Model Theory

J. van Oosten
Department of Mathematics
Utrecht University

Spring 2000

Contents

1 Introduction and Notations

In the course Foundations of Mathematics we have made our acquaintance with the notions of a *language* $\mathcal{L}$ in predicate logic, and a *structure* $\mathfrak{A}$ for the language $\mathcal{L}$. A language is a collection of symbols, divided into three groups: we have *constants*, *function symbols* and *relation (or predicate) symbols*. Given a language $\mathcal{L}$, we have defined the class of *$\mathcal{L}$-terms*.

A *structure* $\mathfrak{A}$ for the language $\mathcal{L}$ is a pair $\mathfrak{A} = (A, (\cdot)^{\mathfrak{A}})$ where A is a *nonempty* set, and $(\cdot)^{\mathfrak{A}}$ is an *interpretation function*, defined on the symbols in $\mathcal{L}$, such that for every constant c of $\mathcal{L}$, $(c)^{\mathfrak{A}}$ is an element of A; for every n-place predicate symbol R, $(R)^{\mathfrak{A}}$ is a subset of the set A^n of n-tuples of elements of A; and for every n-ary function symbol F of $\mathcal{L}$, $(F)^{\mathfrak{A}}$ is a function from A^n to A. The set A is called the *domain*, *universe*, or *underlying set* of the structure $\mathfrak{A}$. In practice we shall often omit brackets, writing $F^{\mathfrak{A}}$ instead of $(F)^{\mathfrak{A}}$ and ditto for R.

In these notes, structures will be denoted by Gothic symbols $\mathfrak{A}, \mathfrak{B}, \ldots$; their domains will be denoted by the corresponding Latin characters $A, B, \ldots$.

Given a structure $\mathfrak{A}$ for the language $\mathcal{L}$, we consider the language $\mathcal{L}_{\mathfrak{A}}$: for every element a of A, we add a constant a to the language (and we assume, that the new constants a are different from constants that are already in the language $\mathcal{L}$). The structure $\mathfrak{A}$ becomes a $\mathcal{L}_{\mathfrak{A}}$-structure by putting $(a)^{\mathfrak{A}} = a$, for each $a \in A$.

We have also defined the relation $\mathfrak{A} \models \phi$, for $\mathcal{L}_{\mathfrak{A}}$-sentences ϕ: first, for closed $\mathcal{L}_{\mathfrak{A}}$-terms t, we define their meaning $t^{\mathfrak{A}}$ in $\mathfrak{A}$. Then we defined by recursion on the $\mathcal{L}_{\mathfrak{A}}$-sentence ϕ:

- $\mathfrak{A} \models t = s$ if and only if $t^{\mathfrak{A}} = s^{\mathfrak{A}}$;
- $\mathfrak{A} \models R(t_1, \ldots, t_n)$ if and only if $((t_1)^{\mathfrak{A}}, \ldots, (t_n)^{\mathfrak{A}}) \in R^{\mathfrak{A}}$;
- $\mathfrak{A} \models \phi \wedge \psi$ if and only if $\mathfrak{A} \models \phi$ and $\mathfrak{A} \models \psi$;
- $\mathfrak{A} \models \phi \vee \psi$ if and only if $\mathfrak{A} \models \phi$ or $\mathfrak{A} \models \psi$ (or both);
- $\mathfrak{A} \models \neg\phi$ if and only if $\mathfrak{A} \not\models \phi$;
- $\mathfrak{A} \models \exists x \phi(x)$ if and only if for some $a \in A$, $\mathfrak{A} \models \phi(a)$;
- $\mathfrak{A} \models \forall x \phi(x)$ if and only if for all $a \in A$, $\mathfrak{A} \models \phi(a)$

Furthermore we have learnt the notions of $\mathcal{L}$- (or $\mathcal{L}_{\mathfrak{A}}$-)*theory*: this is a collection of $\mathcal{L}$- (or $\mathcal{L}_{\mathfrak{A}}$-) sentences. $\mathfrak{A}$ is a *model* of the theory T if $\mathfrak{A} \models \phi$ for every $\phi \in T$. The relation $\mathfrak{A} \models \phi$ is pronounced as: $\mathfrak{A}$ *satisfies* ϕ, ϕ *is true in* $\mathfrak{A}$, or ϕ *holds in* $\mathfrak{A}$.

We also recall one of the most important theorems we proved in the course Foundations of Mathematics:

Theorem 1.1 (Compactness Theorem) *A theory T has a model if and only if every finite subset of T has a model.*

In this course, we shall give an independent proof of Theorem 1.1; that is, independent of the Completeness Theorem.

Often, we shall say that a theory is *consistent*; in this course, this is synonymous with: has a model.

We shall also use the notation $T \models \phi$, where T is an $\mathcal{L}$-theory, and ϕ and $\mathcal{L}$-sentence; this means that ϕ is true in every model of T.

A theory T is *complete* if for every sentence ϕ in the language of T, either $T \models \phi$ or $T \models \neg\phi$ holds.

2 Homomorphisms, Embeddings, and Diagrams

The purpose of Model Theory is the study of theories by means of their classes of models. A very useful tool of Model Theory is the possibility of varying the language. Suppose we have two languages $\mathcal{L} \subseteq \mathcal{L}'$. If $\mathfrak{B}$ is an $\mathcal{L}'$-structure, then restricting its interpretation function to $\mathcal{L}$ gives an $\mathcal{L}$-structure $\mathfrak{A}$. We say that $\mathfrak{A}$ is the *$\mathcal{L}$-reduct* of $\mathfrak{B}$, and that $\mathfrak{B}$ is an *$\mathcal{L}'$-expansion* of $\mathfrak{A}$. In our definition of $\mathfrak{A} \models \phi$, we have already seen the expansion of $\mathfrak{A}$ to $\mathcal{L}_{\mathfrak{A}}$.

We start by giving some structure on the class of $\mathcal{L}$-structures. Let $\mathfrak{A} = (A, (\cdot)^{\mathfrak{A}})$ and $\mathfrak{B} = (B, (\cdot)^{\mathfrak{B}})$ be $\mathcal{L}$-structures. A function $f : A \to B$ is called a *homomorphism* of $\mathcal{L}$-structures, if it commutes with the interpretation functions. That is:

- For every constant c of $\mathcal{L}$, $f((c)^{\mathfrak{A}}) = (c)^{\mathfrak{B}}$;
- for every n-ary function symbol F of $\mathcal{L}$ and every n-tuple $a_1, \ldots, a_n$ of elements of A, we have $f((F)^{\mathfrak{A}}(a_1, \ldots, a_n)) = (F)^{\mathfrak{B}}(f(a_1), \ldots, f(a_n))$;
- for every n-place predicate symbol R of $\mathcal{L}$ and n-tuple $a_1, \ldots, a_n$ of elements of A, we have: if $(a_1, \ldots, a_n) \in (R)^{\mathfrak{A}}$ then $(f(a_1), \ldots, f(a_n)) \in (R)^{\mathfrak{B}}$.

Examples. If $\mathcal{L}$ is the language of groups and $\mathfrak{A}$ and $\mathfrak{B}$ are groups, a homomorphism of $\mathcal{L}$-structures is nothing but a homomorphism of groups. If $\mathcal{L}$ is the language $\{\leq\}$ of partial orders and $\mathfrak{A}$, $\mathfrak{B}$ are partial orders, a homomorphism is nothing but an order-preserving map. Similar for: rings, graphs, etcetera.

We note immediately that if $f : \mathfrak{A} \to \mathfrak{B}$ and $g : \mathfrak{B} \to \mathfrak{C}$ are homomorphisms of $\mathcal{L}$-structures, then so is the composition $gf : \mathfrak{A} \to \mathfrak{C}$. Moreover, the identity function $A \to A$ is always a homomorphism of $\mathcal{L}$-structures. A homomorphism $f : \mathfrak{A} \to \mathfrak{B}$ is an *isomorphism* if there is a homomorphism $g : \mathfrak{B} \to \mathfrak{A}$ inverse to f (so the compositions gf and fg are the identity functions on A and B, respectively).

Exercise 1 Prove: if $f : \mathfrak{A} \to \mathfrak{B}$ is an isomorphism of $\mathcal{L}$-structures then for any $\mathcal{L}$-formula $\phi(x_1, \ldots, x_n)$ and any n-tuple $a_1, \ldots, a_n$ from A we have:

$$\mathfrak{A} \models \phi(a_1, \ldots, a_n) \Leftrightarrow \mathfrak{B} \models \phi(f(a_1), \ldots, f(a_n))$$

A consequence of this exercise is, that isomorphic $\mathcal{L}$-structures satisfy the same $\mathcal{L}$-sentences (check this!). Two $\mathcal{L}$-structures that satisfy the same $\mathcal{L}$-sentences are called *elementarily equivalent*. Notation: $\mathfrak{A} \equiv \mathfrak{B}$. The notation for isomorphic structures is $\cong$. Summarizing:

$$\mathfrak{A} \cong \mathfrak{B} \Rightarrow \mathfrak{A} \equiv \mathfrak{B}$$

The converse implication does not hold!

Exercise 2 Prove: an $\mathcal{L}$-theory T is complete if and only if every pair of models of T is elementarily equivalent.

One sees that the notion of completeness can be characterized by a property of the class of models of the theory.

A homomorphism $f : \mathfrak{A} \to \mathfrak{B}$ is called injective if the function f is. f is an *embedding* if f is injective and moreover for every relation symbol R of $\mathcal{L}$ and every n-tuple $a_1, \dots, a_n$ from A: if $(f(a_1), \dots, f(a_n)) \in (R)^{\mathfrak{B}}$, then $(a_1, \dots, a_n) \in (R)^{\mathfrak{A}}$. Clearly, if $\mathcal{L}$ contains no relation symbols, every injective homomorphism is an embedding, but in the general case the notions are different.

A *substructure* (or *submodel*) of an $\mathcal{L}$-structure $\mathfrak{B}$ is an $\mathcal{L}$-structure $\mathfrak{A}$ such that A is a subset of B and the interpretation function of $\mathfrak{A}$ is the restriction of the one of $\mathfrak{B}$ to A. Hence: an embedding $f : \mathfrak{A} \to \mathfrak{B}$ defines an isomorphism between $\mathfrak{A}$ and a submodel of $\mathfrak{B}$.

If $\mathfrak{B}$ is an $\mathcal{L}$-structure and $X \subseteq B$, there is a least subset A of B which contains X and the elements $c^{\mathfrak{B}}$ (c a constant of $\mathcal{L}$), and is closed under the interpretations in $\mathfrak{B}$ of the function symbols of $\mathcal{L}$; if $A \neq \emptyset$, A is the domain of a submodel $\mathfrak{A}$ of $\mathfrak{B}$, the submodel *generated by* the set X, which we denote by $\langle X \rangle$.

Exercise 3 Show that the domain of $\langle X \rangle$ is the set

$$\{t^{\mathfrak{B}}(x_1, \dots, x_n) \mid n \in \mathbb{N}, t(v_1, \dots, v_n) \text{ an } \mathcal{L}\text{-term}, x_1, \dots, x_n \in X\}$$

A structure $\mathfrak{B}$ is *finitely generated* if $\mathfrak{B} = \langle X \rangle$ for some finite subset X of B.

Recall that a formula is called *atomic* if it contains no connectives or quantifiers; in other words if it is of from $t_1 = t_2$ or $R(t_1, \dots, t_n)$.

Exercise 4 Let $f : \mathfrak{A} \to \mathfrak{B}$ an embedding. Then for every atomic $\mathcal{L}$-formula $\phi(x_1, \dots, x_n)$ and every n-tuple $a_1, \dots, a_n$ from A:

$$\mathfrak{A} \models \phi(a_1, \dots, a_n) \Leftrightarrow \mathfrak{B} \models \phi(f(a_1), \dots, f(a_n))$$

Prove also that this equivalence in fact holds for every quantifier-free formula.

Exercise 5 Let $\mathcal{L} = \{\leq\}$ be the language of partial orders. Find out when a monotone map between partial orders is an injective $\mathcal{L}$-homomorphism, and when it is an embedding. Give an example of an injective homomorphism which is not an embedding.

Let $\mathfrak{A}$ be an $\mathcal{L}$-structure with its natural expansion to $\mathcal{L}_{\mathfrak{A}}$. We define a few $\mathcal{L}_{\mathfrak{A}}$-theories in connection to $\mathfrak{A}$.

- The *positive diagram* $\Delta^+_{\mathfrak{A}}$ of $\mathfrak{A}$ is the collection of all atomic $\mathcal{L}_{\mathfrak{A}}$-sentences that are true in $\mathfrak{A}$.
- The *diagram* $\Delta_{\mathfrak{A}}$ of $\mathfrak{A}$ is the collection of all atomic $\mathcal{L}_{\mathfrak{A}}$-sentences and all negations of atomic $\mathcal{L}_{\mathfrak{A}}$-sentences, that are true in $\mathfrak{A}$.

The notion of *positive diagram* generalizes the idea of *multiplication table* of a group.

Exercise 6 Show that giving an $\mathcal{L}$-homomorphism $\mathfrak{A} \to \mathfrak{B}$ is equivalent to giving an expansion of $\mathfrak{B}$ to $\mathcal{L}_{\mathfrak{A}}$, which is a model of the positive diagram of $\mathfrak{A}$. Show also, that giving an embedding $\mathfrak{A} \to \mathfrak{B}$ is equivalent to giving an expansion of $\mathfrak{B}$ to $\mathcal{L}_{\mathfrak{A}}$ which is a model of the diagram of $\mathfrak{A}$.

We now give an example, in order to apply our definitions to a little mathematical theorem.

Theorem 2.1 (Order Extension Principle) *Let $(A, \leq)$ be a partially ordered set. Then there is a linear order $\leq'$ on A which extends $\leq$; equivalently, there is an injective, monotone function from A into a linearly ordered set $(B, \leq'')$.*

Proof. The equivalence stated in the theorem is clear, because if $(B, \leq'')$ is a linear order and $f : A \to B$ an injective, order-preserving map, then defining $\leq'$ on A by: $x \leq' y$ iff $f(x) \leq'' f(y)$, gives a linear order on A which extends $\leq$.

We prove the theorem in two steps.

a) First, we do it for A finite. This is easy; induction on the number of elements of A. If $|A| = 1$ we are done. Now if $A = \{a_1, \ldots, a_{n+1}\}$, choose $a_i \in A$ such that a_i is maximal w.r.t. the order on A. By induction hypothesis $A \backslash \{a_i\}$ has a linear order which extends the one of A relativized to $A \backslash \{a_i\}$; put a_i back in, as greatest element.

b) The general case. Let $\mathcal{L} = \{\leq\}$. We use the idea of diagrams, and the Compactness Theorem. Let T be the $\mathcal{L}_{\mathfrak{A}}$-theory (where $\mathfrak{A}$ is the partial order A, seen as $\mathcal{L}$-structure) consisting of the axioms:

1. The positive diagram of $\mathfrak{A}$;
2. The axioms of a linear order;
3. The axioms $\neg(a = b)$ for $a \neq b \in A$.

In every finite subset D of T, there is only a finite number of constants from the set A; say $a_1, \ldots, a_n$. Let $\{a_1, \ldots, a_n\}$ be the sub-partial order of A on these elements. This has, by a), an extension to a linear order, so D has a model. By the Compactness Theorem, T has a model $(B, \leq'')$; since B is a model of 1), 2) and 3), there is an injective, monotone map from A to B. ■

Exercise 7 As an example of theorem 2.1, define an injective, order-preserving map from $\mathcal{P}(\mathbb{N})$ to $\mathbb{R}$, where $\mathcal{P}(\mathbb{N})$ is the powerset of $\mathbb{N}$, ordered by the subset relation.

3 Elementary Embeddings and Elementary Diagrams

An embedding $f : \mathfrak{A} \to \mathfrak{B}$ of $\mathcal{L}$-structures is called *elementary* if for every $\mathcal{L}$-formula $\phi(x_1, \ldots, x_n)$ with free variables $x_1, \ldots, x_n$ and every n-tuple $a_1, \ldots, a_n$ of elements of A, we have:

$$\mathfrak{A} \models \phi(a_1, \ldots, a_n) \Leftrightarrow \mathfrak{B} \models \phi(f(a_1), \ldots, f(a_n))$$

This means that the elements of A have the same properties with respect to A as to B. For example, if we consider $\mathbb{Q}$ and $\mathbb{R}$ as fields (or rings), the embedding is not elementary since the element 2 of $\mathbb{Q}$ is a square in $\mathbb{R}$ but not in $\mathbb{Q}$. However, if we consider $\mathbb{Q}$ and $\mathbb{R}$ just as ordered structures, the embedding is elementary, as we shall see later.

The *elementary diagram* $E(\mathfrak{A})$ of $\mathfrak{A}$ is the collection of all $\mathcal{L}_{\mathfrak{A}}$-sentences which are true in $\mathfrak{A}$.

Exercise 8 Giving an elementary embedding $f : \mathfrak{A} \to \mathfrak{B}$ is equivalent to giving an $\mathcal{L}_{\mathfrak{A}}$-expansion of $\mathfrak{B}$ which is a model of $E(\mathfrak{A})$.

$\mathfrak{A}$ is called an *elementary submodel* of $\mathfrak{B}$, notation: $\mathfrak{A} \preceq \mathfrak{B}$, if $\mathfrak{A}$ is a substructure of $\mathfrak{B}$, and the inclusion is an elementary embedding. Therefore, an elementary embedding $\mathfrak{A} \to \mathfrak{B}$ is an isomorphism between $\mathfrak{A}$ and an elementary submodel of $\mathfrak{B}$. We also say that $\mathfrak{B}$ is an *elementary extension* of $\mathfrak{A}$.

Note, that $\mathfrak{A} \preceq \mathfrak{B}$ implies $\mathfrak{A} \equiv \mathfrak{B}$. The converse implication, however, is by no means true, not even if $\mathfrak{A}$ is a submodel of $\mathfrak{B}$. Example: let, for $\mathcal{L} = \{\leq\}$, $\mathfrak{A} = \mathbb{N}\backslash\{0\}$ and $\mathfrak{B} = \mathbb{N}$. Then $\mathfrak{A} \cong \mathfrak{B}$, but the inclusion is not an isomorphism, and not an elementary embedding (check!).

Exercise 9 $\mathfrak{A} \preceq \mathfrak{B}$ if and only if $\mathfrak{A}$ is a submodel of $\mathfrak{B}$ and for every $\mathcal{L}_{\mathfrak{A}}$-sentence of the form $\exists y \phi(y)$ which is true in $\mathfrak{B}$, there is a $b \in A$ such that $\phi(b)$ is true in $\mathfrak{B}$.
Hint: use induction on formulas.

The notion of elementary submodel gives rise to the following important theorems, called the *Löwenheim-Skolem-Tarski Theorems*; roughly, together they say that to infinite structures there are elementarily equivalent structures of almost arbitrary infinite cardinality. That is: predicate logic has nothing to say about infinite cardinalities!

Theorem 3.1 (Upward Löwenheim-Skolem-Tarski Theorem) *Every infinite $\mathcal{L}$-structure has arbitrarily large elementary extensions.*

Proof. Let $\mathfrak{A}$ be infinite (i.e., A is infinite). By "arbitrarily large" we mean: for every set X there is an elementary extension $\mathfrak{B}$ of $\mathfrak{A}$ such that there is an injective function from X into B. The proof is a simple application of the Compactness Theorem (1.1).

Given X, we choose for every $x \in X$ a new constant c_x, not in $\mathcal{L}_{\mathfrak{A}}$. Let $\mathcal{L}' = \mathcal{L}_{\mathfrak{A}} \cup \{c_x \mid x \in X\}$. Consider the $\mathcal{L}'$-theory Γ:

$$E(\mathfrak{A}) \cup \{\neg(c_x = c_y) \mid x, y \in X, x \neq y\}$$

Since $\mathfrak{A}$ is infinite, every finite subset of Γ has an interpretation in $\mathfrak{A}$ (simply interpret the constants c_x by different elements of A) and hence has a model. By the Compactness Theorem, Γ has a model $\mathfrak{B}$. Since $\mathfrak{B}$ is a model of $E(\mathfrak{A})$, there is an elementary embedding $\mathfrak{A} \to \mathfrak{B}$, so we can identify $\mathfrak{A}$ with an elementary submodel of $\mathfrak{B}$. And the assignment $x \mapsto (c_x)^{\mathfrak{B}}$ is an injective function from X into B. ∎

The Downward Löwenheim-Skolem-Tarski Theorem is a little bit more involved. First, we recall that $\|\mathcal{L}\|$ is defined as the maximum of the cardinal numbers ω and $|\mathcal{L}|$ (ω is the cardinality of the set $\mathbb{N}$). Recall also, that $\|\mathcal{L}\|$ is the cardinality of the set of all $\mathcal{L}$-formulas.

Theorem 3.2 (Downward Löwenheim-Skolem-Tarski Theorem) *Let $\mathfrak{B}$ be an $\mathcal{L}$-structure such that $\|\mathcal{L}\| \leq |B|$ and suppose $X \subset B$ is a set with $\|\mathcal{L}\| \leq |X|$. Then there is an elementary submodel $\mathfrak{A}$ of $\mathfrak{B}$ such that $X \subseteq A$ and $|X| = |A|$.*

Proof. Given X, let $\mathcal{L}_X = \mathcal{L} \cup X$ (elements of X as new constants; as usual, we take this union to be disjoint). Note, that $\|\mathcal{L}_X\| = |X|$. For every $\mathcal{L}_X$-sentence $\phi = \exists x \psi(x)$ which is true in $\mathfrak{B}$, we choose an element $b \in B$ such that $\mathfrak{B} \models \psi(c_b)$. Let X_1 be the union of X and all the b's so chosen. Again, since we add at most $\|\mathcal{L}_X\|$ many new elements, $|X_1| = |X|$. Now repeat this, with X_1 in the place of X (and $\mathcal{L}_{X_1}$, etcetera), obtaining X_2, and this infinitely often, obtaining a chain

$$X = X_0 \subset X_1 \subset X_2 \subset \ldots$$

Let $A = \bigcup_{n \in \mathbb{N}} X_n$. By induction one proves easily that $|X_n| = |X|$ for all n, so $|A| = |X|$ since X is infinite.

Now if $\phi = \exists x \psi(x)$ is an $\mathcal{L}_A$-formula which is true in $\mathfrak{B}$, then for some n, ϕ is an $\mathcal{L}_{X_n}$-formula, so by construction there is a $b \in X_{n+1} \subset A$ such that $\mathfrak{B} \models \psi(b)$. In particular, this holds for formulas $\exists x(x = F(a_1, \ldots, a_n))$, so A is closed under the interpretations in $\mathfrak{B}$ of the function symbols of $\mathcal{L}$, hence A is the domain of a submodel $\mathfrak{A}$ of $\mathfrak{B}$. But by the same token, $\mathfrak{A}$ is an elementary submodel (exercise 9). ∎

An immediate application of the Löwenheim-Skolem-Tarski theorems is the so-called *Łos-Vaught test.* A theory is called *α-categorical* for some cardinal number α, if for every pair $\mathfrak{A}, \mathfrak{B}$ of models of T with $|A| = |B| = \alpha$, we have $\mathfrak{A} \cong \mathfrak{B}$.

Theorem 3.3 (Łos-Vaught Test) *Suppose T is a theory which has only infinite models, and T is α-categorical for some cardinal number $\alpha \geq \|\mathcal{L}\|$ ($\mathcal{L}$ is the language of T). Then T is complete.*

Proof. Suppose $\mathfrak{A}$ and $\mathfrak{B}$ are models of T. Then both A and B are infinite by hypothesis, so if β is the maximum of $\{\alpha, |A|, |B|\}$ then by 3.1 both $\mathfrak{A}$ and $\mathfrak{B}$ have elementary extensions $\mathfrak{A}'$, $\mathfrak{B}'$ respectively, with $|A'| = |B'| = \beta$. By 3.2, $\mathfrak{A}'$ and $\mathfrak{B}'$ have elementary submodels $\mathfrak{A}''$ and $\mathfrak{B}''$ respectively, of cardinality α. Since T is α-categorical, we have:

$$\mathfrak{A} \equiv \mathfrak{A}' \equiv \mathfrak{A}'' \cong \mathfrak{B}'' \equiv \mathfrak{B}' \equiv \mathfrak{B}$$

Hence, $\mathfrak{A} \equiv \mathfrak{B}$ for any two models $\mathfrak{A}, \mathfrak{B}$ of T; hence, T is complete (exercise 2).
■

Examples of categorical theories

a) Let $\mathcal{L} = \{\leq\}$ and T the theory of dense linear orders without end-points. By the familiar back-and-forth construction, T is ω-categorical. T is not 2^ω-categorical, because $(0, 1)$ and $(0, 1) \cup (1, 2)$ are nonisomorphic models of T.

b) "Torsion-free divisible abelian groups".
Let $\mathcal{L}$ be the language of groups, say $\mathcal{L} = \{0; +; -\}$ and T the $\mathcal{L}$-theory with axioms:

$$\begin{array}{ll} & \text{Axioms of an abelian group} \\ \text{“torsion-free”} & \{\forall x(\underbrace{x + \cdots + x}_{n \text{ times}} = 0 \rightarrow x = 0) \mid n \geq 1\} \\ \text{“divisible”} & \{\forall x \exists y(\underbrace{y + \cdots + y}_{n \text{ times}} = x) \mid n \geq 1\} \end{array}$$

T is *not* ω-categorical, but T is categorical for uncountable cardinalities α. Sketch of proof: a divisible torsion-free abelian group is a $\mathbb{Q}$-vector space. If such a space has uncountable cardinality α, then since $\mathbb{Q}$ is countable, it must have a basis of cardinality α; but any two vector spaces over the same field with bases of the same cardinality are isomorphic as vector spaces, hence as abelian groups.

