\mathcal{F}$ to $\mathcal{A}_{n+1}$. Then $\mathcal{A}_{n+1}$ and $\mathcal{F}_{n+1}$ have the desired properties.

The desired model $\mathcal{A}$ of T is the union of the chain $(\mathcal{A}_n \mid n \in \mathbb{N})$. Note that by construction every automorphism of $\mathcal{A}_n$ that is a member of $\mathcal{F}_n$ extends to an automorphism of $\mathcal{A}$. □

12.6. **Theorem.** *Any two κ-saturated L-structures that are both of cardinality κ and that are elementarily equivalent are isomorphic.*

Proof. Let $\mathcal{A}, \mathcal{B}$ be κ-saturated structures with $\mathcal{A} \equiv \mathcal{B}$ and $\text{card}(A) = \text{card}(B) = \kappa$. By induction on $\alpha < \kappa$ we construct an increasing chain $(X_\alpha)_{\alpha<\kappa}$ of subsets of A having cardinality $< \kappa$ and an increasing chain $(f_\alpha)_{\alpha<\kappa}$ of functions $f_\alpha \colon X_\alpha \to B$ that are elementary with respect to $\mathcal{A}, \mathcal{B}$ and such that

$$A = \cup\{X_\alpha \mid \alpha < \kappa\} \text{ and } B = \cup\{f_\alpha(X_\alpha) \mid \alpha < \kappa\}.$$

We then obtain the desired isomorphism from $\mathcal{A}$ onto $\mathcal{B}$ by setting $f = \cup\{f_\alpha \mid \alpha < \kappa\}$.

To begin the inductive construction set $X_0 = \emptyset$ and let f_0 be the empty map. Since $\mathcal{A} \equiv \mathcal{B}$, f_0 is elementary with respect to $\mathcal{A}, \mathcal{B}$. Suppose $\alpha < \kappa$ and we have already constructed X_β and f_β for all $\beta < \alpha$. If α is a limit ordinal, we set $X_\alpha = \cup\{X_\beta \mid \beta < \alpha\}$ and $f_\alpha = \cup\{f_\beta \mid \beta < \alpha\}$. If α is a successor ordinal, we may write $\alpha = \lambda + n$ where λ is a limit ordinal and $n \in \mathbb{N}$. Since α is a successor ordinal, $n \geq 1$ and we may set $\beta = \lambda + n - 1$, the predecessor of α. If n is even we extend f_β to f_α using the κ-saturation of $\mathcal{B}$ so that its domain X_α contains a specified element of A; if n is odd we extend f_β to f_α using the κ-saturation of $\mathcal{A}$ so that its range $f(X_\alpha)$ contains a specified element of B. ("Specified elements" are considered in order, relative to some well ordered listing of the universes of $\mathcal{A}$ and $\mathcal{B}$ of order type κ.) The cases are similar and we indicate only how to proceed when n is even. Suppose b is the designated element of A. Since f_β is elementary with respect to $\mathcal{A}, \mathcal{B}$, we have $(\mathcal{A}, a)_{a \in X_\beta} \equiv (\mathcal{B}, f_\beta(a))_{a \in X_\beta}$; also, both of these structures are κ-saturated. Hence there exists $c \in B$ such that $(\mathcal{A}, b, a)_{a \in X_\beta} \equiv (\mathcal{B}, c, f_\beta(a))_{a \in X_\beta}$. Then we set $X_\alpha = X_\beta \cup \{c\}$ and extend f_β to f_α on X_α by setting $f_\alpha(b) = c$. □

12.7. **Corollary.** *If the structure $\mathcal{A}$ is κ-saturated and has cardinality κ, then $\mathcal{A}$ is strongly κ-homogeneous.*

Proof. Suppose $\mathcal{A}$ is κ-saturated and has cardinality κ. Let $S \subseteq A$ have cardinality $< \kappa$ and suppose $f\colon S \to A$ is an elementary map, with respect to $\mathcal{A}$. Then the structures $(\mathcal{A}, a)_{a \in S}$ and $(\mathcal{A}, f(a))_{a \in S}$ are elementarily equivalent, and both of them are κ-saturated and have cardinality κ. The previous results yields that these two structures are isomorphic; any isomorphism between them is an automorphism of $\mathcal{A}$ that extends f. Hence $\mathcal{A}$ is strongly κ-isomorphic. $\square$

In the rest of this chapter we explore the relations among several notions of "richness" for L-structures.

12.8. **Definition.** Let L be a first order language, $\mathcal{A}$ an L-structure, and κ an infinite cardinal number.
(1) $\mathcal{A}$ is *κ-homogeneous* if it has the following property for every subset S of A of cardinality $< \kappa$: any elementary mapping of S into $\mathcal{A}$ can be extended to an elementary mapping of $S \cup \{b\}$ into $\mathcal{A}$, for each $b \in A$.
(2) $\mathcal{A}$ is *κ-universal* if every structure $\mathcal{B}$ that satisfies $\mathrm{card}(\mathcal{B}) < \kappa$ and $\mathcal{B} \equiv \mathcal{A}$ can be elementarily embedded into $\mathcal{A}$.

12.9. **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 $\mathrm{card}(L) < \kappa$. *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 $\mathcal{A}$ be κ-saturated. Corollary 4.4 shows that $\mathcal{A}$ is κ^+-universal. To show that $\mathcal{A}$ is κ-homogeneous, consider a subset S of A whose cardinality is less than κ and let $f\colon S \to A$ be an elementary map with respect to the structure $\mathcal{A}$. Then $(\mathcal{A}, a)_{a \in S} \equiv (\mathcal{A}, f(a))_{a \in S}$ and both of these structures are κ-saturated. Therefore, for any $b \in A$ there exists $c \in A$ such that $(\mathcal{A}, b, a)_{a \in S} \equiv (\mathcal{A}, c, f(a))_{a \in S}$. The desired extension of f can be obtained by setting $f(b) = c$.

(c) Let $\mathcal{A}$ be κ-homogeneous and κ-universal, and suppose $\mathrm{card}(L) < \kappa$. Let $S \subseteq A$ with $\mathrm{card}(S) < \kappa$, and consider a 1-type $\Gamma(x)$ in $L(S)$ that is finitely satisfiable in $\mathrm{Th}((\mathcal{A}, a)_{a \in S})$. There is an $L(S)$ structure $(\mathcal{B}, f(a))_{a \in S}$ and an element $b \in B$ such that b realizes $\Gamma(x)$ in $(\mathcal{B}, f(a))_{a \in S}$. Since the cardinality of $L(S)$ is $< \kappa$, the Downward Löwenheim-Skolem Theorem implies that we may assume $\mathrm{card}(B) < \kappa$. Since $\mathcal{A}$ is κ-universal, there exists an elementary embedding g of $\mathcal{B}$ into $\mathcal{A}$. The composition $g \circ f$ maps S into A and is an elementary map with respect to $\mathcal{A}$. Since $\mathrm{card}(S) < \kappa$

and $\mathcal{A}$ 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 $c \in \mathrm{dom}(h)$ satisfies $h(c) = g(b)$, then c must realize $\Gamma(x)$ in $(\mathcal{A}, a)_{a \in S}$.

(d) Let $\mathcal{A}$ be a κ-homogeneous structure, and let $f\colon S \to A$ be an elementary map with $\mathrm{card}(S) < \kappa$ and $S \subseteq A$. Then we have $(\mathcal{A}, a)_{a \in S} \equiv (\mathcal{A}, f(a))_{a \in S}$. By an argument similar to the one used to prove part Theorem 12.6, we can inductively extend f to an increasing chain of elementary mappings whose union is an automorphism of $\mathcal{A}$. Thus $\mathcal{A}$ is strongly κ-homogeneous. $\square$

12.10. Let $\mathcal{A}, \mathcal{B}$ be L-structures that are elementarily equivalent. Show that there exist elementary extensions $\mathcal{A}'$ of $\mathcal{A}$ and $\mathcal{B}'$ of $\mathcal{B}$ such that $\mathcal{A}' \cong \mathcal{B}'$.

12.11. Let $\mathcal{A}, \mathcal{B}$ be L-structures and let f be a nonempty elementary map from a subset of A into B. Show that there exist elementary extensions $\mathcal{A}'$ of $\mathcal{A}$ and $\mathcal{B}'$ of $\mathcal{B}$ and an isomorphism g of $\mathcal{A}'$ onto $\mathcal{B}'$ such that g is an extension of f.

12.12. Let L be the language whose nonlogical symbols consist of infinitely many constant symbols $\{c_n \mid n \in \mathbb{N}\}$. Let T be the L-theory whose axioms are $c_m \neq c_n$ for all distinct $m, n \in \mathbb{N}$. It follows from Example 3.16(ii) 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 5.5(2).
• Which countable model of T is ω-saturated?
• Which countable models of T are strongly ω-homogeneous.

12.13. 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. From Problem 2.2 we know that T admits QE and is complete. From Exercise 10.2 we know that T is strongly minimal and we understand the meaning of the dimension of a model of T. Note that $(\mathbb{Z}, S)$ is a model of T, where $S(a) = a + 1$ for all $a \in \mathbb{Z}$; therefore $T = \mathrm{Th}(\mathbb{Z}, S)$ and this model of T obviously has dimension 1.
• Which countable models of T are strongly ω-homogeneous?

12.14. Let $\mathcal{A}$ be an L-structure and $B \subseteq A$. Recall that $R \subseteq A^m$ is called *B-definable in $\mathcal{A}$* if there is an L-formula $\varphi(x_1, \ldots, x_m, y_1, \ldots, y_n)$ and parameters $b_1, \ldots, b_n$ from B such that

$$R = \{(a_1, \ldots, a_m) \in A^m \mid \mathcal{A} \models \varphi[a_1, \ldots, a_m, b_1, \ldots, b_m]\}.$$

Now suppose $\mathcal{A}$ is κ-saturated and strongly κ-homogeneous and $B \subseteq A$ has $\mathrm{card}(B) < \kappa$. Suppose further that $R \subseteq A^m$ is A-definable in $\mathcal{A}$.
• Show that R is B-definable in $\mathcal{A}$ iff R is fixed setwise by every automorphism of $\mathcal{A}$ that fixes B pointwise.

13. Omitting Types

The main result of this chapter (Theorem 13.3) gives a sufficient condition for a given countable family of partial types to be omitted from some model of a theory T in a countable language. If T is a complete theory, this result yields a condition on an n-type $\Gamma(x_1, \ldots, x_n)$ that is necessary and sufficient for the existence of a model of T that omits $\Gamma(x_1, \ldots, x_n)$. (Theorem 13.8) This result can be used to study the countable models of T. Most of the results in this chapter require that the language be countable.

13.1. **Definition.** Let T be a satisfiable theory in the language L and let $\Sigma(x_1, \ldots, x_n)$ be a partial n-type in L. We say that T *locally omits* $\Sigma(x_1, \ldots, x_n)$ if for any formula $\varphi(x_1, \ldots, x_n)$ that is consistent with T, there is a formula $\sigma(x_1, \ldots, x_n) \in \Sigma(x_1, \ldots, x_n)$ such that $\varphi \wedge \neg\sigma$ is consistent with T.

13.2. **Remark.** The notion of "local omitting" may seem more natural when rephrased topologically: T locally omits $\Sigma(x_1, \ldots, x_n)$ iff the closed subset $\{\Gamma \in S_n(T) \mid \Sigma \subseteq \Gamma\}$ has empty interior in the space $S_n(T)$.

13.3. **Theorem** (Omitting Types Theorem)**.** *Suppose L is a countable language, and let T be a satisfiable theory in L. For each $k \geq 1$ let Σ_k be a partial n_k-type in L that is locally omitted by T. Then there is a countable model $\mathcal{A}$ of T such that for all k, Σ_k is omitted in $\mathcal{A}$.*

Proof. In order to keep the notation simpler, we first consider the case of a single partial 1-type $\Sigma(x)$. Let L' be the language obtained from L by adding a countable set of new constants $\{c_n \mid n \geq 1\}$. Let $\varphi_1, \varphi_2, \ldots$ list the sentences of L'. Starting with T we construct an increasing sequence $T = T_0 \subseteq T_1 \subseteq \ldots$ of satisfiable sets of sentences in L' such that each T_{n+1} is a finite extension of T_n and the following conditions are satisfied:

(a) For all $m \geq 1$, T_m contains φ_m or $\neg\varphi_m$;

(b) If $m \geq 1$ and $\varphi_m = \exists y \psi(y) \in T_m$, then $\psi(c_p) \in T_m$ for some $p \geq 1$;

(c) For each $m \geq 1$, there is some $\sigma(x) \in \Sigma(x)$ such that $\neg\sigma(c_m) \in T_m$.

First we show that it is sufficient to construct such an increasing chain of L'-theories. Let $T' = \bigcup_{m=1}^{\infty} T_m$. Then $T \subseteq T'$ and T' is a maximal satisfiable set of formulas in L'. Also, for any formula $\psi(y) \in L'$, if $\exists y \psi(y)$ is in T', then $\psi(c_p) \in T'$ for some p. Note that these are the conditions that also appear in the usual proof of the Gödel Completeness Theorem. As in that proof, we define an L'-prestructure $\mathcal{A}$ with $A = \{c_n \mid n \geq 1\}$ by interpreting the nonlogical symbols of L as follows:

(1) If P is a k-ary predicate symbol in L, let

$$P^{\mathcal{A}}(c_{i_1}, \ldots, c_{i_k}) \Leftrightarrow P(c_{i_1}, \ldots, c_{i_k}) \in T'.$$

(2) If F is a k-ary function symbol in L, take $F^{\mathcal{A}}(c_{i_1}, \ldots, c_{i_k})$ to be the earliest c_p for which $F(c_{i_1}, \ldots, c_{i_k}) = c_p$ is in T'. (Note that the sentence $\exists y(F(c_{i_1}, \ldots, c_{i_k}) = y)$ is in T' since it is a valid sentence and T' is maximal satisfiable; therefore by (b) above there exists c_p such that $F(c_{i_1}, \ldots, c_{i_k}) = c_p$ is in T'.)

(3) If c is a constant in L, take $c^{\mathcal{A}}$ to be the earliest c_p for which $c = c_p$ is in T'. (As in (2), condition (b) above ensures that some such c_p exists.)

It is routine to show (by induction on formulas) that for any L'-sentence φ, one has

$$\mathcal{A} \models \varphi \quad \Leftrightarrow \quad \varphi \in T'.$$

In particular, this shows that $\mathcal{A}$ is a prestructure, since any instance of an equality axiom (indeed, any valid sentence) is a member of T'.

Furthermore, no element of A satisfies all formulas in $\Sigma(x)$; this is ensured by condition (c) above. Thus $\mathcal{A}$ nearly satisfies the conclusion of the Theorem; the only problem is that $\mathcal{A}$ will generally be a prestructure rather than a structure. ($\mathcal{A} \models c_m = c_n$ if and only iff the sentence $c_m = c_m$ is in T', and this may happen even when $m \neq n$.)

By applying the quotient construction discussed in Appendix 2 of Section 1, we obtain an L' structure $\mathcal{B}$ as a quotient of the prestructure $\mathcal{A}$ such that $\mathcal{B}$ is a model of T' (and thus a model of T) that omits $\Sigma(x)$.

This argument shows that it suffices to construct the increasing chain $T = T_0 \subseteq T_1 \subseteq \ldots$ satisfying the conditions above, including (a),(b),(c). We set $T_0 = T$ and define T_m for $m \geq 1$ by induction. Given T_{m-1}, with $m \geq 1$, we construct T_m as follows:

(a) Let

$$T'_{m-1} = \begin{cases} T_{m-1} \cup \{\varphi_m\}, & \text{if this is satisfiable} \\ T_{m-1} \cup \{\neg\varphi_m\}, & \text{otherwise.} \end{cases}$$

Note that T'_{m-1} is satisfiable.

