CHAPTER 6

Types

1. Terminology

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

DEFINITION 6.1. Fix $n \in \mathbb{N}$ and let $x_1, \ldots, x_n$ be a fixed sequence of distinct variables. If A is an L-structure and $a_1, \ldots, a_n \in A$, then the *type of* $(a_1, \ldots, a_n)$ *in* A is the set of L-formulas

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

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

Some observations:

- If $i \colon A \to B$ is an elementary embedding and $a_1, \ldots, a_n \in A$, then $(a_1, \ldots, a_n)$ and $(f(a_1), \ldots, f(a_n))$ have the same type.
- Two n-tuples $(a_1, \ldots, a_n)$ from A and $(b_1, \ldots, b_n)$ satisfy the same n-type precisely when $(A, a_1, \ldots, a_n) \equiv (B, b_1, \ldots, b_n)$. (This is supposed to mean: add new constants $c_1, \ldots, c_n$ 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 properties: this is a partial type.

DEFINITION 6.2. A *partial n-type in L* is a collection of formulas $\varphi(x_1, \ldots, x_n)$ in L.

- If $p(x_1, \ldots, x_n)$ is a partial n-type in L, we say $(a_1, \ldots, a_n)$ *realizes p in A* if every formula in p is true of $a_1, \ldots, a_n$ in A.
- 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 *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 types among the partial types? Clearly, they can be realized in some model, and they have to be complete: they contain $\varphi(x_1, \ldots, x_n)$ or $\neg\varphi(x_1, \ldots, x_n)$ for every L-formula φ whose free variables are among the fixed variables $x_1, \ldots, x_n$. This is sufficient: for if a partial n-type p is realized by $(a_1, \ldots, a_n)$, we must have $p \subseteq \mathrm{tp}(a_1, \ldots, a_n)$. If p is also complete, then $p \supseteq \mathrm{tp}(a_1, \ldots, a_n)$ follows as well. (For if $\varphi \notin p$, then $\neg\varphi \in p$, so $\neg\varphi \in \mathrm{tp}(a_1, \ldots, a_n)$, hence $\varphi \notin \mathrm{tp}(a_1, \ldots, a_n)$.)

2. Types and theories

DEFINITION 6.3. 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 6.4. *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 6.5. 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 6.6. *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 partial types that are finitely satisfiable have some properties that will often be used. The next lemma summarises a number of them.

LEMMA 6.7. *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 $T = \mathrm{ElDiag}(M) \cup \{\varphi(c) \colon \varphi \in p\}$, 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. Isolated types

The set $S_n(T)$ is not just a set, but it is in fact a topological space. To see this, consider sets in $S_n(T)$ of the form

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

where $\varphi(x_1, \ldots, x_n)$ is some formula. Since

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

these sets form a basis: the topology generated from the sets is called the *logic topology* and we have:

THEOREM 6.8. *The space $S_n(T)$ with the logic topology is a compact Hausdorff space.*

PROOF. 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

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

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 be inconsistent with T, by the compactness theorem. But then any $p \in S_n(T)$ must contain at least one of the φ_{i_k}: for it contains either φ_{i_k} or $\neg\varphi_{i_k}$ and cannot contain all $\neg\varphi_{i_k}$. Therefore

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

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 6.4). 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. □

Definition 6.9. 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

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

for all $\sigma(\overline{x}) \in p(\overline{x})$. A formula $\varphi(\overline{x})$ is called *complete* or *isolating* over T in case 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 $\varphi(\overline{x})$.

Proposition 6.10. *Let T be a theory and p be a complete type of T. Then the following are equivalent:*

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

$$\psi(x_1, \ldots, x_n) \in p \Leftrightarrow T \models \varphi(x_1, \ldots, x_n) \to \psi(x_1, \ldots, x_n).$$

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

Proposition 6.11. *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 any 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

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

for all $\sigma(\overline{x}) \in p(\overline{x})$. 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 $M \models \sigma(\overline{m})$ for every $\sigma \in p$, so $\overline{m}$ realizes p. □

CHAPTER 7

ω-saturated models

1. Definition

In this chapter we study an important class of models: ω-saturated models. The idea is that such models are "rich": in particular, if M is an ω-saturated model of a complete theory T, then any type of T will be realized in M (see below). Roughly speaking this means that any type of object (or type of a sequence of objects) that occurs in some model of T already occurs in M: so any type of thing whose existence is compatible with T lives in M.