If the cardinality of the vector space is ω however, then its basis may be finite or countably infinite.

c) "Algebraically closed fields of a given fixed characteristic".
Again, this is categorical in every uncountable cardinality but not ω-categorical; consider the fields $\mathbb{Q}$ and $\mathbb{Q}(X)$. Their algebraic closures $\overline{\mathbb{Q}}$ and $\overline{\mathbb{Q}(X)}$ are both countable, but not isomorphic (WhyΓ).

d) Let $\mathcal{L}$ consist of one 1-place function symbol F. Let T be the theory with axioms:

$$\begin{array}{l} \forall xy(F(x) = F(y) \rightarrow x = y) \\ \forall x \exists y(F(y) = x) \\ \{\forall x \neg(F^n(x) = x) \mid n \geq 1\} \end{array}$$

A model of T is nothing but a set X with a $\mathbb{Z}$-action on it, which action is *free*. So X naturally decomposes into a number of orbits, which are all in

1-1 correspondence with $\mathbb{Z}$. Therefore if $|X| > \omega$, the number of orbits is equal to $|X|$. So T is categorical in every uncountable cardinality. However if X is countable, the number of orbits may be either finite or countably infinite.

In this list we have seen examples of theories that are ω-categorical but not categorical in higher cardinalities, or the converse: categorical in every uncountable cardinal, but not in ω (of course, there are also theories which are categorical in every cardinality: the empty theory). That this is not a coincidence, is the content of the famous *Morley Categoricity Theorem*: if a theory (in a countable language) is categorical in an uncountable cardinality, it is categorical in *every* uncountable cardinality. This is a deep result of Model Theory, and the starting point of a whole branch of Model Theory, *Stability Theory*.

4 Directed Systems of $\mathcal{L}$-structures

Let $(K, \leq)$ be a partially ordered set (or *poset* for short). K is called *directed* if K is nonempty and every pair of elements of K has an upper bound in K, that is K satisfies $\forall xy \exists z (x \leq z \wedge y \leq z)$.

A *directed system* of $\mathcal{L}$-structures consists of a family $(\mathfrak{A}_k)_{k \in K}$ of $\mathcal{L}$-structures indexed by K, together with homomorphisms $f_{kl} : \mathfrak{A}_k \to \mathfrak{A}_l$ for $k \leq l$. These homomorphisms should satisfy:

- f_{kk} is the identity homomorphism on $\mathfrak{A}_k$
- if $k \leq l \leq m$, then $f_{km} = f_{lm} f_{kl}$

Given such a directed system, we define its *colimit* as follows.

First, take the disjoint union of all the sets A_k:

$$\bigsqcup_{k \in K} A_k = \{(k, a) \mid k \in K, a \in A_k\}$$

Then define an equivalence relation $\sim$ on this by putting $(k, a) \sim (l, b)$ if there is $m \geq k, l$ in K such that $f_{km}(a) = f_{lm}(b)$ (see for yourself how the directedness of K is used to show that this relation is transitive!). Let $A = (\bigsqcup_{k \in K} A_k)/\sim$ be the set of equivalence classes; we denote the equivalence class of (k, a) by $[k, a]$.

We define an $\mathcal{L}$-structure $\mathfrak{A}$ with underlying set A as follows. For an n-ary function symbol F of $\mathcal{L}$ we put

$$F^{\mathfrak{A}}([k_1, a_1], \ldots, [k_n, a_n]) = [k, F^{\mathfrak{A}_k}(f_{k_1 k}(a_1), \ldots, f_{k_n k}(a_n))]$$

where, by directedness of K, k is chosen so that $k \geq k_1, \ldots, k \geq k_n$ all hold. Of course we must show that this definition makes sense: that it does not depend on the choice of representatives, or the choice of k. This is done in the following exercise:

Exercise 10 Prove: if $(k_1, a_1) \sim (k'_1, a'_1), \ldots, (k_n, a_n) \sim (k'_n, a'_n)$ and $k' \geq k'_1, \ldots, k'_n$, then

$$(k, F^{\mathfrak{A}_k}(f_{k_1 k}(a_1), \ldots, f_{k_n k}(a_n))) \sim (k', F^{\mathfrak{A}_{k'}}(f_{k'_1 k'}(a'_1), \ldots, f_{k'_n k'}(a'_n)))$$

If R is an n-place relation symbol of $\mathcal{L}$ and $[k_1, a_1], \ldots, [k_n, a_n]$ are n elements of A, we let $([k_1, a_1], \ldots, [k_n, a_n]) \in R^{\mathfrak{A}}$ if and only if there is $k \geq k_1, \ldots, k_n$ in K such that

$$(f_{k_1 k}(a_1), \ldots, f_{k_n k}(a_n)) \in R^{\mathfrak{A}_k}$$

Again, check that this does not depend on representatives.

Lemma 4.1 *The function $f_k : A_k \to A$ defined by $f_k(a) = [k, a]$, is a homomorphism of $\mathcal{L}$-structures. If $k \leq l$, the diagram*

$$\begin{tikzcd}
\mathfrak{A}_k \arrow[rr, "f_{kl}"] \arrow[dr, "f_k"'] & & \mathfrak{A}_l \arrow[dl, "f_l"] \\
& \mathfrak{A} &
\end{tikzcd}$$

of $\mathcal{L}$-structures and homomorphisms, is commutative. Moreover, if $(g_k : \mathfrak{A}_k \to \mathfrak{B})_{k \in K}$ is another K-indexed system of homomorphisms satisfying $g_l f_{kl} = g_k$ whenever $k \leq l$, there is a unique homomorphism $f : \mathfrak{A} \to \mathfrak{B}$ such that $f f_k = g_k$ for all $k \in K$.

Exercise 11 Prove lemma 4.1.

Exercise 12 In the notation of lemma 4.1 and above: if all the maps $f_{kl} : \mathfrak{A}_k \to \mathfrak{A}_l$ are injective then all the maps $f_k : \mathfrak{A}_k \to \mathfrak{A}$ are injective. The same holds with "injective" replaced by "an embedding".

Lemma 4.2 (Elementary System Lemma) *If all $f_{kl} : \mathfrak{A}_k \to \mathfrak{A}_l$ are elementary embeddings, so are all the maps f_k.*

Proof. So we suppose all $f_{kl} : \mathfrak{A}_k \to \mathfrak{A}_l$ are elementary embeddings. Now by induction we prove for an $\mathcal{L}$-formula $\phi(x_1, \ldots, x_n)$:

For all $k \in K$ and all $a_1, \ldots, a_n \in A_k$,

$$\mathfrak{A}_k \models \phi(a_1, \ldots, a_n) \Leftrightarrow \mathfrak{A} \models \phi(f_k(a_1), \ldots, f_k(a_n))$$

Note, that the universal quantifier "for all $k \in K$" occurs *inside* the induction hypothesis!

I give only the quantifier step (in fact, the step for atomic ϕ is included in exercise 12, and the steps for the propositional connectives are easy). So let $\phi \equiv \exists x \psi$.

If $\mathfrak{A}_k \models \exists x \psi(x, a_1, \ldots, a_n)$ so for some $a \in A_k$, $\mathfrak{A}_k \models \psi(a, a_1, \ldots, a_n)$, then by induction hypothesis $\mathfrak{A} \models \psi(f_k(a), f_k(a_1), \ldots, f_k(a_n))$ whence $\mathfrak{A} \models \exists x \psi(x, f_k(a_1), \ldots, f_k(a_n))$.

Conversely, suppose $\mathfrak{A} \models \exists x \psi(x, f_k(a_1), \ldots, f_k(a_n))$. Then for some $[k', a] \in A$ we have $\mathfrak{A} \models \psi([k', a], f_k(a_1), \ldots, f_k(a_n))$. Take, by directedness of K, a $k'' \in K$ with $k'' \geq k, k'$. Since $f_{k''} f_{k'k''} = f_{k'}$, $f_{k''} f_{kk''} = f_k$ and $[k', a] = [k'', f_{k'k''}(a)]$, we have

$$\mathfrak{A} \models \psi(f_{k''}(f_{k'k''}(a)), f_{k''}(f_{kk''}(a_1)), \ldots, f_{k''}(f_{kk''}(a_n)))$$

By induction hypothesis, we have:

$$\mathfrak{A}_{k''} \models \psi(f_{k'k''}(a), f_{kk''}(a_1), \ldots, f_{kk''}(a_n))$$

whence $\mathfrak{A}_{k''} \models \exists x \psi(x, f_{kk''}(a_1), \ldots, f_{kk''}(a_n))$. Now we use the fact that $f_{kk''}$ is an elementary embedding, to deduce that $\mathfrak{A}_k \models \exists x \psi(x, a_1, \ldots, a_n)$, and we are done. ∎

5 Theorems of Robinson, Craig and Beth

We shall see an application of the Elementary System Lemma 4.2 below, in the proof of Theorem 5.2. First a lemma whose proof is another application of the Compactness Theorem:

Lemma 5.1 *Let $\mathcal{L} \subseteq \mathcal{L}'$, $\mathfrak{A}$ an $\mathcal{L}$-structure and $\mathfrak{B}$ an $\mathcal{L}'$-structure. Suppose moreover that $\mathfrak{A}$ is elementarily equivalent to the $\mathcal{L}$-reduct of $\mathfrak{B}$. Then there is an $\mathcal{L}'$-structure $\mathfrak{C}$ and a diagram of embeddings:*

$$\begin{tikzcd} \mathfrak{A} \arrow[dr, "f"] & \\ \mathfrak{B} \arrow[r, "g"'] & \mathfrak{C} \end{tikzcd}$$

where g is an elementary embedding of $\mathcal{L}'$-structures, and f is an elementary embedding of $\mathcal{L}$-structures.

Proof. We consider the language $\mathcal{L}'_{\mathfrak{A}\mathfrak{B}} = \mathcal{L}_{\mathfrak{A}} \cup \mathcal{L}'_{\mathfrak{B}}$, where we take the constants from $\mathfrak{A}$ and $\mathfrak{B}$ disjoint. In $\mathcal{L}'_{\mathfrak{A}\mathfrak{B}}$ consider the theory $T = E(\mathfrak{A}) \cup E(\mathfrak{B})$ ($E(\mathfrak{A})$ is an $\mathcal{L}_{\mathfrak{A}}$-theory, and $E(\mathfrak{B})$ is an $\mathcal{L}'_{\mathfrak{B}}$-theory, so both are $\mathcal{L}'_{\mathfrak{A}\mathfrak{B}}$-theories). Then any model $\mathfrak{C}$ of T gives us the required diagram, by exercise 8.

So for contradiction, suppose T has no model. By Compactness, some finite subset of T has no model, so taking conjunctions we may assume that for some $\phi(a_1, \ldots, a_n) \in E(\mathfrak{A})$ and $\psi(b_1, \ldots, b_m) \in E(\mathfrak{B})$, $\phi \wedge \psi$ has no model. But then, no $\mathcal{L}'_{\mathfrak{B}}$-structure which satisfies $\psi(b_1, \ldots, b_m)$ can be expanded with interpretations for $a_1, \ldots, a_n$ such that the expansion satisfies ϕ; this means that every $\mathcal{L}'_{\mathfrak{B}}$-structure which satisfies $\psi(b_1, \ldots, b_m)$ will satisfy $\forall x_1 \cdots \forall x_n \neg\phi(x_1, \ldots, x_n)$. This is a contradiction, since this is an $\mathcal{L}$-sentence, $\mathfrak{A} \models \exists x_1 \cdots x_n \phi(x_1, \ldots, x_n)$, and $\mathfrak{A}$ is elementarily equivalent to the $\mathcal{L}$-reduct of $\mathfrak{B}$. ∎

Exercise 13 Prove that for two $\mathcal{L}$-structures $\mathfrak{A}$ and $\mathfrak{B}$, the following are equivalent:

i) $\mathfrak{A} \equiv \mathfrak{B}$

ii) $\mathfrak{A}$ and $\mathfrak{B}$ have a common elementary extension.

Theorem 5.2 (Robinson's Consistency Theorem) *Let $\mathcal{L}_1$ and $\mathcal{L}_2$ be two languages, $\mathcal{L} = \mathcal{L}_1 \cap \mathcal{L}_2$. Suppose that T is a complete $\mathcal{L}$-theory and let T_1 be an $\mathcal{L}_1$-theory, T_2 an $\mathcal{L}_2$-theory which both extend T. If both T_1 and T_2 have a model, so has $T_1 \cup T_2$.*

Proof. Let $\mathfrak{A}_0$ be a model of T_1, $\mathfrak{B}_0$ a model of T_2. Then $\mathfrak{A}_0$ and $\mathfrak{B}_0$ are also models of T and hence, since T is complete, their $\mathcal{L}$-reducts are elementarily equivalent. By lemma 5.1, there is a diagram

$$\begin{tikzcd}
\mathfrak{A}_0 \arrow[dr, "f_0"] & \\
\mathfrak{B}_0 \arrow[r, "h_0"'] & \mathfrak{B}_1
\end{tikzcd}$$

with h_0 an elementary extension of $\mathcal{L}_2$-structures, and f_0 elementary as a map between $\mathcal{L}$-structures. The $\mathcal{L}$-reducts of $\mathfrak{A}_0$ and $\mathfrak{B}_1$ are still elementarily equivalent, so applying the same lemma in the other direction gives a diagram

$$\begin{tikzcd}
\mathfrak{A}_0 \arrow[r, "k_0"] \arrow[dr, "g_0"'] & \mathfrak{A}_1 \\
\mathfrak{B}_0 \arrow[r, "h_0"'] & \mathfrak{B}_1 \arrow[u, "g_0"']
\end{tikzcd}$$

with g_0 an elementary map of $\mathcal{L}$-structures and k_0 elementary of $\mathcal{L}_1$-structures. We can proceed in this way, obtaining a directed system

$$\begin{tikzcd}
\mathfrak{A}_0 \arrow[r, "k_0"] \arrow[dr, "f_0"'] & \mathfrak{A}_1 \arrow[r, "k_1"] \arrow[dr, "f_1"'] & \mathfrak{A}_2 \arrow[r] & \cdots \\
\mathfrak{B}_0 \arrow[r, "h_0"'] & \mathfrak{B}_1 \arrow[r, "h_1"'] \arrow[u, "g_0"'] & \mathfrak{B}_2 \arrow[r] \arrow[u, "g_1"'] & \cdots
\end{tikzcd}$$

in which the k's are elementary maps of $\mathcal{L}_1$-structures, the f's and g's of $\mathcal{L}$-structures, and the h's of $\mathcal{L}_2$-structures. Since the colimit $\mathfrak{C}$ of this system is equally the colimit of the $\mathfrak{A}$'s and the colimit of the $\mathfrak{B}$'s, it has an $\mathcal{L}_1 \cup \mathcal{L}_2$-structure into which both $\mathfrak{A}_0$ and $\mathfrak{B}_0$ embed elementarily, by lemma 4.2. So $\mathfrak{C}$ is the required model of $T_1 \cup T_2$. ∎

The following theorem states a property that is commonly known as "the amalgamation property for elementary embeddings".

Theorem 5.3 *Every diagram*

$$\begin{tikzcd}
& \mathfrak{B} \\
\mathfrak{A} \arrow[ur, "f"] \arrow[dr, "g"'] & \\
& \mathfrak{C}
\end{tikzcd}$$

consisting of elementary embeddings between $\mathcal{L}$-structures, can be completed to a commutative diagram

$$\begin{tikzcd}
& \mathfrak{B} \arrow[dr] & \\
\mathfrak{A} \arrow[ur, "f"] \arrow[dr, "g"'] & & \mathfrak{D} \\
& \mathfrak{C} \arrow[ur] &
\end{tikzcd}$$

of elementary embeddings between $\mathcal{L}$-structures.

Proof. Formulate the elementary diagram $E(\mathfrak{B})$ in the language

$$\mathcal{L}_1 = \mathcal{L}_{\mathfrak{A}} \cup \{b \mid b \notin f[A]\}$$

(simply replacing the constants $f(a)$ by a); and similarly, write $E(\mathfrak{C})$ in

$$\mathcal{L}_2 = \mathcal{L}_{\mathfrak{A}} \cup \{c \mid c \notin g[A]\}$$

Then $\mathcal{L}_1 \cap \mathcal{L}_2 = \mathcal{L}_{\mathfrak{A}}$. In $\mathcal{L}_{\mathfrak{A}}$ we have the complete theory $E(\mathfrak{A})$, and $E(\mathfrak{B})$ and $E(\mathfrak{C})$ are extensions of it in $\mathcal{L}_1$, $\mathcal{L}_2$, respectively. By 5.2, $E(\mathfrak{B}) \cup E(\mathfrak{C})$ has a model $\mathfrak{D}$. Check that this gives a diagram of the required form. ∎

Exercise 14 Give also a direct proof of theorem 5.3, not using theorem 5.2 but along the lines of the proof of lemma 5.1.

A more serious application of Robinson's Consistency Theorem is Craig's Interpolation Theorem.

Theorem 5.4 (Craig Interpolation Theorem) *Suppose ϕ and ψ are $\mathcal{L}$-sentences such that $\phi \models \psi$, that is: every $\mathcal{L}$-structure in which ϕ is true, satisfies also ψ. Then there is an $\mathcal{L}$-sentence θ with the properties:*

i) $\phi \models \theta$ *and* $\theta \models \psi$;

ii) *Every non-logical symbol (function symbol, constant or relation symbol) which occurs in θ, also occurs both in ϕ and ψ*

Proof. Such a θ as in the theorem is called a *Craig interpolant* for ϕ and ψ. We assume for a contradiction, that no such Craig interpolant exists.

Then ϕ has a model, for otherwise $\exists x \neg(x = x)$ is a Craig interpolant. Similarly, $\neg\psi$ has a model, for otherwise $\forall x(x = x)$ were a Craig interpolant.

Let $\mathcal{L}_\phi$ the set of non-logical symbols in ϕ, and $\mathcal{L}_\psi$ of those in ψ; let $\mathcal{L}' = \mathcal{L}_\phi \cap \mathcal{L}_\psi$. Let Γ be the set of $\mathcal{L}'$-sentences σ such that either $\phi \models \sigma$ or $\neg\psi \models \sigma$.

Suppose $\Gamma \cup \{\phi\}$ has no model. Then by compactness, there are sentences σ_1 and σ_2 such that $\phi \models \sigma_1$ and $\neg\psi \models \sigma_2$, and $\{\sigma_1, \sigma_2, \phi\}$ has no model. Then since $\phi \models \sigma_1$, $\{\phi, \sigma_2\}$ has no model, whence

$$\phi \models \neg\sigma_2 \text{ and } \neg\sigma_2 \models \psi$$

contradicting our assumption that no Craig interpolant exists. In a similar way, $\Gamma \cup \{\neg\psi\}$ is consistent.

Now apply Zorn's Lemma to find a set Γ' of $\mathcal{L}'$-sentences which is maximal with respect to the properties that $\Gamma \subseteq \Gamma'$ and both $\Gamma' \cup \{\phi\}$ and $\Gamma' \cup \{\neg\psi\}$ are consistent. I claim that Γ' is a complete $\mathcal{L}'$-theory.

For, suppose Γ' is not complete, then there is a sentence σ such that $\sigma \notin \Gamma'$ and $\Gamma' \cup \{\sigma\}$ is consistent. Then by maximality of Γ', either $\Gamma' \cup \{\sigma, \phi\}$, or $\Gamma' \cup \{\sigma, \neg\psi\}$ is inconsistent. In the first case, there is a sentence $\tau \in \Gamma'$ such that $\phi \models \tau \to \neg\sigma$. Then $\tau \to \neg\sigma$ is, by definition of Γ, in Γ'; since also $\tau \in \Gamma'$ we see that $\Gamma \cup \{\sigma\}$ is *in*consistent, contradicting the choice of σ. In the other case one obtains a similar contradiction. So Γ' is complete.

We are now in the position to apply Robinson's Consistency Theorem 5.2 to the complete $\mathcal{L}'$-theory Γ' and its consistent extensions $\Gamma' \cup \{\phi\}$ and $\Gamma' \cup \{\neg\psi\}$. The conclusion is that $\Gamma' \cup \{\phi, \neg\psi\}$ is consistent, but this contradicts the assumption $\phi \models \psi$. ∎

Our last theorem in this section is due to Beth. We consider a language $\mathcal{L}$, and a new n-place relation symbol P not in $\mathcal{L}$. Let T be an $\mathcal{L} \cup \{P\}$-theory. T is said to *define P implicitly* if for any $\mathcal{L}$-structure $\mathfrak{A}$, there is at most one way to expand $\mathfrak{A}$ to an $\mathcal{L} \cup \{P\}$-structure which is a model of T. This can be said in a different way: let P' be another, new n-place relation symbol and consider the theory $T(P')$ where all P's are replaced by P''s. Then T defines P implicitly if and only if

$$T \cup T(P') \models \forall x_1 \cdots x_n (P(x_1, \ldots, x_n) \leftrightarrow P'(x_1, \ldots, x_n))$$

On the other hand we say that T defines P *explicitly* if there is an $\mathcal{L}$-formula $\varphi(x_1, \ldots, x_n)$ such that

$$T \models \forall x_1 \cdots x_n (P(x_1, \ldots, x_n) \leftrightarrow \varphi(x_1, \ldots, x_n))$$

Clearly, if T defines P explicitly, then T defines P implicitly; the converse is known as *Beth's Definability Theorem.*

Theorem 5.5 (Beth Definability Theorem) *If T defines P implicitly, then T defines P explicitly.*

Proof. Add new constants $c_1, \ldots, c_n$ to the language. By the remark above, we have

$$T \cup T(P') \models P(c_1, \ldots, c_n) \to P'(c_1, \ldots, c_n)$$

By Compactness we can find finite $\Delta \subseteq T$ and $\Delta' \subseteq T(P')$ such that

$$\Delta \cup \Delta' \models P(c_1, \ldots, c_n) \to P'(c_1, \ldots, c_n)$$

Taking conjunctions we can in fact find an $\mathcal{L} \cup \{P\}$-formula ψ, such that $T \models \psi$ and

$$\psi \wedge \psi(P') \models P(c_1, \ldots, c_n) \to P'(c_1, \ldots, c_n)$$

Taking the P's to one side and the P''s to another, we get

$$\psi(P) \wedge P(c_1, \ldots, c_n) \models \psi(P') \to P'(c_1, \ldots, c_n)$$

By the Craig interpolation theorem, there is an $\mathcal{L}$-formula θ, such that $\psi(P) \wedge P(c_1, \ldots, c_n) \models \theta(c_1, \ldots, c_n)$ and $\theta(c_1, \ldots, c_n) \models \psi(P') \to P'(c_1, \ldots, c_n)$. Replacing P' by P again in this second entailment and using that $T \models \psi(P)$, we find that $\theta(c_1, \ldots, c_n)$ is, in T, equivalent to $P(c_1, \ldots, c_n)$; so since the $c_1, \ldots, c_n$ are arbitrary new constants,

$$T \models \forall x_1 \cdots x_n (\theta(x_1, \ldots, x_n) \leftrightarrow P(x_1, \ldots, x_n))$$

and we are done. ∎

Exercise 15 In this section we have proved the Craig Interpolation Theorem from Robinson's Consistency Theorem. Now assume the Craig Interpolation Theorem and use it to give another proof of the Robinson Consistency Theorem.

