Syllabus Model Theory 2016/2017

Benno van den Berg

CHAPTER 1

Basic definitions

1. On language and interpretation

Definition 1.1. A *language* or *signature* L consists of:

(1) a set of constants.
(2) a set of function symbols, each with an arity $n \in \mathbb{N}$.
(3) a set of relation symbols, each with an arity $n \in \mathbb{N}$.

Once and for all, we fix a countably infinite set of variables.

Definition 1.2. The *terms* in a signature L are the smallest set of expressions such that:

(1) all constants are terms.
(2) all variables are terms.
(3) if $t_1, \ldots, t_n$ are terms and f is an n-ary function symbol, then also $f(t_1, \ldots, t_n)$ is a term.

Terms which do not contain any variables are called *closed*.

Definition 1.3. An *atomic formula* is an expression of the form

(1) $s = t$, where s and t are terms, or
(2) $P(t_1, \ldots, t_n)$, where $t_1, \ldots, t_n$ are terms and P is a n-ary relation symbol.

Definition 1.4. The set of *formulas* is the smallest set of expressions which:

(1) contains the atomic formulas.
(2) contains $\varphi \wedge \psi, \varphi \vee \psi, \varphi \to \psi, \neg\varphi$ whenever φ and ψ are formulas.
(3) contains $\exists x\, \varphi$ and $\forall x\, \varphi$, if φ is a formula.

A formula which does not contain any quantifiers, so can be obtained by applying rules (1) and (2) only, is called *quantifier-free*. A *sentence* is a formula which does not contain any free variables. A set of sentences is called a *theory*.

We will often write $\varphi(x_1, \ldots, x_n)$ instead of φ. The notation $\varphi(x_1, \ldots, x_n)$ is meant to indicate that φ is a formula whose free variables are contained in $\{x_1, \ldots, x_n\}$.

Definition 1.5. A *structure* or *model* M in a language L consists of:

(1) a non-empty set M (the *domain* or the *universe*).
(2) interpretations $c^M \in M$ of all the constants in L,
(3) interpretations $f^M \colon M^n \to M$ of all n-ary function symbols in L,
(4) interpretations $R^M \subseteq M^n$ of all n-ary relation symbols in L.

If $A \subseteq M$, then we will write L_A for the language obtained by adding to L fresh constants $\{c_a : a \in A\}$. In this case M could also be considered an L_A-structure in which c_a is interpreted as a. We will often just write a instead of c_a (!!).

If M is a model then the interpretation in M of constants in the language L_M can be extended to all closed terms in the language L_M by putting:

$$f(t_1, \ldots, t_n)^M = f^M(t_1^M, \ldots, f_n^M).$$

Definition 1.6. If M is a model in in the language L and φ is a sentence in the language L_M, then we will write:

- $M \models s = t$ if $s^M = t^M$;
- $M \models P(t_1, \ldots, t_n)$ if $(t_1^M, \ldots, t_n^M) \in P^M$;
- $M \models \varphi \wedge \psi$ if $M \models \varphi$ and $M \models \psi$;
- $M \models \varphi \vee \psi$ if $M \models \varphi$ or $M \models \psi$;
- $M \models \varphi \to \psi$ if $M \models \varphi$ implies $M \models \psi$;
- $M \models \neg\varphi$ if not $M \models \varphi$;
- $M \models \exists x\, \varphi(x)$ if there is an $m \in M$ such that $M \models \varphi(m)$;
- $M \models \forall x\, \varphi(x)$ if for all $m \in M$ we have $M \models \varphi(m)$.

If $M \models \varphi$ we say that φ *holds in* M or *is true in* M.

Definition 1.7. If M is a model in a language L, then $\mathrm{Th}(M)$ is the collection of all L-sentences true in M. If N is another model in the language L, then we write $M \equiv N$ and call M and N *elementarily equivalent*, whenever $\mathrm{Th}(M) = \mathrm{Th}(N)$.

Definition 1.8. Let Γ and Δ be theories. If $M \models \varphi$ for all $\varphi \in \Gamma$, then M is called a *model* of Γ. We will write $\Gamma \models \Delta$ if every model of Γ is a model of Δ as well. We write $\Gamma \models \varphi$ for $\Gamma \models \{\varphi\}$ and $\varphi \models \psi$ for $\{\varphi\} \models \{\psi\}$.

Definition 1.9. If $L \subseteq L'$ and M is an L'-structure, then we can obtain an L-structure N by taking the universe of M and forgetting the interpretations of the symbols which do not occur in L. In that case, M is an *expansion* of N and N is the *L-reduct* of M.

Lemma 1.10. *If $L \subseteq L'$ and M is an L'-structure and N is its L-reduct, then we have $N \models \varphi(m_1, \ldots, m_n)$ iff $M \models \varphi(m_1, \ldots, m_n)$ for all formulas $\varphi(x_1, \ldots, x_n)$ in the language L and all elements $m_1, \ldots, m_n$ from M.*

2. Morphisms

Any structure in mathematics comes with a notion of homomorphism: a mapping preserving that structure.

Definition 1.11. Let M and N be two L-structures. A *homomorphism* $h\colon M \to N$ is a function $h\colon M \to N$ such that:

(1) $h(c^M) = c^N$ for all constants c in L;
(2) $h(f^M(m_1, \ldots, m_n)) = f^N(h(m_1), \ldots, h(m_n))$ for all function symbols f in L and elements $m_1, \ldots, m_n \in M$;
(3) $(m_1, \ldots, m_n) \in R^M$ implies $(h(m_1), \ldots, h(m_n)) \in R^N$.

A homomorphism h which is bijective and whose inverse h^{-1} is a homomorphism as well is called an *isomorphism*. If there exists an isomorphism between structures M and N, then M and N are called *isomorphic*. An isomorphism from a structure to itself is called an *automorphism*.

Actually, in model theory the general notion of homomorphism turns out to of limited usefulness. More important are the embeddings.

Definition 1.12. A homomorphism $h: M \to N$ is an *embedding* if

(1) h is injective;
(2) $(h(m_1), \ldots, h(m_n)) \in R^N$ implies $(m_1, \ldots, m_n) \in R^M$.

Lemma 1.13. *The following are equivalent for a homomorphism $h: M \to N$:*

(i) *h is an embedding.*
(ii) *$M \models \varphi(m_1, \ldots, m_n) \Leftrightarrow N \models \varphi(h(m_1), \ldots, h(m_n))$ for all $m_1, \ldots, m_n \in M$ and atomic formulas $\varphi(x_1, \ldots, x_n)$.*
(iii) *$M \models \varphi(m_1, \ldots, m_n) \Leftrightarrow N \models \varphi(h(m_1), \ldots, h(m_n))$ for all $m_1, \ldots, m_n \in M$ and quantifier-free formulas $\varphi(x_1, \ldots, x_n)$.*

Definition 1.14. If M and N are two models and the inclusion $M \subseteq N$ is an embedding, then M is a *substructure* of N and N is an *extension* of M.

But the most important notion of morphism in model theory is that of an elementary embedding.

Definition 1.15. An embedding $h: M \to N$ is called *elementary*, if

$$M \models \varphi(m_1, \ldots, m_n) \Leftrightarrow N \models \varphi(h(m_1), \ldots, h(m_n))$$

for all $m_1, \ldots, m_n \in M$ and *all* formulas $\varphi(x_1, \ldots, x_n)$.

Remark 1.16. In the definition of an elementary embedding the equivalence

$$M \models \varphi(m_1, \ldots, m_n) \Leftrightarrow N \models \varphi(h(m_1), \ldots, h(m_n))$$

holds as soon as the implication from left to right or from right to left holds. (Why? *Hint:* Negation!) A similar remark applies to point (iii) of Lemma 1.13.

Lemma 1.17. *Any isomorphism $h: M \to N$ is also an elementary embedding. If $h: M \to N$ is an elementary embedding, then $M \equiv N$.*

3. Exercises

Exercise 1. A theory T is *consistent* if it has a model and *complete* if it is consistent and for any formula φ we have

$$T \models \varphi \quad \text{or} \quad T \models \neg\varphi.$$

Show that the following are equivalent for a consistent theory T:

(1) T is complete.
(2) All models of T are elementarily equivalent.
(3) There is a structure M such that T and $\mathrm{Th}(M)$ have the same models.

Exercise 2. An element a in an L-structure M is *definable* if there is an L-formula $\varphi(x)$ such that for any $m \in M$

$$M \models \varphi(m) \Leftrightarrow a = m.$$

(a) What are the definable elements in $(\mathbb{N}, +)$? And in $(\mathbb{Z}, +)$? Justify your answers.
(b) Is the embedding $(\mathbb{N}, +) \subseteq (\mathbb{Z}, +)$ elementary? And the embedding $(\mathbb{N}, \cdot) \subseteq (\mathbb{Z}, \cdot)$? And the embedding $(\mathbb{Z}, \cdot) \subseteq (\mathbb{Q}, \cdot)$? And the embedding $(\mathbb{Q}, \cdot) \subseteq (\mathbb{R}, \cdot)$? And the embedding $(\mathbb{R}, \cdot) \subseteq (\mathbb{C}, \cdot)$?

EXERCISE 3. (For the algebraists.) Let $L_r = \{0, 1, +, -, \cdot\}$ be the language of (unital) rings with binary operations $+$ and $\cdot$, a unary operation $-$ and constants $0, 1$. Let CR be the theory of commutative rings, saying that both $+$ and $\cdot$ are associative and commutative with units 0 and 1, respectively, plus an axiom saying that $-x$ is an additive inverse for x and the distributive law $x \cdot (y + z) = x \cdot y + x \cdot z$. The theory ID of integral domains is the theory CR together with the axioms $0 \neq 1$ and $\forall x \forall y\, (\, x \cdot y = 0 \to x = 0 \vee y = 0)$, while the theory F of fields is the theory CR together with $0 \neq 1$ and $\forall x\, (\, x \neq 0 \to \exists y\, x \cdot y = 1\,)$.

(a) A *universal sentence* is one of the form $\forall x_1, \ldots, x_n \varphi(x_1, \ldots, x_n)$ where $\varphi(x_1, \ldots, x_n)$ is quantifier-free. A theory T can be *axiomatised using universal sentences* if there is a collection of universal sentences S such that S and T have the same models.

Show that CR and ID can be axiomatised using universal sentences, while this is impossible for F. *Hint:* Check that universal sentences are preserved by substructures.

(b) Write $T_\forall = \{\varphi \colon T \models \varphi$ and φ is universal$\}$. Show that $F_\forall$ and ID have the same models. *Hint:* Use that any integral domain can be embedded into a field (its field of fractions) by mimicking the construction of $\mathbb{Q}$ out of $\mathbb{Z}$.

EXERCISE 4. Let L be signature and M and N be two L-structures. Show that if M is finite and M and N are elementarily equivalent, then M and N are isomorphic. *Hint:* You may find it helpful to first think about the special case where the language L is finite.

CHAPTER 2

Compactness theorem

The most important result in model theory is:

THEOREM 2.1. *Let T be a theory in language L. If every finite subset of T has a model, then T has a model.*

I suspect many of you have seen a proof of this already. In fact, it is often obtained as a direct corollary of the completeness theorem for first-order logic. But one can give a purely model-theoretic proof (without any proof calculus in sight) and such a proof will be sketched below.

1. A proof

For convenience let us temporarily call a theory T *finitely consistent* if any finite subset of T has a model. The goal is to show that finitely consistent theories are consistent (that is, have a model). The first step is to reduce the problem to showing that maximal finitely consistent theories have models.

DEFINITION 2.2. A theory T in a language L is maximal finitely consistent if there is no finitely consistent L-theory T' with $T \subset T'$ (in other words, adding a new sentence to T destroys its finite consistency).

The following is a direct consequence of Zorn's Lemma (see below).

LEMMA 2.3. *Any finitely consistent L-theory T can be extended to a maximal finitely consistent L-theory T'.*

PROOF. Consider the collection P of all finitely consistent L-theories which extend T and order P by inclusion. Since every chain X in P has an upper bound (simply take the union of all theories in X), Zorn's Lemma tells us that P has a maximal element. Such a maximal element is a maximal finitely consistent theory T' extending T. □

LEMMA 2.4. *Let T be maximal finitely consistent L-theory.*

(1) *For any sentence φ the theory T contains either φ or $\neg\varphi$.*
(2) *If T_0 is a finite subset of T and $T_0 \models \varphi$, then $\varphi \in T$.*

PROOF. (i): Suppose T is a maximal finitely consistent L-theory and $\varphi \notin T$. Since T was maximal, $T \cup \{\varphi\}$ cannot be finitely consistent, so there is a finite subset $T_2 \subseteq T$ such that $T_2 \cup \{\varphi\}$ has no models.

We want to show that $\neg\varphi \in T$. For this it suffices to prove that $T \cup \{\neg\varphi\}$ is finitely consistent; indeed, this can only be compatible with the maximality of T if $T \cup \{\neg\varphi\} = T$, or, in other words, if $\neg\varphi \in T$.

To see that $T \cup \{\neg\varphi\}$ is finitely consistent, let $T_0 \subseteq T \cup \{\neg\varphi\}$ be finite. Then T_0 is a subset of a set of form $T_1 \cup \{\neg\varphi\}$ with T_1 a finite subset of T.

Consider $T_1 \cup T_2$. This is a finite subset of T and since T is finitely consistent, the set $T_1 \cup T_2$ has a model M. Because M is a model of T_2, it cannot be a model of φ. So $M \models T_1$ and $M \models \neg\varphi$. Hence M is a model of T_0 and since T_0 was an arbitrary finite subset of $T \cup \{\neg\varphi\}$, we have shown that $T \cup \{\neg\varphi\}$ is finitely consistent, as desired.

(ii): Assume T_0 is a finite subset of a maximal finitely consistent L-theory T and $T_0 \models \varphi$. I claim that $\varphi \in T$. For if $\varphi \notin T$, then $\neg\varphi \in T$ by (i). But then $T_0 \cup \{\neg\varphi\}$ is a finite subset of T, so has a model M. But then M is a model of T_0 in which φ does not hold, contradiction. $\square$

Proposition 2.5. *Suppose T is a finitely consistent theory in a language L and C is a set of constants in L. If for any formula $\psi(x)$ in the language L there is a constant $c \in C$ such that*

$$\exists x\, \psi(x) \to \psi(c) \in T,$$

then T has a model whose universe consists entirely of interpretations of constants in C.

Proof. In view of Lemma 2.3 it suffices to prove the statement for maximal finitely consistent T. In this case we construct a model M by taking the closed terms in L and identifying closed terms s and t whenever the expression $s = t$ belongs to T: it follows from part (ii) of the previous lemma that this is an equivalence relation.

We have to show how to interpret constants as well as function and relation symbols in M. If c is any constant in L, then we put $c^M = [c]$, whilst if f is any n-ary function symbol and $t_1, \ldots, t_n$ are closed L-terms, then we set

$$f^M([t_1], \ldots, [t_n]) := [f(t_1, \ldots, t_n)].$$

Another appeal to part (ii) of the previous lemma is needed to show that this is well-defined.

Finally, if R is an n-ary relation symbol, then we will say that $([t_1], \ldots, [t_n]) \in R^M$ in case $R(t_1, \ldots, t_n) \in T$. Part (ii) of the previous lemma should again to be used to justify this definition.

Now one can easily show by induction on the structure of the term t that $t^M = [t]$ and the structure of the formula φ that $M \models \varphi$ if and only if $\varphi \in T$. In short, M is a model of T.

It remains to verify that any element in M is an interpretation of a constant $c \in C$. We know that any element in M is of the form $[t]$ for some closed term t. By assumption there exists an element $c \in C$ for which the sentence

$$\exists x\, (x = t) \to c = t$$

belongs to T. Since $\exists x\, (x = t)$ is a tautology, it also belongs to T and therefore we have $c = t \in T$ as well. So $M \models c = t$ and $c^M = t^M = [t]$. $\square$

Lemma 2.6. *Suppose T is a finitely consistent L-theory. Then L can be extended to a language L' and T to a finitely consistent L'-theory T' such that for any L'-formula $\varphi(x)$ there is a constant c in L' such that*

$$T' \models \exists x\, \varphi(x) \to \varphi(c).$$

Proof. We define by induction a sequence of languages L_n and L_n-theories T_n. We start by putting $L_0 = L$ and $T_0 = T$.

If L_n and T_n have been defined, we obtain L_{n+1} by adding to L_n a fresh constant c_φ for any L_n-formula $\varphi(x)$. Moreover, T_{n+1} is obtained by adding to T_n for any L_n-formula $\varphi(x)$ the sentence

$$\exists x\, \varphi(x) \to \varphi(c_\varphi).$$

One easily proves by induction on n that each T_n is finitely consistent.

Finally, we put $L' = \bigcup_{n\in\mathbb{N}} L_n$ and $T' = \bigcup_{n\in\mathbb{N}} T_n$. Then T' is finitely consistent (see exercise 5 below). Moreover, any L'-formula $\varphi(x)$ is already an L_n-formula for some n (see again exercise 5 below). So

$$\exists x\, \varphi(x) \to \varphi(c_\varphi) \in T_{n+1} \subseteq T,$$

as desired. □

Theorem 2.7. (Compactness Theorem) *Let T be a theory in language L. If every finite subset of T has a model, then T has a model.*

Proof. Let T be a finitely consistent L-theory. Combining the previous lemma with the previous proposition, one sees that L can be extended to a language L' and T to an L'-theory T' such that T' has a model M. So if N is the reduct of M to L, then N is a model of T by Lemma 1.10. □

2. Appendix: statement of Zorn's Lemma

Definition 2.8. A *partial order* is a set P together with a binary relation $\leq$ which is

(i) reflexive, so $x \leq x$ for any $x \in P$.
(ii) anti-symmetric, so $x \leq y$ and $y \leq x$ imply $x = y$.
(iii) transitive, so $x \leq y$ and $y \leq z$ imply $x \leq z$.

A subset $X \subseteq P$ is called a *chain* if for any two elements $x, y \in X$ we have either $x \leq y$ or $y \leq x$. An *upper bound* for a set $X \subseteq P$ is an element $y \in P$ such that $x \leq y$ for all $x \in X$. An element $x \in P$ is *maximal* if $x \leq y$ implies $x = y$.

Lemma 2.9. (Zorn's Lemma) *Let $(P, \leq)$ be a partial order and assume that any chain in P has an upper bound. Then P contains at least one maximal element.*

Proof. A proof can be found in most textbooks on set theory (for example, on page 114 of Moschovakis, *Notes on Set Theory*, second edition, Springer-Verlag, 2006). □

3. Exercises

Exercise 5. (a) Let $A_0 \subseteq A_1 \subseteq A_2 \subseteq \ldots$ be an increasing sequence of sets, and write $A := \bigcup_{n\in\mathbb{N}} A_n$. Show that any finite subset of A is already a finite subset of some A_n.

(b) Suppose that $L_0 \subseteq L_1 \subseteq L_2 \subseteq \ldots$ is an increasing sequence of languages and $L = \bigcup_{n\in\mathbb{N}} L_n$. Show that any L-formula is also an L_n-formula for some n.

(c) Suppose that $T_0 \subseteq T_1 \subseteq T_2 \subseteq \ldots$ is an increasing sequence of finitely consistent theories. Prove that $\bigcup_{n\in\mathbb{N}} T_n$ is finitely consistent as well.

EXERCISE 6. A class of models $\mathcal{K}$ in some fixed signature is called an *elementary class* if there is a first-order theory such that $\mathcal{K}$ consists of precisely those L-structures that are models of T.

Show that if $\mathcal{K}$ is a class of L-structures and both $\mathcal{K}$ and its complement (in the class of all L-structures) are elementary, then there is a sentence φ such that M belongs to $\mathcal{K}$ if and only if $M \models \varphi$.

EXERCISE 7. We work over the empty language L (no constants, function or relations symbols). Show that the class of infinite L-structures is elementary, but the class of finite L-structures is not. Deduce that there is no sentence φ that is true in an L-structure if and only if the L-structure is infinite.

CHAPTER 3

Method of diagrams

This chapter is devoted to applications of the compactness theorem. One application is to show the dramatic failure of first-order logic to distinguish between different cardinalities: we will show, for instance, that if a first-order theory T in some countable language has an infinite model, then T has models of all infinite sizes. To show this, we use the method of diagrams.

1. Diagrams

Definition 3.1. If M is a model in a language L, then the collection of quantifier-free L_M-sentences true in M is called the *diagram* of M and written $\mathrm{Diag}(M)$. The collection of all L_M-sentences true in M is called the *elementary diagram* of M and written $\mathrm{ElDiag}(M)$.