(b) Suppose $\varphi_m \in T'_{m-1}$ and $\varphi_m = \exists y \psi(y)$. Choose c_p to be the first new constant not in ψ or T'_{m-1}. Now $T'_{m-1} \cup \{\psi(c_p)\}$ is satisfiable (else $T'_{m-1} \models \neg\psi(c_p)$ so $T'_{m-1} \models \forall y \neg\psi(y)$; i.e. $T'_{m-1} \models \neg\varphi_m$, which is a contradiction). Let $T''_{m-1} = T'_{m-1} \cup \{\psi(c_p)\}$.

(c) Suppose $\neg\sigma(c_m)$ is not consistent with T''_{m-1} for all $\sigma \in \Sigma$. Then $T''_{m-1} \models \sigma(c_m)$ for all $\sigma \in \Sigma$. There are L-formulas $\psi_1, \ldots, \psi_k$ and constants $c_1, \ldots, c_N$ of L' such that $T''_{m-1} = T \cup \{\psi_j(c_1, \ldots, c_N) \mid j = 1, \ldots, k\}$ (Choose $N \geq m$ so the new constants in T''_{m-1} are among $c_1, \ldots, c_N$. Choose variables $z_1, \ldots, z_N$ not occurring in $T''_{m-1} \setminus T$. Let ψ_j be the result of replacing c_i by z_i in the jth sentence of $T''_{m-1} \setminus T$.). So $T \cup \{\psi_j(c_1, \ldots, c_N) \mid j = 1, \ldots, k\} \models \sigma(c_m)$ for all $\sigma \in \Sigma$. Consider $\varphi(z_m) = \exists z_1 \ldots \exists z_{m-1} \exists z_{m+1} \ldots \exists z_N \bigwedge_{j=1}^{k} \psi_j(z_1, \ldots, z_N)$. Now $\varphi(z_m)$ is consistent with T, and $T \cup \{\varphi(z_m)\} \models \sigma(z_m)$ for all $\sigma \in \Sigma$. Hence $\varphi(x)$ is consistent with T and $T \cup \{\varphi(x)\} \models \sigma(x)$ for all $\sigma \in \Sigma$, contradicting the hypothesis that Σ is locally omitted. So, we can finally define T_m to be equal to the set $T''_{m-1} \cup \{\neg\sigma(c_m)\}$, where $\sigma \in \Sigma$ is chosen so that this set is satisfiable. It is clear that this set T_m satisfies all of the conditions (a), (b), and (c). (Note that if $m \geq 1$ and φ_m is in T_m, then φ_m must be in T'_{m-1}, since otherwise $\neg\varphi_m \in T'_m \subseteq T_m$.)

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 $\Sigma_k(x_1, \ldots, x_{n_k})$, $k \geq 1$ (each locally omitted). Requirements (a) and (b) of the construction remain unchanged. For (c), enumerate all finite sequences $\alpha = (k, c_{i_1}, \ldots, c_{i_{n_k}})$ where $k \geq 1$, and $c_{i_1}, \ldots, c_{i_{n_k}}$ are new constants. Then condition (c) becomes:

For all $m \geq 1$, T_m contains a formula $\neg\sigma(c_{i_1}, \ldots, c_{i_{n_k}})$ where $\sigma \in \Sigma_k$ and $\alpha = (k, c_{i_1}, \ldots, c_{i_{n_k}})$ is the m^{th} sequence in the enumeration of all such sequences. □

13.4. **Remark.** The Omitting Types Theorem can be rephrased topologically as follows: For each positive integer k, let K_k be a closed subset of $S_{n_k}(T)$ that has empty interior in the logic topology. Then there is a countable model $\mathcal{A}$ of T that omits every type in the union $\bigcup\{K_k \mid k \geq 1\}$. Note that if we set $K'_n = \bigcup\{K_k \mid n_k = n\}$, then we are omitting the union $\bigcup\{K'_n \mid n \geq 1\}$. Moreover, K'_n is a typical *meager* subset of $S_n(T)$ for each $n \geq 1$.

13.5. **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 theory, but not omitted in any model of that theory is the following: Let I be an uncountable set and let L be the language whose nonlogical symbols are the distinct constants $\{c_i \mid i \in I\} \cup \{d_n \mid n \in N\}$. T has axioms $\neg c_i = c_j$ for all $i, j \in I, i \neq j$. Now $\Sigma(x) = \{\neg x = d_n \mid n \in N\}$ is locally omitted, but not omitted. Indeed, every model of T is uncountable, while any structure that omits Σ must be countable.

13.6. **Definition.** Let T be a satisfiable theory and let $\Sigma(x_1, \ldots, x_n)$ be an n-type consistent with T. We say $\Sigma(x_1, \ldots, x_n)$ is *principal* (relative to T) if there is a formula $\varphi(x_1, \ldots, x_n) \in \Sigma(x_1, \ldots, x_n)$ such that $T \models \varphi(x_1, \ldots, x_n) \to \sigma(x_1, \ldots, x_n)$ holds for every formula $\sigma(x_1, \ldots, x_n) \in \Sigma(x_1, \ldots, x_n)$. We say that $\varphi(x_1, \ldots, x_n)$ is *a complete formula in* T if for every formula $\sigma(x_1, \ldots, x_n)$ in L, exactly one of the conditions $T \models \varphi \to \sigma, T \models \varphi \to \neg\sigma$ holds.

13.7. **Remarks.** Note that the n-type $\Sigma(x_1, \ldots, x_n)$ is principal iff the singleton $\{\Sigma\}$ is an open set in $S_n(T)$. That is, principal n-types correspond to isolated points of $S_n(T)$.

If $\varphi(x_1, \ldots, x_n)$ is a complete formula in T, then there is a unique n-type $\Sigma(x_1, \ldots, x_n)$ that is consistent with T and contains $\varphi(x_1, \ldots, x_n)$, and this n-type is necessarily principal relative to T. Conversely, every principal type relative to T is determined in this way by a formula that is complete in T.

For complete theories T in a countable language, the next result gives a characterization of those types that can be omitted in some model of T.

13.8. **Theorem.** *Let T be a complete theory in a countable language and let $\Sigma(\bar{x})$ be an n-type consistent with T. Then $\Sigma(\bar{x})$ is principal if and only if it is realized in every model of T (if and only if it is realized in every countable model of T).*

Proof. The second equivalence is a consequence of the Downward Löwenheim-Skolem Theorem. We prove the first equivalence. Let T be complete, L be countable and $\Sigma(\bar{x})$ an n-type consistent with T. If Σ is not principal, then T locally omits Σ (since Σ is maximal satisfiable), so that Σ is omitted in some model of T by Theorem 13.3. Conversely, we need to prove that every principal type (relative to T) is realized in every model of T. Let Σ be a principal n-type, φ the complete formula determining Σ and let $\mathcal{A} \models T$. We know that $\varphi(\bar{x})$ is consistent with T, hence $\exists x_1 \dots \exists x_n \varphi(x_1, \dots, x_n)$ is true in some model of T. Since T is complete, this implies that $\mathcal{A} \models \exists x_1 \dots \exists x_n \varphi(x_1, \dots, x_n)$. Let $a_1, \dots, a_n$ satisfy φ in $\mathcal{A}$. We have $T \cup \{\varphi\} \models \sigma$ for all $\sigma \in \Sigma$, so that $a_1, \dots, a_n$ realizes $\Sigma(x_1, \dots, x_n)$ in $\mathcal{A}$. $\square$

13.9. **Definition.** We say that a structure $\mathcal{A}$ is *atomic* if every n-tuple in $\mathcal{A}$ satisfies a complete formula in $\mathrm{Th}(\mathcal{A})$. (Equivalently: $\mathcal{A}$ is an atomic model of T if every n-type realized in $\mathcal{A}$ is principal relative to $\mathrm{Th}(\mathcal{A})$.)

13.10. **Fact.** (a) If $\varphi(\bar{x}, y)$ is a complete formula in T, then so is $\exists y \varphi(\bar{x}, y)$; (b) if $\mathcal{A}$ is an atomic model and $a \in A$, then $(\mathcal{A}, a)$ is also atomic.

13.11. **Theorem.** *(a) If $\mathcal{A}$ and $\mathcal{B}$ are countable, atomic models, and $\mathcal{A} \equiv \mathcal{B}$, then $\mathcal{A}$ and $\mathcal{B}$ are isomorphic.*
(b) A countable atomic model is strongly ω-homogeneous.

Proof. (a) Let $\mathcal{A}$ and $\mathcal{B}$ be countable atomic models of T. The proof that $\mathcal{A} \cong \mathcal{B}$ is done by a back-and-forth argument, where the key ingredient is homogeneity:

Claim: If $(\mathcal{A}, a_1, \dots, a_k) \equiv (\mathcal{B}, b_1, \dots, b_k)$ and if $a \in A$ then there exists $b \in B$ such that $(\mathcal{A}, a_1, \dots, a_k, a) \equiv (\mathcal{B}, b_1, \dots, b_k, b)$.

Using a familiar back-and-forth inductive argument, this is sufficient to build an isomorphism (noting that the roles of $\mathcal{A}$ and $\mathcal{B}$ can be interchanged.)

Proof of the claim. $\mathrm{tp}_{\mathcal{A}}(a_1, \dots, a_k, a)$ is a $(k+1)$-type consistent with T. Therefore it is principal (because $\mathcal{A}$ is atomic); say it is generated by the complete formula $\varphi(x_1, \dots, x_k, y)$. We see that $\mathcal{A} \models$

$\exists y\varphi(\bar{x}, y)[a_1, \ldots, a_k]$; our elementary equivalence hypothesis therefore implies $\mathcal{B} \models \exists y\varphi(\bar{x}, y)[b_1, \ldots, b_k]$, so we can find $b \in B$ such that $\mathcal{B} \models \varphi[b_1, \ldots, b_k, b]$. Hence $\mathrm{tp}_{\mathcal{B}}(b_1, \ldots, b_k, b) = \mathrm{tp}_{\mathcal{A}}(a_1, \ldots, a_k, a)$ and thus $(\mathcal{A}, a_1, \ldots, a_k, a) \equiv (\mathcal{B}, b_1, \ldots, b_k, b.)$ This completes the proof of the claim, and therefore the proof of part (a).

Part (b) follows immediately from part (a) and the fact that the structure $(\mathcal{A}, a_1, \ldots, a_n)$ is atomic whenever $\mathcal{A}$ is atomic. (See Fact 13.10.) $\square$

13.12. **Definition.** A model of a theory T is called a *prime model* of T if it can be elementarily embedded into each model of T.

13.13. **Theorem.** *Let T be a complete theory in a countable language and let $\mathcal{A}$ be a model of T. Then $\mathcal{A}$ is a prime model of T if and only if $\mathcal{A}$ is countable and atomic.*

Proof. We assume that the language considered here is countable.

($\Leftarrow$) Suppose the model $\mathcal{A}$ of the complete theory T is countable and atomic. To show that $\mathcal{A}$ is prime we use the "forth" part of a "back-and-forth" construction: Let $\mathcal{A} = \{a_0, a_1, a_2, \ldots\}$, and $\mathcal{B}$ be any model of T. Since $\mathcal{A}$ is atomic, there exists a complete formula $\varphi_0(x)$ satisfied in $\mathcal{A}$ by a_0. Since $\varphi_0(x)$ is consistent with T, there exists a $b_0 \in B$ such that $\mathcal{B} \models \varphi_0(b_0)$. Because $\varphi_0(x)$ is a complete formula (with respect to T) we know that a_0 has the same type (in $\mathcal{A}$) that b_0 has (in $\mathcal{B}$). Next, let $\varphi_1(x, y)$ be a complete formula satisfied in $\mathcal{A}$ by a_0, a_1. Then $T \models \varphi_0(x) \to \exists y\varphi_1(x, y)$, since φ_0 is complete; hence there exists $b_1 \in B$ such that $\mathcal{B} \models \varphi_1(b_0, b_1)$.

Continuing inductively in this way, we get a map $f : A \to B$ with $f(a_n) = b_n$ for each $n \geq 0$. At each stage of the construction we can ensure that f is an elementary map on the set $\{a_0, \ldots, a_n\}$. Therefore it is an elementary map on its entire domain, which is all of A. Hence f is the desired elementary embedding of $\mathcal{A}$ into $\mathcal{B}$.

($\Rightarrow$) Suppose $\mathcal{A}$ is a prime model of the complete theory T. Then $\mathcal{A}$ is necessarily countable as $\mathcal{A}$ can be elementarily embedded in every model of T, and T has countable models by the Downward Löwenheim-Skolem Theorem. Let $\bar{a} \in A$ be an n-tuple. For any model $\mathcal{B}$ of T, we have an elementary embedding $f : A \to B$. Hence $\mathcal{B}$ realizes $\mathrm{tp}_{\mathcal{A}}(\bar{a})$. Since $\mathrm{tp}_{\mathcal{A}}(\bar{a})$ is realized in every model of T it is principal (by Theorem 13.8). It follows that $\mathcal{A}$ is an atomic model. $\square$

For the rest of this chapter we consider the *existence* of atomic models.

Note that if $\mathcal{A}$ is atomic, then so is every elementary substructure of $\mathcal{A}$. Therefore, if T is a complete theory in a countable language and T has an atomic model, then it has a countable atomic model, by the Downward Löwenheim-Skolem Theorem.

13.14. **Theorem.** *Let T be a complete theory in a countable language L. Then T has an atomic model if and only if every L-formula that is consistent with T is an element of some principal type consistent with T.*

Proof. ($\Rightarrow$) Let $\mathcal{A}$ be an atomic model of T. If $\varphi(x_1, \dots, x_n)$ is a formula consistent with T, then there exist $a_1, \dots, a_n \in A$ such that $\mathcal{A} \models \varphi[a_1, \dots, a_n]$. That is, $\varphi(x_1, \dots, x_n)$ is a member of $\mathrm{tp}_{\mathcal{A}}(a_1, \dots, a_n)$, which is principal, since $\mathcal{A}$ is an atomic model.

($\Leftarrow$) This direction uses the Omitting Types Theorem. For each $n \geq 1$ let

$$\Sigma_n(x_1, \dots, x_n) = \{\neg\varphi(x_1, \dots, x_n) \mid \varphi \text{ is a complete formula relative to } T\}.$$

A model of T is atomic iff it omits all of these partial types Σ_n. We use the Omitting Types Theorem to prove that there exists a model of T that omits all of these partial types. To apply this Theorem, we must show that Σ_n is locally omitted by T for each $n \geq 1$. Fix n and assume Σ_n is not locally omitted. Then there exists a formula $\psi(x_1, \dots, x_n)$ consistent with T such that for no $\sigma \in \Sigma_n$ is $T \cup \{\psi \wedge \neg\sigma\}$ satisfiable; that is, $T \cup \{\psi\} \models \sigma$ for all $\sigma \in \Sigma_n$. This means that for no complete formula $\varphi(x_1, \dots, x_n)$ (relative to T) can the set $T \cup \{\psi \wedge \varphi\}$ be satisfiable (by the way Σ_n was defined). This violates our assumption on T; this contradiction shows that each Σ_n is locally omitted by T. $\square$

The next result gives a useful sufficient condition for the existence of an atomic model of a theory T in a countable language.

13.15. **Theorem.** *Let T be a complete theory in a countable language, and suppose that for all n there are strictly fewer than continuum many different n-types consistent with T. Then T has a countable atomic model.*

Proof. We prove the contrapositive: If T has no atomic model, then for some $n \geq 1$, T has at least 2^ω distinct n-types. The proof is by a tree argument.