6 Preservation Theorems

Let T be an $\mathcal{L}$-theory. We consider the class of models of T, as subclass of the class of $\mathcal{L}$-structures; if the class of models of T is closed under certain operations of $\mathcal{L}$-structures, we also say that T is *preserved* under those operations. In this section we shall see that preservation under certain operations can be related to the syntactical structure of axioms for T. For example, the Łos-Tarski Theorem below (Theorem 6.2) says that a theory is preserved under submodels if and only if it is equivalent to a theory consisting solely of sentences of the form $\forall x_1 \cdots \forall x_n \varphi$ with φ quantifier-free.

We shall now define some notions precisely. A *set of axioms* for T is a theory T' which has exactly the same models as T (hence, is logically equivalent to T).

An $\mathcal{L}$-formula of the form $\forall x_1 \cdots \forall x_n \varphi$, with φ quantifier-free, is called a Π_1-formula; a formula $\exists x_1 \cdots \exists x_n \varphi$, with φ quantifier-free, is called a Σ_1-formula. Generally, a formula is Π_{n+1} if it is of form $\forall x_1 \cdots \forall x_m \varphi$ with φ a Σ_n-formula, and dually it is Σ_{n+1} if it is of form $\exists x_1 \cdots \exists x_m \varphi$ with φ a Π_n-formula. One also says $\varphi \in \Sigma_n$, meaning φ is a Σ_n-formula, etc.

Lemma 6.1 *Let T be an $\mathcal{L}$-theory and Δ a set of $\mathcal{L}$-sentences which is closed under disjunction. Suppose that the following holds: whenever $\mathfrak{A}$ is a model of T, and $\mathfrak{B}$ an $\mathcal{L}$-structure which satisfies all sentences in Δ which are true in $\mathfrak{A}$, then $\mathfrak{B}$ is a model of T. Then T has a set of axioms which is a subset of Δ.*

Proof. Let Γ be the set $\{\delta \in \Delta \,|\, T \models \delta\}$. We show that every model of Γ is a model of T, so that Γ is a set of axioms for T.

Suppose $\mathfrak{B}$ is a model of Γ. Let $\Sigma = \{\neg\delta \,|\, \delta \in \Delta, \mathfrak{B} \models \neg\delta\}$. I claim that $T \cup \Sigma$ is consistent. For suppose not, then for some $\neg\delta_1, \ldots, \neg\delta_n \in \Sigma$ we would have $T \models \neg(\neg\delta_1 \wedge \cdots \wedge \neg\delta_n)$, i.e. $T \models \delta_1 \vee \cdots \vee \delta_n$. Since Δ is closed under disjunctions, $(\delta_1 \vee \cdots \vee \delta_n) \in \Gamma$. We obtain a contradiction, since on the one hand $\mathfrak{B}$ is a model of Γ, on the other hand $\mathfrak{B} \models \neg(\delta_1 \vee \cdots \vee \delta_n)$.

Let $\mathfrak{A}$ be a model of $T \cup \Sigma$. Then for $\delta \in \Delta$ we have: if $\mathfrak{A} \models \delta$ then $(\neg\delta) \notin \Sigma$, hence $\mathfrak{B} \not\models \neg\delta$, so $\mathfrak{B} \models \delta$. By the assumption in the lemma, and the fact that $\mathfrak{A}$ is a model of T, we get that $\mathfrak{B}$ is a model of T. Since we started with an arbitrary model $\mathfrak{B}$ of Γ, we see that Γ is a set of axioms for T. ∎

A theory T is said to be *preserved under substructures* if any substructure of a model of T is again a model of T. For example, any substructure of a group is a group; but not every subring of a field is a field; hence the theory of groups is preserved under substructures but the theory of fields isn't.

Theorem 6.2 (Łos-Tarski) *A theory is preserved under substructures if and only if it has a set of axioms consisting of Π_1-sentences.*

Proof. Every Π_1-sentence which is true in $\mathfrak{A}$, is true in every substructure of $\mathfrak{A}$ (check!), so one direction is obvious.

For the other, we shall apply Lemma 6.1, observing that the set of $\mathcal{L}$-sentences which are equivalent to Π_1-sentences, is closed under disjunction.

So let $\mathfrak{A}$ be a model of T, and suppose every Π_1-sentence which is true in $\mathfrak{A}$ also is true in $\mathfrak{B}$. Now consider the theory $T \cup \Delta_{\mathfrak{B}}$ in $\mathcal{L}_{\mathfrak{B}}$. If this theory is inconsistent, then there is a quantifier-free sentence $\phi(b_1, \dots, b_n)$ in $\Delta_{\mathfrak{B}}$ such that $T \cup \{\phi\}$ has no model. This means, that $\mathfrak{A}$, which is a model of T, cannot be expanded to an $\mathcal{L} \cup \{b_1, \dots, b_n\}$-structure which models ϕ; from which it follows that $\mathfrak{A} \models \forall x_1 \cdots \forall x_n \neg\phi(x_1, \dots, x_n)$. But this is a Π_1-sentence, so by assumption we must have $\mathfrak{B} \models \forall x_1 \cdots \forall x_n \neg\phi(x_1, \dots, x_n)$; which is a contradiction, since $\phi(b_1, \dots, b_n)$ is an element of $\Delta_{\mathfrak{B}}$.

We conclude that $T \cup \Delta_{\mathfrak{B}}$ has a model $\mathfrak{C}$, and $\mathfrak{B}$ embeds into $\mathfrak{C}$. Because T is preserved under submodels, $\mathfrak{B}$ is a model of T.

Lemma 6.1 now tells us that T has a set of axioms consisting of Π_1-formulas.
$\blacksquare$

Exercise 16 Strengthen the argument in the proof of Theorem 6.2 to prove: if $\mathfrak{A}$ and $\mathfrak{B}$ are $\mathcal{L}$-structures such that every Π_1-sentence which is true in $\mathfrak{A}$, is true in $\mathfrak{B}$, then $\mathfrak{B}$ is a substructure of an elementary extension of $\mathfrak{A}$.

Exercise 17 Use the previous exercise to prove the following result. Let $T_\forall$ be the set of Π_1-sentences which are consequences of T. Then every model of $T_\forall$ is a substructure of a model of T. Conversely, every substructure of a model of T is a model of $T_\forall$.

Exercise 18 Relativise Theorem 6.2 to: suppose T is a subtheory of T' and for every pair $\mathfrak{A}, \mathfrak{B}$ of models of T with $\mathfrak{B}$ a submodel of $\mathfrak{A}$, we have that if $\mathfrak{A}$ is a model of T', then so is $\mathfrak{B}$. Then there is a set of Π_1-sentences Γ such that $T \cup \Gamma$ is a set of axioms for T'.
Hint: consider the set Δ of sentences ϕ for which there is a Π_1-sentence ψ such that $T \models \phi \leftrightarrow \psi$

Our next theorem characterizes theories which have a set of Π_2-axioms. We say that a theory T is *preserved under directed unions* if for any directed system $(\mathfrak{A}_k)_{k \in K}$ of $\mathcal{L}$-structures as in section 4 such that all homomorphisms f_{kl} (for $k \leq l$) are embeddings: if all $\mathfrak{A}_k$ are models of T, then so is the colimit of the system.

Theorem 6.3 (Chang-Łos-Suszko) *A theory T is preserved under directed unions if and only if T has a set of axioms consisting of Π_2-sentences.*

Proof. First, we do the easy direction: let $\mathfrak{A}$ be the colimit. Then every Π_2-sentence which is true in every $\mathfrak{A}_k$, is also true in $\mathfrak{A}$, for let $\forall x_1 \cdots \forall x_n \exists y_1 \cdots \exists y_m \psi$ be such a sentence, with ψ quantifier-free; and take $a_1, \dots, a_n \in A$. Then by construction of the colimit, since K is directed there is $k \in K$ such that $a_1 = f_k(a'_1), \dots, a_n = f_k(a'_n)$ for suitable $a'_1, \dots, a'_n \in A_k$. Then in A_k we can find $b_1, \dots, b_m$ such that $\mathfrak{A}_k \models \psi(a'_1, \dots, a'_n, b_1, \dots, b_m)$. Because $f_k : \mathfrak{A}_k \to \mathfrak{A}$ is an embedding, we have

$$\mathfrak{A} \models \psi(a_1, \dots, a_n, f_k(b_1), \dots, f_k(b_m))$$

Hence, $\mathfrak{A} \models \forall x_1 \cdots \forall x_n \exists y_1 \cdots \exists y_m \psi$.

For the other direction, we use again Lemma 6.1, observing that the disjunction of two Π_2-sentences is equivalent to a Π_2-sentence. So, let $\mathfrak{A}$ be a model of T, T be preserved under directed unions, and $\mathfrak{B}$ satisfies every Π_2-sentence which is true in $\mathfrak{A}$. We construct a directed system of embeddings, in fact a chain:

$$\mathfrak{B} = \mathfrak{B}_0 \to \mathfrak{A}_1 \to \mathfrak{B}_1 \to \mathfrak{A}_2 \to \mathfrak{B}_2 \to \cdots$$

with the properties:

i) $\mathfrak{A}_n \equiv \mathfrak{A}$

ii) the composed embedding $\mathfrak{B}_n \to \mathfrak{B}_{n+1}$ is elementary;

iii) every Π_1-sentence in the language $\mathcal{L}_{\mathfrak{B}_n}$ which is true in $\mathfrak{B}_n$, is also true in $\mathfrak{A}_{n+1}$ (regarding $\mathfrak{A}_{n+1}$, via the embedding, as an $\mathcal{L}_{\mathfrak{B}_n}$-structure in the obvious way)

The construction will proceed inductively, and we shall assume each of the conditions i)–iii) as induction hypothesis, as we go along. We start by putting $\mathfrak{B}_0 = \mathfrak{B}$; this was easy.

If $\mathfrak{B}_n$ has been constructed, in order to make $\mathfrak{A}_{n+1}$ we consider the theory $\mathrm{Th}(\mathfrak{A}) \cup \Pi_{\mathfrak{B}_n}$, where $\mathrm{Th}(\mathfrak{A})$ is the set of all $\mathcal{L}$-sentences which are true in $\mathfrak{A}$, and $\Pi_{\mathfrak{B}_n}$ is the set of Π_1-sentences in the language $\mathcal{L}_{\mathfrak{B}_n}$ which are true in $\mathfrak{B}_n$ (note, that every quantifier-free sentence is trivially a Π_1-sentence, so $\Delta_{\mathfrak{B}_n} \subseteq \Pi_{\mathfrak{B}_n}$).

Suppose this theory is inconsistent; then for some sentence

$$\psi = \forall y_1 \cdots \forall y_n \varphi(b_1, \ldots, b_m, y_1, \ldots, y_n)$$

with $\psi \in \Pi_{\mathfrak{B}_n}$, we have that $\mathrm{Th}(\mathfrak{A}) \cup \{\psi\}$ has no model. As before, we see then that

$$\mathfrak{A} \models \forall x_1 \cdots \forall x_m \exists y_1 \cdots \exists y_n \neg\varphi$$

Since this is a Π_2-sentence, it must, by assumption on $\mathfrak{A}$ and $\mathfrak{B}$, be true in $\mathfrak{B}$, which by induction hypothesis is an elementary submodel of $\mathfrak{B}_n$; but this is in contradiction with the fact that ψ is true in $\mathfrak{B}_n$ (check!).

Therefore, $\mathrm{Th}(\mathfrak{A}) \cup \Pi_{\mathfrak{B}_n}$ is consistent, and we let $\mathfrak{A}_{n+1}$ be a model of it. Clearly, $\mathfrak{A}_{n+1} \equiv \mathfrak{A}$, and condition iii) also holds, for $\mathfrak{B}_n$ and $\mathfrak{A}_{n+1}$.

If $\mathfrak{A}_{n+1}$ has been constructed, we consider the theory $E(\mathfrak{B}_n) \cup \Delta_{\mathfrak{A}_{n+1}}$ in the language $\mathcal{L}_{\mathfrak{A}_{n+1}}$ (since $\mathfrak{B}_n$ is a submodel of $\mathfrak{A}_{n+1}$ we can take $E(\mathfrak{B}_n)$ to be a theory in this language).

Suppose this theory is inconsistent. Then for some formulas $\phi \in E(\mathfrak{B}_n)$ and $\psi \in \Delta_{\mathfrak{A}_{n+1}}$ we have $\phi \models \neg\psi$. Now ψ is of form $\psi(b_1, \ldots, b_k, a_1, \ldots, a_l)$ with $a_1, \ldots, a_l \in A_{n+1} \backslash B_n$ and $b_1, \ldots, b_k \in B_n$. We see, that

$$\mathfrak{B}_n \models \forall x_1, \ldots, x_l \neg\psi(b_1, \ldots, b_k, x_1, \ldots, x_l)$$

which is a Π_1-sentence in the language $\mathcal{L}_{\mathfrak{B}_n}$. By induction hypothesis iii) then,

$$\mathfrak{A}_{n+1} \models \forall x_1, \dots, x_l \neg\psi(b_1, \dots, b_k, x_1, \dots, x_l)$$

which is, in a now familiar way, a contradiction.

We let $\mathfrak{B}_{n+1}$ be a model of $E(\mathfrak{B}_n) \cup \Delta_{\mathfrak{A}_{n+1}}$. Then the embedding $\mathfrak{B}_n \to \mathfrak{B}_{n+1}$ is elementary, which is the only induction hypothesis we have to check at this stage. This finishes the construction of the sequence

$$\mathfrak{B}_0 \to \mathfrak{A}_1 \to \mathfrak{B}_1 \to \mathfrak{A}_2 \to \cdots$$

In order to finish the argument: let $\mathfrak{C}$ be the colimit of this chain. Then $\mathfrak{C}$ is equally the colimit of the chain

$$\mathfrak{A}_1 \to \mathfrak{A}_2 \to \cdots$$

which is a chain of models of T; so since T is preserved by directed unions, $\mathfrak{C}$ is a model of T. On the other hand, $\mathfrak{C}$ is also colimit of the chain

$$\mathfrak{B}_0 \to \mathfrak{B}_1 \to \mathfrak{B}_2 \to \cdots$$

which is a chain of elementary embeddings; by the Elementary System Lemma (4.2), $\mathfrak{B} \to \mathfrak{C}$ is an elementary embedding. It follows that $\mathfrak{B}$ is a model of T, which was to be proved. A final application of Lemma 6.1 now yields the result. ∎

Exercise 19 Use Theorem 6.3 to prove the following equivalence: T has a set of axioms consisting of Π_2-sentences if and only if the following property holds: whenever $\mathfrak{A}$ is an $\mathcal{L}$-structure and for every finite subset $A' \subseteq A$ there is a substructure $\mathfrak{B}$ of $\mathfrak{A}$ such that $A' \subseteq B$ and $\mathfrak{B}$ is a model of T, then $\mathfrak{A}$ is a model of T.

Exercise 20 Suppose that T satisfies the following property: whenever $\mathfrak{A}$ is a model of T, and $\mathfrak{B}, \mathfrak{C}$ are substructures of $\mathfrak{A}$ which are also models of T, and $B \cap C \neq \emptyset$, then $\mathfrak{B} \cap \mathfrak{C}$ (the substructure with domain $B \cap C$) is also a model of T.

Show that T has a set of axioms consisting of Π_2-sentences.

I mention one more preservation theorem, without proof. Call a formula *positive* if it does not contain the symbols $\neg$ or $\to$.

A theory T is said to be *preserved under homomorphic images* if whenever $\mathfrak{A}$ is a model of T and $f : \mathfrak{A} \to \mathfrak{B}$ a surjective homomorphism, then $\mathfrak{B}$ is a model of T.

Theorem 6.4 *A theory is preserved under homomorphic images if and only if it has a set of positive axioms.*

Exercise 21 Prove the easy part of theorem 6.4, that is: every positive sentence is preserved under homomorphic images.

7 Filters and Ultraproducts

Let I be a set. A *filter* over I is a subset $\mathcal{U}$ of $\mathcal{P}(I)$, the powerset of I, satisfying the properties:

i) $I \in \mathcal{U}$;

ii) if $U \in \mathcal{U}$ and $U \subseteq V \subseteq I$, then $V \in \mathcal{U}$;

iii) $U, V \in \mathcal{U}$ implies $U \cap V \in \mathcal{U}$, for any $U, V \subseteq I$;

iv) $\emptyset \notin \mathcal{U}$.

Examples. For every nonempty subset J of I, there is the *principal filter*

$$\mathcal{U}_J = \{A \subseteq I \mid J \subseteq A\}$$

If I is infinite, there is the *Fréchet filter* or *cofinite filter*

$$\mathcal{U}_F = \{A \subseteq I \mid I \backslash A \text{ is finite}\}$$

Exercise 22 Let $\mathcal{A}$ be a set of nonempty subsets of I, such that for every finite subcollection $\{A_1, \ldots, A_n\}$ of $\mathcal{A}$, the intersection

$$A_1 \cap \cdots \cap A_n$$

is nonempty. Show, that the collection

$$\{U \subseteq I \mid \exists A_1 \cdots A_n \in \mathcal{A}(A_1 \cap \cdots \cap A_n \subseteq U)\}$$

is a filter (it is said to be the filter *generated by* $\mathcal{A}$).

Exercise 23 (Filters and Congruence Relations) A *congruence relation* on $\mathcal{P}(I)$ is an equivalence relation $\sim$ on $\mathcal{P}(I)$, such that $\sim$ satisfies the two properties:

$$\begin{array}{l} A \sim B \ \Rightarrow C \cap A \sim C \cap B \\ A \sim B \ \Rightarrow I \backslash A \sim I \backslash B \end{array}$$

Show: every filter $\mathcal{U}$ over I determines a congruence relation on $\mathcal{P}(I)$ by:

$$A \sim B \ \text{ iff } \ (((I \backslash A) \cup B) \cap (A \cup (I \backslash B))) \in \mathcal{U}$$

Conversely, show that for every congruence relation on $\mathcal{P}(I)$ such that $\emptyset \not\sim I$, $\{A \mid A \sim I\}$ is a filter over I.

Let $\mathcal{U}$ be a filter over the set I, and suppose that $(A_i)_{i \in I}$ is a family of sets. We define the *reduced product* modulo $\mathcal{U}$, written $\Pi_{\mathcal{U}} A_i$, as follows. The set $\Pi_{\mathcal{U}} A_i$ will be

$$(\prod_{i \in I} A_i)/\sim$$

where $\prod_{i\in I} A_i$ is the product of all the sets A_i (that is the set of all I-indexed sequences $(x_i)_{i\in I}$ such that $x_i \in A_i$ for all i), and $\sim$ is the equivalence relation given by:

$$(x_i)_{i\in I} \sim (y_i)_{i\in I} \text{ iff } \{i \in I \mid x_i = y_i\} \in \mathcal{U}$$

This is an equivalence relation: since $I \in \mathcal{U}$ we have $(x)_i \sim (x)_i$; and obviously, the relation is symmetric. If $(x)_i \sim (y)_i$ and $(y)_i \sim (z_i)$, then if $U = \{i \in I \mid x_i = y_i\}$ and $V = \{i \in I \mid y_i = z_i\}$ then clearly $U \cap V \subseteq \{i \in I \mid x_i = z_i\}$, so $(x)_i \sim (z)_i$.

Exercise 24 Show that if $\mathcal{U}$ is the principal filter $\mathcal{U}_J$ for $J \subseteq I$, then $\Pi_{\mathcal{U}} A_i$ is in bijective correspondence with $\prod_{i\in J} A_i$.

An *ultrafilter* is a maximal filter, i.e. a filter that cannot be extended to a larger filter.

Exercise 25 Show that for an ultrafilter $\mathcal{U}$ over I:

a) $U \in \mathcal{U}$ or $I\backslash U \in \mathcal{U}$, for every $U \subseteq I$;

b) if $U \cup V \in \mathcal{U}$, then $U \in \mathcal{U}$ or $V \in \mathcal{U}$.

Exercise 26 Show that in the correspondence of Exercise 23, an ultrafilter corresponds to a congruence relation with exactly two classes.

Exercise 27 Show:

a) If an ultrafilter is a principal filter $\mathcal{U}_J$, J is a singleton set;

b) if an ultrafilter is not principal, it contains every cofinite set.

Lemma 7.1 *Every filter is contained in an ultrafilter.*

Exercise 28 Prove Lemma 7.1, using Zorn's Lemma.

Exercise 29 Refine Lemma 7.1 to the following statement: suppose $\mathcal{U}$ is a filter, and $\mathcal{A}$ a collection of subsets of I such that no element of $\mathcal{U}$ is a subset of a finite union of elements of $\mathcal{A}$. Then $\mathcal{U}$ is contained in an ultrafilter that is disjoint from $\mathcal{A}$.

If $\mathcal{U}$ is an ultrafilter, the reduced product $\prod_{\mathcal{U}} A_i$ is called an *ultraproduct*.

Exercise 30 Let $I = \mathbb{N}$, and $A_i = \{0, 1\}$ for all $i \in \mathbb{N}$. Show that for every ultrafilter $\mathcal{U}$ over $\mathbb{N}$, the ultraproduct $\prod_{\mathcal{U}} A_i$ has cardinality 2.

Exercise 31 Let again $I = \mathbb{N}$, and $A_i = \{0, \ldots, i\}$. Let $\mathcal{U}$ be an ultrafilter over $\mathbb{N}$. Show:

a) If $\mathcal{U}$ is principal, $\prod_{\mathcal{U}} A_i$ is finite;

b) if $\prod_{\mathcal{U}} A_i$ is non-principal, $\prod_{\mathcal{U}} A_i$ is uncountable.

We shall now extend the reduced product construction to $\mathcal{L}$-structures.

First, if we are given, for each i, an n-ary function $f_i : (A_i)^n \to A_i$, we have a function $f : (\prod_{i\in I} A_i)^n \to \prod_{i\in I} A_i$ defined by

$$f((a_i^1)_i, \dots, (a_i^n)_i) = (f_i(a_i^1, \dots, a_i^n)_i$$

Now if $\mathcal{U}$ is a filter over I and $(a_i^1)_i \sim (b_i^1)_i, \dots, (a_i^n)_i \sim (b_i^n)_i$ in the equivalence relation defining the reduced product $\prod_{\mathcal{U}} A_i$, then for each $k = 1, \dots, n$ we have

$$\{i \mid a_i^k = b_i^k\} \in \mathcal{U}$$

so since $\mathcal{U}$ is closed under finite intersections,

$$\{i \mid \forall k (1 \leq k \leq n \to a_i^k = b_i^k)\} \in \mathcal{U}$$

This clearly implies $\{i \mid f_i(a_i^1, \dots, a_i^n) = f_i(b_i^1, \dots, b_i^n)\} \in \mathcal{U}$, so

$$f((a_i^1)_i, \dots, (a_i^n)_i) \sim f((b_i^1)_i, \dots, (b_i^n)_i)$$

Hence f determines a function: $(\prod_{\mathcal{U}} A_i)^n \to \prod_{\mathcal{U}} A_i$, which we shall also denote by f. By the above, we may put

$$f([(a_i^1)_i], \dots, [(a_i^n)_i]) = [(f_i(a_i^1, \dots, a_i^n))_i]$$

Secondly, if in every A_i an element c_i is chosen, this determines an obvious element $[(c_i)_i]$ of $\prod_{\mathcal{U}} A_i$.

Thirdly, if for every i we have an n-ary relation $R_i \subseteq (A_i)^n$, we define an n-ary relation R on $\prod_{\mathcal{U}} A_i$ by

$$\begin{aligned} R = \{([(a_i^1)_i], \dots, [(a_i^n)_i]) \mid \\ \{i \mid (a_i^1, \dots, a_i^n) \in R_i\} \in \mathcal{U}\} \end{aligned}$$

Exercise 32 Show that R is well-defined.

In this way, we see that if every A_i is the domain of an $\mathcal{L}$-structure $\mathfrak{A}_i$, $\prod_{\mathcal{U}} A_i$ is the domain of an $\mathcal{L}$-structure $\prod_{\mathcal{U}} \mathfrak{A}_i$.