Lemma 3.2. *The following amount to the same thing:*

- *A model N of* $\mathrm{Diag}(M)$.
- *An embedding $h: M \to N$.*

As do the following:

- *A model N of* $\mathrm{ElDiag}(M)$.
- *An elementary embedding $h: M \to N$.*

Proof. I suspect that a genuine proof of this lemma would only obscure the main point. The task is to reflect on the question what it would mean to give a model of $\mathrm{Diag}(M)$. It would involve finding a model N and assigning to each constant c_m an interpretation in N in such a way that if φ is quantifier-free and $\varphi(c_{m_1}, \ldots, c_{m_n})$ is true in M, then it is true in N as well. This is the same thing as giving an embedding $h: M \to N$ (see also Lemma 1.13). A similar reflection should make the second point of the lemma clear. $\square$

2. The Łoś-Tarski Theorem

As a first indication of the usefulness of the method of diagrams, we will prove a characterisation theorem for universal theories.

Definition 3.3. A sentence is *universal* if it starts with a string of universal quantifiers followed by a quantifier-free formula. A theory is *universal* if it consists of universal sentences. A theory has a *universal axiomatisation* if it has the same class of models as a universal theory in the same language.

Theorem 3.4. (The Łoś-Tarski Theorem) *T has a universal axiomatisation iff models of T are closed under substructures.*

PROOF. It is easy to see that models of a universal theory are closed under substructures, so we concentrate on the other direction. So let T be a theory such that its models are closed under substructures. Write

$$T_\forall = \{\varphi\colon T \models \varphi \text{ and } \varphi \text{ is universal }\}.$$

Clearly, $T \models T_\forall$. We need to prove the converse.

So suppose M is a model of $T_\forall$. Now it suffices to show that $T \cup \mathrm{Diag}(M)$ is consistent. Because once we do that, it will have a model N. But since N is a model of $\mathrm{Diag}(M)$, it will be an extension of M; and because N is a model of T and models of T are closed under substructures, M will be a model of T.

So the theorem will follow once we show that $T \cup \mathrm{Diag}(M)$ is consistent. We argue by contradiction: so suppose $T \cup \mathrm{Diag}(M)$ would be inconsistent. Then, by the compactness theorem, there are quantifier-free formulas $\psi_1(\overline{m}_1), \ldots, \psi_n(\overline{m}_n) \in \mathrm{Diag}(M)$ which are inconsistent with T. Write $\psi(\overline{m}) := \psi_1(\overline{m}_1) \wedge \psi_2(\overline{m}_2) \wedge \ldots \wedge \psi_n(\overline{m}_n)$. Then $\psi(\overline{m})$ is a single formula from $\mathrm{Diag}(M)$ inconsistent with T.

Replace the constants $\overline{m}$ from M in ψ by variables $\overline{x}$ and consider the sentence $\exists \overline{x}\, \psi(\overline{x})$; because the constants from M do not appear in T, the theory T is already inconsistent with $\exists \overline{x}\, \psi(\overline{x})$ (see Exercise 8 below). Therefore $T \models \neg\exists \overline{x}\, \psi(\overline{x})$ and $T \models \forall \overline{x}\, \neg\psi(\overline{x})$; in other words, $\forall \overline{x} \neg\psi(\overline{x}) \in T_\forall$. Since M is a model of $T_\forall$, it follows that $M \models \forall \overline{x}\, \neg\psi(\overline{x})$ and $M \models \neg\psi(\overline{m})$. This contradicts $\psi(\overline{m}) \in \mathrm{Diag}(M)$. □

3. The theorems of Skolem and Löwenheim

As another application of the compactness theorem we can show that first-order logic is unable to see the difference between different infinite cardinalities. Two theorems due to Skolem and Löwenheim make this point in a very clear way.

DEFINITION 3.5. The *cardinality* of a model is the cardinality of its underlying domain. The *cardinality* of a language L is the sum of the cardinalities of its sets of constants, function symbols and relation symbols.

We will write:

- $|X|$ for the cardinality of a set X,
- $|M|$ for the cardinality of a model M, and
- $|L|$ for the cardinality of a language L.

3.1. Downward.

To prove the first theorem due to Skolem and Löwenheim we need a test for recognising elementary embeddings.

THEOREM 3.6. (Tarski-Vaught Test) *An embedding $h\colon M \to N$ is elementary if and only if for any L_M-formula $\varphi(x)$: if $N \models \exists x\, \varphi(x)$, then there is an element $m \in M$ such that $N \models \varphi(h(m))$.*

PROOF. Let us first check the necessity of the condition: if $h\colon M \to N$ is an elementary embedding and $\varphi(x)$ is an L_M-formula such that $N \models \exists x\, \varphi(x)$, then $M \models \exists x\, \varphi(x)$ as well. So there is an element $m \in M$ such that $M \models \varphi(m)$ and hence $N \models \varphi(h(m))$, because h is elementary.

Conversely, suppose that the condition is satisfied and we wish to prove that

$$M \models \varphi(\overline{m}) \Leftrightarrow N \models \varphi(h(\overline{m}))$$

for any L-formula φ and any tuple $\overline{m}$ of parameters from M. The idea is to prove this bi-implication by induction on the structure of φ. To make our lives easier we will assume that the only logical connectives appearing in φ are $\wedge, \neg$ and $\exists$: since every first-order formula is logically equivalent to one only containing these connectives, we may do this without loss of generality.

Let us start by noting that the desired equivalence is valid for atomic formulas, since h is an embedding (see Lemma 1.13). The induction cases for $\wedge$ and $\neg$ are trivial, so we are left with the case of $\exists x\, \psi(x, \overline{m})$. The induction hypothesis is

$$M \models \psi(m, \overline{m}) \Leftrightarrow N \models \psi(h(m), h(\overline{m}))$$

for all $m, \overline{m} \in M$. Then:

$$\begin{array}{c} M \models \exists x\, \psi(x, \overline{m}) \Leftrightarrow (\exists m \in M)\, M \models \exists x\, \psi(m, \overline{m}) \Leftrightarrow \\ (\exists m \in M)\, N \models \psi(h(m), h(\overline{m})) \Leftrightarrow N \models \exists x\, \varphi(x, \overline{m}). \end{array}$$

(Here we have used the condition in the right to left direction of the last bi-implication.) □

THEOREM 3.7. (Downwards Skolem-Löwenheim Theorem) *Suppose M is an L-structure and $X \subseteq M$. Then there is an elementary substructure N of M with $X \subseteq N$ and $|N| \leq |X| + |L| + \aleph_0$.*

PROOF. We construct N as $\bigcup_{i \in \mathbb{N}} N_i$ where the N_i are defined inductively as follows: we start by putting $N_0 = X$, while

- if i is even, then N_{i+1} is obtained from N_i by adding the interpretations of the constants and closing under f^M for every function symbol f (that is, we add all elements of the form $f^M(n_1, \ldots, n_k)$ with f an k-ary function symbol in L and $n_1, \ldots, n_k \in N_i$).
- if i is odd, we look at all L_{N_i}-sentences of the form $\exists x\, \varphi(x)$. If such a sentence is true in M, then we pick a witness $n \in M$ such that $M \models \varphi(n)$ and put it in N_{i+1}.

Then the first item guarantees that N is a substructure, while the second item ensures that it is an elementary substructure (using the Tarski-Vaught test). □

3.2. Upward. To find larger models we again use the method of diagrams.

THEOREM 3.8. (Upwards Skolem-Löwenheim Theorem) *Suppose M is an infinite L-structure and κ is a cardinal number with $\kappa \geq |M|, |L|$. Then there is an elementary embedding $i \colon M \to N$ with $|N| = \kappa$.*

PROOF. Let Γ be the elementary diagram of M and Δ be the set of sentences $\{c_i \neq c_j : i \neq j \in \kappa\}$ where the c_i are κ-many fresh constants. M is a model of any finite subset of $\Gamma \cup \Delta$: indeed, in any finite subset of $\Gamma \cup \Delta$ only finitely many fresh constants c_i occur; the idea is to interpret the c_i as different elements in M (which we can always do since the model M is infinite). Therefore, by the Compactness Theorem, the theory $\Gamma \cup \Delta$ has a model A. By construction M is an elementary substructure of A and $|A| \geq \kappa$. By the downward Downwards Skolem-Löwenheim Theorem A has an L_M-elementary substructure N of cardinality κ. Since N is still a model of the elementary diagram of M, there is an elementary embedding $i \colon M \to N$. □

4. Exercises

EXERCISE 8. Assume T is a theory and $\varphi(x)$ is a formula in which the constant c does not occur.

(a) Prove: $T \models \varphi(c)$ iff $T \models \forall x\, \varphi(x)$.
(b) Prove: T is consistent with $\varphi(c)$ iff T is consistent with $\exists x\, \varphi(x)$.

EXERCISE 9. A class $\mathcal{K}$ of L-structures is a PC_Δ*-class*, if there is an extension L' of L and an L'-theory T' such that $\mathcal{K}$ consists of all reducts to L of models of T'.

Show that a PC_Δ-class of L-structures is L-elementary if and only if it is closed under L-elementary substructures.

EXERCISE 10. (Challenging!) An *existential sentence* is a sentence which consists of a string of existential quantifiers followed by a quantifier-free formula.

Show that a theory T can be axiomatised using existential sentences if and only if its models are closed under extensions.

CHAPTER 4

Directed systems and Craig interpolation

In this chapter we will introduce an important method for creating new models from old ones: colimits of directed systems. This we will then use to prove a fundamental property of first-order logic: the Craig interpolation theorem.

1. Directed systems

DEFINITION 4.1. A partially ordered set $(K, \leq)$ is called *directed*, if K is non-empty and for any two elements $x, y \in K$ there is an element $z \in K$ such that $x \leq z$ and $y \leq z$.

Note that non-empty linear orders (*aka* chains) are always directed.

DEFINITION 4.2. A *directed system of L-structures* consists of a family $(M_k)_{k \in K}$ of L-structures indexed by a directed partial order K, together with homomorphisms $f_{kl} \colon M_k \to M_l$ for $k \leq l$, satisfying:

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

If K is a chain, we call $(M_k)_{k \in K}$ a *chain of L-structures*

If we have a directed system, then we can construct its *colimit*, another L-structure M with homomorphisms $f_k \colon M_k \to M$. To construct the underlying set of the model M, we first take the disjoint union of all the universes:

$$\sum_{k \in K} M_k = \{(k, a) \colon k \in K, a \in M_k\},$$

and then we define an equivalence relation on it:

$$(k, a) \sim (l, b) \colon\Leftrightarrow (\exists m \geq k, l)\, f_{km}(a) = f_{lm}(b).$$

The underlying set of M will be the set of equivalence classes, where denote the equivalence class of (k, a) by $[k, a]$.

M has an L-structure: if c is some constant symbol, then we put

$$c^M = [k_0, c^{M_{k_0}}],$$

where k_0 is some arbitrary element from K. If R is a relation symbol in L, we put

$$R^M([k_1, a_1], \ldots, [k_n, a_n])$$

if there is a $k \geq k_1, \ldots, k_n$ such that

$$(f_{k_1 k}(a_1), \ldots, f_{k_n k}(a_n)) \in R^{M_k}.$$

And if g is a function symbol in L, we put

$$g^M([k_1, a_1], \dots, [k_n, a_n]) \quad = \quad [k, g^{M_k}(f_{k_1 k}(a_1), \dots, f_{k_n k}(a_n))],$$

where k is an element $\geq k_1, \dots, k_n$. (Check that this makes sense!) In addition, the homomorphisms $f_k \colon M_k \to M$ are obtained by sending a to $[k, a]$.

The following theorem collects the most important facts about colimits of directed systems. Especially useful is part 5, often called the *elementary system lemma*.

THEOREM 4.3. (1) *All f_k are homomorphisms.*

(2) *If $k \leq l$, then $f_l f_{kl} = f_k$.*
(3) *If N is another L-structure for which there are homomorphisms $g_k \colon M_k \to N$ such that $g_l f_{kl} = g_k$ whenever $k \leq l$, then there is a unique homomorphisms $g \colon M \to N$ such that $g f_k = g_k$ for all $k \in K$ (this is the universal property of the colimit).*
(4) *If all maps f_{kl} are embeddings, then so are all f_k.*
(5) *If all maps f_{kl} are elementary embeddings, then so are all f_k.*

PROOF. Exercise! □

The following fact about colimits of directed systems is also very useful:

LEMMA 4.4. *Let $(K, \leq)$ be a directed poset and $(M_k)_{k \in K}$ be a directed system. If J is a cofinal subset of K (meaning that for each $k \in K$ there is a $j \in J$ such that $k \leq j$), then $(M_j)_{j \in J}$ is a directed system as well and the colimits of the directed systems $(M_k)_{k \in K}$ and $(M_j)_{j \in J}$ are isomorphic.*

2. Robinson's Consistency Theorem

The aim of this section is to prove the statement:

> (Robinson's Consistency Theorem) Let L_1 and L_2 be two languages and $L = L_1 \cap L_2$. Suppose T_1 is an L_1-theory, T_2 an L_2-theory and both extend a complete L-theory T. If both T_1 and T_2 are consistent, then so is $T_1 \cup T_2$.

We first need two lemmas.

LEMMA 4.5. *Let $L \subseteq L'$ be languages and suppose A is an L-structure and B is an L'-structure. Suppose moreover $A \equiv B \restriction L$. Then there is an L'-structure C and a diagram of elementary embeddings (f in L and f' in L')*

$$\begin{tikzcd} A \arrow[dr, "f"'] & & B \arrow[dl, "f'"] \\ & C. & \end{tikzcd}$$

PROOF. Consider $T = \mathrm{Diag}^L_{\mathrm{el}}(A) \cup \mathrm{Diag}^{L'}_{\mathrm{el}}(B)$ (making sure we use different constants for the elements from A and B!). We need to show T has a model; so suppose T is inconsistent. Then, by compactness, a finite subset of T has no model; taking conjunctions, we have sentences $\varphi(\overline{a}) \in \mathrm{Diag}^L_{\mathrm{el}}(A)$ and $\psi(\overline{b}) \in \mathrm{Diag}^{L'}_{\mathrm{el}}(B)$ that are contradictory. But as the a_j do not occur in L'_B, we must have that $B \models \neg \exists \overline{x}\, \varphi(\overline{x})$. This contradicts $A \equiv B \restriction L$. □

LEMMA 4.6. *Let $L \subseteq L'$ be languages, suppose A and B are L-structures and C is an L'-structure. Any pair of L-elementary embeddings $f\colon A \to B$ and $g\colon A \to C$ fit into a commuting square*

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

where D is an L'-structure, h is an L-elementary embedding and k is an L'-elementary embedding.

PROOF. Without loss of generality we may assume that L contains constants for all elements of A. Then simply apply Lemma 4.5. □

THEOREM 4.7. (Robinson's Consistency Theorem) *Let L_1 and L_2 be two languages and $L = L_1 \cap L_2$. Suppose T_1 is an L_1-theory, T_2 an L_2-theory and both extend a complete L-theory T. If both T_1 and T_2 are consistent, then so is $T_1 \cup T_2$.*

PROOF. Let A_0 be a model of T_1 and B_0 be a model of T_2. Since T is complete, their reducts to L are elementary equivalent, so, by Lemma 4.5, there is a diagram

$$\begin{tikzcd} A_0 \arrow[dr, "f_0"] & \\ B_0 \arrow[r, "h_0"'] & B_1 \end{tikzcd}$$

with h_0 an L_2-elementary embedding and f_0 an L-elementary embedding. Now by applying Lemma 4.6 to f_0 and the identity on A_0, we obtain

$$\begin{tikzcd} A_0 \arrow[r, "k_0"] \arrow[dr, "f_0"'] & A_1 \\ B_0 \arrow[r, "h_0"'] & B_1 \arrow[u, "g_0"'] \end{tikzcd}$$

where g_0 is L-elementary and k_0 is L_1-elementary. Continuing in this way we obtain a diagram

$$\begin{tikzcd} A_0 \arrow[r, "k_0"] \arrow[dr, "f_0"'] & A_1 \arrow[r, "k_1"] \arrow[dr, "f_1"] & A_2 \arrow[r] & \dots \\ B_0 \arrow[r, "h_0"'] & B_1 \arrow[u, "g_0"'] \arrow[r, "h_1"'] & B_2 \arrow[u, "g_1"'] \arrow[r] & \dots \end{tikzcd}$$

where the k_i are L_1-elementary, the f_i and g_i are L-elementary and the h_i are L_2-elementary. Let C be the L-structure which is the colimit of the entire diagram. By Lemma 4.4, C is also the colimit of the A_i and of the B_i. Therefore C can also be equipped with an $L_1 \cup L_2$-structure, with A_0 and B_0 both embedding elementarily into C by the elementary systems lemma; hence C is a model of both T_1 and T_2, as desired. □

3. Craig Interpolation

THEOREM 4.8. *Let φ and ψ be sentences in some language such that $\varphi \models \psi$. Then there is a sentence θ, a "(Craig) interpolant", such that*

(1) *$\varphi \models \theta$ and $\theta \models \psi$;*
(2) *every predicate, function or constant symbol that occurs in θ occurs also in both φ and ψ.*

PROOF. Let L be the common language of φ and ψ. We will show that $T_0 \models \psi$ where $T_0 = \{\sigma \colon \sigma$ is an L-sentence and $\varphi \models \sigma\}$. Let us first check that this suffices for proving the theorem: for then there are $\theta_1, \ldots, \theta_n \in T_0$ such that $\theta_1, \ldots, \theta_n \models \psi$ by compactness. So $\theta \colon= \theta_1 \wedge \ldots \wedge \theta_n$ is an interpolant.

So we need to prove the following claim: If $\varphi \models \psi$, then $T_0 \models \psi$ where $T_0 = \{\sigma \in L \colon \varphi \models \sigma\}$ and L is the common language of φ and ψ. *Proof of claim:* Suppose not. Then $T_0 \cup \{\neg\psi\}$ has a model A. Write $T = \mathrm{Th}_L(A)$. Observe that we now have $T_0 \subseteq T$ and:

(1) T is a complete L-theory.
(2) $T \cup \{\neg\psi\}$ is consistent (because A is a model).
(3) $T \cup \{\varphi\}$ is consistent. (*Proof:* Suppose not. Then, by the compactness theorem, there would a sentence $\sigma \in T$ such that $\varphi \models \neg\sigma$. But then $\neg\sigma \in T_0 \subseteq T$. Contradiction!)

This means we can apply Robinson's Consistency Theorem to deduce that $T \cup \{\neg\psi, \varphi\}$ is consistent. But that contradicts $\varphi \models \psi$. □

4. Exercises

EXERCISE 11. The aim of this exercise is to prove the Chang-Łoś-Suszko Theorem. To state it we need a few definitions.

A $\forall\exists$-*sentence* is a sentence which consists first of a sequence of universal quantifiers, then a sequence of existential quantifiers and then a quantifier-free formula. A theory T can be axiomatised by $\forall\exists$-sentences if there is a set T' of $\forall\exists$-sentences such that T and T' have the same models.

In addition, we will say that a theory T is *preserved by directed unions* if for any directed system consisting of models of T and embeddings between them, also the colimit is a model T. And T is *preserved by unions of chains* if for any chain of models of T and embeddings between them, also the colimit is a model of T.

Show that the following statements are equivalent:

(1) T is preserved by directed unions.
(2) T is preserved by unions of chains.
(3) T can be axiomatised by $\forall\exists$-sentences.

Hint: To show (2) $\Rightarrow$ (3), suppose T is preserved by unions of chains and let

$$T_{\forall\exists} = \{\varphi \colon \varphi \text{ is a } \forall\exists\text{-sentence and } T \models \varphi\}.$$

Then prove that starting from any model B of $T_{\forall\exists}$ one can construct a chain of embeddings

$$B = B_0 \to A_0 \to B_1 \to A_1 \to B_2 \to A_2 \ldots$$

such that:

(1) Each A_n is a model of T.
(2) The composed embeddings $B_n \to B_{n+1}$ are elementary.
(3) Every universal sentence in the language L_{B_n} true in B_n is also true in A_n (when regarding A_n is an L_{B_n}-structure via the embedding $B_n \to A_n$).

Exercise 12. Use Robinson's Consistency Theorem to prove the following Amalgamation Theorem: Let L_1, L_2 be languages and $L = L_1 \cap L_2$, and suppose A, B and C are structures in the languages L, L_1 and L_2, respectively. Any pair of L-elementary embeddings $f\colon A \to B$ and $g\colon A \to C$ fit into a commuting square

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

where D is an $L_1 \cup L_2$-structure, h is an L_1-elementary embedding and k is an L_2-elementary embedding.