First, 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 7.1. Let A be an L-structure. We say that A is *ω-saturated* if the following condition holds:

> if X is any finite subset of A having cardinality 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 $Y \subseteq A$ is finite and A is ω-saturated, then so is $(A, y)_{y \in Y}$. The reason for this is that any 1-type over a finite set of paramaters X in $(A, y)_{y \in Y}$ is also a 1-type over the finite set of parameters $X \cup Y$ in A.
(2) The definition of ω-saturation only talks about 1-types; however, if $p(x_1, \ldots, x_n)$ is an n-type over a finite set of parameters X that is finitely satisfiable in an ω-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 ω-saturated; moreover, if $a_1, \ldots, a_i$ realize p_i, then p_{i+1} is finitely satisfied in $(A, y)_{y \in X \cup \{a_1, \ldots, a_i\}}$, by Lemma 7.2 below, and hence realized by some a_{i+1} by the previous remark. So we see that each p_i is realized, which includes $p = p_n$.
(3) In the same vein we can observe that the fact that the definition of ω-saturation only takes about *complete* types is not a genuine restriction: for, by Lemma 6.7, any partial type that is finitely satisfied can be extended to a complete type that is finitely satisfied.

LEMMA 7.2. *Let $p(x_1, \ldots, x_n, y)$ be an $(n+1)$-type and let $q(x_1, \ldots, x_n)$ 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 $(a_1, \ldots, a_n)$ realizes p in M, then also $p(a_1, \ldots, a_n, y)$ is finitely satisfiable in M.*

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

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

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:

LEMMA 7.3. *Let M be an ω-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 6.6. So if M is ω-saturated, then p will be realized. $\square$

2. Existence

Examples of ω-saturated structures are: any dense linear order, any random graph. A non-example is $(\mathbb{N}, <)$. However, we do have:

THEOREM 7.4. *Every structure has an ω-saturated elementary extension. So any consistent theory has an ω-saturated model.*

The proof relies on the following lemma:

LEMMA 7.5. *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 7.4.) 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. $\square$

3. Properties of ω-saturated models

In this section we establish some special properties of ω-saturated models.

DEFINITION 7.6. Let A and B be L-structures and $X \subseteq A$. A map $f\colon X \to B$ will be called an *elementary map* if

$$A \models \varphi(a_1, \ldots, a_n) \Leftrightarrow B \models \varphi(f(a_1), \ldots, f(a_n))$$

for all L-formulas φ and $a_1, \ldots, a_n \in X$. Note that this is equivalent to

$$(A, x)_{x \in X} \equiv (B, fx)_{x \in X}.$$

A model M is called *ω-homogeneous*, if for any finite subset X of M, elementary map $f\colon X \to M$ and $a \in M$, the map f can be extended to an elementary map g whose domain includes a.

THEOREM 7.7. *If M is ω-saturated, then M is ω-homogeneous.*

PROOF. Suppose we are given a finite subset X of M, an elementary map $f\colon X \to M$ and an element $a \in M$. Let p be the type of a. Since $(M, x)_{x \in X} \equiv (M, fx)_{x \in X}$, we have by Lemma 6.7 that p is finitely satisfiable in $(M, fx)_{x \in X}$. Since this model is ω-saturated, there is an element b realizing p in $(M, fx)_{x \in X}$, so that

$$(M, a, x)_{x \in X} \equiv (M, b, fx)_{x \in X}$$

and we can extend f by putting $g(a) = b$. □

THEOREM 7.8. *Let M be an ω-saturated model. If A is countable and elementary equivalent to M, then there is an elementary embedding from A to M.*

PROOF. Let $\{a_1, a_2, \ldots\}$ be an enumeration of A. We construct an increasing sequence of elementary maps $f_n\colon \{a_1, \ldots, a_n\} \to M$. The desired embedding will then be $f = \bigcup_n f_n$. We start by putting $f_0 = \emptyset$, which is an elementary map because A and M are elementary equivalent.