Exercise 33 If t is an $\mathcal{L}$-term with variables $x_1, \dots, x_n$ and $\mathfrak{A} = \prod_{\mathcal{U}} \mathfrak{A}_i$, and $[(a_i^1)_i], \dots, [(a_i^n)_i]$ are elements of $\prod_{\mathcal{U}} A_i$, then

$$t^{\mathfrak{A}}([(a_i^1)_i], \dots, [(a_i^n)_i]) = [(t^{\mathfrak{A}_i}(a_i^1, \dots, a_i^n))_i]$$

From this exercise, and the definition of the interpretation of relation symbols of $\mathcal{L}$ in the reduced product, we see:

Proposition 7.2 *Let $\mathfrak{A} = \prod_{\mathcal{U}} \mathfrak{A}_i$ as above. Then every atomic $\mathcal{L}$-formula φ has the following property: if φ has n free variables $x_1, \dots, x_n$, then for every n-tuple $([(a_i^1)_i], \dots, [(a_i^n)_i])$ of elements of $\prod_{\mathcal{U}} A_i$,*

$$\begin{aligned} \mathfrak{A} \models \varphi([(a_i^1)_i], \dots, [(a_i^n)_i]) \Leftrightarrow \\ \{i \mid \mathfrak{A}_i \models \varphi(a_i^1, \dots, a_i^n)\} \in \mathcal{U} \end{aligned}$$

Lemma 7.3 *The collection Φ of $\mathcal{L}$-formulas which have the property in Proposition 7.2, is closed under conjunction and existential quantification.*

Proof. Let us drop the constants from the notation. Since always $U \cap V \in \mathcal{U}$ if and only if both $U \in \mathcal{U}$ and $V \in \mathcal{U}$, we have for $\varphi, \psi \in \Phi$:

$$\begin{array}{l} \mathfrak{A} \models \varphi \wedge \psi \Leftrightarrow \\ \{i \,|\, \mathfrak{A}_i \models \varphi\} \in \mathcal{U} \text{ and } \{i \,|\, \mathfrak{A}_i \models \psi\} \in \mathcal{U} \Leftrightarrow \\ \{i \,|\, \mathfrak{A}_i \models \varphi \wedge \psi\} \in \mathcal{U} \end{array}$$

so $\varphi \wedge \psi$ is in Φ.

For existential quantification, suppose $\varphi(x) \in \Phi$ (again, φ may have more free variables, which we suppress). If $\mathfrak{A} \models \exists x\varphi$ so $\mathfrak{A} \models \varphi([(a_i)_i])$ then by assumption $\{i \,|\, \mathfrak{A}_i \models \varphi(a_i)\} \in \mathcal{U}$ so certainly $\{i \,|\, \mathfrak{A}_i \models \exists x\varphi\} \in \mathcal{U}$ since this is a bigger set.

Conversely if $\{i \,|\, \mathfrak{A}_i \models \exists x\varphi\} = U \in \mathcal{U}$, pick for each $i \in U$ an $a_i \in A_i$ such that $\mathfrak{A}_i \models \varphi(a_i)$. For $i \notin U$, let $a_i \in A_i$ arbitrary. Then $(a_i)_i$ satisfies: $\{i \,|\, \mathfrak{A}_i \models \varphi(a_i)\} \in \mathcal{U}$, so by assumption $\mathfrak{A} \models \varphi([(a_i)_i])$; hence $\mathfrak{A} \models \exists x\varphi$. ∎

Theorem 7.4 (Łos; Fundamental Theorem for Ultraproducts)

a) *If $\mathcal{U}$ is an ultrafilter, every $\mathcal{L}$-formula has the property of proposition 7.2;*

b) *hence, for $\mathcal{L}$-sentences ϕ,*

$$\mathfrak{A} \models \phi \Leftrightarrow \{i \in I \,|\, \mathfrak{A}_i \models \phi\} \in \mathcal{U}$$

Proof. a). From Proposition 7.2 and Lemma 7.3 we know that the collection Φ of formulas satisfying the mentioned property, contains all atomic formulas and is closed under $\wedge$ and $\exists$. Because every formula is logically equivalent to a formula that contains only $\{\wedge, \exists, \neg\}$, it suffices to show that when $\mathcal{U}$ is an ultrafilter, the collection Φ is closed under $\neg$.

But we know for an ultrafilter $\mathcal{U}$, that $U \in \mathcal{U}$ if and only if $I \backslash U \notin \mathcal{U}$, so $\varphi \in \Phi$ implies:

$$\begin{array}{rl} \mathfrak{A} \models \neg\varphi & \Leftrightarrow \\ \{i \,|\, \mathfrak{A}_i \models \varphi\} \notin \mathcal{U} & \Leftrightarrow \\ \{i \,|\, \mathfrak{A}_i \models \neg\varphi\} \in \mathcal{U} & \end{array}$$

so $\neg\varphi \in \Phi$.

b) This is an immediate consequence of a). ∎

Exercise 34 a) Suppose $\mathfrak{A}_i = \mathfrak{A}$ for all $i \in I$, and $\mathcal{U}$ is an ultrafilter over I. The ultraproduct $\prod_{\mathcal{U}} \mathfrak{A}$ is then called an *ultrapower* of $\mathfrak{A}$. Show, that the map $\mathfrak{A} \to \prod_{\mathcal{U}} \mathfrak{A}$, given by $a \mapsto [(a)_i]$, is an elementary embedding.

b) Suppose $(\mathfrak{A}_i)_i$, $(\mathfrak{B}_i)_i$ are two I-indexed collections of $\mathcal{L}$-structures, and let $f_i : \mathfrak{A}_i \to \mathfrak{B}_i$ be a homomorphism for all i. Given an ultrafilter $\mathcal{U}$ over I, define a homomorphism $f : \prod_{\mathcal{U}} \mathfrak{A}_i \to \prod_{\mathcal{U}} \mathfrak{B}_i$ and show: if every f_i is an elementary embedding, so is f.

We turn now to some applications of the Fundamental Theorem for Ultraproducts, and the Compactness Theorem.

Let $\mathcal{K}$ be a class of $\mathcal{L}$-structures. We say that the class $\mathcal{K}$ is *definable* if there is an $\mathcal{L}$-sentence ϕ such that $\mathcal{K}$ is the class of models of ϕ. $\mathcal{K}$ is *elementary* if there is an $\mathcal{L}$-theory T such that $\mathcal{K}$ is the class of models of T. The following theorem characterizes definable and elementary classes.

Theorem 7.5 *a) A class $\mathcal{K}$ of $\mathcal{L}$-structures is definable if and only if both $\mathcal{K}$ and its complement are elementary;*

b) $\mathcal{K}$ is elementary if and only if $\mathcal{K}$ is closed under elementary equivalence and ultraproducts.

Proof. a) If $\mathcal{K}$ is the class of models of ϕ, then the complement of $\mathcal{K}$ is the class of models of $\neg\phi$, so both classes are elementary. Conversely if $\mathcal{K}$ is the class of models of an $\mathcal{L}$-theory T, and its complement is the class of models of another theory T', then clearly $T \cup T'$ is inconsistent; so, assuming T and T' to be closed under conjunctions, by Compactness there are sentences $\phi \in T$, $\psi \in T'$ such that $\phi \wedge \psi$ has no model. Clearly then, every model of ϕ is a model of T (and conversely, since $\phi \in T$), and every model of ψ is a model of T' and vice versa. So $\mathcal{K}$ is definable.

b) One direction follows directly from the Fundamental Theorem: if $\mathcal{K}$ is the class of models of T, $(\mathfrak{A}_i)_{i\in I}$ is a family of elements of $\mathcal{K}$ and $\mathcal{U}$ is an ultrafilter over I, then the ultraproduct $\prod_{\mathcal{U}} \mathfrak{A}_i$ is a model of T, hence in $\mathcal{K}$. And obviously, $\mathcal{K}$ is closed under elementary equivalence.

Conversely, suppose $\mathcal{K}$ is closed under ultraproducts and elementary equivalence. Let T the set of $\mathcal{L}$-sentences which are true in all elements of $\mathcal{K}$. We show that $\mathcal{K}$ is the class of models of T. So let $\mathfrak{B}$ be a model of T. Let $\Delta = \{\delta_1, \ldots, \delta_n\}$ be a finite subset of $\mathrm{Th}(\mathfrak{B})$ (the set of $\mathcal{L}$-sentences true in $\mathfrak{B}$). Then there is an element of $\mathcal{K}$ in which every element of Δ is true; for otherwise, $\neg(\delta_1 \wedge \cdots \wedge \delta_n)$ would be an element of T. But this cannot be, since $\mathfrak{B}$ is a model of T. Choose for every such Δ a model $\mathfrak{A}_\Delta$ of Δ, from $\mathcal{K}$. So, I is the set of all finite subsets of $\mathrm{Th}(\mathfrak{B})$. For every $\phi \in \mathrm{Th}(\mathfrak{B})$ let U_ϕ be $\{\Delta \,|\, \mathfrak{A}_\Delta \models \phi\}$. Then the collection $\{U_\phi \,|\, \phi \in \mathrm{Th}(\mathfrak{B})\}$ is closed under finite intersections, hence generates a filter $\mathcal{U}$ which can be extended to an ultrafilter $\mathcal{F}$; let $\mathfrak{C}$ be the ultraproduct $\prod_{\mathcal{F}} \mathfrak{A}_\Delta$. Just as in the ultraproduct proof of the Compactness theorem, one sees that $\mathfrak{C}$ is a model of $\mathrm{Th}(\mathfrak{B})$, hence is elementarily equivalent to $\mathfrak{B}$. Moreover, $\mathfrak{C}$ is an element of $\mathcal{K}$ since $\mathcal{K}$ is closed under ultraproducts; so $\mathfrak{B}$ is an element of $\mathcal{K}$ since $\mathcal{K}$ is closed under elementary equivalence. Since we started with an arbitrary model $\mathfrak{B}$ of T, we see that $\mathcal{K}$ contains every model of T. ∎

In a way quite similar to the proof of Theorem 7.5b), we can use the Fundamental Theorem for Ultraproducts to prove the Compactness Theorem.

Suppose Γ is a set of $\mathcal{L}$-sentences such that every finite subset Δ of Γ has a model $\mathfrak{A}_\Delta$. We let $I = \{\Delta \subseteq \Gamma \,|\, \Delta \text{ finite}\}$, and for each $\phi \in \Gamma$ let $U_\phi = \{\Delta \in I \,|\, \phi \in \Delta\}$. Since $\{\phi, \psi\} \in U_\phi \cap U_\psi$ for all $\phi, \psi \in \Gamma$, the collection $\{U_\phi \,|\, \phi \in \Gamma\}$ is contained in an ulttrafilter $\mathcal{F}$ over I.

Exercise 35 Prove yourself that the ultraproduct $\prod_{\mathcal{F}} \mathfrak{A}_{\Delta}$ is a model of Γ.

The following exercise gives some examples of the use of Theorem 7.5.

Exercise 36 (Examples) a) Let $I = \mathbb{N}$, $A_i = \{0, \ldots, i\}$ and $\mathcal{F}$ a nonprincipal ultrafilter over I. Use the ultraproduct $\prod_{\mathcal{F}} A_i$ to show:

i) The class of cyclic abelian groups is not elementary;

ii) the class of well-founded linear orders is not elementary (a linear order is *well-founded* if there is no infinite chain $a_0 > a_1 > a_2 > \cdots$);

iii) the class of connected graphs is not elementary (a graph is *connected* if for every pair of points a, b there is a finite path $a = a_0 \to a_1 \to \cdots \to a_n = b$)

b) Let $I \subset \mathbb{N}$ be the set of primes, $A_i = \{0, \ldots, i-1\}$, $\mathcal{F}$ a nonprincipal ultrafilter over I. Show that $\prod_{\mathcal{F}} A_i$ is a field of characteristic zero, and conclude that the class of fields of characteristic $\neq 0$ is not elementary;

c) An abelian group G is *divisible* if for every $n \geq 2$,

$$G \models \forall x \exists y (\underbrace{y + \cdots + y}_{n \text{ times}} = x)$$

Construct yourself an ultraproduct example to show that the class of non-divisible abelian groups is not elementary (hence, the theory of divisible groups not finitely axiomatizable).

A strengthening of Exercise 13 is given by the following theorem, which we state without proof.

Theorem 7.6 (Keisler-Shelah) *Let $\mathfrak{A}$ and $\mathfrak{B}$ be $\mathcal{L}$-structures. Then $\mathfrak{A} \equiv \mathfrak{B}$ if and only if there exist a set I and an ultrafilter $\mathcal{F}$ over I such that the ultrapowers $\prod_{\mathcal{F}} \mathfrak{A}$ and $\prod_{\mathcal{F}} \mathfrak{B}$ are isomorphic.*

Exercise 37 Use Theorem 7.6 to obtain the following refinement of Theorem 7.5: A class $\mathcal{K}$ of $\mathcal{L}$-structures is elementary if and only if $\mathcal{K}$ satisfies the properties:

a) $\mathcal{K}$ is closed under isomorphism and under ultraproducts;

b) Whenever some ultrapower of $\mathfrak{A}$ is in $\mathcal{K}$, then $\mathfrak{A}$ is in $\mathcal{K}$.

I finish this chapter with a theorem whose proof is an application of the idea of ultrafilters. It is an important theorem in the field of *infinite combinatorics* and finds applications in advanced Model Theory, though possibly not in this course.

Theorem 7.7 (Ramsey) *Let I be an infinite set. Write $\mathcal{P}_n(I)$ for the collection of subsets of I with exactly n elements. If $\mathcal{P}_n(I) = A \cup B$ then there is an infinite subset J of I such that either $\mathcal{P}_n(J) \subseteq A$ or $\mathcal{P}_n(J) \subseteq B$.*

Proof. If I is uncountable, we may take any countable subset of I, so without loss of generality we may assume $I = \mathbb{N}$. The theorem is trivial for $n = 1$ (check!), so assume $n > 1$. For a finite subset α of $\mathbb{N}$ we write $\alpha < k$ for $\forall m \in \alpha(m < k)$.

Let $\mathcal{F}$ be a nonprincipal ultrafilter over $\mathbb{N}$. We define, for $1 \leq r \leq n$, sets $A^r, B^r \subseteq \mathcal{P}_r(\mathbb{N})$ as follows: let $A^n = A$ and $B^n = B$. If, for $1 \leq r < n$, A^{r+1} and B^{r+1} have been defined we put

$$\begin{array}{rcl} A^r & = & \{\alpha \in \mathcal{P}_r(\mathbb{N}) \mid \{m > \alpha \mid \alpha \cup \{m\} \in A^{r+1}\} \in \mathcal{F}\} \\ B^r & = & \{\alpha \in \mathcal{P}_r(\mathbb{N}) \mid \{m > \alpha \mid \alpha \cup \{m\} \in B^{r+1}\} \in \mathcal{F}\} \end{array}$$

Then if $\mathcal{P}_{r+1}(\mathbb{N}) \subseteq A^{r+1} \cup B^{r+1}$, also $\mathcal{P}_r(\mathbb{N}) \subseteq A^r \cup B^r$: if $\alpha \in \mathcal{P}_r(\mathbb{N})$, $\alpha \notin A^r$ then $\{m > \alpha \mid \alpha \cup \{m\} \in A^{r+1}\} \notin \mathcal{F}$. Since $\mathcal{F}$ is a nonprincipal ultrafilter and $\{m \mid \alpha \not< m\}$ is finite (so not in $\mathcal{F}$), we must have $\{m > \alpha \mid \alpha \cup \{m\} \in B^{r+1}\} \in \mathcal{F}$.

In particular we have $\{\{n\} \mid n \in \mathbb{N}\} \subseteq A^1 \cup B^1$ and therefore either $\{n \mid \{n\} \in A^1\} \in \mathcal{F}$ or $\{n \mid \{n\} \in B^1\} \in \mathcal{F}$; assume $\{n \mid \{n\} \in A^1\} \in \mathcal{F}$; the other case is dealt with in a symmetric way.

We define J as follows. Let j_0 be the least n such that $\{n\} \in A^1$. Inductively, suppose $j_0 < j_1 < \cdots < j_k$ have been defined such that for all $1 \leq r \leq n$ and all $\alpha \in \mathcal{P}_r(\{j_0, \ldots, j_k\})$, $\alpha \in A^r$.

Then for all $1 \leq r < n$ and all $\alpha \in \mathcal{P}_r(\{j_0, \ldots, j_k\})$ we have

$$U_{r,\alpha} = \{m > \alpha \mid \alpha \cup \{m\} \in A^{r+1}\} \in \mathcal{F}$$

Now there are only finitely many pairs (r, α) with $1 \leq r < n$ and $\alpha \in \mathcal{P}_r(\{j_0, \ldots, j_k\})$, so the set

$$\{n \mid \{n\} \in A^1\} \cap \bigcap_{\substack{1 \leq r < n, \\ \alpha \in \mathcal{P}_r(\{j_0, \ldots, j_k\})}} U_{r,\alpha}$$

is an element of $\mathcal{F}$. Let j_{k+1} be the least element in this set which is $> j_k$. Convince yourself that again, for each $1 \leq r \leq n$ and $\alpha \in \mathcal{P}_r(\{j_0, \ldots, j_{k+1}\})$, $\alpha \in A^r$.

The set $J = \{j_0, j_1, \ldots\}$ thus constructed, has the property that $\mathcal{P}_n(J) \subseteq A$. ■

8 Quantifier Elimination and Model Completeness

We say that an $\mathcal{L}$-theory T *admits elimination of quantifiers*, or T *has quantifier elimination*, if for every $\mathcal{L}$-formula $\varphi(x_1, \ldots, x_n)$ there is a quantifier-free $\mathcal{L}$-formula ψ with at most the variables $x_1, \ldots, x_n$ free, such that

$$T \models \forall x_1 \cdots x_n(\varphi(\vec{x}) \leftrightarrow \psi(\vec{x}))$$

In particular, every $\mathcal{L}$-sentence is, in T, equivalent to a quantifier-free sentence. Therefore, if T has quantifier elimination and $T \models \phi$ or $T \models \neg\phi$ for every quantifier-free $\mathcal{L}$-sentence ϕ, then T is complete.

Exercise 38 a) Show: T has quantifier elimination if and only if every formula ϕ of form $\exists x A$, where A is quantifier-free, is equivalent to a quantifier-free formula having at most the same free variables;

b) show that we can simplify further to: ϕ of form $\exists x A$, where A is a conjunction of atomic formulas and negations of atomic formulas.

[Hint: for a), use induction on the number of quantifiers; for b), by disjunctive normal form, every formula of form in a) is equivalent to a disjunction of formulas of the form in b)]

An immediate consequence of the definition is the following proposition:

Proposition 8.1 *Suppose T has quantifier elimination. Then for any two models $\mathfrak{A}, \mathfrak{B}$ of T: if $\mathfrak{A}$ is a substructure of $\mathfrak{B}$ then $\mathfrak{A}$ is an elementary substructure of $\mathfrak{B}$.*

Proof. Let $\varphi(x_1, \ldots, x_n)$ an $\mathcal{L}$-formula and $a_1, \ldots, a_n \in A$. Let $\psi(\vec{x})$ be quantifier-free such that $T \models \forall\vec{x}(\varphi(\vec{x} \leftrightarrow \psi(\vec{x}))$. Then $\mathfrak{A} \models \varphi(a_1, \ldots, a_n)$ iff $\mathfrak{A} \models \psi(a_1, \ldots, a_n)$ (since $\mathfrak{A}$ is a model of T), iff $\mathfrak{B} \models \psi(a_1, \ldots, a_n)$ (since $\mathfrak{A}$ is a substructure of $\mathfrak{B}$ and ψ is quantifier-free), iff $\mathfrak{B} \models \varphi(a_1, \ldots, a_n)$ (since $\mathfrak{B}$ is a model of T). ∎

The property stated in Proposition 8.1 is weaker than quantifier elimination, and is called *model completeness* of the theory T: a theory T is model complete if for any two models $\mathfrak{A}, \mathfrak{B}$ of T, $\mathfrak{A} \subseteq \mathfrak{B}$ implies $\mathfrak{A} \preceq \mathfrak{B}$.

Examples. The theory of fields (or even fields of characteristic zero) is not model complete: $\mathbb{Q}$ is a subfield of $\mathbb{R}$ but not elementary, as we have seen. The theory of torsion-free divisible abelian groups is model complete.

Exercise 39

a) The following two statements are equivalent for an $\mathcal{L}$-theory T:

   i) For every $\mathcal{L}$-formula $\varphi(\vec{x})$ there is a Σ_1-formula $\psi(\vec{x})$ such that $T \models \forall\vec{x}(\varphi(\vec{x}) \leftrightarrow \psi(\vec{x}))$;

ii) for every $\mathcal{L}$-formula $\varphi(\vec{x})$ there is a Π_1-formula $\psi(\vec{x})$ such that $T \models \forall\vec{x}(\varphi(\vec{x}) \leftrightarrow \psi(\vec{x}))$

b) Show that T is model complete if and only if T satisfies the equivalent conditions of a).
[Hint: use the refinement of the Łos-Tarski Theorem in Exercise 18]

Exercise 40 (Robinson's Test) Show that T is model complete if and only if for every embedding $\mathfrak{A} \subseteq \mathfrak{B}$ of models of T, every Σ_1-sentence of the language $\mathcal{L}_{\mathfrak{A}}$ which holds in $\mathfrak{B}$, also holds in $\mathfrak{A}$.

Exercise 41 Use the Chang-Los-Suszko Theorem (Theorem 6.3) and the Elementary System Lemma (Lemma 4.2) to show that every model complete theory has a set of Π_2-axioms.

Exercise 42 Show that T is model complete if and only if for every model $\mathfrak{A}$ of T, $T \cup \Delta_{\mathfrak{A}}$ is a complete $\mathcal{L}_{\mathfrak{A}}$-theory.

We shall return to model completeness later; for the moment, we focus on quantifier elimination for a while. We give a general lemma which characterizes quantifier elimination in a model-theoretic way, and we shall prove for two theories that they admit quantifier elimination: the theories of *dense linear orders without end-points*, and *real closed fields*. Before starting, however, we have to clear up a triviality about the logic.

Linguistic detail From now on, we assume that in the predicate logic we are using, there is an atomic sentence $-$, which is never true in a model. Of course, we have then that $\models - \leftrightarrow \exists x\neg(x = x)$ so $-$ is redundant in a sense, but without it there may be no quantifier-free sentences at all (if $\mathcal{L}$ has no constants or 0-ary relation symbols).

Lemma 8.2 *The following three conditions are equivalent for an $\mathcal{L}$-theory T:*

i) For any model $\mathfrak{B}$ of T and any finitely generated substructure $\mathfrak{A}$ of $\mathfrak{B}$, $T \cup \Delta_{\mathfrak{A}}$ is a complete $\mathcal{L}_{\mathfrak{A}}$-theory;

ii) T has quantifier elimination;

iii) For every pair $\mathfrak{A} \xrightarrow{f} \mathfrak{B}$, $\mathfrak{A} \xrightarrow{g} \mathfrak{C}$ of embeddings of $\mathcal{L}$-structures, where $\mathfrak{B}$ and $\mathfrak{C}$ are models of T, and $\mathfrak{A}$ is finitely generated, there is a commutative diagram

$$\begin{tikzcd} \mathfrak{A} \arrow[r, "f"] \arrow[d, "g"'] & \mathfrak{B} \arrow[d, "k"] \\ \mathfrak{C} \arrow[r, "h"'] & \mathfrak{D} \end{tikzcd}$$

where h, k are elementary embeddings.

Proof. We prove iii)⇒i)⇒ii)⇒i)⇒iii).

iii)⇒i): this is easy. A model $\mathfrak{B}$ of $T \cup \Delta_{\mathfrak{A}}$ is nothing but an embedding of $\mathfrak{A}$ into a model of T. By iii), every two such models have a common elementary extension such that the diagram commutes; this means that they are elementarily equivalent $\mathcal{L}_{\mathfrak{A}}$-structures. Hence, $T \cup \Delta_{\mathfrak{A}}$ is complete by Exercise 2.

i)⇒ii). This is somewhat similar to the proof of the Los-Tarski Theorem (6.2). Let $\varphi(x_1, \ldots, x_n)$ be an $\mathcal{L}$-formula. Pick new constants $c_1, \ldots, c_n$ and let Δ be the set of all quantifier-free $\mathcal{L} \cup \{c_1, \ldots, c_n\}$-sentences δ such that $T \models \varphi(c_1, \ldots, c_n) \rightarrow \delta$.

Suppose $\mathfrak{B}$ is a model of $T \cup \Delta$. Let $\mathfrak{A}$ be the substructure of $\mathfrak{B}$ generated by $c_1^{\mathfrak{B}}, \ldots, c_n^{\mathfrak{B}}$; so, A is the set

$$\{t^{\mathfrak{B}}(c_1^{\mathfrak{B}}, \ldots, c_n^{\mathfrak{B}}) \mid t(x_1, \ldots, x_n) \text{ an } \mathcal{L}\text{-term}\}$$

Suppose $\mathfrak{B} \models \neg\varphi(c_1, \ldots, c_n)$. We regard $\varphi(c_1, \ldots, c_n)$ as an $\mathcal{L}_{\mathfrak{A}}$-sentence. Since $T \cup \Delta_{\mathfrak{A}}$ is complete, we have $T \cup \Delta_{\mathfrak{A}} \models \neg\varphi(c_1, \ldots, c_n)$. It follows that for some sentence $\psi(a_1, \ldots, a_m, c_1, \ldots, c_n) \in \Delta_{\mathfrak{A}}$,

$$T \models \psi(a_1, \ldots, a_m, c_1, \ldots, c_n) \rightarrow \neg\varphi(c_1, \ldots, c_n)$$

(Here the constants a_i are the constants from A different from the c_i) It follows that $T \models \forall y_1 \cdots y_m(\psi(y_1, \ldots, y_m, c_1, \ldots, c_n) \rightarrow \neg\varphi(c_1, \ldots, c_n))$.