Exercise 13. The aim of this exercise is to prove Beth's Definability Theorem.

Let L be a language a P be a predicate symbol not in L, and let T be an $L \cup \{P\}$-theory. T *defines* P *implicitly* if any L-structure M has at most one expansion to an $L \cup \{P\}$-structure which models T. There is another way of saying this: let T' be the theory T with all occurrences of P replaced by P', another predicate symbol not in L. Then T *defines* P *implicitly* iff

$$T \cup T' \models \forall x_1, \dots x_n \left(P(x_1, \dots, x_n) \leftrightarrow P'(x_1, \dots, x_n) \right).$$

T *defines* P *explicitly*, if there is an L-formula $\varphi(x_1, \dots, x_n)$ such that

$$T \models \forall x_1, \dots, x_n \left(P(x_1, \dots, x_n) \leftrightarrow \varphi(x_1, \dots, x_n) \right).$$

Show that T defines P implicitly if and only if T defines P explicitly.

CHAPTER 5

Back and forth

1. Categoricity and Vaught's Test

Certain theories return again and again in model theory, because from a model-theoretic perspective they have many desirable properties. In this chapter we will discuss two of them.

One property both theories in this chapter share is that they are *complete*. (Recall that an L-theory T is complete if it is consistent and for any L-sentence φ we have either $T \models \varphi$ or $T \models \neg\varphi$.) Not many theories occurring in mathematics have this property, so if one can find a natural example then this is something special.

But how could one show that a theory is complete? For this one often applies Vaught's Test.

DEFINITION 5.1. Let κ be an infinite cardinal and let T be a theory with models of size κ. We say that T is *κ-categorical* if any two models of T of cardinality κ are isomorphic.

THEOREM 5.2. (Vaught's Test) *Let T be a consistent L-theory with no finite models that is κ-categorical for some infinite cardinal $\kappa \geq |L|$. Then T is complete.*

PROOF. Suppose T is not complete; then there is a sentence φ such that $T \not\models \varphi$ and $T \not\models \neg\varphi$. This means that there are models M and N of T such that $M \models \varphi$ and $N \models \neg\varphi$. Since $\kappa \geq |L|$ we can use the upward and downwards Skolem-Löwenheim theorems to arrange that both M and N have cardinality κ. But this contradicts the κ-categoricity of T. $\square$

Vaught's Test reduces the problem of showing completeness to the problem of showing categoricity. For the latter purpose we often use a technique called *back and forth*: the idea is to construct an isomorphism between two models of the same size by some inductive procedure. This is best illustrated through the examples.

2. Dense linear orders

The theory DLO of dense linear orders without endpoints is the theory in the language $<$ saying that:

(1) $<$ defines an ordering: if $x < y$ then not $x = y$ and not $y < x$, and if $x < y$ and $y < z$ then $x < z$.
(2) The order $<$ is linear: $x < y$ or $x = y$ or $y < x$.
(3) It is dense: this says that $x < y$ implies that there is a z with $x < z < y$.
(4) It has no endpoints: for every x there are y and z such that $y < x < z$.

Examples are $(\mathbb{Q}, <)$ and $(\mathbb{R}, <)$.

Definition 5.3. Let M and N be two L-structure. A function $f\colon A \subseteq M \to N$ with $A = \{a_1, \ldots, a_n\}$ a finite subset of M is called a *local isomorphism* if

$$M \models \varphi(a_1, \ldots, a_n) \Leftrightarrow N \models \varphi(f(a_1), \ldots, f(a_n))$$

holds for every atomic (or, equivalently, quantifier-free) L-formula $\varphi(x_1, \ldots, x_n)$.

By considering the formula $x_i = x_j$ we see that local isomorphisms are injective.

Proposition 5.4. *Let $f\colon A \subseteq M \to N$ be a local isomorphism between two models M and N of DLO. For any $m \in M$ there is a local isomorphism $g\colon A \cup \{m\} \subseteq M \to N$ with $g \restriction A = f$.*

Proof. Let M and N be two dense linear orders without endpoint and $f\colon A \subseteq M \to N$ be a local isomorphism. For DLO the latter just means that f preserves and reflects the order relation $<$.

Our task is to show that for any $m \in M$ we can extend the local isomorphism f to one whose domain includes m. For this we put $A_0 := \{a \in A : a < m\}$ and $A_1 := \{a \in A : a > m\}$ and make some case distinctions:

(i) $m \in A$. In this case we can simply put $g := f$.
(ii) $A_0 = A$. In this case m is larger than any element in A and we use that N has no endpoints to find an element $n \in N$ which is larger than any element in $f(A)$. Then we put $g(m) := n$ (and on all elements in A the function g is defined in the same way as f).
(iii) $A_1 = A$. In this case m is smaller than any element in A and we use that N has no endpoints to find an element $n \in N$ which is smaller than any element in $f(A)$. Then we put $g(m) := n$.
(iv) Neither A_0 nor A_1 is the whole of A or empty. Let a_0 be the largest element of A_0 and a_1 be the smallest element of A_1. Using that N is dense we find an element $n \in N$ such that $f(a_0) < n < f(a_1)$. Then we put $g(m) := n$.

$\square$

Theorem 5.5. *The theory DLO is ω-categorical.*

Proof. Let M and N be two countable dense linear orders without endpoints. Fix enumerations $M = \{m_0, m_1, \ldots\}$ and $N = \{n_0, n_1, \ldots\}$. We will construct an increasing sequence of local isomorphisms f_k from some subset of M to N such that m_i belongs to the domain of f_{2i+2} and n_i belongs to the codomain of f_{2i+1}. Then $f = \bigcup_i f_i$ will be the desired isomorphism between M and N. We start with $f_0 = \emptyset$.

So suppose we have constructed f_k and we want to construct f_{k+1}. If $k + 1 = 2i + 2$, then we apply the previous proposition on m_i and f_k to construct a local isomorphism f_{k+1} which extends f_k and whose domain includes m_i (this is the *forth* in back and forth).

If $k + 1 = 2i + 1$, then we consider f_k^{-1}, which is a local isomorphism from some finite subset of N to M. So by the previous proposition there is a local isomorphism g whose domain includes both n_i and the image of f_k. Then we put $f_{k+1} = g^{-1}$, which is a local isomorphism as desired. $\square$

Corollary 5.6. *The theory DLO is complete.*

3. Algebraically closed fields

Recall that a field K is called algebraically closed if every non-constant polynomial has a root in K. Throughout this section we will fix some characteristic, which could be either 0 or some prime p. We will write ACF_0 for the theory of fields of characteristic 0, while ACF_p is the theory of algebraically closed fields of characteristic p.

3.1. Recap on fields. Consider an inclusion $K \subseteq L$ of fields. Recall that L can be considered as a K-vector space and that we write $[K\colon L]$ for its dimension.

PROPOSITION 5.7. *If we have two field extensions $K \subseteq L \subseteq M$, then $[M\colon K] = [M\colon L][L\colon K]$.*

If $K \subseteq L$ and $\xi \in L$, then there are two possibilities:

(1) ξ is algebraic over K. This means that there is a non-zero polynomial $p(x)$ with coefficients from K such that $p(\xi) = 0$. In this case we can consider the monic polynomial $m(x) \in K[x]$ with $m(\xi) = 0$ which has least possible degree: this is called the *minimal polynomial* of ξ. This polynomial has to be irreducible and $K(\xi)$, the smallest subfield of L which contains both K and ξ, is isomorphic to $K[x]/(m(x))$. In this case $[K(\xi)\colon K]$ is finite.
(2) ξ is transcendental over K. In this case $K(\xi)$ is isomorphic to the quotient field $K(x)$ and $[K(\xi)\colon K]$ is infinite.

An extension $K \subseteq L$ is called *algebraic* if all elements in L are algebraic over K. From Proposition 5.7 it follows that:

(1) $K(\xi)$ is algebraic over K precisely when ξ is algebraic over K.
(2) If $K \subseteq L$ and $L \subseteq M$ are two field extensions and they are both algebraic, then so is $K \subseteq M$.

3.2. Algebraic closure.

DEFINITION 5.8. If $K \subseteq L$ is a field extension, then L is an *algebraic closure* of K, if L is algebraic over K, but no proper extension of L is algebraic over K.

THEOREM 5.9. *Algebraic closures are algebraically closed.*

PROOF. Let L be the algebraic closure of K and $p(x)$ be a non-constant polynomial with coefficients from L without any roots in L. Without loss of generality we may assume that $p(x)$ is irreducible (otherwise replace $p(x)$ with one of its irreducible factors); but then $L[x]/(p(x))$ is a proper algebraic extension of L and K, which is a contradiction. □

THEOREM 5.10. *Every field K has an algebraic closure.*

PROOF. Let X the collection of algebraic field extensions of K and order by embedding of fields. We restrict attention to those fields whose cardinality is bounded by the maximum of $|K|$ and $\aleph_0$, and therefore X is a set (essentially). Clearly, every chain of embeddings has an upper bound in X, so by Zorn's Lemma X has a maximal element L. This field is an algebraic closure of X: for if $L \subset M$ is a proper extension of fields and $\xi \in M - L$, then ξ cannot be algebraic over K. For otherwise $L \subset L(\xi) \in X$, contradicting maximality of L. □

THEOREM 5.11. *Algebraic closures are unique up to (non-unique) isomorphism.*

Proof. By a back and forth argument. Let L and M be algebraic closures of K. Since L and M must have the same infinite cardinality $\kappa = \max(|K|, \aleph_0)$, we can fix enumerations $\{l_i : i \in \kappa\}$ and $\{m_i : i \in \kappa\}$ of L and M, respectively. By induction on $i \in \kappa$ we will construct an increasing sequence of isomorphisms $f_i \colon L_i \to M_i$ between subfields of L and M such that $\bigcup L_i = L$ and $\bigcup M_i = M$. We start by declaring f_0 to be isomorphism between the isomorphic copies of K inside L and M; and at limit stages we simply take the union.

At successor stage $i+1$, we can write $i = \lambda + k$ with λ a limit ordinal and k a finite ordinal. If $k = 2j$, then look at the minimal polynomial $m(x) = a_n x^n + a_{n-1}x^{n-1} + \ldots + a_0$ of $l_{\lambda+j}$ over L_i: such a thing exists because L is algebraic over K and hence over L_i. Because M is algebraically closed, there exists a root $m \in M$ of the polynomial $n(x) = f_i(a_n)x^n + f_i(a_{n-1})x^{n-1} + \ldots + f(a_0)$; since f_i is an isomorphism, the polynomial $n(x)$ is irreducible over M_i and since M is algebraically closed, $n(x)$ must be the minimal polynomial of $m \in M$ over M_i. So we can extend the isomorphism by sending $l_{\lambda+j}$ to m:

$$f_{i+1} \colon L_i(l_{\lambda+j}) \cong L_i[x]/(m(x)) \cong M_i[x]/(n(x)) \cong M_i(m).$$

If $k = 2j+1$, then we can use a similar argument to show that the isomorphism f_i can be extended to one whose codomain includes $m_{\lambda+j}$. □

3.3. Categoricity. A similar argument shows:

Theorem 5.12. *The theories ACF_0 and ACF_p are λ-categorical for any uncountable λ.*

Proof. Let L and M be two algebraically closed fields of the same uncountable cardinality λ and fix enumerations $\{l_i : i \in \lambda\}$ and $\{m_i : i \in \lambda\}$ of L and M, respectively. By induction on $i \in \lambda$ we will construct an increasing sequence of isomorphisms $f_i \colon L_i \to M_i$ between subfields of L and M *of cardinality strictly less than* λ such that $\bigcup L_i = L$ and $\bigcup M_i = M$. We start by declaring f_0 to be isomorphism between the isomorphic copies of $\mathbb{Q}$ (if the characteristic is 0) or $\mathbb{F}_p$ (if the characteristic is p) inside L and M; and at limit stages we simply take the union.

At successor stage $i+1$, we can again write $i = \lambda + k$ with λ a limit ordinal and k a finite ordinal. If $k = 2j$, then there are two possibilities for $l_{\lambda+j}$ *vis-à-vis* L_i: it can either be algebraic or transcendental. If it is algebraic, we proceed as in the proof of the previous theorem. We look at the minimal polynomial $m(x) = a_n x^n + a_{n-1}x^{n-1} + \ldots + a_0$ of $l_{\lambda+j}$ over L_i and use that M is algebraically closed to find an element $m \in M$ with minimal polynomial $n(x) = f_i(a_n)x^n + f_i(a_{n-1})x^{n-1} + \ldots + f(a_0)$ over M_i. And we extend the isomorphism by sending $l_{\lambda+j}$ to m:

$$f_{i+1} \colon L_i(l_{\lambda+j}) \cong L_i[x]/(m(x)) \cong M_i[x]/(n(x)) \cong M_i(m).$$

If, one the other hand, $l_{\lambda+j}$ is transcendental over L_i, we use the fact that $|M_i| < |M|$ to deduce that M also contains an element $m \in M$ which transcendental over M_i. And the isomorphism can be extended by sending $l_{\lambda+j}$ to m:

$$f_{i+1} \colon L_i(l_{\lambda+j}) \cong L_i(x) \cong M_i(x) \cong M_i(m).$$

If $k = 2j+1$, then we can use a similar argument to show that the isomorphism f_i can be extended to one whose codomain includes $m_{\lambda+j}$. □

Corollary 5.13. *The theories ACF_0 and ACF_p are complete.*

4. Exercises

EXERCISE 14. Show that DLO is not λ-categorical for any $\lambda > \omega$.

EXERCISE 15. Show that the embedding $(\mathbb{Q}, <) \subseteq (\mathbb{R}, <)$ is elementary.

EXERCISE 16. By a *graph* we will mean a pair (V, E) where V is a non-empty set and E is a binary relation on V which is both symmetric and irreflexive. We will refer to the elements of V as the *vertices* and the elements of E as the *edges*. If xEy holds for two $x, y \in V$, we say that x and y are *adjacent*.

A graph (V, E) will be called *random* if for any two finite sets of vertices X and Y which are disjoint there is a vertex $v \notin X \cup Y$ which adjacent to all of the vertices in X and to none of the vertices in Y. We will write RG for the theory of random graphs.

Show that the theory RG is ω-categorical, and hence complete.

EXERCISE 17. Show that the theory ACF_0 is not ω-categorical.

EXERCISE 18. Let φ be a sentence in the language of rings. Show that the following are equivalent:

(i) φ is true in the complex numbers.
(ii) φ is true in every algebraically closed field of characteristic 0.
(iii) φ is true in some algebraically closed field of characteristic 0.
(iv) There are arbitrarily large primes p such that φ is true in some algebraically closed field of characteristic p.
(v) There is an m such that for all $p > m$, the sentence φ is true in all algebraically closed fields of characteristic p.

CHAPTER 6

Ehrenfeucht-Fraïssé games

This chapter will be devoted to a game-theoretic characterisation of the notion of elementary equivalence. This interpretation in terms of games are not just fun: it can often be applied in situations where other methods fail.

Throughout this chapter we will, for simplicity, be working in a finite language without function symbols.

1. Definition of the game

Definition 6.1. Given two models M and N and a natural number $n \in \mathbb{N}$ we define a game as follows. It is a two-player game in which two players, player I (who is male) and player II (who is female), move in turn. Player I starts and the game ends after n rounds, so after both players have played n moves. A move by a player consists of picking an element from one of the two structures. Player I has complete freedom and can pick an element from whichever structures he likes, but player II always has to reply by picking an element from the other structure (that is, player II is not allowed to respond by picking an element from the same structure as the one player I just played in). So if in round i player I chooses an element $a_i \in M$, player II replies by picking an element $b_i \in N$, and if in round i player I chooses an element $b_i \in N$, then player II replies by picking an element $a_i \in M$. After n rounds the two players have constructing two sequences $(a_1, \ldots, a_n)$ and $(b_1, \ldots, b_n)$ of elements from M and N, respectively. Player II wins if $\{(a_i, b_i) \colon 1 \leq i \leq n\}$ is a well-defined injective function $f \colon \{a_1, \ldots, a_n\} \to N$ and, moreover, this function is a local isomorphism; otherwise player I wins. We denote this game by $\mathcal{G}_n(M, N)$ and we call it an *Ehrenfeucht-Fraïssé game.*

Let us first remark that:

Proposition 6.2. *One of the players has a winning strategy in $\mathcal{G}_n(M, N)$.*

This is a consequence of a general result in game theory:

Theorem 6.3. (Zermelo) *In a two-player game of perfect information in which there are no infinite plays and no ties, one of the two players has a winning strategy.*

Proof. (Sketch) Let us say that a position in the game is *losing* for a player if in that position the other player has a winning strategy. The idea is that if none of the two players has a winning strategy, both can play in such a way that they avoid any losing positions. That is, if player II does not have a winning strategy, player I can play a move after which the position is not lost for him. But if player I also does not have a winning strategy, the position after he has played this move is also not lost for player II. That means that player II can reply by playing a move after which the position is not lost for her. But after player II has played such a move, the position is also not lost for player I: because otherwise the position just before player II played her move must have been lost for player I. This means that player I can reply by playing a move after which the position is not lost for him; actually, it will not be lost for either of the two players. By proceeding in this vein both players end up playing a game of infinite length, which contradicts the assumption that any possible way of playing the game ends after a finite number of moves in a win for one of the two players. □

The reason we are interested in Ehrenfeucht-Fraïssé games is that they allow us to characterise elementary equivalence.

THEOREM 6.4. *Let L be a finite language without function symbols and let M and N be L-structures. Then $M \equiv N$ if and only if the player II has a winning strategy in $\mathcal{G}_n(M, N)$ for all n.*

A more refined statement is true. To formulate it, we need the following definition.

DEFINITION 6.5. The *quantifier-depth* $\mathrm{dp}(\varphi)$ of a formula φ is defined inductively as follows:

- $\mathrm{dp}(\varphi) = 0$ if φ is atomic.
- $\mathrm{dp}(\varphi \Box \psi) = \max\{\mathrm{dp}(\varphi), \mathrm{dp}(\psi)\}$ for $\Box \in \{\wedge, \vee, \to\}$.
- $\mathrm{dp}(\neg\varphi) = \mathrm{dp}(\varphi)$,
- $\mathrm{dp}(\exists x\, \varphi) = \mathrm{dp}(\forall x\, \varphi) = \mathrm{dp}(\varphi) + 1$.

We will write $M \equiv_n N$ if

$$M \models \varphi \Leftrightarrow N \models \varphi$$

for any sentence φ with quantifier-depth at most n.

THEOREM 6.6. *Let L be a finite language without function symbols and let M and N be L-structures. Then $M \equiv_n N$ if and only if the player II has a winning strategy in $\mathcal{G}_n(M, N)$.*

The proof of this theorem is a bit finicky: we will give it in Section 3. But before we give this proof, let us first discuss an application.

2. An application

We give one application of this theorem. Let $L = \{<\}$ and let T be the L-theory asserting that $<$ is a discrete linear order without greatest or smallest element. Discreteness means:

$$\forall x\, \exists y_0\, \exists y_1\, \big(\, y_0 < x < y_1 \wedge \forall z\, (\, z < x \to z \leq y_0 \wedge x < z \to y_1 \leq z\,)\,\big),$$

where $x \leq y$ abbreviates $x < y \vee x = y$. In other words, discreteness means that each element has an immediate successor and predecessor. For example, $(\mathbb{Z}, <)$ is a model of T.

We claim that T is a complete theory, or, equivalently that any model N of T is elementarily equivalent to $(\mathbb{Z}, <)$.

PROPOSITION 6.7. *The theory T of discrete linear orders with no top or bottom element is a complete theory. In particular, $(\mathbb{Z}, <) \models \varphi$ if and only if $T \models \varphi$ for all L-sentences φ.*

PROOF. We are going to use games. But before we do this, we should first try to understand how a model N of T looks like.

For elements $a, b \in N$ let us write aEb if b is the nth successor or predecessor of a for some natural number n. Then E is an equivalence relation and each E-class is a linear order that looks like $(\mathbb{Z}, <)$. In addition, the collection of E-classes is linearly ordered as well (by saying that $[a] < [b]$ if $\neg E(a,b) \wedge a < b)$. This means that every model of T is of the form $(L \times \mathbb{Z}, <)$, where L is a linear order and $<$ is the lexicographic order on $L \times \mathbb{Z}$ (that is, $(i,a) < (j,b)$ if $i < j$, or both $i = j$ and $a < b$). Conversely, every linear order of this form is a model of T.

So let M be $(\mathbb{Z}, <)$, and let N be $L \times \mathbb{Z}$ with the lexicographic order, where L is any linearly ordered set. We wish to show that $M \equiv N$ and we do this by supplying for each natural number n a winning strategy for player II in the game $\mathcal{G}_n(M,N)$.