Suppose T has no atomic model. Using Theorem 13.14, there exists $n \geq 1$ and a formula $\varphi(x_1, \dots, x_n)$ consistent with T such that $\varphi(\bar{x})$ is not in any principal n-type of T. That is, if $\tau(\bar{x})$ is any formula such that $\varphi(\bar{x}) \wedge \tau(\bar{x})$ is consistent with T, then τ is not complete relative to T. From this it follows that if $\varphi(\bar{x}) \wedge \tau(\bar{x})$ is consistent with T then there is some further formula $\chi(\bar{x})$ such that $\varphi(\bar{x}) \wedge \tau(\bar{x}) \wedge \chi(\bar{x})$ and $\varphi(\bar{x}) \wedge \tau(\bar{x}) \wedge \neg\chi(\bar{x})$ are each consistent with T.

This can be used to build a binary tree of formulas (with free variables among $x_1, \dots, x_n$) with the property that each branch of the tree is consistent with T, and distinct branches are inconsistent with each other. Each branch yields a nonempty closed subset of $S_n(T)$, and these closed sets partition $S_n(T)$. Since there are 2^ω many branches, it will follow that $S_n(T)$ has cardinality at least 2^ω.

The partially ordered set used to index this binary tree will be denoted $(\Lambda, \leq)$. Here $\Lambda = \{0,1\}^{<\omega}$ is the set of all finite sequences from $\{0,1\}$, including the empty sequence, which we denote by $\emptyset$. For $s, t \in \Lambda$, we write $s \leq t$ to mean that s is an initial subsequence of t. A typical position in the tree is indexed by an element s of Λ, and its immediate successors under $\leq$ are indexed by the sequences $s0$ and $s1$.

We associate to each $s \in \Lambda$ an L-formula $\psi_s(\bar{x})$, proceeding by induction on the length of s. Each $\psi_s(\bar{x})$ will be consistent with T. Further, if $s \leq t$, then $\psi_t(\bar{x})$ will be the conjunction of $\psi_s(\bar{x})$ with some other formula. Finally, for any $s \in \Lambda$, the formulas $\psi_{s0}(\bar{x})$ and $\psi_{s1}(\bar{x})$ will be logically inconsistent with each other; indeed, one of them will contain a conjunct whose negation is a conjunct of the other.

At the top of the tree, $s = \emptyset$ and we set $\psi_\emptyset(\bar{x})$ equal to the formula $\varphi(\bar{x})$ that was identified at the beginning of this proof. Proceeding inductively, suppose we have defined $\psi_s(\bar{x})$ and that this formula is consistent with T and is the conjunction of $\varphi(\bar{x})$ with some other formula. As discussed above, there exists a formula $\chi_s(\bar{x})$ for which both $\psi_s(\bar{x}) \wedge \chi_s(\bar{x})$ and $\psi_s(\bar{x}) \wedge \neg\chi_s(\bar{x})$ are consistent with T. We then define $\psi_{s0}(\bar{x}) = \psi_s(\bar{x}) \wedge \chi_s(\bar{x})$ and $\psi_{s1}(\bar{x}) = \psi_s(\bar{x}) \wedge \neg\chi_s(\bar{x})$. Evidently this construction produces formulas satisfying the conditions in the previous paragraph.

As we go down a branch of the tree, we get a list of formulas $\psi_s(\bar{x})$ corresponding to the nodes on that branch. More precisely:

For each function $\alpha : \mathbb{N} \to \{0,1\}$ set $\Sigma_\alpha(\bar{x}) = \{\psi_{\alpha(0)\alpha(1)\dots\alpha(k)} \mid k \geq 0\}$. Each such set of formulas is a partial n-type consistent with T. Therefore, for each such α we may choose $p_\alpha \in S_n(T)$ such that $\Sigma_\alpha \subseteq p_\alpha$.

If $\alpha \neq \beta$ then $\Sigma_\alpha \cup \Sigma_\beta$ is not consistent with T. Indeed, if k is the least argument at which α and β disagree and if $s = \alpha(0)\dots\alpha(k-1)$, then $\Sigma_\alpha \cup \Sigma_\beta$ will contain both $\psi_{s0}(\bar{x})$ and $\psi_{s1}(\bar{x})$, which are logically contradictory. It follows that $p_\alpha \neq p_\beta$ whenever α, β are distinct functions from $\mathbb{N}$ into $\{0,1\}$. Hence the cardinality of $S_n(T)$ is at least 2^ω, as desired. □

13.16. **Corollary.** *Let T be a complete theory in a countable language. If T has a countable ω-saturated model, then T also has a countable atomic model.*

Proof. If there exists a countable ω-saturated model of T, then by Theorem 6.2 there are only countably many types consistent with T. Hence T has a countable atomic model by Theorem 13.15 and the Downward Löwenheim-Skolem Theorem. □

Exercises

13.17. Let T be a complete L-theory and let Σ be a partial n-type in L. If T has a model that omits Σ, show that Σ is locally omitted by T.

13.18. Let T be a complete L-theory and $p \in S_n(T)$. Show that p is locally omitted by T iff p is not an isolated point in the compact space $S_n(T)$.

13.19. Let T be a complete theory in a countable language, with no finite models. Show that T has a countable atomic model iff for each $n \geq 1$ the set of isolated points is dense in the space $S_n(T)$.

13.20. Let T be a complete theory in a countable language. For each positive integer k let $\Sigma_k(\bar{x})$ be a partial n_k-type in L that is omitted in some model $\mathcal{A}_k$ of T. Show that there is a single countable model $\mathcal{A}$ of T that omits $\Sigma_k(\bar{x})$ for all k.

13.21. Let L be a countable language and let L' be the result of adding countably many new predicate symbols $\{P_1, P_2, \ldots\}$ to L. Let T be a complete theory in the language L' and let $\Gamma(x_1, \ldots, x_n)$ be a set of formulas in L. Let T_m be the set of sentences in T that contain P_n only for $n = 1, \ldots, m$. Assume that for each m, T_m has a model that omits $\Gamma(x_1, \ldots, x_n)$. Show that T has a model that omits $\Gamma(x_1, \ldots, x_n)$.

13.22. Let T be a complete theory in a countable language and assume T has no finite models. Show that T is ω-categorical iff T has a countable model that is both atomic and ω-saturated.

13.23. 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. (Example 3.16(ii), Exercise 10.1)
+ Infinite vector spaces over a field K. (Exercises 3.6, 5.4, 9.3, 10.3)
+ ACF_p for a fixed characteristic p.
+ Bijections without a finite cycle. (Problem 2.2, Exercise 10.2)
+ Discrete linear orderings without endpoints. (Example 5.6, Exercises 5.3, 9.4)
+ Discrete linear orderings with minimum but no maximum. (Problem 3.1)
+ Descending equivalence relations with infinite splitting of classes. (Problem 4.1)
+ Dense linear orderings with increasing sequence of elements. (Problem 4.2)
For each of these theories, do the following:
• Show that T has a countable (infinite) atomic model.
• Try to describe the countable atomic model of T as a clear, specific mathematical structure. (According to Theorem 13.11, the countable infinite atomic model of a complete theory is unique up to isomorphism, if such a model exists.)
• For each principal n-type p that is consistent with T, try to give explicitly a complete formula contained in p.

In this chapter we consider ω-categorical theories in a countable language.

14.1. **Theorem** (Engeler, Ryll-Nardzewski, Svenonius)**.** *Let T be a complete theory in a countable language and suppose T has only infinite models. The following conditions are equivalent:*
(1) T is ω-categorical;
(2) For each positive integer n, every n-type consistent with T is principal;
(3) For each positive integer n, there are only finitely many n-types consistent with T;
(4) For each positive integer n and $x = x_1, \ldots, x_n$, there are finitely many formulas $\varphi_1(x), \ldots, \varphi_{k_n}(x)$ such that each formula $\sigma(x)$ is equivalent in T to $\varphi_j(x)$ for some $j = 1, \ldots, k_n$.

Proof. Let T be a complete theory in a countable language. We will show (1) $\Rightarrow$ (2) $\Rightarrow$ (3) $\Rightarrow$ (4) $\Rightarrow$ (1).

(1) $\Rightarrow$ (2): If T is ω-categorical, then any two countable models of T are isomorphic and hence realize the same types. Therefore, no type consistent with T is omitted from any model of T. This implies that all types are principal.

(2) $\Rightarrow$ (3): Fix $n \geq 1$. We argue by contradiction. Assume that all n-types consistent with T are principal and that there are infinitely many n-types consistent with T. Each of them is determined by a complete formula, so there exist infinitely many complete formulas $\varphi_1(x_1, \ldots, x_n), \varphi_2(x_1, \ldots, x_n), \ldots$ in L that are pairwise inequivalent in T. Since each $\varphi_i(x_1, \ldots, x_n)$ is a complete formula, we have $T \cup \{\varphi_i\} \models \neg\varphi_j$ whenever $i \neq j$.

Take $\{\varphi_i\}$ to be a maximal list of inequivalent complete formulas; we have $T \cup \{\neg\varphi_i \mid i \geq 1\}$ is inconsistent. (Every type is principal, so every tuple in every model of T satisfies a complete formula.) By the Compactness Theorem, there is some $N \in \mathbb{N}$ such that $T \cup \{\neg\varphi_1, \ldots, \neg\varphi_N\}$ is inconsistent, i.e. $T \models \varphi_1 \vee \ldots \vee \varphi_N$. But $T \cup \{\varphi_{N+1}\}$ is satisfiable and implies $\neg(\varphi_1 \vee \ldots \vee \varphi_N)$. This contradiction implies that there can only be finitely many φ_i.

(3) $\Rightarrow$ (4): Fix $n \geq 1$. Let $t_1, \ldots, t_N$ be a list of all the n-types consistent with T. It is an elementary fact that if φ_1, φ_2 are formulas and if for all types t consistent with T we have $\varphi_1 \in t \iff \varphi_2 \in t$, then $T \models \varphi_1 \leftrightarrow \varphi_2$. For each formula $\varphi(x_1, \ldots, x_n)$, let $F(\varphi) = \{j \mid \varphi \in t_j, j = 1, \ldots, N\}$. If $F(\varphi_1) = F(\varphi_2)$ then $T \models \varphi_1 \leftrightarrow \varphi_2$. There are only finitely many distinct sets of the form $F(\varphi)$ and hence there are only finitely many inequivalent formulas (relative to T) in the variables $x_1, \ldots, x_n$.

(4) $\Rightarrow$ (1): We use (4) to show that every model of T is atomic. The ω-categoricity of T follows using Theorem 13.11.

Assume $\varphi_1(x_1, \ldots, x_n), \ldots, \varphi_k(x_1, \ldots, x_n)$ are the finitely many formulas given by condition (4) (where $x_1, \ldots, x_n = x_1, \ldots, x_n$ and k depends on n.) Given any $\mathcal{A} \models T$ and $a_1, \ldots, a_n \in A$, we want to find a complete formula φ such that $\mathcal{A} \models \varphi[\bar{a}]$. Consider

$$\psi(x_1, \ldots, x_n) = \bigwedge_{\mathcal{A} \models \varphi_j[\bar{a}]} \varphi_j(x_1, \ldots, x_n) \wedge \bigwedge_{\mathcal{A} \models \neg\varphi_j[\bar{a}]} \neg\varphi_j(x_1, \ldots, x_n)$$

Clearly $\mathcal{A} \models \psi(x_1, \ldots, x_n)[a_1, \ldots, a_n]$ (by the way ψ was defined). We claim that ψ is a complete formula relative to T. To see this, consider an arbitrary formula $\varphi(x_1, \ldots, x_n)$; there is a j such that $T \models \varphi \leftrightarrow \varphi_j$. So either $\varphi(x_1, \ldots, x_n)$ is implied by $T \cup \{\psi\}$ or $\neg\varphi(x_1, \ldots, x_n)$ is implied by $T \cup \{\psi\}$. But this implies that the type of $(a_1, \ldots, a_n)$ in $\mathcal{A}$ is principal, and hence that every model of T is atomic. $\square$

14.2. **Remark.** Let $\mathcal{A}$ be a countable structure for a countable first order language and suppose $T = \mathrm{Th}(\mathcal{A})$ is ω-categorical. Let G be the automorphism group of $\mathcal{A}$, acting coordinatewise on A^n for each $n \geq 1$. Then G has only finitely many distinct orbits on A^n for each n. This is an immediate consequence of condition (3) in Theorem 14.1 and the fact that the unique countable model of an ω-categorical theory is strongly ω-homogeneous. (See Theorems 6.2 and 12.5 or, alternatively, Theorem 13.11 and the existence results for atomic models later in Section 13.)

Infinite permutation groups of this kind have turned out to be very interesting; they are treated in the book *Oligomorphic Permutation Groups* by Peter Cameron.

The next result gives a sufficient condition for $\mathrm{Th}(\mathcal{A})$ to be ω-categorical that is based on automorphism group considerations of this kind.

14.3. **Theorem.** *Let L be a countable language, let $\mathcal{A}$ be any L-structure, and $T = \mathrm{Th}(\mathcal{A})$. If $G = \mathrm{Aut}(\mathcal{A})$ has only finitely many orbits on A^n for each $n \geq 1$, then T is ω-categorical.*

Proof. Let $\mathcal{A}$ satisfy the given hypotheses. We will show that $\mathcal{A}$ realizes every type that is consistent with T. The automorphism condition on $\mathcal{A}$ implies that $\mathcal{A}$ can only realize finitely many n-types for each n. Therefore T is ω-categorical since it satisfies condition (3) of Theorem 14.1.

So, let $\Sigma(x_1, \ldots, x_n)$ be any n-type consistent with T. Given $\sigma(x_1, \ldots, x_n) \in \Sigma(x_1, \ldots, x_n)$, $T \models \exists x_1 \ldots \exists x_n \sigma(x_1, \ldots, x_n)$ (since T is complete), hence there is $\bar{a} \in A^n$ with $\mathcal{A} \models \sigma[\bar{a}]$. Let $F \subseteq A^n$ be a finite set that selects one n-tuple from each orbit under the action of G. The $\bar{a}$ satisfying σ in $\mathcal{A}$ can be taken from F. Suppose $\Sigma(x_1, \ldots, x_n)$ is not realized in $\mathcal{A}$. For each $\bar{a} \in F$ we get $\sigma_{\bar{a}} \in \Sigma(x_1, \ldots, x_n)$ such that $\mathcal{A} \models \neg\sigma_{\bar{a}}[\bar{a}]$. Consider $\sigma(x_1, \ldots, x_n) = \bigwedge_{\bar{a} \in F} \sigma_{\bar{a}} \in \Sigma$. There is some $\bar{b} \in F$ satisfying $\sigma(x_1, \ldots, x_n)$ in $\mathcal{A}$. However, this implies that $\bar{b}$ satisfies $\neg\sigma_{\bar{b}}$ since it is a conjunct of σ. This is a contradiction. $\square$

14.4. **Example.** Let $\mathcal{A} = (\mathbb{Q}, <)$. We will show directly (using additional structure on $\mathbb{Q}$) that $G = \mathrm{Aut}(\mathcal{A})$ has only finitely many orbits on $\mathbb{Q}^n$ for each n. The main idea is the following: given $q_0 < q_1 < \ldots < q_n$ in $\mathbb{Q}$, there is $g \in G$ such that $g(k) = q_k$ for $k = 0, \ldots, n$. 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 14.3 proves that $\mathrm{Th}(\mathbb{Q}, <)$ 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.)