Now pick for each a_i an $\mathcal{L}$-term t_i such that $a_i = t_i^{\mathfrak{B}}(\vec{c})$. Then

$$T \models \psi(t_1(\vec{c}), \ldots, t_m(\vec{c}), \vec{c}) \rightarrow \neg\varphi(\vec{c})$$

so $\neg\psi(t_1(\vec{c}), \ldots, t_m(\vec{c}), \vec{c})$ is an element of Δ and therefore true in $\mathfrak{B}$; but this is a contradiction since $a_i = t_i^{\mathfrak{B}}(\vec{c}^{\mathfrak{B}})$. We conclude that $\mathfrak{B} \models \varphi(\vec{c})$; since $\mathfrak{B}$ was an arbitrary model of $T \cup \Delta$, we have $T \cup \Delta \models \varphi(\vec{c})$ and hence, for some $\delta(\vec{c}) \in \Delta$, $T \models \delta(\vec{c}) \rightarrow \varphi(\vec{c})$, so $T \models \delta(\vec{c}) \leftrightarrow \varphi(\vec{c})$, so $T \models \forall\vec{x}(\delta(\vec{x}) \leftrightarrow \varphi(\vec{x}))$, as required.

ii)⇒i). Let $\varphi(a_1, \ldots, a_n)$ be an $\mathcal{L}_{\mathfrak{A}}$-sentence, where $\mathfrak{A}$ is a substructure of a model $\mathfrak{B}$ of T. Since T admits quantifier elimination,

$$T \models \varphi(a_1, \ldots, a_n) \leftrightarrow \psi(a_1, \ldots, a_n)$$

for some quantifier-free $\psi(a_1, \ldots, a_n)$. If $\psi(a_1, \ldots, a_n) \in \Delta_{\mathfrak{A}}$ then $T \cup \Delta_{\mathfrak{A}} \models \varphi(a_1, \ldots, a_n)$, and if $\psi(a_1, \ldots, a_n) \notin \Delta_{\mathfrak{A}}$ then $T \cup \Delta_{\mathfrak{A}} \models \neg\varphi(a_1, \ldots, a_n)$. So $T \cup \Delta_{\mathfrak{A}}$ is complete.

i)⇒iii). Since both $\mathfrak{B}$ and $\mathfrak{C}$ are models of the complete theory $T \cup \Delta_{\mathfrak{A}}$, they are $\mathcal{L}_{\mathfrak{A}}$-elementarily equivalent and have therefore a common $\mathcal{L}_{\mathfrak{A}}$-elementary extension by Exercise 13. The extension is then also $\mathcal{L}$-elementary, and the fact that it is a common $\mathcal{L}_{\mathfrak{A}}$-elementary extension entails that the diagram in iii) commutes. ■

We shall now use this lemma to prove that the theory of dense linear orders without end-points has quantifier elimination. Let's recall the axioms of this theory: the language $\mathcal{L}$ is $\{<\}$, and the axioms are:

irreflexivity	$\neg(x < x)$
transitivity	$(x < y \wedge y < z) \rightarrow x < z$
linearity	$x < y \vee x = y \vee y < x$
density	$x < y \rightarrow \exists w(x < w \wedge w < y)$
no end-points	$\exists wz(w < x \wedge x < z)$

We shall call this theory DLO.

Lemma 8.3 *Let $\mathfrak{A}$ be a finite linear order $a_1 < \cdots < a_m$. Suppose $\mathfrak{A}$ is embedded in both $\mathfrak{B}$ and $\mathfrak{C}$, where $\mathfrak{B}$ and $\mathfrak{C}$ are models of* DLO. *Consider $\mathfrak{B}, \mathfrak{C}$ as $\mathcal{L}_{\mathfrak{A}}$-structures. Then for any $\mathcal{L}_{\mathfrak{A}}$-formula $\varphi(x_1, \ldots, x_n)$ we have: if two n-tuples $b_1, \ldots, b_n \in B$, $c_1, \ldots, c_n \in C$ satisfy the conditions:*

$$\forall 1 \leq i \leq m, 1 \leq j \leq n\,((a_i < b_j \Leftrightarrow a_i < c_j) \wedge (b_j < a_i \Leftrightarrow c_j < a_i))$$
$$\forall 1 \leq i, j \leq n\,(b_i < b_j \Leftrightarrow c_i < c_j)$$

then $\mathfrak{B} \models \varphi(b_1, \ldots, b_n)$ if and only if $\mathfrak{C} \models \varphi(c_1, \ldots, c_n)$.

Proof. We use induction on φ. The case for atomic φ is left to you, as are the induction steps for $\neg$ and $\wedge$. Now suppose the lemma is true for $\varphi(x, x_1, \ldots, x_n)$, the tuples $b_1, \ldots, b_n \in B$, $c_1, \ldots, c_n \in C$ satisfy the condition in the lemma, and $\mathfrak{B} \models \exists x \varphi(x, b_1, \ldots, b_n)$. Then for some $b \in B$, $\mathfrak{B} \models \varphi(b, b_1, \ldots, b_n)$. It suffices now to find $c \in C$ such that the $n+1$-tuples $b, b_1, \ldots, b_n \in B$, $c, c_1, \ldots, c_n \in C$ satisfy the conditions in the lemma; but it is easy to see that thanks to the axioms of DLO, this can always be achieved. ∎

Corollary 8.4 DLO *is model complete.*

Proof. If $\mathfrak{B} \subseteq \mathfrak{C}$ are models of DLO, apply Lemma 8.3 with $\mathfrak{A}$ empty. ∎

Theorem 8.5 DLO *has quantifier elimination.*

Proof. We prove property iii) of Lemma 8.2 for DLO.

So suppose $\mathfrak{A}$ is a substructure of $\mathfrak{B}$ and of $\mathfrak{C}$, where $\mathfrak{B}, \mathfrak{C}$ are models of DLO. We regard $\mathfrak{B}$ and $\mathfrak{C}$ as $\mathcal{L}_{\mathfrak{A}}$-structures. Suppose DLO $\cup \Delta_{\mathfrak{B}} \cup \Delta_{\mathfrak{C}}$ is inconsistent (as $\mathcal{L}_{\mathfrak{A},\mathfrak{B},\mathfrak{C}}$-theory). Then for some quantifier-free $\varphi(b_1, \ldots, b_n)$ and $\psi(c_1, \ldots, c_m)$ such that $\mathfrak{B} \models \varphi(b_1, \ldots, b_n)$ and $\mathfrak{C} \models \psi(c_1, \ldots, c_m)$, we have

$$\text{DLO} \models \varphi(\vec{b}) \rightarrow \neg\psi(\vec{c})$$

hence DLO $\models \forall \vec{y}(\varphi(\vec{b}) \rightarrow \neg\psi(\vec{y}))$. Now in φ and ψ together there are finitely many constant from $\mathfrak{A}$, say $a_1, \ldots, a_k$. By the axioms of DLO, pick in $\mathfrak{B}$ elements $b'_1, \ldots, b'_m$, such that the m-tuples $b'_1, \ldots, b'_m \in B$, $c_1, \ldots, c_m \in C$ satisfy the conditions in Lemma 8.3, with respect to the linear order on $\{a_1, \ldots, a_k\}$. Then by that lemma we must have $\mathfrak{B} \models \psi(b'_1, \ldots, b'_m)$. Clearly, a contradiction is obtained.

Therefore, DLO $\cup \Delta_{\mathfrak{B}} \cup \Delta_{\mathfrak{C}}$ is consistent as $\mathcal{L}_{\mathfrak{A},\mathfrak{B},\mathfrak{C}}$-theory, and has a model $\mathfrak{D}$. This means we have a commutative diagram

$$\begin{tikzcd} \mathfrak{A} \arrow[r] \arrow[d] & \mathfrak{B} \arrow[d, "k"] \\ \mathfrak{C} \arrow[r, "h"] & \mathfrak{D} \end{tikzcd}$$

with h, k embeddings. But we have already seen that DLO is model complete (8.4), so the embeddings h, k are elementary. This proves property iii) of Lemma 8.2 for DLO, which therefore has quantifier elimination. ∎

8.1 Quantifier Elimination: Real Closed Fields

In this section we review a famaous theorem by Tarski, that the theory of *real closed fields* has quantifier elimination. The proof is not extremely difficult, but rather long (if one doesn't want to assume deep results from algebra); and the techniques used in it have little to do with the rest of these notes. Therefore I have decided to put it in a separate section, which may be skipped without disturbing one's reading of these notes.

There are a number of equivalent ways to formulate the theory of real closed fields. We shall use the language $\mathcal{L} = \{0, 1; +, \cdot; <\}$, which we fix for this entire section. The theory RCF is given by the following axioms:

1) The axioms for a field;

2) the axioms for a linear order;

3) $\forall x(0 < x \to \exists y(y^2 = x))$

4) $\forall x_1 \cdots x_n \, (x_1^2 + \cdots + x_n^2 = 0 \to x_1 = 0 \wedge \cdots \wedge x_n = 0)$
for all n;

5) $\forall y_1 \cdots y_n \exists x(x^n + y_1 x^{n-1} + \cdots + y_{n-1}x + y_n = 0)$
for all odd n.

A *real closed field* is a model of RCF. Examples are $\mathbb{R}$ and various subfields of $\mathbb{R}$, such as the algebraic closure of $\mathbb{Q}$ in $\mathbb{R}$, the algebraic closure of $\mathbb{Q}(e)$ in $\mathbb{R}$, ... Another example is the field of *recursive reals*, that is the real numbers r which are the limit of a Cauchy sequence in $\mathbb{Q}$ that is a recursive function from $\mathbb{N}$ into $\mathbb{Q}$.

Exercise 43 a) Show that every real closed field is an *ordered field*; this means that the following sentences are true in it:

$$\forall xyz(x < y \to x + z < y + z)$$
$$\forall xyz(x < y \wedge 0 < z \to xz < yz)$$

b) Show that in a real closed field, the ordering is dense and has no endpoints;

c) Every real closed field has characteristic zero.

Definable functions. Suppose $\varphi(x_1, \ldots, x_n, y)$ is an $\mathcal{L}$-formula for which

$$\text{RCF} \models \forall\vec{x}\exists!y\,\varphi(x_1, \ldots, x_n, y)$$

Then φ defines an n-ary function on any real closed field. We may introduce a function symbol f (or f_φ) and consider the theory RCF_f in the language $\mathcal{L} \cup \{f\}$ which is RCF together with the axiom

$$\forall x_1 \cdots x_n \varphi(x_1, \ldots, x_n, f(x_1, \ldots, x_n))$$

The theory RCF_f is *conservative* over RCF. This means: every $\mathcal{L}$-sentence which is a consequence of RCF_f is also a consequence of RCF. This follows from the easy observation that every real closed field has a unique expansion to an $\mathcal{L} \cup \{f\}$-structure which is a model of RCF_f.

We shall call f a *definable function* (when φ is understood). We can have more than one definable function: if $f_1, \ldots, f_n$ are definable functions we shall have the theory $\text{RCF}_{f_1,\ldots,f_n}$ which extends RCF by all the defining axioms for $f_1, \ldots, f_n$. The process can be iterated: if $\varphi(x_1, \ldots, x_m, y)$ is now an $\mathcal{L} \cup \{f_1, \ldots, f_n\}$-formula such that

$$\text{RCF}_{f_1,\ldots,f_n} \models \forall\vec{x}\exists!y\varphi(\vec{x}, y)$$

we can have $(\text{RCF}_{f_1,\ldots,f_n})_f$ (where f is defined by φ in $\text{RCF}_{f_1,\ldots,f_n}$). However, it is an easy exercise that f can already be defined in RCF. We express this as follows: *definable functions are closed under composition.*

Let f be an n-ary definable function. We shall call f *eliminable* if for every quantifier-free $\mathcal{L}$-formula $B(y, u_1, \ldots, u_m)$ there exists another quantifier-free $\mathcal{L}$-formula $C(x_1, \ldots, x_n, u_1, \ldots, u_m)$ such that

$$\text{RCF}_f \models \forall\vec{x}\vec{u}\,(B(f(\vec{x}), \vec{u}) \leftrightarrow C(\vec{x}, \vec{u}))$$

We note, that also the eliminable functions are closed under composition: if $f, f_1, \ldots, f_n$ are eliminable, $g(\vec{x})$ is defined as $f(f_1(\vec{x}), \ldots, f_n(\vec{x}))$, and $B(y, \vec{u})$ is quantifier-free, then, in $\text{RCF}_{g,f,f_1,\ldots,f_n}$,

$$\begin{array}{rcl} B(g(\vec{x}), \vec{u}) & \leftrightarrow & C(f_1(\vec{x}), \ldots, f_n(\vec{x}), \vec{u}) \\ & \leftrightarrow & C_1(f_2(\vec{x}), \ldots, f_n(\vec{x}), \vec{x}, \vec{u}) \\ & \vdots & \\ & \leftrightarrow & C_n(\vec{x}, \vec{u}) \end{array}$$

for suitable quantifier-free formulas $C, C_1, \ldots, C_n$.

Polynomial relations. Every term $t(x_1, \ldots, x_n)$ of $\mathcal{L}$ denotes a polynomial in indeterminates $x_1, \ldots, x_n$ and integer coefficients.

A formula $t(\vec{x}) > 0$, where t is an $\mathcal{L}$-term, is called a *polynomial relation*. Every quantifier-free formula is, in RCF, equivalent to a propositional combination of polynomial relations: e.g., $t = s$ is equivalent to $\neg(t - s > 0) \wedge \neg(s - t > 0)$.

Therefore, for testing whether a definable function is eliminable, it suffices to look at polynomial relations.

We shall make use of the following theorem, familiar from elementary analysis; the proof is omitted.

Theorem 8.6 (Rolle's Theorem for Real Closed Fields) *Let K be a real closed field and $P \in K[X]$ be a polynomial with coefficients in K. Let P' denote its derivative; then*

$$K \models \forall xy(x < y \wedge P(x) = P(y) \rightarrow \exists z(x < z < y \wedge P'(z) = 0))$$

The following lemma will take up most of this section.

Lemma 8.7 *Let $t(x, x_1, \dots, x_m)$ be an $\mathcal{L}$-term, regarded as a polynomial whose degree in x is n. Then:*

a) *There are eliminable functions $\xi_1(x_1, \dots, x_m), \dots, \xi_{n-1}(x_1, \dots, x_m)$ such that the following statements are consequences of* $\mathrm{RCF}_{\xi_1,\cdots,\xi_m}$*:*

- $\forall x_1 \cdots x_m\, (\xi_1(\vec{x}) < \cdots < \xi_{n-1}(\vec{x}))$*;*
- *for all $x_1, \dots x_m$, the function $x \mapsto t(x, \vec{x})$ is either constant, or it is strictly monotonic (increasing or decreasing) on each interval* $(-\infty, \xi_1(\vec{x})), (\xi_1(\vec{x}), \xi_2(\vec{x})), \dots, (\xi_{n-2}(\vec{x}), \xi_{n-1}(\vec{x})), (\xi_{n-1}(\vec{x}), \infty)$

b) *There are eliminable functions $k(\vec{x}), \eta_1(\vec{x}), \dots, \eta_n(\vec{x})$ such that the following statements are consequences of* $\mathrm{RCF}_{k,\eta_1,\cdots,\eta_n}$*:*

- $k(\vec{x}) = 0 \vee \cdots \vee k(\vec{x}) = n + 1$*;*
- $\eta_1(\vec{x}) < \cdots \eta_n(\vec{x})$*;*
- *For each $j \in \{1, \dots, n\}$: $k(\vec{x}) = j$ implies that $\eta_1(\vec{x}), \dots, \eta_j(\vec{x})$ are exactly the zeros of $t(x, \vec{x})$;*
- *$k(\vec{x}) = 0$ implies that $t(x, \vec{x})$ has no zeros;*
- *$k(\vec{x}) = n + 1$ implies that $t(x, \vec{x})$ is the constant zero polynomial.*

Proof. I have given the statements in lemma 8.7 in informal language, and leave it to the reader to see that these statements can be expressed by $\mathcal{L}$-formulas. Of course, for a natural number j, j is the $\mathcal{L}$-term $\underbrace{1 + \cdots + 1}_{j \text{ times}}$.

We shall prove the lemma by induction on n. We write $t(x, \vec{x})$ as

$$u_0(\vec{x})x^n + \cdots + u_{n-1}(\vec{x})x + u_n(\vec{x})$$

For $n = 0$, $t(x, \vec{x}) = u_0(\vec{x})$ and there is nothing to prove for a); for b), we let

$$k(\vec{x}) = \begin{cases} 0 & \text{if } u_0(\vec{x}) \neq 0 \\ 1 & \text{if } u_0(\vec{x}) = 0 \end{cases}$$

Then $k(\vec{x})$ is eliminable: for a polynomial relation $p(y, \vec{v}) > 0$ we have

$$\begin{array}{rcl} p(k(\vec{x}), \vec{v}) > 0 & \leftrightarrow & (u_0(\vec{x}) \neq 0 \wedge p(0, \vec{v}) > 0) \\ & & \vee(u_0(\vec{x}) = 0 \wedge p(1, \vec{v}) > 0) \end{array}$$

For $n = 1$, $t(x, \vec{x}) = u_1(\vec{x})x + u_0(\vec{x})$. Induction hypothesis a) means now that the statement that either $u_1(\vec{x}) = 0$ or $t(x, \vec{x})$ is monotonic on $(-\infty, \infty)$, is a consequence of RCF. This is easy to see and left to you. For induction hypothesis b) we let

$$k(\vec{x}) = \begin{cases} 1 & \text{if } u_1(\vec{x}) \neq 0 \\ 0 & \text{if } u_1(\vec{x}) = 0 \wedge u_0(\vec{x}) \neq 0 \\ 2 & \text{if } u_1(\vec{x}) = 0 \wedge u_0(\vec{x}) = 0 \end{cases}$$

Again, $k(\vec{x})$ is eliminable. We let

$$\eta_1(\vec{x}) = \begin{cases} 0 & \text{if } u_1(\vec{x}) = 0 \\ u_0(\vec{x})/u_1(\vec{x}) & \text{otherwise} \end{cases}$$

And also $\eta_1(\vec{x})$ is eliminable and has the right properties.

Now suppose $n > 1$ and we have proved the lemma for all $n' < n$. Let $t'(x, \vec{x})$ be the derivative of $t(x, \vec{x})$ with respect to x. Since the degree of t' in x is $n-1$, by induction hypothesis b) we have eliminable functions $k'(\vec{x}), \eta'_1(\vec{x}), \ldots, \eta'_{n-1}(\vec{x})$ satisfying b) for $t'(x, \vec{x})$.

To prove a) for $t(x, \vec{x})$, we take the η''s for the ξ's. Now either $k'(\vec{x}) = n$ ($t(x, \vec{x})$ is constant), or $t(x, \vec{x})$ is monotonic on each interval of form as in a); this follows from Rolle's Theorem (8.6) for RCF.

To prove b) we define formulas $C_0(\vec{x}), \ldots, C_{n-1}(\vec{x})$ as follows:

$$\begin{array}{rcl} C_0(\vec{x}) & \equiv & (t(\eta'_1(\vec{x}), \vec{x}) - t(\eta'_1(\vec{x}) - 1, \vec{x}))(t(\eta'_1(\vec{x}), \vec{x})) > 0 \\ C_i(\vec{x}) & \equiv & t(\eta'_i(\vec{x}), \vec{x}) = 0 \vee \\ & & t(\eta'_i(\vec{x}), \vec{x})t(\eta'_{i+1}(\vec{x}), \vec{x}) < 0 \\ & & \text{for } i = 1, \ldots, n-2 \\ C_{n-1}(\vec{x}) & \equiv & t(\eta'_{n-1}(\vec{x}), \vec{x}) = 0 \vee \\ & & t(\eta'_{n-1}(\vec{x}), \vec{x})(t(\eta'_{n-1}(\vec{x}) + 1, \vec{x}) - t(\eta'_{n-1}(\vec{x}), \vec{x})) < 0 \end{array}$$

Now $C_0(\vec{x})$ means that either $t(\eta'_1(\vec{x}), \vec{x}) > 0$ and $t(\eta'_1(\vec{x}), \vec{x}) > t(\eta'_1(\vec{x}) - 1, \vec{x})$, or $t(\eta'_1(\vec{x}), \vec{x}) < 0$ and $t(\eta'_1(\vec{x}), \vec{x}) < t(\eta'_1(\vec{x}) - 1, \vec{x})$; by the induction hypothesis and the axioms of RCF, this means either $t(x, \vec{x})$ is constant or $t(x, \vec{x})$ has a zero in the interval $(-\infty, \eta'_1(\vec{x}))$. In a similar way, for $i = 1, \ldots, n-2$, $C_i(\vec{x})$ expresses that $t(x, \vec{x})$ has a zero in the half-open interval $[\eta'_i(\vec{x}), \eta'_{i+1}(\vec{x}))$, and $C_{n-1}(\vec{x})$ expresses that $t(x, \vec{x})$ has a zero in $(\eta'_{n-1}(\vec{x}), \infty)$. Note, that all C_i are quantifier-free formulas in the η''s.

Again using Rolle's theorem, one sees that if $t(x, \vec{x})$ is non-constant, each of the intervals contains at most one zero of $t(x, \vec{x})$. Therefore we can define a formula $K(\vec{x}, j)$, expressing that $t(x, \vec{x})$ has exactly j zeros, in a quantifier-free way by:

$$K(\vec{x}, j) \equiv \bigvee_{\substack{A \subseteq \{0, \ldots, n-1\}, \\ |A| = j}} ((\bigwedge_{k \in A} C_k(\vec{x})) \wedge (\bigwedge_{k \notin A} \neg C_k(\vec{x})))$$

Let $L(\vec{x})$ be the quantifier-free formula $u_0(\vec{x}) = \cdots = u_n(\vec{x}) = 0$, then we define $k(\vec{x})$ as $n+1$ if $L(\vec{x})$, and j for the unique j such that $K(\vec{x}, j)$, otherwise. Every polynomial relation $p(k(\vec{x}), \vec{v}) > 0$ is now equivalent to a disjunction

$$\bigvee_{j=0}^{n+1} (k(\vec{x}) = j \wedge p(j, \vec{v}) > 0)$$

so k is eliminable.

We have to define the functions $\eta_1, \ldots, \eta_n$ for $t(x, \vec{x})$. I do this in words. $\eta_1(\vec{x})$ is defined as 0 if $k(\vec{x}) = 0$ or $k(\vec{x}) = n+1$, and as the least zero of $t(x, \vec{x})$ otherwise.

Suppose $\eta_j(\vec{x})$ has been defined; then $\eta_{j+1}(\vec{x})$ is defined as $\eta_j(\vec{x}) + 1$ if $k(\vec{x}) \leq j$ or $k(\vec{x}) = n+1$; otherwise it is the least zero which is greater than $\eta_j(\vec{x})$.

We are left to prove that the functions η_j are eliminable. Every polynomial relation $p(\eta_j(\vec{x}), \vec{v}) > 0$ is equivalent to the disjunction

$$\bigvee_{i=0}^{n+1} (k(\vec{x}) = i \wedge p(\eta_j(\vec{x}), \vec{v}) > 0)$$

Now for $i = 0$ or $i = n+1$, $k(\vec{x}) = i \wedge p(\eta_j(\vec{x}), \vec{v}) > 0$ is equivalent to $k(\vec{x}) = i \wedge p(j-1, \vec{v}) > 0$.

Suppose we have shown that for $j' < j$, $\eta_{j'}$ is eliminable. Then if $0 \leq i < j$, $k(\vec{x}) = i \wedge p(\eta_j(\vec{x}), \vec{v}) > 0$ is equivalent to $k(\vec{x}) = i \wedge p(\eta_i(\vec{x}) + j - i, \vec{v}) > 0$. So we are left with the case that $\eta_j(\vec{x})$ is a real zero of $t(x, \vec{x})$. For simplicity we assume $j = 1$; the other cases involve bigger formulas, but are essentially similar. So, $t(\eta_1(\vec{x}), \vec{x}) = 0$, t is not constant zero, and we have to consider $p(\eta_1(\vec{x}), \vec{v}) > 0$.