If $a, b \in \mathbb{Z}$, we define the distance between a and b to be $\operatorname{dist}(a,b) = |b - a|$, and if $x = (i,a), y = (j,b) \in L \times \mathbb{Z}$, we define the distance to be $\operatorname{dist}(x,y) = |b-a|$ if $i = j$ and $\operatorname{dist}(a,b) = \infty$ if $i \neq j$. The problem for player II is that player I can play elements that are infinitely far apart in N and force player II to play elements that are finitely far apart in M. The crux is that the number of rounds that the game will last has been fixed in advance (and player II knows this number) and if player II can play elements that are sufficiently far apart to avoid conflicts, then she can win the game. Indeed, we claim that if the game lasts i more rounds then Player II only has to ensure that distances $< 2^i$ are preserved. More precisely, player II can win by ensuring the following condition:

> (†) After m rounds of $\mathcal{G}_n(M,N)$ we have $a_i < a_j$ iff $b_i < b_j$ and $a_i = a_j$ iff $b_i = b_j$ and $\min(\operatorname{dist}(a_i,a_j), 2^{n-m}) = \min(\operatorname{dist}(b_i,b_j), 2^{n-m})$.

Clearly, if player II can actually achieve this, she will win because after n rounds there will be a local isomorphism.

So it remains to argue that player II can always choose a move to preserve (†). In round 1, player II chooses an arbitrary element and (†) holds. Suppose that we have played m rounds and (†) holds, and the moves played so far have been $a_1, \ldots, a_m$ in M and $b_1, \ldots, b_m$ in N. Suppose that player I plays $b \in L \times \mathbb{Z}$. There are several cases to consider.

(1) $b < b_i$ for all i. Suppose b_j is the smallest element of the b_i. Then choose $a = a_j - \min(\operatorname{dist}(b,b_j), 2^{n-m-1})$.
(2) $b_i < b < b_j$ for some i and j. Choose i and j such that $b_i < b < b_j$ and there are no b_k such that $b_i < b_k < b_j$.
   (a) If $\operatorname{dist}(b,b_i) < 2^{n-m-1}$, then put $a = a_i + \operatorname{dist}(b,b_i)$.
   (b) If $\operatorname{dist}(b,b_j) < 2^{n-m-1}$, then put $a = a_j - \operatorname{dist}(b,b_j)$.
   (c) If $\operatorname{dist}(b,b_i) \geq 2^{n-m-1}$ and $\operatorname{dist}(b,b_j) \geq 2^{n-m-1}$, then $\operatorname{dist}(b_i,b_j) \geq 2^{n-m}$ and $\operatorname{dist}(a_i,a_j) \geq 2^{n-m}$. Put $a = a_i + 2^{n-m-1}$.
(3) If $b > b_i$ for all i. Suppose b_j is the biggest element of the b_i. Then choose $a = a_j + \min(\operatorname{dist}(b,b_j), 2^{m-n-1})$.

This explains the strategy if player I plays $b \in L \times \mathbb{Z}$. The case where player I plays $a \in \mathbb{Z}$ is simpler and left to the reader. □

3. A proof

The aim of this section give a proof (sketch) for Theorem 6.6. We start off with some syntactic considerations.

A formula is a *boolean combination* of formulas in S if it can be obtained from S by applying conjunction, disjunction, implication and negation (that is, all the possible propositional operations). In addition, let us say that a collection of formulas S is *finite up to logical equivalence* if there is a finite set of formulas $S_0 \subseteq S$ such that each element in S is logically equivalent to some element in S_0.

LEMMA 6.8. *Let A be a collection of formulas and assume B consists of all boolean combinations of A.*

(i) *If M and N make the same formulas in A true, then they also make the same formulas in B true.*
(ii) *If A is finite up to logical equivalence, then so is B.*

PROOF. (i) is proved by induction on the logical complexity of the formulas in B.

(ii): Suppose each element in A is logically equivalent to some element of $\{\varphi_0, \dots, \varphi_{n-1}\}$. Then each element $\psi \in B$ is equivalent to a formula of the form

$$\bigvee_{\sigma \in S} \Big(\bigwedge_{\{i \,:\, \sigma(i)=1\}} \varphi_i \wedge \bigwedge_{\{i \,:\, \sigma(i)=0\}} \neg\varphi_i \Big)$$

for some $S \subseteq \{0,1\}^n$. Details are left to the reader. □

DEFINITION 6.9. Let us say that for a set of variables $\{x_1, \dots, x_m\}$ and a natural number n a *formula* φ is special if:

(1) either $n = 0$ and φ is an atomic formula with free variables among $\{x_1, \dots, x_m\}$, or
(2) $n > 0$ and φ is of the form $\exists x_{m+1} \varphi$ where φ is a formula with quantifier-depth at most $n-1$ and free variables among $\{x_1, \dots, x_{m+1}\}$.

LEMMA 6.10. *Let $\mathcal{L}$ be a finite language without function symbols.*

(i) *Every formula with quantifier-depth at most n and free variables among $\{x_1, \dots, x_m\}$ is logically equivalent to a boolean combination of special formulas with respect to $\{x_1, \dots, x_m\}$ and n.*
(ii) *The collection of formulas with quantifier-depth at most n and free variables among $\{x_1, \dots, x_m\}$ is finite up to logical equivalence.*

PROOF. (i) Each quantifier-free formula is a boolean combination of atomic formulas, and each formula of quantifier-depth at most $n+1$ is a boolean combination of formulas of the form $\exists x\, \varphi$ and $\forall x\, \varphi$ with φ having quantifier-depth at most n. Up to logical equivalemce, we can rename variables so that x becomes x_{m+1} and eliminate $\forall x\, \varphi$ in favour of $\neg\exists x\, \neg\varphi$.

(ii) is proved by induction. For $n = 0$ observe that the number of atomic formulas with free variables among $\{x_1, \dots, x_m\}$ is finite. So then this follows from point (i) in this lemma and point (ii) in the previous lemma. For the induction step we argue in the same way, with special formulas instead of atomic formulas. □

THEOREM 6.11. *Let L be a finite language without function symbols and let M and N be L-structures. Then $M \equiv_n N$ if and only if the player II has a winning strategy in $\mathcal{G}_n(M, N)$.*

PROOF. (Sketch) $\Rightarrow$: Suppose $M \equiv_n N$. We will outline a winning strategy for player II. Suppose player I plays a move $a \in M$ (the case that player I plays an element $b \in N$ is similar). We claim that player II can choose an element $b \in N$ in such a way that $(M, a) \equiv_{n-1} (N, b)$. (This statement should be understood in the following way: if we would add a constant c to the language and interpret it as a in M and b in N, then M and N are $(n-1)$-elementary equivalent in the extended language.)

Indeed, suppose that $\{\varphi_i(x_0) : i \in I\}$ is a finite set of formulas of quantifier-depth at most $n-1$ and with free variable $\{x_0\}$ and suppose that each such formula is equivalent to an element in this set. Put

$$\psi(x_0) := \bigwedge\{\varphi_i(x_0) : M \models \varphi_i(a)\} \wedge \bigwedge\{\neg\varphi_i(x_0) : M \not\models \varphi_i(a)\}.$$

Then $M \models \exists x_0\, \psi(x_0)$, because $M \models \psi(a)$. Since $M \equiv_n N$ and $\exists x_0\, \psi(x_0)$ is of quantifier-depth n, there is an element $b \in N$ such that $N \models \psi(a)$. But then $(M, a) \equiv_{n-1} (N, b)$, as desired.

If player II continues in this way, the players will end up producing sequences $(a_1, \ldots, a_n)$ in M and $(b_1, \ldots, b_n)$ in N such that

$$(M, a_1, \ldots, a_n) \equiv_0 (N, b_1 \ldots, b_n).$$

But then the function $f(a_i) = b_i$ is a local isomorphism, so player II will win the game.

$\Leftarrow$: Suppose $M \not\equiv_n N$. Now we will outline a winning strategy for player I.

If $M \not\equiv_n N$ there must a sentence of the form $\exists x\, \psi(x)$ where ψ has quantifier-depth $n-1$ which is true in one structure, but not in the other. Suppose it is true in M (the case where it is true in N is similar). Then player I starts by picking an element $a \in M$ such that $M \models \psi(a)$. Player II has to respond by picking an element $b \in N$. But then $N \not\models \psi(b)$, so $(M, a) \not\equiv_{n-1} (N, b)$.

If player I continues in this way, the players will end up producing sequences $(a_1, \ldots, a_n)$ in M and $(b_1, \ldots, b_n)$ in N such that

$$(M, a_1, \ldots, a_n) \not\equiv_0 (N, b_1 \ldots, b_n).$$

But then $f(a_i) = b_i$, even when this defines a function, cannot be a local isomorphism. Therefore player I wins the game. □

4. Exercises

EXERCISE 19. Give a direct proof of Proposition 6.2, that is, without using Theorem 6.3. *Hint:* Simply use induction on n.

EXERCISE 20. Let L be the first-order language of linear ordering. Show that if $h < 2^k$ then there is a formula $\varphi(x, y)$ of L of quantifier depth $\leq k$ which expresses (in any linear ordering) "$x < y$ and there are at least h elements strictly between x and y".

EXERCISE 21. The circle of length $N \in \mathbb{N}$ is the structure $\mathcal{C}_N := (C_N, R)$, where $C_N = \{0, \ldots, N-1\}$ and $R = \{(i, j) \in C_N \times C_N : j = i + 1 \mod N\}$.

(a) Give a function $f : \mathbb{N} \to \mathbb{N}$ such that $\mathcal{C}_N \equiv_n \mathcal{C}_{N'}$ whenever $N, N' \geq f(n)$.
(b) Is there a first-order formula φ such that $\mathcal{C}_N \models \varphi$ if and only if N is even?

EXERCISE 22. Show that there is no formula of first-order logic which expresses "(a, b) is in the transitive closure of R", even on finite structures. (For infinite structures it is easy to show there is no such formula.)

CHAPTER 7

Types

1. Terminology

One of the most important notions in model theory is that of a *type*. Intuitively, a type is the complete list of formulas $\varphi(x_1, \dots, x_n)$ satisfied by some tuple $(a_1, \dots, a_n)$.

DEFINITION 7.1. Fix $n \in \mathbb{N}$ and let $x_1, \dots, x_n$ be a fixed sequence of distinct variables. If A is an L-structure and $a_1, \dots, a_n \in A$, then the *type of* $(a_1, \dots, a_n)$ *in* A is the set of L-formulas

$$\{\, \varphi(x_1, \dots, x_n) \colon A \models \varphi(a_1, \dots, a_n) \,\};$$

we denote this set by $\mathrm{tp}_A(a_1, \dots, a_n)$ or simply by $\mathrm{tp}(a_1, \dots, a_n)$ if A is understood. An *n-type in L* is a set of formulas of the form $\mathrm{tp}_A(a_1, \dots, a_n)$ for some L-structure A and some $a_1, \dots, a_n \in A$. (I will sometimes call types *complete types* to distinguish them from the partial types defined below.)

Some observations:

- If $i \colon A \to B$ is an elementary embedding and $a_1, \dots, a_n \in A$, then $(a_1, \dots, a_n)$ and $(f(a_1), \dots, f(a_n))$ have the same type.
- Two n-tuples $(a_1, \dots, a_n)$ from A and $(b_1, \dots, b_n)$ from B satisfy the same n-type precisely when $(A, a_1, \dots, a_n) \equiv (B, b_1, \dots, b_n)$. (This is supposed to mean: add new constants $c_1, \dots, c_n$ to the language and regard A and B as $(L \cup C)$-structures by interpreting c_i as a_i in A and as b_i in B.)

It will occasionally be useful to also consider "incomplete" (or even inconsistent) lists of formulas: this is a partial type.

DEFINITION 7.2. Fix $n \in \mathbb{N}$ and let $x_1, \dots, x_n$ be a fixed sequence of distinct variables. A *partial n-type in L* is a collection of formulas $\varphi(x_1, \dots, x_n)$ in L.

- If $p(x_1, \dots, x_n)$ is a partial n-type in L, we say $(a_1, \dots, a_n)$ *realizes p in A* if every formula in p is true of $a_1, \dots, a_n$ in A.
- If $p(x_1, \dots, x_n)$ is a partial n-type in L and A is an L-structure, we say that p is *realized or satisfied in A* if there is some n-tuple in A that realizes p in A. If no such n-tuple exists, then we say that A *omits p*.

What distinguishes the (complete) types among the partial types? Essentially, the types are the maximally consistent partial types. This follows from the fact that they can be realized in some model, and that they contain either $\varphi(x_1, \dots, x_n)$ or $\neg\varphi(x_1, \dots, x_n)$ for any L-formula φ whose free variables are among the fixed variables $x_1, \dots, x_n$. And, indeed, if a partial type has these two properties it must be a complete type: for if a partial n-type p is realized by $(a_1, \dots, a_n)$, we must have $p \subseteq \mathrm{tp}(a_1, \dots, a_n)$. If p is also complete, then $p \supseteq \mathrm{tp}(a_1, \dots, a_n)$ follows as well. (For if $\varphi \notin p$, then $\neg\varphi \in p$, so $\neg\varphi \in \mathrm{tp}(a_1, \dots, a_n)$, hence $\varphi \notin \mathrm{tp}(a_1, \dots, a_n)$.)

2. Types and theories

Definition 7.3. Let T be a theory in L and let $p = p(x_1, \ldots, x_n)$ be a partial n-type in L. If T has a model realizing p, then we say that *p is consistent with T* or that *p is a type of T*. The set of all complete n-types consistent with T is denoted by $S_n(T)$.

Observe:

Lemma 7.4. *Let T be a theory and p be a partial n-type consistent with T. Then p can be extended to a complete n-type q which is still consistent with T.*

Proof. If $p(\overline{x})$ is some partial n-type consistent with T then, by definition, there is some model M of T in which there is some n-tuple of elements $\overline{a}$ realizing $p(\overline{x})$. Then $q = \mathrm{tp}_M(\overline{a})$ is a complete type consistent with T and extending p. □

Suppose p is consistent with T and M is a model of T: does this mean that p will be realized in M? The answer is *no*: the types consistent with T are those types that are realized in *some* model of T. It may very well happen that M is a model of T and p is an n-type consistent with T, but p is not realized in M, even when the theory T is complete. So what can we say?

Definition 7.5. If $p(x_1, \ldots, x_n)$ is a partial n-type in L and A is an L-structure, we say that p is *finitely satisfiable in A* if any finite subset of p is realized in A.

Proposition 7.6. *Let M be a model of a complete theory T. Then a partial type p is consistent with T if and only if it is finitely satisfiable in M.*

Proof. First suppose that p is consistent with T. To show that p is finitely satisfiable in M, let $\varphi_1(x), \ldots, \varphi_n(x)$ be finitely many formulas in p. We must have

$$T \models \exists x\big(\varphi_1(x) \wedge \ldots \varphi_n(x)\big);$$

for if this is not true, then $T \models \neg\exists x\big(\varphi_1(x) \wedge \ldots \varphi_n(x)\big)$ by completeness of T. But then p cannot be satisfied in any model of T, contradicting the fact that p is consistent with T. So, if M is a model of T, we must have

$$M \models \exists x\big(\varphi_1(x) \wedge \ldots \varphi_n(x)\big);$$

since $\varphi_1(x), \ldots, \varphi_n(x)$ were arbitrary, the type p is finitely satisfiable in M.

Conversely, suppose that p is finitely satisfiable in M. Add a fresh constant c to the language and look at the theory

$$T' = T \cup \{\varphi(c) : \varphi \in p\}.$$

If p is finitely satisfiable in M, then M is a model for every finite subset of T'. So, by the compactness theorem, T' has a model N: this is a model of T in which p is realized, showing that p is consistent with T. □

The next lemma formulates some useful properties of finitely satisfiable partial types.

Lemma 7.7. *Let M be a model and p be a partial type.*

(1) *If $M \equiv N$ and p is finitely satisfiable in M, then p is also finitely satisfiable in N.*
(2) *p is finitely satisfiable in M if and only if p is realized in some elementary extension of M.*

(3) *If p is finitely satisfiable in M, then p can be extended to a complete type q which is still finitely satisfiable in M.*

PROOF. (1) If $M \equiv N$ then M and N are models of the same complete theory T. So if p is finitely satisfiable in M, then it is consistent with T and hence finitely satisfiable in N (using the previous proposition twice, once for M and once for N).

(2) Consider the theory $T = \mathrm{ElDiag}(M) \cup \{\varphi(c) : \varphi \in p\}$, where c is a fresh constant which does not occur in L. If p is finitely satisfiable in M, then M is a model of every finite subset of T, so, by the compactness theorem, T has a model N. This, by construction, is a model in which M embeds and in which p is realized.

Conversely, if p is realized in some elementary extension of M, then this extension is a model which is elementary equivalent to M and in which p is (finitely) satisfied, so p is finitely satisfiable in M by (1).

(3) By (2) p is realized in some elementary extension, by some element a say. Then the type of a in this elementary extension is a complete type extending p.

□

3. Type spaces

Crucially, the set $S_n(T)$ can be given the structure of a topological space. To see this, consider sets in $S_n(T)$ of the form

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

where $\varphi(x_1, \ldots, x_n)$ is some formula. The following lemma states some basic properties of sets of the form $[\varphi]$: they are not hard to prove (in fact, they are direct consequences of the completeness properties of types).

LEMMA 7.8.

$$\begin{aligned}
[\varphi] \subseteq [\psi] &\Leftrightarrow T \models \varphi \to \psi \\
[\varphi] = [\psi] &\Leftrightarrow T \models \varphi \leftrightarrow \psi \\
[\bot] &= \emptyset \\
[\top] &= S_n(T) \\
[\varphi] \cap [\psi] &= [\varphi \wedge \psi] \\
[\varphi] \cup [\psi] &= [\varphi \vee \psi] \\
[\varphi]^c &= [\neg\varphi]
\end{aligned}$$

Since

$$[\varphi \wedge \psi] = [\varphi] \cap [\psi] \text{ and } [\top] = S_n(T)$$

sets of the form $[\varphi]$ constitute a basis. The topology generated from these sets is called the *logic topology* and we have:

THEOREM 7.9. *The set $S_n(T)$ with the logic topology is a compact Hausdorff space with a basis of clopens.*

PROOF. Since $[\varphi]^c = [\neg\varphi]$ it is clear that each basic open set is also closed. In addition, if p and q are two n-types and $p \neq q$, then there is some formula φ such that $\varphi \in p$ and $\varphi \notin q$ (or vice versa). But the latter means that $\neg\varphi \in q$, so $[\varphi]$ and $[\neg\varphi]$ are two disjoint open sets with p being an element of the first set and q being an element of the second. So $S_n(T)$ is Hausdorff.

To see that $S_n(T)$ is compact, let $(U_i)_{i\in I}$ be a collection of opens such that $\bigcup_{i\in I} U_i$. The task is to find a finite subset $I_0 \subseteq I$ such that $\bigcup_{i\in I_0} U_i = S_n(T)$. Since every open set is a union of basis elements, we may just as well assume that each U_i is of the form $[\varphi_i]$. Now suppose that $\bigcup_{i\in I}[\varphi_i] = S_n(T)$ but there is no finite subset I_0 such that $\bigcup_{i\in I_0}[\varphi_i] = S_n(T)$.

Consider the partial type

$$p(\overline{x}) = \{\neg\varphi_i(\overline{x}) : i \in I\}.$$

We claim that $p(\overline{x})$ is consistent with T: for if not, there would be $i_1, \ldots, i_n \in I$ such that

$$\{\neg\varphi_{i_1}, \ldots, \neg\varphi_{i_n}\}$$

would already be inconsistent with T, by the compactness theorem. But then

$$[\neg\varphi_{i_1} \wedge \ldots \wedge \neg\varphi_{i_n}] = [\neg\varphi_{i_1}] \cap \ldots \cap [\neg\varphi_{i_n}] = \emptyset,$$

and hence

$$[\varphi_{i_1} \vee \ldots \vee \varphi_{i_n}]^c = [\neg(\varphi_{i_1} \vee \ldots \vee \varphi_{i_n})] = [\neg\varphi_{i_1} \wedge \ldots \wedge \neg\varphi_{i_n}] = \emptyset.$$

Therefore

$$[\varphi_{i_1} \vee \ldots \vee \varphi_{i_n}] = [\varphi_{i_1}] \cup \ldots \cup [\varphi_{i_n}] = S_n(T),$$

contradicting our assumption.