If f_n has been defined, write $m_i = f(a_i)$. Consider $p = \mathrm{tp}(a_{i+1})$ in $(A, a_1, \ldots, a_n)$. Because f_n is an elementary map, the models $(A, a_1, \ldots, a_n)$ and $(M, m_1, \ldots, m_n)$ are elementarily equivalent, so the type p, which is realized in $(A, a_1, \ldots, a_n)$ is finitely satisfied in $(M, m_1, \ldots, m_n)$. Because $(M, m_1, \ldots, m_n)$ is ω-saturated there is an element m_{n+1} realizing p and we can extend f_n by putting $f_{n+1}(a_{n+1}) = m_{n+1}$. □

CHAPTER 8

Countable ω-saturated models

1. Small theories

The model $(\mathbb{Q}, <)$ is not just ω-saturated, but also countable. We have seen that every consistent theory has an ω-saturated model (Theorem 7.4), but it is not true that every theory has a *countable* ω-saturated model, like the theory of dense linear orders.

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 $T = \mathrm{Th}(M)$. 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. (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 $f(n) = 0$.)

Indeed, 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 8.1. A theory T is *nice* (this is not standard terminology), if it is complete, formulated in a countable language and has infinite models. (Note that nice theories cannot have finite models.) A theory is *small* if all its type spaces are countable.

Theorem 8.2. *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 7.4 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 7.5 may be taken to be countable as well. And that in turn means that in the proof of Theorem 7.4 we may consider a countable chain of countable models: but then its colimit, which was an ω-saturated model, is countable as well. $\square$

2. Properties

THEOREM 8.3. *Any two countable ω-saturated models that are elementarily equivalent are isomorphic.*

PROOF. Let $\{a_1, a_2, a_3, \ldots\}$ and $\{b_1, b_2, b_3, \ldots\}$ be enumerations of two countable ω-saturated models A and B. The idea is to construct elementary maps f_n from some subset of A to B, such that the domain of f_n includes $a_1, \ldots, a_n$ and the range includes all of $b_1, \ldots, b_n$. The union $f = \bigcup_n f_n$ will then be an isomorphism between A and B. We start by putting $f_n = \emptyset$, which is an elementary map if A and B are elementarily equivalent. If $f_n \colon X_n \to B$ has been constructed, then we start by finding an element b which realizes in $(B, x)_{x \in X_n}$ the type of a_n over X_n and then by finding an element a which realizes in $(A, a_n, x)_{x \in X_n}$ the type of b_n over $\{b\} \cup f(X_n)$. Then we put $f_{n+1} = f_n \cup \{(a_n, b), (a, b_n)\}$. $\square$

DEFINITION 8.4. A model M is called *strongly ω-homogeneous* if any elementary map $f \colon X \to M$, where X is a finite subset of M, can be extended to an automorphism of M.

THEOREM 8.5. *A countable ω-saturated model is strongly ω-homogeneous.*

PROOF. The proof is a variation on the previous one. $\square$

CHAPTER 9

The number of countable models

1. The omitting types theorem

In the previous chapter we have seen that we can read off from the type spaces of a nice theory T whether a countable ω-saturated model exists: it exists if and only if all the type spaces are countable. We will see that there is a similar characterisation for when the theory T is ω-categorical: this happens precisely when all the type spaces are finite (this is the Ryll-Nardzewski Theorem). This result relies on another theorem, the Omitting Types Theorem. As we have seen, realizing a type is easy: just go to an ω-saturated model. Omitting a type is much harder, and we know that it is impossible in case the type is isolated. But if the type is not isolated and the language is countable, then we can manage to omit it.