By division with remainder in polynomial rings, there are polynomials $f(x, \vec{x}, \vec{v})$ and $g(x, \vec{x}, \vec{v})$ such that

- $p(x, \vec{v}) = f(x, \vec{x}, \vec{v})t(x, \vec{x}) + g(x, \vec{x}, \vec{v})$
- the degree of g in x is less than n.

Since η_1 is a real zero of t, we have that $p(\eta_1(\vec{x}), \vec{v}) = g(\eta_1(\vec{x}), \vec{x}, \vec{v})$. Let the degree of g in x be $r < n$; then we may apply induction hypothesis b) to g, and assume there are eliminable functions

$$l(\vec{x}, \vec{v}), \zeta_1(\vec{x}, \vec{v}), \ldots, \zeta_r(\vec{x}, \vec{v})$$

for g as in b), i.e. giving number of zeros and list of possible zeros.

We also have the eliminable functions $\xi_1(\vec{x}), \ldots, \xi_{n-1}(\vec{x})$, satisfying a) for $t(x, \vec{x})$. Let us, from now on, suppress the extra variables and just write $\zeta_1, \ldots, \zeta_r, \xi_1, \ldots, \xi_{n-1}, \eta_1, p(u), g(u)$.

The statement $p(\eta_1) > 0$ is, as we have seen, equivalent to $g(\eta_1) > 0$. This is equivalent to a disjunction, distinguishing cases according to the relative position of η_1 among the ζ's: for example if $\zeta_1 < \eta_1 < \zeta_2$ then $g(\eta_1) > 0$ is equivalent to $g(\frac{1}{2}(\zeta_1 + \zeta_2)) > 0$. So,

$$\begin{aligned} g(\eta_1) > 0 \quad \leftrightarrow \quad & (\eta_1 < \zeta_1 \wedge g(\zeta_1 - 1) > 0) \\ & \vee (\eta_1 = \zeta_1 \wedge g(\zeta_1) > 0) \\ & \vee (\zeta_1 < \eta_1 < \zeta_2 \wedge g(\tfrac{1}{2}(\zeta_1 + \zeta_2)) > 0 \\ & \vdots \\ & \vee (\zeta_r < \eta_1 \wedge g(\zeta_r + 1) > 0 \end{aligned}$$

By eliminability of the ζ's, all parts $g(\zeta_1 - 1) > 0$, $g(\frac{1}{2}(\zeta_1 + \zeta_2)) > 0$, etc. are equivalent to quantifier-free formulas not involving the ζ's. So we are left with the formulas $\eta_1 < \zeta_1, \ldots$ We do the case $\eta_1 < \zeta_1$. This time we distinguish cases according to the relative position of ζ_1 among the ξ's; we have $\zeta_1 < \xi_1 \vee \cdots$ Take for example the case $\eta_1 < \zeta_1 < \xi_1$. Recall that η_1 is the first zero of t. Now this zero occurs $< \zeta_1$ if and only if:

$$\begin{aligned} \text{either} \quad & t(\xi_1) > 0 \wedge t(\xi_1) > t(\xi_1 - 1) \wedge t(\zeta_1) > 0 \\ \text{or} \quad & t(\xi_1) < 0 \wedge t(\xi_1) < t(\xi_1 - 1) \wedge t(\zeta_1) < 0 \end{aligned}$$

By eliminability of the ξ's and ζ's, these formulas are equivalent to quantifier-free formulas not involving these symbols.

All other cases are dealt with in a similar way. This completes the proof of lemma 8.7. ∎

Theorem 8.8 (Tarski) RCF *has quantifier elimination.*

Proof. By exercise 38 it suffices to consider $\mathcal{L}$-formulas ϕ of form $\exists x A(x, \vec{x})$ where A is a conjunction of atomic formulas and negations of atomic formulas. By the axioms for a linear order we can eliminate the negations: $\neg(t = s)$ is equivalent to $t < s \vee s < t$. So A is equivalent to a disjunction $\bigvee_i A_i$ where every A_i is a conjunction of formulas of form $p(x, \vec{x}) = 0$ or $p(x, \vec{x}) > 0$; hence it suffices to consider the formulas $\exists x A_i$, so we may assume that A is of this form.

For each p, we have from lemma 8.7 eliminable functions $k_p(\vec{x})$ and $\xi_1^p(\vec{x})$, $\ldots$,$\xi_{n_p}^p(\vec{x})$ such that $p(x, \vec{x}) = 0$ is equivalent to the disjunction

$$\bigvee_{i=1}^{n_p} (k_p(\vec{x}) \geq i \wedge x = \xi_i^p(\vec{x}))$$

and $p(x, \vec{x}) > 0$ is equivalent to the disjunction

$$\begin{aligned} & (x < \xi_1^p(\vec{x}) \wedge p(\xi_1^p(\vec{x}) - 1, \vec{x}) > 0) \vee \\ & \textstyle\bigvee_{i=1}^{n_p - 1} (\xi_i^p(\vec{x}) < x < \xi_{i+1}^p(\vec{x}) \wedge p(\frac{1}{2}(\xi_i^p(\vec{x}) + \xi_{i+1}^p(\vec{x})), \vec{x}) > 0) \\ & \vee (x > \xi_{n_p}^p(\vec{x}) \wedge p(\xi_{n_p}^p(\vec{x}) + 1, \vec{x}) > 0) \\ & \vee \textstyle\bigvee_{i+1} n_p (i > k_p(\vec{x}) \wedge x = \xi_i^p(\vec{x}) \wedge p(\xi_i^p(\vec{x}), \vec{x}) > 0) \end{aligned}$$

So again, we can reduce to the case of a formula $\exists x A$ where A is a conjunction of statements of form: x is in a certain interval bounded by ξ_i^p's, and some other side conditions which don't depend on x.

Then $\exists x A$ is equivalent to the statement that the side conditions hold (which, by eliminability of the ξ's, is equivalent to a quantifier-free $\mathcal{L}$-formula, and that the intersection of the intervals is nonempty; this is equivalent to a quantifier-free $\mathcal{L}$-formula. ∎

8.2 Model Completeness and Model Companions

We now return to model complete theories. Robinson's Test (Exercise 40) suggests to look at so-called *existentially closed* models of a theory T. Let $\mathfrak{A} \subseteq \mathfrak{B}$ be an embedding of $\mathcal{L}$-structures. $\mathfrak{A}$ is called *existentially closed* in $\mathfrak{B}$, if for every Σ_1-sentence ϕ in the language $\mathcal{L}_{\mathfrak{A}}$, if $\mathfrak{B} \models \phi$ then $\mathfrak{A} \models \phi$. $\mathfrak{A}$ is existentially closed for T if whenever $\mathfrak{A} \subseteq \mathfrak{B}$ and $\mathfrak{B}$ is a model of T, $\mathfrak{A}$ is existentially closed in $\mathfrak{B}$.

Exercise 44 Show that T is model complete if and only if every model of T is existentially closed for T.

Now let T be an arbitrary $\mathcal{L}$-theory, not necessarily model complete. Suppose that the class of $\mathcal{L}$-structures which are models of T and existentially closed for T, is elementary, so equal to the class of models of a theory T'. Then $T \subseteq T'$ and T' is model complete.

There are several examples of mathematical theories T for which the class of existentially closed models of T is indeed the class of models of such an extension T', and moreover, every model of T can be *completed*, that is: embedded in a model of T':

- Every torsion-free abelian group can be embedded in a divisible torsion-free abelian group;
- every integral domain can be embedded in an algebraically closed field;
- every distributive lattice can be embedded in an atomless Boolean algebra;
- every ordered field can be embedded in a real closed field.

In Model Theory, this situation is described with the notion of *model companion*.

Let T and T' be $\mathcal{L}$-theories. T' is called a *model companion* of T, if the following conditions hold:

i) Every model of T can be embedded in a model of T', and vice versa;

ii) T' is model complete.

Exercise 45 Show that condition i) above is equivalent to: $T_\forall = T'_\forall$, where $T_\forall$ is the set of all Π_1-sentences which are consequences of T
[Hint: use Exercise 17]

In this section we shall prove for two pairs of theories T and T' that T' is a model companion of T: distributive lattices–atomless Boolean algebras, and integral domains–algebraically closed fields.

Definition 8.9 A *lattice* is a partial order which has a least and greatest element, and in which for each pair of elements x, y the infimum (or *meet*) $x \sqcap y$ and the supremum (or *join*) $x \sqcup y$ exist; these elements are defined by the conditions:

$$\begin{array}{l} \forall z(x \sqcup y \leq z \leftrightarrow x \leq z \wedge y \leq z) \\ \forall z(z \leq x \sqcap y \leftrightarrow z \leq x \wedge z \leq y) \end{array}$$

A lattice is called *distributive* if moreover the distributive law holds:

$$\forall xyz(x \sqcap (y \sqcup z) = (x \sqcap y) \sqcup (x \sqcap z))$$

Exercise 46 a) Show that the conditions which $x \sqcap y$ and $x \sqcup y$ are required to satisfy, indeed determine these elements uniquely;

b) show, that the distributive law implies its dual:

$$\forall xyz(x \sqcup (y \sqcap z) = (x \sqcup y) \sqcap (x \sqcup z))$$

The most immediate examples of distributive lattices are: collections of subsets of a given set X which contain $\emptyset$ and X, and are closed under union and intersection (with inclusion of subsets as the partial order). We shall soon see, that every distributive lattice is isomorphic to one of this form.

A lattice which is not distributive is the following partial order:

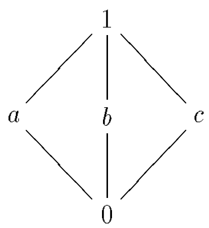

We now define formally the theory of distributive lattices.

Definition 8.10 The theory of distributive lattices, DL, is formulated in the language $\{0, 1; \sqcap, \sqcup\}$ and has the following axioms:

$$\begin{array}{ll} \forall xy(x \sqcap y = y \sqcap x) & \forall xy(x \sqcup y = y \sqcup x) \\ \forall xyz(x \sqcap (y \sqcap z) = (x \sqcap y) \sqcap z) & \forall xyz(x \sqcup (y \sqcup z) = (x \sqcup y) \sqcup z) \\ \forall x(1 \sqcap x = x) & \forall x(0 \sqcup x = x) \\ \forall x(0 \sqcap x = 0) & \forall x(1 \sqcup x = 1) \\ \forall x(x \sqcap x = x) & \forall x(x \sqcup x = x) \\ \forall xy(x \sqcap (y \sqcup x) = x) & \forall xy(y \sqcup (x \sqcap y) = y) \end{array}$$

together with the distributive law:

$$\forall xyz(x \sqcap (y \sqcup z) = (x \sqcap y) \sqcup (x \sqcap z))$$

A distributive lattice is a model of DL.

Exercise 47 Show, that definitions 8.9 and 8.10 agree. In particular, given a model of DL, if we define

$$x \leq y \text{ iff } x \sqcap y = x$$

we get a partial order, such that $x \sqcap y$ is the meet of x and y, and $x \sqcup y$ the join; 0 is the least element, and 1 the greatest.

Let $(A; 0, 1; \sqcap, \sqcup)$ be a distributive lattice, $a \in A$. A *complement* of a in A is an element b satisfying $a \sqcap b = 0 \wedge a \sqcup b = 1$. If a has a complement, it is unique. This follows by distributivity: if both b and b' are complements of a, then $b = b \sqcap 1 = b \sqcap (a \sqcup b') = (b \sqcap a) \sqcup (b \sqcap b') = 0 \sqcup (b \sqcap b') = b \sqcap b'$ so $b \leq b'$; similarly, $b' \leq b$. Note that this also implies that complements are preserved by any homomorphism of distributive lattices.

A distributive lattice in which every element has a complement is called a *Boolean algebra*. Note, that a Boolean algebra is a model of DL together with the axiom:

$$\forall x \exists y (x \sqcap y = 0 \wedge x \sqcup y = 1)$$

We call this the *theory of Boolean algebras*. Since DL has a set of Π_1-axioms, the theory of Boolean algebras has a set of Π_2-axioms.

Examples of Boolean algebras are: the power-set of any set (where complement is "real" complement); if T is a theory, we can consider the set of equivalence classes of sentences in the language of T, with $\phi \sim \psi$ iff $T \models \phi \leftrightarrow \psi$; for equivalence classes $[\phi], [\psi]$ we have $[\phi] \sqcap [\psi] = [\phi \wedge \psi]$, $[\phi] \sqcup [\psi] = [\phi \vee \psi]$, and the complement of $[\phi]$ is $[\neg\phi]$. This Boolean algebra is called the *Lindenbaum algebra* of the theory T. Note, that T is complete if and only if its Lindenbaum algebra has exactly two elements. An important case is, where we just consider equivalence classes of formulas in propositional logic, with propositional variables $p_0, p_1, \ldots$. The resulting Lindenbaum algebra is the *free Boolean algebra on countably many generators.*

Let A be a distributive lattice, and $a \in A$. a is called an *atom* in A if $a \neq 0$ and for every $b \leq a$ we have $b = 0$ or $b = a$. A Boolean algebra which contains no atoms is called *atomless.* Note that an atomless Boolean algebra is a Boolean algebra satisfying the axiom

$$\forall x \exists y (x \neq 0 \rightarrow y \neq 0 \wedge y \neq x \wedge y = y \sqcap x)$$

So also the theory of atomless Boolean algebras has a set of Π_2-axioms. The free Boolean algebra on countably many generators is an example of an atomless Boolean algebra.

Theorem 8.11 *Every distributive lattice can be embedded in a Boolean algebra, and every Boolean algebra can be embedded in an atomless Boolean algebra. Therefore, every existentially closed distributive lattice is an atomless Boolean algebra.*

Proof. Let us first show the last statement, assuming the first: if $\mathfrak{A}$ is a distributive lattice and $a \in A$, the statement "a has a complement" can be expressed by a Σ_1-sentence of $\mathcal{L}_{\mathfrak{A}}$. So since every distributive lattice can be embedded in a Boolean algebra, if $\mathfrak{A}$ is existentially closed it is a Boolean algebra. Similarly, the statement "a is not an atom" can be expressed by a Σ_1-sentence of $\mathcal{L}_{\mathfrak{A}}$, so if $\mathfrak{A}$ is existentially closed, it must be atomless.

For the first statement we use the notion of filter and prime filter. In a distributive lattice $\mathfrak{A}$, a *filter* is a subset U of A with the properties: $0 \notin U$, $a \in U, a \leq b \Rightarrow b \in U$, $a, b \in U \Rightarrow a \sqcap b \in U$, and $1 \in U$. Note that if $a \neq 0$, the set $\uparrow(a) = \{b \in A \mid a \leq b\}$ is a filter.

A *prime filter* is a filter U which moreover has the property that whenever $a \sqcup b \in U$, $a \in U$ or $b \in U$.

First we prove: if, in a distributive lattice, $a \not\leq b$, there is a prime filter U which contains a but not b. This is done with the help of Zorn's Lemma: since $a \not\leq b$, $a \neq 0$ so $\uparrow(a)$ is a filter containing a but not b. So the partially ordered set of all filters which contain a but not b (ordered by inclusion), is nonempty. One easily sees that it is closed under unions of chains. By Zorn's Lemma, it has a maximal element U. Suppose that $c \sqcup d \in U$ but $c \notin U$, $d \notin U$. Then by maximality of U, there must be $u_1, u_2 \in U$ such that $c \sqcap u_1 \leq b$ and $d \sqcap u_2 \leq b$. Then $u = u_1 \sqcap u_2 \in U$, and we have

$$u \sqcap (c \sqcup d) = (u \sqcap c) \sqcup (u \sqcap d) \leq b$$

so $b \in U$; contradiction.

Now let X be the set of all prime filters of the distributive lattice $\mathfrak{A}$. Define $f : A \to \mathcal{P}(X)$ by

$$f(a) = \{U \in X \mid a \in U\}$$

By what we just proved, it follows that f is 1-1, and that $f(a \sqcup b) = f(a) \cup f(b)$. Clearly, $f(1) = X$, $f(0) = \emptyset$, and $f(a \sqcap b) = f(a) \cap f(b)$, so f is an embedding of distributive lattices; and $\mathcal{P}(X)$ is a Boolean algebra.

Next, we show that every Boolean algebra can be embedded in an atomless Boolean algebra. Since every Boolean algebra is a distributive lattice, it can be embedded in a Boolean algebra of the form $\mathcal{P}(X)$, so it suffices to see that every $\mathcal{P}(X)$ can be embeeded in an atomless Boolean algebra. For this we observe that for every function $f : Y \to X$ of sets, the "inverse image" function $f^{-1} : \mathcal{P}(X) \to \mathcal{P}(Y)$ is a homomorphism of distributive lattices. Moreover, if f is a surjective function, then f^{-1} is 1-1.

We define a chain of Boolean algebras and embeddings:

$$\mathcal{P}(X_0) \stackrel{\pi_0^{-1}}{\to} \mathcal{P}(X_1) \stackrel{\pi_1^{-1}}{\to} \cdots$$

by letting $X_0 = X$, $X_{i+1} = X_i \times \{0, 1\}$ and $\pi_i : X_{i+1} \to X_i$ the projection. We may assume $X \neq \emptyset$, so all π_i are surjective and the π_i^{-1} therefore injective. The atoms in $\mathcal{P}(X_i)$ are the singleton subsets of X_i. Now for $y \in X_i$, $\pi_i^{-1}(\{y\}) = \{(y, 0), (y, 1)\}$ so π_i^{-1} sends every atom to a non-atom.

Let $\mathfrak{B}$ be the colimit of the chain above. Since the theory of Boolean algebras has a set of Π_2-axioms, it is preserved under unions of chains so $\mathfrak{B}$ is a Boolean algebra, in which all $\mathcal{P}(X_i)$ embed. Let $b \in B$. Then by the construction of the colimit, b comes from an element of some $\mathcal{P}(X_i)$, say $Y \subseteq X_i$. If Y is not an atom in $\mathcal{P}(X_i)$ then b is certainly not an atom in $\mathfrak{B}$, but if Y is an atom, $\pi_i^{-1}(Y)$ is a non-atom, and gives the same element b in the colimit. So, b cannot be an atom, and $\mathfrak{B}$ is an atomless Boolean algebra in which all $\mathcal{P}(X_i)$ embed. This completes the proof. ∎

Next, we wish to prove that every atomless Boolean algebra is existentially closed for the theory of distributive lattices. Equivalently, that the theory of atomless Boolean algebras is model complete. This will follow from Lindstrøm's Test (Lemma 8.15) below, once we have proved that the theory of atomless Boolean algebras is ω-categorical. So we do that first.

Lemma 8.12 *a) Let $\mathfrak{A}$ be an atomless Boolean algebra and $a \in A$, $a \neq 0$. Then for each $n \in \mathbb{N}$ there are elements $b_1, \ldots, b_n$ in A such that $b_i \neq 0$, $b_i \sqcap b_j = 0$ for $i \neq j$, and $b_1 \sqcup \cdots \sqcup b_n = a$.*

b) Let $\mathfrak{A}$ be an atomless Boolean algebra, and $\mathfrak{C} \subseteq \mathfrak{D}$ be an embedding of finite Boolean algebras. Then every embedding $\mathfrak{C} \to \mathfrak{A}$ extends to an embedding $\mathfrak{D} \to \mathfrak{A}$.

Proof. a) is proved by induction on n. For $n = 1$ take $b_1 = a$. For $n+1$: since a is not an atom, there is $0 < b_{n+1} < a$ in A; then if c is the complement of b_{n+1}, $a \sqcap c \neq 0$. Moreover, $b_{n+1} \sqcup (a \sqcap c) = a$. So apply the induction hypothesis to find $b_1, \ldots, b_n$ for $a \sqcap c$.

For b), we note that every finite Boolean algebra is isomorphic to $\mathcal{P}(X)$ for a finite set X. Now every embedding $f : \mathcal{P}(X) \to \mathcal{P}(Y)$ (for finite X, Y) is determined by its values on the singleton subsets of X, since it must preserve unions and every subset of X is a finite union of singletons. If $X = \{x_1, \ldots, x_n\}$ we see that $f(\{x_i\}) \cap f(\{x_j\}) = \emptyset$ for $i \neq j$ (since f preserves $\emptyset$ and intersection), $f(\{x_i\}) \neq \emptyset$ (since f is an embedding) and $\bigcup_{i=1}^n f(\{x_i\}) = Y$ since f preserves 1 and unions. So f divides Y into n parts. Now suppose $g : \mathcal{P}(X) \to \mathfrak{A}$ is an embedding. For each $y \in Y$ there is a unique i with $y \in f(\{x_i\})$; suppose $f(\{x_i\}) = \{y_1, \ldots, y_m\}$. By a), choose $b_1, \ldots, b_m$ in A for $g(\{x_i\})$ and send $\{y_j\}$ to b_j. This extends to an embedding $h : \mathcal{P}(Y) \to \mathfrak{A}$, such that $g = h \circ f$. ∎

Corollary 8.13 *The theory of atomless Boolean algebras is ω-categorical.*

Proof. Let $A = \{a_0, a_1, a_2, \ldots\}$ and $B = \{b_0, b_1, b_2, \ldots\}$ be two countable atomless Boolean algebras. We define a chain of isomorphisms $f_i : A_i \to B_i$ such that $A_0 \subseteq A_1 \subseteq A_2 \subseteq \cdots$ is a chain of finite Boolean algebras with union A, and $B_0 \subseteq B_1 \subseteq B_2 \subseteq \cdots$ is a chain of finite Boolean algebras with union B, and $f_{i+1} : A_{i+1} \to B_{i+1}$ extends f_i for each i. The construction is very similar to Cantor's "back-and-forth" construction for dense linear orders.

Let A_0 be the sub-Boolean algebra of A containing just the elements $0, 1$, and B_0 likewise the sub-Boolean algebra of B with two elements, and $f_0 : A_0 \to B_0$ the obvious isomorphism.

Suppose $f_i : A_i \to B_i$ is constructed. Let C_i be the sub-Boolean algebra of A generated by $A_i \cup \{a_i\}$. Then C_i is finite, and $A_i \subseteq C_i$. Since f_i gives an embedding of A_i into the atomless Boolean algebra B, by lemma 8.12 f_i extends to an embedding $g_i : C_i \to B$. Let D_i be the image of g_i and B_{i+1} the sub-Boolean algebra of B generated by $D_i \cup \{b_i\}$. Since g_i is an embedding, its inverse is an embedding of D_i in A, which again by lemma 8.12 extends to an embedding $h_i : B_{i+1} \to A$. Let A_{i+1} be the image of h_i. Then the inverse of h_i is an isomorphism $f_{i+1} : A_{i+1} \to B_{i+1}$ which extends f_i. Since $a_i \in A_{i+1}$, the union of the A_i's is A, and similarly the union of the B_i's is B; and the union of the f_i's is an isomorphism $A \to B$. ∎

The following two general lemmas find frequent application in the study of model completeness.

Lemma 8.14 *Let T be an $\mathcal{L}$-theory with a set of Π_2-axioms. Then for any model $\mathfrak{A}$ of T, there is an embedding of $\mathfrak{A}$ into a model $\mathfrak{B}$ of T which is existentially closed for T. Moreover, if $|A| \geq \|\mathcal{L}\|$ we may assume that $|B| = |A|$.*

Proof. We construct a chain of models of T:

$$\mathfrak{A} = \mathfrak{A}_0 \subseteq \mathfrak{A}_1 \subseteq \mathfrak{A}_2 \subseteq \cdots$$

as follows: if $\mathfrak{A}_k$ has been defined, let Γ_k be a maximal set of Σ_1-sentences in the language $\mathcal{L}_{\mathfrak{A}_k}$, such that $T \cup \Delta_{\mathfrak{A}_k} \cup \Gamma_k$ is consistent (such a set exists by Zorn's Lemma). Let $\mathfrak{A}_{k+1}$ be a model of $T \cup \Delta_{\mathfrak{A}_k} \cup \Gamma_k$; then $\mathfrak{A}_k \subseteq \mathfrak{A}_{k+1}$.