So the type $p(\overline{x})$ is consistent with T. But that means that p can be extended to a complete type $q(\overline{x})$ which is still consistent with T (see Lemma 7.4). So $q \in S_n(T)$, but $q \notin [\varphi_i]$ for any i as q extends p. This contradicts our assumption that $\bigcup_{i\in I}[\varphi_i] = S_n(T)$. We conclude that $S_n(T)$ is compact. □

REMARK 7.10. Compact Hausdorff spaces with a basis of clopens are called *Stone spaces*, after Marshall Stone who established a duality between these spaces and Boolean algebras.

4. Exercises

EXERCISE 23. Suppose M is an L-structure and $\sigma\colon M \to M$ is an automorphism of M. Show that for any n-tuple $\overline{m} = (m_1, \ldots, m_n)$ of elements from M, the types of $\overline{m}$ and $\sigma(\overline{m}) = (\sigma m_1, \ldots, \sigma m_n)$ are the same.

EXERCISE 24. Let κ be an infinite cardinal with $\kappa \geq |L|$, and let T be a κ-categorical L-theory without finite models. Show that if M is a model of T of cardinality κ, then M realizes all n-types over T.

EXERCISE 25. Use the previous two exercises to determine all $S_n(T)$ for

(a) $T = DLO$, the theory of dense linear orders without endpoints.
(b) $T = RG$, the theory of the random graph.
(c) $T = ACF_0$, the theory of algebraically closed fields of characteristic 0.

EXERCISE 26. In this exercise we look at the theory $VS_{\mathbb{Q}}$ of vector spaces over $\mathbb{Q}$ of positive dimension. The language of this theory contains symbols $+$ and 0, for vector addition and the null vector, as well as unary operations m_q, one for every $q \in \mathbb{Q}$, for scalar multiplication with q. The theory $VS_{\mathbb{Q}}$ has axioms expressing that $(+, 0)$ is an infinite Abelian group on which $\mathbb{Q}$ acts as a set of scalars.

(a) For which infinite κ is $VS_{\mathbb{Q}}$ κ-categorical?
(b) Show that $VS_{\mathbb{Q}}$ is complete.
(c) Determine all type spaces $S_n(T)$ for $T = VS_{\mathbb{Q}}$.

Exercise 27. Show that the theory of $(\mathbb{R}, 0, +)$ has exactly two 1-types and $\aleph_0$ many 2-types. *Hint:* Think of the previous exercise.

Exercise 28. We work in the language consisting of a single binary relation symbol E. Let T be the theory expressing that E is an equivalence relation, that all the equivalence classes are infinite and that there are infinitely many equivalence classes.

(a) Convince yourself that there is such a first-order theory T.
(b) For which infinite κ is T κ-categorical?
(c) Give a complete description of all $S_n(T)$.

Exercise 29.

(a) Consider $M = (\mathbb{Z}, +)$ and $T = \mathrm{Th}(M)$. Determine for any pair of elements $a, b \in M$ whether they realize the same or different 1-types. Are there 1-types consistent with T that are not realized in M?
(b) Idem dito for $M = (\mathbb{Z}, \cdot)$.

CHAPTER 8

Isolated types and the omitting types theorem

Types can either be isolated or not: this is the most important distinction one can make between different kinds of types. A type is isolated if it is an isolated point in the type space: this turns out to be equivalent to saying that it is generated by a single formula (for this reason isolated types are also often called principal types).

Isolated types and non-isolated types behave very differently. Indeed, suppose T is a complete theory formulated in a countable language. Then every isolated type will be realized in every model of T, while for any non-isolated type there will be at least one model in which it is omitted. The aim of this chapter is to prove these facts.

1. Isolated types

Definition 8.1. A formula $\varphi(\overline{x})$ is called *complete* or *isolating* over a theory T if $\exists\overline{x}\,\varphi(\overline{x})$ is consistent with T and we have

$$T \models \varphi(\overline{x}) \to \psi(\overline{x}) \text{ or } T \models \varphi(\overline{x}) \to \neg\psi(\overline{x})$$

for any formula $\psi(\overline{x})$.

Note that if a formula $\varphi(\overline{x})$ is complete, then

$$p(\overline{x}) = \{\, \psi(\overline{x}) : T \models \varphi(\overline{x}) \to \psi(\overline{x}) \,\}$$

is a type. Indeed, we will have $\{p\} = [\varphi]$, showing that p is isolated point in the type space. In general, we have:

Proposition 8.2. *Let T be a theory and p be a complete type of T. Then the following are equivalent:*

1. *The type p is an isolated point in the space $S_n(T)$.*
2. *The type p contains a complete formula.*
3. *There is a formula $\varphi(x_1, \ldots, x_n) \in p$ such that*

$$T \models \varphi(\overline{x}) \to \psi(\overline{x})$$

for every $\psi(\overline{x}) \in p(x)$.

Proof. These are all different ways of saying that $\{p\} = [\varphi]$ for some formula φ. $\square$

A type will be called *isolated* if it satisfies any of the equivalent conditions in the previous proposition. It will be useful to extend the notion of isolatedness to partial types, which we do as follows:

DEFINITION 8.3. Let T be an L-theory and $p(\overline{x})$ be a partial type. Then $p(\overline{x})$ is *isolated in* T if there is a formula $\varphi(\overline{x})$ such that $\exists\overline{x}\,\varphi(\overline{x})$ is consistent with T and

$$T \models \varphi(\overline{x}) \to \psi(\overline{x})$$

for all $\psi(\overline{x}) \in p(\overline{x})$.

PROPOSITION 8.4. *Let T be a complete theory and p be a partial type which is consistent with T. If p is isolated, then p is realized in every model of T.*

PROOF. Let M be a model of T and suppose that $\varphi(\overline{x})$ is a formula such that $\exists\overline{x}\,\varphi(\overline{x})$ is consistent with T and

$$T \models \varphi(\overline{x}) \to \psi(\overline{x})$$

for all $\psi(\overline{x}) \in p(\overline{x})$. If $\exists\overline{x}\,\varphi(\overline{x})$ is consistent with T and T is complete, we must have

$$T \models \exists\overline{x}\,\varphi(\overline{x}),$$

and therefore

$$M \models \exists\overline{x}\,\varphi(\overline{x}).$$

So we have some n-tuple $\overline{m}$ such that $M \models \varphi(\overline{m})$. This implies that $M \models \psi(\overline{m})$ for every $\psi \in p$, so p is realized in M. □

2. The omitting types theorem

Our next task it to prove a kind of converse to Proposition 8.4, showing that non-isolated types can be omitted. For this we need the following result, which was Proposition 2.5:

PROPOSITION 8.5. (=Proposition 2.5) *Suppose T is a consistent theory in a language L and C is a set of constants in L. If for any formula $\psi(x)$ in the language L there is a constant $c \in C$ such that*

$$\exists x\,\psi(x) \to \psi(c) \in T,$$

then T has a model whose universe consists entirely of interpretations of elements of C.

THEOREM 8.6. (Omitting types theorem) *Let T be a consistent theory in a countable language. If a partial type $p(x)$ is not isolated in T, then there is a countable model of T which omits $p(x)$.*

PROOF. Let $C = \{c_i\,;\, i \in \mathbb{N}\}$ be a countable collection of fresh constants and L_C be the language L extending with these constants. Let $\{\psi_i(x) \colon i \in \mathbb{N}\}$ be an enumeration of the formulas with one free variable in the language L_C.

We will now inductively create a sequence of sentences $\varphi_0, \varphi_1, \varphi_2, \ldots$, and then apply Proposition 8.5 to $T' = T \cup \{\varphi_0, \varphi_1, \ldots\}$ and the set of constants C.

If $n = 2i$, we take a fresh constant $c \in C$ (one that does not occur in φ_m with $m < n$) and put

$$\varphi_n = \exists x\,\psi_i(x) \to \psi_i(c).$$

This makes sure that the witnessing condition from Proposition 8.5 will be satisfied.

If $n = 2i+1$ we make sure that c_i omits $p(x)$, as follows. Consider $\delta = \bigwedge_{m<n} \varphi_m$, and write δ as $\delta(c_i, \overline{c})$ where $\overline{c}$ is a sequence of constants not containing c_i. Since $p(x)$ is not isolated, there must be a formula $\sigma(x) \in p(x)$ such that

$$T \not\models \exists\overline{y}\,\delta(x, \overline{y}) \to \sigma(x);$$

in other words, there is a formula $\sigma(x) \in p(x)$ such that $T \cup \{\exists \overline{y}\, \delta(x, \overline{y})\} \cup \{\neg\sigma(x)\}$ is consistent. Put $\varphi_n = \neg\sigma(c_i)$.

The proof is now finished by showing by induction that each $T \cup \{\varphi_0, \dots, \varphi_n\}$ is consistent and then applying Proposition 8.5. $\square$

3. Exercises

Exercise 30. Consider all the type space $S_n(T)$ from the exercises in the previous chapter. Determine for each type in $S_n(T)$ whether it is isolated or not. Also, if the type is isolated, find a complete formula in it; and if the type is not isolated, find a model in which it is omitted.

Exercise 31. Prove the generalised omitting types theorem: Let T be a consistent theory in a countable language and let $\{p_i : i \in \mathbb{N}\}$ be a sequence of partial n_i-types (for varying n_i). If none of the p_i is isolated in T, then T has a countable model which omits all p_i.

Exercise 32. Prove that the omitting types theorem is specific to the countable case: give an example of a consistent theory T in an uncountable language and a partial type in T which is not isolated, but which is nevertheless realised in every model of T.

CHAPTER 9

Prime models

Type spaces provide a lot of information about a theory. In fact, certain model-theoretic properties of a theory turn out to correspond precisely with certain topological properties of the type spaces. We will see a first example of this phenomenon here: we will prove that a theory has what is called a prime model if and only if the isolated points are dense in every type space of this theory. In order to prove this we exploit quite heavily the properties of isolated and non-isolated types that we established in the previous chapter (that is, they rely on the fact that isolated types are realized in every model of a theory, while non-isolated types can be omitted).

I should add that what I wrote in the previous paragraph is true only for sufficiently nice theories. In fact, from now on we will often assume that a theory T

- is complete,
- has infinite models, and
- is formulated in a countable language.

If T satisfies these conditions, I will call T *nice* (this is not standard terminology). Note that nice theories have models of every infinite cardinality κ, do not have finite models and are such that every type over T is already realized in a countable model of T.

1. Atomic models

Before we embark on a study of prime models, we will first look at atomic models.

Definition 9.1. A model A is *atomic* if it only realises isolated types in $S_n(\mathrm{Th}(A))$; put differently, a model is atomic if it omits all non-isolated types in $S_n(\mathrm{Th}(A))$.

Before we proceed, let us unwind this definition. Suppose A is an atomic model and $\overline{a}$ is a tuple of elements from A. Then, by definition, $p := \mathrm{tp}_A(\overline{a})$ is an isolated type over $\mathrm{Th}(A)$. This means that it contains a complete formula $\varphi(\overline{x})$ such that

$$\mathrm{Th}(A) \models \varphi(\overline{x}) \to \psi(\overline{x})$$

if and only if $\psi(\overline{x}) \in p$. What this does is reducing the "local question" whether $\overline{a}$ satisfies a formula $\psi(\overline{x})$ to the "global question" whether $A \models \varphi(\overline{x}) \to \psi(\overline{x})$. In other words, a model A is atomic if for any tuple $\overline{a}$ of elements from A there is formula $\varphi(\overline{x})$ such that for any formula $\psi(\overline{x})$ we have $A \models \psi(\overline{a})$ if and only if

$$A \models \varphi(\overline{x}) \to \psi(\overline{x}).$$

Proposition 9.2. *If A is atomic and $\overline{a} \in A$, then $(A, \overline{a})$ is atomic as well.*

PROOF. Let $\bar{b}$ be a tuple of elements from $(A, \overline{a})$. Look at $(\overline{a}, \bar{b})$. Since A is atomic there is a formula $\varphi(\overline{y}, \overline{x})$ with $A \models \varphi(\overline{a}, \bar{b})$ and

$$A \models \varphi(\overline{y}, \overline{x}) \to \psi(\overline{y}, \overline{x})$$

for every $\psi(\overline{y}, \overline{x})$ with $A \models \psi(\overline{a}, \bar{b})$. But then $\varphi(\overline{a}, \overline{x})$ is a formula satisfied by $\bar{b}$ such that

$$(A, \overline{a}) \models \varphi(\overline{a}, \overline{x}) \to \chi(\overline{x})$$

for every $\chi(\overline{x})$ with $(A, \overline{a}) \models \chi(\bar{b})$ (because each such $\chi(\overline{x})$ can be obtained from a formula $\psi(\overline{y}, \overline{x})$ with $\overline{a}$ substituted for $\overline{y}$). $\square$

For the further study of atomic models we need the notion of an elementary map.

DEFINITION 9.3. Let M and N be two L-structures. A partial function $f: X \subseteq M \to N$ from a subset X of M to N will be called an *elementary map* if

$$M \models \varphi(m_1, \ldots, m_n) \Leftrightarrow N \models \varphi(f(m_1), \ldots, f(m_n))$$

for all L-formulas $\varphi(x_1, \ldots, x_n)$ and elements $m_1, \ldots, m_n \in X$. Note that this is equivalent to saying that $(M, x)_{x \in X} \equiv (N, fx)_{x \in X}$.

PROPOSITION 9.4. *Let $f: \{a_1, \ldots, a_n\} \subseteq A \to M$ be an elementary map whose domain is a finite subset of an atomic model A. Then for any $a \in A$ there is an elementary map $g: \{a_1, \ldots, a_n\} \cup \{a\} \to M$ which extends f.*

PROOF. Suppose $f: \{a_1, \ldots, a_n\} \subseteq A \to M$ is an elementary map whose domain is a finite subset of an atomic model A. Let us write $\overline{a}$ for the n-tuple $\langle a_1, \ldots, a_n \rangle$ and $f\overline{a}$ for the n-tuple $\langle fa_1, \ldots, fa_n \rangle$. The fact that f is an elementary map is equivalent to saying that $(A, \overline{a}) \equiv (M, f\overline{a})$.

So let $a \in A$. Since $(A, \overline{a})$ is atomic by the previous proposition, there is a formula $\varphi(x)$ such that $(A, \overline{a}) \models \varphi(a)$ and

$$(A, \overline{a}) \models \varphi(x) \to \psi(x)$$

for any formula $\psi(x)$ such that $(A, \overline{a}) \models \psi(a)$. Because $(A, \overline{a}) \models \varphi(a)$ and $(A, \overline{a}) \equiv (M, f\overline{a})$, we have $(A, \overline{a}) \models \exists x\, \varphi(x)$ and $(M, f\overline{a}) \models \exists x\, \varphi(x)$. So let $m \in M$ be such that $(M, f\overline{a}) \models \varphi(m)$. Then the type of a over $(A, \overline{a})$ and the type of m over $(M, f\overline{a})$ both contain the formula $\varphi(x)$, which is complete over $\mathrm{Th}(A, \overline{a}) = \mathrm{Th}(M, f\overline{a})$. This implies that these types are identical and we have $(M, \overline{a}, a) \equiv (M, f\overline{a}, m)$. So if we put $g(a) = m$ and $g(a_i) = f(a_i)$, then g is an elementary map extending f. $\square$

THEOREM 9.5. *Suppose A and M are two L-structures. If A is countable and atomic and $A \equiv M$, then A embeds elementarily into M.*

PROOF. Suppose $\{a_0, a_1, a_2, \ldots\}$ is an enumeration of A. Using the previous proposition one can construct an increasing sequence of elementary maps $f_n: \{a_0, \ldots, a_n\} \to M$, starting with $f_0 = \emptyset$ (which is an elementary map as $A \equiv M$). But then $f = \bigcup_{n \in \mathbb{N}} f_n$ is an elementary embedding A into M. $\square$

THEOREM 9.6. *Suppose A and B are two L-structures which are both countable and atomic. If $A \equiv B$, then $A \cong B$.*

Proof. We use the back and forth method. So suppose $\{a_0, a_1, a_2, \ldots\}$ and $\{b_0, b_1, b_2, \ldots\}$ are enumerations of A and B, respectively. Using Proposition 9.4 one can construct an increasing sequence of elementary maps $f_n : X \subseteq A \to B$ such that $a_n \in \mathrm{dom}(f_{2n+1})$ and $b_n \in \mathrm{ran}(f_{2n+2})$, starting with $f_0 = \emptyset$. Then $f = \bigcup_{n \in \mathbb{N}} f_n$ is an isomorphism between A and B. □

2. Prime models

Definition 9.7. Let T be a theory. A model M of T is called *prime* if it can be elementarily embedded into any model of T.

Theorem 9.8. *A model of a nice theory T is prime iff it is countable and atomic.*

Proof. $\Rightarrow$: Let A be a prime model of a nice theory T. As a nice theory has countable models and A embeds in any model, A has to be countable as well. Moreover, if p is a non-isolated type of T, then there is a model B of T in which it is omitted, by the Omitting Types Theorem. Since A embeds elementarily into B, the type p will be omitted in A as well.

$\Leftarrow$: Let A be a countable and atomic model of a nice theory T and M be any other model of T. Since T is complete, we have $A \equiv M$, so A embeds elementarily into M by Theorem 9.5. □

Corollary 9.9. *Any two prime models of a nice theory T are isomorphic.*

Proof. This follows from Theorem 9.6 and Theorem 9.8. □

Theorem 9.10. *A nice theory T has a prime model iff the isolated n-types are dense in $S_n(T)$ for all n.*

Proof. Let us first translate the statement that that isolated n-types are dense in $S_n(T)$ in more logical terms. To say that the isolated types are dense means that every non-empty (basic) open set contains at least one isolated type: so any $[\varphi]$ which is not empty contains at least one isolated type p. But if p is isolated there is a complete formula ψ such that $\{p\} = [\psi] \subseteq [\varphi]$. So the isolated types are dense in $S_n(T)$ if every consistent formula $\varphi(\overline{x})$ is the consequence over T of some complete formula $\psi(\overline{x})$.

$\Rightarrow$: Let A be a prime model of T. Because a consistent formula $\varphi(\overline{x})$ is realised in *all* models of a complete theory, it is realized in A as well, by $\overline{a}$ say. Since A is atomic, $\varphi(\overline{x})$ belongs to the isolated type $\mathrm{tp}_A(\overline{a})$, so is the consequence over $\mathrm{Th}(A) = T$ of some complete formula $\psi(\overline{x})$.

$\Leftarrow$: Suppose isolated types are dense in every type space $S_n(T)$. Then we define for each natural number n a partial n-type

$$p_n(x_1, \ldots, x_n) = \{\, \neg\varphi(x_1, \ldots, x_n) : \varphi \text{ is complete} \,\},$$

and claim that these are not isolated. Because if p_n would be isolated there would be a consistent formula $\psi(\overline{x})$ such that

$$T \models \psi(\overline{x}) \to \neg\varphi(\overline{x})$$

for any complete formula $\varphi(\overline{x})$. But this would mean that $\psi(\overline{x})$ could not be a consequence of any complete formula, contradicting the fact that the isolated types are dense. So by the generalised omitting types theorem there is a countable model A omitting all p_n. But a structure omitting all p_n has to be atomic. □

3. Exercises

Exercise 33. Let T be the theory of $(\mathbb{R}, <, Q)$ where Q is a predicate for the rational numbers. Does T have a prime model?

CHAPTER 10

ω-categoricity

In this chapter we will see another example of a model-theoretic properties of a theory which corresponds precisely with certain topological properties of the type spaces. Indeed, we will prove that a nice theory is ω-categorical if and only if all its type spaces are finite. (Recall that a theory is ω-categorical if all its countably infinite models are isomorphic.)

1. ω-categorical theories

THEOREM 10.1. (Ryll-Nardzewski Theorem) *For a nice theory T the following are equivalent:*

1. *T is ω-categorical;*
2. *all n-types are isolated;*
3. *all models of T are atomic;*
4. *all countable models of T are prime.*

PROOF. (1) $\Rightarrow$ (2): If $S_n(T)$ contains a non-isolated type p then there is a countable model where p is realized and a countable model where p is omitted (by the Omitting Types Theorem). So T cannot be ω-categorical.

(2) $\Rightarrow$ (3): If all types of a theory T are isolated, then any model of T can only realize isolated types. So all models of T are atomic.

(3) $\Rightarrow$ (4) follows from Theorem 9.8.

(4) $\Rightarrow$ (1) follows from Corollary 9.9. $\square$

So a nice theory T is ω-categorical iff all types over T are isolated. But to say that every type is isolated means that there are only finitely many types.