Theorem 9.1. *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 $p(x)$.*

The proof relies on a lemma:

Lemma 9.2. *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*

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

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

Proof. Extend T to a maximally consistent theory using the Lemma on page 4 of the slides for week 2 and then apply the Lemma on page 3 of the slides for week 2. □

Proof. (Of Theorem 9.1.) 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 $\varphi_0, \varphi_1, \varphi_2, \ldots$. The idea is to apply to previous lemma to $T \cup \{\varphi_0, \varphi_1, \ldots\}$.

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

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

This makes sure the condition in Lemma 9.2 is satisfied and we can create a model from the constants in C.

If $n = 2i + 1$ we make sure that c_i omits $p(x)$, as follows. Consider $\delta = \bigwedge_{m<n} \varphi_m$. δ is really of the form $\delta(c_i, \overline{c})$ 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, such that $T \cup \{\exists \overline{y}\, \delta(x,y)\} \cup \{\neg\sigma(x)\}$ is consistent. Put $\varphi_{2n} = \neg\sigma(c_i)$.

The proof is now finished by showing by induction that each $T \cup \{\varphi_0, \dots, \varphi_n\}$ is consistent and then applying the previous lemma. □

2. ω-categoricity

PROPOSITION 9.3. *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 $\varphi(x_1, \dots, x_n)$ 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

$$p_1(\underline{x}), \dots, p_m(\underline{x}),$$

then each of these isolated, so there are formulas $\psi_1(x_1, \dots, x_n), \dots, \psi_m(x_1, \dots, x_n)$ isolating $p_1, \dots, p_m$. We claim that each formula is equivalent over T to a disjunction of the ψ_i.

Let $\varphi(x_1, \dots, x_n)$ be arbitrary and put

$$I = \{i \in \{1, \dots, m\} : \varphi \in p_i\}.$$

If $i \in I$, then $T \models \psi_i \to \varphi$, because ψ_i isolates p_i, so $T \models \bigvee_{i \in I} \psi_i \to \varphi$. Conversely, if $T \not\models \varphi \to \bigvee_{i \in I} \psi_i$, then

$$\{\varphi(\underline{x})\} \cup \{\neg\psi_i(x) : i \in I\}$$

is a partial type consistent with T, so can be extended to a complete type $p \in S_n(T)$. Since $\varphi \in p$ we must have that $p = p_i$ for some $i \in I$; on the other hand, p also contains $\neg\psi_i(\underline{x})$, so cannot be equal to p_i. Contradiction.

(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 which it does not contain. □

THEOREM 9.4. *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 ω-saturated;*
(4) *all countable models of T are ω-saturated.*

PROOF. (1) $\Rightarrow$ (2) If T contains a non-isolated type then there is a countable model where it is realized and a countable model where it is not realized (by the Omitting Types Theorem).

(2) $\Rightarrow$ (3) Let $p(\underline{a}, x)$ be 1-type with n parameters from a model A and assume that $p(\underline{a}, x)$ is finitely satisfied in A. Since $p(\underline{y}, x)$ is isolated, there is a formula $\varphi(\underline{y}, x)$ isolating it: but then $\varphi(\underline{a}, x)$ isolates $p(\underline{a}, x)$. Moreover, since $p(\underline{a}, x)$ is finitely satisfied in A we have an element $a \in A$ satisfying $\varphi(\underline{a}, x)$: but then a realizes the type $p(\underline{a}, x)$.

(3) $\Rightarrow$ (4) is clear.

(4) $\Rightarrow$ (1) holds, because elementarily equivalent countable ω-saturated models of cardinality are isomorphic (this was Theorem 8.3). $\square$

COROLLARY 9.5. *If A is a model in a countable language and $a_1, \ldots, a_n$ are elements from A, then* $\mathrm{Th}(A)$ *is ω-categorical iff* $\mathrm{Th}(A, a_1, \ldots, a_n)$ *is ω-categorical.*

PROOF. Every m-type $p(\underline{a}, \underline{x})$ of $\mathrm{Th}(A, a_1, \ldots, a_n)$ determines an $(n+m)$-type $p(\underline{y}, \underline{x})$ 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 $\mathrm{Th}(A, a_1, \ldots, a_n)$.