Let $\mathfrak{B}$ be the colimit of this chain. Since T has a set of Π_2-axioms, $\mathfrak{B}$ is a model of T. Suppose ϕ is a Σ_1-sentence of $\mathcal{L}_{\mathfrak{B}}$ which holds in some extension of $\mathfrak{B}$ which is a model of T. By construction of the colimit, ϕ is already an $\mathcal{L}_{\mathfrak{A}_k}$-sentence for some k; then $T \cup \Delta_{\mathfrak{A}_k} \cup \Gamma_k \cup \phi$ is consistent, whence by maximality of Γ_k, $\phi \in \Gamma_k$ and ϕ holds in $\mathfrak{A}_{k+1}$, so ϕ holds in $\mathfrak{B}$. So $\mathfrak{B}$ is existentially closed for T.

To prove the final statement: if $|A| \geq \|\mathcal{L}\|$ we may take all $\mathfrak{A}_k$ such that $|A_k| = |A|$. Then $|B| = |A|$. ∎

Lemma 8.15 (Lindstrøm's Test) *Suppose T is an $\mathcal{L}$-theory with a set of Π_2-axioms, which only has infinite models, and is α-categorical for some cardinal number $\alpha \geq \|\mathcal{L}\|$. Then T is model complete.*

Proof. If $\mathfrak{A}$ is a model of T of cardinality α, then by the previous lemma $\mathfrak{A}$ may be extended to an existentially closed model $\mathfrak{B}$ of cardinality α. Since T is α-categorical, $\mathfrak{A} \cong \mathfrak{B}$ so $\mathfrak{A}$ is existentially closed for T. So every model of T of cardinality α is existentially closed for T.

Now let $\mathfrak{A} \subseteq \mathfrak{B}$ be an arbitrary embedding of models of T. We have to show that $\mathfrak{A}$ is existentially closed in $\mathfrak{B}$.

First suppose $|A| \geq \alpha$. Then for any Σ_1-sentence ϕ of $\mathcal{L}_{\mathfrak{A}}$ which holds in $\mathfrak{B}$, there is by downward Löwenheim-Skolem-Tarski (3.2) an elementary submodel $\mathfrak{A}'$ of $\mathfrak{A}$ of cardinality α, which contains all constants occurring in ϕ; then $\mathfrak{A}'$ is existentially closed so ϕ holds in it; so ϕ holds in $\mathfrak{A}$ too. So $\mathfrak{A}$ is existentially closed for T.

If $|A| < \alpha$ we view $\mathfrak{B}$ as a model for the language $\mathcal{L} \cup \{X\}$ where X is a new 1-ary relation symbol, putting $X^{\mathfrak{B}} = A$.

Consider the language $\mathcal{L}' = \mathcal{L} \cup \{X\} \cup \{c_\lambda \,|\, \lambda \in \Lambda\}$ where Λ is a set of cardinality α and the c_λ are new constants; let Γ be the $\mathcal{L}'$-theory consisting of the elementary diagram of $\mathfrak{B}$ (as $\mathcal{L}_{\mathfrak{B}} \cup \{X\}$-structure) together with the set of axioms

$$\{X(c_\lambda) \,|\, \lambda \in \Lambda\} \cup \{\neg(c_\lambda = c_{\lambda'} \,|\, \lambda \neq \lambda'\}$$

Since $\mathfrak{A}$ is infinite (because T has only infinite models), every finite subset of Γ has an interpretation in $\mathfrak{B}$; hence, Γ is consistent and has a model $\mathfrak{C}$. Then $\mathfrak{C}$ is a model of T and an elementary extension of $\mathfrak{B}$; and $X^{\mathfrak{C}}$ is the domain of a submodel $\mathfrak{A}'$ of $\mathfrak{C}$ which is an elementary extension of $\mathfrak{A}$. This is seen as follows: for an $\mathcal{L}_{\mathfrak{A}}$-sentence ϕ, let ϕ^X be obtained by replacing each quantifier $\forall x$ by $\forall x(X(x) \to \cdots)$ and each $\exists x$ by $\exists x(X(x) \wedge \cdots)$. Then $\mathfrak{A} \models \phi$ iff $\mathfrak{B} \models \phi^X$ iff $\mathfrak{C} \models \phi^X$ iff $\mathfrak{A}' \models \phi$.

Note, that $\mathfrak{A}'$ has cardinality at least α, so by what we have already proved, $\mathfrak{A}'$ is existentially closed in $\mathfrak{C}$. But now it is easy to deduce that $\mathfrak{A}$ is existentially closed in $\mathfrak{B}$. ■

Corollary 8.16 *The theory of atomless Boolean algebras is model complete; hence, it is a model companion of the theory of distributive lattices.*

Proof. Every atomless Boolean algebra is infinite (this follows at once from lemma 8.12), and the theory of atomless Boolean algebras is ω-categorical, as we have seen (8.13), so this follows from Lindstrøm's Test. ■

As a final example of a model companion, we show that the theory of algebraically closed fields is a model companion of the theory of integral domains (commutative rings with $1 \neq 0$ having no zero divisors). By elementary algebra, every integral domain is embedded in a field and every field is embedded in an algebraically closed field.

Since, for an integral domain $\mathfrak{A}$, the statement that $a \in A$ is either 0 or has a multiplicative inverse, and the statement that the polynomial $a_0X^n + \cdots + a_{n-1}X + a_n$ has a root (where $a_0, \ldots, a_n \in A$), are Σ_1-sentences of $\mathcal{L}_{\mathfrak{A}}$ which hold in some extension of $\mathfrak{A}$, every existentially closed integral domain is an algebraically closed field.

Conversely, let $\mathfrak{A}$ be an algebraically closed field, and $\phi \equiv \exists x_1 \cdots x_n \psi$ where ψ is a conjunction of atomic $\mathcal{L}_{\mathfrak{A}}$-formulas and negations of such; by Exercise 38 it suffices to consider sentences of this form. First we use the axioms of a field, to reduce formulas $\neg(t = s)$ to $\exists y(y(t - s) - 1 = 0)$, so ϕ is equivalent to $\exists x_1 \cdots x_m \chi$ where χ is a conjunction of statements $P(\vec{x}) = 0$, P a polynomial with coefficients in $\mathfrak{A}$.

In order to show that $\mathfrak{A}$ is existentially closed we need to show for such ϕ that if ϕ holds in some extension of $\mathfrak{A}$, it holds in $\mathfrak{A}$. This follows from the following elementary theorem from commutative algebra:

Theorem 8.17 (Hilbert Nullstellensatz) *Let K be an algebraically closed field and $f_1, \ldots, f_n$ be polynomials in $X_1, \ldots, X_m$ with coefficients in K. If there is no m-tuple $a_1, \ldots, a_m \in K$ such that*

$$f_1(a_1, \ldots, a_m) = \cdots = f_n(a_1, \ldots, a_m) = 0$$

then there are polynomials $g_1, \ldots, g_n$ in $X_1, \ldots, X_m$ and coefficients in K, such that

$$f_1 g_1 + \cdots f_n g_n = 1 \text{ in } K[X_1, \ldots, X_m]$$

Let $P_1(X_1, \ldots, X_m), \ldots, P_n(X_1, \ldots, X_m)$ be the system of polynomials in the formula χ, and suppose that this system has a common zero $b_1, \ldots, b_m$ in an extension L of $\mathfrak{A}$. Then the $b_1, \ldots, b_m$ induce a ring homomorphism $\mathfrak{A}[X_1, \ldots, X_m] \to L$ by $f \mapsto f(b_1, \ldots, b_m)$.

On the other hand, if the system has no common zero in $\mathfrak{A}$, by the Nullstellensatz there are $Q_1, \ldots, Q_n$ with $\sum_{i=1}^n P_i Q_i = 1$ in $\mathfrak{A}[X_1, \ldots, X_m]$. Combining, this gives $0 = 1$ in L, a contradiction.

Conclusion:

Corollary 8.18 *The theory of algebraically closed fields is a model companion of the theory of integral domains.*

8.3 Model Completions, Amalgamation, Quantifier Elimination

In this section we collect some miscellaneous facts and definitions about model completeness and model companions.

Let us first observe that if a theory T has a model companion it is unique up to equivalence of theories. For, suppose T_1 and T_2 are model companions of T. Then T_1, T_2 are model complete, and every model of T_1 can be embedded in a model of T_2, and vice versa. So given any model $\mathfrak{A}$ of T_1 we can form a chain

$$\mathfrak{A} = \mathfrak{A}_0 \subseteq \mathfrak{A}_1 \subseteq \mathfrak{A}_2 \subseteq \cdots$$

where $\mathfrak{A}_0, \mathfrak{A}_2, \mathfrak{A}_4, \ldots$ are models of T_1, and $\mathfrak{A}_1$, $\mathfrak{A}_3$, $\ldots$ are models of T_2. If $\mathfrak{B}$ is the colimit, then since both theories are model complete we have:

$$\begin{array}{c} \mathfrak{A}_0 \preceq \mathfrak{A}_2 \preceq \mathfrak{A}_4 \preceq \cdots \preceq \mathfrak{B} \\ \mathfrak{A}_1 \preceq \mathfrak{A}_3 \preceq \mathfrak{A}_5 \preceq \cdots \preceq \mathfrak{B} \end{array}$$

So $\mathfrak{A}$ is also a model of T_2. By symmetry, T_1 and T_2 have the same models.

Definition 8.19 Let T be an $\mathcal{L}$-theory.

a) A model $\mathfrak{A}$ of T is said to be *algebraically prime* if $\mathfrak{A}$ can be embedded in every model of T;

b) T is said to have the *joint embedding property* if for every two models $\mathfrak{A}$ and $\mathfrak{B}$ of T, there is a model $\mathfrak{C}$ of T and embeddings $\mathfrak{A} \to \mathfrak{C}$, $\mathfrak{B} \to \mathfrak{C}$;

c) T is said to have the *amalgamation property* if every diagram

$$\begin{tikzcd} \mathfrak{A} \arrow[r] \arrow[d] & \mathfrak{B} \\ \mathfrak{C} & \end{tikzcd}$$

of embeddings of models of T, can be completed to a commutative diagram

$$\begin{tikzcd} \mathfrak{A} \arrow[r] \arrow[d] & \mathfrak{B} \arrow[d] \\ \mathfrak{C} \arrow[r] & \mathfrak{D} \end{tikzcd}$$

of embeddings between models of T.

Exercise 48 Let T be model complete.

i) **[Prime Model Test]** If T has an algebraically prime model, T is complete;

ii) T is complete if and only if T has the joint embedding property.

The notion of *model completion* is a refinement of that of model companion. If T' is a model companion of T, every model of T can be embedded in a model of T'. If T' is a model completion, such an extension is unique up to elementary equivalence (in parameters of the model one starts with). The formal definition is:

Definition 8.20 T' is called a *model completion* of T if T' is a model companion of T and for every model $\mathfrak{A}$ of T, $T' \cup \Delta_{\mathfrak{A}}$ is complete.

Exercise 49 Let T' be a model companion of T. Show that the following two statements are equivalent:

i) T' is a model completion of T;

ii) T has the amalgamation property.

Exercise 50 Let T be a model complete theory. Show that the following are equivalent:

i) T is a model completion of $T_\forall$;

ii) $T_\forall$ has the amalgamation property;

iii) T has quantifier elimination.

Exercise 51 a) Prove that the theory of integral domains has the amalgamation property.

b) Deduce from a), that the theory of algebraically closed fields has quantifier elimination.

9 Countable Models

In this section we shall be concerned with countable languages and countable structures. So, $\mathcal{L}$ is countable in this chapter (and, if $\mathfrak{A}$ is a countable $\mathcal{L}$-structure, of course $\mathcal{L}_{\mathfrak{A}}$ is countable too).

If Γ is a set of $\mathcal{L}$-formulas with at most the variables $x_1, \ldots, x_n$ free, we write $\Gamma(x_1, \ldots, x_n)$ or $\Gamma(\vec{x})$. An $\mathcal{L}$-structure $\mathfrak{A}$ *realizes* $\Gamma(\vec{x})$ if there is an n-tuple $a_1, \ldots, a_n \in A$ such that $\mathfrak{A} \models \varphi(a_1, \ldots, a_n)$ for every $\varphi(x_1, \ldots, x_n) \in \Gamma(\vec{x})$. If $\mathfrak{A}$ does not realize Γ, $\mathfrak{A}$ is said to *omit* Γ.

Let T be an $\mathcal{L}$-theory. T is said to *realize* Γ *locally*, if there is an $\mathcal{L}$-formula $\theta(x_1, \ldots, x_n)$ such that $T \cup \{\exists x_1 \cdots x_n \theta(\vec{x})\}$ is consistent, and $T \models \forall \vec{x}(\theta(\vec{x}) \to \varphi(\vec{x}))$ for every $\varphi(\vec{x}) \in \Gamma$.

Theorem 9.1 (Omitting Types Theorem) *Let T be a consistent $\mathcal{L}$-theory, and for each $m \in \mathbb{N}$ let Γ_m be a set of formulas in variables $x_1, \ldots, x_{k_m}$. If T does not realize any Γ_m locally, then T has a countable model which omits each Γ_m.*

Proof. First we add a countable set $C = \{c_1, c_2, \ldots\}$ of new constants to the language $\mathcal{L}$; let $\mathcal{L}' = \mathcal{L} \cup C$. Fix an enumeration $\phi_0, \phi_1, \ldots$ of all $\mathcal{L}'$-sentences, in such a way that every $\mathcal{L}'$-sentence occurs infinitely often in this enumeration.

We shall build a chain of $\mathcal{L}'$-theories

$$T = T_0 \subseteq T_1 \subseteq T_2 \subseteq \cdots$$

such that the following hold:

- Every T_n is a consistent extension of T by finitely many $\mathcal{L}'$-sentences;
- if $T_n \cup \{\phi_n\}$ is consistent, then $\phi_n \in T_{n+1}$
- if $\phi_n \equiv \exists x \psi$ and $\phi_n \in T_{n+1}$, then $\psi(c) \in T_{n+1}$ for some $c \in C$
- for each $m \leq n$ and each k_m-tuple $\vec{c}$ of elements of C which occur in T_n, there is a formula $\varphi(x_1, \ldots, x_{k_m}) \in \Gamma_m$ such that $\neg\varphi(\vec{c}) \in T_{n+1}$.

Now suppose we have constructed $T_0 \subseteq T_1 \subseteq \cdots$ with these properties; let $T_\omega = \bigcup_n T_n$. Then T_ω is consistent and has a model $\mathfrak{A}$; but by construction, $\mathfrak{A}$ has a submodel $\mathfrak{B}$ with underlying set $\{c^{\mathfrak{A}} \mid c \in C\}$; and $\mathfrak{B}$ is actually an elementary submodel, so is a countable model of T. And $\mathfrak{B}$ omits each Γ_m by construction.

To construct our chain, we start by putting $T_0 = T$; T was assumed consistent, which is all there is to check at this stage.

Suppose T_n has been constructed. We build T_{n+1} in stages:

Stage 1. We check ϕ_n. If $T \cup \{\phi_n\}$ is consistent, we put ϕ_n in T_{n+1}. If not, we do nothing in this stage.

Stage 2. If $\phi_n \equiv \exists x \psi$ was put into T_{n+1} at stage 1, let c be the first constant in the enumeration of C which doesn't occur in T_n (which contains only finitely many constants from C by induction hypothesis), and put $\psi(c)$ into T_{n+1}.

Stage 3. Let C' be the finite set of constants from C which so far occur in T_{n+1}.

For each $m \leq n$ let $L_m = \{\vec{c}_0, \ldots, \vec{c}_{U(m)}\}$ be a list of all k_m-tuples of elements of C'. We work through each m and each j, $0 \leq j \leq U(m)$, as follows: we start with $m = 0, j = 0$. At substage (m, j) we have added a finite number of $\mathcal{L}'$-formulas to T; let $C_{m,j}$ be the conjunction of these. Write $C_{m,j}$ as

$$\theta(c_1, \ldots, c_{k_m}, d_1, \ldots, d_u)$$

where $(c_1, \ldots, c_{k_m})$ is the tuple $\vec{c}_j$, and $d_1, \ldots, d_u$ are the other constants from C. Since T does not locally realize Γ_m, there is a formula $\varphi(x_1, \ldots, x_{k_m}) \in \Gamma_m$ such that

$$T \not\models \forall x_1 \cdots x_{k_m}(\exists y_1 \cdots y_u \theta(x_1, \ldots, x_{k_m}, y_1, \ldots, y_u) \to \varphi(x_1, \ldots, x_{k_m}))$$

Now add $\neg\varphi(c_1, \ldots, c_{k_m})$ to T_{n+1} and proceed to substage $(m, j+1)$ if $j < U(m)$; otherwise to $(m+1, 0)$ if $m < n$; otherwise, stage 3 is completed and the construction of T_{n+1} too. ∎

The Omitting Types Theorem is often applied in order to construct (countable) models which have to be 'small' in some sense: there are no elements (or tuples) realizing any of countably many sets of formulas.

An example of this is the construction of *end-extensions*; another one is in the construction of *atomic models*.

9.1 End extensions of models of PA

The theory of *Peano Arithmetic* (PA) is formulated in the language $\mathcal{L} = \{0, 1; +, \cdot\}$, and its axioms are:

$$\begin{array}{ll} \forall x \neg(x+1=0) & \forall xy(x+1=y+1 \to x=y) \\ \forall x(x+0=x) & \forall x(x \cdot 0 = 0) \\ \forall xy(x+(y+1)=(x+y)+1) & \forall xy(x \cdot (y+1) = x \cdot y + x) \end{array}$$
$$\forall \vec{x}[(\varphi(0, \vec{x}) \wedge \forall y(\varphi(y, \vec{x}) \to \varphi(y+1, \vec{x}))) \ \to \ \forall y \varphi(y, \vec{x})]$$

The last axiom is meant to be an axiom for every $\mathcal{L}$-formula $\varphi(y, \vec{x})$. These axioms are called *induction axioms.*

Any model $\mathfrak{A}$ of PA has the properties that $+$ and $\cdot$ are commutative and associative, that $\cdot$ is distributive over $+$; that the formula $\exists z(x + (z+1) = y)$ defines a linear order $x < y$, for which 0 is the least element, and such that every element x has a successor, that is a least element greater than x.

Let $\mathfrak{A} \subseteq \mathfrak{B}$ be an embedding of models of PA. $\mathfrak{B}$ is called an *end extension* of $\mathfrak{A}$, or $\mathfrak{A}$ an *initial segment* of $\mathfrak{B}$, if for $a \in A$ and $b \in B$: if $b \leq a$ then $b \in A$. We say that the embedding is proper, if $A \neq B$.

We shall use the Omitting Types Theorem to show that every countable model of PA has a countable proper elementary end extension.

Exercise 52 Let $\mathfrak{A}$ be a model of PA and $\varphi(u,x)$ be an $\mathcal{L}_{\mathfrak{A}}$-formula such that

$$\mathfrak{A} \models \forall z \exists y \forall x u\, (x > y \wedge \varphi(u,x) \to \neg(u = z))$$

Show, using the induction axioms of PA, that

$$\mathfrak{A} \models \forall z \exists y \forall x u\, (x > y \wedge \varphi(u,x) \to u > z)$$

Proposition 9.2 *Every countable model of* PA *has a countable proper elementary end extension.*

Proof. Let $\mathfrak{A}$ be a countable model of PA; let $\mathcal{L}' = \mathcal{L}_{\mathfrak{A}} \cup \{c\}$ where c is a new constant. Consider the $\mathcal{L}'$-theory

$$E(\mathfrak{A}) \cup \{c > a \mid a \in A\}$$

Clearly, T is consistent by the Compactness Theorem, and every model of T is a proper elementary extension of $\mathfrak{A}$.

Now consider for each $a \in A$ the set of formulas

$$\Gamma_a(x) \;=\; \{x < a\} \cup \{\neg(x = b) \mid b \in A\}$$

Convince yourself that a proper elementary end extension of $\mathfrak{A}$ is nothing but a model of T which omits each $\Gamma_a(x)$.

Since there are only countably many sets $\Gamma_a(x)$ because $\mathfrak{A}$ is assumed countable, the Omitting Types Theorem gives us such a (countable) model, provided we can show that no $\Gamma_a(x)$ is locally realized by T.

So suppose $\theta'(x)$ is an $\mathcal{L}'$-formula such that

1) $T \models \forall x(\theta'(x) \to x < a)$

2) $T \models \forall x(\theta'(x) \to \neg(x = b))$ for all $b \in A$

Write $\theta'(x) = \theta(x,c)$ where θ is an $\mathcal{L}_{\mathfrak{A}}$-formula. From 1) we deduce, using the Compactness Theorem, that there is some $a' \in A$ such that

3) $\mathfrak{A} \models \forall x u(u > a' \wedge \theta(x,u) \to x < a)$

From 2) we deduce in the same way that for each $b \in A$ there is $b' \in A$ such that $\mathfrak{A} \models \forall x u(u > b' \wedge \theta(x,u) \to \neg(x = b))$; in other words,

$$\mathfrak{A} \models \forall z \exists w \forall x u(u > w \wedge \theta(x,u) \to \neg(x = z))$$

Applying Exercise 52, we find that

$$\mathfrak{A} \models \forall z \exists w \forall x u(u > w \wedge \theta(x,u) \to x > z)$$

However, combining this with 3), we get a contradiction unless $\mathfrak{A} \models \neg\exists x u \theta(x,u)$; but this means that $T \cup \{\exists x \theta'(x)\}$ is inconsistent.

Therefore, no $\Gamma_a(x)$ is locally realized by T, and we are done. ∎

9.2 Atomic Theories and Atomic Models

If $\mathfrak{A}$ is an $\mathcal{L}$-structure and $a_1, \dots, a_n \in A$, the set of $\mathcal{L}$-formulas

$$\Gamma_{a_1 \cdots a_n} = \{\varphi(x_1, \dots, x_n) \mid \mathfrak{A} \models \varphi(a_1, \dots, a_n)\}$$

is maximal w.r.t. the property that it is realized in some $\mathcal{L}$-structure.

We call $\Gamma_{a_1 \cdots a_n}$ the *type* of the tuple $a_1, \dots, a_n$.

In general, if T is an $\mathcal{L}$-theory and $\Gamma(x_1, \dots, x_n)$ a set of $\mathcal{L}$-formulas in variables $x_1, \dots, x_n$ which is maximal w.r.t. the property that it is realized in some model of T, we call Γ a *type* of T.

Exercise 53 If a type $\Gamma(\vec{x})$ of T is locally realized by T, there is a formula $\theta(\vec{x})$ such that

$$\Gamma(\vec{x}) = \{\varphi(\vec{x}) \mid T \models \forall\vec{x}\,(\theta(\vec{x}) \to \varphi(\vec{x}))\}$$

In this case we say that the type Γ is *principal* for T, and that Γ is *generated* by $\theta(\vec{x})$.

Exercise 54 Let $L_T(\vec{x})$ be the set of equivalence classes of $\mathcal{L}$-formulas in variables $\vec{x}$, where $\varphi(\vec{x}) \sim \psi(\vec{x})$ if and only if $T \models \forall\vec{x}\,(\varphi(\vec{x}) \leftrightarrow \psi(\vec{x}))$.

Show that $L_T(\vec{x})$ is a Boolean algebra. Show that $\Gamma(\vec{x})$ is realized in some model of T if and only if the set $\{[\gamma] \mid \gamma \in \Gamma\}$ is contained in a filter on $L_T(\vec{x})$. Show that $\Gamma(\vec{x})$ is a type of T if and only if $\Gamma = \{\gamma \mid [\gamma] \in U\}$ for some ultrafilter U on $L_T(\vec{x})$.

We say that T is an *atomic* theory if every $\mathcal{L}$-formula $\varphi(\vec{x})$ such that $T \cup \{\exists\vec{x}\varphi(\vec{x})\}$ is consistent, is an element of a principal type (in $\vec{x}$) of T.

An $\mathcal{L}$-structure $\mathfrak{A}$ is called *atomic* if for every n-tuple $a_1, \dots, a_n$, the type $\Gamma_{a_1 \cdots a_n}$ is principal for $\mathrm{Th}(\mathfrak{A})$.

Exercise 55 Call a Boolean algebra B *atomic* if for every $b \in B$, if $b \neq 0$ there is some atom in B which is $\leq b$.

Show that in a Boolean algebra, b is an atom if and only if the filter $\uparrow b$ is an ultrafilter.

Show that a theory T is atomic if and only if for each tuple $\vec{x}$ of variables, the Boolean algebra $L_T(\vec{x})$ is atomic.

Exercise 56 Let $\mathfrak{A}$ be an $\mathcal{L}$-structure such that for each $a \in A$ there is an $\mathcal{L}$-formula $\varphi(x)$ such that

$$\mathfrak{A} \models \forall x\,(\varphi(x) \leftrightarrow x = a)$$

Show that $\mathfrak{A}$ is atomic.