14.5. **Fact.** Let T be an ω-categorical complete theory in a countable language.
(a) If $\mathcal{A}$ is any model of T and $a_1, \ldots, a_n \in A$, then $\mathrm{Th}(\mathcal{A}, a_1, \ldots, a_n)$ is ω-categorical.
(b) If $\mathcal{A}$ is any model of T and if $\mathcal{B}$ is any structure for any countable language such that the universe of $\mathcal{B}$ and all of its interpretations of predicate symbols and function symbols are 0-definable in $\mathcal{A}$, then $\mathrm{Th}(\mathcal{B})$ is also ω-categorical. (The universe of $\mathcal{B}$ is allowed to be a set of n-tuples from the universe of $\mathcal{A}$.) In particular, any restriction of T to a smaller language is ω-categorical.

14.6. **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 (possibly infinite). Then T has at most countably many n-types for all $n \geq 1$ (as countably many countable models realize all the types that can be realized.) Then by Theorem 13.15, T has a countable atomic model $\mathcal{A}$ and by Theorem 6.2, T has a countable ω-saturated model $\mathcal{B}$. If $\mathcal{A} \cong \mathcal{B}$ then every type consistent with T is realized in $\mathcal{A}$ and hence is principal, so T is ω-categorical, by Theorem 14.1.

So we assume that $\mathcal{A} \not\cong \mathcal{B}$. For some $n \geq 1$ there exists at least one n-type $\Sigma(x_1, \ldots, x_n)$ that is consistent with T and is not principal. Then $\mathcal{B}$ realizes $\Sigma(x_1, \ldots, x_n)$ and $\mathcal{A}$ doesn't. Let $b_1, \ldots, b_n$ realize $\Sigma(x_1, \ldots, x_n)$ in $\mathcal{B}$ and let $T' = \mathrm{Th}(\mathcal{B}, b_1, \ldots, b_n)$ in $L(b_1, \ldots, b_n)$.

Note that T' cannot satisfy condition (4) of Theorem 14.1, 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 $(\mathcal{C}, c_1, \ldots, c_n)$ be a model of T' that is not isomorphic to $(\mathcal{B}, b_1, \ldots, b_n)$. Then $\mathcal{A} \not\cong \mathcal{C}$, since $\mathcal{C}$ realizes Σ and $\mathcal{A}$ doesn't. Moreover, $\mathcal{B} \not\cong \mathcal{C}$, since $\mathcal{B}$ is ω-saturated and $\mathcal{C}$ isn't; otherwise, by Theorem 12.6, $(\mathcal{B}, b_1, \ldots, b_n)$ and $(\mathcal{C}, c_1, \ldots, c_n)$ 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 continuum many nonisomorphic countable models.

15. Skolem Hulls

Let L be a first order language and T an L-theory. We say that T *has Skolem functions* if for every L-formula $\varphi(x_1, \ldots, x_n, y)$ there is an n-ary function symbol f in L such that

$$T \models \forall x_1 \ldots \forall x_n \big(\exists y \varphi(x_1, \ldots, x_n, y) \to \varphi(x_1, \ldots, x_n, f(x_1, \ldots, x_n))\big).$$

Note that if T has Skolem functions, then so does any extension of T in the same language. In particular, every completion of T has Skolem functions. Note also that L must contain a constant symbol (apply the definition to the formula $\exists y(y = y)$).

15.1. **Proposition** (Skolemization)**.** *Let L be a first order language and T an L-theory. There exists a first order language $L' \supseteq L$ and an L'-theory $T' \supseteq T$ with the following properties:*
(a) T' has Skolem functions;
(b) Every model of T has an expansion to a model of T';
(c) L' 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 of theories $(T_k \mid k \in \mathbb{N})$ with $T_0 = T$, such that T_k an L_k-theory for each $k \geq 1$. To obtain L_{k+1} from L_k we add a new function symbol f_φ for each L_k-formula $\varphi(x_1, \ldots, x_n, y)$; to obtain T_{k+1} from T_k add all sentences of the form

$$\forall x_1 \ldots \forall x_n \big(\exists y \varphi(x_1, \ldots, x_n, y) \to \varphi(x_1, \ldots, x_n, f_\varphi(x_1, \ldots, x_n))\big).$$

Finally, let L' be the union of all the languages L_k and let T' be the union $\cup\{T_k \mid k \in \mathbb{N}\}$.

To prove (a), we note that each formula of L' is an L_k-formula for some k.

To prove (b), we see that each L_k-structure that is a model of T_k can be expanded to an L_{k+1}-structure that is a model of T_{k+1}, using the axiom of choice to interpret each new function symbol appropriately.

To prove (c), we note that in constructing L_{k+1} from L_k, we added one new symbol for each L_k-formula; hence $\text{card}(L_{k+1}) = \text{card}(L_k)$ for each $k \in \mathbb{N}$. It follows that $\text{card}(L') = \text{card}(L_0) = \text{card}(L)$. $\square$

15.2. **Remark.** Any theory T' satisfying the conditions in Proposition 15.1 will be called a *Skolemization of T*.

From condition 15.1(b) it follows that any Skolemization of T is a conservative extension of T. (Proof: Suppose σ is an L-sentence such that $T' \models \sigma$. If $\mathcal{A}$ is any L-structure that is a model of T, then $\mathcal{A}$ has an expansion to an L'-structure $\mathcal{A}'$ that is a model of T'. It follows that $\mathcal{A}' \models \sigma$ and hence also $\mathcal{A} \models \sigma$. Thus $T \models \sigma$.)

15.3. **Proposition.** *Let T be a theory that has Skolem functions. Then T admits quantifier elimination. Moreover, if $\mathcal{A}$ is a model of T and $X \subseteq A$,* *then the substructure of $\mathcal{A}$ generated by X is an elementary substructure of $\mathcal{A}$.*

Proof. Suppose $\varphi(x_1, \dots, x_n, y)$ is any formula in the language of T. Since T has Skolem functions there is a function symbol f such that

$$T \models \forall x_1 \dots \forall x_n(\exists y \varphi(x_1, \dots, x_n, y) \to \varphi(x_1, \dots, x_n, f(x_1, \dots, x_n))).$$

But this implies that $\exists y \varphi(x_1, \dots, x_n, y)$ and $\varphi(x_1, \dots, x_n, f(x_1, \dots, x_n))$ are equivalent in all models of T. Lemma 8.4 yields that T admits QE.

Let $\langle X \rangle_{\mathcal{A}}$ be the substructure of $\mathcal{A}$ generated by X (it is nonempty since L has constants). The Tarski-Vaught test yields that $\langle X \rangle_{\mathcal{A}}$ is an elementary substructure of $\mathcal{A}$. Indeed, suppose $\psi(x_1, \dots, x_n, y)$ is a formula in the language of T and $a_1, \dots, a_n$ are in $\langle X \rangle_{\mathcal{A}}$ and satisfy $\mathcal{A} \models \exists y \psi(x_1, \dots, x_n, y)[a_1, \dots, a_n]$. Since T has Skolem functions, its language has a function symbol f such that $\mathcal{A} \models \psi(x_1, \dots, x_n, y)[a_1, \dots, a_n, f^{\mathcal{A}}(a_1, \dots, a_n)]$. Since $f^{\mathcal{A}}(a_1, \dots, a_n)$ is in $\langle X \rangle_{\mathcal{A}}$, we see that this substructure satisfies the criterion in Theorem 3.8. $\square$

15.4. **Definition.** Let T be a theory that has Skolem functions and let $\mathcal{A}$ be a model of T. For any subset X of A, the substructure of $\mathcal{A}$ generated by X (denoted $\langle X \rangle_{\mathcal{A}}$) is called the *Skolem hull of X in $\mathcal{A}$*.

Note that if T is a theory that has Skolem functions, $\mathcal{A} \models T$, and $X \subseteq A$, then $\langle X \rangle_{\mathcal{A}} = \mathrm{acl}_{\mathcal{A}}(X)$. The containment $\subseteq$ is true for every structure $\mathcal{A}$ and $X \subseteq A$. The opposite containment follows from the fact that $\langle X \rangle_{\mathcal{A}}$ is an elementary substructure of $\mathcal{A}$ and contains X.

15.5. **Proposition.** *Let T be a theory that has Skolem functions and let $\mathcal{A}, \mathcal{B}$ be models of T. Suppose $f\colon X \to Y$ is an elementary map with respect to $\mathcal{A}, \mathcal{B}$ (so $X \subseteq A$ and $Y \subseteq B$). Then f extends to an elementary embedding from the Skolem hull $\langle X \rangle_{\mathcal{A}}$ into $\langle Y \rangle_{\mathcal{B}}$, and such an extension is unique. Moreover, if $f(X) = Y$ then this extension is an isomorphism.*

Proof. The existence of an extension of f to an embedding of $\langle X \rangle_{\mathcal{A}}$ into $\langle Y \rangle_{\mathcal{B}}$ follows from Lemma 3.14. This is an elementary embedding because $\langle X \rangle_{\mathcal{A}}$ and $\langle Y \rangle_{\mathcal{B}}$ are models of T, which admits QE. Uniqueness is immediate from the fact that every element of $\langle X \rangle_{\mathcal{A}}$ is the value of an $L(X)$ term in $\mathcal{A}$. If $f(X) = Y$, then we may reverse the roles of X and Y and apply the same argument to $f^{-1}\colon Y \to X$, and the extensions of f and f^{-1} will be inverses of each other. $\square$

16.1. **Definition.** Let $\mathcal{A}$ be an L-structure, $(I, <)$ a linear ordering, and $(a_i \mid i \in I)$ a family of elements of A. We say $(a_i \mid i \in I)$ is *a family of ordered indiscernibles* (with respect to the ordering $<$ on I and the structure $\mathcal{A}$) if it has the following property: for any L-formula $\varphi(x_1, \ldots, x_n)$ and any two increasing sequences $i_1 < \cdots < i_n$ and $j_1 < \cdots < j_n$ from $(I, <)$, we have the equivalence

$$\mathcal{A} \models \varphi[a_{i_1}, \ldots, a_{i_n}] \quad \Leftrightarrow \mathcal{A} \models \varphi[a_{j_1}, \ldots, a_{j_n}].$$

Note that if $(a_i \mid i \in I)$ is indiscernible in $\mathcal{A}$ 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 $x = y$.)

16.2. **Proposition.** *Let T be a theory with an infinite model. There exists a model $\mathcal{A}$ of T and a nonconstant sequence $(a_k \mid k \in \mathbb{N})$ that is an sequence of ordered indiscernibles in $\mathcal{A}$.*

Proof. Let L be the language of T; let L' be the language obtained from L by adding a family $(c_k \mid k \in \mathbb{N})$ of distinct constants. Let T' be the L'-theory consisting of T together with all the sentences $c_k \neq c_l$ for distinct $k, l \in \mathbb{N}$, and all the sentences

$$\varphi(c_{i_1}, \ldots, c_{i_n}) \leftrightarrow \varphi(c_{j_1}, \ldots, c_{j_n})$$

where $\varphi(x_1, \ldots, x_n)$ is any L-formula and $i_1 < \cdots < i_n$ and $j_1 < \cdots < j_n$ are sequences from $\mathbb{N}$. (We will refer to these last sentences as the "indiscernibility axioms" in T'.)

If $\mathcal{B}$ is a model of T' and $\mathcal{A}$ is the reduct of $\mathcal{B}$ to L, then $(c_k^{\mathcal{B}} \mid k \in \mathbb{N})$ is evidently a nonconstant sequence of ordered indiscernibles in $\mathcal{A}$. Hence it suffices to show that T' has a model, which we do using the Compactness Theorem.

Let $\mathcal{A}$ be any infinite model of T and let $\alpha \colon \mathbb{N} \to A$ be any 1-1 function. Let Σ be any finite subset of T'. Let F be the set of $k \in \mathbb{N}$ such that c_k occurs in some member of Σ. Let $\psi_1, \ldots, \psi_m$ be all the indiscernibility axioms that occur in Σ. We may assume that there exist L-formulas $\varphi_j(x_1, \ldots, x_n)$ such that for each $j = 1, \ldots, m$ the sentence ψ_j is logically equivalent to

$$\varphi_j(c_{i_1}, \ldots, c_{i_n}) \leftrightarrow \varphi_j(c_{j_1}, \ldots, c_{j_n}).$$

for some sequences $i_1 < \cdots < i_n$ and $j_1 < \cdots < j_n$ from F.

We now define a coloring function $C \colon [\mathbb{N}]^n \to \mathcal{P}(\{1, \ldots, m\})$, to which we will apply Ramsey's Theorem. Namely, for each $i_1 < \cdots < i_n$ from $\mathbb{N}$ we take $C(i_1, \ldots, i_n)$ to be the set of all $j \in \{1, \ldots, m\}$ such that

$$\mathcal{A} \models \varphi_j[\alpha(i_1), \ldots, \alpha(i_n)].$$

By Ramsey's Theorem there is an infinite set $H \subseteq \mathbb{N}$ that is homogeneous for C; that is, $C(i_1, \ldots, i_n) = C(j_1, \ldots, j_n)$ whenever $i_1 < \cdots < i_n$ and $j_1 < \cdots < j_n$ are sequences from H.

Let $g \colon F \to H$ be any increasing function. We obtain a model of Σ by using $\mathcal{A}$ to interpret the symbols of L and by interpreting c_k as $\alpha(g(k))$ for each $k \in F$. (The other c_k do not occur in Σ.) This shows that every finite subset of T' 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 $\mathcal{A}$ is a model of T with an infinite subset $S \subset A$ such that $\{(a, b) \in S^2 \mid \mathcal{A} \models \varphi[a, b]\}$ is a linear ordering on S. By taking an elementary extension if necessary, we may assume that there is a function $\alpha\colon \mathbb{N} \to A$ such that $\mathcal{A} \models \varphi[\alpha(k), \alpha(l)]$ for all $k < l$ in $\mathbb{N}$. Using this function in the above proof yields a nonconstant sequence of ordered indiscernibles $(a_k \mid k \in \mathbb{N})$ in a model of T such that $\varphi(a_k, a_l)$ holds for $k, l \in \mathbb{N}$ if and only if $k < l$.

16.4. **Definition.** Let $(a_i \mid i \in I)$ be a sequence of ordered indiscernibles in $\mathcal{A}$, with I infinite, 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 $\mathcal{A}$ is the set of all L-formulas $\varphi(x_1, \ldots, x_n)$ such that $\mathcal{A} \models \varphi[a_{i_1}, \ldots, a_{i_n}]$ for every (equivalently, some) sequence $i_1 < \cdots < i_n$ from $(I, <)$.

16.5. **Proposition.** *Let $(a_i \mid i \in I)$ be a sequence of ordered indiscernibles in $\mathcal{A}$, with I infinite, and let $(J, <)$ be another infinite linear ordering. There exists $\mathcal{B} \equiv \mathcal{A}$ and a sequence of ordered indiscernibles $(b_j \mid j \in J)$ in $\mathcal{B}$ having the same type as $(a_i \mid i \in I)$.*

Proof. An easy application of the Compactness Theorem. $\square$

Ehrenfeucht-Mostowski Models

We combine the construction of indiscernible sequences with the Skolem hulls that were discussed in the previous chapter to produce models that have a large group of automorphisms and models that realize few types.