Conversely, every m-type p consistent with $\mathrm{Th}(A)$ can be extended to an n-type consistent with $\mathrm{Th}(A, a)_{a \in A}$ (by Robinson's Consistency Theorem). Since these extensions have to be different for different types, $\mathrm{Th}(A)$ cannot have more n-types than $\mathrm{Th}(A, a)_{a \in A}$. So if the latter has only finitely many n-types, then so does the former. $\square$

THEOREM 9.6. (Vaught's Theorem) *A nice theory cannot have exactly two countable models (up to isomorphism).*

PROOF. Suppose T is a nice theory with precisely two countable models. Clearly, T cannot be ω-categorical, but it has to be small: because T is a theory in a countable language every n-type can be realized in some countable model. Moreover, since every countable model can realize only countably many n-types and there are only two countable models, the total number of n-types has to be countable.

We will now show that T has at least three models. First of all, there is a countable ω-saturated model A, because T is small. In addition, because T is not ω-categorical, there is a non-isolated type p which is omitted in some countable model B. Of course, it is realized in A by some tuple $\overline{a}$. Since $\mathrm{Th}(A, \overline{a})$ also not ω-categorical by the previous corollary, it has a model C different from A. Since C realizes p, it must be different from B as well. Contradiction. $\square$

CHAPTER 10

Prime models

1. Existence of prime models

In this chapter we will study *prime models*: these are models of a theory that can be elementarily embedded into any other model of the theory. We will see that also the existence of prime models can be "read off" from the type spaces.

DEFINITION 10.1. Let T be a theory. A model M of T is called *prime* if it can be elementarily embedded into any model of T. A model M of T is called *atomic* if it only realises isolated types (or, put differently, omits all non-isolated types) in $S_n(T)$.

THEOREM 10.2. *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. Let $\{a_1, a_2, \ldots\}$ be an enumeration of A; by induction on n we will construct an increasing sequence of elementary maps $f_n\colon \{a_1, \ldots, a_n\} \to M$. We start with $f_0 = \emptyset$, which is elementary as A and M are elementarily equivalent. (They are both models of a complete theory T.)

Suppose f_n has been constructed. The type of $a_1, \ldots, a_{n+1}$ in A is isolated, hence generated by a single formula $\varphi(x_1, \ldots, x_{n+1})$. In particular, $A \models \exists x_{n+1}\, \varphi(a_1, \ldots, a_n, x_{n+1})$, and since f_n is elementary, $M \models \exists x_{n+1}\, \varphi(f_n(a_1), \ldots, f_n(a_n), x_{n+1})$. So choose $m \in M$ such that $M \models \varphi(f_n(a_1), \ldots, f_n(a_n), m)$ and put $f(a_{n+1}) = m$. □

THEOREM 10.3. *A nice theory T has a prime model iff the isolated n-types are dense in $S_n(T)$ for all n.*

PROOF. Recall that a formula $\varphi(\overline{x})$ is called *complete* in T if it generates an isolated type in $S_n(T)$: that is, it is consistent and for any other formula $\psi(\overline{x})$ we have either $T \models \varphi(\overline{x}) \to \psi(\overline{x})$ or $T \models \varphi(\overline{x}) \to \neg\psi(\overline{x})$. To say that the isolated types are dense means that every non-empty (basic) open set contains at least one isolated type: so the n-types are dense iff every consistent formula $\varphi(\overline{x})$ follows from some complete formula.

$\Rightarrow$: 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 $\mathrm{tp}_A(\overline{a})$.

$\Leftarrow$: 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

$$T \models \psi(\overline{x}) \to \neg\varphi(x_1, \ldots, x_n)$$

for any complete formula $\varphi(x_1, \ldots, x_n)$. 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. □

2. More on small theories

The aim of this section is to prove that small theories have prime models. In view of the previous theorem this basically amounts to showing that the isolated types are dense in all the type spaces of a small theory. In order to prove this we need to understand these small theories a bit better.