Theorem 9.3 *Let T be a complete $\mathcal{L}$-theory.*

a) If for every $n \in \mathbb{N}$, T has only countably many types in $x_1, \dots, x_n$, T has a countable atomic model.

b) *If $\mathfrak{A}$ is a countable atomic model of T, $\mathfrak{A}$ is elementarily embedded in every model of T.*

c) *If $\mathfrak{A}$ and $\mathfrak{B}$ are countable atomic models of T, they are isomorphic.*

Proof. For a), we use the Omitting Types Theorem. Since there are only countably many types altogether, there are certainly only countably many non-principal types. So by the Omitting Types Theorem, let $\mathfrak{A}$ be a countable model of T which omits each non-principal type (we cannot omit a principal type; whyΓ). Then for each $a_1, \ldots, a_n \in A$, the type $\Gamma_{a_1 \cdots a_n}$ is principal for $T = \mathrm{Th}(\mathfrak{A})$, so $\mathfrak{A}$ is atomic.
For b), suppose $\mathfrak{A}$ is a countable atomic model of T, and let $\mathfrak{B}$ be any model of T. List A as $a_0, a_1, \ldots$. Since Γ_{a_0} is principal for $\mathrm{Th}(\mathfrak{A}) = T$, there is a formula $\theta_0(x_0)$ such that $\Gamma_{a_0} = \{\varphi(x_0) \,|\, T \models \forall x_0(\theta_0(x_0) \rightarrow \varphi(x_0))\}$. Since T is complete, $T \models \exists x_0 \theta_0(x_0)$, so $\mathfrak{B} \models \theta_0(b_0)$ for some $b_0 \in B$.

Now suppose we have defined $b_0, \ldots, b_n \in B$ such that $b_0, \ldots, b_n$ realizes $\Gamma_{a_0 \cdots a_n}$ in B. The type $\Gamma_{a_0 \ldots a_{n+1}}$ is generated by some formula $\theta_{n+1}(x_0, \ldots, x_{n+1})$. Then the formula $\exists x_{n+1} \theta_{n+1}$ is an element of $\Gamma_{a_0 \cdots a_n}$, so $\mathfrak{B} \models \theta_{n+1}(b_0, \ldots, b_{n+1})$ for some b_{n+1}. Then $b_0, \ldots, b_{n+1}$ realizes $\Gamma_{a_0 \cdots a_{n+1}}$.

The map $a_n \mapsto b_n$ is then an elementary embedding: $\mathfrak{A} \rightarrow \mathfrak{B}$.
c) is proved by performing the construction of b) in two directions; this is similar to the back-and-forth method in the proof of Corollary 8.13, and is left to you.
∎

Exercise 57 Work out the proof of c) above.

Exercise 58 Theorem 9.3b) has a converse: if $\mathfrak{A}$ is a countable structure such that $\mathfrak{A}$ is elementarily embedded in every structure that is a model of $\mathrm{Th}(\mathfrak{A})$, then $\mathfrak{A}$ is atomic.
[Hint: for $a_1, \ldots, a_n \in A$, show that $\Gamma_{a_1 \cdots a_n}$ must be locally realized by $\mathrm{Th}(\mathfrak{A})$, *using the Omitting Types Theorem. So it must be principal]*

Exercise 59 Show, for a complete theory T, that T has a countable atomic model if and only if T is an atomic theory.
[Hint: in one direction, if $\mathfrak{A}$ is an atomic model of T and $\exists \vec{x} \varphi(\vec{x})$ consistent with T, show that $\varphi(\vec{x})$ is realized by a tuple $\vec{a}$ in A; so $\varphi(\vec{x}) \in \Gamma_{\vec{a}}$, which is a principal type.
For the other direction, let, for each n, Γ_n be the set of formulas $\neg\psi(x_1, \ldots, x_n)$ for ψ such that ψ generates a principal type for T. Use the Omitting Types Theorem to show that T has a countable model which omits each Γ_n, and hence is atomic]

Exercise 60 Let $\mathcal{L}$ be the language of partial orders, together with a countably infinite set of new constants $c_0, c_1, \ldots$. Let T be the $\mathcal{L}$-theory

$$\text{DLO} \cup \{c_i < c_{i+1} \,|\, i \in \mathbb{N}\}$$

Describe (up to isomorphism) the countable models of T. Does T have an atomic modelΓ

Exercise 61 For those of you who are familiar with Gödel's Incompleteness Theorems.

a) Show that Peano Arithmetic has an atomic model: the *standard model* $\mathcal{N}$. Show that the model $\mathcal{N}$ is *not* elementarily embedded in every other model of PA. Why is this not in conflict with Theorem 9.3Γ

b) The theory RCF of real closed fields has an atomic model. What is itΓ Show, that RCF has uncountably many types in one variable. So the converse to 9.3c) is false.

I finish this chapter with a theorem giving a characterization of ω-categorical theories in terms of the number of types. Its proof uses atomic models.

Theorem 9.4 *Let T be a complete theory which has infinite models. Then T is ω-categorical if and only if for every $n \in \mathbb{N}$, T has only finitely many types in $x_1, \ldots, x_n$.*

Proof. First suppose T is ω-categorical; let $\mathfrak{A}$ be a countably infinite model of T. Every model of T is infinite (check!), hence has a countable elementary submodel by downwards Löwenheim-Skolem-Tarski; therefore, since T is ω-categorical, $\mathfrak{A}$ is elementarily embedded in every model of T. So $\mathfrak{A}$ is atomic by exercise 58.

Let $\Gamma(\vec{x})$ be a type of T. Then $\Gamma(\vec{x})$ is realized in some countable model of T, so it is realized in $\mathfrak{A}$ which is atomic; it follows that $\Gamma(\vec{x}) = \Gamma_{\vec{a}}$ for some $\vec{a} \in A$, so $\Gamma(\vec{x})$ is principal. So every type of T is principal.

Consider now the set

$$\Sigma(\vec{x}) = \{\neg\psi(\vec{x}) \mid \psi(\vec{x}) \text{ generates a type of } T\}$$

Then $\Sigma(\vec{x})$ cannot be extended to a type (because every type is principal). By the compactness theorem, there are finitely many formulas $\psi_1(\vec{x}), \ldots, \psi_k(\vec{x})$, generators of types, such that

$$T \models \forall\vec{x}\,(\psi_1(\vec{x}) \vee \cdots \vee \psi_k(\vec{x}))$$

We see that the types generated by the $\psi_i(\vec{x})$ are the *only* types of T, which therefore has only finitely many types.

Conversely, if T has only finitely many types in $\vec{x}$ (for each tuple $\vec{x}$), then each Boolean algebra $L_T(\vec{x})$ must be finite (for, if $T \not\models \forall\vec{x}\,(\varphi(\vec{x}) \to \psi(\vec{x}))$, there is a type which contains $\varphi(\vec{x})$ but not $\psi(\vec{x})$). It follows that every type of T is principal, so that every model of T is atomic. Then T is ω-categorical by Theorem 9.3c). ∎

Exercise 62 Show that a Boolean algebra B has only finitely many ultrafilters if and only if B itself is finite. Show moreover, that every infinite Boolean algebra has a nonprincipal ultrafilter.

Exercise 63 Describe the types of DLO in variables $x_1, \ldots, x_n$.

9.3 ω-Saturated Models

The notion of a *saturated* model is dual to the notion of atomic model: saturated models are 'large', they realize as many types as possible. In this chapter, where we focus on countable models, we restrict ourselves to the notion of ω-saturated model.

Definition 9.5 Let $\mathfrak{A}$ be an $\mathcal{L}$-structure. We say that $\mathfrak{A}$ is *ω-saturated* if for every finitely generated substructure $\mathfrak{B}$ of $\mathfrak{A}$ and every set of $\mathcal{L}_{\mathfrak{B}}$-formulas $\Gamma(x_1, \dots, x_n)$ in free variables $x_1, \dots, x_n$ the following holds: if Γ is realized in an $\mathcal{L}_{\mathfrak{B}}$-elementary extension of $\mathfrak{A}$, then Γ is realized in $\mathfrak{A}$.

Exercise 64 Show that every finite structure is ω-saturated.

Let's recall that if $\mathfrak{B}$ is generated by $a_1, \dots, a_m$, we may replace $\mathcal{L}_{\mathfrak{B}}$ by $\mathcal{L}_{\{a_1,\dots,a_m\}}$.

Exercise 65 Suppose $\mathfrak{A}$ and $\mathfrak{B}$ are $\mathcal{L}$-structures and $a_1, \dots, a_n \in A$, $b_1, \dots, b_n \in B$. We can then see $\mathfrak{A}$ and $\mathfrak{B}$ as $\mathcal{L} \cup \{c_1, \dots, c_n\}$-structures by putting $(c_i)^{\mathfrak{A}} = a_i$ and $(c_i)^{\mathfrak{B}} = b_i$ (of course, the constants c_i are new).

We write $\mathfrak{A}_{a_1 \cdots a_n} \equiv \mathfrak{B}_{b_1 \cdots b_n}$ if the $\mathcal{L} \cup \{c_1, \dots, c_n\}$-structures thus defined, are elementarily equivalent.

Prove: $\mathfrak{A}$ is ω-saturated if and obly if for each $\mathcal{L}$-structure $\mathfrak{B}$ the following hold:

1) If $\mathfrak{A} \equiv \mathfrak{B}$ then for every $b \in B$ there is an $a \in A$ such that $\mathfrak{A}_a \equiv \mathfrak{B}_b$;

2) for tuples $a_1, \dots, a_n \in A$ and $b_1, \dots, b_n \in B$ such that $\mathfrak{A}_{a_1 \cdots a_n} \equiv \mathfrak{B}_{b_1 \cdots b_n}$, and every $b_{n+1} \in B$, there is $a_{n+1} \in A$ such that $\mathfrak{A}_{a_1 \cdots a_{n+1}} \equiv \mathfrak{B}_{b_1 \cdots b_{n+1}}$.

Theorem 9.6 *Let T be a complete $\mathcal{L}$-theory.*

a) If $\mathfrak{A}$ is an ω-saturated model of T, every countable model of T is elementarily embedded in $\mathfrak{A}$;

b) T has a countable ω-saturated model if and only if for each $n \in \mathbb{N}$, T has only countably many types in $x_1, \dots, x_n$;

c) if T has a countable ω-saturated model, T has an atomic model.

Proof. a) Suppose $\mathfrak{A}$ is an ω-saturated model of T, and $\mathfrak{B}$ a countable model of T. Enumerate B as $\{b_1, b_2, \dots\}$. Since T is complete, $\mathfrak{A} \equiv \mathfrak{B}$. Therefore, applying Exercise 65, we can find a sequence $(a_n)_n$ in A such that for each n, $\mathfrak{A}_{a_1 \cdots a_n} \equiv \mathfrak{B}_{b_1 \cdots b_n}$. The assignment $b_n \mapsto a_n$ now defines an elementary embedding of $\mathfrak{B}$ into $\mathfrak{A}$.

b) If T has a countable ω-saturated model $\mathfrak{A}$, then there can be at most countably many types since every type is realized in $\mathfrak{A}$.

Conversely, suppose T has only countably many types in $x_1, \dots, x_n$ for each n. Let $\mathfrak{A}$ be a countable model of T. Then we can enumerate all pairs $(\vec{a}_i, \Gamma_i)$, where $\vec{a}_i$ is a finite (possibly empty) tuple $(a_{n_1}, \dots, a_{n_i})$ of elements of A, and Γ_i is a type in variables $(x_{n_1}, \dots, x_{n_i}, y_1, \dots, y_{k_i})$, such that the set

$$\Gamma_i(\vec{a}_i) = \{\gamma(a_{n_1}, \dots, a_{n_i}, y_1, \dots, y_{k_i}) \mid \gamma(\vec{x}, \vec{y}) \in \Gamma_i\}$$

is realized in some $\mathcal{L}_{\vec{a}_i}$-elementary extension of $\mathfrak{A}$. We construct a chain of models

$$\mathfrak{A} = \mathfrak{A}_0 \preceq \mathfrak{A}_1 \preceq \mathfrak{A}_2 \preceq \cdots$$

as follows: suppose $\mathfrak{A}_n$ is constructed. Let $\mathfrak{B}_{n+1}$ be a countable $\mathcal{L}_{\vec{a}_n}$-elementary extension of $\mathfrak{A}$ such that $\Gamma_n(\vec{a}_n)$ is realized in $\mathfrak{B}_{n+1}$. Then certainly $\mathfrak{A} \preceq \mathfrak{B}_{n+1}$ so $\mathfrak{A}_n \equiv \mathfrak{B}_{n+1}$; let $\mathfrak{A}_{n+1}$ be a countable common elementary extension of them.

One sees by induction that for each n, $\mathfrak{A}_{n+1}$ realizes $\Gamma_i(\vec{a}_i)$ for all $i \leq n$.

Let $\mathfrak{A}^{(1)}$ be the colimit of the chain $\mathfrak{A}_0 \preceq \mathfrak{A}_1 \preceq \cdots$. Then $\mathfrak{A}^{(1)}$ is countable, $\mathfrak{A} \preceq \mathfrak{A}^{(1)}$ and moreover the following property holds:

- For every finitely generated substructure $\mathfrak{B}$ of $\mathfrak{A}$ and every set $\Gamma(x_1, \dots, x_n)$ of $\mathcal{L}_{\mathfrak{B}}$-formulas which is realized in some $\mathcal{L}_{\mathfrak{B}}$-elementary extension of $\mathfrak{A}$, Γ is realized in $\mathfrak{A}^{(1)}$.

Now we iterate this procedure infinitely often:

$$\mathfrak{A} \preceq \mathfrak{A}^{(1)} \preceq \mathfrak{A}^{(2)} \preceq \cdots$$

where each $\mathfrak{A}^{(n)}$ is countable, and each $\mathfrak{A}^{(n+1)}$ has the property $\bullet$ w.r.t. $\mathfrak{A}^{(n)}$.

Let $\mathfrak{C}$ be the colimit of the chain $\mathfrak{A}^{(n)}$. Now it is easy to see that $\mathfrak{C}$ is countable and ω-saturated.

c) follows at once from b) and Theorem 9.3a). ■

Exercise 66 Show that if $\mathfrak{A}$ and $\mathfrak{B}$ are two countable ω-saturated models of a complete theory T, then $\mathfrak{A} \cong \mathfrak{B}$.

10 Stone Duality

In this chapter we shall establish a natural correspondence between Boolean algebras and a certain type of topological spaces, the so-called *Stone spaces*. The correspondence, when applied to the Lindenbaum algebra L_T of a theory T, or to the algebras $L_T(\vec{x})$, gives another perspective on the Compactness Theorem.

Let us first recall and collect the facts about filters on a Boolean algebra that we have already seen at different places, or that are easy to derive:

- A filter on a Boolean algebra B is an upwards closed, nonempty proper subset of B which is closed under finite meets (infima);
- an ultrafilter is a maximal filter;
- in a Boolean algebra, the notions of prime filter and ultrafilter coincide;
- if $a \not\leq b$ in a Boolean algebra B, there is an ultrafilter on B which contains a but not b;
- a filter U is an ultrafilter if and only if for each $b \in B$, exactly one element of $\{b, b^c\}$ belongs to U (here b^c denotes the complement of b);
- if $f : A \to B$ is a homomorphism of Boolean algebras and U is an ultrafilter on B, then $f^{-1}(U)$ is an ultrafilter on A;
- if A is a subset of a Boolean algebra B, there is an ultrafilter on B which contains A, if and only if for each finite subset $\{a_1, \ldots, a_n\}$ of A, $a_1 \sqcap \cdots \sqcap a_n \neq 0$.

Let B be a Boolean algebra. We denote the set of ultrafilters on B by $S(B)$. We give $S(B)$ a topology.

Let, for $b \in B$, U_b be $\{\mathcal{F} \in S(B) \mid b \in \mathcal{F}\}$.

I claim that the collection $\{U_b \mid b \in B\}$ forms a basis for a topology on $S(B)$. Indeed, it is clear that $S(B) = \bigcup_{b \in B} U_b$, and you can check for yourself that $U_b \cap U_{b'} = U_{b \sqcap b'}$.

We now examine the topological space $S(B)$ (with the topology generated by the sets U_b).

Proposition 10.1 *The space $S(B)$ is a compact Hausdorff space which has a basis of clopen (closed and open) sets.*

Proof. Every set U_b is clopen, since its complement is U_{b^c}. So $S(B)$ has a basis of clopen sets.

If $\mathcal{F} \neq \mathcal{G}$ in $S(B)$, there is $b \in B$ with $b \in \mathcal{F}$ and $b \notin \mathcal{G}$; so $\mathcal{F} \in U_b$, $\mathcal{G} \in U_{b^c}$, which are disjoint open neighborhoods of $\mathcal{F}$ and $\mathcal{G}$, respectively. So $S(B)$ is Hausdorff.

Finally, suppose the collection $\{U_{b_i} \mid i \in I\}$ covers $S(B)$; then every ultrafilter on B contains some b_i. It follows, that the set $\{(b_i)^c \mid i \in I\}$ is not a subset of any ultrafilter. But then we must have, for some $i_1, \ldots, i_n \in I$, that $b_{i_1} \sqcap \cdots \sqcap b_{i_n} = 0$ in B; but this implies that $\{U_{b_{i_1}}, \ldots, U_{b_{i_n}}\}$ covers $S(B)$. So $S(B)$ is compact. ∎

At this generality, Proposition 10.1 says all there is to say about $S(B)$. This follows from:

Proposition 10.2 *Let X be a nonempty, compact Hausdorff space which has a basis of clopen sets. Then there is a Boolean algebra $B(X)$, unique up to isomorphism, such that X is homeomorphic to $S(B(X))$. Consequently, if B is a Boolean algebra, then $B(S(B))$ is isomorphic to B.*

Proof. For any topological space X, the set $B(X)$ of clopen subsets of X contains $\emptyset$ and X and is closed under finite unions and intersections as well as under taking complements. Therefore it is a Boolean algebra.

Now consider an ultrafilter $\mathcal{F}$ on $B(X)$. In particular this is a collection of closed subsets of X which is closed under finite intersections, so if X is compact, the intersection $\bigcap \mathcal{F}$ is nonempty, and if X is Hausdorff, the intersection is a singleton $\{x\}$, by maximality of the filter. So, we have a function $h : S(B(X)) \to X$ which is easily seen to be a bijection. It is also continuous, for if $h(\mathcal{F}) = x$ then $\mathcal{F} = \{U \subseteq X \,|\, U \text{ clopen}, x \in U\}$; so for a clopen basiselement A we have

$$h^{-1}(A) = \{\mathcal{F} \in S(B(X)) \,|\, A \in \mathcal{F}\} = U_A$$

which is a basic open in $S(B(X))$. Since both X and $S(B(X))$ are compact Hausdorff, we are done and h is an homeomorphism.

Now let's look at $B(S(B))$. Clearly, there is a homomorphism of Boolean algebras: $B \to B(S(B))$ given by $b \mapsto U_b$. This is an injective map. Now take any clopen subset W of $S(B)$. Then both W and its complement are unions of clopen basis elements. Since $S(B)$ is compact, we must have $W = U_{b_1} \cup \cdots \cup U_{b_n}$ for some $b_1, \ldots, b_n \in B$; so $W = U_{b_1 \sqcup \cdots \sqcup b_n}$ and we see that the map $B \to B(S(B))$ is also surjective. It is easy to see that its inverse is also a homomorphism of Boolean algebras. ∎

Definition 10.3 A *Stone space* is a compact Hausdorff space with a basis of clopen sets.

Exercise 67 (Examples) i) Show that Cantor space is a Stone space. Show, that under the correspondence of proposition 10.2, it corresponds to the Boolean algebra $\mathcal{P}(\mathbb{N})$.

ii) Show that the set $\{0\} \cup \{\frac{1}{n} \,|\, n \in \mathbb{N}, n > 0\}$, viewed as a subspace of $\mathbb{R}$, is a Stone space. To which Boolean algebra does it correspondΓ

We shall now extend the correspondence of proposition 10.2 also to maps.

Suppose $f : X \to Y$ is a continuous function between topological spaces. Then $f^{-1} : \mathcal{P}(Y) \to \mathcal{P}(X)$ maps clopen sets to clopen sets, and commutes with finite unions and intersections, as well as complements. So it is a homomorphism of Boolean algebras: $B(Y) \to B(X)$.

Now if X and Y are Stone spaces, *every* homomorphism $\varphi : B(Y) \to B(X)$ is f^{-1} for a unique continuous $f : X \to Y$. For, we have seen that points of a Stone space are in 1–1 correspondensce with ultrafilters on its associated Boolean algebra. Since φ^{-1} sends ultrafilters on $B(X)$ to ultrafilters on $B(Y)$, it determines a map $f : X \to Y$ such that $\varphi = f^{-1}$, and it is left to you to see that f is continuous.

Summing up: we have established that for Stone spaces X and Y there is a 1-1 correspondence between continuous maps $X \to Y$ and Boolean homomorphisms $B(Y) \to B(X)$. If we take $Y = X$, the the identity map on X corresponds to the identity map on $B(X)$; and given $X \xrightarrow{f} Y \xrightarrow{g} Z$, the composition of f and g corresponds to the composition of the maps corresponding to g and f. In the language of Category Theory, we say that the categories of Stone spaces and Boolean algebras are *dually equivalent*, or *dual to each other* (the word 'dual' refers to the fact that the direction of the arrows is reversed). For the record:

Theorem 10.4 (Stone Duality Theorem) *The category of Stone spaces is dual to the category of Boolean algebras.*

Exercise 68 The Boolean algebra $\mho = \{0, 1\}$ has very special features:

a) for every Boolean algebra B, ultrafilters on B are in 1-1 correspondence with Boolean homomorphisms $B \to \mho$;

b) for every Boolean algebra B there is exactly one Boolean homomorphism $\mho \to B$.

Applying Stone Duality, interpret these facts for the corresponding Stone space.

Exercise 69 Let T be a theory and consider the Boolean algebra L_T. Describe the points of the Stone space corresponding to L_T, in terms of T.

Exercise 70 Show, without using Zorn's Lemma or any of its equivalents, that the Compactness Theorem is equivalent to the statement that each of the spaces $S(L_T)$ is compact.

Exercise 71 Let X be an arbitrary topological space. Show that there is a Stone space $T(X)$ and a continuous map $\eta : X \to T(X)$ such that the following holds: for every continuous function f from X to a Stone space Y, there is a unique continuous function $\tilde{f} : T(X) \to Y$ such that $\tilde{f} \circ \eta = f$.

11 Literature

There are two standard reference works for general Model Theory:

C.C. Chang and H.J. Keisler, *Model Theory*, Amsterdam, North Holland, 3rd edition 1990

W. Hodges, *Model Theory*, Cambridge, Cambridge University Press, 1993

Both books are expensive and not very suitable for a first course. The first is moreover, despite the additions in the 1990 edition, rather old-fashioned.

A nice, compact introduction is

K. Doets, *Basic Model theory*, Stanford, CSLI Publications, 1996

However, its emphasis is quite different from that of these notes. Has nothing on quantifier elimination or model completeness.

General Mathematical Logic books, such as Shoenfield's *Mathematical Logic* (a classic), or Bell & Machover's *A Course in Mathematical logic*, usually have a chapter on Model Theory where basic notions are given.

Material on Real Closed Fields can be found in

N. Jacobson, *Lectures in Abstract Algebra*, vol. III, Princeton, Van Nostrand, 1964

The proof of Quantifier Elimination for Real Closed Fields in section 8.1 is based on the paper

P.J. Cohen, *Decision Procedures for Real and p-Adic Fields*, in: *Communications on Pure and Applied Mathematics* XXII (1969), pp. 131–151

For a detailed proof of Hilbert's Nullstellensatz, look at

H. Matsumura, *Commutative Ring Theory*, Cambridge University Press, 1989

There are several directions in which to study Model Theory further. The hottest fields are *Stability Theory* and *o-Minimality Theory*. A few introductory texts:

A. Pillay, *Geometric Stability Theory*, New York, Clarendon Press, 1996

D. Marker, M. Messmer and A. Pillay, *Model Theory of Fields*, Berlin, Springer, 1996

L. van den Dries, *Tame Topology and o-minimal structures*, Cambridge University Press 1998

However, without some further introduction in Model Theory it is not advisable to start on these books right away.

A nice and accessible specialization of Model Theory is the topic of models of Peano Arithmetic. Two texts:

R. Kaye, *Models of Peano Arithmetic*, Oxford, Clarendon, 1991

J. van Oosten, *Introduction to Peano Arithmetic – Gödel Incompleteness and Nonstandard Models*, Communications of the Mathematical Institute 21–1999, Utrecht University, 1999

Index