PROPOSITION 10.2. *The following are equivalent for any theory T:*

1. *All n-types are isolated.*
2. *Every $S_n(T)$ is finite.*
3. *For for every n there are only finite many formulas $\varphi(x_1, \ldots, x_n)$ up to equivalence relative to T.*

PROOF. (1) $\Leftrightarrow$ (2) holds because $S_n(T)$ is a compact Hausdorff space.

(2) $\Rightarrow$ (3) If there are only finitely many n-types $p_1, \ldots, p_m$, then each p_i is isolated by some complete formula ψ_i. We claim that each formula with free variables among $x_1, \ldots, x_n$ is equivalent over T to some disjunction of the ψ_i, showing that up to logical equivalence there are only finitely many formulas with free variables among $x_1, \ldots, x_n$.

If φ is any formula with free variables among $x_1, \dots, x_n$, then $[\varphi] \subseteq S_n(T)$, so

$$[\varphi] = \{p_i : i \in I\}$$

for some $I \subseteq \{1, \dots, m\}$. But then

$$[\varphi] = \bigcup_{i \in I} \{p_i\} = \bigcup_{i \in I} [\psi_i] = [\bigvee_{i \in I} \psi_i],$$

so $T \models \varphi \leftrightarrow \bigvee_{i \in I} \psi_i$.

(3) $\Rightarrow$ (2): If every formula $\varphi(x_1, \dots, x_n)$ is equivalent modulo T to one of

$$\psi_1(x_1, \dots, x_n), \dots, \psi_m(x_1, \dots, x_n),$$

then every n-type is completely determined by saying which ψ_i it does and does not contain. $\square$

COROLLARY 10.3. *Suppose A is a model in a countable language and $\overline{a}$ is some tuple of elements from A. Then* $\mathrm{Th}(A)$ *is ω-categorical iff* $\mathrm{Th}(A, \overline{a})$ *is ω-categorical.*

PROOF. Every m-type $p(\overline{a}, \overline{x})$ of $\mathrm{Th}(A, \overline{a})$ determines an $(n+m)$-type $p(\overline{y}, \overline{x})$ of $\mathrm{Th}(A)$: so if there are only finitely many $(n+m)$-types consistent with $\mathrm{Th}(A)$, then there are only finitely many m-types consistent with $\mathrm{Th}(A, \overline{a})$.

Conversely, an m-type p consistent with $\mathrm{Th}(A)$ will be realized by some elements $\overline{c}$ in some elementary extension B of A. If $i\colon A \to B$ is the elementary embedding and $\overline{b} = i\overline{a}$, then $(B, \overline{b})$ is an elementary extension of $(A, \overline{a})$. Then $q = \mathrm{tp}_{(B, \overline{b})}(\overline{c})$ is a type over $\mathrm{Th}(A, \overline{a})$ extending p. Since $p \subseteq q$, these extensions q have to be different for different types p, and therefore the theory $\mathrm{Th}(A)$ cannot have more n-types than $\mathrm{Th}(A, \overline{a})$. So if the latter has only finitely many n-types, then so does the former. $\square$

2. Vaught's Theorem

All of this has the following odd consequence. There are nice theories T_n having, up to isomorphism, n models, for $n = 1, 3, 4, 5, 6, \dots$ (see Exercise 34 below). But the case $n = 2$ is impossible.

THEOREM 10.4. (Vaught's Theorem) *A nice theory cannot have exactly two countable models (up to isomorphism).*

PROOF. If T is a nice theory which has more than one model (up to isomorphism), then T is not ω-categorical, so there must be some type p over T which is not isolated. But then there is a countable model A in which p is realized, by $\overline{a}$ say, and a model B in which p is omitted. Clearly, A and B cannot be isomorphic.

Since $T = \mathrm{Th}(A)$ is not ω-categorical, also $\mathrm{Th}(A, \overline{a})$ is not ω-categorical, by the previous corollary. So, again, there is a type q over $\mathrm{Th}(A, \overline{a})$ which is not isolated. Now we make a case distinction:

(1) If q is realized in $(A, \overline{a})$, let $(C, \overline{c})$ be a countable model in which it is omitted. Then C cannot be isomorphic to A; but C can also not be isomorphic to B because C realizes p, while B omits it.
(2) If q is omitted in $(A, \overline{a})$, let $(C, \overline{c})$ be a countable model in which it is realized. Then C cannot be isomorphic to A; but C can also not be isomorphic to B because C realizes p, while B omits it.

We conclude that any nice theory which is not ω-categorical must have at least three non-isomorphic models. $\square$

3. Exercises

EXERCISE 34. Let $L_3 = \{<, c_0, c_1, c_2, \ldots\}$, where $c_0, c_1, \ldots$ are constant symbols. Let T_3 be the theory of dense linear orders with sentences added asserting $c_0 < c_1 < \ldots$.

(a) Show that T_3 is a nice theory which has exactly three countable models up to isomorphism. *Hint:* Consider the questions: Does $c_0, c_1, c_2, \ldots$ have an upper bound? A least upper bound?
(b) Let $L_4 = L_3 \cup \{P\}$, where P is a unary predicate. Let T_4 be T_3 with the added sentences $P(c_i)$ and
$$\forall x\, \forall y\, \big(\, x < y \to \exists z\, \exists w\, (\, x < z < y \wedge x < w < y \wedge P(z) \wedge \neg P(w))\,\big).$$
In other words, P is a dense-codense subset. Show that T_4 is a nice theory with exactly four countable models.
(d) Generalise (c) to give examples of nice theories which have exactly n countable models for $n = 5, 6, \ldots$

EXERCISE 35. A theory T has *quantifier elimination* if for any formula $\varphi(\overline{x})$ there is a quantifier-free formula $\psi(\overline{x})$ such that
$$T \models \varphi(\overline{x}) \leftrightarrow \psi(\overline{x}).$$

(a) Suppose T is a nice ω-categorical theory and each $p \in S_n(T)$ contains a complete formula which is also quantifier-free. Deduce that T has quantifier elimination.
(b) Use (a) to show that $T = DLO$ and $T = RG$ have quantifier elimination.

CHAPTER 11

κ-saturated models

In the previous chapters we have looked at prime models and models of nice ω-categorical theories: these models realize few types (indeed, as few types as possible). In this and the next chapters we look at rich models which will realize as many types as possible: these are called *κ-saturated models*. In fact, κ-saturated models of a complete theory T realise all types over T. For this reason κ-saturated models are useful for computing type spaces: the type space of a complete theory T can be exhaustively analysed by looking at configurations of elements in a single κ-saturated model.

1. κ-saturated models: definition

To define κ-saturated models we need to introduce some notational conventions. Let A be an L-structure and X a subset of A. We often refer to the elements in X as *parameters*. In addition, we will use the following notation:

- We write L_X for the language L extended with constants for all elements of X.
- We write $(A, a)_{a \in X}$ for the L_X-expansion of A where we interpret the constant $a \in X$ as itself.

Definition 11.1. Let A be an infinite L-structure and κ be an infinite cardinal. We say that A is *κ-saturated* if the following condition holds:

> if X is any subset of A with $|X| < \kappa$ and $p(x)$ is any 1-type in L_X that is finitely satisfiable in $(A, a)_{a \in X}$, then $p(x)$ can be realized in $(A, a)_{a \in X}$.

We first make a number of observations:

(1) If A is κ-saturated, then $|A| \geq \kappa$ and A is also λ-saturated for any infinite $\lambda \leq \kappa$.
(2) If Y is a subset of a κ-saturated model A and $|Y| < \kappa$, then $(A, y)_{y \in Y}$ is κ-saturated as well. The reason for this is that any 1-type over a set of parameters X with $|X| < \kappa$ in $(A, y)_{y \in Y}$ is also a 1-type over the set of parameters $X \cup Y$ in A, and $|X \cup Y| < \kappa$.
(3) The definition of κ-saturation only talks about 1-types; however, if $p(x_1, \ldots, x_n)$ is an n-type over a set of parameters X with $|X| < \kappa$ and p is finitely satisfiable in an κ-saturated model A, then it is realized. To see this, consider the types
$$p_1(x_1), p_2(x_1, x_2), \ldots, p_n(x_1, \ldots, x_n)$$
which are the types obtained from p by considering only those formulas that contain $x_1, \ldots, x_i$ free. Then p_1 is realized, because it is finitely satisfiable in A and A is κ-saturated; moreover, if $a_1, \ldots, a_i$ realize p_i, then $p_{i+1}(a_1, \ldots, a_i, x_{i+1})$ is finitely satisfied in $(A, y)_{y \in X \cup \{a_1, \ldots, a_i\}}$, by Lemma 11.2 below, and hence realized by some a_{i+1} by the previous remark. So each p_i is realized, including $p = p_n$.

(4) The definition only talk about complete types, but this is not a genuine restriction. Indeed, Lemma 7.7(3) tells us that any partial type that is finitely satisfied in a model can be extended to a complete type that is finitely satisfied in that model.

LEMMA 11.2. *Let $p(x_1, \dots, x_n, y)$ be an $(n+1)$-type and let $q(x_1, \dots, x_n)$ be the n-type obtained from p by taking only those $\varphi \in p$ that do not contain y free. If p is finitely satisfiable in M and $(a_1, \dots, a_n)$ realizes q in M, then also $p(a_1, \dots, a_n, y)$ is finitely satisfiable in M.*

PROOF. Let $\varphi_1(\underline{x}, y), \dots, \varphi_n(\underline{x}, y)$ be finitely many formulas in p. The formula

$$\psi(\underline{x}) := \exists y \left(\varphi_1(\underline{x}, y) \wedge \dots \wedge \varphi_n(\underline{x}, y) \right)$$

has to belong to p: if it would not, its negation would have to belong to p, and p could not be finitely satisfiable. This means that $\psi \in q$, by definition, so $M \models \psi(\underline{a})$. We conclude that $p(\underline{a}, y)$ is finitely satisfiable. $\square$

As promised, we have:

PROPOSITION 11.3. *Let M be an κ-saturated model of a complete theory T. Then M realizes any type over T.*

PROOF. Let M be a model of a complete theory T. If p belongs to $S_n(T)$ then p is finitely satisfiable in M by Proposition 7.6. So if M is κ-saturated, then p will be realized in M. $\square$

2. κ-saturated models: existence

It can hard to determine whether a concrete model is κ-saturated or not: we will see some criteria later in this chapter. However, it is not so hard to prove that they exist. In fact, we have:

THEOREM 11.4. *Every structure has an κ-saturated elementary extension. So any consistent theory has κ-saturated models for each κ.*

The proof relies on the following lemma:

LEMMA 11.5. *Let A be an L-structure. There exists an elementary extension B of A such that for every subset $X \subseteq A$, every 1-type in L_X which is finitely satisfied in $(A, a)_{a \in X}$ is realized in $(B, a)_{a \in X}$.*

PROOF. Let $(p_i(x_i))_{i \in I}$ be the collection of all such 1-types and b_i be new constants. Consider:

$$T := \bigcup_{i \in I} p_i(b_i).$$

Since the p_i are finitely satisfiable in $(A, a)_{a \in A}$, every finite subset of T can be satisfied in $(A, a)_{a \in A}$. So, by the compactness theorem, T has a model B. Since T contains $\mathrm{ElDiag}(A)$, the model A embeds into B. $\square$

PROOF. (Of Theorem 11.4.) Let us first look at the case $\kappa = \omega$. Let A be an L-structure. We will build an elementary chain of L-structures $(A_i : i \in \mathbb{N})$. We set $A_0 = A$ and at successor stages we apply the previous lemma. Now let B be the colimit of the entire chain.

We claim B is ω-saturated: for if $X \subseteq B$ is a finite subset, then X is already a finite subset of some A_i and any 1-type p with parameters from X will be realized in A_{i+1}, by construction, say by $a \in A_{i+1}$. Since the embedding from A_{i+1} in B is elementary, the type p will also be realized by a in B.

Note that in the previous argument we relied on the following property of ω: if $(A_i)_{i\in\omega}$ is an increasing sequence of sets and X is a subset of $\bigcup_{i\in\omega} A_i$ with $|X| < \omega$, then $X \subseteq A_i$ for some $i \in \omega$. An infinite cardinal κ is called *regular* if for any increasing sequence of sets $(A_i)_{i\in\kappa}$ and any subset X of $\bigcup_{i\in\kappa} A_i$ with $|X| < \kappa$ there is an $i \in \kappa$ with $X \subseteq A_i$. It is not hard to see that the argument we just gave works for every regular cardinal: if κ is a regular cardinal and A is any model, then we can create by transfinite recursion an elementary chain $(A_i : i \in \kappa)$ of models, starting with $A_0 = A$; at successor stages we apply Lemma 11.5 and at limit stages we take colimits. The colimit of the entire chain will be a model in which A embeds elementarily and it will be κ-saturated, because κ is regular.

At this point the proof would be finished once we know that there are arbitrarily large regular cardinals, that is, if for every cardinal κ there is a regular cardinal λ with $\lambda \geq \kappa$. According to the set theorists this is true: indeed, $\lambda = \kappa^+$ is always regular. □

3. Tests for κ-saturation

In this section we give an equivalent characterisation of κ-saturation which is often easier to verify. For this we need a lemma and a definition.

LEMMA 11.6. *Let $f: X \subseteq M \to N$ be an elementary map, and $m \in M$. If κ is an infinite cardinal such that N is κ-saturated and $|X| < \kappa$, then f can be extended to an elementary map whose domain includes m.*

PROOF. If $f: X \subseteq M \to N$ is an elementary map, then $(M, x)_{x\in X} \equiv (N, fx)_{x\in X}$. So if $p = \mathrm{tp}_{(M,x)_{x\in X}}(m)$, then p is finitely satisfied in $(N, fx)_{x\in X}$ by Lemma 7.7(1). Since $(N, fx)_{x\in X}$ is also κ-saturated, we find an element $n \in N$ realizing p in this model. This means that we can extend f to an elementary map g whose domain includes m by putting $g(x) = f(x)$ for every $x \in X$ and $g(m) = n$. □

DEFINITION 11.7. A model M is called *κ-homogeneous*, if for any subset X of M with $|X| < \kappa$, any elementary map $f: X \subseteq M \to M$ and any element $m \in M$, the map f can be extended to an elementary map g whose domain includes m. A model M is called *κ-universal*, if for any model N with $N \equiv M$ and $|N| < \kappa$ there is an elementary embedding $N \preceq M$.

THEOREM 11.8. *Let M be an infinite L-structure and κ be an infinite cardinal with $\kappa \geq |L|$. Then M is κ-saturated if and only if M is κ-homogeneous and κ^+-universal.*

PROOF. Assume M is a κ-saturated L-structure with $\kappa \geq |L|$. Lemma 11.6 immediately implies that M is κ-homogeneous, so it suffices to prove that M is also κ^+-universal. To this purpose let N be a model with $N \equiv M$ and $|N| \leq \kappa$. Choose an enumeration $N = (n_\alpha)_{\alpha\in\kappa}$; the idea is to construct by transfinite recursion on α an increasing sequence of elementary maps $f_\alpha: \{n_\beta : \beta < \alpha\} \subseteq N \to M$. (Note that $|\{n_\beta : \beta < \alpha\}| < \kappa$ for each $\alpha \in \kappa$.) Since $N \equiv M$, we can start by putting $f_0 = \emptyset$; at successor stages we use Lemma 11.6 and at limit stages we take unions.

Conversely, suppose M is an infinite model which is κ-homogeneous and κ^+-universal and p is a complete 1-type with parameters $X \subseteq M$ and $|X| < \kappa$ which is finitely satisfied in M. We know by the Skolem-Löwenheim Theorems that p is realized by some element n in some L_X-structure N with $|N| \leq \kappa$. Since M is a κ^+-universal L-structure there is an L-elementary embedding $i: N \to M$. Note that i need not be an L_X-elementary embedding, so that for any $x \in X$ there may be a difference between $i(x^N)$ and x^M. But by the completeness of p we know that the partial map f from M to itself defined by sending $i(x^N)$ to x^M is elementary. Since M is κ-homogeneous, we know that this elementary map can be extended by one whose domain includes $i(n)$; write g for such an extension and $m = g(i(n))$. Then m realizes p in M. $\square$

4. Exercises

EXERCISE 36. Let L be a language and κ be an infinite cardinal with $\kappa \geq |L|$. An infinite L-structure M is called *strongly κ-homogeneous* if every elementary map $f: X \subseteq M \to M$ with $|X| < \kappa$ can be extended to an automorphism of M.

(a) Show that a κ-homogeneous model of cardinality κ is strongly κ-homogeneous.
(b) Show that a saturated model of cardinality κ is strongly κ-homogeneous.
(c) Show that prime models of nice theories are strongly ω-homogeneous.
(d) Give an example of a model which is ω-saturated but not strongly ω-homogeneous.

EXERCISE 37. Let M be an infinite L-structure and κ be an infinite cardinal with $\kappa > |L| + \aleph_0$. Show that M is κ-saturated if and only if it is κ-homogeneous and κ-universal.

EXERCISE 38. Suppose U is an non-principal ultrafilter on $\mathbb{N}$. Let $(M_i)_{i \in \mathbb{N}}$ be a sequence of L-structures, and let ${}^*M = \prod M_i / U$.

Let $A \subseteq {}^*M$ be arbitrary, and choose for each $a \in A$ an $f_a \in \prod M_i$ such that $a = [f_a]$. Let $p(x) = \{\varphi_i(x) : i < \omega\}$ be a set of L_A-formulas such that $p(x)$ is finitely satisfiable in ${}^*M$. By taking conjunctions, we may, withour loss of generality, assume that $\varphi_{i+1}(x) \to \varphi_i(x)$ for $i < \omega$. Let $\varphi_i(x)$ be $\theta_i(x, a_{i,1}, \ldots, a_{i,m_i})$, where θ_i is an L-formula.

(a) Let
$$D_i = \{\, n < \omega : M_n \models \exists x\, \theta_i(x, f_{a_{i,1}}(n), \ldots, f_{a_{i,m_i}}(n)) \,\}.$$
Show that $D_i \in U$.
(b) Find $g \in \prod M_i$ such that if $i \leq n$ and $n \in D_i$, then
$$M_n \models \theta_i(g(n), f_{a_{i,1}}(n), \ldots, f_{a_{i,m_i}}(n)).$$
(c) Show that g realizes $p(x)$. Where do you use the fact that U is non-principal?
(d) Assume that L is countable. Conclude that ${}^*M$ is $\aleph_1$-saturated.
(e) Show that if the Continuum Hypothesis holds then every nice theory has a saturated model with size $\aleph_1$.

EXERCISE 39. Let κ be an infinite cardinal and suppose T is a κ-categorical theory in a countable language. Show that if M is an ω-homogeneous model of cardinality κ, then M is ω-saturated.

CHAPTER 12

Saturated models and small theories

In the previous chapter we have defined the notion of a κ-saturated model, and we have seen that κ-saturated models have size at least κ, and that every consistent theory has κ-saturated models for each κ. A model M which is κ-saturated and has size κ is often simply called *saturated*. Now it is not true that every consistent theory has saturated models of every possible size κ.

For example, take a language L consisting of a countable number of unary predicates $P_0, P_1, P_2, \ldots$, and consider the following L-structure M: its elements are the finite subsets of the natural numbers and for such an $m \in M$ we will say that it has the property P_n precisely when $n \in m$. Let $T = \mathrm{Th}(M)$ (note that T is a nice theory). For each function $f\colon \mathbb{N} \to \{0,1\}$ we have a partial type

$$p_f = \{P_i(x)\colon f(i) = 1\} \cup \{\neg P_i(x)\colon f(i) = 0\}.$$

These are finitely satisfiable in M, so consistent with T, meaning that an ω-saturated model would have to realize all p_f. But an element realizing p_f cannot also realize p_g when $g \neq f$, hence an ω-saturated model of T would have to have size at least that of the continuum. In particular, T does not have countable saturated models. (A fancier version of this example would take the theory $T = \mathrm{Th}(\mathbb{N}, +, \cdot, 0, 1)$ and consider partial types p_f containing formulas saying that x is divisible by the nth prime number if $f(n) = 1$, and not divisible by that prime number if $f(n) = 0$.)

In this chapter we will look at saturated models and isolate a necessary and sufficient condition for nice theories to have a countable saturated model. We will also show that if a nice theory has a countable saturated model, it must also have a prime model.

1. Saturated models

Definition 12.1. An infinite model M is called *saturated* if it is $|M|$-saturated.