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'. Take any infinite model $\mathcal{A}$ of T' with a nonconstant sequence of ordered indiscernibles $(a_k \mid k \in \mathbb{N})$ and let Φ be the type of this sequence in $\mathcal{A}$. We will refer to such a Φ as *the type of a nonconstant sequence of ordered indiscernibles 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 $\mathcal{B}$ be a model of T' and $(b_i \mid i \in I)$ a sequence of ordered indiscernibles that has type Φ in $\mathcal{B}$. Take $\Phi(I, <)$ to be the Skolem hull of $\{b_i \mid i \in I\}$ in $\mathcal{B}$. 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. Suppose $(c_i \mid i \in I)$ is another sequence of ordered indiscernibles that has type Φ in another model $\mathcal{C}$. Let $X = \{b_i \mid i \in I\}$ and $Y = \{c_i \mid i \in I\}$ and consider the map $f\colon X \to Y$ defined by $f(b_i) = c_i$ for all $i \in I$. Since these indiscernible sequences have the same type, f is an elementary map with respect to $\mathcal{B}, \mathcal{C}$. Therefore, by Proposition 15.5, f extends to an isomorphism from $\langle X \rangle_{\mathcal{B}}$ onto $\langle Y \rangle_{\mathcal{C}}$.

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 $\mathcal{B}$ and that $(c_i \mid j \in J)$ has type Φ in $\mathcal{C}$. For each order preserving function $F\colon I \to J$ we consider F as a map from $\{b_i \mid i \in I\}$ to $\{c_j \mid j \in J\}$ by taking each b_i to $c_{F(i)}$. Since these indiscernible sequences have the same type, this defines an elementary map with respect to $\mathcal{B}, \mathcal{C}$. Therefore it extends in a unique way to an elementary embedding of the Skolem hulls, by Proposition 15.5. We denote this extension by $\Phi(F)$.

To summarize, our construction yields the following: for each infinite linear ordering $(I,<)$ we have a model $\Phi(I,<)$ of the Skolemization T' of T; this model is generated as an L'-structure by the set I, and $(I,<)$ itself is a sequence of ordered indiscernibles in $\Phi(I,<)$. Moreover, for each order preserving map $F\colon (I,<) \to (J,<)$ 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.

We give two applications of this construction.

16.6. **Corollary.** *Let T be an L-theory with infinite models. For any cardinal κ such that* $\operatorname{card}(L) \leq \kappa$ *there is a model $\mathcal{A}$ of T such that A has cardinality κ and $\mathcal{A}$ has 2^κ automorphisms (which is the maximum possible number).*

Proof. Let Φ be the type of a nonconstant sequence of ordered indiscernibles in a model of a Skolemization T' of T, such that the language of T' has the same cardinality as L. 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 that $(I,<)$ has 2^κ many automorphisms. (For each function $\varphi\colon \kappa \to \{0,1\}$, 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 infinite cardinal κ, there is a model $\mathcal{A}$ of T such that A has cardinality κ but for every countable subset $C \subseteq A$ and every $n \geq 1$, only countably many n-types are realized in $(\mathcal{A}, a)_{a \in C}$.*

Proof. Let Φ be the type of a nonconstant sequence of ordered indiscernibles in a model of a Skolemization T' of T (so the language of T' is countable).

Let $\mathcal{A} = \Phi(\kappa, <)$. Then the reduct of $\mathcal{A}$ to L is a model of T of cardinality κ. We will show it satisfies the condition in this Corollary.

Let C be a countable subset of A. For each $a \in C$ 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 $\mathcal{A}$. Let S be the subset of κ consisting of all ordinals that occur in s_a for some $a \in C$. Since C 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 $f\colon X \to Y$ is order preserving and is the identity on S. By Proposition 15.5, f has a unique extension to an elementary map from $\langle X\rangle_{\mathcal{A}}$ to $\langle Y\rangle_{\mathcal{A}}$, which we denote by $\tilde{f}$. Both of these Skolem hulls are elementary substructures of $\mathcal{A}$. Since f is the identity on S, its extension is the identity on $\langle S\rangle_{\mathcal{A}}$, which contains C. Therefore, for any tuple $a_1, \dots, a_n$ in $\langle X\rangle_{\mathcal{A}}$, the types realized by $(a_1, \dots, a_n)$ and by $(\tilde{f}(a_1), \dots, \tilde{f}(a_n))$ in $(\mathcal{A}, a)_{a\in C}$ are the same.

Suppose $\alpha_1, \dots, \alpha_n$ and $\beta_1, \dots, \beta_n$ are 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 $\langle X\rangle_{\mathcal{A}}$ is countable.

Let $(a_1, \dots, a_n)$ be any n-tuple from A. 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 $\mathcal{A}$. Let $\alpha_1, \dots, \alpha_p$ be the ordinals that occur in the sequences $s_1, \dots, s_n$. Let $\beta_1, \dots, \beta_p$ 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 $\langle X\rangle_{\mathcal{A}}$ 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 $(\mathcal{A}, a)_{a\in C}$.

Since $\langle X\rangle_{\mathcal{A}}$ is countable, this shows that for any countable $C \subseteq A$, only countably many n-types are realized in the $L'(C)$-structure $(\mathcal{A}, a)_{a\in C}$. Hence the same is true if we replace $\mathcal{A}$ by its reduct to L. $\square$

In this chapter T is a complete L-theory. Also, x and y denote finite tuples of variables, $x = x_1, \ldots, x_m$ and $y = y_1, \ldots, y_n$; we will write $\forall x$ instead of $\forall x_1 \ldots \forall x_m$, 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 $(\mathcal{A}, a)_{a \in A}$).

17.1. **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) $\mathcal{B} \models \forall x \big(\varphi_k(x) \to \varphi(x)\big)$ for all $k \in \mathbb{N}$;
(b) $\mathcal{B} \models \forall x \neg\big(\varphi_k(x) \wedge \varphi_l(x)\big)$ 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 $\alpha < \lambda$.

17.2. **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) $S = \{\alpha \mid \alpha$ 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 17.1. 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 $RM(\mathcal{A}, \varphi(x)) \geq \beta + 1$. Finally, if β is a limit ordinal $< \alpha$, we have (from the induction hypothesis and the successor ordinal case just treated) $RM(\mathcal{A}, \varphi(x)) \geq \delta$ for every $\delta + 1 < \beta$. But this implies $RM(\mathcal{A}, \varphi(x)) \geq \beta$, 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 17.1(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* $\varphi(x)$.

17.3. **Definition** (Morley rank)**.** Let $\mathcal{A}$ be a model of T and let $\varphi(x)$ be an $L(A)$-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 $L(A)$-formula $\varphi(x)$ *is ranked.*

17.4. **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 A, then the value of $RM_x(\mathcal{A}, \varphi(x, a))$ depends only on* $\mathrm{tp}_{\mathcal{A}}(a)$.

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 $\mathrm{tp}_{\mathcal{A}}(a)$. 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 A_j, and assume that $\mathrm{tp}_{\mathcal{A}_1}(a_1) = \mathrm{tp}_{\mathcal{A}_2}(a_2)$. 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 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) $\mathcal{B}_1 \models \forall x\big(\varphi_k(x, b_k) \to \varphi(x, a_1)\big)$ for all $k \in \mathbb{N}$;
(b) $\mathcal{B}_1 \models \forall x\, \neg\big(\varphi_k(x, b_k) \wedge \varphi_l(x, b_l)\big)$ 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 $\mathrm{tp}_{\mathcal{B}_1}(a_1) = \mathrm{tp}_{\mathcal{B}_2}(a_2)$. Since $\mathcal{B}_2$ is ω-saturated, we may construct inductively a sequence $(c_k \mid k \in \mathbb{N})$ of finite tuples from B_2 such that for all $k \in \mathbb{N}$

$$\mathrm{tp}_{\mathcal{B}_2}(a_2 c_0 \dots c_k) = \mathrm{tp}_{\mathcal{B}_1}(a_1 b_0 \dots b_k).$$

It follows that

(a) $\mathcal{B}_2 \models \forall x\big(\varphi_k(x, c_k) \to \varphi(x, a_2)\big)$ for all $k \in \mathbb{N}$;
(b) $\mathcal{B}_2 \models \forall x\, \neg\big(\varphi_k(x, c_k) \wedge \varphi_l(x, c_l)\big)$ 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 $RM_x(\mathcal{A}_2, \varphi(x, a_2)) \geq \beta + 1$. $\square$

17.5. **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 $\mathrm{tp}_{\mathcal{A}}(a)$ and the tuple of variables x are understood.

17.6. **Lemma.** *Let $\mathcal{A}$ be an ω-saturated model of T and let $\varphi(x)$ be an $L(A)$-formula. In applying Definition 17.1, 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 17.4. □

17.7. **Lemma** (Properties of Morley rank)**.** *Let $\mathcal{A}$ be a model of $T$and let $\varphi(x), \psi(x)$ be $L(A)$-formulas.*
(1) $RM(\varphi(x)) = 0$ iff the number of tuples $u \in A$ for which $\mathcal{A} \models \varphi(u)$ is finite and > 0;
(2) if $\mathcal{A} \models \forall x\big(\varphi(x) \to \psi(x)\big)$, 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 $L(B)$-formula $\chi(x)$ such that $\mathcal{B} \models \chi(x) \to \varphi(x)$ 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 $\mathcal{A} \models \forall x\big(\varphi(x) \to \psi(x)\big)$ and $RM(\varphi(x)) \geq \alpha$, then $RM(\psi(x)) \geq \alpha$.

(3) Since $\mathcal{A} \models \forall x\big(\varphi(x) \to (\varphi(x) \vee \psi(x))\big)$, 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 $\psi(x) \to \varphi(x)$ 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 $RM(\varphi(x)) \geq \beta$, then one has $RM(\psi(x)) \geq \alpha$. 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. □

17.8. **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, $\kappa \leq 2^{\mathrm{card}(L)}$. Therefore, there exists an ordinal $\alpha_T < (2^{\mathrm{card}(L)})^+$ such that the set of ordinal Morley ranks is exactly the set of ordinals $< \alpha_T$.

17.9. **Lemma** (Morley degree). *Let $\mathcal{A}$ be a model of T and $\varphi(x)$ a ranked $L(A)$-formula. There exists a finite bound on the integers k such that there exists an elementary extension $\mathcal{B}$ of $\mathcal{A}$ and $L(B)$-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) $\mathcal{B} \models \forall x\big(\varphi_j(x) \to \varphi(x)\big)$ for all $j < k$;

(c) $\mathcal{B} \models \forall x\, \neg\big(\varphi_i(x) \wedge \varphi_j(x)\big)$ for all distinct $i, j < k$.

Moreover, the maximum value of k depends only on $\mathrm{tp}_{\mathcal{A}}(a)$. 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 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 $L(B)$-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 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 $a, b_0, \ldots, b_{k-1}$ in $\mathcal{B}$. (For clause (a) we apply Lemma 17.4.) 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 $b_0, \ldots, b_{k-1}$ 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 $L(A)$-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 $(\Lambda, \subseteq)$ 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 $L(A)$-formula φ_σ of Morley rank α. This is done by induction on the binary tree $(\Lambda, \subseteq)$. 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 17.7(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 $f \colon \mathbb{N} \to \{0, 1\}$ such that for all $k \in \mathbb{N}$ the sequence $f|k = f(0), \ldots, f(k-1)$ 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 $d = \mathrm{card}(S_0)$ and let $\chi_0, \ldots, \chi_{d-1}$ enumerate the formulas φ_σ with $\sigma \in S_0$. Our construction has ensured that $(\chi_j | 0 \leq j < d)$ 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_j for $0 \leq j < d$. Suppose now that $(\varphi_j(x) \mid 0 \leq j < k)$ is any sequence of $L(A)$-formulas that satisfy conditions (a),(b),(c) and that $k > d$. Consider any j with $0 \leq j < d$ and distinct r, s with $0 \leq r, s < k$. By our construction, χ_j is φ_σ for some σ that is a leaf in S. Using Lemma 17.7 and the fact that φ_r and φ_s are contradictory in $\mathcal{A}$, it follows that at most one of $\chi_j \wedge \varphi_r$ and $\chi_j \wedge \varphi_s$ can have Morley rank $= \alpha$. Since $d < k$, 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_j for $0 \leq j < d$. Therefore, φ_r is equivalent to the disjunction of the formulas $\chi_j \wedge \varphi_r$ with $0 \leq j < d$. Using Lemma 17.7(3) it follows that φ_r itself has Morley rank $< \alpha$. This contradiction proves the Lemma. $\square$

17.10. **Definition.** Given a ranked $L(A)$-formula $\varphi(x)$, the greatest integer whose existence is proved in Lemma 17.9 is called the *Morley degree* of $\varphi(x)$ and it is denoted $dM(\varphi(x))$.

17.11. **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* $\mathcal{A} \models \varphi(x) \to \psi(x)$*, 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 17.7 and 17.9 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 17.7(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 17.7(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 $(\varphi_0(x), \ldots, \varphi_{k-1}(x), \psi_0(x), \ldots, \psi_{l-1}(x))$ 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 $k = d$. □

17.12. **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* $\mathrm{Th}((\mathcal{A}, a)_{a \in C})$*. 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:*

$p(x)$ consists exactly of the $L(C)$-formulas $\psi(x)$ such that
$RM(\psi(x) \wedge \varphi_p(x)) = RM(\varphi_p(x))$ *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 17.7. A similar argument using Lemma 17.11 proves $dM(\psi(x) \wedge \varphi_p(x)) = dM(\varphi(x))$.

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 17.11 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. □

17.13. **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 $\mathrm{Th}((\mathcal{A}, a)_{a \in C})$. 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 $RM(\varphi(x)) = RM(p(x))$.

17.14. **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 $\mathrm{Th}((\mathcal{A}, a)_{a \in C})$.

17.15. **Theorem.** *Let L be countable. The following conditions are equivalent:*
(1) T is ω-stable;
(2) for any $\mathcal{A} \models T$ and any $L(A)$-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 $L(A)$-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 17.7(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 $L(A)$-formula whose Morley rank is $+\infty$. Then $RM(\varphi(x)) \geq \alpha_T + 1$, so there exist two $L(A)$-formulas $\psi_1(x), \psi_2(x)$ that are contradictory in $\mathcal{A}$ and have Morley rank $\geq \alpha_T$, and such that $\psi_j(x) \to \varphi(x)$ 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 17.7 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 17.9 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 many 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 $L(A)$-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 A that occur in $\varphi_\sigma(x)$ for some $\sigma \in \Lambda$; note that C is countable. For each function $f \colon \mathbb{N} \to \{0,1\}$, 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 $\{p_f(x) \mid f \colon \mathbb{N} \to \{0,1\}\}$ is a family of uncountably many types of finite tuples consistent with $\mathrm{Th}((\mathcal{A}, a)_{a \in C})$. This contradicts (1).

$(2 \Rightarrow 3)$: Let $\mathcal{A} \models T$ and $C \subseteq A$ with $\mathrm{card}(C) \leq \lambda$. 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 $\mathrm{Th}(\mathcal{A}, a)_{a \in C})$. Given such $p(x)$, condition (2) ensures that Lemma 17.12 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 $p(x)$. □

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.

17.16. **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 $\mathrm{tp}_{\mathcal{A}}(b/B)$ for the type (in $L(B)$) realized by b in $(\mathcal{A}, a)_{a \in B}$.

17.17. **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 $p_k(x) = \mathrm{tp}_{\mathcal{A}}(a_k/C \cup \{a_0, \ldots, a_{k-1}\})$. 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 $i_0 < \cdots < i_n$ are in $\mathbb{N}$, $\mathrm{tp}_{\mathcal{A}}(a_{i_0} \ldots a_{i_n}/C) = \mathrm{tp}_{\mathcal{A}}(a_0 \ldots a_n/C)$.