Definition 10.4. Let $\{0, 1\}^*$ be the set of finite sequences consisting of zeros and ones. A *binary tree* of formulas in variables $\overline{x} = x_1, \ldots, x_n$ (in 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 10.5. *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 $x_1, \ldots, x_n$.*
(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

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

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

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

By assumption this collection contains a formula $\psi(\overline{x})$ or its negation, but not both. In addition, if p contains both $\psi_0 \vee \psi_1$, 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 $\psi_1, \ldots, \psi_n$ 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) $\Rightarrow$ (1): This is clear, because $\omega < 2^\omega$. $\square$

Corollary 10.6. *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. $\square$

APPENDIX A

Topology

Definition A.1. A *topological space* is a pair (X, Ω) consisting a set X and a collection Ω of subsets of X, where the subsets in Ω have the following properties:

(i) Both $\emptyset$ and X belong to Ω.
(ii) If $U, V \in \Omega$, then also $U \cap V \in \Omega$.
(iii) If $U_i \in \Omega$ for some collection of subsets $(U_i)_{i \in I}$ of X, then also $\bigcup_{i \in I} U_i \in \Omega$.

Such a collection Ω is also called a *topology on* X. The elements in X are called the *points* and the elements in Ω the *opens* of the topological space.

Some more terminology: a set whose complement is open is called *closed* and a set which is both open and closed is called *clopen.* If U is open and $x \in U$, then U is called a *neighborhood* of x. A subset $A \subseteq X$ is called *dense* if any non-empty open set contains at least one element from A.

Definition A.2. Let X be a set. A collection $\mathcal{B}$ of subsets of X that are is under binary intersections (that is, if $U, V \in \mathcal{B}$, then $U \cup V \in \mathcal{B}$) and such that $\bigcup \mathcal{B} = X$ is called a *basis.*

Proposition A.3. *Let X be a set and $\mathcal{B}$ be a basis. Then we can define a topology on X by saying that a subset U is open precisely when it arises as a union of elements in $\mathcal{B}$, that is, can be written as:*

$$U = \bigcup S, \qquad \textit{for some } S \subseteq \mathcal{B}.$$

The topology defined in this way is called the topology generated by the basis *$\mathcal{B}$.*

Definition A.4. A topological space (X, Ω) is called *Hausdorff*, if for any two points $x, y \in X$ with $x \neq y$ there are disjoint open sets U and V with $x \in U$ and $y \in V$.

Proposition A.5. *Singleton sets are closed in a Hausdorff space.*

Proof. Let (X, Ω) be a Hausdorff space and fix $x \in X$. Then for every $y \in X$ which is distinct from x one can define an open set N_y which contains y but does not contain x. But then

$$\bigcup_{y \in X, y \neq x} N_y = X \setminus \{x\}$$

is open, so $\{x\}$ is closed. □

Definition A.6. A point $x \in X$ is *isolated* in a topological space (X, Ω) in case $\{x\}$ is open.

Definition A.7. A topological space (X, Ω) is called *compact* if for any collection of subsets $(U_i)_{i \in I}$ such that $\bigcup_{i \in I} U_i = X$ there is a finite subset $I_0 \subseteq I$ such that $\bigcup_{i \in I_0} U_i = X$.

Proposition A.8. *Let (X, Ω) be a compact Hausdorff space. Then every point is isolated in (X, Ω) iff X is finite.*

Proof. $\Rightarrow$: If every point in X is isolated then we can write X as a union of open subsets, as follows:

$$X = \bigcup_{x \in X} \{x\}$$

So if (X, Ω) is compact, then there is a finite subset $X_0 \subseteq X$ such that

$$X = \bigcup_{x \in X_0} \{x\}.$$

But then $X = \bigcup_{x \in X_0} \{x\} = X_0$ is finite.

$\Leftarrow$: Suppose X is finite and let $x \in X$ be arbitrary. For any $y \in X$ different from x we can find an open subset U_y which contains x but does not contain y. Then

$$\{x\} = \bigcap_{y \in Y, y \neq x} U_y$$

is a finite intersection of open sets and hence open. $\square$