Theorem 12.2. *Suppose A and B are two saturated models having the same cardinality. If A and B are elementarily equivalent, then they are isomorphic.*

Proof. Suppose $|A| = |B| = \kappa$ and $A = (a_\alpha)_{\alpha\in\kappa}$ and $B = (b_\alpha)_{\alpha\in\kappa}$ are enumerations of A and B respectively. Assume also that $A \equiv B$. We will use back and forth to show $A \cong B$: indeed, we will create by transfinite recursion an increasing sequence of elementary maps $f_\alpha\colon X \subseteq A \to B$ with $|X| < \kappa$, such that for any limit ordinal $\lambda < \kappa$ and natural number n we have $a_{\lambda+n} \in \mathrm{dom}(f_{\lambda+2n+2})$ and $b_{\lambda+n} \in \mathrm{ran}(f_{\lambda+2n+1})$. Then $f = \bigcup_{\alpha\in\kappa} f_\alpha$ is the desired isomorphism.

Recall that Lemma 11.6 told us that for any $m \in M$ and any elementary map $f\colon X \subseteq M \to N$, where $|X| < \kappa$ and N is κ-saturated, there is an elementary map $g\colon X \cup \{m\} \subseteq M \to N$

extending f. So we can create the increasing sequence of elementary maps by starting with $f_0 = \emptyset$, applying this lemma at the successor stages and taking unions at limit stages. □

Corollary 12.3. *For a nice theory T the following are equivalent:*

(1) *T is ω-categorical;*
(2) *all models of T are ω-saturated;*
(3) *all countable models of T are saturated.*

Proof. (1) $\Rightarrow$ (2): Assume T is a nice ω-categorical theory and $\overline{a}$ is a finite tuple of parameters from a model A of T. Assume moreover that $p(x)$ is a 1-type which is finitely satisfiable in $(A, \overline{a})$. Then $\mathrm{Th}(A, \overline{a})$ is ω-categorical as well by Corollary 10.3 and therefore p is isolated over this theory by Theorem 10.1. So p is realized in $(A, \overline{a})$ by Proposition 8.4.

(2) $\Rightarrow$ (3) is obvious, while (3) $\Rightarrow$ (1) follows from the previous theorem. □

2. Small theories

In this section we will characterise those nice theories which have countable saturated models. We will also show that nice theories which have countable saturated models have prime models as well.

Intuitively, a countable ω-saturated model has to harmonize two antagonistic tendencies: on the one hand such models are rich, because ω-saturated; on the other hand, they are small, because only countable. You may suspect that theories can only have such models if their type spaces are not too big, and you would be right.

Definition 12.4. A theory T is *small* if all its type spaces are countable.

Theorem 12.5. *A nice theory T has a countable ω-saturated model if and only if it is small.*

Proof. If T is complete and has an ω-saturated model M, then every n-type is realized in M. So if M is countable, there can be at most countably many n-types for any n.

For the other direction, we take a closer look at the proof of Theorem 11.4 and assume that A is a model of small theory T. First of all, we may assume that A is countable (by downward Löwenheim-Skolem). In that case how many 1-types $p(\underline{a}, x)$ are there where $\underline{a}$ is a finite set of parameters from A? The answer is that there at most countably many, because the collection of finite sequences with parameters from A is countable and there are countably many types of the form $p(\underline{y}, x)$. This means that the model B in the proof of Lemma 11.5 may be taken to be countable as well. And that in turn means that in the proof of Theorem 11.4 we have to consider a countable chain of countable models: but then its colimit, which was an ω-saturated model, is countable as well. □

To prove that nice and small theories have prime models, we need to understand these small theories a bit better.

Definition 12.6. Let $\{0,1\}^*$ be the set of finite sequences consisting of zeros and ones. A *binary tree* of formulas in variables $\overline{x} = x_1, \ldots, x_n$ over T is a collection $\{\varphi_s(\overline{x}) \colon s \in \{0,1\}^*\}$ such that $T \models \big(\varphi_{s0}(\overline{x}) \vee \varphi_{s1}(\overline{x})) \to \varphi_s(\overline{x})\big)$ and $T \models \neg\big(\varphi_{s0}(\overline{x}) \wedge \varphi_{s1}(\overline{x})\big)$.

Theorem 12.7. *The following are equivalent for a nice theory T:*

(1) $|S_n(T)| < 2^\omega$.
(2) *There is no binary tree of consistent formulas in $x_1, \dots, x_n$ over T.*
(3) $|S_n(T)| \leq \omega$.

Proof. (1) ⇒ (2): We show that the existence of a binary tree of consistent formulas implies that the type space has size at least that of the continuum. If $\{\varphi_s(\overline{x}) \colon s \in \{0,1\}^*\}$ is a binary tree of consistent formulas, then

$$p_\alpha = \{\varphi_s \colon s \subseteq \alpha\}$$

is a consistent partial type for every $\alpha \colon \mathbb{N} \to \{0,1\}$. Since consistent partial types can be extended to complete types and nothing can realize both p_α and p_β when α and β are distinct, we see that the existence of a binary tree of consistent formulas implies that there are at least 2^ω many types.

(2) ⇒ (3): We show that the uncountability of $S_n(T)$ implies that there must exist a binary tree of consistent formulas. If $|S_n(T)| > \omega$, then we have $|[\varphi]| > \omega$ for any tautology φ. So we can construct a binary tree of consistent formulas by repeated application of the following claim.

Claim: If $|[\varphi]| > \omega$, then there is a formula $\psi(\overline{x})$ such that $|[\varphi \wedge \psi]| > \omega$ and $|[\varphi \wedge \neg\psi]| > \omega$. *Proof:* Suppose not. Define

$$p(\overline{x}) \colon = \{\psi(\overline{x}) \colon |[\varphi \wedge \psi]| > \omega\}.$$

By assumption this collection contains a formula $\psi(\overline{x})$ or its negation, but not both. In addition, if p contains $\psi_0 \vee \psi_1$, then

$$|[\varphi \wedge (\psi_0 \vee \psi_1)]| = |[\varphi \wedge \psi_0] \cup [\varphi \wedge \psi_1]| > \omega,$$

so p will contain either ψ_0 or ψ_1. This implies that if p contains $\psi_1, \dots, \psi_n$ then it also contains $\psi_1 \wedge \ldots \wedge \psi_n$: for if $\psi_1 \wedge \ldots \wedge \psi_n \notin p$, then $\neg(\psi_1 \wedge \ldots \wedge \psi_n) \in p$, hence $\neg\psi_i \in p$ for some i. Since each $\psi \in p$ is consistent, this implies that each finite subset of p is consistent; hence p is consistent and therefore a complete type.

But now we arrive at a contradiction, as follows: if $\psi \notin p$, then $|[\varphi \wedge \psi]| \leq \omega$, by definition. In addition, the language is countable, so

$$[\varphi] = \bigcup_{\psi \notin p} [\varphi \wedge \psi] \cup \{p\}$$

is a countable union of countable sets and hence countable, contradicting our assumption for φ.

(3) ⇒ (1): This is clear, because $\omega < 2^\omega$. □

Corollary 12.8. *If T is nice and small, then isolated types are dense. So T has a prime model.*

Proof. If isolated types are not dense, then there is a consistent $\varphi(\overline{x})$ which is not a consequence of a complete formula. Call such a formula *perfect*. We claim that perfect formulas can be “decomposed” into two consistent formulas which are jointly inconsistent. Repeated application of this claim leads to a binary tree of consistent formulas, so T cannot be small, by the previous theorem.

To see that any perfect formula φ can be decomposed into two perfect formulas, note that perfect formulas cannot be complete, so there is a formula ψ such that both $\varphi \wedge \psi$ and $\varphi \wedge \neg\psi$

are consistent. But as these formulas imply φ and φ is not a consequence of a complete formula, these formulas have to be perfect as well. □

3. Exercises

EXERCISE 40. Let T be a theory in a countable language without a binary tree of consistent formulas. Show that T is small.

CHAPTER 13

Quantifier elimination

In this penultimate chapter we take a closer look at quantifier elimination. We have seen the concept already a few times, but to remind you: a theory T has *quantifier elimination* if, over T, any formula is equivalent to one without quantifiers. As one can imagine, this makes models of T much easier to understand. Indeed, quantifier elimination is so useful that even when a theory T does not have quantifier elimination, model theorists will typically search for natural extensions of T which do have quantifier elimination. We will see a couple of examples of that.

1. Consequences of quantifier elimination

Definition 13.1. A theory T in a language L has *quantifier elimination* if for any L-formula $\varphi(x_1, \ldots, x_n)$ there is quantifier-free L-formula $\psi(x_1, \ldots, x_n)$ such that

$$T \models \varphi \leftrightarrow \psi.$$

An easy, but crucial observation is the following:

Proposition 13.2. *Assume T is an L-theory with quantifier elimination and A and B are two L-structures which model T. Then any embedding $h\colon A \to B$ is elementary, and any local isomorphism between A and B is an elementary map.*

Corollary 13.3. *Suppose T is an L-theory with quantifier elimination and κ is an infinite cardinal with $\kappa \geq |L|$. If M is an infinite model of T such that:*

(1) *every model N of T with $|N| \leq \kappa$ embeds into M, and*
(2) *for any element $m \in M$ and any local isomorphism $f\colon X \subseteq M \to M$ where X is a subset of M with $|X| < \kappa$, the map f can be extended to a local isomorphism whose domain includes m.*

Then M is κ-saturated.

Proof. Immediate from Theorem 11.8 and the previous proposition. □

Corollary 13.4. *If a theory T has quantifier elimination and there is a model M of T that can be embedded into every other model of T, then T is complete and M is prime.*

Proof. If N is any model of T, then M can be embedded into it. By Proposition 13.2 this embedding is elementary. Therefore M is prime and all models of T are elementarily equivalent. So T is complete. □

2. Tests for quantifier elimination

Clearly, in order to use the results from the previous section we need some tests for quantifier elimination. In this section we give two. The first one is simple:

Definition 13.5. A *literal* is an atomic formula or a negated atomic formula. A formula will be called *primitive* if it is of the form

$$\exists x\, \varphi(\underline{y}, x)$$

where φ is a conjunction of literals.

Proposition 13.6. *A theory T has quantifier elimination if and only if any primitive formula is equivalent over T to a quantifier-free formula.*

Proof. Suppose every primitive formula is equivalent over T to a quantifier-free formula. Then every formula of the form

$$\exists x\, \varphi(\underline{y}, x)$$

with φ quantifier-free is also equivalent to a quantifier-free formula: for we can write $\varphi(\underline{y}, x)$ in disjunctive normal form, that is, as a disjunction $\bigvee_i \varphi_i(\underline{y}, x)$, where each $\varphi_i(\underline{y}, x)$ is a conjunctions of literals. Then we can push the disjunction through the existential quantifier, using that

$$\exists x \bigvee_i \varphi_i(\underline{y}, x) \leftrightarrow \bigvee_i \exists x\, \varphi_i(\underline{y}, x),$$

so that we are left with a disjunction of primitive formulas, which is equivalent to a quantifier-free formula, by assumption.

Now let φ be an arbitrary formula. We can rewrite φ into an equivalent formula using only $\neg, \wedge$ and $\exists$, and then, working inside out, eliminate all the existential quantifiers using the previous observation. $\square$

The second is a bit more complicated, but generally easier to apply.

Theorem 13.7. *Let κ be an infinite cardinal. A theory T has quantifier elimination if and only if, given*

(1) *two models M and N of T, where N is κ-saturated,*
(2) *a local isomorphism $f\colon \{a_1, \ldots, a_n\} \subseteq M \to N$, and*
(3) *an element $m \in M$,*

there is a local isomorphism $g\colon \{a_1, \ldots, a_n, m\} \subseteq M \to N$ which extends f.

Proof. Necessity is clear: if T has quantifier elimination, then any local isomorphism is an elementary map, so this follows from Lemma 11.6.

Conversely, let L be the language of T and suppose $\exists x\, \varphi(\underline{y}, x)$ is a primitive formula not equivalent over T to a quantifier-free formula in L. Extend the language with constants $\underline{c}$ and work in the extended language. Now let T_0 be the collection of all quantifier-free sentences which are a consequence over T of $\neg\exists x\, \varphi(\underline{c}, x)$. Then the union of T, T_0 and $\exists y\, \varphi(\underline{c}, y)$ has a model M.

Next, consider T_1, which consists of the theory T, all quantifier-free sentences in the extended language which are true in M, as well as the sentence $\neg\exists y\, \varphi(\underline{c}, y)$. This theory T_1 is consistent: for if not, there would be a quantifier-free sentence $\psi(\underline{c})$ which is false in M and

and which is a consequence of $\neg\exists x\, \varphi(\underline{c}, x)$ over T. But such a sentence must belong to T_0 and therefore be true in M. Contradiction!

So T_1 has a model N and we may assume that N is κ-saturated. Now let f be the map which sends the interpretation of c_i in M to its interpretation in N and let m be such that $M \models \varphi(\underline{c}, m)$. This f is a local isomorphism, but cannot be extended to one whose domain includes m, because $\exists y\, \varphi(\underline{c}, y)$ fails in N. □

3. Exercises

Exercise 41. Let $L = \{E\}$ where E is a binary relation symbol. For each of the following theories either prove that they have quantifier elimination, or give an example showing that they do not have quantifier elimination; in the latter case, also formulate a natural extension $T' \supseteq T$ in an extended language $L' \supseteq L$ in which they do have quantifier elimination.

(a) E is an equivalence relation with infinitely many equivalence classes, each having size 2.
(b) E is an equivalence relation with infinitely many equivalence classes, each having infinite size.
(c) E is an equivalence relation with infinitely many equivalence classes of size 2, infinitely many equivalence classes of size 3, and each equivalence class has size 2 or 3.

Exercise 42. Let $M = (\mathbb{Z}, s)$, where $s(x) = x + 1$, and let $T = \mathrm{Th}(M)$.

(a) Show that T has quantifier elimination.
(b) Give a concrete description of a countable ω-saturated model of T.
(c) Describe the type spaces of T.
(d) Show that $\mathrm{Th}(\mathbb{N}, s)$ does not have quantifier elimination.

Exercise 43. (a) Show that the theory of $(\mathbb{Z}, <)$ has quantifier elimination in the language where we add a function symbol s for the function $s(x) = x + 1$.
(b) Give a concrete description of a countable ω-saturated model of $\mathrm{Th}(\mathbb{Z}, <)$.
(c) Describe the type spaces of $\mathrm{Th}(\mathbb{Z}, <)$

Exercise 44. Let T be the theory of infinite vector spaces over $\mathbb{Q}$.

(a) Show that T has quantifier elimination.
(b) Which models of T are κ-saturated?
(c) Describe the type spaces of T.

CHAPTER 14

Examples

In this final chapter of this syllabus we discuss two more examples of nice theories to illustrate some of the concepts from this course.

1. Atomless Boolean algebras

DEFINITION 14.1. A *(bounded) lattice* L is a partial order in which every finite subset $A \subseteq L$ has a least upper bound (a *supremum* or *join*, written $\bigvee A$) and a greatest lower bound (an *infimum* or *meet*, written $\bigwedge A$). More concretely this means that L has a smallest element 0, a largest element 1 and that for any two elements $p, q \in L$ there are elements $p \wedge q$ and $p \vee q$ such that:

$$\begin{aligned} x \leq p \wedge q &\quad \Leftrightarrow \quad x \leq p \text{ and } x \leq q, \\ p \vee q \leq x &\quad \Leftrightarrow \quad p \leq x \text{ and } p \leq x. \end{aligned}$$

EXERCISE 45. Show that in any lattice $\wedge$ and $\vee$ are associative, commutative and idempotent (that is, $x \wedge x = x$ and $x \vee x = x$ hold). In addition, show that the absorbative laws $x = x \wedge (x \vee y)$ and $x = x \vee (x \wedge y)$ hold, as well as $0 \wedge x = 0$ and $1 \vee y = y$.

EXERCISE 46. Conversely, show that if L is a set equipped with two binary operations $\wedge$ and $\vee$ and in which there are elements $0, 1 \in L$ such that all the properties from the previous exercise hold, then there is a unique ordering on L turning L into a lattice. (Hint: observe that in a lattice we have $x \leq y$ iff $x = x \wedge y$ iff $y = x \vee y$.)

DEFINITION 14.2. A lattice L is called *distributive* if both distributive laws

$$\begin{aligned} x \wedge (y \vee z) &= (x \wedge y) \vee (x \wedge z), \\ x \vee (y \wedge z) &= (x \vee y) \wedge (x \vee z) \end{aligned}$$

are satisfied. A distributive lattice L is called a *Boolean algebra* if for any element $x \in L$ there is an element $\neg x \in L$ (its *complement*) for which both

$$x \wedge \neg x = 0 \quad \text{and} \quad x \vee \neg x = 1$$

hold.

EXAMPLE 14.3. For any set X the powerset $\mathcal{P}(X)$ is a Boolean algebra with order given by inclusion, meets and joins given by intersection and union, complements given by set-theoretic complement and smallest and largest elements $\emptyset$ and X.

EXAMPLE 14.4. If X is a topological space, then the clopens in X also form a Boolean algebra with the same operations as in the previous example.

EXERCISE 47. Show that in any lattice one distributive law implies the other.

Exercise 48. Let L be a distributive lattice and suppose $x \in L$ is a complemented element, meaning that there is an element $y \in L$ such that $x \wedge y = 0$ and $x \vee y = 1$. Show that for any other element $p \in L$, we have

$$x \wedge p = 0 \Longrightarrow p \leq y \quad \text{and} \quad x \vee p = 1 \Longrightarrow y \leq p.$$

Deduce that complements are unique.

Exercise 49. Show that if B is a Boolean algebra, then B^{op}, which is B with the order reversed, is a Boolean algebra as well. In fact, B and B^{op} are isomorphic with the isomorphism given by negating (taking complements). Deduce the De Morgan laws: $\neg(p \wedge q) = \neg p \vee \neg q$ and $\neg(p \vee q) = \neg p \wedge \neg q$.

For what follows we need to understand finitely generated Boolean algebras. Recall that a Boolean algebra B is finitely generated if there are elements $b_1, \ldots, b_n \in B$ such that B has proper Boolean subalgebra also containing the elements $b_1, \ldots, b_n$.

Theorem 14.5. *Finitely generated Boolean algebras are finite.*

Proof. Suppose B is generated by $b_1, b_2, \ldots, b_n$. Let C be the collection of elements in B that can be written as "conjunctions" of the form $c_1 \wedge c_2 \wedge \ldots \wedge c_n$ where c_i is either b_i or its complement, and let D the collection of elements in B that can be written as "disjunctions" of elements in C. The collections C and D are finite, because they contains at most 2^n and $2^{(2^n)}$ elements, respectively. But D is a Boolean subalgebra of B, because it contains 0 (no disjuncts), 1 (all disjuncts) and is closed under disjunction (clear), conjunction (by the distributive laws) and negation (by the De Morgan laws). So $B = D$ is finite; in fact, it contains at most $2^{(2^n)}$ many elements. □

So we need to understand finite Boolean algebras. But these are always of the form $\mathcal{P}(X)$, where X is finite. To show this, we need some definitions.

Definition 14.6. An element a in a Boolean algebra B is called an *atom* if $a > 0$ and there are no elements strictly in between a and 0. A Boolean algebra in which for any element $x > 0$ there is an atom a such that $x \geq a$ is called *atomic*. A Boolean algebra in which there are no atoms is called *atomless*.

Proposition 14.7. *Finite Boolean algebras are atomic.*

Proof. Let B is a finite Boolean algebra. Suppose $x_0 \in B$ is an element different from 0 and there are no atoms a with $x_0 \geq a$. This means that x_0 itself is no atom, so there is an element x_1 with $x_0 > x_1 > 0$. Of course, x_1 cannot be atom, by our assumption on x_0, so there must be an element x_2 such that $x_0 > x_1 > x_2 > 0$. Continuing in this way we create an infinitely descending sequence of elements in B, which contradicts its finiteness. □

Proposition 14.8. *If B is an atomic Boolean algebra and $x < y$, then there is an atom $a \in B$ which lies below y, but not below x.*

Proof. If $x < y$, then $y \wedge \neg x \neq 0$ (for if $y \wedge \neg x = 0$, then $\neg x \leq \neg y$ and $x \geq y$ by the exercises). So there is an atom a with $y \wedge \neg x \geq a$. So we have $y \geq a$ and $\neg x \geq a$; but the latter implies that $x \not\geq a$, for if also $x \geq a$, then $0 = x \wedge \neg x \geq a$. □

Theorem 14.9. *All finite Boolean algebras B are of the form $\mathcal{P}(X)$ for a finite set X. In fact, X can be chosen to be the collection of atoms in B.*

Proof. Let B be a finite Boolean algebra and let A be its collection of atoms. Then we define maps $f\colon B \to \mathcal{P}(A)$ by sending $b \in B$ to the set $f(b) = \{a \in A\colon a \leq b\}$ and $g\colon \mathcal{P}(A) \to B$ by sending a set $X \subseteq A$ to $g(X) = \bigvee X$. It will suffice to prove that f and g are order preserving and each other's inverses (since all operations in a Boolean algebra are uniquely determined in terms of its order, any order isomorphism between Boolean algebras must be an isomorphism of Boolean algebras). That they are order preserving is clear, so we only check that they are each other's inverses.