In the basis case, $n = 0$. Take any $i \in \mathbb{N}$. Since $\varphi(x) \in \mathrm{tp}_{\mathcal{A}}(a_i/C)$ we have $(RM(\mathrm{tp}_{\mathcal{A}}(a_i/C)), dM(\mathrm{tp}_{\mathcal{A}}(a_i/C)) \leq (\alpha, d)$ lexicographically. On the other hand, $\mathrm{tp}_{\mathcal{A}}(a_i/C) \subseteq p_i(x)$; thus our assumptions yield $(RM(\mathrm{tp}_{\mathcal{A}}(a_i/C)), dM(\mathrm{tp}_{\mathcal{A}}(a_i/C)) \geq (\alpha, d)$. It follows that $(RM(\mathrm{tp}_{\mathcal{A}}(a_i/C)), dM(\mathrm{tp}_{\mathcal{A}}(a_i/C)) = (\alpha, d)$. Lemma 17.12 implies $\mathrm{tp}_{\mathcal{A}}(a_i/C)) = \mathrm{tp}_{\mathcal{A}}(a_0/C))$.

For the induction step $n > 0$. Consider any $i_0 < \cdots < i_n$ from $\mathbb{N}$. As argued in the previous paragraph, $\mathrm{tp}(a_{i_n}/C \cup \{a_{i_0}, \ldots, a_{i_{n-1}}\})$ has Morley rank α and degree d, as does $p_n(x)$. Both of these types contain the formula $\varphi(x)$. Applying Lemma 17.12 we conclude that for any $L(C)$-formula $\psi(x, y_0, \ldots, y_{n-1})$ we have:
(i) $\mathcal{A} \models \psi[a_n, a_0, \ldots, a_{n-1}]$ if and only if the formula $\varphi(x) \wedge \psi(x, a_0, \ldots, a_{n-1})$ has Morley rank α and degree d;
(ii) $\mathcal{A} \models \psi[a_{i_n}, a_{i_0}, \ldots, a_{i_{n-1}}]$ if and only if the formula $\varphi(x) \wedge \psi(x, a_{i_0}, \ldots, a_{i_{n-1}})$ has Morley rank α and degree d.

The induction hypothesis states that $\mathrm{tp}_{\mathcal{A}}(a_{i_0} \ldots a_{i_{n-1}}/C) = \mathrm{tp}_{\mathcal{A}}(a_0 \ldots a_{n-1}/C)$. This implies that the right hand sides of statements (i) and (ii) are equivalent to each other. Therefore we conclude $\mathrm{tp}_{\mathcal{A}}(a_{i_0} \ldots a_{i_n}/C) = \mathrm{tp}_{\mathcal{A}}(a_0 \ldots a_n/C)$ as claimed. $\square$

17.18. **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* $\mathrm{card}(C) < \mathrm{card}(A)$. *Then there exists a nonconstant sequence of ordered indiscernibles in $(\mathcal{A}, a)_{a \in C}$.*

Proof. We may assume C is infinite. Let $\lambda = \mathrm{card}(C)$. 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 $i \in \mathbb{N}$, the Morley rank and degree of $\mathrm{tp}_{\mathcal{A}}(a_k/C \cup \{a_0, \ldots, a_{k-1}\})$ is exactly (α, d).

First we obtain a_0 with this property. If no such element of A exists, we have $(RM(\mathrm{tp}_{\mathcal{A}}(a/C)), dM(\mathrm{tp}_{\mathcal{A}}(a/C))) < (\alpha, d)$ 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 $a_0, \ldots, a_{k-1}$ and seek a_k. This is handled by the same argument as in the previous paragraph, replacing C by $C \cup \{a_0, \ldots, a_{k-1}\}$.

Finally, by Lemma 17.17 the resulting sequence $(a_k \mid k \in \mathbb{N})$ is indiscernible over C in $\mathcal{A}$. $\square$

17.19. 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 $n \in \mathbb{N}$.

17.20. Let $\mathcal{A}$ be an L-structure, let $\varphi(x)$ be an $L(A)$-formula and let $t(x)$ an L-term, with $x = x_1, \ldots, x_m$. 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.)

17.21. 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 $G \supseteq H_0 \supseteq H_1 \supseteq \ldots$ 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.)

17.22. Let L be the language whose nonlogical symbols are the unary predicate symbols $P_1, \ldots, P_n$. Let T be the L-theory whose axioms express that the sets $P_1, \ldots, P_n$ 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.

17.23. 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.

17.24. 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 $\varphi(x_1, \ldots, x_n)$, one has $RM(\varphi(x_1, \ldots, x_m)) < +\infty$. (Hint: use the preceding exercise together with a careful reading of the proof of Theorem 17.15.)

The goal 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.

In this chapter L is a countable first order language and T is a complete L-theory with infinite models.

18.1. **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. It's use is justified by the following result.

18.2. **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 $\text{card}(A) = \kappa$. 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 17.15. □

18.3. **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 17.15, 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 $\mathrm{Th}((\mathcal{A}_\alpha, a)_{a \in A_\alpha})$ is realized in $\mathcal{A}_{\alpha+1}$. Let $\mathcal{A}$ be the union of this chain, so $\mathcal{A} \models T$ and $\mathrm{card}(A) = \kappa$. 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 $\mathrm{Th}((\mathcal{A}, a)_{a \in C}$ 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. □

18.4. **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 18.2, 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. □

18.5. **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.

18.6. **Definition.** Suppose $\mathcal{A}$ is a model of T and $C \subseteq A$. Let a be a finite tuple from A and take $p(x) = \mathrm{tp}_{\mathcal{A}}(a/C)$ 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 $\mathrm{Th}((\mathcal{A}, c)_{c \in C})$. (Here $p(x)$ is a complete type in $L(C)$ and $\mathrm{Th}((\mathcal{A}, c)_{c \in C})$ 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 $\varphi(x, d) \to \psi(x)$ 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.

18.7. **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 $\mathrm{tp}_{\mathcal{A}}(a/C)$ is principal and every coordinate of b is either a coordinate of a or a member of C, then $\mathrm{tp}_{\mathcal{A}}(b/C)$ is principal;
(2) $\mathrm{tp}_{\mathcal{A}}(ab/C)$ is principal if and only if $\mathrm{tp}_{\mathcal{A}}(a/C)$ and $\mathrm{tp}_{\mathcal{A}}(b/C \cup \{a\})$ are principal;
(3) $\mathrm{tp}_{\mathcal{A}}(a_0 \dots a_n/C)$ is principal if and only if $\mathrm{tp}_{\mathcal{A}}(a_j/C \cup \{a_0, \dots, a_{j-1}\})$ is principal for each $0 \le j \le n$.

Proof. (1) Suppose $\mathrm{tp}_{\mathcal{A}}(a/C)$ 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 $\mathrm{tp}_{\mathcal{A}}(a/C)$.

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

$$\exists x_1 \dots \exists x_m \big(\varphi(x_1, \dots, x_m) \wedge y_1 = x_{\pi(1)} \wedge \dots \wedge y_n = x_{\pi(n)}\big)$$

is a complete formula in $\mathrm{tp}_{\mathcal{A}}(b/C)$.

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 $\mathrm{tp}_{\mathcal{A}}(ac/C)$ 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 $\mathrm{tp}_{\mathcal{A}}(ab/C)$ contains the complete formula $\varphi(x, y)$. Then $\exists y \varphi(x, y)$ is a complete formula in $\mathrm{tp}_{\mathcal{A}}(a/C)$ and $\varphi(a, y)$ is a complete formula in $\mathrm{tp}_{\mathcal{A}}(b/C \cup \{a\})$. Conversely, suppose $\varphi(x)$ is a complete formula in $\mathrm{tp}_{\mathcal{A}}(a/C)$ and $\psi(x, y)$ is an $L(C)$-formula such that $\psi(a, y)$ is a complete formula in $\mathrm{tp}_{\mathcal{A}}(b/C \cup \{a\})$. Then $\psi(x, y) \wedge \psi(y)$ is a complete formula in $\mathrm{tp}_{\mathcal{A}}(ab/C)$.

(3) This is proved by induction on n using part (2). $\square$

18.8. **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 $\mathrm{tp}_{\mathcal{A}}(a_\beta / C \cup \{a_\alpha \mid \alpha < \beta\})$ is principal for all $\beta < \gamma$.

18.9. **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 18.8 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.)

18.10. **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, $\mathrm{tp}_{\mathcal{A}}(a/C)$ is principal for each finite tuple a from A.*

Proof. Let $(a_\alpha \mid \alpha < \gamma)$ satisfy the conditions in Definition 18.8 and the preceding remark. That is,
$A = C \cup \{a_\alpha \mid \alpha < \gamma\}$;
$\mathrm{tp}_{\mathcal{A}}(a_\beta / C \cup \{a_\alpha \mid \alpha < \beta\})$ 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 $\mathrm{tp}_{\mathcal{A}}(a_{\beta_j} / C \cup C_{\beta_j})$ is principal over $C \cup \{\alpha_{\beta_1}, \dots, \alpha_{\beta_{j-1}}\}$ for each $j = 1, \dots, n$. Lemma 18.7 implies that $\mathrm{tp}_{\mathcal{A}}(b/C)$ 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 $\mathrm{tp}_\mathcal{A}(a_\beta / C \cup C_\beta)$ 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 $\mathrm{tp}_\mathcal{A}(b/C)$ is principal. By Lemma 18.7 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, $\mathrm{tp}_\mathcal{A}(b'/C)$ is principal. Hence Lemma 18.7 yields that $\mathrm{tp}_\mathcal{A}(b/C)$ is also principal, as desired. □

18.11. **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 universe 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). Chose 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 $\mathrm{Th}((\mathcal{A}, a)_{a \in C})$. Otherwise there is an $L(C)$-formula $\psi(x)$ such that $\varphi(x) \wedge \psi(x)$ and $\varphi \wedge \neg\psi(x)$ are both consistent with $\mathrm{Th}((\mathcal{A}, a)_{a \in C})$. 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, $\mathrm{tp}_\mathcal{A}(a_0/C)$ is principal and $a_0 \notin C$. Continue inductively as long as possible to construct a sequence of distinct elements a_α in $A \backslash C$ for an initial segment of ordinals α such that whenever a_α is defined, we have that $\mathrm{tp}_\mathcal{A}(a_\alpha / C \cup \{a_\delta \mid \delta < \alpha\})$ 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 universe of an elementary substructure of $\mathcal{A}$ that is constructible over A. □

Next we prove the main technical result of this chapter, from which Morley's Theorem is an easy consequence.

18.12. **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 $\lambda = \mathrm{card}(A)$ 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 $\mathrm{Th}((\mathcal{A}, a)_{a \in C})$ but is not realized in $(\mathcal{A}, a)_{a \in C}$.

By Proposition 17.18 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 $\mathrm{Th}((\mathcal{A}, a)_{a \in C'})$ 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 18.11 there is $\mathcal{B}' \preceq \mathcal{B}$ such that $C' \cup \{b_\alpha \mid \alpha < \kappa\} \subseteq B'$ 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 18.10, we have that $\mathrm{tp}_{\mathcal{B}'}(b/C' \cup \{b_\alpha \mid \alpha < \kappa\}$ 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 $(\mathcal{A}, a_0, \dots, a_0, a)_{a \in C'}$ 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 $C' \subseteq B''$ and $\mathrm{card}(B'') = \kappa$. 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 18.1 (Morley's Theorem). Suppose κ is an uncountable cardinal and T is κ-categorical. By Proposition 18.2, T is ω-stable. By Corollary 18.4, every model of T of cardinality κ is κ-saturated. Let λ be any uncountable cardinal. By Theorem 18.12, every model of T of cardinality λ is λ-saturated. Using Theorem 12.9(d) 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. 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.

Our first results give necessary and sufficient conditions for a relation to be definable in a given structure:

19.1. **Theorem.** *Let L be any first order language and let P be an n-ary predicate symbol that is not in L. Suppose $(\mathcal{A}, R)$ is an ω-saturated $L(P)$-structure and that $\mathcal{A}$ is strongly ω-homogeneous. The following conditions are equivalent:*
(1) R is 0-definable in $\mathcal{A}$;
(2) every automorphism of $\mathcal{A}$ leaves R setwise invariant.

Proof. (1) $\Rightarrow$ (2) is immediate, since automorphisms are elementary maps.

(2) $\Rightarrow$ (1): Assume every automorphism of $\mathcal{A}$ leaves R setwise invariant. By way of contradiction, assume that R is not 0-definable in $\mathcal{A}$. It suffices to find n-tuples a and b from A that realize the same n-type in $\mathcal{A}$ but such that $R(a)$ is true and $R(b)$ is false. We would then have an automorphism of $\mathcal{A}$ taking a to b, since $\mathcal{A}$ is strongly ω-homogeneous; this automorphism would not leave R invariant, contradicting our hypothesis.

Since $(\mathcal{A}, R)$ is ω-saturated, it suffices to show there exists a in R so that $\mathrm{tp}_{\mathcal{A}}(a) \cup \{\neg P(v)\}$ is consistent with $\mathrm{Th}(\mathcal{A}, R)$. If b realizes this partial type in $(\mathcal{A}, R)$, then a and b have the desired properties.

Let $\Sigma(v) = \{\varphi(v) \in L \mid (\mathcal{A}, R) \models \forall v(\neg P(v) \rightarrow \varphi(v))\}$. If a realizes the partial type $\Sigma(v) \cup \{P(v)\}$ in $(\mathcal{A}, R)$, then obviously $\mathrm{tp}_{\mathcal{A}}(a) \cup \{\neg P(v)\}$ is consistent with $\mathrm{Th}(\mathcal{A}, R)$. Therefore, since $(\mathcal{A}, R)$ is ω-saturated it suffices to prove that $\Sigma(v) \cup \{P(v)\}$ is consistent with $\mathrm{Th}(\mathcal{A}, R)$.

Arguing by contradiction, suppose $\Sigma(v) \cup \{P(v)\}$ is not consistent with $\mathrm{Th}(\mathcal{A}, R)$. Then there is a formula $\varphi(v)$ in $\Sigma(v)$ such that $(\mathcal{A}, R) \models \varphi(v) \rightarrow \neg P(v)$. Thus $(\mathcal{A}, R) \models \forall v(\neg P(v) \leftrightarrow \varphi(v))$ and hence $\neg\varphi(v)$ defines R in $\mathcal{A}$. This contradicts the assumption that R is not definable in $\mathcal{A}$. $\square$

19.2. **Corollary** (Svenonius's Theorem)**.** *Let $\mathcal{A}$ be any L-structure and let R be any n-ary relation on A. The following conditions are equivalent:*
(1) R is not 0-definable in $\mathcal{A}$;
(2) there is an elementary extension $(\mathcal{B}, S)$ of $(\mathcal{A}, R)$ and an automorphism of $\mathcal{B}$ that does not leave S setwise invariant.

Proof. (2) $\Rightarrow$ (1): If R were 0-definable in $\mathcal{A}$, then S would be 0-definable in $\mathcal{A}$, by the same L-formula, and thus S would be invariant under every automorphism of $\mathcal{B}$.

(1) $\Rightarrow$ (2): By Theorem 12.3 we may take $(\mathcal{B}, S)$ to be an elementary extension of $(\mathcal{A}, R)$ such that $(\mathcal{B}, S)$ is ω-saturated and $\mathcal{B}$ is strongly ω-homogeneous. Note that S is definable in $\mathcal{B}$ if and only if R is definable in $\mathcal{A}$. Now apply Theorem 19.1 to $(\mathcal{B}, S)$. $\square$

19.3. **Definition.** Let T be a satisfiable theory in a language that contains $L(P)$ where P is an n-ary predicate symbol.
(a) We say T *defines* P *explicitly* over L if there is a formula $\varphi(x_1, \ldots, x_n)$ in L such that $T \models \forall x_1 \ldots \forall x_n (P(x) \leftrightarrow \varphi(x))$.
(b) We say T *defines* P *implicitly* over L if for any models $\mathcal{A}$ and $\mathcal{B}$ of T that have the same reduct to L, $P^{\mathcal{A}}$ and $P^{\mathcal{B}}$ are identical.