So if b is an element in B and $X = \{a \in A\colon a \leq b\}$, then b is an upper bound for X, so $b \geq \bigvee X$. Here we must have equality: for if $b > \bigvee X$, then the previous two results imply that there is an atom a' such that $b \geq a'$ but not $\bigvee X \geq a'$. But the former implies that $a' \in X$ so we should have $\bigvee X \geq a'$ after all. Contradiction! We deduce $g(f(b)) = b$.

Conversely, let X be a set of atoms in B and $b = \bigvee X$. Clearly, all atoms in X are below b, but the converse is true as well: for suppose a' is an atom and $b \geq a'$. Then

$$0 < a' = (a' \wedge b) = a' \wedge \bigvee_{a \in X} a = \bigvee_{a \in X} (a' \wedge a).$$

So there must be an element $a \in X$ such that $a' \wedge a$ is not zero. But since a and a' are atoms and $a' \wedge a$ is below each of them, we must have $a = a \wedge a' = a'$. We deduce $f(g(X)) = X$, which finishes the proof. □

Lemma 14.10. *Suppose M and N are atomless Boolean algebras and $f\colon A \subseteq M \to N$ is a local isomorphism with A finite. Then for any $m \in M$ there is a local isomorphism $g\colon A' \subseteq M \to N$ with $A \cup \{m\} \subseteq A'$ and $g \restriction A = f$.*

Proof. Since finitely generated Boolean algebras are finite, we may assume, without loss of generality, that A is a Boolean subalgebra of M. Let us write $B = f(A)$: since f is a local isomorphism, f induces an isomorphism between A and B, which is a Boolean subalgebra of N. Let us also write $a_0, \ldots, a_{k-1}$ for the atoms in the Boolean algebra A, and $b_i = f(a_i)$; clearly, the b_i are the atoms of the Boolean algebra B.

For any $m \in M$, there are three possibilities for $m \wedge a_i$: it can be 0, or a_i or something in between. Let us call the function which says for every i which of these three scenarios happens, the *profile* of m. Similarly, we can define the profile of elements $n \in N$, but then with respect to the b_i instead of the a_i.

The proof will be finished once I show:

(1) For any $m \in M$ there is an $n \in N$ which has the same profile, and vice versa.
(2) If $m \in M$ and $n \in N$ have the same profile, then the local isomorphism f can be extended to one which sends m to n.

I will only sketch the argument: as for (1), let $I = \{i < k\colon m \wedge a_i = a_i\}$ and $J = \{j < k\colon 0 < (m \wedge a_j) < a_j\}$. For any $j \in J$ we consider b_j: since it is not an atom in N, we can choose an element $y_j \in N$ with $0 < y_j < b_j$.

Now put $n\colon= \bigvee_{i \in I} b_i \vee \bigvee_{j \in J} y_j$. Using that the b_i are atoms in B and we therefore have that $b_i \wedge b_j = 0$ whenever $i \neq j$, we see that n has the same profile as m.

As for (2): the crucial observation here is that if $J = \{j < k\colon 0 < (m \wedge a_j) < a_j\}$, then the atoms of the Boolean subalgebra generated by $a_0, \ldots, a_{k-1}$ and m are the a_i with $i \in J^c$ together with $a_j \wedge m$ and $a_j \wedge \neg m$ for every $j \in J$. Sending these to b_i, $b_j \wedge n$ and $b_j \wedge \neg n$, respectively, we have a function sending atoms to atoms which extends uniquely to a morphism of Boolean algebras: this morphism extends the original map and sends m to n. □

COROLLARY 14.11. *The theory ABA of atomless Boolean algebras is ω-categorical and complete and has quantifier elimination.*

PROOF. First of all, observe that any atomless Boolean algebra has to be infinite (by Proposition 14.7) and that there is a countable and atomless Boolean algebra (consider the clopens in Cantor space). Using the previous lemma and Theorem 13.7 one can show that ABA has quantifier elimination. The same lemma in combination with a standard back-and-forth argument shows that this theory is also ω-categorical and therefore complete. □

EXERCISE 50. Show that all Boolean algebras of the form $\mathcal{P}(X)$ are atomic, but that there are atomic Boolean algebras which are not of this form.

EXERCISE 51. Not so easy: show that ABA is not λ-categorical for any $\lambda > \omega$.

2. Real closed ordered fields

2.1. Ordered fields.

DEFINITION 14.12. An *ordered field* is a field equipped with a linear order $\leq$ satisfying

(1) if $x \leq y$, then $x + z \leq y + z$,
(2) if $x \leq y$ and $0 \leq z$, then $xz \leq yz$.

Let us call elements x for which $x \geq 0$ *positive*; otherwise x is called negative. Note that if x is negative, then $x < 0$ and

$$-x = 0 - x \geq x - x = 0,$$

so $-x$ is positive. Using property (2) and the observation that $x^2 = (-x)^2$, it follows that $1 = 1^2$ is positive and also $2, 3, 4, \ldots$ are positive. But -1 is negative and hence ordered fields always have characteristic 0.

DEFINITION 14.13. If K is a field, then we call a subset $P \subseteq F$ a *positive cone*, if:

(1) P is closed under sums and products.
(2) $-1 \notin P$.
(3) for any x, either x or $-x$ belongs to P.

PROPOSITION 14.14. *If K is an ordered field, then the elements $x \in K$ satisfying $x \geq 0$ form a positive cone. Conversely, if P is a positive cone on a field K, then K can be ordered by putting $x \leq y$ iff $y - x \in P$.*

In ordered fields sums of squares have to be positive. In fact, we have:

PROPOSITION 14.15. *Let K be a field and $r \in K$. If both -1 and r cannot be written as a sum of squares, then K can be ordered in such a way that r becomes negative.*

PROOF. Let S be the collection of those elements in K that can be written as sums of squares. This set has the following properties:

(1) it is closed under sums and products,
(2) it contains all squares,

(3) and it does not contain -1.

Such a set is called a *semipositive cone*. We use two properties of such sets: first, if X is a semipositive cone and $s \in X - \{0\}$, then $(\frac{1}{s})^2 \in X$ and hence also $\frac{1}{s} \in X$. And if X is a semipositive cone and $s \notin X$, then $X - sX$ is also semipositive cone. For if there would be $x_0, x_1 \in X$ such that $x_0 - sx_1 = -1$, then $x_1 \neq 0$ and

$$s = \frac{1 + x_0}{x_1} \in X.$$

So put $Y := S - rS$. This is a semipositive cone, and, using Zorn's Lemma, we can extend Y to a maximal semipositive cone Y_{max}. Then Y_{max} is a positive cone, for if $x \notin Y_{max}$, then $-x \in Y_{max} - xY_{max} = Y_{max}$. □

2.2. Some analysis in ordered fields.

Now suppose that K is an ordered field.

PROPOSITION 14.16. *Let $p(x) = x^d + a_{d-1}x^{d-1} + \ldots + a_0$ and $m = \max(|a_{d-1}|, \ldots, |a_0|) + 1$. Then all roots of $p(x)$ lie between $-m$ and m.*

PROOF. If $|x| \geq m$, then

$$|P(x) - x^d| \leq (|m| - 1)\,(|x|^{d-1} + |x|^{d-2} + \ldots + 1) \leq (|m| - 1)\,\frac{|x|^d - 1}{|x| - 1} \leq |x|^d - 1$$

so $P(x) \neq 0$. □

PROPOSITION 14.17. *If $p(x) \in K[x]$ and $p(0) > 0$, then there is an $\epsilon > 0$ such that $P(x) > 0$ for all $x \in [-\epsilon, +\epsilon]$.*

PROOF. Let $p(x) = a_d x^d + a_{d-1}x^{d-1} + \ldots + a_0$. Then put $m = \max(|a_d|, |a_{d-1}|, \ldots, |a_0|)$ and $\epsilon = \min(1, \frac{P(0)}{2md})$. Then $x \in [-\epsilon, +\epsilon]$ implies

$$\begin{aligned} |p(x) - p(0)| &\leq |a_d x^d + a_{d-1}x^{d-1} + \ldots + a_0 - a_0| \\ &\leq m\epsilon^d + m\epsilon^{d-1} + \ldots + m\epsilon \\ &\leq md\epsilon \\ &\leq \frac{1}{2}p(0) \end{aligned}$$

and hence $p(x) > 0$. □

PROPOSITION 14.18. *If $p(x) \in K[x]$ and $p'(a) > 0$, then there is an $\epsilon > 0$ such that $p(x) > p(a)$ for every $x \in (a, a + \epsilon]$ and $p(x) < p(a)$ for every $x \in [a - \epsilon, a)$.*

PROOF. Write $p(x) = (x-a)q(x) + p(a)$. Then $p'(x) = q(x) + (x-a)q'(x)$, so $q(a) = p'(a) > 0$. Then choose ϵ such that $q(x) > 0$ for all $x \in [a - \epsilon, a + \epsilon]$ using the previous result. □

2.3. Real closed ordered fields.

DEFINITION 14.19. An ordered field will be called *real closed* if it satisfies the intermediate value theorem for polynomials: if for any polynomial $P(x)$ and elements $a < b$ such that $P(a) < 0$ and $P(b) > 0$ there is an element $c \in (a, b)$ such that $P(c) = 0$.

For example, the field $\mathbb{R}$ is real closed, but $\mathbb{Q}$ is not.

Proposition 14.20. *In a real closed field an element is positive iff it can be written as a square.*

Proof. We already know that squares are positive. So suppose $a > 0$ and consider $p(x) = x^2 - a$. Then $p(a+1) = (a+1)^2 - a = a^2 + 2a + 1 - a = a^2 + a + 1 > 0$ and $p(0) < 0$, so there is an element r such that $p(r) = 0$ and hence $r^2 = a$. □

Exercise 52. Use the previous proposition to write down a set of first-order sentences T in the language of fields (without order!) such that M is a model T if and only if M is a field which can be ordered in such a way that it becomes real closed.

Theorem 14.21. *Let K be a real closed field and $p(x)$ be a polynomial over K. If $a < b \in K$ and $p'(x) > 0$ for all $x \in (a, b)$, then $p(a) < p(b)$.*

Proof. First suppose that $p'(a) > 0$ and $p'(b) > 0$. Then we can use Proposition 14.18 to find c, d with $a < c < d < b$ such that $p(a) < p(c)$ and $p(d) < p(b)$. So if $p(a) \geq p(b)$, then $p(c) > p(b) > p(d)$ and there is an $e_0 \in (c, d)$ such that $p(e_0) = p(b)$. By repeating this argument for e_i and b instead of a and b we find for every $i \in \mathbb{N}$ an $e_{i+1} \in (e_i, b)$ such that $p(e_{i+1}) = p(b)$, contradicting the fact that a polynomial can have only finitely many zeros.

In the general case choose arbitrary c, d such that $a < c < d < b$. We have $p(c) < p(d)$ by the previous argument. In addition, we have $p(a) \leq p(c)$, for if $p(a) > p(c)$, then there is an $e \in (a, c)$ such that $p(e) > p(c)$ by Proposition 14.17. But that would again lead to a contradiction by an argument as in the previous paragraph. Similary, $p(c) \leq p(d)$, so $p(a) < p(b)$. □

Corollary 14.22. (Rolle's Theorem for real closed ordered fields) *Let K be a real closed ordered field and $p(x)$ be a polynomial over K. If $p(a) = p(b)$ for $a < b$, then there exists $c \in (a, b)$ with $P'(c) = 0$.*

Proof. For if $P'(c) \neq 0$ for all $c \in (a, b)$, then P' is either strictly positive or strictly negative on (a, b), by real closure. □

2.4. Real closure.

Definition 14.23. Let $K \subseteq L$ be an order preserving embedding between ordered fields. L is a *real closure* of K, if L is algebraic over K and no ordered field properly extending L is algebraic over K.

Note, by the way, that an inclusion of ordered fields $K \subseteq L$ is order preserving iff it is order reflecting, because ordered fields are linearly ordered.

Theorem 14.24. *If L is a real closure of K, then L is real closed.*

Proof. Suppose there are polynomials in $L[x]$ for which the intermediate value theorem for polynomials fails. Let p be a counterexample of minimal degree: so the intermediate value theorem holds for polynomials in $L[x]$ with degree smaller than p, but there are $a < b \in L$ with $p(a) < 0$ and $p(b) > 0$ for which no $\xi \in (a, b)$ with $p(\xi) = 0$ exists.

In that case p has to be irreducible so $L[x]/(p(x))$ is a field extending L, still algebraic over K. So once we show that $L[x]/(p(x))$ can be ordered in a way which extends to the order on L, we have obtained our desired contradiction.

Let $A = \{x \in [a,b] : (\exists y \geq x)\, p(y) < 0\}$ and $B = [a,b] - A = \{x \in [a,b] : (\forall y \geq x)\, p(y) > 0\}$. Since polynomials are continuous, both A and B are open and have no greatest or least element, respectively. So if $q(x)$ is any non-zero polynomial, then q has only finitely many roots, so there are $a_0 \in A$ and $b_0 \in B$ such that q has no roots in the interval $[a_0, b_0]$. If $q(x)$ has a degree strictly smaller than $p(x)$, then the intermediate value theorem holds for $q(x)$ and $q(x)$ is either strictly positive or strictly negative on $[a_0, b_0]$. If the former holds we declare $q(x)$ positive. It is easy to see that this defines a positive cone on $L[x]/(p(x))$ extending the one on L. So we have our desired contradiction. □

Theorem 14.25. *Real closures exist and are unique up to unique isomorphism.*

Proof. The existence of real closures follows from Zorn's Lemma: consider all ordered extensions of a field K which are still algebraic over K and all field embeddings between them which preserve the ordering. Since fields algebraic over K have the same infinite cardinality as K, this is essentially a set. Since chains have upper bounds given by unions, a maximal element must exist, which is a real closure of K.

Now suppose both L_0 and L_1 are real closures of an ordered field K. By Zorn's Lemma, again, there are subfields $K_0 \subseteq L_0$ and $K_1 \subseteq L_1$ between which there exists an order preserving isomorphism f which leaves K invariant and which is maximal with these properties. If either $L_0 - K_0$ or $L_1 - K_1$ is non-empty, then we may assume, without loss of generality, that there is an element $\xi \in L_0 - K_0$ with minimal polynomial $p(x)$ over K such that all other elements $\xi' \in L_i - K_i$ have a minimal polynomial over K whose degree is at least that of p.

Since p is minimal, we have $p'(\xi) \neq 0$, so p changes sign in ξ. Moreover, in L_1 and L_2 it holds that in between any two roots of $p(x)$ lies a root of $p'(x)$, by Rolle's Theorem. Since roots of $p'(x)$ have a minimal polynomial whose degree is strictly smaller than that of $p(x)$, these roots of $p'(x)$ lie already in K_0 and K_1. So for ξ there are three possibilities:

(1) ξ lies in between two roots of $p'(x)$, call them x_0 and x_1, and it is the only root lying in this interval. In that case p has different signs in x_0 and x_1. So the same applies to $f(x_0)$ and $f(x_1)$ and the polynomial p can have only one root in K_1 in between these points. Then ξ should be sent to this root.
(2) ξ is bigger than the largest root of $p'(x)$. Let x_0 be this largest root and let x_1 be a number in K bounding the zeros of p from above (using Proposition 14.16). Then again p changes sign between x_0 and x_1 and ξ should be sent to the unique root of p in K_1 between $f(x_0)$ and $f(x_1)$.
(3) ξ is smaller than the smallest root of $p'(x)$. Then the same argument as in (2) applies.

This determines a field isomorphism between $K(\xi) \cong K[x]/(p(x)) \cong K(\xi')$. The question now is why this field isomorphism should be order preserving. But this follows from the following observation: if $q(x)$ is any non-zero polynomial of degree strictly smaller than $p(x)$, then q is strictly positive or negative on some interval $[x_2, x_3]$ with $x_2, x_3 \in K_0$ and $x_0 < x_2 < \xi < x_3 < x_1$. So the sign of $q(\xi)$ in L_0 can be determined by checking the sign of $q(x_2)$ and the sign of $q(\xi')$ in L_1 can be determined by checking to sign of $q(f(x_2))$. But both answers should agree because f is an order preserving isomorphism.

So we have an isomorphism between L_0 and L_1. This isomorphism is necessarily unique because it should send the nth root from the left of the polynomial $p(x) \in K[x]$ in L_0 to the nth root from the left of $p(x)$ in L_1. □

2.5. Quantifier elimination.

Theorem 14.26. *The theory RCOF of real closed ordered fields has quantifier elimination.*

Proof. We use Theorem 13.7. So let K, L be two real closed ordered fields, where L in addition is ω_1-saturated, and suppose $f\colon \{k_1, \ldots, k_n\} \to L$ is a local isomorphism and $k \in K$. Then $\mathbb{Q}(k_1, \ldots, k_n)$, considered as an ordered subfield of K, and $\mathbb{Q}(f(k_1), \ldots, f(k_n))$, considered as an ordered subfield of L, are isomorphic. So we can use the previous theorem to extend f to an isomorphism $\overline{f}$ of ordered fields between the real closure $\overline{K}$ of $\mathbb{Q}(k_1, \ldots, k_n)$ inside K and the real closure $\overline{L}$ of $\mathbb{Q}(f(k_1), \ldots, f(k_n))$ inside L. If $k \in \overline{K}$, then we send k to $\overline{f}(k)$. So the interesting case is where k is transcendental over $\overline{K}$. To simplify notation, we will assume $\overline{K} = \overline{L}$.

In that case we should send k to an element $l \in L$ which is transcendental over the subfield $\overline{K}$ and for which

$$(\forall x \in \overline{K})\, x \leq k \Leftrightarrow x \leq l$$

holds. Such an element certainly exists because $|\overline{K}| = \omega$ and L is assumed to be ω^+-saturated. And this is enough, for to see that the composite isomorphism

$$\overline{K}(k) \cong \overline{K}(x) \cong \overline{K}(l)$$

is order preserving it suffices to check that $p(k)$ and $p(l)$ have the same sign for every irreducible polynomial $p \in \overline{K}[x]$. This is true for irreducible polynomials of degree one (by construction), and if p has degree greater than one, then p has no roots in K or L (since $\overline{K}$ is maximal as an algebraic extension over $\mathbb{Q}(k_1, \ldots, k_n)$ inside K or L). So p does not change sign inside K or L and $p(k)$ and $p(l)$ have the same sign as $p(0)$. □

Corollary 14.27. *The theory RCOF is complete.*

Proof. Since the theory of real closed ordered fields has quantifier elimination and has a model which can be embedded into any other model (to wit, the real numbers which are algebraic over $\mathbb{Q}$), this theory is complete by Corollary 13.4. □

Remark 14.28. The theory RCOF is not λ-categorical for any infinite λ, but that is not so easy to prove!

2.6. Hilbert's 17th Problem.

Theorem 14.29. (Hilbert's 17th Problem) *Let K be a real closed field. If $f \in K(x_1, \ldots, x_n)$ is such that $f(a_1, \ldots, a_n) \geq 0$ for all $a_1, \ldots, a_n \in K$, then f can be written as*

$$f = g_1^2 + \ldots + g_n^2$$

for suitable $g_i \in K(x_1, \ldots, x_n)$.

Proof. Suppose f cannot be written as a sum of squares in $K(x_1, \ldots, x_n)$. The same applies to -1, because -1 cannot be written as a sum of squares in K. So we can order $K(x_1, \ldots, x_n)$ in such a way that f becomes negative. This order extends the original order on K because K is real closed and hence positive elements in K can be written as squares (see Proposition 14.20). Now embed $K(x_1, \ldots, x_n)$ with this order into a real closed field L. So we have embeddings of fields

$$K \subseteq K(x_1, \ldots, x_n) \subseteq L,$$

all of which preserve and reflect the ordering. So the inclusion $K \subseteq L$ reflects truth of atomic sentences, and hence of quantifier-free sentences and hence, as the theory of real closed fields has quantifier elimination, of *all* sentences. Therefore the sentence

$$\exists x_1 \ldots \exists x_n \, f(x_1, \ldots, x_n) < 0,$$

which is true in L, must be true in K as well. □

REMARK 14.30. Hilbert's 17th Problem asked whether Theorem 14.29 holds in case K is the reals. This was settled by Artin in 1927, who proved the result for general real closed fields. The model-theoretic proof we just gave is due to Robinson.