19.4. **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 defines P explicitly over L if and only if T defines P implicitly over L.*

Proof. ($\Rightarrow$) Obvious.

($\Leftarrow$) Assume that T defines P implicitly over L but that T does not define P explicitly. Let x be the sequence $x_1, \ldots, x_n$ of distinct variables.

Consider the L'-theory T' consisting of T together with all sentences of the form $\neg\forall x\big(P(x) \leftrightarrow \varphi(x)\big)$ where $\varphi(x)$ is an L-formula. We claim that T' is unsatisfiable. Note that if $(\mathcal{A}, R)$ is the reduct to $L(P)$ of a model of T', then R cannot be 0-definable in $\mathcal{A}$.

If T' is satisfiable, use Theorem 12.3 to get an ω-saturated model $\mathcal{A}'$ of T' such that every reduct of $\mathcal{A}'$ to a sublanguage of $L(P)$ is strongly ω-homogeneous. Let $\mathcal{A}$ be the reduct of $\mathcal{A}'$ to L and $R = P^{\mathcal{A}'}$. Since R is not 0-definable in $\mathcal{A}$, by Theorem 19.1 there is an automorphism σ of $\mathcal{A}$ such that $\sigma R \neq R$. Let $\mathcal{B}'$ be the unique L'-structure with underlying set A that is determined by requiring that the function σ is an isomorphism from $\mathcal{A}'$ onto $\mathcal{B}'$. Then $\mathcal{B}'$ is a model of T, the reduct of $\mathcal{B}'$ to L is $\mathcal{A}$ (since σ is an automorphism of $\mathcal{A}$), and $P^{\mathcal{B}'} = \sigma(R) \neq R = P^{\mathcal{A}'}$. This contradicts the assumption that T defines P implicitly over L.

Therefore T' is unsatisfiable. Consequently, there exist L-formulas $\varphi_1(x), \ldots, \varphi_k(x)$ such that T together with the sentences $\neg\forall x\big(P(x) \leftrightarrow \varphi_j(x)\big)$ $(j = 1, \ldots, k)$ is unsatisfiable. Therefore

$$T \models \forall x\big(P(x) \leftrightarrow \varphi_1(x)\big) \vee \cdots \vee \forall x\big(P(x) \leftrightarrow \varphi_k(x)\big).$$

Let T_0 be the $L(P)$-theory consisting of all $L(P)$-sentences that are provable in T. In particular the sentence

$$\forall x\big(P(x) \leftrightarrow \varphi_1(x)\big) \vee \cdots \vee \forall x\big(P(x) \leftrightarrow \varphi_k(x)\big)$$

is in T_0. Therefore we see that T_0 defines P implicitly over L.

For each $j = 1, \ldots, k$ 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. Finally, let $T_j = \{\psi^{(j)} \mid \psi$ is an $L(P)$-sentence and $T_0 \models \psi\}$. Note that if $\mathcal{A}$ is any model of T_j and we set $R = \{a \in A^n \mid \mathcal{A} \models \varphi_j[a]\}$, then $(\mathcal{A}, R) \models T_0$.

Fix $j \in \{1, \ldots, k\}$. Note that the $L(P)$-theories $T_0 \cup T_j$ and $T_0 \cup \{\forall x\big(P(x) \leftrightarrow \varphi_k(x)\big)\}$ are equivalent; indeed, their models are exactly the $L(P)$-structures $(\mathcal{A}, R)$ where $\mathcal{A} \models T_j$ and $R = \{a \in A^n \mid \mathcal{A} \models \varphi_j[a]\}$. Therefore there is an $L(P)$-sentence ψ_j such that $T_0 \models \psi_j$ and

$$T_0 \models \psi_j^{(j)} \leftrightarrow \forall x\big(P(x) \leftrightarrow \varphi_j(x)\big).$$

Therefore

$$T_0 \models \forall x\big(P(x) \leftrightarrow \bigwedge_{j=1}^{k} (\psi_j^{(j)} \to \varphi_j(x))\big).$$

But this sentence is in $L(P)$ and hence we have

$$T \models \forall x\big(P(x) \leftrightarrow \bigwedge_{j=1}^{k} (\psi_j^{(j)} \to \varphi_j(x))\big).$$

showing that T explicitly defines P over L. □

19.5. **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 *defines* F *explicitly over* L if there is an L-formula $\varphi(x_1, \ldots, x_n, y)$ such that

$$T \models \forall x_1 \ldots \forall x_n \forall y(F(x_1, \ldots, x_n) = y \leftrightarrow \varphi(x_1, \ldots, x_n, y)).$$

Further, T *defines* F *implicitly over* L if for any models $\mathcal{A}, \mathcal{B}$ of T that have the same reduct to L, $F^{\mathcal{A}}$ and $F^{\mathcal{B}}$ are identical. Beth's Definability Theorem then holds also for functions: T defines F explicitly over L if and only if T defines F implicitly over L.

19.6. **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.

19.7. **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 $P(v_1, \ldots, v_n)$ or $f(v_1, \ldots, v_n) = w$ or $c = w$, where $v_1, \ldots, v_n$ and w are distinct variables.)

19.8. **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 $\mathcal{A}$ 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 $\mathcal{A}$ 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 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 $\mathcal{A}$ is any model of T, then it has a (unique) expansion $\mathcal{A}'$ that is a model of T'. Since σ is true in $\mathcal{A}'$, and σ is an L-sentence, it must be true in $\mathcal{A}$. 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 19.7 and Beth's Definability Theorem (using Fact 19.5 in the case of function symbols and constant symbols). □

19.9. **Remark.** See pages 57-61 of Shoenfield, *Mathematical Logic*, for a useful discussion of extensions by definitions.

19.10. **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.

19.11. **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 prove Craig's Interpolation Theorem, which gives another characteristic important property of first order logic.

19.12. **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:

19.13. **Lemma.** *Let L be a first order language and let L_1 and L_2 be extensions of L whose intersection is L. Suppose $\mathcal{A}_j$ is an L_j structure for $j = 1, 2$, and $\mathcal{A}_1 \mid L \equiv \mathcal{A}_2 \mid L$. Then there exists an elementary extension $\mathcal{A}_1'$ of $\mathcal{A}_1$ and $f : \mathcal{A}_2 \to \mathcal{A}_1'$ such that f is an elementary embedding of $\mathcal{A}_2 \mid L$ into $\mathcal{A}_1' \mid L$.*

Proof. For any L-structure $\mathcal{A}$ we let $\mathrm{EDiag}(\mathcal{A})$ denote the first order theory of the $L(A)$-structure $(\mathcal{A}, a)_{a \in A}$. It is easy to see that $\mathcal{A}$ can be elementarily embedded into an L-structure $\mathcal{B}$ iff $\mathcal{B}$ has an expansion that is a model of $\mathrm{EDiag}(\mathcal{A})$.

We first show that to prove the Lemma it suffices to prove that $\Sigma = \mathrm{EDiag}(\mathcal{A}_1) \cup \mathrm{EDiag}(\mathcal{A}_2|L)$ is satisfiable. If so, let $\mathcal{A}_1'$ be the reduct of a model of Σ to L_1. Without loss of generality we may assume that for each $a \in A_1$, the interpretation of a in $\mathcal{A}_1'$ is a itself. This implies that $\mathcal{A}_1'$ is an elementary extension of $\mathcal{A}_1$. Moreover, there is an elementary embedding of $\mathcal{A}_2|L$ into $\mathcal{A}_1'|L$ because $\mathcal{A}_1'$ has an expansion that is a model of $\mathrm{EDiag}(\mathcal{A}_2|L)$.

Arguing by contradiction, suppose $\mathrm{EDiag}(\mathcal{A}_1) \cup \mathrm{EDiag}(\mathcal{A}_2|L)$ is not satisfiable. By the Compactness Theorem and the fact that $\mathrm{EDiag}(\mathcal{A}_1)$ and $\mathrm{EDiag}(\mathcal{A}_2|L)$ are closed under conjunction, there exist $\sigma_1 \in \mathrm{EDiag}(\mathcal{A}_1)$ and $\sigma_2 \in \mathrm{EDiag}(\mathcal{A}_2|L)$ such that $\{\sigma_1, \sigma_2\}$ has no model.

There exist $a_1, \ldots, a_m \in A_1, b_1, \ldots, b_n \in A_2$, and L-formulas τ_1, τ_2, such that $\sigma_1 = \tau_1(a_1, \ldots, a_m)$ and $\sigma_2 = \tau_2(b_1, \ldots, b_m)$ We may assume that $A_1 \cap A_2$ is empty, and also that τ_1 and τ_2 have no free variables in common. Let the free variables in τ_1 and τ_2 be $z_1, \ldots, z_{m+n}$. Since $\{\sigma_1, \sigma_2\}$ has no model, it follows that $\exists z \tau_1 \wedge \exists z \tau_2$ has no model. But $\exists z \tau_2$ is true in $\mathcal{A}_2$, because $\sigma_2 \in \mathrm{EDiag}(\mathcal{A}_2|L)$, and $\mathcal{A}_1|L \equiv \mathcal{A}_2|L$, so $\exists z \tau_2$ is true in $\mathcal{A}_1$ as well. Moreover $\exists z \tau_1$ is true in $\mathcal{A}_1$, because $\sigma_1 \in \mathrm{EDiag}(\mathcal{A}_1)$. Therefore $\mathcal{A}_1$ is a model of $\exists z \tau_1 \wedge \exists z \tau_2$. This is a contradiction. □

Proof of Theorem 19.12. 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 $\mathcal{A}_1$ be a model of T_1,and let $\mathcal{B}_1$ be a model of T_2. Now $T = \mathrm{Th}(\mathcal{A}_1|L) = \mathrm{Th}(\mathcal{B}_1|L)$ since T is complete; therefore $\mathcal{A}_1|L \equiv \mathcal{B}_1|L$. By the preceding Lemma there is a model $\mathcal{B}_2 \succ \mathcal{B}_1$ and a map $f_1 : A_1 \to B_2$ that elementarily embeds $\mathcal{A}_1|L$ into $\mathcal{B}_2|L$. Next we apply the Lemma to $(\mathcal{A}_1, a)_{a \in A_1}$ in the language $L_1(A_1)$ and $(\mathcal{B}_2, f_1(a))_{a \in A_1}$ in the language $L_2(A_1)$. We then have

$$(\mathcal{A}_1|L, a)_{a \in A_1} \equiv (\mathcal{B}_2|L, f_1(a))_{a \in A_1},$$

and these two structures are the reducts to $L(A_1)$ of $(\mathcal{A}_1, a)_{a \in A_1}$ and $(\mathcal{B}_2, f_1(a))_{a \in A_1}$ (respectively). Using the lemma again, we see that there exists an elementary extension $(\mathcal{A}_2, a)_{a \in A_1}$ of $(\mathcal{A}_1, a)_{a \in A_1}$ and an elementary embedding g_1 of $(\mathcal{B}_2|L, f_1(a))_{a \in A_1}$ into $(\mathcal{A}_2|L, a)_{a \in A_1}$. Note that these last two structures are reducts to the language $L(A_1)$, which is the intersection of $L_1(A_1)$ and $L_2(A_1)$. So we have $\mathcal{B}_1 \preceq \mathcal{B}_2, \mathcal{A}_1 \preceq \mathcal{A}_2$ and maps $f_1 : \mathcal{A}_1 \to \mathcal{B}_2$ and $g_1 : \mathcal{B}_2 \to \mathcal{A}_2$ 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 $a \in A_1$.

We continue inductively in this way. The result is a pair of elementary chains $\mathcal{A}_1 \preceq \mathcal{A}_2 \preceq \mathcal{A}_3 \preceq \dots$ and $\mathcal{B}_1 \preceq \mathcal{B}_2 \preceq \mathcal{B}_3 \preceq \dots$ and mappings $f_n : \mathcal{A}_n \to \mathcal{B}_{n+1}$ and $g_n : \mathcal{B}_{n+1} \to \mathcal{A}_{n+1}$ that are elementary embeddings with respect to formulas of L and that satisfy $g_n(f_n(x)) = x$ for all $x \in A_n$ and $f_{n+1}(g_n(y)) = y$ for all $y \in B_{n+1}$. Note that for all $n \geq 1, f_{n+1} = f_n$ on A_n and $g_{n+1} = g_n$ on B_{n+1}.

Now let $\mathcal{A} = \cup \mathcal{A}_n, \mathcal{B} = \cup \mathcal{B}_n, f = \cup f_n, g = \cup g_n$. We have $\mathcal{A} \models T_1$ since $\mathcal{A}_1 \models T_1$ and $\mathcal{A}_1 \preceq \mathcal{A}$; similarly $\mathcal{B} \models T_2$. Moreover we see that f is an isomorphism of $\mathcal{A}|L$ onto $\mathcal{B}|L$ whose inverse is g. We can replace $\mathcal{B}$ by an isomorphic copy $\mathcal{B}'$ such that $\mathcal{B}'|L = \mathcal{A}|L$, 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 $\mathcal{C}|L_1 = \mathcal{A}$ and $\mathcal{C}|L_2 = \mathcal{B}'$. The fact that $\mathcal{B}'|L = \mathcal{A}|L$ 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 $T = T_1 \cup T_2$: $\mathcal{C}|L_1 = \mathcal{A} \models T_1$ and $\mathcal{C}|L_2 = \mathcal{B}' \models T_2$, completing the proof. □

19.14. **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 19.12 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. □

19.15. **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.

Finally, we give an alternate proof of Beth's Definability Theorem that uses Craig's Theorem:

Assume that T 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 occurence of any other symbol 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 $(\mathcal{A}, R, R')$ denote its reduct to $L(P, P')$. Then $(\mathcal{A}, R)$ and $(\mathcal{A}, R')$ are both reducts of models of T. It follows from our hypothesis (that T implicitly defines P over L) that $R = R'$. Therefore there exist finite subsets $\Sigma \subseteq T$ and $\Sigma' \subseteq T'$ such that $\Sigma \cup \Sigma' \models \forall\bar{x}(P(\bar{x}) \leftrightarrow P'(\bar{x}))$. 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 $c_1, ...c_n$ to the language of $T \cup T'$. Evidently $\Sigma \cup \Sigma' \models P(c_1, ...c_n) \rightarrow P'(c_1, ...c_n)$. Let σ be the conjunction of all the sentences in Σ and let σ' be the conjunction of the sentences in Σ'. Then $\sigma \wedge P(c_1, ...c_n) \models (\sigma' \rightarrow P'(c_1, ...c_n))$. Obviously the common language of the sentences $\sigma \wedge P(c_1, ...c_n)$ and $(\sigma' \rightarrow P'(c_1, ...c_n))$ is $L(c_1, ...c_n)$. Now apply Craig's Theorem to the above: there is an L-formula $\theta(x_1, ...x_n)$ such that $\sigma \wedge P(c_1, ...c_n) \models \theta(c_1, ...c_n)$ and $\theta(c_1, ...c_n) \models (\sigma' \rightarrow P'(c_1, ...c_n))$. That is to say:

(a) $\Sigma \models (P(c_1, ...c_n) \rightarrow \theta(c_1, ...c_n))$ and

(b) $\Sigma' \models (\theta(c_1, ...c_n) \rightarrow P'(c_1, ...c_n))$.

From (b) we conclude that $\Sigma \models (\theta(c_1, ...c_n) \rightarrow P(c_1, ...c_n))$. (To see this, consider a formal derivation of $(\theta(c_1, ...c_n) \rightarrow P'(c_1, ...c_n))$ 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 $(\theta(c_1, ...c_n) \rightarrow P(c_1, ...c_n))$ from Σ.) Therefore $\Sigma \models (\theta(c_1, ...c_n) \leftrightarrow P(c_1, ...c_n))$. But Σ does not in fact contain the new constants c_i, so we conclude $\Sigma \models \forall \bar{x}(P(\bar{x}) \leftrightarrow \theta(\bar{x}))$. This shows that T explicitly defines P, since Σ is a subset of T.

Math 571 takes 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 appendix we explore the collections of definable sets and characterize them using simple "geometric" properties. This discussion can be read at any time in the course.

Let L be a first order language and $\mathcal{M}$ an L-structure; let M be the underlying set of $\mathcal{M}$.

20.16. **Definition.** A set $A \subseteq M^m$ is *definable in* $\mathcal{M}$ if there is an L-formula $\varphi(x_1, \ldots, x_m, y_1, \ldots, y_n)$ and elements $b_1, \ldots, b_n$ of M such that

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

If $S \subseteq M$ and the above equation holds for some φ and some $b_1, \ldots, b_n \in S$, then we say that A is *S-definable in* $\mathcal{M}$. Let $A \subseteq M^m$ and $B \subseteq M^n$; a function $f\colon A \to B$ is *S-definable in* $\mathcal{M}$ if the graph of f is S-definable in $\mathcal{M}$. (We regard the graph of f as a subset of M^{m+n}.)

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 familar 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 A, B (respectively), and let $\varphi(x, y)$ and $\psi(x, y)$ denote conditions on (x, y) defining subsets Φ and Ψ (respectively) of $A \times B$. 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 $A \times B$,
$\varphi(x, y) \vee \psi(x, y)$	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 $f\colon A \to B$. The image $f(A)$ of A under f can be defined by the equivalence

$$y \in f(A) \Longleftrightarrow \exists x[f(x) = y].$$

Let Γ be the graph of f, which is defined as a subset of $A \times B$ by the condition $f(x) = y$. The displayed equivalence exhibits the fact that $f(A)$ 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:

$\to$	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, $\varphi \to \psi$ 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

$$\forall y\varphi(x,y) \Longleftrightarrow \neg\exists y\neg\varphi(x,y)$$

shows that the set defined by $\forall y\varphi(x,y)$ can be obtained from Φ by first taking the complement in $A \times B$, then projecting onto the first coordinate, and then taking the complement of that set in A. 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 $A \times B$ is sometimes viewed as defining a condition on triples (x,y,z), where z ranges over a nonempty set C; in that case $\varphi(x,y)$ defines a subset of $A \times B \times C$. 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 $f\colon A \to B$ be a function and let Γ be the graph of f. The condition $\varphi(x, f(x))$ defines a subset S of A. This condition is equivalent to

$$\exists y[f(x) = y \wedge \varphi(x,y)].$$

Therefore S can be obtained by applying the projection π' to $\Gamma \cap \Phi$, where $\pi'(x,y) = x$ is the projection of $A \times B$ onto the first coordinate.

We will show that the essential features of the collection of S-definable sets in $\mathcal{M}$ are captured by the following definition:

20.17. **Definition.** Let X be a nonempty set. A *definability system on X* 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 X^m that contains $\emptyset$ and X^m as elements;

(2) if $A \in \mathcal{S}_m$, then $X \times A$ and $A \times X$ belong to $\mathcal{S}_{m+1}$;

(3) $\{(a_1, \ldots, a_m) \in X^m \mid a_1 = a_m\} \in \mathcal{S}_m$;

(4) if $A \in \mathcal{S}_{m+1}$, then $\pi(A) \in \mathcal{S}_m$, where $\pi\colon X^{m+1} \to X^m$ is the projection map on the first m coordinates.

If $A \subseteq X^m$ we say *A belongs to $\mathcal{S}$* if $A \in \mathcal{S}_m$. If $A \subseteq X^m$ and $B \subseteq X^n$ and if $f\colon A \to B$ is a function, then we say *f belongs to $\mathcal{S}$* if the graph of f belongs to $\mathcal{S}$.

20.18. **Notation.** Let L be a first order language and $\mathcal{M}$ an L-structure with underlying set M; let S be any subset of M. We write $\mathcal{D}(\mathcal{M}, S)$ 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 S-definable in M.

20.19. **Proposition.** *Let L be a first order language and $\mathcal{M}$ an L-structure with underlying set M; let S be any subset of M. Then $\mathcal{D}(\mathcal{M}, S)$ 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 20.19. It states that each definability system is *closed under definability.* If X is a nonempty set and $M \subseteq X^k$, then for each $m \geq 0$ we regard M^m as a subset of X^{km}.

20.20. **Theorem.** *Let X be a nonempty set and let $\mathcal{S}$ be a definability system on X. Let L be a first order language and $\mathcal{M}$ an L-structure whose underlying set is M; let S be a subset of M. Suppose all of the following sets belong to $\mathcal{S}$:*

(a) M;

(b) $\{c^{\mathcal{M}}\}$ for each constant symbol c in L;

(c) $\{s\}$ for each $s \in S$;

(d) $R^{\mathcal{M}}$ for each relation symbol R in L;

(e) the graph of $f^{\mathcal{M}}$ for each function symbol f in L.

Then every set that is S-definable in $\mathcal{M}$ belongs to $\mathcal{S}$.

We first give a series of basic results about definability systems that will be used in the proof of Theorem 20.20. For these lemmas we fix a nonempty set X and a definability system $\mathcal{S}$ on X.

20.21. **Lemma.** *If A and B belong to $\mathcal{S}$, then $A \times B$ belongs to $\mathcal{S}$.*

Proof. Suppose $A \subseteq X^m$ and $B \subseteq X^n$. Then

$$A \times B = (A \times X^n) \cap (X^m \times B).$$

Condition (2) of Definition 20.17 (used repeatedly) followed by condition (1) yields that this set belongs to $\mathcal{S}$. $\square$

20.22. **Lemma.** *For all $1 \le i < j \le m$ 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 20.17 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 20.17. (See also Lemma 20.21.) $\square$

20.23. **Lemma.** *Let $B \in \mathcal{S}_n$ and let $i(1), \dots, i(n) \in \{1, \dots, m\}$. Then the set $A \subseteq X^m$ defined by*

$$A \colon = \{(a_1, \dots, a_m) \in X^m \mid (a_{i(1)}, \dots, a_{i(n)}) \in B\}$$

belongs to $\mathcal{S}$.

Proof. Note that for any $a_1, \dots, a_m \in X$, the tuple $(a_1, \dots, a_m)$ is in A iff

$$\exists y_1 \dots \exists y_n (x_{i(1)} = y_1 \wedge \dots \wedge x_{i(n)} = y_n \wedge (y_1, \dots, y_n) \in B\}.$$

Let D_j denote the diagonal set $\Delta^{m+n}(i(j), m+j)$ and let $\pi_j \colon X^{m+j} \to X^{m+j-1}$ denote the projection map onto the first coordinates, for each $j = 1, \dots, n$. The displayed condition shows that

$$A = \pi_1(\dots \pi_n(D_1 \cap \dots \cap D_n \cap (X^m \times B)) \dots).$$

Using Definition 20.17 and Lemma 20.22 we see that $A \in \mathcal{S}_m$. $\square$

20.24. **Lemma.** *Suppose $A \subseteq X^m$, $B \subseteq X^n$, and $C \subseteq X^p$ belong to $\mathcal{S}$. Let $f \colon A \to B$ and $g \colon B \to C$ be functions that belong to $\mathcal{S}$. Then their composition $g \circ f \colon A \to C$ also belongs to $\mathcal{S}$.*

Proof. Let x be a tuple of variables that ranges over X^m and let y range over X^n and z range over X^p similarly. Use Lemma 20.23 and the equivalence

$$(x,z) \in \Gamma(g \circ f) \Leftrightarrow \exists y((x,y) \in \Gamma(f) \wedge (y,z) \in \Gamma(g)).$$

$\square$

20.25. **Lemma.** *Suppose $A \subseteq X^m$ belongs to $\mathcal{S}$. Let $f = (f_1, \dots, f_n) \colon A \to X^n$ be a function with coordinate functions $f_j \colon A \to X$. 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 $\pi_j \colon X^n \to X$ be the projection map defined by $\pi_j(x_1, \ldots, x_n) :\, = x_j$. Using Lemma 20.22 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 20.24 completes the proof of this direction.

($\Leftarrow$) Let x range over X^m and $y = (y_1, \ldots, y_n)$ range over X^n. The graph of f is defined by the equivalence

$$(x, y) \in \Gamma(f) \Leftrightarrow ((x, y_1) \in \Gamma(f_1) \wedge \ldots \wedge (x, y_n) \in \Gamma(f_n)).$$

If the functions $f_1, \ldots, f_n$ all belong to $\mathcal{S}$, this equivalence together with Lemma 20.23 and condition (1) of Definition 20.17 show that f belongs to $\mathcal{S}$. $\square$

Proof of Theorem 20.20.

Let L, $\mathcal{M}$, and S 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 S-definable set in $\mathcal{M}$ belongs to $\mathcal{S}$.

Let k be such that $M \subseteq X^k$. As noted above, we consider M^m as a subset of X^{km} for each $m \geq 0$.

First we prove the following statement by induction on the complexity of terms:

Let t be an $L(S)$-term, and let $x_1, \ldots, x_m$ be a sequence of distinct variables that includes all variables of t; the function $t^{\mathcal{M}} \colon M^m \to M$ defined by interpreting t in $\mathcal{M}$ 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 S, or one of the variables x_i. In the first case the graph of the function $t^{\mathcal{M}}$ is $M^m \times \{c^{\mathcal{M}}\}$, and the second case is similar; in the third case it is the intersection of M^{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 $f(t_1, \ldots, t_n)$ where f is an n-ary function symbol of L and $t_1, \ldots, t_n$ are $L(S)$-terms of which the statement being proved is true. Let $G \colon M^m \to M^n$ be the function with coordinate functions $t_j^{\mathcal{M}}$, $j = 1, \ldots, n$. Lemma 20.25 shows that G belongs to $\mathcal{S}$; Lemma 20.24 shows that $t^{\mathcal{M}} = f^{\mathcal{M}} \circ G$ 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 20.20 follows immediately; the proof is by induction on formulas:

Let φ be an $L(S)$-formula and let $x_1, \ldots, x_m$ be a sequence of distinct variables that includes all free variables of φ; the set

$$\varphi^{\mathcal{M}} :\, = \{(a_1, \ldots, a_m) \in M^m \mid \mathcal{M} \models \varphi[a_1, \ldots, a_m]\}$$

belongs to the definability system $\mathcal{S}$.

In the basic step of this induction φ is an atomic formula of the form $R(t_1, \ldots, t_n)$ where R is an n-ary relation symbol of L and $t_1, \ldots, t_n$ are $L(S)$-terms. Let $G\colon M^m \to M^n$ be the function defined above using the terms $t_1, \ldots, t_n$. As shown there, G belongs to $\mathcal{S}$. We see that $\varphi^{\mathcal{M}}$ is defined by the equivalence

$$(a_1, \ldots, a_m) \in \varphi^{\mathcal{M}} \Leftrightarrow$$
$$\exists y_1 \ldots \exists y_n((a_1, \ldots, a_m, y_1, \ldots, y_n) \in \Gamma(G) \wedge (y_1, \ldots, y_n) \in R^{\mathcal{M}}).$$

(Strictly speaking note that each $\exists y_j$ stands for a sequence of k existential quantifiers over X.) This shows that $\varphi^{\mathcal{M}}$ is the result of applying a sequence of kn projections to the set

$$\Gamma(G) \cap (M^m \times R^{\mathcal{M}})$$

which shows that $\varphi^{\mathcal{M}}$ 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 $x_1, \ldots, x_m$ 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 $\alpha^{\mathcal{M}}$ and $\beta^{\mathcal{M}}$, which are both subsets of M^m, belong to $\mathcal{S}$. It follows immediately from condition (1) of Definition 20.17 that $\varphi^{\mathcal{M}}$ 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 $x_1, \ldots, x_m$. (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 $x_1, \ldots, x_m$, then we apply the induction hypothesis to the formula ψ and the list of variables $x_1, \ldots, x_m, y$. Evidently $\varphi^{\mathcal{M}} = \pi(\psi^{\mathcal{M}})$, where $\pi\colon M^{m+1} \to M^m$ is the projection on the first m coordinates. Condition (4) of Definition 20.17 yields that $\varphi^{\mathcal{M}}$ belongs to $\mathcal{S}$. This completes the proof of Theorem 20.20. $\square$

Let X be a nonempty set. Given two definability systems $\mathcal{S}(1)$ and $\mathcal{S}(2)$ on X, 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 X; 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 X^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 X^m for all $m \geq 0$. It follows that there exists a smallest definability system $\mathcal{S}$ on X that contains $\mathcal{F}$. Namely, let $\mathcal{S}$ be the greatest lower bound (intersection) of all definability systems on X that contain $\mathcal{F}$. We call this *the definability system on X generated by $\mathcal{F}$.*

20.26. **Corollary.** *Let L be a first order language and let $\mathcal{M}$ be an L-structure with underlying set M; let S be a subset of M. Then $\mathcal{D}(\mathcal{M}, S)$ is the definability system on M generated by the sets listed in (b)–(e) of Theorem 20.20.*

Proof. Exercise. $\square$

Exercises

In the following Exercises, let $\mathcal{S}$ be a definability system on the nonempty set X.

20.27. Show that there is a language L and an L-structure $\mathcal{M}$ based on the set X such that $\mathcal{S} = \mathcal{D}(\mathcal{M}, \emptyset)$.

20.28. Let I be a finite index set and let $A \in \mathcal{S}_m$ be the union of the sets $A_i \in \mathcal{S}_m$, i ranging over I. Show that a function $f\colon A \to X^n$ belongs to $\mathcal{S}$ if and only if all of its restrictions $f|A_i$ belong to $\mathcal{S}$.

20.29. Let $A \subseteq X^{m+n}$ and $x \in X^m$, and put $A_x\colon = \{y \in X^n \mid (x, y) \in A\}$. Show that if A belongs to $\mathcal{S}$ and $k \in \mathbb{N}$, then the set $\{x \in X^m \mid \mathrm{card}(A_x) \leq k\}$ also belongs to $\mathcal{S}$.

20.30. Let the sets A, B, C and the function $f\colon A \times B \to C$ belong to $\mathcal{S}$. Show that the sets $\{a \in A \mid f(a, \cdot)\colon B \to C$ is injective $\}$ and $\{a \in A \mid f(a, \cdot)\colon B \to C$ is surjective $\}$ also belong to $\mathcal{S}$.

20.31. Suppose $X = \mathbb{R}$ and the order relation $\{(x, y) \in \mathbb{R}^2 \mid x < y\}$ belongs to $\mathcal{S}$. Suppose $A \subseteq \mathbb{R}^m$ belongs to $\mathcal{S}$. Show that the topological closure $\mathrm{cl}(A)$ of A and the interior $\mathrm{int}(A)$ of A in $\mathbb{R}^m$ also belong to $\mathcal{S}$.

20.32. Suppose $X = \mathbb{R}$ and the order relation $\{(x, y) \in \mathbb{R}^2 \mid x < y\}$ belongs to $\mathcal{S}$. Suppose that the function $f\colon \mathbb{R}^{m+1} \to \mathbb{R}$ belongs to $\mathcal{S}$. Show that the set

$$A\colon = \{a \in \mathbb{R}^m \mid f(a, t) \text{ tends to a limit } \ell(a) \text{ as } t \to +\infty\}$$

belongs to $\mathcal{S}$, and the limit function $\ell\colon A \to \mathbb{R}$ so defined also belongs to $\mathcal{S}$.